On Perturbed Best Response Learning and Equilibrium Selection in Homogeneous and Bayesian Evolutionary Games

by

Sayan Mukherjee

A dissertation submitted to

The Economic Research Unit

in partial fulfillment of the requirements

for the Degree of

Doctor of Philosophy

in the subject of

QUANTITATIVE ECONOMICS

Indian Statistical Institute

Kolkata, West Bengal, India

2025

DEDICATION

To my family and friends.

ACKNOWLEDGEMENTS

I express my heartfelt gratitude to my Ph.D supervisor, Prof. Souvik Roy, for his invaluable guidance, support, and encouragement throughout my doctoral journey. I am deeply grateful for his patience, mentorship and unwavering belief in my potential, which not only provided me with the freedom to think independently but also inspired me to push my limits and work on fundamental problems besides introducing me to the field of evolution and learning in games. This thesis would not have been possible without his generous investment of time and effort. It goes without saying that several people are involved in the journey of a successful Ph.D candidate. The same is true for mine. But if I were to thank a single person in this regard, it would undoubtedly be my supervisor.

I sincerely thank Prof. Ratul Lahkar for his collaboration during my Ph.D. From the very beginning, Prof. Lahkar demonstrated remarkable efficiency and dedication to the projects we worked on together. I am pleased to share that over the past five years, we successfully completed three projects, all of which form integral parts of my Ph.D thesis. I have always been captivated by the application of probability theory to real-world problems. I am deeply thankful to Dr. Soumendu Sundar Mukherjee for collaborating with me on a project exploring mean-field stochastic evolution in games. Additionally, Dr. Mukherjee's advanced course on probability theory, 'Weak Convergence and Empirical Processes' at the Indian Statistical Institute, was pivotal in shaping my understanding and proved instrumental in the development of nearly all the papers that constitute my Ph.D. thesis. I also extend my warm regards to Dr. Yannick Viossat for hosting me at the Centre de Recherche en Mathématiques de la Décision (CEREMADE) during my visit to Université Paris Dauphine-PSL in France, from May to July 2022. This visit facilitated collaborative research in Adaptive Cancer Therapy, a sub-field of mathematical oncology that utilizes evolutionary game theory to model cancer cell dynamics and develop innovative strategies for sustainable cancer treatment.

I would like to take a moment to thank the faculty members of the Economic Research Unit (ERU) for the numerous courses that laid the foundation for my understanding of Economics. Coming from a background in Statistics, these courses were instrumental in equipping me with the knowledge and skills essential for pursuing a Ph.D in Quantitative Economics. I am also deeply thankful to my fellow Ph.D colleagues in ERU for enriching my academic journey. From engaging discussions after seminars and workshops to non-academic moments during departmental lunches, picnics, and other events, these interactions have been truly memorable.

A heartfelt cheer to my academic brother, Abhigyan, for his unwavering camaraderie and steadfast presence throughout the journey of my Ph.D. Your friendship, support, and shared experiences made this challenging path much more fulfilling. I would also like to express my gratitude to Soumyarup Da, Ujjwal Da, and Madhuparna Di for their unconditional support during the early stages of my Ph.D. Their passion for research, dedication, and willingness to guide me were invaluable and served as a constant source of motivation. Although we were technically "colleagues," our bond went far beyond professional interactions, evolving into a deep and lasting relationship that continues to thrive even though they completed their Ph.D years ago. I am truly fortunate to have such remarkable individuals as peers, and seniors in my academic journey. These experiences and bonds will remain cherished for lifetime.

Countless tea and coffee sessions at Laxman Da and Sanjib Da’s tea stalls with the “Core Group" kept me sane during the Ph.D journey. What made these moments truly special was that our discussions rarely touched academics. Instead, much of our time was spent pushing boundaries with playful teasing—of each other and ourselves. Thank you all for embracing the jokes and satires with such good humor. I feel incredibly lucky to have shared these moments with such amazing seniors, peers, and juniors, and I will treasure these memories forever. This list is far from exhaustive, and I sincerely apologize if I’ve missed anyone unintentionally. However, I make no apologies for the jokes I cracked—they were absolutely worth it!

Now, I would like to acknowledge the individuals who have been an essential part of my life since childhood. Above all, I am grateful to my friend Indranil. For the past 25 years, we have been the closest of companions, sharing almost every experience together. Thank you for being someone who has never judged my actions, even when I sometimes felt I should have been. I thank you for your constant support and for always standing by me, even in the most challenging times. As unprecedented and unexpected as the current health situation with your mother is, I pray for her steady recovery and wish your father continued good health. Special thanks to Deep for being there listening to my problems whenever I had one. A toast to my movie partner Tata for cooking delicious meals and sharing delightful moments enjoying good food and watching classic films together. I would also like to thank Aniket and Ankita didi for the wonderful memories since our school days. In conclusion, I would like to thank Pritam Sir for being the only person with zero self-interests who keeps in touch with me on a regular basis to know about my whereabouts and my well-being, and who I am sure is ready to jump in help if any dire situation arises.

I feel incredibly fortunate to have a supportive family who has always believed in me and been a constant source of encouragement. I consider myself blessed to have spent so much quality time with all my grandparents (Dada, Amma, Neno, and Dimma), an experience that feels truly rare. I lost my Dada in November 2023, and the moment of receiving my Ph.D. degree would have been even more fulfilling if he were here. However, I am certain he is watching over me from wherever he is. Despite living outside India, my uncle and aunt have always been my well-wishers, supporting me with thoughtful gestures and essential resources that have been pivotal in my educational journey. Much love to Adi and Tia from Babu dada! Thanks are also due to my maternal aunts and uncles for their constant love and affection. Lastly, thanks to my cousin brother, Shaan, and my cousin sister, June, for enduring my quirks and standing by me during tough times.

In a time where discovering a true soulmate feels like a daunting task, you have stood by me through every high and low over the past decade. Words fall short in capturing the depth of your presence in my life, and honestly, I don’t feel the need to try. Through all the chaos and challenges, my only wish is for us to remain together for the rest of our lives. Thank you, Shrena!

Words cannot convey the profound impact my parents have had on my Ph.D journey and my life as a whole. My mother played a pivotal role in shaping my education, mentoring me through my studies

from childhood to the completion of my schooling. Her guidance during the beginning years might perhaps be the reason why I am here today pursuing a Ph.D. If there is one person who inspires me everyday and quietly motivates me to keep moving forward, it is my father. His discipline and work ethic continue to stand out to this day. If I can adopt even a small fraction of his approach to work, I would consider myself successful. I will forever remain grateful for their unconditional support and trust—they have always stood by me in every way possible and above all, they have never questioned the choices I've made in life.

Sayan Mukherjee

TABLE OF CONTENTS

LIST OF FIGURES

AUTHOR LIST

- The following authors contributed to Chapter 2: Ratul Lahkar, Sayan Mukherjee, and Souvik Roy.
- The following authors contributed to Chapter 3: Ratul Lahkar, Sayan Mukherjee, and Souvik Roy.
- The following authors contributed to Chapter 4: Ratul Lahkar, Sayan Mukherjee, and Souvik Roy.
- The following authors contributed to Chapter 5: Sayan Mukherjee, and Souvik Roy.
- The following authors contributed to Chapter 6: Sayan Mukherjee, and Souvik Roy.

LIST OF PUBLICATION(S)/SUBMITTED ARTICLE(S)

- **Published Papers**

  1. Regularized Bayesian best response learning in finite games. **Mukherjee, S.**, Roy, S.; *Games and Economic Behavior, Volume 149, January 2025, Pages 1-31*; https://doi.org/10.1016/j.geb.2024.11.005.

  2. Perturbed Bayesian best response dynamic in continuum games. **Mukherjee, S.**, Roy, S.; *SIAM Journal on Control and Optimization, Vol 62, No. 6, pp. 3091-3120, 2024*; https://doi.org/10.1137/23M1591931.

  3. The logit dynamic in supermodular games with a continuum of strategies: A deterministic approximation approach. Lahkar, R., **Mukherjee, S.**, Roy, S.; *Games and Economic Behavior, Volume 139, May 2023, Pages 133-160*; https://doi.org/10.1016/j.geb.2023.02.003.

  4. Generalized perturbed best response dynamics with a continuum of strategies. Lahkar, R., **Mukherjee, S.**, Roy, S.; *Journal of Economic Theory, Volume 200, March 2022, 105398*; https://doi.org/10.1016/j.jet.2021.105398.

- **Submitted Articles**

  1. Equilibrium selection via stochastic evolution in continuum potential games. Lahkar, R., **Mukherjee, S.**, Roy, S.

CHAPTER

1 INTRODUCTION

Evolutionary game theory is a sub-field of game theory concerned with the behavior of large populations of strategically interacting players who recurrently revise their strategies based on the availability of better payoff opportunities. In such games, rather than directly focusing on the equilibrium analysis (as in the case of static/one-shot games), the strategic behavior of the population is modeled via an *evolutionary game dynamic* (a differential equation on a suitable state space of the population). Consequently, the theory aims to study the *asymptotic properties* (such as stability, convergence) of these evolutionary dynamics.

An important class of evolutionary dynamic studied in the literature is the class of *"perturbed best response dynamic."* In such a dynamic, the expected payoff of the population is perturbed using a “perturbation function" (in most cases, the Shannon entropy) and the population *“best responds"* to these perturbed expected payoffs. This thesis is primarily dedicated towards extending the literature of perturbed best response dynamics and deriving asymptotic stability and equilibrium selection results of the dynamic for various classes of games with a continuum of strategies and under incomplete information.

To be more precise, the theory of evolutionary games can broadly be classified into four categories:

- *homogeneous population (players are of the same type) finite strategy games*,
- *homogeneous population continuum strategy games*,
- *heterogeneous population (players are of different types) finite strategy games*, and
- *heterogeneous population continuum strategy games.*[1]

[1] Heterogeneous population games are often referred to as Bayesian population games in the literature.

As far as perturbed best response dynamic is concerned, the existing literature of evolutionary games has mostly been confined to the following cases where: the underlying population is homogeneous; the games are defined on finitely many strategies; and the perturbation function used is the Shannon entropy. Some works in this direction include Hofbauer (1995); Hofbauer and Sandholm (2002); Ely and Sandholm (2005); Hofbauer and Hopkins (2005); Hofbauer and Sorin (2006); Hofbauer and Sandholm (2007); Zusai (2023) to name a few. This leads us to the following set of questions:

1. Is it possible to extend the theory of perturbed best response dynamic to homogeneous population games with a *continuum of strategies*; and perhaps also for *arbitrary perturbation functions*?
2. Suppose that the underlying game has multiple equilibria. Is there a way to characterize *which equilibrium among the many is most likely to be selected* by the population in the long run?
3. Can we develop a theory of perturbed best response dynamic for *Bayesian population games* with *finitely many strategies* as well as for games with a *continuum of strategies*?

This thesis aims to solve the aforementioned problems in a series of five chapters. A brief description of the chapters is provided below.

1.1 Chapter 2: Generalized perturbed best response dynamics with a continuum of strategies

In this chapter, we consider a generalization of perturbed best response dynamics in population games with a continuum of strategies. The previous literature has considered the logit dynamic generated through the Shannon entropy as a deterministic perturbation. We consider a wider class of deterministic perturbations satisfying lower semicontinuity and strong convexity. Apart from the Shannon entropy, Tsallis entropy and Burg entropy are other perturbations that satisfy these conditions. We thereby generate the generalized perturbed best response dynamic with a continuum of strategies. We establish fundamental properties of the dynamic and show convergence in potential games and negative semidefinite games.

1.2 Chapter 3: The logit dynamic in supermodular games with a continuum of strategies: A deterministic approximation approach

In this chapter, we consider large population supermodular games with pairwise interaction and a continuous strategy set. Our objective is to establish the convergence of the logit dynamic in such games to logit equilibria. For this purpose, we apply the deterministic approximation approach, which interprets a deterministic dynamic as an approximation of a stochastic process. We first establish the closeness of this dynamic with a step–wise approximation. We then show that the logit stochastic process is close to the step–wise logit dynamic in a discrete approximation of the original game. Combining the two results, we obtain our deterministic approximation result. We then apply this result to supermodular games. Over finite but sufficiently long time horizons, the logit stochastic process converges to logit equilibria in a discrete approximation of the supermodular game. By the deterministic approximation approach, so does the logit dynamic in the continuum supermodular game.

1.3 Chapter 4: Equilibrium selection via stochastic evolution in continuum potential games

In this chapter, we consider stochastic evolution in large population potential games with a continuous strategy set. Examples include aggregative potential games and doubly symmetric games. We approximate the continuous strategy game with a finite strategy game. We then define a stochastic process based on a noisy exponential strategy revision protocol in the finite strategy game. Existing results imply the existence of a stationary distribution of the process on the finite strategy game. We then characterize this stationary distribution as the number of strategies go to infinity and the noise level goes to zero. The resulting distribution puts all its mass on the Nash equilibrium that maximizes the potential function of the continuous strategy game. This provides us with an exact history–independent equilibrium selection result in our large population potential game with a continuum of strategies.

1.4 Chapter 5: Regularized Bayesian best response learning in finite games

In this chapter, we introduce the notion of regularized Bayesian best response (RBBR) learning dynamic in heterogeneous population games. We obtain such a dynamic via perturbation by an arbitrary lower semicontinuous, strongly convex regularizer in Bayesian population games with finitely many strategies. We provide a sufficient condition for the existence of rest points of the RBBR learning dynamic, and hence the existence of regularized Bayesian equilibrium in Bayesian population games. These equilibria are shown to approximate the Bayesian equilibria of the game for vanishingly small regularizations. We also explore the fundamental properties of the RBBR learning dynamic, which includes the existence of unique solutions from arbitrary initial conditions, as well as the continuity of the solution trajectories thus obtained with respect to the initial conditions. Finally, as applications to the above theory, we introduce the notions of Bayesian potential and Bayesian negative semidefinite games and provide convergence results for such games.

1.5 Chapter 6: Perturbed Bayesian best response dynamics in continuum games

In this chapter, the notion of perturbed Bayesian best response dynamic for continuum strategy Bayesian population games is introduced. Fundamental properties of the dynamic such as existence of perturbed equilibrium, convergence of the perturbed equilibrium to the Bayesian equilibrium of the underlying game, as well as existence, uniqueness and continuity of solutions from arbitrary initial conditions is established. As applications to the theory, convergence of solutions to the perturbed equilibria is shown to hold for two classes of games, namely, Bayesian potential games and Bayesian negative semidefinite games.

In conclusion, Chapter 2 and Chapter 3 address Question 1; Chapter 4 addresses Question 2; and Chapter 5 and Chapter 6 address Question 3.

CHAPTER

2

GENERALIZED PERTURBED BEST RESPONSE DYNAMICS WITH A CONTINUUM OF STRATEGIES

2.1 Introduction

Perturbed best response dynamics are one of the most important classes of dynamics in evolutionary game theory. Best response is, of course, the fundamental behavioral norm in game theory. But due to multiplicity of best responses and the fact that the best response can change abruptly, the best response dynamic arising from this behavioral protocol is technically difficult to analyze (Gilboa and Matsui (1991), Hofbauer (1995)). Perturbed best response dynamics avoid this problem by considering a uniquely defined smoothed version of the best response which can then be analyzed using the standard theory of differential equations. Even if best responses are uniquely defined, perturbed best responses are important in their own right because there may be errors in agents' perception of payoffs or in implementing the best response. In such situations, perturbed best response may be a better model of behavior than the classical best response. For this reason, the notion of perturbed best response has also found applications in rationalizing data in experimental economics (Goeree et al. (2016)).

The logit dynamic (Fudenberg and Levine (1998), Chapter 4) is the prototypical representative of this class of dynamics obtained by perturbing the best response through the Shannon entropy function. A number of authors have since generalized the logit dynamic to the class of perturbed best response dynamics (Benaım and Hirsch (1999), Hofbauer and Hopkins (2005), Hofbauer and Sandholm (2002, 2007)). As with most of evolutionary game theory, most of this literature on perturbed best response dynamics has been in the context of large population games with finite strategy sets.[1] However, similar to classical game theory, various economic applications of evolutionary game theory are also more conveniently done with games with a continuum of strategies. Despite introducing certain measure theoretic complications related to, for example, defining the state space, continuous strategy models greatly simplify the characterization of Nash equilibrium and Pareto optimal states.[2] Therefore, to make evolutionary game theory a more relevant technique for economic analysis, there is a need to develop this theory in the context of games with continuous strategy sets. This has been the main motivation behind the literature on extending the prominent evolutionary game dynamics from finite strategy to continuous strategy games. Papers in this area include Oechssler and Riedel (2001, 2002) and Cheung (2016) on the replicator dynamic, Hofbauer et al. (2009) on the Brown–von Neumann–Nash (BNN) dynamic, Cheung (2014) on the pairwise comparison dynamic, and Perkins and Leslie (2014) and Lahkar and Riedel (2015) on the logit dynamic.

The present paper is a contribution to this literature on continuous strategy evolutionary dynamics. Specifically, it focuses on the class of perturbed best response dynamics for such games. The earlier literature on this topic has been confined to the logit dynamic (Perkins and Leslie (2014), Lahkar and Riedel (2015)). The logit dynamic is generated when agents' payoffs are perturbed using the Shannon entropy and agents best respond with respect to these perturbed payoffs. Here, we generalize the logit dynamic by considering perturbations other than the Shannon entropy. Thus, instead of assuming any particular form of perturbation, we consider a broader class of perturbations that satisfy certain lower semicontinuity and strong convexity conditions. For example, apart from the Shannon entropy, the Tsallis entropy and the Burg entropy are other admissible entropies under these conditions. As far as we know, this is the first paper that considers such a generalization of perturbed best response dynamics in continuous strategy games. We do note, though, that our generalisation is entirely with respect to deterministic perturbations unlike in the finite strategy literature where both deterministic and stochastic perturbations have been considered (Hofbauer and Sandholm (2002)).[3]

Similar to the logit best response, we first perturb an agent's payoff using the general perturbation function and allow the agent to best respond. This generates the generalized perturbed best response (GPBR) and the associated dynamic. We establish the fundamental properties of the dynamic. Thus, we show that rest points of the dynamic, which are fixed points of the perturbed best response function and which we call perturbed equilibria, exist and converge to Nash equilibria of the underlying game when the perturbation becomes small. We also show that unique solution trajectories that are continuous

[1]See Sandholm (2010) for an extensive review of the literature on finite strategy evolutionary game theory. In addition to perturbed best response dynamics, there have been other extensions of the finite dimensional best response dynamic. For example, Balkenborg et al. (2013) consider *refined best reply dynamics*. Suppose there is a continuum of Nash equilibria but, due to the multiplicity of best responses, no Nash equilibrium is Lyapunov stable under the classical best response dynamic. The refined best reply dynamic can then be used to analyze the evolutionary stability of certain faces of the simplex of mixed strategies.

[2]See, for example, Cheung and Lahkar (2018), Lahkar (2020) and Lahkar and Mukherjee (2019, 2021) for such applications to continuum strategy economic models like Cournot competition and implementation of efficiency in problems like public goods, public bads and the tragedy of the commons.

[3]As far as we know, stochastic perturbations have not been studied in the context of continuous strategy games.

with respect to initial states exist for this class of dynamics. Results of a similar nature have, of course, already been established for the logit dynamic (Lahkar and Riedel (2015)). But by extending these results beyond the logit dyanamic, this paper establishes the idea of perturbed best response on sounder foundations in continuous strategy games.

This is important because one of main motivations behind looking for broader and more diverse classes of evolutionary dynamics is to ensure that the results of evolutionary game theory are not dependent upon the specific properties of a particular dynamic. Otherwise, those results will not be robust. Our findings are in line with this general quest in evolutionary game theory. Conclusions obtained from the application of the logit dynamic in continuous strategy games now become more credible since such results now generalize to the wider class of perturbed best response dynamics. This applies, for example, immediately to the results on evolutionary implementation in Lahkar and Mukherjee (2019) where the logit dynamic has been applied. These results now also hold for the entire class of deterministic perturbed best response dynamics.

Another ubiquitous pursuit in evolutionary game theory has been classes of games in which a wide variety of evolutionary dynamics converge to Nash equilibria. Nash equilibrium prediction in such games then becomes more credible even when agents are myopic as is the case in evolutionary game theory. Two such classes of games in which almost all well known evolutionary dynamics, both finite strategy and continuous strategy, do converge are potential games (Monderer and Shapley (1996), Sandholm (2001)) and negative semidefinite games (Hofbauer and Sandholm (2009)).[4] Potential games are defined by a real valued function which summarises payoffs in the game. Negative definite games are characterized by the property of "self–defeating externalities" due to which small groups of agents changing strategy find themselves at a relative payoff disadvantage. We show that our GPBR dynamic also converges in these two classes of games when their strategy sets are continuous. Convergence in this case is, of course, not to Nash equilibria but to their approximation, perturbed equilibria. Once again, such convergence have already been established for the continuous strategy logit dynamic. The present results show that such convergence arises from fundamental properties of perturbations like lower semicontinuity and convexity. This generalization is also in line with similar conclusions that arise in the case of finite strategy perturbed best response dynamics (Hofbauer and Sandholm (2007)).

Analyzing evolutionary dynamics for continuous strategy dynamics raises certain significant mathematical complications which makes the exercise a non–trivial extension of finite strategy dynamics. For example, an important question to resolve is the choice of topology in the space of population states, which takes a measure theoretic form. Following much of the literature (for example, Oechssler and Riedel (2001, 2002), Lahkar and Riedel (2015)) we apply the strong and weak topologies in our analysis. The weak topology is helpful in resolving questions like the existence of a perturbed equilibrium and convergence of solution trajectories of the perturbed best response dynamics to perturbed equilibria in

[4] Monderer and Shapley (1996) originally introduced the notion of potential games in finite player games. Sandholm (2001) extended the concept to large population games with finite strategies. Further generalizations to large population games with continuous strategy sets can be found in, for example, Cheung (2014), Lahkar and Riedel (2015) and Cheung and Lahkar (2018). Negative semidefinite games are also called stable games (Hofbauer and Sandholm (2009)) and they generalize the fundamental idea of an evolutionary stable state (Maynard Smith (1982)). Continuous strategy extensions of such games have been considered in Hofbauer et al. (2009), Cheung (2014) and Lahkar and Riedel (2015).

potential and negative definite games. The strong topology, on the other hand, is required for results like the existence of solution trajectories of the dynamic.

Nor is the present analysis a straightforward extension of earlier papers on the continuous strategy logit dynamic. The absence of any specific functional form for the perturbation raises certain technical challenges that requires us make appropriate assumptions to make the analysis tractable. For example, we need assumptions like lower semicontinuity, strong convexity and the existence of a continuous density function for the perturbation to establish the existence of a perturbed equilibria. Showing convergence of perturbed equilibria to Nash equilibria requires us to impose some restrictions (Conditions A1 and A2) on the perturbation function to ensure that a small change in the population state does not lead to a large change in the perturbation. Such assumptions also ensure that, like the logit dynamic, our general class of perturbed best response dynamics are forward invariant not just in the state space, but in the space of states with bounded density functions. Such results then allow us to construct appropriate Lyapunov functions that establish convergence in the classes of potential and negative semidefinite games.

Beyond its theoretical merit, is there any importance of our exercise from the point of view of applications? To answer this question, we consider a large population model of Cournot competition with a linear cost function. Such a game is a potential game. Moreover, due to the linearity of the cost function, the game has a convex set of Nash equilibria. Hence, it is possible that GPBR dynamics generated by the Shannon, Tsallis and Burg entropies converge to perturbed equilibria that, as the perturbation becomes small, approximate different Nash equilibria. We present certain simulations that suggest that is indeed the case. Therefore, the pattern of social behavior can depend upon the manner in which best responses are perturbed, particularly in situations where there is a continuum of Nash equilibria. Hence, that our generalization of perturbed best response dynamics can have meaningful consequences for applications.

The rest of the paper is as follows. Section 6.2 defines preliminaries of the model. In Section 2.3, we analyze the perturbed best response and perturbed equilibria. This section also provides examples of perturbations beyond the Shannon entropy to which our analysis applies; namely, the Tsallis entropy and the Burg entropy. Section 2.4 establishes fundamental properties of the perturbed best response dynamic. In Sections 6.4 and 6.5, we show convergence in potential games and negative definite games respectively. Section 2.7 presents the application to the Cournot model and Section 6.6 concludes.

2.2 Preliminaries

We consider a population consisting of a continuum of agents of mass 1. Let S be a compact and convex subset of a finite dimensional Euclidean space with unit Lebesgue measure. We denote the collection of all probability measures on S by $\mathscr{P}(S)$ and the collection of all signed measures on S by $\mathscr{M}(S)$. Let $\mathscr{M}_0(S)$ be the collection of all signed measures on S whose measure of the whole space S is 0, that is, $\mathscr{M}_0(S) := \{\mu \in \mathscr{M}(S) : \mu(S) = 0\}$. We also use the notation $\mathscr{P}_0(S)$ to describe the set of all probability measures on S that admit a probability density function or, equivalently, is absolutely continuous with respect to the Lebesgue measure. Finally, let $\mathscr{P}_0^c(S)$ be the set of probability measures which admits a continuous probability density function on S. We will denote the density of a generic probability measure $\mathbb{P}$ in $\mathscr{P}_0(S)$ by $p(\cdot)$.

A probability measure $\mathbb{P} \in \mathscr{P}(S)$ denotes the state of the population. This means that for any Borel set $A \in \mathscr{B}(S)$, $\mathbb{P}(A)$ denotes the proportion of the population using strategies in A. Let $\mathscr{L}^\infty(S)$ be the set of bounded measurable functions on S endowed with the sup-norm[5] metric. We define a *population game* as a map $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ such that $\pi_x(\mathbb{P})$ is the payoff of an agent who plays strategy x at population state $\mathbb{P}$. Throughout the paper, we assume that the population game π is bounded in the sup-norm and weakly continuous with respect to the Prohorov metric.[6] A population state $\mathbb{P}^\circ$ is said to be a Nash equilibrium of the underlying population game π if for all $x \in \mathrm{Supp}(\mathbb{P}^\circ)$ and all $y \in S$, we have $\pi_x(\mathbb{P}^\circ) \geq \pi_y(\mathbb{P}^\circ)$.[7]

We now seek to define the generalized deterministic perturbed best response dynamic. This in turn requires us to introduce the notions of a perturbed expected payoff and the perturbed best response. Suppose the current population state is $\mathbb{P} \in \mathscr{P}(S)$. We then define the expected payoff at $\mathbb{P}$ with respect to another probability measure $\mathbb{Q} \in \mathscr{P}(S)$ as $\mathbb{E}^{\mathbb{Q}}[\pi(\mathbb{P})] := \int_S \pi(\mathbb{P}) d\mathbb{Q}$. We further define a deterministic perturbation function as a map $v : \mathscr{P}(S) \to \mathbb{R} \cup \{\infty\}$ with the restriction that $v(\mathbb{Q}) = \infty$ if and only if $\mathbb{Q} \notin \mathscr{P}_0(S)$. Let $\mathscr{L}^1(S)$ denote the collection of Lebesgue integrable functions on S endowed with the $\|\,.\,\|_{\mathscr{L}^1(S)}$[8] norm. Also, define the following subset of $\mathscr{L}^1(S)$:

$$\mathscr{L}_+^{1,1}(S) := \left\{ f \in \mathscr{L}^1(S) \,|\, f(x) \geq 0 \text{ for all } x \in S \text{ and } \int_S f(x) dx = 1 \right\}. \tag{2.1}$$

It is a well-known fact that for any $f \in \mathscr{L}_+^{1,1}(S)$, there exists a *unique* probability measure $\mathbb{Q}_f \in \mathscr{P}_0(S)$ such that the probability density function of $\mathbb{Q}_f$ is f. We define the *counterpart* of v as the mapping $\tilde{v} : \mathscr{L}^1(S) \to \mathbb{R} \cup \{\infty\}$ such that

$$\tilde{v}(f) = \begin{cases} v(\mathbb{Q}_f) & \text{if } f \in \mathscr{L}_+^{1,1}(S) \\ \infty & \text{otherwise.} \end{cases} \tag{2.2}$$

The fact that $\mathbb{Q}_f$ is unique for every $f \in \mathscr{L}_+^{1,1}(S)$ makes $\tilde{v}$ well-defined. Throughout, we assume that $\tilde{v}$ is lower semicontinuous[9] and strongly convex[10] with respect to the $\|\,.\,\|_{\mathscr{L}^1(S)}$ norm. One interpretation of

[5]The sup-norm of $f \in \mathscr{L}^\infty(S)$ is defined as $\|f\|_\infty := \sup_{x \in S} |f(x)|$.

[6]See (2.9) for the definition of Prohorov metric.

[7]We denote the support of a probability measure $\mathbb{P}$ by $\mathrm{Supp}(\mathbb{P})$.

[8]The $\|\,.\,\|_{\mathscr{L}^1(S)}$ norm of $f \in \mathscr{L}^1(S)$ is $\|f\|_{\mathscr{L}^1(S)} := \int_S |f(x)| dx$.

[9]A mapping $h : \mathscr{L}^1(S) \to \mathbb{R} \cup \{\infty\}$ is said to be lower semicontinuous with respect to the $\|\,.\,\|_{\mathscr{L}^1(S)}$ norm if for all sequences $(f_n)_{n \geq 1}$ in $\mathscr{L}^1(S)$, $\|f_n - f\|_{\mathscr{L}^1(S)} \to 0$ implies

$$\liminf_{n \to \infty} h(f_n) \geq h(f).$$

[10]An extended function T on $(\mathscr{L}^1(S), \|\,.\,\|_{\mathscr{L}^1(S)})$ to $\mathbb{R} \cup \{\infty, -\infty\}$ is said to be strongly convex if there exists $L > 0$ such that

$$T(\alpha f + (1-\alpha) g) \leq \alpha T(f) + (1-\alpha) T(g) - \frac{L\alpha(1-\alpha)}{2} \|f - g\|^2_{\mathscr{L}^1(S)}$$

for all $f, g \in \mathrm{dom}\, T$ and $\alpha \in [0,1]$ where $\mathrm{dom}\, T := \{f \in \mathscr{L}^1(S) : T(f) < \infty\}$.

$\mathbb{Q}$ is that it is a mixed strategy and $\mathbb{E}^{\mathbb{Q}}[\pi(\mathbb{P})]$ is the expected payoff of an agent who plays strategy $\mathbb{Q}$ at the population state $\mathbb{P}$. The perturbed expected payoff for the agent with respect to v is then $\mathbb{E}^{\mathbb{Q}}[\pi(\mathbb{P})]-v(\mathbb{Q})$. Given $\eta>0$, we define the generalized perturbed best response map $\mathscr{G}_\eta:\mathscr{P}(S)\to\mathscr{P}(S)$ as

$$\mathscr{G}_\eta(\mathbb{P})=\left[\arg\max_{\mathbb{Q}\in\mathscr{P}(S)}\left(\mathbb{E}^{\mathbb{Q}}[\pi(\mathbb{P})]-\eta v(\mathbb{Q})\right)\right]. \tag{2.3}$$

It now follows that the extended functional $\mathbb{P}\mapsto\left(\mathbb{E}^{\mathbb{Q}}[\pi(\mathbb{P})]-\eta v(\mathbb{Q})\right)$ has a unique maximizer (see Héliou et al. (2020), Pg. 12, Lemma B.1). Therefore, $\mathscr{G}_\eta$ is well-defined. Furthermore, due to the restriction that $v(\mathbb{Q})=\infty$ if $\mathbb{Q}\notin\mathscr{P}_0(S)$, it must be that $\mathscr{G}_\eta(\mathbb{P})\in\mathscr{P}_0(S)$ for all $\mathbb{P}\in\mathscr{P}(S)$.

We refer to $\mathscr{G}_\eta(\mathbb{P})$ as the generalized perturbed best response (GPBR) to population state $\mathbb{P}$ with respect to the perturbation function v. Intuitively, it is a perturbation of the more usual game theoretic notion of the best response, with the extent of perturbation declining as the perturbation factor $\eta\to 0$. But it is analytically more tractable than the best response because for every $\mathbb{P}\in\mathscr{P}(S)$, it is uniquely defined. The best response, on the other hand, may not even exist at certain state in games with continuum strategy sets or, even if it exists, may be multi–valued. Using the generalized perturbed best response, we can define the generalized perturbed best response (GPBR) dynamic as

$$\dot{\mathbb{P}}(A)=\mathscr{G}_\eta(\mathbb{P})(A)-\mathbb{P}(A) \tag{2.4}$$

for $A\in\mathscr{B}(S)$. Thus, under this dynamic, a population state $\mathbb{P}$ moves in the direction of its generalized perturbed best response $\mathscr{G}_\eta(\mathbb{P})$. Since both $\mathscr{G}_\eta(\mathbb{P}),\mathbb{P}\in\mathscr{P}(S)$, it must be that $\dot{\mathbb{P}}\in\mathscr{M}_0(S)$.

The reason behind calling $\mathscr{G}_\eta(\mathbb{P})$ as the generalized perturbed best response is that its construction arises as a generalization of the logit best response for games with continuous strategy sets (Mattsson and Weibull (2002)). They consider the perturbation function

$$v(\mathbb{Q})=\begin{cases}\int_S q(x)\log q(x)dx & \text{if } \mathbb{Q}\in\mathscr{P}_0(S)\\ \infty & \text{otherwise,}\end{cases} \tag{2.5}$$

which is, of course, the Shannon entropy of $\mathbb{Q}\in\mathscr{P}(S)$. Maximizing the resulting perturbed expected payoff as in (2.3), we obtain the logit best response (also called logit choice measure in Lahkar and Riedel (2015)) $\mathscr{G}_\eta=\mathscr{L}_\eta$ (Mattsson and Weibull (2002)) defined as

$$\mathscr{L}_\eta(\mathbb{P})(A)=\frac{\int_A\exp(\eta^{-1}\pi_y(\mathbb{P}))dy}{\int_S\exp(\eta^{-1}\pi_x(\mathbb{P}))dx} \tag{2.6}$$

for all $A\in\mathscr{B}(S)$. The logit best response is, therefore, a particular form of the generalized perturbed best response $\mathscr{G}_\eta$, which does not assume any structural form of v except that $\tilde{v}$ is lower semicontinuous and strongly convex with respect to the $\|.\|_{\mathscr{L}^1(S)}$ norm and that $v(\mathbb{Q})=\infty$ for all $\mathbb{Q}\notin\mathscr{P}_0(S)$.[11] Applying $\mathscr{L}_\eta(\mathbb{P})$ defined by (2.6) to the general form of the perturbed best response dynamic (2.4) then gives us

[11] The counterpart of Shannon entropy also satisfies lower semicontinuity and strong convexity as we will show later.

the continuous strategy logit dynamic (Perkins and Leslie (2014), Lahkar and Riedel (2015))

$$\dot{\mathbb{P}}(A)=\mathscr{L}_\eta(\mathbb{P})(A)-\mathbb{P}(A)=\frac{\int_A \exp(\eta^{-1}\pi_y(\mathbb{P}))dy}{\int_S \exp(\eta^{-1}\pi_x(\mathbb{P}))dx}-\mathbb{P}(A). \tag{2.7}$$

Recall that $\mathscr{G}_\eta(\mathbb{P})\in\mathscr{P}_0(S)$ for every $\mathbb{P}\in\mathscr{P}(S)$ and, therefore, always has a density function. In addition, if the existing population state $\mathbb{P}$ also has a density function, then (2.4) implies that $\dot{\mathbb{P}}\in\mathscr{M}_0(S)$ also a density function. A rest point of the dynamic (2.4) is a state $\mathbb{P}$ such that $\dot{\mathbb{P}}(A)=0$ for all $A\in\mathscr{B}(S)$. It is evident that any such rest point satisfies $\mathscr{G}_\eta(\mathbb{P})=\mathbb{P}$ and is, therefore, also a fixed point of the perturbed best response map $\mathscr{G}_\eta$. We call such a fixed point of $\mathscr{G}_\eta$ or, equivalently, a rest point of (2.4) a *perturbed equilibrium*. In the case of the logit dynamic, such an equilibrium is called a logit equilibrium.

2.3 Perturbed Best Response and Perturbed Equilibria

We now establish certain fundamental properties of the GPBR dynamic like the existence of a perturbed equilibrium and existence of solutions to the dynamic. To do so, we first define certain topological concepts that are standard in the analysis of continuous strategy evolutionary dynamics.

Two topologies arise in the context of such dynamics. One is the strong topology induced by the variational norm on $\mathscr{M}(S)$ given by $\|\mu\|_{\mathbf{TV}}=\sup_g|\int_S g\,d\mu|$ where g is a measurable function $g:S\to\mathbb{R}$ such that $\sup_{x\in S}|g(x)|\le 1$. If $\mathbb{P},\mathbb{Q}$ are two probability measures, then the distance between them under this norm is (see, e.g., Shiryaev (1995), Pg. 360)

$$\|\mathbb{P}-\mathbb{Q}\|_{\mathbf{TV}}:=2\sup_{A\in\mathscr{B}(S)}|\mathbb{P}(A)-\mathbb{Q}(A)|, \tag{2.8}$$

where $\|.\|_{\mathbf{TV}}:\mathscr{P}(S)\times\mathscr{P}(S)\to\mathbb{R}_+$ is called the total variation distance.

The other is the weak topology, which is induced by convergence in distribution. On way to metrize the weak topology on $\mathscr{P}(S)$ is through the Prohorov metric $\rho:\mathscr{P}(S)\times\mathscr{P}(S)\to\mathbb{R}_+$ defined by

$$\rho(\mathbb{P},\mathbb{Q}):=\inf\left\{\epsilon>0:\mathbb{P}(A)\le\mathbb{Q}(A^\epsilon)+\epsilon \text{ and } \mathbb{Q}(A)\le\mathbb{P}(A^\epsilon)+\epsilon \text{ for all } A\in\mathscr{B}(S)\right\} \tag{2.9}$$

for all $\mathbb{P},\mathbb{Q}\in\mathscr{P}(S)$ (see, e.g., Billingsley (2013), Pg. 72). Thus, a sequence of probability measures $(\mathbb{P}_n)_{n\ge1}\subseteq\mathscr{P}(S)$ converges weakly to $\mathbb{P}\in\mathscr{P}(S)$ if $\rho(\mathbb{P}_n,\mathbb{P})\to 0$. We will be using the strong topology to establish results on convergence under perturbed best response dynamics, which would imply convergence under the weak topology. The weak topology, on the other hand, is useful for establishing results like existence of equilibrium since it renders the state space $\mathscr{P}(S)$ compact.

We begin with the existence of a perturbed equilibrium. For that we will require few preliminary notations and definitions. Let $\mathscr{C}(S)$ denote the collection of continuous functions on S endowed with the sup-norm (see Footnote (5)). Recall that the perturbation function v is not continuous in the strong topology. Hence, it is not Fréchet differentiable, which is the more general notion of differentiablity in Banach spaces. Therefore we apply a weaker notion of differentiability called the Gâteaux derivative which is defined as follows. Let $\phi : \mathscr{M}(S) \to \mathbb{R}$. The Gateaux derivative $\mathrm{d}\phi(\mu) : \mathscr{M}_0(S) \to \mathbb{R}$ of ϕ at a point $\mu \in \mathscr{M}(S)$ towards the direction $\nu \in \mathscr{M}_0(S)$ is defined as

$$\mathrm{d}\phi(\mu)(\nu) := \lim_{\epsilon \to 0} \frac{\phi(\mu + \epsilon\nu) - \phi(\mu)}{\epsilon}$$

provided the limit exists.

If ϕ is Gâteaux differentiable at $\mu \in \mathscr{M}(S)$, we will assume that there exists a measurable map $\nabla^{\mathbf{G}}\phi(\mu)$ called *Gâteaux gradient of* ϕ *at* μ such that $\mathrm{d}\phi(\mu)(\nu) = \langle \nabla^{\mathbf{G}}\phi(\mu), \nu \rangle := \int_S \nabla^{\mathbf{G}}\phi_x(\mu)\nu(dx)$.

Our next theorem (Theorem 1) requires us to show that the mapping $\mathbb{P} \mapsto \mathscr{G}_\eta(\mathbb{P})$ is continuous in the weak toplogy on $\mathscr{P}(S)$. For the purpose of this theorem, we introduce the notion of admissible perturbations which ensure lower semicontinuity of the perturbation function on $\mathscr{P}_0^c(S)$ with respect to the weak topology.

Definition 1 (Admissible Perturbation) *Let $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ be a population game. We call a perturbation function v admissible if Range$(\pi) \subseteq \mathscr{C}(S)$ implies $\mathscr{G}_\eta(\mathbb{P}) \in \mathscr{P}_0^c(S)$ for all $\eta > 0$ and all $\mathbb{P} \in \mathscr{P}_0^c(S)$.*[12]

We now show that a perturbed equilibrium always exists subject to some restrictions on the perturbation function v. The following theorem states the relevant result. The proof follows from the Brouwer-Schauder-Tychonoff fixed point theorem (see Aliprantis and Border (2013), Pg. 583, Corollary 17.56).

Theorem 1 *Suppose $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ is population game such that Range$(\pi) \subseteq \mathscr{C}(S)$. Suppose that the perturbation function v is admissible and Gâteaux differentiable on $\mathscr{P}_0(S)$ such that the Gateaux gradient $\nabla^{\mathbf{G}}v(\mathbb{Q}) \in \mathscr{C}(S)$ for all $\mathbb{Q} \in \mathscr{P}_0^c(S)$. Then, for every $\eta > 0$, the GPBR map $\mathscr{G}_\eta : \mathscr{P}(S) \to \mathscr{P}(S)$ admits a fixed point, i.e. there exists $\mathbb{P}_\eta^\circ \in \mathscr{P}(S)$ such that $\mathscr{G}_\eta(\mathbb{P}_\eta^\circ) = \mathbb{P}_\eta^\circ$. Any such $\mathbb{P}_\eta^\circ$ is a perturbed equilibrium of π.*

The proof of Theorem 1 is in Appendix A.1. The proof requires us to show that the mapping $\mathscr{G}_\eta : \mathscr{P}(S) \to \mathscr{P}(S)$ is continuous in the weak topology (Lemma 9). We show this in two steps. Firstly, under the given conditions of Theorem 1, we show that the perturbation function is *weakly lower semicontinuous*[13] on $\mathscr{P}_0^c(S)$, or, equivalently lower semicontinuous on $\mathscr{P}_0^c(S)$ in the weak topology. This in turn calls for the use of the Argmax Continuous Mapping Theorem (Ferger (2004), Van Der Vaart and Wellner (1996), Pg. 286, Theorem 3.2.2). It states that if one has a sequence of stochastic processes indexed by some metric space and if the sample paths of these processes are almost surely upper semicontinuous with unique maximizers, then the sequence of maximizers of these processes converges in distribution to the unique maximizer of the limiting process. Once we establish this lemma, the result follows from the Brouwer-Schauder-Tychonoff fixed point theorem and the compactness of $\mathscr{P}(S)$. Lahkar and Riedel (2015) establish a similar result on the existence of logit equilibria, which are fixed points of the logit best response. The proof of Theorem 1 is, however, much more general as it does not assume any functional

[12]The fact that Shannon entropy (see (2.5)) is an admissible perturbation function follows from (2.6) and Definition 1.

[13]A mapping $h : (\mathscr{P}(S), \rho) \to \mathbb{R}$ is weakly lower semicontinuous if for all sequences $(\mathbb{P}_n)_{n\geq 1}$ in $\mathscr{P}(S)$, $\rho(\mathbb{P}_n, \mathbb{P}) \to 0$ implies $\liminf_{n\to\infty} h(\mathbb{P}_n) \geq h(\mathbb{P})$.

form of the perturbation function v. The only requirements that Lemma 9 imposes upon v are strong convexity and lower semicontinuity.

We now consider the relationship between a perturbed equilibrium of π and a Nash equilibrium of the game. Recall from (2.3) that the extent of perturbation in payoffs while calculating the perturbed best response is determined by the parameter η. It, therefore, seems plausible that lower is η, the closer is a perturbed best response to a best response and, therefore, closer is a perturbed equilibrium to a Nash equilibrium. Our next theorem verifies this intuition subject to certain conditions. We need to establish certain preliminaries before stating the theorem.

Let $\mathrm{GE}(\pi,\eta)$ be the set of perturbed equilibria of π with respect to the GPBR map $\mathscr{G}_\eta : \mathscr{P}(S) \to \mathscr{P}(S)$. That is, $\mathrm{GE}(\pi,\eta) := \{\mathbb{P}^\circ_\eta \in \mathscr{P}(S) : \mathscr{G}_\eta(\mathbb{P}^\circ_\eta) = \mathbb{P}^\circ_\eta\}$. We will require the notion of the convex conjugate of $\tilde{v}$. The general definition of the convex conjugate of a function is as follows.

Definition 2 (Convex conjugate of a function) *Let X be a real topological vector space and let X^* be the dual space of X. For a function $f : X \to \mathbb{R} \cup \{\infty, -\infty\}$, the convex conjugate $f^* : X^* \to \mathbb{R} \cup \{\infty, -\infty\}$ is defined by*

$$f^*(x^*) := \sup_{x \in X} (x^*(x) - f(x)).$$

Following Definition 2, we denote the convex conjugate of $\tilde{v}$ as $\tilde{v}^*$. We will also use the notation v_η to denote ηv.[14] In that case, $\tilde{v}^*_\eta$ is the convex conjugate of the perturbation counterpart $\tilde{v}_\eta$.

We will also require the notion of the Fréchet derivative, which is a generalization of the usual derivative to Banach spaces. Let V and W be normed linear spaces. An operator $\phi : V \to W$ is Fréchet differentiable at x if there exists a bounded linear operator $A : V \to W$ such that

$$\lim_{\|h\|_V \to 0} \frac{\|\phi(x+h) - \phi(x) - Ah\|_W}{\|h\|_V} = 0.$$

Thus, if A is the Fréchet derivative of ϕ, then $\phi(x+h) = \phi(x) + A(h) + \mathbf{o}(\|h\|_V)$ for all h near zero in V. In this paper, we consider Fréchet differentiability of functions $\phi : \mathscr{M}(S) \to \mathbb{R}$. We denote the Fréchet derivative of such a function by $\mathrm{D}\phi$.

If ϕ is Fréchet differentiable, we will assume that there exists a bounded measurable map $\nabla^{\mathbf{F}}\phi(\mu) \in \mathscr{L}^\infty(S)$ such that for all $\nu \in \mathscr{M}_0(S)$, we have

$$\mathrm{D}\phi(\mu)\nu := \int_S \nabla^{\mathbf{F}}\phi(\mu) d\nu = \int_S \nabla^{\mathbf{F}}\phi(\mu)(x)\nu(dx) = \langle \nabla^{\mathbf{F}}\phi(\mu), \nu \rangle. \tag{2.10}$$

We then call the map $\nabla^{\mathbf{F}}\phi$ the *Fréchet gradient* of ϕ.

With these notations, we now define the following two important conditions, Condition A1 and Condition A2, that will be required for establishing convergence of perturbed equilibria to Nash equilibrium. Intuitively, these conditions impose some restrictions on the perturbation function so that small changes in the population state does not lead to large changes in the perturbation. Recall that $\tilde{v}$

[14] Notice from, for example, (2.5) that the original perturbation function v is defined independent of η. This additional piece of notation is required to capture the effects as $\eta \to 0$.

(Definition 2.2) denotes the perturbation counterpart of v, and $\tilde{v}^*$ denotes the convex conjugate of $\tilde{v}$ (Definition 2).

Definition 3 (Condition A1) *A perturbation function v satisfies Condition A1 if for all $\eta > 0$, all disjoint $E, F \in \mathscr{B}(S)$ and all $g \in \mathscr{L}^{\infty}(S)$ with $\inf_{s\in F} g(s) > \sup_{t\in E} g(t)$, there exists a function $u : (0,\infty) \to \mathbb{R}$ with $\lim_{x\to 0} u(x) = \infty$ such that*

$$D\tilde{v}^*_{\eta}(g)(\mathbb{I}_F) \geq u(\eta) D\tilde{v}^*_{\eta}(g)(\mathbb{I}_E),$$

where $D\tilde{v}^*_{\eta}(g)$ is the Fréchet derivative of $\tilde{v}^*_{\eta}(g)$, and $\mathbb{I}_E, \mathbb{I}_F$ are indicator functions on E and F respectively.

Before defining Condition A2, we required another definition; that of the *shift* of a probability measure. It is defined as follows.

Definition 4 *Let $\mathbb{P} \in \mathscr{P}_0(S)$ and let $E, F \subseteq S$ be disjoint. We say that $\tilde{\mathbb{P}}$ is a shift of $\mathbb{P}$ from E to F if*

1. *$\tilde{\mathbb{P}}(E) < \mathbb{P}(E)$,*
2. *$\tilde{\mathbb{P}}(F) = \mathbb{P}(F) + (\mathbb{P}(E) - \tilde{\mathbb{P}}(E))$, and*
3. *$\tilde{\mathbb{P}}(G) = \mathbb{P}(G)$ for all $G \subseteq S \setminus (E \cup F)$.*

Condition A2 is now defined as follows.

Definition 5 (Condition A2) *We say that a perturbation function satisfies Condition A2 with respect to a sequence of probability measures $(\mathbb{P}_n)_{n\geq 1}$ if for all disjoint $E, F \subseteq S$ with $\lim_{n\to\infty} \mathbb{P}_n(E) \neq 0$ and $\inf_{z\in F} \pi_z(\mathbb{P}_n) > \sup_{w\in E} \pi_w(\mathbb{P}_n)$ for sufficiently large n, there exist $(A_n)_{n\geq 1} \subseteq E$ and $(\tilde{\mathbb{P}}_n)_{n\geq 1}$ where $\tilde{\mathbb{P}}_n$ is a shift of $\mathbb{P}_n$ from A_n to F, and a constant $\kappa > 0$ such that for sufficiently large n,*

$$|v(\mathbb{P}_n) - v(\tilde{\mathbb{P}}_n)| \leq \kappa(\mathbb{P}_n(A_n) - \tilde{\mathbb{P}}_n(A_n)). \tag{2.11}$$

Condition A2 is a mild continuity requirement that calls for the existence of a κ independent of n such that (2.11) is satisfied. The existence of such κ also implies that for small η, small changes in the population state cannot cause an abrupt change in the perturbation. The condition does look technical, but it is not very demanding. We show in Section 2.3.1 that the Tsallis and Burg entropies satisfy this condition. Moreover, it is also satisfied for all population games π that are weakly continuous and where Range $(\pi) \subseteq \mathscr{C}(S)$. With these preliminaries, we now state our desired theorem on convergence of perturbed equilibria to a Nash equilibria.

Theorem 2 *Suppose that the conditions of Theorem 1 are satisfied. Let $(\eta_n)_{n\geq 1}$ be any positive sequence such that $\eta_n \downarrow 0$ and let $(\mathbb{P}^{\circ}_n)_{n\geq 1}$ denote the corresponding sequence of perturbed equilibria. Suppose also that the perturbation function v satisfies Condition A1, or Condition A2 with respect to $(\mathbb{P}^{\circ}_n)_{n\geq 1}$, as stated in Definitions 3 and 5 respectively. Then the family $GE(\pi, \eta_n)_{n\geq 1}$ of perturbed equilibria has accumulation points under the weak topology and any such accumulation point is a Nash equilibrium of the original game.*

The proof of this theorem is in Appendix A.2. We provide some intuition as to how Condition A1 and Condition A2 are relevant for the proof. Since the space of signed measures is infinite dimensional, assuming lower semicontinuity and strong convexity under the strong topology does not guarantee continuity of the perturbation function. However, some type of continuity of the perturbation function is required to ensure the convergence of the perturbed equilibria to the Nash equilibrium of the underlying game. To ensure such weaker form of continuity of the perturbation function, we impose two conditions, namely A1 and A2. The weaker form of continuity we consider implies that when a small proportion of agents switch their strategies from one set to another, the change in the perturbation function will also be small. This guarantees that the change in the perturbation function caused by small shift of the agents' strategies will be under control, and as a consequence, the payoff from the shifted strategies can be made to dominate that change. Therefore, when perturbation decreases gradually, agents switch to exercising strategies in those subsets which have a greater payoff and thereby will weakly converge to a Nash equilibrium of the underlying game.

2.3.1 Examples of Perturbation Functions

We now provide some examples of perturbation functions $v : \mathscr{P}(S) \to \mathbb{R} \cup \{\infty\}$ that are admissible and that satisfy Conditions A1 or A2. Theorem 2 would then be valid for with respect to these perturbation functions.

Example 1 (Shannon entropy) *The first example is the Shannon entropy itself as defined in (2.5) and upon which the canonical logit dynamic is based.*

We now establish that Shannon entropy satisfies Condition A1 (Definition 3). Let $g \in \mathscr{L}^{\infty}(S)$ *and let* $E, F \in \mathscr{B}(S)$ *be disjoint such that* $\inf_{s\in F} g(s) > \sup_{t\in E} g(t)$. *Then*

$$\frac{D\tilde{v}_\eta^*(g)(\mathbb{I}_F)}{D\tilde{v}_\eta^*(g)(\mathbb{I}_E)} = \frac{\int_F \exp(\eta^{-1} g(x))dx}{\int_E \exp(\eta^{-1} g(y))dy} \geq \exp((\inf_{s\in F} g(s) - \sup_{t\in E} g(t))/\eta)\frac{\lambda(F)}{\lambda(E)}$$

where λ *denotes the Lebesgue measure. Therefore Shannon entropy satisfies Condition (A1) with* $u(\eta) = \exp((\inf_{s\in F} g(s) - \sup_{t\in E} g(t))/\eta)\frac{\lambda(F)}{\lambda(E)}$ *for all* $\eta > 0$.

We know from Lahkar and Riedel (2015) that under the Shannon entropy, a sequence of the relevant perturbed equilibria , which is the logit equilibria, converges to Nash equilibria as $\eta \to 0$. The exercise in Example 12 provides another demonstration, albeit indirect, of this fact. Such convergence is a consequence of the more general Condition A1. We now present two additional entropy functions which, as far as we know, has not been explored much in the literature on evolutionary game theory. These are the Tsallis (γ) entropy and the Burg entropy. In Appendix A.3, we show that these two entropy functions satisfy Condition A2 (Definition 5).

Example 2 (Tsallis(γ) entropy) *For $0<\gamma<1$, the Tsallis(γ) entropy is defined as*

$$v^{\boldsymbol{t}}(\mathbb{Q})=\begin{cases}\int_S \dfrac{q(x)-q(x)^\gamma}{\gamma(1-\gamma)}\,dx & \text{if } \mathbb{Q}\in\mathscr{P}_0(S)\\ \infty & \text{otherwise.}\end{cases}$$

We now show that Tsallis(γ) entropy is an admissible perturbation function. It follows from Héliou et al. (2020) that under perturbation by Tsallis(γ) entropy,

$$\mathscr{G}_\eta(\mathbb{P})(A)=\left[\frac{\eta}{1-\gamma}\right]^{1/(1-\gamma)}\int_A \frac{1}{(\theta_\eta^{\boldsymbol{t}}-\pi_z(\mathbb{P}))^{1/(1-\gamma)}}\,dz \tag{2.12}$$

for all $A\in\mathscr{B}(S)$, where θ_η is such that the RHS in the above equation is a valid probability density function on S. It follows from the above expression that Tsallis(γ) entropy is admissible.

It is well known in the literature (see Tsallis (1988)) that the Shannon entropy (Example 12) can be obtained from the Tsallis(γ) entropy (Example 13) by taking $\gamma\to 1$. However, the analytical properties of these entropies are quite different. Tsallis entropy is used in Héliou et al. (2020) in the context of online non-convex optimization problems.

Example 3 (Burg entropy) *The Burg entropy is defined as*

$$v^{\boldsymbol{b}}(\mathbb{Q})=\begin{cases}-\int_S \log q(x)\,dx & \text{if } \mathbb{Q}\in\mathscr{P}_0(S)\\ \infty & \text{otherwise.}\end{cases}$$

It also follows from Héliou et al. (2020) that under perturbation by the Burg entropy,

$$\mathscr{G}_\eta(\mathbb{P})(A)=\eta\int_A \frac{dx}{(\theta_\eta^{\boldsymbol{b}}-\pi_x(\mathbb{P}))}, \tag{2.13}$$

for all $A\in\mathscr{B}(S)$, where $\theta_\eta^{\boldsymbol{b}}$ is such that the integral if the above density is 1. The fact that Burg entropy is admissible then follows from arguments similar to the Tsallis entropy.

The definition of Burg entropy for games with continuum strategy spaces (Example 5.3.1) is motivated from the definition of Burg entropy for games on finite strategy spaces. Suppose that the strategy space $S=\{1,\ldots,n\}$ is finite and let $\Delta:=\{\mathbf{x}=(x_1,\ldots,x_n)\in[0,1]^n \mid x_1+\ldots+x_n=1\}$ be the collection of population states. In such a scenario, the Burg entropy is defined as $v^{\mathbf{b}}(\mathbf{x}):=-\sum_{i=1}^{n}\log x_i$, for all $\mathbf{x}\in\Delta$ (Hofbauer and Sandholm (2002)). Héliou et al. (2020) use the continuous version of the Burg entropy for online non-convex optimization problems.

The Shannon entropy has, of course, been widely used to model perturbations in both learning and evolutionary game theory.[15] In comparison, the Tsallis and Burg entropies are not so popular in game

[15] Fudenberg and Levine (1998), Chapter 4, applied the Shannon entropy to model stochastic fictitious play. The logit

theory. Nevertheless, they are also legitimate models for providing microfoundations to a perturbed best response. Like the Shannon entropy, we can also interpret the Tsallis and Burg entropies as "control costs" that become large when too much probability is put on some particular pure strategy. There is no *a priori* reason why such control costs can only take the form of Shannon entropy. In fact, if we wish to generalize the notion of control costs, then the Tsallis entropy is a natural candidate because, as mentioned earlier, the Shannon entropy is the limiting form of the Tsallis (γ) entropy as $\gamma \to 1$. As far as we know, the Burg entropy doesn't have any such relationship with the Shannon entropy. But there are precedents to its use in game theory. For example, Proposition 2.2 in Hofbauer and Sandholm (2002) consider the finite strategy version of the Burg entropy to define a perturbed best response. Harsanyi (1973) also applied such a logarithmic perturbation to establish oddness of the number of Nash equilibria. This earlier application was not in the contexts of learning or evolution but it does gives the Burg entropy some game theoretic pedigree.[16]

It is known that all the three entropies mentioned in this section (namely, Shannon, Tsallis, and Burg) satisfy strong convexity and lower semicontinuity (see Héliou et al. (2020), Pg. 7, Definition 1, for details).[17] We also show in Appendix A.3 that the Tsallis and Burg entropies satisfy Condition A2. Therefore, it is not just the Shannon entropy to which Theorem 2 applies. Instead, we can meaningfully define perturbed best responses for other entropies as well and approximate Nash equilibria with their fixed points. This makes our exercise of generalizing perturbed best response dynamics beyond the logit dynamic a substantial one.

2.4 Fundamental Properties of Perturbed Best Response Dynamics

We now seek to establish the fundamental properties of the GPBR dynamic (2.4). These are (i) the existence of a unique solution trajectory from every initial state in $\mathscr{P}(S)$ and (ii) the continuity of such solution trajectories with respect to initial state. We will require the following definitions for establishing the relevant result.

Definition 6 (Lipschitz population game) *A population game $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ is $\beta-$Lipschitz continuous if there exists a real number $\beta > 0$ such that*

$$\|\pi(\mathbb{P}) - \pi(\mathbb{Q})\|_\infty \leq \beta \|\mathbb{P} - \mathbb{Q}\|_{TV}$$

for all $\mathbb{P}, \mathbb{Q} \in \mathscr{P}(S)$, where $\|\,.\,\|_{TV}$ is the total variation distance as defined in (2.8).

Definition 7 (Semiflow of a dynamic) *Consider an abstract differential equation*

$$\Gamma'(t) = \mathscr{H}(\Gamma(t))$$

dynamic arose as the dynamics of mixed strategies in this model of learning. This dynamic was subsequently extended to evolutionary game theory where the state variable is the current population state rather than the mixed strategies of a finite group of players.

[16] Of course, (Harsanyi (1973)) did not use the term "Burg entropy". He only applied the functional form of this entropy.

[17] Héliou et al. (2020) show that the perturbation functions we consider, satisfy the property called regularizers, which, by definition is a combination of lower semicontinuity, strong convexity and some other properties.

on a Banach space $(X, \|.\|_X)$*. Suppose that a unique solution to the above differential equation exists for every initial condition* $\Gamma(0) = \gamma$ *(say). Denote the solution of the above differential equation with initial condition* γ *by* $\Gamma_\gamma(t)$*. Then the semiflow of the dynamic is the map* $\zeta : X \times [0, \infty) \to X$ *defined by* $\zeta(\gamma, t) = \Gamma_\gamma(t)$ *for all* $\gamma \in X$ *and* $t \in [0, \infty)$*.*

We will apply Definitions 6 and 7 in establishing the fundamental properties of the GPBR dynamic. Before that, we note from (2.4) that if indeed the GPBR dynamic admits a solution, then $\mathscr{P}(S)$ must be forward invariant under that dynamic. This means from any initial state in $\mathscr{P}(S)$, the solution trajectory must lie in $\mathscr{P}(S)$ for all times $t \geq 0$. This is immediate because from the definition of the perturbed best response (2.3), if $\mathbb{P} \in \mathscr{P}(S)$, then $\mathscr{G}_\eta(\mathbb{P}) \in \mathscr{P}(S)$.

While forward invariance of $\mathscr{P}(S)$ is obvious, we can actually establish a stronger result which will also be useful for later in establishing the fundamental properties of the dynamic. For this purpose, recall that $\mathscr{P}_0(S) \subset \mathscr{P}(S)$ is the set of absolutely continuous probability measures. Let $D > 1$. We then define the set

$$\mathscr{P}_D(S) := \Bigg\{ \mathbb{P} \in \mathscr{P}_0(S) : \frac{1}{D} \leq p(x) \leq D \text{ for } all\, x \in S$$

$$\text{and } |p(x) - p(y)| \leq D|x - y| \text{ for all } x, y \in S \Bigg\}. \tag{2.14}$$

Thus, $\mathscr{P}_D(S)$ is the set of absolutely continuous probability measures with D-Lipschitz continuous density functions and whose lower and upper bounds are given by $\frac{1}{D}$ and D respectively. It is known (see Perkins and Leslie (2014), Hanche-Olsen and Holden (2010)) that $\mathscr{P}_D(S)$ is a compact subset of $\mathscr{P}(S)$ under the total variation norm (2.8).

We now show that the GPBR dynamic is forward invariant not just in $\mathscr{P}(S)$ but also $\mathscr{P}_D(S)$. Thus, if $\mathbb{P} \in \mathscr{P}_D(S)$, then the GPBR $\mathscr{G}_\eta(\mathbb{P}) \in \mathscr{P}_D(S)$ as well.[18] Thus, if the GPBR dynamic admits a solution, then the entire trajectory originating in $\mathscr{P}_D(S)$ must lie in $\mathscr{P}_D(S)$ at all time $t \geq 0$. To arrive at this result, we will require a further piece of notation. Recall the set $\mathscr{L}_+^{1,1}(S) \subset \mathscr{L}^1(S)$ from (2.1). We now define the following subset of $\mathscr{L}_+^{1,1}(S)$ by imposing bounds $\frac{1}{D}$ and D on the functions in $\mathscr{L}_+^{1,1}(S)$ and which are D–Lispchitz continuous, where $D > 1$. Thus,

$$\mathscr{L}_{+,D}^{1,1}(S) := \Bigg\{ f \in \mathscr{L}_+^{1,1}(S) : \frac{1}{D} \leq f(x) \leq D \text{ for all } x \in S$$

$$\text{and } |f(x) - f(y)| \leq D|x - y| \text{ for all } x, y \in S \Bigg\}. \tag{2.15}$$

It is easy to note that ξ restricted to $\mathscr{L}_{+,D}^{1,1}(S)$ is a bijection onto $\mathscr{P}_D(S)$. The following lemma states the result on forward invariance. Recall that $\tilde{v}$ is the counterpart of v (see 2.2), $\tilde{v}^*$ is the convex conjugate of $\tilde{v}$ (see Definition 2) and $\nabla^{\mathbf{F}} \tilde{v}^*$ is the gradient. The proof is in Appendix A.4.

[18] Perkins and Leslie (2014) establish this result for the specific case of the logit best response. Our result is more general.

Lemma 1 *Let $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ and suppose that $\nabla^{\mathbf{F}}\tilde{v}^*(\mathbb{P}) \in \mathscr{L}^{1,1}_{+,D}(S)$ for all $\mathbb{P} \in \mathscr{P}_D(S)$. Consider the GPBR dynamic (2.4). If $\mathbb{P} \in \mathscr{P}_D(S)$, then $\mathscr{G}_\eta(\mathbb{P}) \in \mathscr{P}_D(S)$.*

Thus, Lemma 1 would imply that if indeed the generalized PBR dynamic admits a solution, then $\mathscr{P}_D(S)$ would be invariant. The following theorem shows that indeed such solutions exist. Moreover, any such solution must be unique and continuous with respect to initial states. This is the main result of this section. The proof, which follows from the Picard–Lindelöf theorem, is in Appendix A.4.

Theorem 3 *Suppose that $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ is a β-Lipschitz population game and v is a deterministic perturbation function. Suppose also that the counter part $\tilde{v}$ of v is closed*[19] *and L-strongly convex (see footnote 10). Then, given any initial condition $\mathbb{P}_0 \in \mathscr{P}(S)$, there exists a unique solution to the GPBR dynamic $\dot{\mathbb{P}} = \mathscr{G}_\eta(\mathbb{P}) - \mathbb{P}$. Furthermore, the semiflow of the dynamic is continuous in the initial condition $\mathbb{P}_0$ with respect to the total variation norm.*

2.5 Convergence in Potential Games

We now establish convergence of the GPBR dynamic to perturbed equilibria in potential games (Monderer and Shapley (1996), Cheung (2014), Lahkar and Riedel (2015)). In order to define such games, recall from (5.18) our assumption that a Fréchet differentiable function ϕ has a Fréchet gradient $\nabla^{\mathbf{F}}\phi$.

Definition 8 (Potential Games) *A population game $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ is said to be a potential game if there exists a Fréchet Differentiable functional $\varphi : \mathscr{M}(S) \to \mathbb{R}$ such that*

$$\nabla^{\mathbf{F}}\varphi(\mathbb{P}) = \pi(\mathbb{P}) \qquad \textit{for all } \mathbb{P} \in \mathscr{P}(S).$$

The functional φ is called the potential function of the game π.

Applying the Frèchet derivatinve to φ, we obtain

$$\mathrm{D}\varphi(\mathbb{P})(\nu) := \int_S \nabla^{\mathbf{F}}\varphi(\mathbb{P})d\nu = \int_S \nabla^{\mathbf{F}}\varphi(\mathbb{P})(x)\nu(dx) = \int_S \pi_x(\mathbb{P})\nu(dx) =: \langle \pi(\mathbb{P}), \nu \rangle.$$

In order to establish the desired convergence result, we introduce the *entropy adjusted potential function* defined as follows.

Definition 9 (Entropy adjusted potential function) *For a potential game $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$, we define the entropy adjusted potential function $\tilde{\varphi} : \mathscr{P}_0(S) \to \mathbb{R}$ as $\tilde{\varphi}(\mathbb{P}) := \varphi(\mathbb{P}) - v_\eta(\mathbb{P})$. For some $D > 1$, the restriction of $\tilde{\varphi}$ to $\mathscr{P}_D(S)$ is denoted by $\tilde{\varphi}_D$.*

Recall that since the perturbation function v is not Fréchet differentiable, we defined the Gateaux derivative which is a generalization of the directional derivative and which will allow us to consider

[19] We call a function $f : (X, \tau) \to (\mathbb{R}, |\,.\,|)$ closed if the set $\{(x,t) : x \in X, t \in \mathbb{R}, t \geq f(x)\}$ is a closed subset of $X \times \mathbb{R}$.

movements along only specific directions; for example, directions that have density functions that are bounded and bounded away from zero (as in $\mathscr{P}_D(S)$). Therefore, we apply the Gateaux derivative to analyze the entropy adjusted potential function.

We need an additional result for establishing convergence in potential games. We say that the GPBR dynamic satisfies *positive correlation* with respect to a function $\Gamma : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ if $\int_S \Gamma_x(\mathbb{P})\dot{\mathbb{P}}(dx) \geq 0$ for all $\mathbb{P} \in \mathscr{P}(S)$ with equality only if $\dot{\mathbb{P}} = 0$. Define the virtual payoff of strategy x at state $\mathbb{P}$ as

$$\tilde{\pi}_x(\mathbb{P}) = \pi_x(\mathbb{P}) - \nabla^{\mathbf{G}} v_\eta(\mathbb{P})(x). \tag{2.16}$$

The following lemma shows that the dynamic satisfies positive correlation with respect to the virtual payoff. In order to state the lemma, we impose a further condition, stated in (2.17), on our perturbation functions. Let $\mathscr{C}^1([0,\infty))$ be the collection of all continuously differentiable functions on $[0,\infty)$. Suppose that there exists a $\mathscr{C}^1([0,\infty))$ function $\Theta : [0,\infty) \to \mathbb{R}$ such that the perturbation function v can be expressed as

$$v(\mathbb{P}) = \begin{cases} \int_S \Theta \circ p(x) dx & \text{if } \mathbb{P} \in \mathscr{P}_0(S) \\ \infty & \text{otherwise.} \end{cases} \tag{2.17}$$

The three entropies we consider, namely the Shannon entropy, the Tsallis entropy and the Burg entropy, do satisfy (2.17). We then have the following lemma. The proof is in Appendix A.5.

Lemma 2 *Let $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ be a potential game. Suppose that the perturbation function v satisfies (2.17). Then the followings hold:*

1. *Suppose that Θ' is strictly increasing on $[0,\infty)$. Then for every $\eta > 0$, the generalized dynamic satisfies positive correlation with respect to the virtual payoff $\tilde{\pi}$ as defined in (5.19). Thus, $\int_S \tilde{\pi}_x(\mathbb{P})\dot{\mathbb{P}}(dx) \geq 0$ for all $\mathbb{P} \in \mathscr{P}_0(S)$ with equality only if $\dot{\mathbb{P}} = 0$.*

2. *Suppose that there exists $0 \leq \alpha < 1 < \beta < \infty$ such that Θ' is strictly increasing on $[\alpha, \beta)$. Then for every $\eta > 0$ such that $[D_\eta^{-1}, D_\eta] \subset [\alpha, \beta)$, the generalized dynamic satisfies positive correlation with respect to the virtual payoff $\tilde{\pi}$ as defined in (5.19). Thus, $\int_S \tilde{\pi}_x(\mathbb{P})\dot{\mathbb{P}}(dx) \geq 0$ for all $\mathbb{P} \in \mathscr{P}(S)$ with equality only if $\dot{\mathbb{P}} = 0$.*

Intuitively, this implies that there is a positive relationship between the virtual payoff of a strategy and it rate of change under the generalized PBR dynamic. It is a generalization of the corresponding result for the logit dynamic in Lahkar and Riedel (2015) to more generalized PBR dynamics. Lemma 2 lemma leads to the following theorem on convergence in potential games. The proof is in Appendix A.5.

Theorem 4 *Let $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ be a potential game with potential function φ which is Fréchet differentiable. Let v be a deterministic perturbation function such that $\tilde{v}$ is closed. Suppose further that the Fréchet derivative of the convex conjugate $\tilde{v}^*$ of $\tilde{v}$ is identified by some element in $\mathscr{L}^{1,1}_{+,D}(S)$ for some $D > 1$. Let (2.17) be satisfied so that Lemma 2 holds. Then the following hold.*

1. *$\tilde{\varphi}_D$ increases weakly along every solution trajectory of the GPBR dynamic (2.4) arising in $\mathscr{P}_D(S)$ and increases strictly across every non stationary solution trajectory.*

2. *The set of ω-limit points (in the strong topology on $\mathscr{P}_D(S)$) of any trajectory of the GBPR dynamic is a non-empty connected compact set of perturbed equilibria. Moreover, such limit points are local maximizers of the entropy-adjusted potential function in $\mathscr{P}_D(S)$.*

A significant feature of Theorem 4 is that solution trajectories of GPBR dynamics ascend the entropy adjusted potential function and, therefore, converge to a local maximizer of the function. Such local maximizers of the entropy adjusted potential function are the perturbed equilibria of the underlying potential game. This is a feature common to other standard evolutionary game dynamics like the replicator dynamic, the BNN dynamic and the pairwise comparison dynamic. The subtle distinction is that these dynamics locally optimize the potential function while the GPBR dynamics do so for the entropy adjusted potential function. Other evolutionary dynamics like adaptive dynamics also have similar features in certain circumstances (Hofbauer and Sigmund (1990)). For example, if the relevant payoffs are "frequency independent", then these dynamics also lead to local optimization.

We end this section with an example of continuous strategy potential games. Section 2.7 presents a more elaborate example that is of relevance for economics.

Example 4 *[**Two–Player Symmetric Game with Common Interests**] Let the strategy set be $S = [\underline{x}, \bar{x}] \subset \mathbb{R}$. Consider the population game π defined by the payoff function*

$$\pi_x(\mathbb{P}) = \int_{y \in S} h(x, y)\mathbb{P}(dy), \tag{2.18}$$

where $h : S \times S \to \mathbf{R}$ is a continuous function satisfying $h(x, y) = h(y, x)$. The interpretation here is that players are randomly matched in pairs to play a symmetric two–player game with $h(x, y)$ being the payoff of a player who plays strategy x against an opponent playing y. The payoff in (2.18) is then the expected payoff of a player playing x at the population state $\mathbb{P}$.

If we now impose the further restriction that h is doubly symmetric, i.e. $h(x, y) = h(y, x)$ for all $x, y \in S$, then the game becomes one of common interest. The two matched players always receive the same payoffs. An example is a quadratic game where

$$h(x, y) = -x^2 + axy - y^2 \tag{2.19}$$

for some $a > 0$, $\underline{x} < 0$ and $\bar{x} > 0$ (Hofbauer et al. (2009)). It is known that such doubly symmetric games are *potential games with potential function (see, for instance, Example 2, Cheung (2014))*

$$\begin{aligned}\varphi(\mathbb{P}) &= \frac{1}{2}\int_S\int_S h(x,y)\mathbb{P}(dy)\mathbb{P}(dx)\\ &= \frac{1}{2}\int_S \pi_x(\mathbb{P})(dx)\\ &= \frac{1}{2}\bar{\pi}(\mathbb{P}),\end{aligned} \tag{2.20}$$

where $\bar{\pi}(\mathbb{P}) = \int_S \pi_x(\mathbb{P})(dx)$ is the average payoff in the population at state $\mathbb{P}$. Thus, the potential function (2.20) of a doubly symmetric game is half the average payoff function. Theorem 4 implies that GPBR dynamics converge to perturbed equilibria in such doubly symmetric games.

2.6 Convergence in Negative Semidefinite Games

In this section, we establish convergence of the GPBR dynamic to perturbed equilibria in negative semidefinite games. We first recall the definition of such games from Lahkar and Riedel (2015).

Definition 10 *A population game $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ is said to be negative semidefinite if*

$$\int_S (\pi_x(\mathbb{Q}) - \pi_x(\mathbb{P}))(\mathbb{Q}-\mathbb{P})(dx) \leq 0 \tag{2.21}$$

for all $\mathbb{P}, \mathbb{Q} \in \mathscr{P}(S)$.

If the inequality in (2.21) holds strictly, then we call the game negative definite. Two examples of negative semidefinite games are linear quadratic symmetric normal form games with payoff function $\varphi(x,y) = -x^2 + axy$, for some constant $a \leq 0$ (Hofbauer et al. (2009)) and random matching in a two–player symmetric zero-sum game (Cheung (2014)).

Suppose that the population game $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ is C^1 in the sense of Fréchet differentiability with respect to the appropriate topology under consideration. Then it is well known (see Cheung (2014), Lahkar and Riedel (2015)) that π is negative semidefinite if and only if

$$\langle \mathrm{D}\pi(\mathbb{P})\mu, \mu\rangle \leq 0 \text{ for all } \mathbb{P} \in \mathscr{P}(S) \text{ and } \mu \in \mathscr{M}_0(S).$$

The following lemma shows that such games have a unique perturbed equilibrium. The proof, in Appendix A.6, is a generalization of a similar result in Lahkar and Riedel (2015) for the logit equilibrium. That result in Lahkar and Riedel (2015) is itself an extension of the result about uniqueness of perturbed equilibrium in finite strategy negative semidefinite games as established in Hofbauer and Sandholm (2007).

Lemma 3 *Let $\eta > 0$ and $D > 1$. Let $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ be a negative semidefinite game which C^1 in the sense of Frećhet differentiability.*[20] *Suppose that the perturbation function v be Gâteaux differentiable*

[20] For the purpose of this lemma, we extend the domain of the definition of π to $\mathscr{M}(S)$.

such that for every $\epsilon > 0$, $\langle d\nabla^{\mathbf{G}} v_\eta(\mathbb{P}+\epsilon\mu), \mu\rangle \geq 0$ for all $\mathbb{P} \in \mathscr{P}_0(S)$ and all absolutely continuous $\mu \in \mathscr{M}_0(S)$. Then the game π has a unique perturbed equilibrium.

The three entropies considered in this paper, Shannon, Tsallis, and Burg, satisfy the conditions in Lemma 3.[21] Therefore, in each of these cases, the associated perturbed best response leads to a unique perturbed equilibrium in negative semidefinite games.

It is known that a negative semidefinite game has a convex set of Nash equilibria (Hofbauer and Sandholm (2009)). In fact, if the game is negative definite, then there exists a unique Nash equilibrium. Lemma 3, however, shows that when it comes to perturbed equilibria, the game does not have to be negative definite for the equilibrium to be unique. Even if (2.21) is satisfied weakly so that the game is only negative semidefinite, we still have uniqueness of perturbed equilibrium.

We now establish convergence of the generalized perturbed best response dynamic to the unique perturbed equilibrium, as shown in Lemma 3, in negative semidefinite games. For this purpose, we adopt a similar approach as in the case of potential games and prescribe a Lyapunov function which satisfy certain monotonicity properties. Following Lahkar and Riedel (2015) we define the Lyapunov function for a negative semidefinite game π as[22]

$$\Lambda(\mathbb{P}) := \left(\int_S \pi_x(\mathbb{P})\mathscr{G}_\eta(\mathbb{P})(dx) - v_\eta(\mathscr{G}_\eta(\mathbb{P}))\right) - \left(\int_S \pi_x(\mathbb{P})\mathbb{P}(dx) - v_\eta(\mathbb{P})\right). \tag{2.22}$$

Since the perturbation function v is admissible (Definition 1), $\mathscr{G}_\eta(\mathbb{P}) \in \mathscr{P}_0(S)$ for all $\mathbb{P} \in \mathscr{P}_0(S)$. Also, since the perturbation function v is finite on $\mathscr{P}_0(S)$, we have Λ is well defined.

We show that the Lyapunov function defined in (2.22) decreases along every solution trajectory to the GPBR dynamic and decreases strictly along every non-stationary solution. This is stated in the following theorem, which is the main result of this section. The proof is in Appendix A.6.

Theorem 5 *Let $\pi : \mathscr{P}(S) \to \mathscr{L}^\infty(S)$ be a negative semidefinite population game that is C^1 in the sense of Fréchet differentiability. Consider the GBPR dynamic (2.4) and suppose the perturbation function satisfies the conditions of Lemma 3. Let $\Lambda_D : \mathscr{P}_D(S) \to R$ be the restriction of Λ (see (2.22)) to $\mathscr{P}_D(S)$. This function decreases monotonically to zero along every solution trajectory of the dynamic that originates in $\mathscr{P}_D(S)$, and decreases strictly along nonstationary solutions. Hence, every solution trajectory of the GPBR dynamic in $\mathscr{P}_D(S)$ converges, with respect to the strong topology on $\mathscr{P}_D(S)$, to the unique perturbed equilibrium of the game π.*

[21] Fix $\epsilon > 0$. Let $\mathbb{P} \in \mathscr{P}_0(S)$ and $\mu \in \mathscr{M}_0(S)$ with densities $p(\cdot)$ and $z(\cdot)$ respectively. It is known from Lahkar and Riedel (2015) that for Shannon entropy, $\nabla^{\mathbf{G}} v_\eta(\mathbb{P}+\epsilon\mu)(x) = \eta(1+\log(p(x)+\epsilon\mu(x)))$ for all $x \in S$. It then follows that $d\nabla^{\mathbf{G}} v_\eta(\mathbb{P}+\epsilon\mu)(x) = \dfrac{z(x)}{p(x)+\epsilon z(x)}$ for all $x \in S$. Therefore, $\langle d\nabla^{\mathbf{G}} v_\eta(\mathbb{P}+\epsilon\mu), \mu\rangle = \int_S \dfrac{z^2(x)}{p(x)+\epsilon z(x)} dx \geq 0$. For Tsallis and Burg entropy, $d\nabla^{\mathbf{G}} v_\eta(\mathbb{P}+\epsilon\mu)(x)$ is equal to $\dfrac{\eta\gamma(1-\gamma)}{(p(x)+\epsilon z(x))^{2-\gamma}}$ and $\dfrac{\eta}{(p(x)+\epsilon z(x))^2}$, respectively, for all $x \in S$. The result then follows for these two entropies through similar arguments.

[22] The Lyapunov function in Lahkar and Riedel (2015) is defined with respect to the logit dynamic and is an extension of the Lyapunov function in Hofbauer and Hopkins (2005) for finite strategy negative semidefinite games. The function in (2.22) is for the general class of perturbed best response dynamics.

Theorem 5 generalizes the result on convergence of the logit dynamic (Lahkar and Riedel (2015)) in negative semidefinite games to all continuous strategy perturbed best response dynamics considered in this paper. It is, therefore, the continuous strategy analog of the result obtained by Hofbauer and Sandholm (2007) for all perturbed best response dynamics in finite strategy negative semidefinite games.

2.7 An Application

We have generalized the class of perturbed best response dynamics with a continuum of strategies beyond the Shannon entropy and the logit dynamic that this entropy generates. But what is the relevance of this generalization? Does it lead to significantly different predictions about social behavior? In particular, we know that the limit of a sequence of perturbed equilibria under any such perturbed best response is a Nash equilibrium. Therefore, will different entropies lead to different limiting Nash equilibrium? If the answer is yes, then the generalization is indeed meaningful because that will mean that the nature of perturbation does determine our prediction of limiting behavior in a society.

To answer these questions, we consider a canonical model from economics–the model of Cournot competition extended to a large population. Since the analysis in this section will be based on simulations, we examine a numerical example of such a model. Let the population of agents be a set of firms of mass 1. The strategy set of each firm is $S=[0,1]$ where $x\in S$ represents the output that a firm can produce. Given the population state $\mathbb{P}$, denote by $A(\mathbb{P})=\int_S x\mathbb{P}(dx)\in[0,1]$ the aggregate output produced by the firms at that state. Let $\beta:[0,1]\to\mathbb{R}_+$ be the inverse demand function defined by $\beta(\alpha)=1-\alpha$ where $\alpha=A(\mathbb{P})$ is the prevailing aggregate output. Thus, given the population state $\mathbb{P}$ and the associated aggregate output $\alpha=A(\mathbb{P})$, the market price is $1-\alpha$. Further, let $c:S\to\mathbb{R}_+$ defined by $c(x)=0.2x$ be the cost function of a firm. As we will explain further below, it is important for our results that this cost function be linear and not strictly convex.

The payoff of a firm which produces output x when the population state is $\mathbb{P}$ is then

$$\begin{aligned}\pi_x(\mathbb{P})&=x\beta(A(\mathbb{P}))-c(x)\\&=x(1-A(\mathbb{P}))-0.2x.\end{aligned}\tag{2.23}$$

Thus, the payoff of the firm is the revenue it obtains from output x minus the cost of producing x. Formally, therefore, (2.23) is a model of Cournot competition with a large population of firms. Notice that the payoff in (2.23) depends entirely upon the firm's own strategy x and the aggregate strategy level $A(\mathbb{P})$ in the population. Hence, such games are aggregative games (Corchon (1994)).

We now characterize the set of Nash equilibrium of (2.23). Proposition 1 in Cheung and Lahkar

(2018) imply that the aggregative game (2.23) is a potential game with potential function

$$\begin{aligned}\varphi(\mathbb{P}) &= \int_0^{A(\mathbb{P})} \beta(z)dz - \int_S c(x)\mathbb{P}(dx)\\ &= \int_0^{A(\mathbb{P})} (1-z)dz - 0.2\int_S x\mathbb{P}(dx)\\ &= \int_0^{A(\mathbb{P})} (1-z)dz - 0.2A(\mathbb{P}). \end{aligned} \tag{2.24}$$

Moreover, this function is concave in $\mathscr{P}(S)$ (Lemma 1, Cheung and Lahkar (2018)). It is known that the set of Nash equilibria of potential games with a concave potential function coincide with the convex set of maximizers of the potential function (Sandholm (2001)). Clearly, any $\mathbb{P}$ with expected value (aggregate strategy) $A(\mathbb{P}) = 0.8$ is a maximizer of (2.24). Therefore, any such $\mathbb{P}$ is a Nash equilibrium of (2.23).[23] We summarize this discussion in the following proposition.

Proposition 1 *The large population model of Cournot competition π defined by (2.23) is a potential game with potential function (2.24). Further, a population state $\mathbb{P} \in \mathscr{P}(S)$ is a Nash equilibrium of π if and only if $A(\mathbb{P}) = 0.8$. Therefore, π has a convex set of Nash equilibria.*

Proof 1 *Follows from the preceding discussion.*

Since (2.23) is a potential game, we are assured of convergence of GPBR dynamics in π to a perturbed equilibrium (Theorem 4). Further, by Theorem 2, as $\eta \to 0$, any such perturbed must be a Nash equilibrium and, hence, by Proposition 1 satisfy $A(\mathbb{P}) = 0.8$. It is here that the convexity of the set of Nash equilibria becomes important. We, therefore, have a continuum of Nash equilibria. This creates the possibility that GPBR dynamics generated by different entropies will converge, as $\eta \to 0$, to different Nash equilibria. As $\eta \to 0$, GPBR dynamics generated by different entropies like Shannon, Tsallis or Burg would be expected to have only minute differences. Hence, it seems unlikely that such dynamics would converge to different Nash equilibria from the initial state if those equilibria are sufficiently far apart. But if we have a continuum of Nash equilibria, then even with minute differences, then different dynamics may converge to equilibria that are that are close to each other but still distinct.

Analytical answers to such questions are not feasible. Instead, we consider the game (2.23) and conduct certain numerical simulations of the GPBR dynamics generated by the three entropies we have been considering in this paper–Shannon, Tsallis (0.5) and Burg–in this game.[24] We present these simulations in Figure 2.1 for four different values of η, the highest being $\eta = 0.04$ and the lowest being $\eta = 0.006$. The initial distribution in all our simulations is the uniform distribution on S. In each case, we present the density function to which the GPBR dynamics converge to for that value of η. Thus, in the top left panel, where $\eta = 0.04$, the logit dynamic converges to the blue density function, the Tsallis

[23] Notice from (2.23) that if $A(\mathbb{P}) = 0.8$, $\pi_x(\mathbb{P}) = 0$ for all $x \in S = [0,1]$. Hence, any such $\mathbb{P}$ is a Nash equilibrium. It may be that there are other Nash equilibria as well with $A(\mathbb{P}) \neq 0.8$. But the potential game argument rules out that possibility.

[24] For the Tsallis (γ) entropy, we take $\gamma = 0.5$. Hence, the notation Tsallis (0.5). The GPBR dynamic generated by the Shannon entropy is, of course, the logit dynamic.

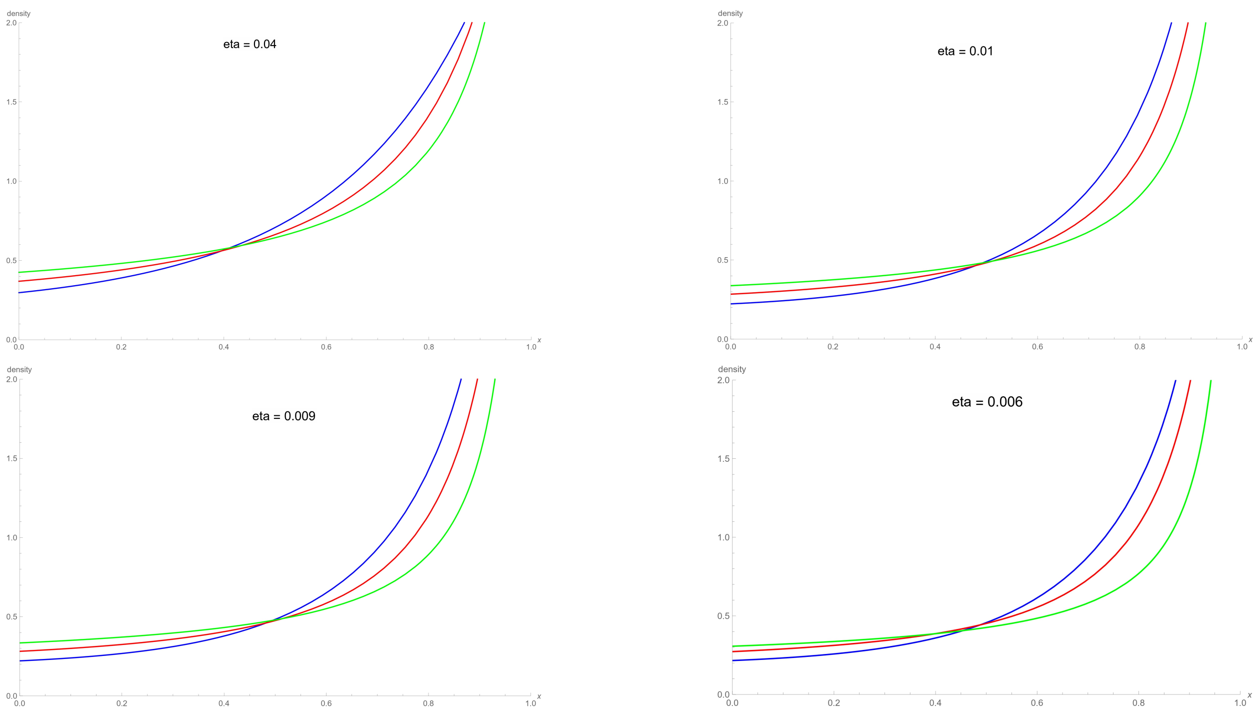

Figure 2.1: Density functions of the perturbed equilibrium under the Shannon entropy (in red), Tsallis (0.5) entropy (in blue) and Burg entropy (in green) for $\eta = 0.04, 0.01, 0.009, 0.006$ in the Cournot competition model (2.23). The horizontal axis represents strategies $x \in S = [0,1]$. In each case, the relevant GBPR dynamics converge from the uniform distribution on S to the distributions shown. As η declines, differences in the density functions persist suggesting convergence to different Nash equilibria as $\eta \to 0$.

(0.5) dynamic to the red and the Burg dynamic to the green one. Each such density function represents a perturbed equilibrium under that particular entropy. For such low values of η, each of them is also close to a Nash equilibrium of π. Therefore, by Proposition 1, the expected value of each of these distributions is close to 0.8. Values of $\eta < 0.006$ generated numerical errors as the perturbed best responses blow up at such low perturbations.

The simulations does suggest that in the limit, the GPBR dynamics generated by the three entropies are converging to different Nash equilibria. For each value of η, the distributions to which these dynamics are converging to are distinct. Of course, that is to be expected since the dynamics are different. But more significantly, these distinctions seem to persist even for low values of η. Hence, it is a reasonable conjecture that even as $\eta \to 0$, the three distributions will remain distinct and, therefore, converge to different Nash equilibria of π. Each of these limiting Nash equilibria will have an aggregate strategy level of 0.8. The simulations also suggest that the equilibria are not too distant (in distribution) from each other as the density functions presented are not drastically dissimilar. But they are not identical Nash equilibria. Hence, the different dynamics select different Nash equilibria in the limit.

We recognize that no simulation can ever provide conclusive evidence. But to the extent possible, the simulations presented in Figure 2.1 highlight the possibility that if there is a continuum of Nash equilibria, then GPBR dynamics generated by different perturbations can converge to different Nash equilibria as the perturbations become small. We emphasise that this is only a possibility and need not always happen whenever there is a continuum of equilibria. We can also now elaborate upon the significance of the linearity of the cost function in our example (2.23). It is this linear cost function that generates a convex set of equilibria. Instead, had the cost function been strictly convex, then we would have had a unique Nash equilibrium (Propositions 3 and 4, Cheung and Lahkar (2018)). In that case, the question of convergence to different Nash equilibria would not have arisen. We could have also considered the potential game defined by the quadratic payoff function (2.19) in Example 4. For $a > 2$, this game has three monomorphic Nash equilibria, $\delta_{\underline{x}}$, δ_0 and $\delta_{\bar{x}}$ (Proposition 4, Hofbauer et al. (2009)). Thus, in this case, there are multiple equilibria but which are disconnected. We also conducted certain simulations with the three GPBR dynamics in this game. The dynamics converge to either $\delta_{\underline{x}}$ or $\delta_{\bar{x}}$ (or, more precisely, the approximating perturbed equilibrium). But from an identical initial condition, all these dynamics converge to the same equilibrium. This supports our earlier hypothesis that with such disconnected equilibria, different GPBR dynamics will converge to the same limiting Nash equilibrium. But in our Cournot model with its simple linear cost function and the resulting continuum of Nash equilibria, even small differences in the perturbed best response can lead to selection of different equilibrium from the same initial state.

2.8 Conclusion

This paper has considered a generalization of the logit dynamic in large population games with a continuum of strategies. We have considered general deterministic perturbations that satisfy lower semicontinuity and strong convexity. Apart from the well known Shannon entropy, other perturbations that satisfy these conditions are the Tsallis entropy and the Burg entropy. Once we perturb payoffs using such perturbation functions and allow agents to best respond, we obtain the generalized perturbed best response (GPBR) and the associated dynamic. This constitutes our generalization of the logit dynamic.

We characterize the rest points of the GPBR dynamic, which we call perturbed equilibria, and show that they approximate Nash equilibria of the underlying population game as the extent of the perturbation becomes small. Establishing this result requires us to impose appropriate lower bounds on the perturbations we consider. We then establish the fundamental properties of the GPBR dynamic; namely, existence and uniqueness of solution trajectories, and their continuity with respect to initial conditions. We also establish convergence of the GPBR dynamic in potential games and negative semidefinite games. Finally, we apply the GPBR dynamics generated by the Shannon, Tsallis and Burg entropies to a large population Cournot model with a continuum of Nash equilibria. Our simulations in this application suggests that when the set of Nash equilibria is a continuum, then as $\eta \to 0$, the GPBR dynamics generated by the different entropies converge to different Nash equilibria. This is significant because it means that differences in perturbations does lead to differences in our prediction of the limiting social behavior.

It would be interesting to try and find other examples where different entropies generate different limiting behavior. One possibility is to consider a situation where a Nash equilibrium is Lyapunov stable (but not asymptotically stable) under the best response dynamic. Thus, there may be cycles around such equilibria under the best response dynamic. Then, it may happen that under certain perturbations, the cycle persist while under other perturbations, we obtain convergence to the Nash equilibrium. These are interesting possibilities to explore for the future. Finding such examples will be challenging because, as far as we know, there has been no general analysis of the best response dynamic in continuous strategy games. But if successful, then along with our present example of the Cournot model, our approach of generating perturbations with different entropies will add to the literature on refinements of Nash equilibria for games with a continuum of strategies (Myerson and Reny (2020)).

This paper has been confined to deterministic perturbations of the best response. Hence, another interesting research question for the future is to derive perturbed best response dynamics using stochastic perturbations. This exercise has led to interesting results in finite strategy large population games (Hofbauer and Sandholm (2002)). But so far, this question has remained unexplored in games with continuous strategy sets. One more possible direction of research is to extend GPBR dynamics from the single population considered here to multiple populations. This would not just be a theoretical exercise but can also lead to interesting applications involving multiple populations or type of agents. For example, Lahkar and Mukherjee (2019) apply continuous strategy evolutionary dynamics, including the logit dynamic, to the problem of implementing social efficiency in a public goods game. It will be interesting to apply GPBR dynamics to assess the implications of the choice of entropy on the spread of the perturbed equilibrium around the Pareto efficient state.[25]

[25] Evolutionary implementation relies on adding externalities as a transfer price to the original game to construct an externality adjusted game. Like the Cournot model in Section 2.7, the externality adjusted public goods game is also an aggregative potential game but with multiple types of agents. It also has a unique Nash equilibrium which is the Pareto efficient state of the original game. As in the Cournot game, GPBR dynamics will converge to a perturbed equilibrium that approximate

the Nash equilibrium. The spread of the perturbed equilibrium may depend upon which entropy we apply. But rigorously analyzing these issues will require an extension of these dynamics to multiple populations.

CHAPTER

3

THE LOGIT DYNAMIC IN SUPERMODULAR GAMES WITH A CONTINUUM OF STRATEGIES: A DETERMINISTIC APPROXIMATION APPROACH

3.1 Introduction

The deterministic approach to evolutionary game theory seeks to provide dynamical foundations to Nash equilibria through the use of evolutionary dynamics, which usually take the form of ordinary differential equations (Sandholm (2010)).[1] A large population of agents play a game and the distribution of strategies of agents in that population constitute the population state. Evolutionary dynamics are defined over such population states. The key focus of most deterministic evolutionary game theoretic models is whether the population state converges to a Nash equilibrium under different evolutionary dynamics.

Two important questions need to be addressed before admitting evolutionary dynamics as valid descriptions of the behavior of a large population. The first is about the behavioral foundations of

[1]Well known evolutionary dynamics include the replicator dynamic, the logit dynamic, the Brown–von Neumann–Nash (BNN) dynamic and the pairwise comparison dynamic. Sandholm (2010) provides an extensive exposition of such dynamics for finite strategy games. Such dynamics have also been extended to games with continuous strategy sets. See Cheung and Lahkar (2018) for a review. Not all evolutionary dynamics need to be differential equations. For example, the best response dynamic is a differential inclusion.

evolutionary dynamics. This question is resolved by appealing to revision protocols which describe how agents revise strategies when they receive a chance to do so. For example, imitative revision protocols generate the replicator dynamic and the logit best response, which is a particular perturbation of the best response, generates the logit dynamic.[2] The second key question is about the mathematical foundations of such dynamics. One justification of a deterministic evolutionary dynamic is through an informal appeal to the law of large numbers. Thus, we envisage a large population of agents where behavior is random but idiosyncratic. Hence, as long as the population size is sufficiently large, such randomness would be averaged away generating the deterministic dynamic (Chapter 4, Sandholm (2010)). However, to make these foundations more robust, we need to make this argument more rigorous.

The present paper, therefore, considers this second question in the context of the logit dynamic with a continuum of strategies (Perkins and Leslie (2014), Lahkar and Riedel (2015)). It does so by taking a *deterministic approximation* approach which we explain more fully below. This approach has been applied earlier for evolutionary dynamics defined on large population games with a finite strategy set (for example, Benaïm and Weibull (2003)). This paper extends that approach to an evolutionary dynamic defined for a continuum of strategies. We then apply the deterministic approximation approach to an important class of large population games; namely, supermodular games. Hofbauer and Sandholm (2007) analyze such games with finite strategy sets. We extend that notion to the context of continuous strategy sets and establish convergence of the continuum logit dynamic to logit equilibria, which are a perturbation of Nash equilibria, in such games. As far as we know, large population supermodular games with a continuum of strategies have not been considered earlier. Along with potential games and negative definite games, supermodular games are one of the three broad classes of large population games for which convergence results have been established under finite strategy evolutionary dynamics (Hofbauer and Sandholm (2007)). For continuous strategy dynamics, such results have been established for potential games and negative definite games independently of the deterministic approximation approach (see, for example, Lahkar and Riedel (2015)). But that approach is required for supermodular games and, hence, constitutes an important application in this paper.

An evolutionary dynamic is defined for a continuum of agents. The deterministic approximation approach to such a dynamic involves introducing a stochastic process that decides how agents revise strategies in a finite population. It is known, at least for games with a finite strategy set, that by making the number of players in the stochastic process sufficiently large, we can approximate its behavior over a finite time horizon by the solution trajectory of a deterministic dynamic.[3] Intuitively, with the number of players being so large, it becomes realistic to assume that the set of agents is a continuum in the evolutionary dynamic. As a purely mathematical result, such a deterministic approximation approach was introduced by Kurtz (1970). Since then, various authors have extended such results to evolutionary game theory in the context of finite strategy games. They include Boylan (1995), Binmore et al. (1995), Börgers and Sarin (1997), Schlag (1998), and Sandholm (2003). The strongest and most general such

[2]See, for example, Lahkar and Sandholm (2008) for some such revision protocols that generate the different important evolutionary dynamics. Also see Cheung (2014) for an extension of the notion of revision protocols to large population games with a continuum of strategies.

[3]See Chapter 10 of Sandholm (2010) for further details of this methodology.

results can be found in Benaïm and Weibull (2003). Thus, under this approach, a deterministic evolutionary dynamic may be interpreted as the approximation of the way a stochastic evolutionary process would behave over finite time periods when there are a large number of agents.

The deterministic approximation approach, therefore, provides rigorous mathematical foundations to evolutionary dynamics. From a practical point of view, this approach is useful because very often, it is simpler to analyze an evolutionary dynamic instead of the underlying stochastic process. However, as far as we know, all such deterministic approximation results are for large population games with a finite number of strategies. But by now, there is a significant literature on evolutionary dynamics with a continuum of strategies. These include the replicator dynamic and the more general class of imitative dynamics (Oechssler and Riedel (2001),Oechssler and Riedel (2002), Cheung (2016)), the BNN dynamic (Hofbauer et al. (2009)), the pairwise comparison dynamic (Cheung (2014)), the logit dynamic (Perkins and Leslie (2014), Lahkar and Riedel (2015)) and the more general class of perturbed best response dynamics to which the logit dynamic belongs (Lahkar et al. (2022a)). Such dynamics have also been applied to a variety of classical models in economics like public goods, common resource problems and Cournot competition (Lahkar and Mukherjee (2019), Lahkar and Mukherjee (2021), Lahkar and Ramani (2021)). Hence, dynamics with a continuum of strategies are now a well–established part of evolutionary game theory, both from the theoretical perspective as well as the applied. Investigating their mathematical foundations through the deterministic approximation approach would, therefore, place them on a more solid footing. This is the main motivation of this paper.

We consider the logit dynamic with a continuous strategy set for this purpose. This dynamic is one of the canonical dynamics of learning and evolutionary game theory. It was introduced by Fudenberg and Levine (1998) for finite strategy games in the context of stochastic fictitious play and has since then been widely used to model evolution in large population games.[4] It is the prototype of the wider class of perturbed best response dynamics (Hofbauer and Sandholm (2002), Hofbauer and Sandholm (2007)) and is generated by perturbing payoffs with the Shannon entropy. Thus, it is an approximation of the best response dynamic but is always uniquely defined and, hence, is a differential equation. Therefore, it is more tractable than the best response dynamic, particulary in large population games with a continuous strategy set where the best response may not even exist.

Our focus on the logit dynamic is because of its importance in evolutionary game theory. A larger goal, of course, would be to have a general deterministic approximation result for all continuous strategy evolutionary dynamics. Right now, however, that is beyond our reach. In fact, even deriving such a result for the broader class of continuous strategy perturbed best response dynamics may not be feasible at present. Even though Lahkar et al. (2022a) present a general class of continuous strategy perturbed best response dynamics using perturbations other than the Shannon entropy, closed form expressions of such dynamics are not easily derivable. Our analysis in the present paper, on the other hand, relies significantly on the closed form expression of the logit best response, upon which the logit dynamic is based. Nevertheless, the paper does illustrate the possibility of obtaining deterministic approximation

[4] See, for example, Hofbauer and Hopkins (2005) and Hofbauer and Sandholm (2007). As noted earlier, this dynamic has been extended to continuous strategy sets by Perkins and Leslie (2014) and Lahkar and Riedel (2015).

results for continuous strategy evolutionary dynamics, even though the technical challenges of extending it more generally to other dynamics may be significant.

We consider a population game generated by pairwise random matching which has a continuum of strategies and is played by a continuum of agents. Our objective is to provide rigorous foundations to the logit dynamic in such a game. We first define a step–wise approximation of the original game and argue that solution trajectories of the original logit dynamic are close to those of the logit dynamic in the step–wise approximation. We then consider a discrete approximation of the original game with a finite number of strategies and a finite number of players and introduce the discrete time logit stochastic process in this game. But in the logit dynamic, each of the three variables, time, players and strategies, is continuous. Hence, our approximation approach has to address each one of them. Therefore, first, we extend the logit stochastic process defined initially over discrete time to continuous time by taking its interpolation. Next, by adapting the deterministic approximation framework of Benaïm and Weibull (2003), we conclude that for a finite number of strategies, as the number of players become large, the logit stochastic process would approximate the step–wise approximation of the logit dynamic. The additional technical challenge that this paper addresses is then to also let the set of strategies approach a continuum. Thus, by taking the double limit as, first, the number of players and then, the number of strategies approach infinity, we show that the behavior of the stochastic process approximates the continuum logit dynamic.

This gives us our desired deterministic approximation result. As per this result, the continuum logit dynamic is an approximation of the logit stochastic process in a finite game as the number of strategies and players in that game go to infinity. We should note that even though we apply the finite strategy deterministic approximation framework of Benaïm and Weibull (2003), its extension to a continuum of strategies is not trivial. For example, as we focus on continuous strategy games, we have to introduce new techniques to relate continuum and discrete strategy games while constructing the logit stochastic process. Further, even determining the notion of closeness in the continuous strategy case raises significant topological considerations like whether to choose the strong or weak topology. All our results are based on the strong topology or the topology induced by the total variation norm. As is typically the case, our deterministic approximation result is a finite horizon result. It holds for finite time intervals but may break down if the time span is of infinite length.

In the existing literature, Perkins and Leslie (2014) has also sought to approximate a stochastic process using the continuous strategy logit dynamic. They consider a finite player learning model in a continuous strategy game where the beliefs of players evolve according to the stochastic fictitious play rule. This learning rule generates a discrete time stochastic process which the authors, using an asymptotic pseudo–trajectory approach, approximate using the continuum logit dynamic. Their approach, however, vary significantly from ours. First, as theirs is a learning model, they approximate the evolution of beliefs using the logit dynamic. But as ours is a large population evolutionary model, we need to approximate evolution of the strategy distribution in such a population through the logit dynamic. Hence, our methodology is closer to the literature on finite strategy deterministic approximation in large population games described earlier. Second, of the three variables–time, players and strategy–their approach only requires approximating the discrete time process with the continuous time logit dynamic. In contrast, as discussed earlier, our deterministic dynamic not only approximates a discrete time process with a deterministic trajectory, but also a finite strategy and finite player game with a game having a continuous strategy set and a continuum of players. Therefore, the foundations we provide to the continuum logit dynamic are very different from that in Perkins and Leslie (2014). We should, however, note one point of similarity. Both papers provide a finite horizon approximation result.

We then apply our deterministic approximation approach to the question of convergence of the continuum logit dynamic in supermodular games. Supermodular games are characterized by strategic complementarities and form an important class of games in economics with applications like search, coordination and modeling of positive externalities (Topkis (1998), Vives (1990), Milgrom and Roberts (1990)). Hofbauer and Sandholm (2007) introduce large population finite strategy supermodular games and establish convergence of perturbed best response dynamics to perturbed equilibria in such games. [5] We consider a continuous strategy supermodular game and apply the approach of Hofbauer and Sandholm (2007) to argue that the logit stochastic process will converge in the medium run, i.e. over a finite time horizon, to a logit equilibrium of the discrete approximation of the continuous strategy game. Our deterministic approximation result then allows us to conclude that the logit dynamic converges to logit equilibria in the continuous strategy supermodular game. Along the way, we have to address significant technical problems like showing that the sequence of logit equilibria in a discrete approximation converges to a logit equilibrium of the original continuous strategy game under the strong topology.

As noted earlier, this is the first paper that defines continuous strategy supermodular games and, therefore, also the first paper that establishes any result on convergence of an evolutionary dynamic in such games. One question that arises is whether it is necessary to apply the deterministic approximation approach to establish such a convergence result. Hofbauer and Sandholm (2007) establish their convergence results both for the perturbed best response stochastic process as well as for the perturbed best response dynamic directly. Can't such results be established directly in our case as well for the continuous strategy logit dynamic? The difficulty is that the direct result in Hofbauer and Sandholm (2007) (Theorem 3.3) is based on the theory of cooperative differential equations (Smith (1995)). As far as we know, that theory has not been extended to differential equations with a continuous set of variables. Hence, analyzing supermodular games directly with the continuous strategy logit dynamic in not feasible for us. We have to rely upon the deterministic approximation approach, due to which our results on supermodular games are an important application of our approach.

To summarize, the paper makes two important contributions. First, it provides a rigorous mathematical foundation to the continuous strategy logit dynamic, one of the canonical dynamics of evolutionary game theory. By doing so, it adds to our understanding of how the logit dynamic describes behavior in a large population of players. Second, it establishes the convergence of logit dynamic to logit equilibria in supermodular games. This result enhances the scope of applying evolutionary game theoretic

[5] A perturbed best response equilibrium is an approximation of a Nash equilibrium. It is a rest point of a perturbed best response dynamic. For the logit dynamic in particular, we refer to its rest points as logit equilibria.

arguments to economic phenomenon that can be modeled as supermodular games.

The rest of the paper is as follows. Section 3.2 introduces the model and establishes the step–wise approximation of the continuum logit dynamic. In Section 3.3, we derive our main deterministic approximation result. Section 3.4 defines large population supermodular games with a continuum of strategies and establishes the medium run convergence of the logit dynamic in such games. Section 6.6 concludes.

3.2 Preliminaries

We consider a population of unit mass with each agent in the population being of measure zero. Let $\mathscr{S} := [0,1] \subset \mathbb{R}$ be the strategy space of the players in the population.[6] We endow $\mathscr{S}$ with the usual Borel sigma-algebra $\mathscr{B}(\mathscr{S})$. Let $\mathscr{P}(\mathscr{S})$ denote the space of probability measures on $\mathscr{S}$. A probability measure $\mu \in \mathscr{P}(\mathscr{S})$ denotes the state of the population which we interpret as the distribution of strategies in the population. Hence, $\mu(A)$ is the proportion of players in the population using strategies in $A \subseteq \mathscr{S}$. Let $\mathscr{L}^{\infty}(\mathscr{S})$ be the collection of bounded measurable functions on $\mathscr{S}$. A *population game* is a weakly continuous mapping $\varphi : \mathscr{P}(\mathscr{S}) \to \mathscr{L}^{\infty}(\mathscr{S})$ such that $\varphi_x(\mu)$ denotes the payoff to a player playing strategy x at population state μ. A population state μ° is said to be a Nash equilibrium of the underlying population game φ if for all $x \in \mathrm{Supp}(\mu^{\circ})$ and all $y \in \mathscr{S}$, we have $\varphi_x(\mu^{\circ}) \geq \varphi_y(\mu^{\circ})$.[7]

In this paper, we focus on games with pairwise interaction. We define such a game using a function $f : \mathscr{S} \times \mathscr{S} \to \mathbb{R}$ such that $f(x,y)$ is the payoff of an agent who plays x when randomly matched with another agent playing y. Players are randomly matched in pairs to play the game. We can then write payoffs in the associated population game φ as the expected value

$$\varphi_x(\mu) = \int_{\mathscr{S}} f(x,y)\mu(dy), \quad \text{for all } x \in \mathscr{S}. \tag{3.1}$$

An example of such a game with pairwise interaction is as follows.

Example 5 *Consider strategies $x, y \in \mathscr{S}$ and let $f(x,y) = m(x)y - c(x)$, where m is an increasing function of x and c is an arbitrary function of x. Applying (3.1), we obtain a population game φ in which the payoff of an agent playing strategy x at a population state μ is*

$$\begin{aligned}\varphi_x(\mu) &= \int_{\mathscr{S}} f(x,y)\mu(dy) \\ &= m(x)\int_{\mathscr{S}} y\mu(dy) - c(x)\int_{\mathscr{S}} \mu(dy) \\ &= m(x)a(\mu) - c(x),\end{aligned} \tag{3.2}$$

[6]The strategy set being $[0,1]$ is a convenient normalization. Our arguments extend to any strategy set that is a compact convex subset of $\mathbb{R}$.

[7]We denote the support of a probability measure μ by $\mathrm{Supp}(\mu)$.

where $a(\mu)=\int_{\mathscr{S}} y\mu(dy)$ is the aggregate strategy level in the population. We note that m being an increasing function in (3.2) is not required to obtain the population game φ from f. But we make this assumption in order to establish supermodularity later.

We focus on the logit dynamic for a large population game φ with a continuous strategy set. To define the dynamic, fix a parameter $\eta>0$. Then the *continuum strategy logit best response* is a mapping $\mathfrak{L}_\eta:\mathscr{P}(\mathscr{S})\to\mathscr{P}(\mathscr{S})$ such that for all $\mu\in\mathscr{P}(\mathscr{S})$,

$$\mathscr{L}_\eta[\mu](A):=\int_A \frac{\exp(\eta^{-1}\varphi_x(\mu))dx}{\int_{\mathscr{S}}\exp(\eta^{-1}\varphi_y(\mu))dy}, \quad \text{for all } A\in\mathscr{B}(\mathscr{S}). \tag{3.3}$$

Intuitively, the logit best response is an approximation of the best response, particularly for η small. In that case, the probability measure $\mathfrak{L}_\eta[\mu]$ puts nearly all its mass on the set of best responses to μ. Unlike best responses, though, the logit best response is uniquely defined and varies smoothly with μ. From (3.3), we then obtain the *continuum strategy logit dynamic* (Perkins and Leslie (2014); Lahkar and Riedel (2015))

$$\dot{\mu}=\mathfrak{L}_\eta(\mu)-\mu \tag{3.4}$$

induced by φ. Like any other evolutionary dynamic, the logit dynamic also measures the rate of change of the population state. Under the logit dynamic, a population state moves in the direction of the logit best response.

The definition of the logit dynamic does not depend upon φ being generated by the underlying pairwise interaction game f. But our subsequent results will depend greatly on the pairwise interaction structure. A probability measure μ° is said to be a *logit equilibrium* if $\mathfrak{L}_\eta(\mu^\circ)=\mu^\circ$. Theorem 3.1 in Lahkar and Riedel (2015) establishes the existence of such equilibria. As $\eta\to 0$, any such logit equilibrium converges weakly to a Nash equilibrium of φ (Theorem 3.2, Lahkar and Riedel (2015)). We use the notation $\mathrm{LE}_\eta(\varphi)$ to denote the collection of logit best response equilibrium.

3.2.1 Step-wise Approximation of the Logit Dynamic

We now introduce a step–wise approximation of the logit dynamic that will be useful for our analysis that follows. Our objective is to show that this step–wise approximation is close to the original logit dynamic (3.4).The methodology we apply has certain similarities with the one employed by Oechssler and Riedel (2002) to establish the closeness of the continuous strategy replicator dynamic with its step–wise approximation.[8]

For this purpose, consider $n\geq 1$, $j\in\{0,1,\ldots,2^n-1\}$ and denote the dyadic rational number $\frac{j}{2^n}$ by $\alpha_{n,j}$. Further, let S_n be the collection of left-closed right-open intervals with dyadic rational endpoints in $\mathscr{S}$ defined in the following way:

$$\mathscr{S}_n:=\left\{[\alpha_{n,j},\alpha_{n,j+1}):j=0,\ldots,2^n-2\right\}\cup[\alpha_{n,2^n-1},1]. \tag{3.5}$$

[8] See Section 7 of Oechssler and Riedel (2002).

Thus, $\mathscr{S}_n$ is a collection of 2^n disjoint intervals which covers the original strategy space $\mathscr{S}$. Recall the payoff $f(x,y)$ in (3.1). Our next step is to construct a step function approximation of f. For $n \geq 1$, let $f_n : \mathscr{S} \times \mathscr{S} \to \mathbb{R}$ is such that for all $(x,y) \in \mathscr{S} \times \mathscr{S}$,

$$f_n(x,y) = \begin{cases} f(\alpha_{n,j}, \alpha_{n,k}) & \text{if } (x,y) \in [\alpha_{n,j}, \alpha_{n,j+1}) \times [\alpha_{n,k}, \alpha_{n,k+1}) \text{ and } j,k = 0,1,\ldots,2^n - 2 \\ f(\alpha_{n,2^n-1}, \alpha_{n,2^n-1}) & \text{if } (x,y) \in [\alpha_{n,2^n-1}, 1] \times [\alpha_{n,2^n-1}, 1]. \end{cases} \tag{3.6}$$

Hence, by definition, $(f_n)_{n\geq 1}$ is a sequence of step functions such that $f_n \to f$ uniformly as $n \to \infty$.

Now, similar to the definition of the population game φ in (3.1), let $\varphi^n : \mathscr{P}(\mathscr{S}) \to \mathscr{L}^\infty(\mathscr{S})$ be such that for all $\mu \in \mathscr{P}(\mathscr{S})$,

$$\varphi_x^n(\mu) := \int_{\mathscr{S}} f_n(x,y)\mu(dy), \quad \text{for all } x \in \mathscr{S}. \tag{3.7}$$

Note from (4.16) that we can write (3.7) as

$$\varphi_x^n(\mu) = \begin{cases} \sum_{k=0}^{2^n-1} f(\alpha_{n,j}, \alpha_{n,k}) \int_{[\alpha_{n,j}, \alpha_{n,j+1})} \mu(dy) \text{ if } x \in [\alpha_{n,j}, \alpha_{n,j+1}) \text{ and } j = 0,1,\ldots,2^n - 2 \\ \sum_{k=0}^{2^n-1} f(\alpha_{n,2^n-1}, \alpha_{n,k}) \int_{[\alpha_{n,2^n-1}, 1]} \mu(dy) \text{ if } x \in [\alpha_{n,2^n-1}, 1]. \end{cases} \tag{3.8}$$

We note that for every $\mu \in \mathscr{P}(\mathscr{S})$, φ^n is also a step function approximation of φ as defined in (3.1). This follows from the definition of f_n in (4.16). For a given $n \geq 1$ and parameter $\eta > 0$, we can then define the *logit best response induced by* φ_n as a mapping $\mathfrak{L}_\eta^n : \mathscr{P}(\mathscr{S}) \to \mathscr{P}(\mathscr{S})$ such that for all $\mu \in \mathscr{P}(\mathscr{S})$,

$$\mathscr{L}_\eta^n[\mu](A) := \int_A \frac{\exp(\eta^{-1}\varphi_x^n(\mu))dx}{\int_{\mathscr{S}} \exp(\eta^{-1}\varphi_y^n(\mu))dy}, \quad \text{for all } A \in \mathscr{B}(\mathscr{S}). \tag{3.9}$$

Comparing (3.3) and (3.9), it is evident that the probability density function of the logit best response induced by φ_n is also a step–wise approximation of the probability density function of the logit best response induced by φ. Finally for $n \geq 1$ and $\eta > 0$, we define the logit dynamic induced by φ_n as

$$\dot{\mu} = \mathfrak{L}_\eta^n(\mu) - \mu. \tag{3.10}$$

Letting $t \geq 0$ be time, we denote as $\mu(t) \in \mathscr{P}(\mathscr{S})$ and $\mu^n(t) \in \mathscr{P}(\mathscr{S})$, $t \geq 0$, the solution trajectories of the two logit dynamics (3.4) and (3.10) respectively from the initial state $\mu(0)$. Theorem 3.4 in Lahkar and Riedel (2015) establishes the existence of such unique solution trajectories of the logit dynamic from every initial state in $\mathscr{P}(\mathscr{S})$ provided the underlying population game is bounded and Lipschitz continuous with respect to the variational norm.[9] Such conditions are automatically satisfied if φ is as

[9]The total variation distance is a mapping $\|\cdot\|_{\mathbf{TV}} : \mathscr{P}(\mathscr{S}) \times \mathscr{P}(\mathscr{S}) \to \mathbb{R}_+$ defined as

$$\|\mu - \nu\|_{\mathbf{TV}} := 2 \sup_{A \in \mathscr{B}(\mathscr{S})} |\mu(A) - \nu(A)|, \quad \text{for all } \mu, \nu \in \mathscr{P}(\mathscr{S}).$$

defined in (3.1) since the underlying two player game f, and hence $(f_n)_{n\geq 1}$ are bounded. Intuitively, it is also reasonable to expect that these two solution trajectories would be close to each other since the logit best response (3.9) is itself a step function approximation of (3.3). We formalize this intuition in the following theorem.

Proposition 2 *Consider the population game φ as defined in (3.1) and its step function approximation φ^n as defined in (3.7). Denote by $(\mu(t))_{t\geq 0}$ the solution to the logit dynamic (3.4) induced by φ and by $(\mu^n(t))_{t\geq 0}$ the solution to the logit dynamic (3.10) induced by φ_n for $n \geq 1$. Suppose that $\mu(0) = \mu^n(0)$ for every $n \geq 1$. Then $(\mu^n(t))_{t\geq 0}$ converges uniformly on compact sets to $(\mu(t))_{t\geq 0}$ in the variational norm, that is, for every finite time horizon $T > 0$,*

$$\lim_{n\to\infty} \sup_{0\leq t\leq T} \|\mu(t) - \mu^n(t)\|_{TV} = 0.$$

Proof 2 *See Appendix B.1.*

Proposition 2 is our desired result on the step–wise approximation of the logit dynamic. It shows that the solution trajectories of the two dynamics are close to each other the strong topology, or the topology induced by the total variation norm. Of course, the result requires that the initial states be identical. Also, n needs to be sufficiently large so that the step function f_n defined in (4.16) is a reasonably close approximation of the original pairwise interaction payoff f in (3.1). Intuitively, the probability mass on each interval of the finite set S_n defined in (4.15) under the two trajectories $\mu(t)$ and $\mu^n(t)$ become increasingly identical as $n \to \infty$.

3.3 Deterministic Approximation

The deterministic approximation approach to the continuum logit dynamic would require us to approximate it as the behavior of a stochastic process in a discrete game where the number of strategies and the number of players are finite but sufficiently large. With this objective, we now define the discrete analogue of our population game φ. Consider the finite set

$$S_n = \left\{\alpha_{n,0}, \alpha_{n,1}, \alpha_{n,2}, \ldots, \alpha_{n,2^n-1}\right\}, \tag{3.11}$$

where, as defined previously, $\alpha_{n,j} = \frac{j}{2^n}$. Thus, S_n is a finite approximation of $\mathscr{S} = [0,1]$. Like $\mathscr{S}_n$ in (4.15), it contains 2^n elements with the difference being it consists of points instead of intervals. Each element of S_n is the leftmost point of the 2^n intervals in $\mathscr{S}_n$. Now, let

$$\Delta(S_n) = \left\{\mathbf{p}^n \in \mathbb{R}_+^{2^n} : \sum_{0\leq k\leq 2^n-1} \mathbf{p}_i^n = 1\right\} \tag{3.12}$$

be the set of probability distributions on S_n. Hence, $\mathbf{p}_j^n$ is the probability on $\alpha_{n,j}$ for $j \in \{0, 1, \ldots, 2^n - 1\}$. Recalling the step function f_n from (4.16), we then define payoffs in the discrete version of φ as follows.

Definition 11 *For $n \geq 1$, the* ***level-n discretization of a population game*** φ *is the mapping* $\mathscr{D}^n_\varphi : \Delta(S_n) \to \mathbb{R}^{2^n}$ *such that for all* $\boldsymbol{p}^n \in \Delta(S_n)$,

$$\mathscr{D}^n_\varphi[\boldsymbol{p}^n](\alpha_{n,j}) := \langle f_n(\alpha_{n,j},\cdot), \boldsymbol{p}^n \rangle = \sum_{0 \leq k \leq 2^n-1} f(\alpha_{n,j},\alpha_{n,k})\boldsymbol{p}^n_k, \quad \textit{for all } j = 0,1,\ldots,2^n-1. \tag{3.13}$$

We can, therefore, envisage $\mathscr{D}^n_\varphi$ as the population game induced by the matrix game whose entries are given by $(f_n(\alpha_{n,j},\alpha_{n,k}))_{0\leq j,k\leq 2^n-1}$. Equivalently, $\mathscr{D}^n_\varphi$ is a finite strategy population game in which players' strategy set is S_n and the proportion of players playing $\alpha_{n,j}$ is $\mathbf{p}^n_j$. Notice from (4.18) and (3.13) that if we have $\mu \in \mathscr{P}(\mathscr{S})$ and $\mathbf{p}^n \in \Delta(S_n)$ such that $\mathbf{p}^n_j = \int_{[\alpha_{n,j},\alpha_{n,j+1})} \mu(dy)$ for $j = 0,1,\ldots,2^n-2$, then the payoffs in φ^n and $\mathscr{D}^n_\varphi$ are identical.

We focus on the discrete game $\mathscr{D}^n_\varphi$ introduced in Definition 11. Further, we now assume that the game is being played by a finite number of players, with $N > 1$ being the number of players . Recall our interpretation of $\mathbf{p}^n_j$ as being the mass of agents playing $\alpha_{n,j} \in S_n$ in $\mathscr{D}^n_\varphi$. For $n \geq 1$, and $\mathbf{p}^n \in \Delta(S_n)$, we also define the atomic probability measure on $\mathscr{S}$ as

$$\mu^N_{\mathbf{p}^n} := \sum_{0 \leq j \leq 2^n-1} \mathbf{p}^n_j \delta_{\alpha_{n,j}}, \tag{3.14}$$

where δ_x denotes the Dirac measure at $x \in \mathscr{S}$ and the superscript N indicates the dependence on the number of players N.[10] Thus, the measure $\mu_{\mathbf{p}^n}$ on $\mathscr{S}$ puts probability $\mathbf{p}^n_j$ on $\alpha_{n,j} = \frac{j}{2^n}$, $j \in \{0,1,\cdots,2^n-1\}$ and 0 on all other points. This atomic measure now allows us to treat this discretized game as being played on $\mathscr{S}$ with the proportion of players using strategy $\alpha_{n,j}$ being $\mathbf{p}^n_j$.

We now introduce the logit best response on the discrete game $\mathscr{D}^n_\varphi$. Given the parameter $\eta > 0$ and $n \geq 1$, this logit best response $L^n_\eta : \Delta(S_n) \to \Delta(S_n)$ at the state $\mathbf{p}^n$ takes the standard finite dimensional form (Fudenberg and Levine (1998))

$$\begin{aligned} L^n_{\eta,i}[\mathbf{p}^n] &= \frac{\exp(\eta^{-1}\mathscr{D}^n_\varphi[\mathbf{p}^n](\alpha_{n,j}))}{\sum_{0\leq k\leq 2^n-1}\exp(\eta^{-1}\mathscr{D}^n_\varphi[\mathbf{p}^n](\alpha_{n,k}))} \\ &= \frac{\exp(\eta^{-1}\langle f_n(\alpha_{n,i},\cdot),\mathbf{p}^n\rangle)}{\sum_{0\leq k\leq 2^n-1}\exp(\eta^{-1}\langle f_n(\alpha_{n,k},\cdot),\mathbf{p}^n\rangle)}. \end{aligned} \tag{3.15}$$

Thus, given η, (3.15) is the probability with which strategy $\alpha_{n,i} \in \mathscr{S}$ would get chosen under the logit best response in the population game $\mathscr{D}^n_\varphi$ at the population state $\mu^N_{\mathbf{p}^n}$.

Let $(\Omega, \mathscr{F}, \mathbb{P})$ be an arbitrary probability space. We now consider the following discrete time stochastic process $(\mathbb{X}^N_n(k))_{k\geq 1}$, where $\mathbb{X}^N_n(k) : \Omega \to \mathscr{P}(\mathscr{S})$ describes the state of the population at time k for every

[10] This dependence will arise from the fact that with N being the number of players, $\mathbf{p}^n_j$ can take values only in $\{0, \frac{1}{N}, \ldots, \frac{N-1}{N}, 1\}$.

$k \geq 1$. We use the discrete logit best response (3.15) to define the transition probability function as

$$\mathbb{P}\left(\mathbb{X}_n^N\left(k+\frac{1}{N}\right)=\mu^N_{\mathbf{q}^n} | \mathbb{X}_n^N(k)=\mu^N_{\mathbf{p}^n}\right)=\begin{cases}\mathbf{p}_i^n L^n_{\eta,j}[\mathbf{p}^n], \text{ if } \mathbf{q}^n=\mathbf{p}^n+\frac{1}{N}(\mathbf{e}_j-\mathbf{e}_i) \text{ and } j\neq i\\ \sum_{i=0}^{2^n-1}\mathbf{p}_i^n L^n_{\eta,i}[\mathbf{p}^n] \text{ if } \mathbf{q}^n=\mathbf{p}^n\\ 0, \text{ otherwise.}\end{cases} \tag{3.16}$$

To understand these transition probabilities, we suppose that at discrete time intervals of length $\frac{1}{N}$, a single agent gets the opportunity to change strategy.[11] If that agent is currently playing strategy $\alpha_{n,i}$ and chooses to shift to $\alpha_{n,j}$, then the social state would change from $\mu_{\mathbf{p}^n}$ to $\mu_{\mathbf{q}^n}$, with $\mathbf{q}^n=\mathbf{p}^n+\frac{1}{N}(\mathbf{e}_j-\mathbf{e}_i)$. The probability of this transition is the mass of agents playing $\alpha_{n,i}$, which is $\mathbf{p}_i^n$, multiplied by the logit probability of playing $\alpha_{n,j}$, which is $L^n_{\eta,j}[\mathbf{p}^n]$. Alternatively, the agent can continue playing the same strategy, the probability of which is given by the second line of (3.16).

In the rest of this section, we show that the behavior of the logit stochastic process in $\mathscr{D}^n_\varphi$ approximates the behavior of the continuum logit dynamic (3.4) in φ over finite time spans as the number of players, N, and the number of strategies, 2^n, becomes large. That would be the interpretation of the continuum logit dynamic as a deterministic approximation of the logit stochastic process in a finite strategy and finite player game. For that, we now extend the process (3.16) defined over discrete time to one over continuous time. Thus, let $(\widehat{\mathbb{X}}_n^N(t))_{t\geq 0}$ denote the continuous time Markov process obtained by affine interpolation of the process $(\mathbb{X}_n^N(k))_{k\geq 1}$ as defined in (3.16). We call this process the **interpolated logit stochastic process**. We establish our deterministic approximation result with respect to this interpolated process. Note that this process is defined on the discrete game $\mathscr{D}^n_\varphi$. Establishing that result requires us to define the notion of smoothing of a probability measure.

Definition 12 *For $n\geq 1$, the **smoothing of a probability measure** is a mapping $\mathfrak{s}_n:\mathscr{P}(S_n)\to\mathscr{P}(\mathscr{S})$ such that the following two conditions are satisfied:*[12]

- *$\mathfrak{s}_n(\mu^N_{\boldsymbol{p}^n})$ is absolutely continuous with respect to the Lebesgue measure λ, for all $\boldsymbol{p}^n\in\Delta(S_n)$.*[13]
- *The Radon-Nikodym derivative of the probability measure $\mathfrak{s}_n(\mu^N_{\boldsymbol{p}^n})$ is expressed as*

$$\frac{\int^n(\mu^N_{\boldsymbol{p}^n})}{d\lambda}(x):=\sum_{0\leq k\leq 2^n-2}2^n\mu^N_{\boldsymbol{p}^n,k}\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x), \quad \textit{for all } x\in\mathscr{S}. \tag{3.17}$$

The smoothing of a probability measure transforms atomic measures on $\mathscr{S}$ into probability measures admitting a probability density function. Suppose that we have an atomic measure $\mu_{\mathbf{p}^n}$ on $\mathscr{S}$, where

[11] Here, $\mathbf{e}_j$, $j\in\{0,1,\ldots,2^n-1\}$, is the standard basis vector on $\mathbb{R}^{2^n}$. Under our atomic measures, the strategies that receive positive probability are $\alpha_{n,j}=\frac{j}{2^n}$, $j\in\{0,1,\ldots,2^n-1\}$. Hence, $\mathbf{e}_j$ assigns value 1 to strategy $\frac{j}{2^n}$ and 0 to all other strategies.

[12] Here, $\mathscr{P}(S_n)$ denotes the collection of atomic probability measures like (3.14) with support S_n.

[13] A probability measure μ is said to be absolutely continuous with respect to the Lebesgue measure λ if for all $A\in\mathscr{B}(\mathscr{S})$, we have $\lambda(A)=0 \Longrightarrow \mu(A)=0$.

the probabilities are concentrated on the points $\alpha_{n,j}$. Then $\mathfrak{s}_n(\mu_{\mathbf{p}^n})$ denotes the probability measure obtained by distributing the probability mass at $\alpha_{n,j}$ uniformly on the interval $[\alpha_{n,j},\alpha_{n,j+1})$. We, thereby, construct a probability measure with a probability density function that is a step function. This density function is given by (3.17). Applying Definition 12 to $(\widehat{\mathbb{X}}_n^N(t))_{t\geq 0}$, we can derive the smoothed interpolated logit stochastic process. As an intermediate step towards our desired result, we first establish a deterministic approximation result with respect to the step–wise logit dynamic (3.10).

Proposition 3 *Suppose that the interpolated logit stochastic process* $(\widehat{\mathbb{X}}_n^N(k))_{k\geq 1}$ *has as initial condition the atomic measure* $\mu_{\boldsymbol{p}^n}^N \in \mathscr{P}(\mathscr{S})$*, for some* $\boldsymbol{p}^n \in \Delta(S_n)$*. Let* $(\mu^n(t))_{t\geq 0}$ *denote the solution to the level-*n *step logit dynamic (3.10) induced by the population game* φ_n *with initial condition* $\mu^n(0)=\mu_{\boldsymbol{p}^n}^N$*. Then for every* $T>0$*, there exists* $\bar{C}>0$ *(independent of* $\mu^n(0)$*) such that for all* $\epsilon>0$ *and for every fixed* n *large enough,*

$$\mathbb{P}_{\mu_{\boldsymbol{p}^n}^N}\left(\sup_{0\leq t\leq T}\left\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))-\mu^n(t)\right\|_{\boldsymbol{TV}}\geq \epsilon\right)\leq 2(2^n+1)e^{-N\bar{C}\epsilon^2},$$

as $N\to\infty$.

Proof 3 *See Appendix B.2.*

Proposition 3, therefore, approximates the interpolated logit stochastic process in $\mathscr{D}_\varphi^n$ with the step–wise deterministic logit dynamic (3.10) over finite time horizons. It does so by fixing the number of strategies, indicated by n, at a sufficiently large level and then allowing the number of players, N, in $\mathscr{D}_\varphi^n$ to increase. Its proof can broadly be divided into two steps. First, we show that for every $\eta>0$, the level-n step logit choice $\mathfrak{L}_\eta^n$ is "approximately Lipschitz continuous" for sufficiently large n in a sense made precise in Appendix B.2. Then, along the lines of Benaïm and Weibull (2003), we bound the total variation norm between the smoothed stochastic evolutionary process and the solution to the level-n step logit dynamic uniformly over finite time horizons using Grönwall's inequality and some other relevant estimates to complete the proof.

Thus, Proposition 3 is a deterministic approximation result for the step–wise logit dynamic. Its importance lies as a key intermediate step towards our main objective, which is establishing an approximation result not just for the step–wise dynamic but the original continuum logit dynamic. We state that result in the following theorem by allowing not just N but also n, which was held fixed in Proposition 3, to go to infinity.

Theorem 6 *Fix* $n,N\geq 1$*. Suppose that the interpolated logit stochastic process* $(\widehat{\mathbb{X}}_n^N(t))_{t\geq 0}$ *has initial condition* $\mu_{\boldsymbol{p}^n}^N$*, for some* $\boldsymbol{p}^n\in\Delta(S_n)$*. Let* $(\mu(t))_{t\geq 0}$ *denote the solution to the continuum strategy logit dynamic induced by the population game* φ *with initial condition* $\mu(0)=\mu_{\boldsymbol{p}^n}^N$*. Then for every* $T>0$*, and every* $\epsilon>0$*, we have that*

$$\lim_{n\to\infty}\lim_{N\to\infty}\mathbb{P}_{\mu_{\boldsymbol{p}^n}^N}\left(\sup_{0\leq t\leq T}\left\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))-\mu(t)\right\|_{\boldsymbol{TV}}\geq \epsilon\right)=0.$$

Proof 4 *See Appendix B.3.*

Theorem 6 is our desired deterministic approximation result and is a consequence of Propositions 2 and 3. By Proposition 3, the step–wise logit dynamic provides a deterministic approximation of the logit stochastic process for sufficiently large but fixed n as $N \to \infty$. But Proposition 2 implies that the step–wise logit dynamic itself is an arbitrarily close approximation of the original continuum logit dynamic as $n \to \infty$. Hence, intuitively, the original logit dynamic should also be a deterministic approximation of the logit stochastic process as both $n, N \to \infty$. Notice that in Proposition 3, we are also able to establish the rate of convergence of the step–wise dynamic to the stochastic process. This is due to the simpler structure of the step–wise dynamic. In comparison, Theorem 6 is a weaker result as it only establishes convergence of the original logit dynamic and not the rate of convergence. Nevertheless, it suffices in establishing the conclusion we desire.

This theorem provides microfoundations to the continuum logit dynamic (3.4). It implies that over finite time horizons, the behavior of the logit stochastic process (3.16) or, equivalently, its smoothed version dynamic $\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))$ in $\mathscr{D}_\varphi^n$, is arbitrarily well approximated by the solution trajectory from the same initial state of the logit dynamic (3.4) in the continuum game φ provided the number of strategies and the number of players in $\mathscr{D}_\varphi^n$ are sufficiently high. Thus, the continuum logit dynamic may be interpreted as approximating evolution in situations where the underlying population game have a finite but large number of strategies and players and where the strategic behavior of players is governed by the logit stochastic process. This is important because the direct analysis of the stochastic process itself may be difficult. In such situations, we can make the problem significantly more tractable by focusing, instead, on the behavior of the continuum logit dynamic.

3.4 Medium Run Convergence in Supermodular Games

As an application of our deterministic approximation approach, we now consider supermodular games. Our objective is to show that over sufficiently large but finite time horizon, the continuum logit dynamic converges to the set of logit equilibria in supermodular games. For finite strategy supermodular games, Hofbauer and Sandholm (2007) establish results on convergence in such games for the class of stochastically perturbed best response dynamics.[14] But since we do not have a general theory of stochastic perturbation of best response in continuous strategy games, we confine ourselves to the logit dynamic.

To define supermodular games with a continuous strategy set, we introduce the notation $\bar{F}_\mu : \mathscr{S} \to [0,1]$ to denote the inverse cumulative distribution function of a probability measure $\mu \in \mathscr{P}(\mathscr{S})$. Thus, $\bar{F}_\mu(x) := \mu((x,1])$ for all $x \in \mathscr{S}$. Let $\preceq$ be a partial order on $\mathscr{P}(\mathscr{S})$ defined as follows:

$$[\mu \preceq \nu] \Longleftrightarrow [\bar{F}_\mu(x) \leq \bar{F}_\nu(x), \quad \text{for all } x \in \mathscr{S}]. \tag{3.18}$$

Therefore, under this partial order, a population state ν is "ranked higher" than another state μ if the proportion of agents playing higher strategies is greater under ν than μ. We then arrive at the following

[14] For other canonical evolutionary dynamics like the replicator dynamic, the Brown–von Neumann–Nash dynamic, payoff comparison dynamics and the best response dynamic, general results on convergence in supermodular games are not available.

definition of a large population supermodular game, which is an extension of the definition provided by Hofbauer and Sandholm (2007) to our setting of a continuous strategy set. Intuitively, the definition implies that in a supermodular game, higher strategy choices by other agents make a higher strategy more attractive for every agent. Such games are, therefore, characterized by strategic complementarities.

Definition 13 *A population game $\varphi : \mathscr{P}(\mathscr{S}) \to \mathscr{L}^\infty(\mathscr{S})$ is said to be a **continuum supermodular game on $\mathscr{S}$** if for all $\mu, \nu \in \mathscr{P}(\mathscr{S})$, the following condition is satisfied:*

$$[\mu \preceq \nu] \Longrightarrow [\varphi_y(\mu) - \varphi_x(\mu) \leq \varphi_y(\nu) - \varphi_x(\nu), \quad \textit{for all}\ 0 \leq x < y \leq 1]. \tag{3.19}$$

The partial order $\preceq$ used in (3.18) is the canonical notion of first order stochastic dominance (FSD). Thus Definition 13 states that if a population state ν stochastically dominates another population state μ, then for a supermodular game φ, the difference in switching from a lower ranked strategy x to a higher ranked strategy y is higher under ν than under μ. As an illustration, the game in Example 5 is a supermodular game provided the function $m(x)$ is strictly increasing. To check that the relevant payoff (3.2) in this example satisfies supermodularity, notice that $\varphi_y(\mu) - \varphi_x(\mu) = (m(y) - m(x))a(\mu) - (c(y) - c(x))$. Further, if $\mu \preceq \nu$, then $a(\mu) \leq a(\nu)$. Hence, with m increasing, it must then be that $\varphi_y(\mu) - \varphi_x(\mu) < \varphi_y(\nu) - \varphi_x(\nu)$ if $x < y$. Thus, (3.19) is satisfied.

Recall from (3.1) that we are considering population games with pairwise interaction. Thus, $\varphi_x(\mu) = \int_{\mathscr{S}} f(x, y)\mu(dx)$. Intuitively, therefore, if φ is supermodular, then that should have some implication on whether f is also supermodular. We show that is indeed the case. Following Topkis (1998), we define supermodularity of the pairwise interaction game in the standard manner of f satisfying increasing differences.

Definition 14 *The pairwise interaction game $f : \mathscr{S} \times \mathscr{S} \to \mathbb{R}$ is supermodular if for all $x, y \in \mathscr{S}$ such that $x < y$, $f(y, z) - f(x, z)$ is non–decreasing in $z \in \mathscr{S}$.*

We then obtain the following result on the equivalence of the supermodularity of φ and f. The result implies that our subsequent analysis of convergence of the logit dynamic in such games could have been stated in terms of either the population game φ or the underlying pairwise interaction game f. However, as our interest is in large population games, we find it more convenient to state our findings in terms of φ.

Proposition 4 *A population game $\varphi : \mathscr{P}(\mathscr{S}) \to \mathscr{L}^\infty(\mathscr{S})$ defined by (3.1) is a supermodular game if and only if the underlying pairwise interaction game $f : \mathscr{S} \times \mathscr{S} \to \mathbb{R}$ which generates φ is supermodular on $\mathscr{S}$.*

Proof 5 *See Appendix B.4.*

Our result on convergence of the logit dynamic in supermodular games will depend upon our key deterministic approximation result as presented in Theorem 6. Since such convergence is over finite time horizons, we follow Hofbauer and Sandholm (2007) and call it medium run convergence. Hofbauer and Sandholm (2007) also extend such convergence to an infinite time horizon by directly computing the stationary distribution of a stochastic process like (3.16). We do not pursue this long run analysis here as that is independent of the deterministic approximation approach to the logit dynamic.

We now consider a supermodular game φ with an underlying pairwise interaction structure f such as Example 5. Recall that $\mathrm{LE}_\eta(\varphi)$ is the set of logit equilibria in a continuum population game φ. For $\epsilon > 0$, let $\mathscr{O}^\epsilon(\mathrm{LE}_\eta(\varphi))$ be the open ball of radius ϵ under the total variation norm around $\mathrm{LE}_\eta(\varphi)$ in a supermodular game φ generated by f. Applying the methodology used to derive Theorem 6, we first discretize our supermodular game as in Definition 15. The following theorem establishes the relationship between a continuum supermodular game and its discrete version. As is to be expected, the two notions are equivalent. Before presenting the theorem, we state the definition of a discrete supermodular game.

Definition 15 *A population game $\mathscr{D}^n_\varphi : \Delta(S_n) \to \mathbb{R}^{2^n}$ as defined in Definition 11 is a supermodular game if it exhibits strategic complementarities, that is, for all $\boldsymbol{p}^n, \boldsymbol{q}^n \in \Delta(S_n)$, the following condition is satisfied:*

$$[\boldsymbol{p}^n \preceq \boldsymbol{q}^n] \Longrightarrow [\mathscr{D}^n_\varphi[\boldsymbol{p}^n](\alpha_{n,j+1}) - \mathscr{D}^n_\varphi[\boldsymbol{p}^n](\alpha_{n,j}) \leq \mathscr{D}^n_\varphi[\boldsymbol{q}^n](\alpha_{n,j+1}) - \mathscr{D}^n_\varphi[\boldsymbol{q}^n](\alpha_{n,j})], \tag{3.20}$$

for all $j = 0, 1, \ldots, 2^n - 2$.

In words, if $\mathbf{q}^n$ stochastically dominates $\mathbf{p}^n$, then for any strategy $\alpha_{n,j}$, the payoff advantage of $\alpha_{n,j+1}$ over $\alpha_{n,j}$ is greater at $\mathbf{q}^n$ than at $\mathbf{p}^n$. We can also write the strategic complementarity condition (3.20) equivalently by using partial derivatives as[15]

$$\frac{\partial \mathscr{D}^n_\varphi[\mathbf{p}^n](\alpha_{n,j+1})}{\partial \mathbf{p}^n_{i+1}} - \frac{\partial \mathscr{D}^n_\varphi[\mathbf{p}^n](\alpha_{n,j+1})}{\partial \mathbf{p}^n_i} \geq \frac{\partial \mathscr{D}^n_\varphi[\mathbf{p}^n](\alpha_{n,j})}{\partial \mathbf{p}^n_{i+1}} - \frac{\partial \mathscr{D}^n_\varphi[\mathbf{p}^n](\alpha_{n,j})}{\partial \mathbf{p}^n_i}, \tag{3.21}$$

for all $i, j = 0, 1, \ldots, 2^n - 2$. Thus, if a small of mass of players switches from strategy $\alpha_{n,i}$ to $\alpha_{n,i+1}$, then the improvement in the payoff of strategy $\alpha_{n,j+1}$ is greater than that of $\alpha_{n,j}$. The result we now seek is as follows.

Proposition 5 *Fix a weakly continuous population game $\varphi : \mathscr{P}(\mathscr{S}) \to \mathscr{L}^\infty(\mathscr{S})$. Then φ is a supermodular if and only if for every $n \geq 1$, the level-n discretization $\mathscr{D}^n_\varphi : \Delta(S_n) \to \mathbb{R}^{2^n}$ of the population game φ is a supermodular game.*

Proof 6 *See Appendix B.5.*

Recall now the step–wise approximation of φ, φ^n, as defined in (3.7) or (4.18). Also recall our comments about the relationship between the payoffs in $\mathscr{D}^n_\varphi$ and φ^n in the paragraph following Definition 11. We then obtain the following corollary from Proposition 5. The proof is obvious and is left to the reader.

Corollary 1 *Fix a population game $\varphi : \mathscr{P}(\mathscr{S}) \to \mathscr{L}^\infty(\mathscr{S})$. Then φ is a supermodular if and only if φ^n is a supermodular game for every $n \geq 1$.*

[15] See Section 3.3 in Hofbauer and Sandholm (2007) and Theorem 3.4.2 in Sandholm (2010).

Next, we introduce the discrete logit process $\mathbb{X}_n^N(k)$ as defined in (3.16) on the discrete supermodular game $\mathscr{D}_\varphi^n$, construct the interpolated logit stochastic process $\widehat{\mathbb{X}}_n^N(t)$ and smoothen the process as described in Definition 12 to obtain $\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))$. In addition, we also need the following definition of an *irreducible discrete supermodular game* from Hofbauer and Sandholm (2007). In stating this definition, we recall the finite strategy set $S_n = \{\alpha_{n,0}, \alpha_{n,1}, \ldots, \alpha_{n,2^n-1}\}$ over which we defined the discrete game $\mathscr{D}_\varphi^n$ in Definition 11. We also recall the notation $\alpha_{n,j}$ to denote the strategy $\frac{j}{2^n}$ in (4.15).

Definition 16 *Let $\widehat{S}_n = S_n \setminus \alpha_{n,2^n-1} = \{\alpha_{n,0}, \alpha_{n,1}, \ldots, \alpha_{n,2^n-2}\}$. For $n \geq 1$, the level-n discretization of a supermodular game φ is said to be irreducible if for all states $\boldsymbol{p}^n \in \Delta(S_n)$ and all $j \in \widehat{S}_n$, there exists at least one $i \in \widehat{S}_n$ such that (3.21) holds with strict inequality.*

Definition 16 is a mild strengthening of the strategic complementarity condition (3.21) for a discrete supermodular game. It requires that at every state $\mathbf{p}^n$, a movement from strategy j to $j+1$ improves the relative performance of at least one strategy. As an illustration, we show that the discrete version of the game in Example 5 does satisfy irreducibility.

Proposition 6 *For every $n \geq 1$, the discrete version $\mathscr{D}_\varphi^n$ of the game φ in Example 5 is irreducible.*

Proof 7 *See Appendix B.6.*

With this irreducibility condition, we can now establish our main result about supermodular games which, following Hofbauer and Sandholm (2007), we call a medium run convergence theorem. Before considering that theorem, we state the following lemma which will be useful to us.

Lemma 4 *Let φ be a population game as defined in (3.1) and $\mathscr{D}_\varphi^n$ be its discretization. Suppose $(\boldsymbol{p}_n^*)_{n\geq 1}$ is a sequence of logit equilibria of the discretized games $(\mathscr{D}_\varphi^n)_{n\geq 1}$. Recall from (3.14) the atomic measure $\mu_{\boldsymbol{p}_n^*}$ derived from $\boldsymbol{p}_n^*$ on $\mathscr{S}$ and the smoothing $\mathfrak{s}_n(\mu_{\boldsymbol{p}_n^*})$ of $\mu_{\boldsymbol{p}_n^*}$ (Definition 12). Then the accumulation points of the sequence $(\mathfrak{s}_n(\mu_{\boldsymbol{p}_n^*}))_{n\geq 1}$ under the total variation norm is non-empty. Furthermore, any such accumulation point is a logit equilibrium of the game φ.*

Proof 8 *See Appendix B.7.*

Lemma 4 is independent of supermodularity. It holds for population game of the form (3.1) that is generated by pairwise random matching. It is well known that the set of probability measures $\mathscr{P}(\mathscr{S})$ is compact under the weak topology. Hence, the sequence $\mathfrak{s}_n(\mu_{\mathbf{p}_n^*})$ that we obtain from the logit equilibria of the discretized game $\mathscr{D}_\varphi^n$ will have an accumulation point in $\mathscr{P}(\mathscr{S})$. We show that this accumulation point is a logit equilibrium of the original game φ. In addition, we also need to show that in case of such logit equilibria, convergence under the weak topology also implies convergence under the strong topology, or the topology induced by the total variation norm. This follows because as $n \to \infty$, the probability mass put by the logit equilibrium $\mathbf{p}_n^*$ on any strategy in the finite strategy set of $\mathscr{D}_\varphi^n$ goes to zero.

We now establish our medium run convergence theorem. Suppose φ is a continuum supermodular game and that its discretized form $\mathscr{D}_\varphi^n$ satisfies irreducibility. Consider the interpolated logit stochastic process on $\mathscr{D}_{\varphi}^{n}$ and let the number of strategies and players in $\mathscr{D}_{\varphi}^{n}$ become large. The theorem states that over sufficiently long but finite time horizons (hence the name “medium run”), the continuum logit dynamic spends most of its time near the set of logit equilibria of φ.

Theorem 7 *Suppose that the two player game $f: \mathscr{S} \times \mathscr{S} \to \mathbb{R}$ be Lipschitz continuous. Suppose the the population game $\varphi: \mathscr{P}(\mathscr{S}) \to \mathscr{L}^{\infty}(\mathscr{S})$ is supermodular game such that for every n, the level-n discretization of φ is irreducible. Then for every $\epsilon > 0$, there exists $T_\epsilon \geq 0$ such that for all $T \geq T_\epsilon$, $(\mu(t))_{t \geq 0} \in \mathscr{O}^{3\epsilon}(LE_{\eta}(\varphi))$ for all $T_\epsilon \leq t \leq T$.*

Proof 9 *See Appendix B.8.*

The proof of Theorem 7 proceeds in several steps. First, we discretize φ into $\mathscr{D}_{\varphi}^{n}$. Since φ is supermodular, Proposition 5 implies that $\mathscr{D}_{\varphi}^{n}$ is also supermodular. We then consider the interpolated logit stochastic process $\widehat{\mathbb{X}}_{n}^{N}(t)$, define it appropriately on $\Delta(S_n)$ and allow it to run in $\mathscr{D}_{\varphi}^{n}$. By Theorem 6, once we smoothen $\widehat{\mathbb{X}}_{n}^{N}(t)$ and $N, n \to \infty$, the behavior of the stochastic process in $\mathscr{D}_{\varphi}^{n}$ is well approximated by the continuum logit dynamic (3.4) on φ. Therefore, it is legitimate to analyze the convergence of the logit dynamic on φ by considering the convergence of the logit stochastic process on $\mathscr{D}_{\varphi}^{n}$. We do so by appealing to results in Hofbauer and Sandholm (2007) for finite strategy supermodular games and applying some of our earlier results.

Since the discretized or finite–strategy game $\mathscr{D}_{\varphi}^{n}$ is supermodular, Theorem 4.1 (part (iii)) in Hofbauer and Sandholm (2007) implies that as the number of players $N \to \infty$, the logit stochastic process converges to the set of logit equilibria of $\mathscr{D}_{\varphi}^{n}$. We then show that once we smoothen a logit equilibrium of $\mathscr{D}_{\varphi}^{n}$, we obtain a logit equilibrium of the step–wise approximation of φ, φ^n. Corollary 1 then implies that φ^n is also supermodular. Hence, the smoothed interpolated process $\mathfrak{s}_n(\widehat{\mathbb{X}}_n(t))$ lies near the set of logit equilibria of φ^n over sufficiently long but finite time horizons as $N \to \infty$. By the deterministic approximation result for the step–wise logit dynamic (Proposition 3), $\mathfrak{s}_n(\widehat{\mathbb{X}}_n(t))$ is itself close to the trajectory of the step–wise logit dynamic, $\mu^n(t)$, over such finite time horizons. Hence, over a finite time horizon, $\mu^n(t)$ converges to the set of logit equilibria of φ^n. Now, as $n \to \infty$, a logit equilibrium of φ^n approaches a logit equilibrium of φ under the strong topology. This follows from Lemma 4. Moreover, as $n \to \infty$, $\mu^n(t)$ itself is close to $\mu(t)$ by Proposition 2. Combining all these arguments, we conclude that as both $N, n \to \infty$, a trajectory $\mu(t)$ of the continuum logit dynamic converges to a logit equilibrium of the supermodular game φ over finite but sufficiently long time periods.

Corollary 2 *Fix $n, N \geq 1$. Suppose that the interpolated logit stochastic process $(\widehat{\mathbb{X}}_{n}^{N}(t))_{t \geq 0}$ has identical initial conditions $\mu_{\boldsymbol{p}^n}^{N} = \mu(0)$. Suppose $\varphi: \mathscr{P}(\mathscr{S}) \to \mathscr{L}^{\infty}(\mathscr{S})$ be a supermodular game such that for every n, the level-n discretization of φ is irreducible. Then for every $\epsilon > 0$, there exists $T_\epsilon \geq 0$ such that for all $T \geq T_\epsilon$, we have*

$$\lim_{n \to \infty} \lim_{N \to \infty} \mathbb{P}_{\mu_{\boldsymbol{p}^n}^{N}} \left(\mathfrak{s}_n(\widehat{\mathbb{X}}_{n}^{N}(t)) \in \mathscr{O}^{\epsilon}(LE_{\eta}(\varphi)) \text{ for all } T_\epsilon \leq t \leq T \right) = 1.$$

Proof 10 *The proof follows from Theorem 7.*

3.5 Conclusion

The deterministic approach to evolutionary game theoretic dynamics seeks to provide mathematical foundations to such dynamics as the limiting outcome of a stochastic process as the number of players become large. In this paper, we have adopted such a deterministic approach to the logit dynamic in a large population game with a continuous strategy set. We first construct a step–wise approximation of such a game and show that the logit dynamic in that approximate game is close to the dynamic in the original game. We then introduce a discrete approximation of the original game with a finite number of strategies and players. Over finite time horizons, the logit stochastic process in this discrete game approximates the behavior of the step–wise logit dynamic, which itself approximates the original continuum logit dynamic. This gives us our deterministic approximation result. The continuum logit dynamic is an approximation of the logit stochastic process of a game in which both the number of players and strategies are finite but large.

We then apply our approach to large population supermodular games with a continuous strategy set. If the discrete version of the supermodular game satisfies a condition called irreducibility, then over time horizons that are finite but sufficiently long, the continuum logit dynamic converges to the set of logit equilibria of the supermodular game. To establish this result, we consider the logit stochastic process in the discrete version of the game. Results from finite strategy supermodular games imply that this process converges. Our deterministic approximation result then show that the logit dynamic must also converge.

This paper is limited to the logit dynamic. An obvious research question for the future is to extend the deterministic approach to other well–known continuous strategy evolutionary dynamics. For example, should it prove feasible to establish such a result for the larger class on continuous strategy perturbed best response dynamics, then it may be possible to extend our result on convergence of the logit dynamic in supermodular games to such dynamics as well.

CHAPTER

4

EQUILIBRIUM SELECTION VIA STOCHASTIC EVOLUTION IN CONTINUUM POTENTIAL GAMES

4.1 Introduction

Evolutionary game theory is broadly divided into two parts; deterministic evolution and stochastic evolution (Sandholm (2010)). In such models, there exists a large number of players in a population who revise their strategies according to an underlying stochastic process. For time horizons that are finite but sufficiently long (medium run), the change in the strategy distribution (the population state) due to the stochastic process can be well approximated by a deterministic mean dynamic called an evolutionary dynamic (Benaïm and Weibull (2003)). Well known examples of such dynamics are the replicator dynamic, the logit dynamic, the BNN dynamic and the pairwise comparison dynamic.[1] Such evolutionary dynamics, being ordinary differential equations, are relatively more convenient to analyze than the stochastic process directly. However, over infinitely long time horizons, the deterministic approximation is no longer valid. Therefore, as perhaps first recognized by Foster and Young (1990), to address questions of long run stability of Nash equilibria under evolution, one must analyze the stochastic process directly.

An important motivation behind studying stochastic evolution is to generate exact equilibrium

[1] Sandholm (2010) provides a comprehensive review of the literature on such evolutionary dynamics with a finite number of strategies. Also see Chapter 10 of Sandholm (2010) for a discussion of the deterministic approximation of a stochastic process with an evolutionary dynamic.

selection results. We wish to predict to which Nash equilibrium the population would converge to irrespective of the initial state. Most deterministic dynamics do not provide such history–independent results, although there are exceptions like the sampling best response dynamic (Oyama et al. (2015)). In contrast, stochastic evolutionary analysis seeks to characterize the stationary distribution of the stochastic process to understand long run stability of a Nash equilibrium. If it can be shown that the stationary distribution puts probability one on a particular equilibrium, then we obtain an exact prediction that the population would spend almost all the time at that equilibrium, although only in the long run. This is independent of the history of the game. By now, there is a significant literature on the stochastic approach to evolutionary game theory. But as far as we know, this entire literature has been on games with a finite set of strategies.[2] In fact, starting with Kandori et al. (1993) and Young (1993), much of this literature has been on symmetric games with two strategies with a common conclusion being that the risk dominant equilibrium is stochastically stable in such games.[3] Other finite strategy games for which results on stochastic stability are available are three–strategy games (Sandholm and Staudigl (2016)) and potential games (Blume (1993), Blume (2018), Alós-Ferrer and Netzer (2010), Section 11.5 of Sandholm (2010)).

In this paper, we consider stochastic evolution in population games with a continuum of strategies. Apart from raising interesting mathematical and technical questions, the evolutionary analysis of continuous strategy games is also important from the point of view of applications. Various problems of substantive economic interest like oligopoly, public goods and common resources where evolutionary analysis can provide insights are more naturally modeled as continuous strategy games (Cheung and Lahkar (2018), Lahkar and Mukherjee (2021)). By now a significant literature on the canonical deterministic evolutionary dynamics in continuous strategy games has developed. These are the replicator dynamic and the more general class of imitative dynamics (Oechssler and Riedel (2001),Oechssler and Riedel (2002), Cheung (2016)), the BNN dynamic (Hofbauer et al. (2009)), the pairwise comparison dynamic (Cheung (2014)), the logit dynamic (Perkins and Leslie (2014), Lahkar and Riedel (2015)) and the more general class of perturbed best response dynamics to which the logit dynamic belongs (Lahkar et al. (2022a)).[4] But as far as we know, the question of stochastic evolution under a continuum of strategies has not been addressed. This paper seeks to contribute towards filling that gap in the literature.

We focus our analysis on large population potential games with a continuum of strategies. Potential games are games in which strategic interests can be summarized using a real valued function called the potential function. Monderer and Shapley (1996) originally defined these games in the context of a finite number of players. Since then, the concept has been extended to large population games with a finite number of strategies (Sandholm (2010)) as well as a continuum of strategies (Lahkar and Riedel (2015), Cheung and Lahkar (2018)). Potential games are of great importance in evolutionary game theory due to their attractive convergence properties. Local maximizers of the potential function

[2] See Chapters 11 and 12 of Sandholm (2010) for a review of this literature.

[3] Other stochastic models with two–strategy games include Binmore et al. (1995), Binmore and Samuelson (1997), Blume (2003), Sandholm (2007), Sandholm (2012), Staudigl (2012).

[4] Also see Lahkar et al. (2023) on the derivation of the continuous strategy logit dynamic as a deterministic approximation of an underlying stochastic process.

are Nash equilibria of the game. All the canonical evolutionary dynamics we have mentioned earlier converge to such local maximizers. Hence, Nash equilibria that are local maximizers of the potential function emerge as evolutionarily stable under such dynamics (Sandholm (2001)).[5]

Therefore, Nash equilibrium prediction is particularly robust in potential games. But if there are multiple such Nash equilibria, we run into the equilibrium selection problem characteristic of such medium run analysis. To which of those locally potential maximizing Nash equilibrium the population converges to depends upon the initial state. Potential games are, however, amenable to stochastic analysis. The reason is that such games satisfy a condition called reversibility which makes it convenient to characterize their stationary distribution. In fact, as noted earlier, this is the reason why finite strategy potential games have been a focus of stochastic evolutionary analysis. Results in such models show that the Nash equilibrium which is the global maximizer of the potential function is the unique stochastically stable equilibrium.

We aim to establish such an exact equilibrium selection for continuous strategy potential games. Again, the question is not merely of technical interest. There are significant examples of continuous strategy potential games with multiple Nash equilibria. These include games with strategic complementarities (Bandhu and Lahkar (2022)) and doubly symmetric games (Hofbauer et al. (2009)), which are extensions of pure coordination games to the continuous strategy case. An exact equilibrium selection result in such games would provide answers to meaningful questions like whether in the long run, it is the Pareto optimal Nash equilibrium that gets selected. If not, then under what circumstances will the Pareto optimal equilibrium be selected. To make our problem tractable, we consider continuous strategy games with an underlying pairwise interaction structure. Such a structure includes doubly symmetric games as well as certain types of potential games with strategic complementarities called aggregative potential games (Lahkar (2017)). In such games, a higher aggregate strategy creates incentives for individual players to play a higher strategy.

The key to our analysis is a sequence of approximation that enables us to move from the continuous strategy potential game to a potential game with a finite set of strategies and a finite number of players. Our motivation behind these approximations is to apply existing results of large population potential games with a finite number of strategies. Accordingly, our first approximation is a large population step approximation where we replace our continuous strategy payoff function with a continuous strategy step approximation. Next, we replace the step approximation with a large population–finite strategy game where the payoff of each strategy corresponds to the payoff at each step of the step approximation. Our final approximation is a finite player–finite strategy game which we construct from the large population finite strategy approximation by ruling out the possibility of self–matching. As our original continuous strategy game is a potential game, it follows that all the three approximations we construct are also potential games,.

We then follow Chapters 11 and 12 in Sandholm (2010) and establish stochastic stability results in our finite player–finite strategy and large population finite–strategy approximations. We apply a class of stochastic processes known as a direct exponential stochastic process on the finite player–finite

[5] In the case of the logit dynamic, convergence is to a perturbed version of the Nash equilibrium.

strategy approximation. A direct exponential stochastic process is a noisy strategy revision protocol wherein agents switch strategies based on how attractive a new strategy is or how dissatisfied they are with their current strategy. The most well–known such revision protocol is the logit protocol (Blume (1993), Blume (2003)) wherein agents switch to the best response with probability nearly one. The noise parameter makes the stochastic process irreducible. Under the additional assumption that agents employ clever payoff evaluation which recognizes that a change of strategy by oneself in a finite population would change the social state, we also obtain a reversible stochastic process with a unique stationary distribution.

While this result is for the finite player–finite strategy approximation, we can take the limit as the number of players become large. This characterizes the stationary distribution in the large population–finite strategy approximation for a fixed noise parameter in the stochastic process. The main conclusion we obtain is that this stationary distribution puts all its weight on the global maximizer of the entropy adjusted potential function of the large population–finite strategy approximation. This is an equilibrium selection result for this game. Maximizers of the entropy adjusted potential function are perturbations of Nash equilibria (Hofbauer and Sandholm (2007)). For any fixed level of noise in the direct exponential stochastic process, the stationary distribution in the large population–finite strategy approximation chooses that perturbed equilibrium that globally maximizes the entropy adjusted potential function.

The above result is akin to the conclusion in Sandholm (2010) on the stationary distribution of a large population potential game with a finite number of strategies. The main contribution of this paper lies in extending this result to a large population potential game with a continuum of strategies. For this, we now need to allow the number of strategies in the stochastic process to also go to infinity. Hence, keeping the noise parameter fixed, we take the double limit by first allowing the number of players and then the number of strategies to become large. This gives us the stationary distribution of a direct exponential stochastic process for a fixed noise parameter in a large population continuous strategy potential game. As is to be expected, this stationary distribution is concentrated on the perturbed equilibrium that is the global maximizer of the continuous strategy version of entropy adjusted potential function. Our main interest, however, is in selection among Nash equilibria and not perturbed equilibria. For this, we establish the final and main result of this paper. This involves taking the triple limit of the stochastic process, where first the number of agents go to infinity, then the number of strategies go to infinity and finally, the noise parameter goes to zero. The stationary distribution then puts probability one on the global maximizer of the potential function of our original large population continuous strategy game.

We, therefore, obtain our desired equilibrium selection result. In a large population continuous strategy potential game, as the noise parameter goes to zero, a direct exponential stochastic process selects that Nash equilibrium that globally maximizes the potential function of the game. The entire mass of the stationary distribution is concentrated on that global maximizer. This is irrespective of the initial state of the stochastic process. The interpretation of this result is that as the time horizon becomes infinitely long, the proportion of time that the process spends on the globally potential maximizing Nash equilibrium goes to one.

Superficially, it might seem that our result is a simply extends finite strategy results by taking the limit as the number of strategies increase. The extension, however, is not trivial and we need to resolve significant technical complications. We need to choose the correct method of approximating our continuous strategy game and the fact that these approximations use the underlying pairwise structure is the reason why we confine ourselves to games with pairwise interaction. The construction of the potential function for the large population step approximation game also involves significant challenges. This step approximation is required in taking the limit as the number of strategies become infinite. We first apply the result for the large population finite strategy game to the large population step approximation. Then, as the number of strategies increase, the size of each step decreases and the step approximation approaches our original continuous strategy game giving us our desired result. Such analytical techniques are new to this paper.

We apply our main result to the two motivating examples of large population aggregative potential games with multiple Nash equilibria and doubly symmetric games. The structure of equilibria in such aggregative potential games is like that of two–strategy symmetric games with multiple equilibria. There are three Nash equilibria, two of which are local maximizers of the potential function. One of them is Pareto optimal and the other Pareto inferior among all Nash equilibria. Whichever of these two Nash equilibria has the larger basin of attraction with respect to the best response strategy revision protocol emerges as the global maximizer of the potential function and, therefore, the one that is selected by stochastic evolution. Thus, if the Pareto optimal equilibrium has a smaller basin of attraction, and there are circumstances when it does, then it is the Pareto inferior equilibrium that is stochastically stable. The intuition behind this result is similar to the risk dominance argument in two–strategy symmetric games. In contrast, in doubly symmetric games, the Pareto optimal equilibrium is always stochastically stable. This is because the potential function of such games is a rescaled version of the aggregate payoff function due to which, its global maximizer must be Pareto optimal.

The rest of the paper is as follows. Section 6.2 introduces the preliminaries of the model including the classes of aggregative potential games and doubly symmetric games. In Section 4.3, we provide specific examples of these two classes of games as motivations for our entire analysis. In particular, we identify conditions under which the global maximizer of the potential function in aggregative potential games is not the Pareto optimal equilibrium. Section 4.4 describes the three approximations of our continuous strategy game that we require for the analysis. Our main results on stochastic evolution and equilibrium selection are in Section 4.5. We conclude in Section 4.6 with a discussion on applying our results to aggregative potential games and doubly symmetric games. Proofs are in the Appendix.

4.2 Preliminaries

We consider a large population of players with mass normalized to unity. Let $\mathscr{S} := [0,1]$ denote the strategy space of the players in the population. The strategy space $\mathscr{S}$ is endowed with the usual Borel sigma-algebra $\mathscr{B}(\mathscr{S})$. Let $\mathscr{P}(\mathscr{S})$ denote the collection of all probability measures on $\mathscr{S}$. A probability measure $\mu \in \mathscr{P}(\mathscr{S})$ denotes the state of the underlying population. Thus, for every $A \in \mathscr{B}(\mathscr{S})$, $\mu(A)$ denotes the proportion of players in the population exercising strategies in A. Let $\mathbb{L}^{\infty}(\mathscr{S})$ be the collection of bounded measurable functions on $\mathscr{S}$ equipped with the sup-norm metric.[6] A *population game* is defined as a weakly continuous mapping $\pi : \mathscr{P}(\mathscr{S}) \to \mathbb{L}^\infty(\mathscr{S})$ such that $\pi_x(\mu)$ denotes the payoff to an agent exercising strategy $x \in \mathscr{S}$ at population state $\mu \in \mathscr{P}(\mathscr{S})$.[7]

A population state μ° is said to be a *Nash equilibrium* of the game π if $\pi_x(\mu^\circ) \geqslant \pi_y(\mu^\circ)$ for all $x \in \operatorname{supp}(\mu^\circ)$ and all $y \in \mathscr{S}$.[8] We shall assume throughout, that the population game π is generated via a continuous pairwise interaction game $f : \mathscr{S} \times \mathscr{S} \to \mathbb{R}$ in the sense that for all $\mu \in \mathscr{P}(\mathscr{S})$,

$$\pi_x(\mu) := \int_{\mathscr{S}} f(x,y)\mu(dy), \quad \text{for all } x \in \mathscr{S}. \tag{4.1}$$

In this paper, we focus on equilibrium selection in continuum strategy potential games.[9] In order to define such games, we will require the notion of the Fréchet derivative, which is a generalization of the usual derivative to Banach spaces. Let V and W be normed linear spaces. An operator $\phi : V \to W$ is Fréchet differentiable at x if there exists a bounded linear operator $A_x : V \to W$ such that

$$\lim_{\|h\|_V \to 0} \frac{\|\phi(x+h) - \phi(x) - A_x(h)\|_W}{\|h\|_V} = 0.$$

Thus, if A_x denotes the the Fréchet derivative of ϕ at $x \in V$, then $\phi(x+h) = \phi(x) + A_x(h) + \mathbf{o}(\|h\|_V)$ for all h near zero in V. Let $\mathscr{M}(\mathscr{S})$ be the collection of all signed measures on $\mathscr{S}$ equipped with the variational norm.[10] In this paper, we shall consider Fréchet differentiability of functions $\phi : \mathscr{M}(\mathscr{S}) \to \mathbb{R}$. We denote the Fréchet derivative of such a function by $\mathbf{D}_{\mathscr{F}}\phi$. If ϕ is Fréchet differentiable, we will assume that for every $\mu \in \mathscr{M}(\mathscr{S})$, there exists a bounded measurable map $\nabla_{\mathscr{F}}\phi(\mu) \in \mathbb{L}^\infty(\mathscr{S})$ such that,

$$\mathbf{D}_{\mathscr{F}}\phi(\mu)(\nu) := \int_{\mathscr{S}} \nabla_{\mathscr{F}}\phi(\mu)\, d\nu = \int_{\mathscr{S}} \nabla_{\mathscr{F}}\phi_x(\mu)\nu(dx) =: \langle \nabla_{\mathscr{F}}\phi(\mu), \nu \rangle, \quad \text{for all } \nu \in \mathscr{M}(\mathscr{S}). \tag{4.2}$$

We then call the map $\nabla_{\mathscr{F}}\phi : \mathscr{M}(\mathscr{S}) \to \mathbb{L}^\infty(\mathscr{S})$ the *Fréchet gradient* of ϕ. We now proceed with the definition of continuum strategy potential games (Lahkar and Riedel (2015), Cheung and Lahkar (2018)).

Definition 17 *A population game $\pi : \mathscr{P}(S) \to \mathbb{L}^\infty(\mathscr{S})$ is said to be a potential game if there exists a Fréchet differentiable mapping $\varphi : \mathscr{M}(\mathscr{S}) \to \mathbb{R}$ such that*

$$\nabla_{\mathscr{F}}\varphi(\mu) = \pi(\mu), \quad \textit{for all } \mu \in \mathscr{P}(\mathscr{S}). \tag{4.3}$$

In other words, a population game π is said to be a potential game if there exists a Fréchet differentiable

[6] The sup-norm of a function $f \in \mathbb{L}^\infty(\mathscr{S})$ is defined as $\|f\|_\infty := \sup_{x \in \mathscr{S}} |f(x)|$.

[7] Weakly continuous implies continuous with respect to the weak topology. The weak topology on $\mathscr{P}(\mathscr{S})$ is the topology induced by the following notion of convergence: $\mu_n \to^w \mu$ if and only if $\int_{\mathscr{S}} f\, d\mu_n \to \int_{\mathscr{S}} f\, d\mu$ for all continuous maps $f : \mathscr{S} \to \mathbb{R}$.

[8] Here $\operatorname{supp}(\mu)$ denotes the support of the probability measure μ.

[9] See, for example, Cheung (2014), Lahkar and Riedel (2015), Lahkar et al. (2022a) for an exposition on continuum strategy potential games.

[10] The variational norm or the total variation norm on $\mathscr{P}(\mathscr{S})$ is such that

$$\|\mu - \nu\|_{\mathbf{TV}} := 2 \sup_{A \in \mathscr{B}(\mathscr{S})} |\mu(A) - \nu(A)|, \quad \text{for all } \mu, \nu \in \mathscr{P}(\mathscr{S}).$$

map φ such that its Fréchet gradient $\nabla_{\mathscr{F}}\varphi$ is identified by π. We then call φ the potential function corresponding to the potential game π. We now discuss two classes of continuous strategy potential games which will be relevant for illustrating the stochastic approach to evolution. These are aggregative potential games and doubly symmetric games. In particular, in both these classes of games, the stochastic approach to evolution can generate exact equilibrium selection even though the deterministic evolution may not. In Section 4.3, we will present examples where that possibility may arise.

4.2.1 Aggregative Potential Games

The first class of examples is the class of aggregative potential games. Suppose that the underlying pairwise structure takes the form $f(x,y)=x(a+by)-c(x)$, where $a,b\in\mathbb{R}$ are constants such that $b\neq 0$ and $c:\mathscr{S}\to\mathbb{R}_+$ is a function that describes the cost of playing a strategy $x\in\mathscr{S}$. Let

$$A(\mu)=\int_{\mathscr{S}} x\mu(dx) \tag{4.4}$$

be the *aggregate strategy* in the population at state μ. Notice that as $\mathscr{S}=[0,1]$, the aggregate strategy $A(\mu)\in[0,1]$.

Applying (4.1), we can then generate the population game with payoffs

$$\begin{aligned}\pi_x(\mu)&=\int_{\mathscr{S}} f(x,y)\mu(dy)\\&=\int_{\mathscr{S}}(x(a+by)-c(x))\mu(dy)\\&=x(a+bA(\mu))-c(x).\end{aligned} \tag{4.5}$$

The population game π defined by (4.5) is an aggregative game as the payoff of an individual agent in this game depends only upon that agent's strategy and the aggregate strategy of the population (Corchon (1994)). In addition, Proposition 1 in Cheung and Lahkar (2018) establish that an aggregative game of the form (4.5) is a potential game with potential function

$$\begin{aligned}\varphi(\mu)&=\int_0^{A(\mu)}(a+bz)dz-\int_{\mathscr{S}} c(x)\mu(dx)\\&=aA(\mu)+\frac{b(A(\mu))^2}{2}-\int_{\mathscr{S}} c(x)\mu(dx).\end{aligned} \tag{4.6}$$

Hence, we refer to a population game π of the form (4.5) as an *aggregative potential game.*[11]

In this paper, we are particularly interested in the case where $b>0$ in (4.5). Under that condition, we may interpret (4.5) as a pairwise model of coordination. For example, it may represent a model of

[11]The result in Cheung and Lahkar (2018) is more general. It applies to aggregative games of the form $\pi_x(\mu)=x\beta(A(\mu))-c(x)$, where β may be a non–linear function of $A(\mu)$. In the current paper, however, we confine ourselves to linear functions of $\beta(A(\mu))=a+bA(\mu)$ in order to generate (4.5) through a pairwise interaction structure.

education or investment choice. The benefit that an agent receives from a particular strategy $x \in \mathscr{S}$ is increasing in the aggregate strategy $A(\mu)$. Hence, such models are also characterized by strategic complementarities in the sense that a higher aggregate strategy makes it more attractive for an individual agent to play a higher strategy x. Models of strategic complementarities typically generate multiple Nash equilibria which may be Pareto ranked. The stochastic approach of this paper is then particularly relevant in addressing questions such as whether in the long run, it is the Pareto equilibrium that gets selected. The example in Section 4.3 will discuss that question in greater detail.

4.2.2 Doubly Symmetric Pairwise Interaction Game

Consider a pairwise interaction game $f : \mathscr{S} \times \mathscr{S} \to \mathbb{R}$. The game f is said to be doubly symmetric if $f(x, y) = f(y, x)$ for all $x, y \in \mathscr{S}$ (Cheung (2014)). The payoff to $x \in \mathscr{S}$ in the resulting population game $\pi : \mathscr{P}(\mathscr{S}) \to \mathbb{L}^{\infty}(\mathscr{S})$ at the population state $\mu \in \mathscr{P}(\mathscr{S})$ is then,

$$\pi_x(\mu) = \int_{\mathscr{S}} f(x, y) dy, \quad \text{for all } x \in \mathscr{S}. \tag{4.7}$$

Then, it is well known (see for eg. Cheung (2014)) that π is a potential game with potential function φ given by

$$\begin{aligned} \varphi(\mu) &= \frac{1}{2} \int_{\mathscr{S}} \int_{\mathscr{S}} f(x, y) \mu(dy) \mu(dx) \\ &= \frac{1}{2} \int_{\mathscr{S}} \pi_x(\mu) \mu(dx) \\ &= \frac{1}{2} \bar{\pi}(\mu), \end{aligned} \tag{4.8}$$

where $\bar{\pi}(\mu) = \int_{\mathscr{S}} \pi_x(\mu) \mu(dx)$ is the average or, equivalently, aggregate payoff in the population at state μ.

4.3 Deterministic versus Stochastic Evolution: A Discussion

We now present examples of the two classes of potential games we discussed in Section 6.2. These are aggregative potential games and doubly symmetric games. We will use these examples to illustrate insights we may obtain from stochastic evolution particularly in comparison to deterministic evolution. The deterministic approach is indeed the more usual approach to evolutionary game theory. It relies on ordinary differential equations called evolutionary dynamics to track the change in the population state. But as noted in the Introduction, deterministic dynamics may not be appropriate to model evolution over infinitely long time horizons. Nor do such dynamics provide history–independent predictions in the presence of multiple Nash equilibria. The examples we present will have multiple equilibria. Hence, the stochastic approach will be useful in obtaining exact equilibrium selection in the long run in these games. The first such example is of an aggregative potential game introduced in Section 4.2.1.

Example 6 *Consider aggregative potential games introduced in Section 4.2.1. Let the cost function be $c(x) = dx + kx^2$, where $d, k > 0$. Two important features of this cost function are that it is strictly convex and that $c'(0) = d > 0$. With this cost function, we can assume, without loss of generality, that $a = 0$ in (4.5) as the effect of a can be subsumed under the parameter d. Therefore, let the pairwise interaction payoff in this game be $f(x, y) = bxy - (dx + kx^2)$, where $b > 0$. Hence, applying (4.5), we obtain the population game π with payoff*

$$\pi_x(\mu) = bxA(\mu) - (dx + kx^2). \tag{4.9}$$

Therefore, the potential function (4.6) for this game takes the form

$$\varphi(\mu) = \frac{b}{2}(A(\mu))^2 - \int_{\mathscr{S}} (dx + kx^2)\mu(dx). \tag{4.10}$$

To calculate Nash equilibria of this game, we write (4.9) as $bx\alpha - (dx + kx^2)$. For any given value of α, we maximize this function over $x \in \mathscr{S} = [0, 1]$. Clearly, this maximizer is

$$BR(\alpha) = \max\left\{0, \min\left\{\frac{b\alpha - d}{2k}, 1\right\}\right\}. \tag{4.11}$$

The interpretation of $BR(\alpha)$ is that it is the unique best response in the population game (4.9) to every state μ such that $A(\mu) = \alpha$.[12] *Notice that because $d > 0$, $BR(0) = 0$. Intuitively, this happens because the marginal cost at 0, $c'(0) = d > 0$. On the other hand, the benefit that an agent obtains from playing a strategy x in (4.9) is $bx\alpha$. Hence, the marginal benefit at $\alpha = 0$ is 0 causing $BR(0) = 0$. In fact, even for α sufficiently close to 0, we would have $BR(\alpha) = 0$.*

We now make another assumption, which is $b - d > 2k$. This assumption implies that $BR(1) = 1$ and even for α sufficiently close to 1, $BR(\alpha) = 1$. Figure 4.1 illustrates the general pattern of such a best response function for suitable parameter values. In both panels, we have $\{b, d\} = \{12, 4\}$. But in the left panel, $k = 3$ while in the right panel $k = 1$. In either case, the assumptions $d > 0$ and $b - d > 2k$ are satisfied. Hence, in both panels, the best response is flat in the vicinity of $\alpha = 0$ and $\alpha = 1$. But otherwise, $BR(\alpha)$ is strictly increasing, which indicates strategic complementarities in our model and which arises from $b > 0$.

Proposition 3.1 in Lahkar (2019) presents a method to characterize Nash equilibria of an aggregative game such as (4.9). Such Nash equilibria are related to solutions of

$$BR(\alpha) = \alpha. \tag{4.12}$$

The intuition is that at a Nash equilibrium, players playing their unique best response to the aggregate strategy level α should leave that aggregate strategy level unchanged. If α^* is a solution to (4.12), then the corresponding Nash equilibrium is one where all agents play α^*. We denote this monomorphic state by the Dirac distribution δ_{α^*}.

Applying (4.12) to the best response function (4.11), we obtain three solutions in our example. They

[12] The strict convexity of the cost function $c(x)$ plays a role in (4.11). It ensures that the payoff function $bx\alpha - c(x)$ is strictly convex and, therefore, has a unique maximizer $BR(\alpha)$ for every $\alpha \in [0, 1]$.

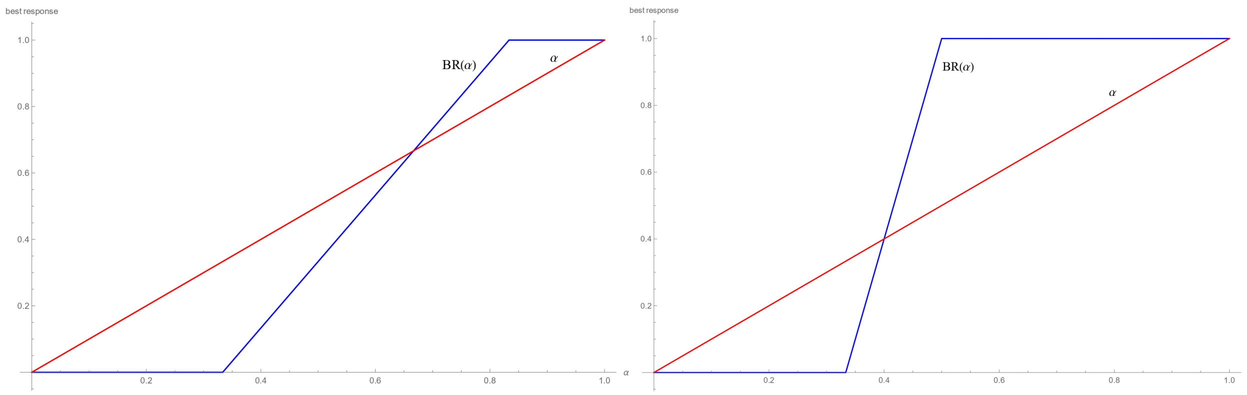

Figure 4.1: In both panels, the blue curve is the best response function $BR(\alpha)$ defined by (4.11) and α is the aggregate strategy in the aggregative potential game (4.9). The parameter values are $\{b,d,k\}=\{12,4,3\}$ in the left panel and $\{b,d,k\}=\{12,4,1\}$ in the right panel. In the left panel, we obtain three intersections, 0, $\frac{2}{3}$ and 1. Notice that this middle intersection $\frac{2}{3}>\frac{1}{2}$ because $b<2(d+k)$ in this panel (part 2 of Proposition 7). In the right panel, the three intersections are 0, $\frac{2}{5}$ and 1. Hence, the middle intersection $\frac{2}{5}<\frac{1}{2}$ as $b>2(d+k)$ in this panel.

are 0, $\frac{d}{b-2k}$ *and* 1. *The solutions* 0 *and* 1 *arise because, as we have already noted,* $B(0)=0$ *and* $B(1)=1$ *by our assumptions. The additional solution arises if from (4.11), we solve* $\frac{b\alpha-d}{2k}=\alpha$, *which gives* $\alpha^*=\frac{d}{b-2k}$. *To be a meaningful solution, we require* $0<\frac{d}{b-2k}<1$, *which follows from our assumptions that* $d>0$ *and* $b-d>2k$. *Figure 4.1 illustrates the solutions to (4.12) for the two sets of parameter values mentioned earlier. In both panels,* 0 *and* 1 *are points of intersection. In addition, there is also an intersection in the middle, which is* $\frac{d}{b-2k}=\frac{2}{3}$ *in the left panel and* $\frac{2}{5}$ *in the right panel.*

Therefore, the aggregative potential game π *defined by (4.9) has three Nash equilibria. We characterize these equilibria in the following proposition. The proposition also identifies the Pareto efficient Nash equilibrium in this game. As all Nash equilibria are in monomorphic states, all agents in a Nash equilibrium will have identical payoffs. The Pareto efficient Nash equilibrium is then the one with the highest equilibrium payoff.*

Proposition 7 *Consider the aggregative potential game* π *defined by (4.9). Suppose* $b,d,k>0$ *and* $b-d>2k$. *Then the following conclusions hold true.*

1. *The set of Nash equilibria of this game is*

$$NE(\pi)=\left\{\delta_0,\delta_{\frac{d}{b-2k}},\delta_1\right\}. \tag{4.13}$$

   The aggregate strategy level at these three Nash equilibria are 0, $\frac{d}{b-2k}$ *and* 1 *respectively.*

2. *If* $b<2(d+k)$, *then* $\frac{d}{b-2k}>\frac{1}{2}$. *If* $b>2(d+k)$, *then* $\frac{d}{b-2k}<\frac{1}{2}$.

3. *The Pareto efficient Nash equilibrium is* δ_1.

Proof 11 *See Appendix C.1.*

The first property in Proposition 7 arises from the fact that (4.12) has three solutions. Hence, in Figure 4.1, the set of Nash equilibria in the left panel is $\left\{\delta_0, \delta_{\frac{2}{3}}, \delta_1\right\}$ while in the right panel, that set is $\left\{\delta_0, \delta_{\frac{2}{5}}, \delta_1\right\}$. The second property follows from straightforward calculation. This property will be relevant in understanding whether a stochastic evolutionary process in an aggregative potential game will converge to a Pareto efficient Nash equilibrium. We observe the implication of this property on the value of the middle intersections in Figure 4.1. The third property in the proposition characterizes the Pareto efficient Nash equilibrium. Given strategic complementarities in our example, it is not surprising that δ_1 is the Pareto efficient equilibrium. A Nash equilibrium with a higher strategy will have a higher payoff. This is true, for example, in both panels of Figure 4.1.

We now present a deterministic evolutionary analysis of Example 6. This analysis will establish that multiple Nash equilibria in this example are locally stable under evolutionary dynamics. We will then contrast this finding with that of the stochastic approach once we complete our analysis of stochastic evolution in the following sections. Our argument will rely on the following result about the potential function (4.10) that characterizes its maximizers.

Proposition 8 *Consider the aggregative potential game π defined by (4.9). Suppose $b, d, k > 0$ and $b - d > 2k$. The following conclusions hold.*

1. *The potential function φ defined by (4.10) has two local maximizers in $\mathscr{P}(\mathscr{S})$ with respect to the weak topology. They are the two Nash equilibria, δ_0 and δ_1 characterized in (4.13).*

2. *Of these, δ_1 is the global maximizer of (4.10) in $\mathscr{P}(\mathscr{S})$ if $b > 2(d+k)$ or, equivalently, if $\frac{d}{b-2k} < \frac{1}{2}$. On the other hand, δ_0 is the global maximizer if $b < 2(d+k)$ or, equivalently, if $\frac{d}{b-2k} > \frac{1}{2}$.*

3. *The other Nash equilibrium characterized in (4.13), $\delta_{\frac{d}{b-2k}}$ is not a local maximizer of the potential function (4.10).*

Proof 12 *See Appendix C.2.*

Thus, two of the three Nash equilibria characterized in (4.13), δ_0 and δ_1, are local maximizers of the potential function (4.10). Part 2 of this proposition identifies the condition under which one of these two local maximizers is the global maximizer of the potential function. Recall from part 3 of Proposition 7 that δ_1 is the Pareto efficient Nash equilibrium of the game (4.9). Hence, part 2 of Proposition 8 shows that the Pareto efficient Nash equilibrium need not always be global maximizer of the potential function. In fact, by comparing part 2 of Propositions 7 and 8, we find that the Pareto efficient Nash equilibrium is the global maximizer only if the aggregate strategy level of the third Nash equilibrium, $\frac{d}{b-2k} < \frac{1}{2}$. Again, this observation will be helpful in understanding the intuition behind our stochastic evolutionary analysis.

The fact that certain Nash equilibria of the game (4.9) are local maximizers of its potential function is not surprising. It is well known that in large population potential games, all local maximizers of the potential function are Nash equilibria. Further, all the canonical evolutionary dynamics mentioned in the Introduction converge to such local maximizers (Sandholm (2001), Cheung and Lahkar (2018)).[13] Therefore, by part 1 of Proposition 8, both δ_0 and δ_1 would be stable under the standard evolutionary dynamics in the game (4.9). The third Nash equilibrium, $\delta_{\frac{d}{b-2k}}$, which is not a local maximizer, would be unstable.

Whether a deterministic dynamic converges to δ_0 or δ_1 would depend upon the initial state. Figure 4.1 provides some intuition. Notice that to the left of $\frac{d}{b-2k}$, $BR(\alpha) < \alpha$ while to the right of $\frac{d}{b-2k}$, $BR(\alpha) > \alpha$. We would expect reasonable evolutionary dynamics to respect the direction of the best response. Hence, we would expect that if the initial state μ is such that $A(\mu) < \frac{d}{b-2k}$, δ_0 would be evolutionarily stable while if $A(\mu) > \frac{d}{b-2k}$, δ_1 would be evolutionarily stable. The fact that δ_0 is also evolutionarily stable means that it is not necessary that an evolutionary dynamic must always converge to the Pareto efficient Nash equilibrium δ_1. It would, therefore, be of interest to see whether any exact equilibrium selection result we obtain under stochastic evolution would imply selection of the Pareto efficient Nash equilibrium in Example 6.

We conclude this discussion of Example 6 with a few more remarks on the importance of some of the assumptions. First, the cost function $c(x)$ needs to be strictly convex, as it is in Example 6. Otherwise, we would have a multiplicity of equilibria which cannot be Pareto ranked. Equilibrium selection then is not a very interesting question.[14] Second, we require $b > 0$ in (4.5) for strategic complementarities. If $b < 0$, we would have strategic substitutability which would generate a unique Nash equilibrium even if $c(x)$ is strictly convex (Section 4.2, Cheung and Lahkar (2018)). Equilibrium selection then becomes irrelevant. Deterministic dynamics in the medium run would also converge to that unique equilibrium.

We now present another example of a continuous strategy potential game with multiple equilibria. This example, adapted from Hofbauer et al. (2009), belongs to the class of doubly symmetric games described in Section 4.2.2.

Example 7 *Let $\mathscr{S} = [0,1]$ and*

$$f(x,y) = -\left(x - \frac{1}{4}\right)^2 + \beta\left(x - \frac{1}{4}\right)\left(y - \frac{1}{4}\right) - \left(y - \frac{1}{4}\right)^2. \tag{4.14}$$

This is a modified version of the game in Proposition 4 of Hofbauer et al. (2009).[15] Clearly, it is a doubly symmetric game. Hence, by (4.8), this is a potential game with potential function $\frac{1}{2}\bar{\pi}(\mu)$, where $\bar{\pi}(\mu)$ is the aggregate payoff of the potential game generated by (4.14).[16] Our interest is in the case where $\beta > 2$.

[13]For the logit dynamic, a caveat applies as it is generated through best response to a perturbed version of payoffs. Hence, convergence happens to a perturbed version of the potential function, called the entropy adjusted potential function. We discuss this function in Section 4.5 as it will also be useful for our stochastic analysis. But for low levels of perturbation, maximizers of the perturbed potential function would be arbitrarily close to maximizers of the original potential function. Also see Lahkar et al. (2022a) for a generalization of the logit dynamic in continuous strategy games and similar convergence results for potential games.

[14]Suppose the cost function is linear, i.e. $c(x) = dx$ with $d < b$. Then the payoff function in (4.5) would be $xbA(\mu) - cx = x(bA(\mu) - d)$, where $c < b$. Then, any μ with $A(\mu) = \frac{d}{b}$ will be a Nash equilibrium. Payoffs in any such equilibrium will be zero.

[15]This modification is necessary because in our model, $\mathscr{S} = [0,1]$ whereas in Hofbauer et al. (2009), the strategy set is $[-A, B]$, for $A, B > 0$. We can rewrite this game equivalently as one where the pairwise interaction payoff is $g(u,v) = -u^2 + \beta uv - v^2$ for $u, v \in \left[-\frac{1}{4}, \frac{3}{4}\right]$, where $u = x - \frac{1}{4}$ and $v = y - \frac{1}{4}$. An application of Proposition 4 in Hofbauer et al. (2009) then generates the conclusions of this example.

[16]The exact calculation of the potential function is cumbersome and is, hence, omitted. Nor is that calculation required for

Proposition 4 of Hofbauer et al. (2009) then implies that this game has three Nash equilibria; δ_0, $\delta_{\frac{1}{4}}$ and δ_1.

That same proposition in Hofbauer et al. (2009) also implies that δ_0 and δ_1 are the two local maximizers of the potential function $\frac{1}{2}\bar{\pi}(\mu)$ in the population game generated by (4.14). Therefore, by our discussion in Example 6, all standard evolutionary dynamics will converge to one of these Nash equilibria depending upon the initial state. But of these two Nash equilibria, it is easy to verify that δ_1 is the Pareto optimal equilibrium. Hence, just as in Example 6, it is possible that even in this example, deterministic evolutionary dynamics converge to the Pareto inferior equilibrium δ_0.

On the other hand, if $\beta \leqslant 2$ in (4.14) , then the stochastic approach may not yield additional insights. If $\beta < 2$, then $\delta_{\frac{1}{4}}$ is the unique Nash equilibrium and, hence, must be the unique maximizer of the potential function. Even deterministic dynamics will globally converge to this equilibrium. If $\beta = 2$, δ_x for any $x \in \mathscr{S}$ is a Nash equilibrium. The equilibrium payoff in all such equilibria is zero and Pareto ranking is not possible.

As in aggregative potential games, our interest will be again in whether it is the Pareto efficient Nash equilibrium that gets uniquely selected by stochastic evolution in doubly symmetric games. In answering this question, we note that Pareto efficiency implies maximization of the aggregate payoff $\bar{\pi}(\mu)$. Hence, in doubly symmetric games, the global maximizer of the potential function, $\frac{1}{2}\bar{\pi}(\mu)$, must always be the Pareto efficient Nash equilibrium. This is unlike in aggregative potential games. Part 2 in Proposition 8 shows that it is possible in such games for the Pareto inferior equilibrium δ_0 to be the global maximizer of the potential function. This happens, for example, in the left panel of Figure 4.1.

4.4 Approximations of the game π

To start our stochastic evolutionary analysis, we define three possible approximations of our original population game π. First, we define another continuous strategy large population game by replacing our original pairwise interaction payoff $f(x, y)$ with its step function approximation. Next, we apply the step function approximation to construct a large population finite strategy approximation of π. These two approximations have also been applied in Lahkar et al. (2022a) for the deterministic approximation of the logit dynamic in supermodular games. The third approximation, which follows from Sandholm (2010), requires us to consider not just a finite discretization of the original continuous strategy set but also a finite number of agents instead of the original continuum of agents. Finally, we establish that if π is a potential game, so are the three approximations.

We require some notation for defining these approximations. For $n \geqslant 1$ and $j \in \{0, 1, \ldots, 2^n - 1\}$, let us denote the dyadic rational number $\frac{j}{2^n}$ by $\alpha_{n,j}$. Further, let $\mathscr{S}_n$ be the collection of left-closed right-open intervals with dyadic rational endpoints in $\mathscr{S}$ defined in the following way:

$$\mathscr{S}_n := \left\{[\alpha_{n,j}, \alpha_{n,j+1}) : j = 0, 1, \ldots, 2^n - 2\right\} \cup [\alpha_{n,2^n-1}, 1]. \tag{4.15}$$

our analysis. Interested readers are referred to equation (8) of Hofbauer et al. (2009).

Thus, $\mathscr{S}_n$ is a collection of 2^n disjoint intervals which covers the original strategy space $\mathscr{S}$. We shall denote the sub-intervals $[\alpha_{n,j},\alpha_{n,j+1})$ by $\mathscr{S}_{n,j}$ for $j=0,1,\dots,2^n-2$ and $[\alpha_{n,2^n-1},1]$ by $\mathscr{S}_{n,2^n-1}$. We begin with the notion of large population step approximation of π. Recall the pairwise interaction payoff $f(x,y)$ in (4.1). Our next step is to construct a step function approximation of f. For $n\geqslant 1$, let $f_n:\mathscr{S}\times\mathscr{S}\to\mathbb{R}$ be such that for all $(x,y)\in\mathscr{S}\times\mathscr{S}$,

$$f_n(x,y)=\begin{cases} f(\alpha_{n,j},\alpha_{n,k}) & \text{if } (x,y)\in\mathscr{S}_{n,j}\times\mathscr{S}_{n,k} \text{ and } j,k=0,1,\dots,2^n-2\\ f(\alpha_{n,2^n-1},\alpha_{n,2^n-1}) & \text{if } (x,y)\in\mathscr{S}_{n,2^n-1}\times\mathscr{S}_{n,2^n-1}.\end{cases} \tag{4.16}$$

In this step function approximation, the payoff at any strategy profile $(x,y)\in\mathscr{S}\times\mathscr{S}$ equals the original pairwise payoff $f(\alpha_{n,j},\alpha_{n,k})$ where $(\alpha_{n,j},\alpha_{n,k})$ are the left-bottom points of the intervals in $\mathscr{S}_n$ to which (x,y) belongs. Since the pairwise interaction game f is assumed to be continuous, it follows by definition that, $(f_n)_{n\geqslant 1}$ is a sequence of step functions such that $f_n\to f$ uniformly as $n\to\infty$. We now obtain the first of our three approximations of π which is defined as follows.

Definition 18 *For $n\geqslant 1$, the* ***large population-step approximation (LPSA)*** *at level-n of our original population game π defined by (4.1) is the mapping $\pi^n:\mathscr{P}(\mathscr{S})\to\mathbb{L}^\infty(\mathscr{S})$ such that for all $\mu\in\mathscr{P}(\mathscr{S})$,*

$$\pi^n_x(\mu):=\int_{\mathscr{S}} f_n(x,y)\mu(dy),\quad \textit{for all } x\in\mathscr{S}, \tag{4.17}$$

where $f_n(x,y)$ is the step function approximation (4.16).

Thus, as the name LPSA suggests, (4.17) is the large population game we would obtain on the original state space $\mathscr{P}(\mathscr{S})$ if we replace the original pairwise interaction payoff $f(x,y)$ with its stepwise approximation. Note from (4.16) that we can equivalently write (4.17) as

$$\pi^n_x(\mu)=\begin{cases} \sum_{k=0}^{2^n-1} f(\alpha_{n,j},\alpha_{n,k})\mu(\mathscr{S}_{n,k}) & \text{if } x\in\mathscr{S}_{n,j} \text{ and } j=0,1,\dots,2^n-2\\ \sum_{k=0}^{2^n-1} f(\alpha_{n,2^n-1},\alpha_{n,k})\mu(\mathscr{S}_{n,k}) & \text{if } x\in\mathscr{S}_{n,2^n-1}.\end{cases} \tag{4.18}$$

Next, we define the notion of large population-finite strategy discretization of π. For this, we require some further notation. For a fixed $n\geqslant 1$, let $S_n:=\{\alpha_{n,j}:j=0,1,\dots,2^n-1\}$. Thus, S_n acts as a finite strategy discretization of $\mathscr{S}$. It consists of the left hand points of each of the finite intervals in the strategy set $\mathscr{S}_n$ defined by (4.15). Let $\Delta(S_n)$ be the collection of all probability distributions on the finite set S_n, that is, $\Delta(S_n):=\{\mathbf{p}^n\in[0,1]^{2^n}:\sum_{0\leqslant j\leqslant 2^n-1}\mathbf{p}^n_j=1\}$. We shall denote a generic element in $\Delta(S_n)$ by $\mathbf{p}^n$. We now define our second approximation of π.

Definition 19 *For $n\geqslant 1$, the* ***large population-finite strategy discretization (LPFSD)*** *at level-n of our original population game π defined by (4.1) is the mapping $\mathscr{D}^n_\pi:\Delta(S_n)\to\mathbb{R}^{2^n}$ such that for all $\boldsymbol{p}^n\in\Delta(S_n)$,*

$$\mathscr{D}^n_\pi[\boldsymbol{p}^n](\alpha_{n,j}):=\langle f_n(\alpha_{n,j},\cdot),\boldsymbol{p}^n\rangle=\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\boldsymbol{p}^n_k,\quad \textit{for all } j=0,1,\dots,2^n-1, \tag{4.19}$$

where $f_n(\alpha_{n,j}, \alpha_{n,k})$ is the step function approximation (4.16) but in this case defined only upon $S_n \times S_n$.

To interpret Definiiton 19, we can view $\mathscr{D}_\pi^n$ as the population game induced by random matching in the matrix game whose entries are given by $(f_n(\alpha_{n,j}, \alpha_{n,k}))_{0 \leqslant j,k \leqslant 2^n - 1}$. Equivalently, $\mathscr{D}_\pi^n$ is a finite strategy population game with $\mathscr{D}_\pi^n[\mathbf{p}^n](\alpha_{n,j})$ being the payoff of an agent playing strategy $\alpha_{n,j} \in S_n$ at the population state $\mathbf{p}^n \in \Delta(S_n)$. The notation $\mathbf{p}_j^n$ denotes the proportion of players in the population using strategy $\alpha_{n,j} \in S_n$. It is also worth noting that if $\mu \in \mathscr{P}(\mathscr{S})$ and $\mathbf{p}^n \in \Delta(S_n)$ such that $\mathbf{p}_j^n = \int_{\mathscr{S}_{n,j}} \mu(dy)$ for $j = 0, 1, \ldots, 2^n - 2$, then the two payoffs in (4.18) and in Definition 19 are equal. This implies the following relationship between the LPFSD and the original game π defined by (4.1). If $\mathbf{p}_j^n = \int_{\mathscr{S}_{n,j}} \mu(dy)$, then for n sufficiently large, $\mathscr{D}_\pi^n[\mathbf{p}^n](\alpha_{n,j}) \approx \int_{\mathscr{S}} f(\alpha_{n,j}, y)\mu(dy) = \pi_{\alpha_{n,j}}(\mu)$.

Our subsequent analysis will require the notion of a *large population finite strategy potential game*. We shall say that $\mathscr{D}_\pi^n$ is a *large population finite strategy potential game* if there exists a differentiable function $\mathscr{D}_\varphi^n : \mathbb{R}^n \to \mathbb{R}$ such that (Sandholm (2001))

$$\nabla \mathscr{D}_\varphi^n(\mathbf{p}^n) = \mathscr{D}_\pi^n[\mathbf{p}^n], \quad \text{for all } \mathbf{p}^n \in \Delta(S_n). \tag{4.20}$$

Formally, (4.20) is an application of Definition 17 to the special case of a finite strategy set. Hence, like in (4.3), we require a real valued function $\mathscr{D}_\varphi^n$ whose gradient equals the payoff function $\mathscr{D}_\pi^n[\mathbf{p}^n]$. Conceptually, however, (4.20) is a simpler condition as it is defined over finite dimensional probability distributions $\mathbf{p}^n$. Hence, it can be interpreted in terms of usual partial derivatives instead of the measure theoretic Frechét derivative in (4.3), Therefore, we can write (4.20) equivalently as $\frac{\partial \mathscr{D}_\varphi^n}{\partial \mathbf{p}_j^n}(\mathbf{p}^n) = \mathscr{D}_\pi^n[\mathbf{p}^n](\alpha_{n,j})$.

Finally, we define our third notion of approximation, which is a finite population-finite strategy discretization (FPFSD) of the population game π. Thus, in this approximation, there would not only be a finite strategy set S_n (see Definition 19) but also a finite number of players N. For $n, N \geqslant 1$, let $\Delta_N(S_n) := \Delta(S_n) \cap \epsilon_N \mathbb{Z}^{2^n}$. Thus, $\Delta_N(S_n)$ consists of all probability distributions on S_n which are integer multiples on ϵ_N. This is the set of population states that is now relevant as the total number of players playing all strategies in S_n would have to add up to N. Note that for every $N \geqslant 1$, $\Delta_N(S_n) \subseteq \Delta(S_n)$. We wish to define a finite strategy population game on $\Delta_N(S_n)$. However, with a finite number of players, we cannot directly define the FPFSD as the restriction of LPFSD (Definition 19) onto $\Delta_N(S_n)$. This is because in finite player population games, self-matching can occur with a positive probability. Hence, following Example 11.4.1 in Sandholm (2010), we define the FPFSD without self-matching as follows.

Definition 20 *For $n, N \geqslant 1$, the **finite player-finite strategy discretization (FPFSD)** at level-(n, N) of our original population game π defined by (4.1) is the mapping $\mathscr{D}_\pi^{n,N} : \Delta_N(S_n) \to \mathbb{R}^{2^n}$ such that for all $\boldsymbol{p}^{n,N} \in \Delta_N(S_n)$, and all $j = 0, 1, \ldots, 2^n - 1$,*

$$\mathscr{D}_\pi^{n,N}[\boldsymbol{p}^{n,N}](\alpha_{n,j}) := \mathscr{D}_\pi^n[\boldsymbol{p}^{n,N}](\alpha_{n,j}) + \frac{1}{N-1}(\mathscr{D}_\pi^n[\boldsymbol{p}^{n,N}](\alpha_{n,j}) - f_n(\alpha_{n,j}, \alpha_{n,j})), \tag{4.21}$$

where $\mathscr{D}_\pi^n[\boldsymbol{p}^{n,N}]$ is the LPFSD defined by (4.19), N is the finite number of players and each player has the strategy set S_n defined in the context of Definition 19.

Some explanation is required to understand (4.21). Recall that $\mathscr{D}_\pi^n[\mathbf{p}^{n,N}]$ is the finite strategy large population game (4.19) generated by the matrix game $(f_n(\alpha_{n,j},\alpha_{n,k}))_{0\leqslant j,k\leqslant 2^n-1}$ at the population state $\mathbf{p}^{n,N}\in\Delta_N(S_n)$. The key difference is that in the finite player game with N players, avoiding self–mathching means a player playing $\alpha_{n,j}$ only faces $N-1$ possible rivals. Of these, $N\mathbf{p}_k^{n,N}$ opponents play strategy $\alpha_{n,k}$, but only $N\mathbf{p}_j^{n,N}-1$ opponents play strategy $\alpha_{n,i}$. Hence, the effective population state is $\frac{1}{N-1}(N\mathbf{p}^{n,N}-\mathbf{e}_j^n)=\frac{N}{N-1}\mathbf{p}^{n,N}-\frac{\mathbf{e}_j^n}{N-1}$, where $\mathbf{e}_j^n$ is the standard basis vector on $\mathbb{R}^{2n}$ whose $\frac{j}{2^n}$–th entry is 1. Hence, when a player playing $\alpha_{n,j}\in S_n$ is randomly matched with another agent in the matrix game $(f_n(\alpha_{n,j},\alpha_{n,k}))_{0\leqslant j,k\leqslant 2^n-1}$, the payoff obtained is as follows:

$$
\begin{aligned}
&\frac{N}{N-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\mathbf{p}_k^n-\frac{1}{N-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\mathbf{e}_j^n\\
=&\frac{N}{N-1}\mathscr{D}_\pi^n[\mathbf{p}^n](\alpha_{n,j})-f_n(\alpha_{n,j},\alpha_{n,j}),\text{ by (4.19)}\\
=&\mathscr{D}_\pi^n[\mathbf{p}^{n,N}](\alpha_{n,j})+\frac{1}{N-1}(\mathscr{D}_\pi^n[\mathbf{p}^{n,N}](\alpha_{n,j})-f_n(\alpha_{n,j},\alpha_{n,j})),
\end{aligned}
$$

which is (4.21). Obviously, $\mathscr{D}_\pi^{n,N}[\mathbf{p}^{n,N}](\alpha_{n,j})\to\mathscr{D}_\pi^n[\mathbf{p}^{n,N}](\alpha_{n,j})$ as $N\to\infty$.

Chapter 11 in Sandholm (2010) defines the condition such that a finite player-finite strategy game is a potential game.[17] Following that definition, we say that the FPFSD $\mathscr{D}_\pi^{n,N}$ defined by (20) is a finite-population potential game if there exists a function $\mathscr{D}_\varphi^{n,N}:\Delta_N(S_n)\to\mathbb{R}$ such that for all $\mathbf{p}^{n,N}\in\Delta_N(S_n)$,

$$
\begin{aligned}
\mathscr{D}_\pi^{n,N}&[\mathbf{p}^{n,N}+N^{-1}(\mathbf{e}_j^n-\mathbf{e}_i^n)](\alpha_{n,j})-\mathscr{D}_\pi^{n,N}[\mathbf{p}^{n,N}](\alpha_{n,i})\\
&=\mathscr{D}_\varphi^{n,N}(\mathbf{p}^{n,N}+N^{-1}(\mathbf{e}_j^n-\mathbf{e}_i^n))-\mathscr{D}_\varphi^{n,N}(\mathbf{p}^{n,N}).
\end{aligned}\tag{4.22}
$$

The left hand side of (4.22) is the difference in payoff to an agent who changes strategy from $\alpha_{n,i}$ to $\alpha_{n,j}$. This causes a change in the population state from $\mathbf{p}^{n,N}$ to $\mathbf{p}^{n,N}+N^{-1}(\mathbf{e}_j^n-\mathbf{e}_i^n)$. Thus, for $\mathscr{D}_\pi^{n,N}$ to be a finite strategy–finite player potential game, we require that there exists a real–valued function $\mathscr{D}_\varphi^{n,N}$ defined on $\Delta_N(S_n)$ such that this difference in payoffs is equal to the difference in the values of $\mathscr{D}_\varphi^{n,N}$ at these population states.

Suppose that our original population game π is a potential game in the sense of Definition 17. It is then natural to expect that for $n,N\geqslant 1$, the three approximations of π, namely; the LPSA π^n at level-n, the LPFSD $\mathscr{D}_\pi^n$ at level-n, and the FPFSD $\mathscr{D}_\pi^{n,N}$ at level-(n,N) are also all potential games in the relevant strategy–player context in which they are defined. Our next proposition verifies that conjecture. In order to establish that result, we need the following notion of an embedding.

Recall the finite strategy set S_n and the set of probability distributions $\Delta(S_n)$ from Definition 19. For $n\geqslant 1$, let $\beta_n:\Delta(S_n)\to\mathscr{P}(\mathscr{S})$ be the mapping such that for every $\mathbf{p}^n\in\Delta(S_n)$,

$$
\beta_n(\mathbf{p}^n)(A):=\sum_{j=0}^{2^n-1}\mathbf{p}_j^n\delta_{\alpha_{n,j}}(A),\quad\text{for all } A\in\mathscr{B}(\mathscr{S}).\tag{4.23}
$$

[17] See equation (11.25) and Theorem 11.5.6 in that chapter.

Thus, for a generic probability vector $\mathbf{p}^n \in \Delta(S_n)$, $\beta_n(\mathbf{p}^n)$ denotes the corresponding Dirac probability measure on our original strategy set $\mathscr{S}$, with support S_n, where the point $\alpha_{n,j}$ is assigned probability $\mathbf{p}_j^n$. We however note from (4.23) that the range of the map β_n is essentially $\mathscr{P}(S_n)$. Further, for $n \geqslant 1$, the mapping $\mathbf{p}^n \mapsto \beta_n(\mathbf{p}^n)$ induces a continuous embedding of $\Delta(S_n)$ into $\mathscr{P}(\mathscr{S})$. This embedding will play an important role in the following proposition in establishing that the LPSA (4.17) and the LPFS (4.19) are potential games. Let $\mathbf{AC}((0,1);\mathscr{P}(\mathscr{S}))$ denote the collection of all absolutely continuous paths from $(0,1)$ to $\mathscr{P}(\mathscr{S})$.[18]

Proposition 9 *Suppose that π defined by (4.1) is a potential game with potential function φ as characterized in Definition 17. Then, the following statements hold true.*

1. *Fix $n \geqslant 1$. Let $\mathscr{D}_\varphi^n : \Delta(S_n) \to \mathbb{R}$ be a mapping such that*

$$\mathscr{D}_\varphi^n(\boldsymbol{p}^n) := \varphi(\beta_n(\boldsymbol{p}^n)), \quad \textit{for all } \boldsymbol{p}^n \in \Delta(S_n). \tag{4.25}$$

   where $\beta_n(\boldsymbol{p}^n)$ is as defined in (4.23). Then, the LPFSD $\mathscr{D}_\pi^n$ defined by (4.19) is a potential game with potential function $\mathscr{D}_\varphi^n$ defined by (4.25).

2. *Fix $n, N \geqslant 1$. Let $\mathscr{D}_\varphi^{n,N} : \Delta_N(S_n) \to \mathbb{R}$ be a mapping such that*

$$\mathscr{D}_\varphi^{n,N}(\boldsymbol{p}^{n,N}) := \frac{N}{2(N-1)}\Big(N\mathscr{D}_\varphi^n(\boldsymbol{p}^{n,N}) - \sum_{k \in S_n} f_n(\alpha_{n,k},\alpha_{n,k})\boldsymbol{p}_k^{n,N}\Big), \quad \textit{for all } \boldsymbol{p}^{n,N} \in \Delta_N(S_n). \tag{4.26}$$

   Then the FPFSD $\mathscr{D}_\pi^{n,N}$ defined by (4.21) is a potential game with potential function $\mathscr{D}_\varphi^{n,N}$ defined by (4.26).

3. *Fix $n \geqslant 1$. Recall the embedding $\beta_n(\boldsymbol{p}^n)$ from (4.23). Let $\varphi^n : \mathscr{P}(\mathscr{S}) \to \mathbb{R}$ be a mapping defined as*

$$\varphi^n(\mu) := \begin{cases} \varphi(\beta_n(\boldsymbol{p}^n)), & \textit{if } \mu = \beta_n(\boldsymbol{p}^n) \textit{ for some } \boldsymbol{p}^n \in \Delta(S_n) \\ \varphi(\delta_0) + \int_0^1 \langle \pi^n(\gamma_t), \dot{\gamma}_t \rangle dt, & \textit{otherwise,} \end{cases} \tag{4.27}$$

   where $\gamma \in \boldsymbol{AC}((0,1);\mathscr{P}(\mathscr{S}))$ such that $\gamma_0 = \delta_0$ and $\gamma_1 = \mu$. Then, the LPSA π^n defined in (4.17) is a potential game with potential function φ^n defined by (4.27).

Proof 13 *See Appendix C.3.*

For part 1, we check that the LPFSD $\mathscr{D}_\pi^n$ is a potential game with potential function (4.25) according to the definition in (4.20). For part 2, we apply the definition in (4.22) to show that the FPFSD $\mathscr{D}_\pi^{n,N}$ is a potential game with potential game with potential function (4.26). Indeed the definition of (4.26)

[18] Let (X,d) be a complete metric space and let $\gamma : (a,b) \to X$ be a curve. We say that $\gamma \in \mathbf{AC}^1((a,b);X)$ if there exists $m \in \mathbb{L}^1(a,b)$ such that

$$d(\gamma(s),\gamma(t)) \leqslant \int_s^t m(u)du, \quad \text{for all } a < s \leqslant t < b. \tag{4.24}$$

follows from Exercise 11.5.2 in Sandholm (2010). Part 3 perhaps requires more explanation. If μ can be expressed as an embedding as defined in (4.23), then the potential function of the LPSA φ^n is simply our original continuous strategy potential function (4.3) applied to that embedding. If, however, μ cannot be written as an embedding, we need to apply a curve $\gamma \in \mathbf{AC}((0,1);\mathscr{P}(\mathscr{S}))$ to define the potential function. We choose any such γ with the only restriction being that $\gamma_1 = \mu$. Even though the initial point $\gamma_0 = \delta_0$ in (4.27), that choice is arbitrary. If we then take the gradient of the second part of the RHS of (4.27), then, intuitively, we obtain a telescopic sum of differences $\pi^n(\delta_0) + \sum(\pi^n(\gamma_{t+\varepsilon}) - \pi^n(\gamma_t))$. After canceling out common terms, we would be left with $\nabla\varphi^n(\mu) = \pi^n(\gamma_1) = \pi^n(\mu)$, which is as required by our original definition of the potential function in (4.3).[19]

4.5 Stochastic Evolution and Equilibrium Selection

We now present our main results on stochastic evolution and equilibrium selection in potential games in this section. We first discuss stochastic convergence along the lines of Sandholm (2010) for finite strategy potential games in the context of both a finite number of players and a large population of players. For this purpose, we define a class of revision protocols called exponential revision protocols which describe how agents revise strategies stochastically. We then establish our results for the class of games we are interested in, which are large population continuous strategy potential games. Our main conclusion would be that a stochastic process generated by a direct exponential revision protocol would converge to the global maximizer of the potential function of the continuous strategy potential game.

Consider a finite strategy game with 2^n strategies like LPFSD $\mathscr{D}^n_\pi$ in Definition 19 or the FPFSD $\mathscr{D}^{n,N}_\pi$ in Definition 20. Let $\mathbf{u}^n$ be a payoff vector and $\alpha_{n,j}, \alpha_{n,k} \in S_n$ be two strategies in such a game. Thus, $\mathbf{u}^n_j$ and $\mathbf{u}^n_k$ are respectively the payoffs at these two strategies. A revision protocol $\rho : \mathbb{R}^{2^n} \to \Delta(S_n)^{2^n}$ then takes the form $\rho_{jk}(\mathbf{u}^n)$, which represents the probability with which an agent currently playing $\alpha_{n,j} \in S_n$ adopts strategy $\alpha_{n,k} \in S_n$ upon receiving a chance to revise strategy. Thus, $\sum_{k\in S_n} \rho_{jk}(\mathbf{u}^n) = 1$. Following Sandholm (2010) (Section 11.5.2), we then call a revision protocol $\rho : \mathbb{R}^{2^n} \to \Delta(S_n)^{2^n}$ an *exponential revision protocol with noise level* $\eta > 0$ if

$$\rho_{jk}(\mathbf{u}^n) = \frac{\exp(\eta^{-1}\psi(\mathbf{u}^n_j, \mathbf{u}^n_k))}{d_{jk}(\mathbf{u}^n)}, \quad \text{for all } \mathbf{u}^n \in \mathbb{R}^{2^n}, \tag{4.28}$$

where the functions $\psi : \mathbb{R}^2 \to \mathbb{R}$ and $d : \mathbb{R}^{2^n} \to (0,\infty)^{2^n \times 2^n}$ satisfy

1. $\psi(u,v) - \psi(v,u) = v - u$, for all $u, v \in \mathbb{R}$ and
2. $d_{jk}(\mathbf{u}^n) = d_{kj}(\mathbf{u}^n)$, for all $\mathbf{u}^n \in \mathbb{R}^{2^n}$ and all $j, k = 0, 1, \ldots, 2^n - 1$.

[19] Even though the choice of γ is arbitrary, the potential function (4.27) is well defined. This is the sense that the potential function is independent of the path traced by the curve γ. To see this, note from part (i) of Proposition 9 that the LPFSD game is a potential game. As a result the vector field $\mathbf{p}^n \mapsto \mathscr{D}^n_\pi(\mathbf{p}^n)$ is a conservative. Hence, the integral in the RHS of (4.27) is independent of the curve γ. The definition of the potential function does depend upon the initial condition of γ. But a change in the initial condition of γ only causes a uniform translation in the value of the potential by a fixed constant. This renders the potential function (4.27) non–unique but does not open it up to any mathematical inconsistency.

Perhaps the most well known example of an exponential revision protocol is the logit protocol. In the logit protocol, we have $\psi(u,v)=v$ and $d_{jk}(\mathbf{u}^n)=\sum_{i\in S_n}\exp\left(\eta^{-1}\mathbf{u}_i^n\right)$. Therefore, in our context, the logit protocol takes the form

$$\rho_{jk}(\mathbf{u}^n)=\frac{\exp\left(\eta^{-1}\mathbf{u}_k^n\right)}{\sum_{i\in S_n}\exp\left(\eta^{-1}\mathbf{u}_i^n\right)}. \tag{4.29}$$

For any $\eta>0$, the logit protocol assigns a strictly positive probability to moving to any strategy $\alpha_{n,k}$ from $\alpha_{n,j}$. But as $\eta\downarrow 0$, the probability assigned to any strategy $\alpha_{n,j}$ that is not a best response goes to zero. The logit protocol is, however, not the only form that a direct exponential revision protocol can take. For example, instead of $\psi(u,v)=v$, we can have $\psi(u,v)=-u$ or $\max(0,v-u)$ or $-\max(0,u-v)$. We refer the reader to Section 11.5.2 of Sandholm (2010) for further details of such revision protocols.[20]

Our analysis will require us to define a stochastic process on the FPFSD $\mathscr{D}_\pi^{n,N}$ using the exponential revision protocol (4.28). Before that, however, we introduce a modification of the state variable $\mathbf{p}^{n,N}$ on which $\mathscr{D}_\pi^{n,N}$ is defined in Definition 20. This modification is required because with a finite number of players, a change of strategy by an agent will have an impact on the social state. For example, suppose an agent playing strategy $\alpha_{n,j}$ at the population state $\mathbf{p}^{n,N}$ receives a revision opportunity and switches to strategy $\alpha_{n,k}$. The population state then changes to $\mathbf{p}^{n,N}+\epsilon_N(\mathbf{e}_k^n-\mathbf{e}_j^n)$. A naive agent, however, wouldn't realize this impact that his strategy change has on the social state. Such an agent would, therefore, believe that the social state continues to be $\mathbf{p}^{n,N}$ and, hence, compare his current payoff $\mathscr{D}_\pi^{n,N}[\mathbf{p}^{n,N}](\alpha_{n,j})$ to $\mathscr{D}_\pi^{n,N}[\mathbf{p}^{n,N}](\alpha_{n,k})$, which is the payoff of strategy $\alpha_{n,k}$ at $\mathbf{p}^{n,N}$. A "clever agent", on the other hand, would realize the impact of one's strategy revision on the social state. Hence, such an agent would compare the current payoff $\mathscr{D}_\pi^{n,N}[\mathbf{p}^{n,N}](\alpha_{n,j})$ to $\mathscr{D}_\pi^{n,N}[\mathbf{p}^{n,N}+\epsilon_N(\mathbf{e}_k^n-\mathbf{e}_j^n)](\alpha_{n,k})$, which is the payoff of strategy $\alpha_{n,k}$ at the new actual state $\mathbf{p}^{n,N}+\epsilon_N(\mathbf{e}_k^n-\mathbf{e}_j^n)$. Following Sandholm (2010) (Section 11.4.2), we refer to this method of evaluating payoffs as *clever payoff evaluation*.

To formalize the notion of a clever payoff, recall from Definition 20 that for fixed $N,n\geqslant 1$, $\Delta_N(S_n)$ is the collection of probability distributions on S_n which are integral multiples of N^{-1}. Thus, $\Delta_N(S_n)$ is the state space of a finite population consisting of N-players with a generic element in that space being $\mathbf{p}^{n,N}$. We then define the collection of *diminished states* by $\Delta_N^-(S_n):=\{\mathbf{p}_-^{n,N}\in\mathbb{R}_+^{2^n}:\sum_{0\leqslant j\leqslant 2^n-1}\mathbf{p}_{-,j}^{n,N}=\frac{N-1}{N}$ and $N\mathbf{p}_-^{n,N}\in\mathbb{Z}^{2^n}\}$. Thus, every diminished state denotes the behavior of the opponents of one member in a population of size N. Given the FPFSD game $\mathscr{D}_\pi^{n,N}$, the *clever payoff function* is defined as a mapping $\check{\mathscr{D}}_\pi^{n,N}:\Delta_N^-(S_n)\to\mathbb{R}^{2^n}$ such that

$$\check{\mathscr{D}}_\pi^{n,N}[\mathbf{p}_-^{n,N}](\alpha_{n,k}):=\mathscr{D}_\pi^{n,N}[\mathbf{p}_-^{n,N}+N^{-1}\mathbf{e}_k^n](\alpha_{n,k}),\quad \text{for all } k=0,1,\ldots,2^n-1. \tag{4.30}$$

Thus, (4.30) is the "clever payoff" of an agent who has moved to a new strategy $\alpha_{n,k}$. If that player had been playing $\alpha_{n,j}$ earlier, then we can also write the RHS of (4.30) as $\mathscr{D}_\pi^{n,N}[\mathbf{p}^{n,N}+\epsilon_N(\mathbf{e}_k^n-\mathbf{e}_j^n)](\alpha_{n,k})$. In that case, the diminished state would have been $\mathbf{p}_-^{n,N}=\mathbf{p}^{n,N}-\epsilon_N\mathbf{e}_j^n$, which describes the behavior of all other agents except the one playing $\alpha_{n,j}$. Comparing (4.30) with Definition 20, we also conclude that clever payoff evaluation incorporates two subtleties introduced by a finite set of players. First, as in

[20]Over finite time horizons, the logit protocol generates the canonical logit dynamic discussed in the Introduction.

Definition 20, payoffs need to be defined by excluding the possibility of self–matching. Second, as in (4.30), agents also have the be conscious of the impact they have on the population state while changing strategy.[21] It is also worth pointing out that the difference between the clever payoff and the naive FPFS payoff $\mathscr{D}_\pi^{n,N}[\mathbf{p}^{n,N}](\alpha_{n,k})$ diminishes asymptotically since a single agent becomes progressively insignificant as the number of players $N \to \infty$.

We now introduce our stochastic process on FPFSD game $\mathscr{D}_\pi^{n,N}$ but under the condition that players revise strategies on the basis of clever payoff evaluation. Let $(\Omega, \mathscr{H}, \mathbb{P})$ be an arbitrary but fixed probability space. For $N, n \geqslant 1$ and $\alpha_{n,j}, \alpha_{n,k} \in S_n$, let $\rho_{jk}(\check{\mathscr{D}}_\pi^{n,N}[\mathbf{p}_-^{n,N}])$ be the direct exponential revision protocol (4.28) wherein players decide on strategy revision based on the clever payoff function (4.30). Also recall that if a player is playing $\alpha_{n,j}$, then from that player's perspective, the reduced population state $\mathbf{p}_-^{n,N} = \mathbf{p}^{n,N} - \epsilon_N \mathbf{e}_j^n$. We then define the *direct exponential clever stochastic process* $(\mathbf{X}_{\eta,k}^{n,N})_{k\geqslant 1}$ on $\mathscr{D}_\pi^{n,N}$ induced by this revision protocol as

$$\mathbb{P}\big(\mathbf{X}_{\eta,\tau+\epsilon_N}^{n,N} = \mathbf{q}^{n,N} | \mathbf{X}_{\eta,\tau}^{n,N} = \mathbf{p}^{n,N}\big) = \begin{cases} \mathbf{p}_j^{n,N} \rho_{jk}(\check{\mathscr{D}}_\pi^{n,N}(\mathbf{p}^{n,N} - \epsilon_N \mathbf{e}_j^n)), \text{ if } \mathbf{q}^{n,N} = \mathbf{p}^{n,N} + \epsilon_N(\mathbf{e}_k^n - \mathbf{e}_j^n),\ k \neq j \\ \displaystyle\sum_{j=0}^{2^n-1} \mathbf{p}_j^{n,N} \rho_{jj}(\check{\mathscr{D}}_\pi^{n,N}(\mathbf{p}^{n,N} - \epsilon_N \mathbf{e}_j^n)), \qquad \text{if } \mathbf{q}^{n,N} = \mathbf{p}^{n,N} \\ 0, \qquad\qquad \text{otherwise.} \end{cases} \tag{4.31}$$

The transition probabilities given by (4.31) induces a discrete time stochastic process $(\mathbf{X}_{\eta,\tau}^{n,N})_{\tau\geqslant 1}$ on $\Delta_N(S_n)$. To understand this process, suppose that the population state at time τ is $\mathbf{p}^{n,N}$. Further, suppose that at time τ, a player playing strategy $\alpha_{n,j}$ receives a revision opportunity and selects $\alpha_{n,k} \neq \alpha_{n,j}$. Then the population state moves to $\mathbf{q}^{n,N} = \mathbf{p}^{n,N} + \epsilon_N(\mathbf{e}_k^n - \mathbf{e}_j^n)$ at the next time interval $\tau + \varepsilon_N$. This transition from $\mathbf{p}^{n,N}$ to $\mathbf{q}^{n,N}$ happens with probability $\mathbf{p}_j^{n,N} \rho_{jk}(\check{\mathscr{D}}_\pi^{n,N}(\mathbf{p}^{n,N} - \epsilon_N \mathbf{e}_j^n))$. Thus, a player playing $\alpha_{n,j}$ is chosen with probability $\mathbf{p}_j^{n,N}$, who then chooses a new strategy $\alpha_{n,k} \in S_n$ with probability $\rho_{jk}(\check{\mathscr{D}}_\pi^{n,N}(\mathbf{p}^{n,N} - \epsilon_N \mathbf{e}_j^n))$. On the other hand, the population remains unchanged between the two time periods, i.e. $\mathbf{q}^{n,N} = \mathbf{p}^{n,N}$ if an agent who receives a revision opportunity doesn't change strategy. This happens with probability $\sum_{j=0}^{2^n-1} \mathbf{p}_j^{n,N} \rho_{jj}(\check{\mathscr{D}}_\pi^{n,N}(\mathbf{p}^{n,N} - \epsilon_N \mathbf{e}_j^n))$. Due to the presence of the noise parameter η in (4.28), (4.31) is an irreducible stochastic process.

Agents in the stochastic process (4.31) apply clever payoff evaluation and follow a direct exponential revision protocol. Moreover, the FPFSD game $\mathscr{D}_\pi^{n,N}$ is a potential game (part (ii) of Proposition 9). Hence, by Theorem 11.5.12 in Sandholm (2010), this stochastic process admits a stationary distribution on $\Delta_N(S_n)$. We wish to characterize the limiting behavior of this stationary distribution as $N \to \infty$. Intuitively, as $N \to \infty$, the FPFSD game converges to the LPFSD game $\mathscr{D}_\pi^n$ (Definitions 19). Hence, the limiting stationary distribution as $N \to \infty$ will inform us of the behavior of the distribution in the LPFSD game.

To characterize this limiting distribution, we adapt Theorem 12.2.3 in Sandholm (2010) to our two finite strategy approximations, the FPFSD $\mathscr{D}_\pi^{n,N}$ and LPFSD $\mathscr{D}_\pi^n$. For establishing the result, we require

[21] Thus, by applying (4.21), we can also write the clever payoff (4.30) as $\mathscr{D}_\pi^n[\mathbf{p}_-^{n,N} + N^{-1}\mathbf{e}_k^n](\alpha_{n,j}) + \frac{1}{N-1}(\mathscr{D}_\pi^n[\mathbf{p}_-^{n,N} + N^{-1}\mathbf{e}_k^n](\alpha_{n,j}) - f_n(\alpha_{n,j}, \alpha_{n,j}))$.

the *entropy adjusted potential function* for the LPFSD game. Recall from (4.25) the potential function $\mathscr{D}^n_\varphi$ for the LPFSD game. Using this potential function, we define the entropy adjusted potential function for the perturbation factor $\eta > 0$ as a mapping $\mathscr{D}^{n,\eta}_\varphi : \Delta(S_n) \to \mathbb{R}$ as

$$\mathscr{D}^{n,\eta}_\varphi(\mathbf{p}^n) := \mathscr{D}^n_\varphi(\mathbf{p}^n) - \eta \sum_{j=0}^{2^n-1} \mathbf{p}^n_j \log \mathbf{p}^n_j, \quad \text{for all } \mathbf{p}^n \in \Delta(S_n).^{22} \tag{4.32}$$

The entropy adjusted potential function is, therefore, the original potential function minus $\eta \sum_{j=0}^{2^n-1} \mathbf{p}^n_j \log \mathbf{p}^n_j$, which is the Shannon entropy of $\mathbf{p}^n$ multiplied by the same perturbation parameter η that appears in the exponential revision protocol (4.28). This function has been used in the literature on deterministic evolution to show convergence of perturbed best response dynamics like the logit dynamic in potential games (Hofbauer and Sandholm (2007)). Our present analysis will illustrate its pivotal importance in stochastic evolution as well. In addition to the entropy adjusted potential function, we define a new piece of notation which will come handy in stating the stochastic evolutionary results. For a bounded continuous map $h : E \to \mathbb{R}$, let $\Delta h : E \to \mathbb{R}$ be defined as

$$\Delta h(u) := h(u) - \max_{y \in E} h(y), \quad \text{for all } u \in E. \tag{4.33}$$

Thus, $\Delta h(u) \leq 0$, with equality only at the maximizer of h. We now state our result on the stationary distribution of the stochastic process (4.31) as $N \to \infty$.

Proposition 10 *Suppose that $\pi : \mathscr{P}(\mathscr{S}) \to \mathbb{L}^\infty(\mathscr{S})$ is a potential game with potential function φ as defined in Definition 17. For perturbation parameter $\eta > 0$ and $n \geqslant 1$ fixed, let $(\boldsymbol{X}^{n,N}_{\eta,\tau})_{\tau \geqslant 1}$ be the direct exponential clever stochastic process on $\Delta_N(S_n)$ induced by the clever payoff evaluation game $\mathscr{D}^{n,N}_\pi$ with stationary distribution $\mu^{n,N}_\eta$ for every $N \geqslant 1$. Then,*

$$\lim_{N \to \infty} \sup_{\boldsymbol{p}^{n,N} \in \Delta_N(S_n)} \left| \frac{\eta}{N} \log \mu^{n,N}_\eta(\boldsymbol{p}^{n,N}) - \Delta \mathscr{D}^{n,\eta}_\varphi(\boldsymbol{p}^{n,N}) \right| = 0, \tag{4.34}$$

where $\mathscr{D}^{n,\eta}_\varphi$ is the entropy adjusted potential function (4.32).

Proof 14 *See Appendix C.4.*

We emphasise that in this result, we hold n fixed. Therefore, this result is only with reference to the two finite strategy games, the FPFSD $\mathscr{D}^{n,N}_\pi$ and the LPFSD $\mathscr{D}^n_\pi$, we are considering. As noted earlier, with clever payoff evaluation, the direct exponential stochastic process (4.31) will admit a stationary distribution in the FPFSD. But as $N \to \infty$, the FPFSD also converges to the LPSFD. Hence, intuitively, as $N \to \infty$ in (4.34), Proposition 10 is characterizing the stationary distribution of a direct exponential stochastic process for the LPFSD $\mathscr{D}^n_\pi$. The proof of this result relies on the potential functions $\mathscr{D}^n_\varphi$ and $\mathscr{D}^{n,N}_\varphi$ as defined in (4.25) and (4.26) for the LPFSD and FPFSD games respectively. It shows that a rescaled

[22] Here we use the standard convention that $0 \log 0 = 0$.

version of $\mathscr{D}_{\varphi}^{n,N}$ converges uniformly to $\mathscr{D}_{\varphi}^{n}$ as $N \to \infty$ and that $\mathscr{D}_{\varphi}^{n}$ is a $\mathscr{C}^1$ function. Hence, by Theorem 12.2.3 in Sandholm (2010), the conclusion follows.

The main implication of Proposition 10 is that as $N \to \infty$, the stationary distribution of the direct exponential stochastic process (4.31) in $\Delta_N(S_n)$ puts all its mass on the state that globally maximizes the entropy adjusted potential function (4.32) in the LPFSD. To understand the result, we note that for a fixed but sufficiently large N, the expression inside $|\cdot|$ in (4.34) implies $\mu_{\eta}^{n,N}(\mathbf{p}^{n,N}) \leqslant \exp(\frac{N}{\eta}\Delta\mathscr{D}_{\varphi}^{n,\eta}(\mathbf{p}^{n,N}))$. By (4.33), $\Delta\mathscr{D}_{\varphi}^{n,\eta}(\mathbf{p}^{n,N}) \leqslant 0$ at all $\mathbf{p}^{n,N}$ with equality only at the global maximizer of (4.32) in $\Delta_N(S_n)$. Hence, for a fixed but sufficiently large N, $\exp(\frac{N}{\eta}\Delta\mathscr{D}_{\varphi}^{n,\eta}(\mathbf{p}^{n,N})) \approx 0$ or $\mu_{\eta}^{n,N}(\mathbf{p}^{n,N}) \approx 0$ for all $\mathbf{p}^{n,N}$ except the global maximizer of (4.32). Equivalently, the stationary distribution $\mu_{\eta}^{n,N}$ must be putting nearly all its mass on the global maximizer (4.32) on $\Delta_N(S_n)$, which is the state space of the FPFSD game. But as $N \to \infty$, the global maximizer of $\Delta_N(S_n)$ would become indistinguishable from that of $\Delta(S_n)$, the state space of the LPFSD game. This leads to our conclusion.

Proposition 10 is an equilibrium selection result based on stochastic evolution. Maximizers, both local and global, of the entropy adjusted potential function (4.32) are perturbed Nash equilibria of the underlying large population finite strategy potential game $\mathscr{D}_{\pi}^{n}$, with the perturbation depending upon the noise parameter $\eta > 0$ (Hofbauer and Sandholm (2007)). Hence, an equivalent interpretation of Proposition 10 is that the stationary distribution $\mu_{\eta}^{n,N}$ puts all its mass on that particular perturbed equilibrium of $\mathscr{D}_{\pi}^{n}$ that globally maximizes the entropy adjusted potential function. Following standard interpretation of the stationary distribution, this mean that in the long run, as $\tau \to \infty$, the stochastic process spends nearly all its time on the global maximizer irrespective of the initial state. This is, therefore, an exact or history–independent equilibrium selection result.

However, Proposition 10 applies only to the LPFSD $\mathscr{D}_{\pi}^{n}$ while our objective is to obtain an equilibrium selection result in our original continuous strategy potential game π. Hence, we now extend Proposition 10 to the continuous strategy case in the following two theorems, which represent the main contribution of this paper. For this purpose, we embed the direct exponential stochastic process (4.31) in the infinite dimensional space $\mathscr{P}(\mathscr{S})$ over which our original game is defined. To do so, we recall $\mathscr{S}_{n,j} = [\alpha_{n,j}, \alpha_{n,j+1})$ (see the discussion following (4.15)). For $n, N \geqslant 1$, we now define the following subset of $\mathscr{P}(\mathscr{S})$ as:

$$\bar{\mathscr{P}}_{n,N}(\mathscr{S}) := \left\{ \mu \in \mathscr{P}(\mathscr{S}) : \mu(A) = \sum_{j=0}^{2^n-1} 2^n \mathbf{p}_j^{n,N} \lambda(A \cap \mathscr{S}_{n,j}) \text{ for all } A \in \mathscr{B}(\mathscr{S}) \text{ for some } \mathbf{p}^{n,N} \in \Delta_N(S_n) \right\}. \quad (4.35)$$

Thus, $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ is the collection of all absolutely continuous probability measures on our continuous strategy set $\mathscr{S}$ with step probability density functions, such that the probability in the sub-intervals $\{\mathscr{S}_{n,j} : j = 0, 1, \ldots 2^n - 1\}$ is identified by some element in $\Delta_N(S_n)$. The multiplicative term 2^n in (4.35) is arising due to the dyadic rational numbers $\frac{j}{2^n}$ we are using for our finite strategy approximations. We note that as N followed by n is allowed to approach infinity, the set $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ approaches the subset of absolutely continuous probability measures in $\mathscr{P}(\mathscr{S})$.[23]

Fix $n, N \geqslant 1$. Another piece of notation that will be useful for us is the mapping $\bar{\beta}_{n,N} : \Delta_N(S_n) \to$

[23]The double limit implies first N and then n approaches infinity.

$\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ such that for all $\mathbf{p}^{n,N} \in \Delta_N(S_n)$,

$$\bar{\beta}_{n,N}[\mathbf{p}^{n,N}](A) := \sum_{j=0}^{2^n-1} 2^n \mathbf{p}_j^{n,N} \lambda(A \cap \mathscr{S}_{n,j}), \quad \text{for all } A \in \mathscr{B}(\mathscr{S}). \tag{4.36}$$

In words, $\bar{\beta}_{n,N}$ maps a genetic probability vector $\mathbf{p}^{n,N} \in \Delta_N(S_n)$ in the LPFSD game into an absolutely continuous step-probability measure $\bar{\beta}_{n,N}(\mathbf{p}^{n,N})$ on $\mathscr{S}$ such that the probability density of $\bar{\beta}_{n,N}(\mathbf{p}^{n,N})$ in each sub-interval $\mathscr{S}_{n,j}$ is constant and is subject to the restriction that $\bar{\beta}_{n,N}[\mathbf{p}^{n,N}](\mathscr{S}_{n,j}) = \mathbf{p}_j^n$ for all $j = 0, 1, \ldots, 2^n - 1$. Thus, in order to embed the direct exponential stochastic process $(\mathbf{X}_{\eta,\tau}^{n,N})_{\tau \geqslant 1}$ defined by (4.31) into $\mathscr{P}(\mathscr{S})$, we consider the new process $\bar{\beta}_{n,N}(\mathbf{X}_{\eta,\tau}^{n,N})$. Since $(\mathbf{X}_{\eta,\tau}^{n,N})_{\tau \geqslant 1} \subseteq \Delta_N(S_n) \subset \Delta(S_n)$, we have by (4.36) that $(\bar{\beta}_{n,N}(\mathbf{X}_{\eta,\tau}^{n,N}))_{\tau \geqslant 1} \subseteq \bar{\mathscr{P}}_{n,N}(\mathscr{S})$. The fact that $\bar{\beta}_{n,N}$ is a measurable map, implies that the process $(\bar{\beta}_{n,N}(\mathbf{X}_{\eta,\tau}^{n,N}))_{\tau \geqslant 1}$ is measurable.

We can now state the counterpart of Proposition 10 for our continuous strategy potential game π. Recall from Definition 19 that the LPFSD $\mathscr{D}_\pi^n$ has 2^n strategies and approximates π as $n \to \infty$. Hence, we need to extend the conclusion of Proposition 10 by taking $n \to \infty$. We require some additional notation. For $n, N \geqslant 1$, let $\Xi_\eta^{n,N}$ denote the measure on $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ obtained via the push-forward of the measure $\mu_\eta^{n,N}$ on $\Delta(S_n)$ by the map $\bar{\beta}_{n,N}$ defined by (4.36). We also need the continuous strategy counterpart of (4.32). Hence, for $\eta > 0$, we define the continuous strategy entropy adjusted potential function as a mapping $\varphi_\eta : \mathscr{P}(\mathscr{S}) \to \mathbb{R} \cup \{-\infty\}$ as

$$\varphi_\eta(\mu) := \begin{cases} \varphi(\mu) - \eta \int_{\mathscr{S}} \mu_x \log \mu_x \, dx & \text{if } \mu \text{ admits a probability density function} \\ -\infty & \text{otherwise.} \end{cases} \tag{4.37}$$

In the above definition of entropy adjusted potential function, with an abuse of notation, we denote the density of a probability measure μ at x by μ_x. Also recall the notation $\Delta h(u)$ from (4.33). The following theorem then characterizes the limiting distribution of a direct exponential stochastic process in our continuous strategy potential game. In this theorem, which is the first of our two main results, we keep the perturbation factor $\eta > 0$ fixed. In the second theorem, we will also take $\eta \to 0$.

Theorem 8 *Suppose that the population game $\pi : \mathscr{P}(\mathscr{S}) \to \mathbb{L}^\infty(\mathscr{S})$ is a potential game with potential function φ as defined in Definition 17. Let $(\mathbf{X}_{\eta,\tau}^{n,N})_{\tau \geqslant 1}$ be the direct exponential clever stochastic process (4.31) induced by the FPFSD game $\mathscr{D}_\pi^{n,N}$. Suppose that $\Xi_\eta^{n,N}$ denotes the stationary distribution of the embedded clever stochastic process $(\bar{\beta}_n(\mathbf{X}_{\eta,\tau}^{n,N}))_{\tau \geqslant 1}$. Then, for a fixed $\eta > 0$,*

$$\lim_{n \to \infty} \lim_{N \to \infty} \sup_{\bar{\mu}^{n,N} \in \bar{\mathscr{P}}_{n,N}(\mathscr{S})} \left| \frac{\eta}{N} \log \Xi_\eta^{n,N}(\bar{\mu}^{n,N}) - \Delta \varphi_\eta(\bar{\mu}^{n,N}) \right| = 0. \tag{4.38}$$

Proof 15 *See Appendix C.5.*

The conclusion of this theorem holds for our continuous strategy game π. This is because although the direct exponential stochastic process (4.31) is defined on the FPFSD game, that game converges to

π under the double limit $\lim_{n\to\infty}\lim_{N\to\infty}$. Essentially, we translate the stationary distribution $\mu_\eta^{n,N}$ in Proposition 10 to the new stationary distribution $\Xi_\eta^{n,N}$ in Theorem 8 through the push–forward $\bar{\beta}_\#^n \mu_\eta^{n,N}$ and apply this stationary distribution to π by taking the double limit. Hence, the main implication of this theorem is that at this double limit, the stationary distribution $\Xi_\eta^{n,N}$ of the embedded clever stochastic process $(\bar{\beta}_n(\mathbf{X}_{\eta,\tau}^{n,N}))_{\tau\geqslant 1}$ puts all its mass on the global maximizer of the entropy adjusted potential function φ_η defined by (4.37) in our original continuous strategy potential game π.

To understand how we obtain this conclusion, it would be useful to apply a similar manipulation as in (4.34) and express (4.38) equivalently as

$$\Xi_\eta^{n,N}(\bar{\mu}^{n,N}) \approxeq \exp\Big[\eta^{-1}N\Big(\Delta\varphi_\eta(\bar{\mu}^{n,N})+\mathscr{E}_{n,N}(\bar{\mu}^{n,N})\Big)\Big], \quad \text{uniformly in } \bar{\mathscr{P}}_{n,N}(\mathscr{S}) \tag{4.39}$$

for sufficiently large n,N. Here $\mathscr{E}_{n,N}$ denotes a term which converges uniformly in $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ to 0 as N followed by n approaches infinity. Recall from (4.33) that $\Delta\varphi_\eta(\bar{\mu}^{n,N})\leq 0$ with equality only at the global maximizer of φ_η in $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$. Hence, as $N\to\infty$, $\exp(\eta^{-1}N\Delta\varphi_\eta(\bar{\mu}^{n,N}))\to 0$ except at the global maximizer. Once we take the second limit $n\to\infty$, the global maximizer on $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ will approach the global maximizer on $\mathscr{P}(\mathscr{S})$ under the total variation topology.[24] In proving the result, we use the large population step approximation (LPSA) game φ^n (Definition 18) and its potential function defined in (4.27). As $N\to\infty$, the FPFSD game $\mathscr{D}_\pi^{n,N}$ approximates the LPFSD game $\mathscr{D}_\pi^n$. We already have the conclusion from Proposition 10 for the LPFSD game. By reversing the process through which we went from Definition 18 to Definition 19, we generate the LPSA game from the LPFSD game. We then apply the conclusion of Proposition 10 to the potential function (4.27) of the LPSA game. This is where the constructions in (4.35) and (4.36) become useful. Then, by taking $n\to\infty$, we obtain our original game π from the LPSA thereby giving us the conclusion of Theorem 8. Of course, this heuristic explanation conceals multiple technical complications which we address in the proof.

Like Proposition 10, Theorem 8 is also an equilibrium selection result but for the actual continuous strategy potential game π. In such games, all maximizers of entropy adjusted potential function $\varphi_\eta(\mu)$ are perturbed Nash equilibria (Lahkar and Riedel (2015)). Theorem 8 implies that among all such perturbed equilibria, the stationary distribution of the direct exponential stochastic process uniquely selects the global maximizer of $\varphi_\eta(\mu)$. This is once again in the sense that in the long run, as $\tau\to\infty$, the stochastic process spends nearly all the time at this global maximizer regardless of the initial state.

However, Theorem 8 is an equilibrium selection result on the set of perturbed Nash equilibria. Our objective is to obtain such a result on the set of Nash equilibria of π. It is again well–known that in a continuous strategy potential game like π, every maximizer of the potential function φ (see Definition 17) is a Nash equilibrium (Cheung and Lahkar (2018)). Therefore, a result like Theorem 8 will suffice for us if we can establish that our stochastic evolutionary process converges globally to one of those maximizers of the potential function. That would require us to extend Theorem 8 by taking the additional

[24]Recall that as N followed by n is allowed to approach infinity, the space $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ approaches the subset of absolutely continuous probability measures in $\mathscr{P}(\mathscr{S})$ under the variational norm. That suffices for our purpose as any maximizer of φ_η must be absolutely continuous by the very definition of entropy adjusted potential function in (4.32). Hence, the global maximizer of φ_η on the set of absolutely continuous probability measures will also be the global maximizer on $\mathscr{P}(\mathscr{S})$.

limit $\eta \downarrow 0$. Notice from (4.37) that as $\eta \downarrow 0$, the entropy adjusted potential function converges to the true potential function. Therefore, it would be natural to expect that the global maximizer of the entropy adjusted potential would also converge to the global maximizer of the true potential. We state the result formally as follows.

Theorem 9 *Suppose that the population game $\pi : \mathscr{P}(\mathscr{S}) \to \mathbb{L}^\infty(\mathscr{S})$ is a potential game with potential function φ as defined in Definition 17. For perturbation parameter $\eta > 0$ and $n, N \geqslant 1$, let $(\boldsymbol{X}^{n,N}_{\eta,\tau})_{\tau \geqslant 1}$ be the direct exponential clever stochastic process (4.31) induced by the FPFSD game $\mathscr{D}^{n,N}_\pi$. Suppose that $\Xi^{n,N}_\eta$ denotes the stationary distribution of the embedded clever stochastic process $(\bar{\beta}_n(\boldsymbol{X}^{n,N}_{\eta,k}))_{k \geqslant 1}$. Then the following statement holds true:*

$$\lim_{\eta \downarrow 0} \lim_{n \to \infty} \lim_{N \to \infty} \sup_{\bar{\mu}^{n,N} \in \bar{\mathscr{P}}_{n,N}(\mathscr{S})} \frac{1}{n \log 2} \left| \frac{\eta}{N} \log \Xi^{n,N}_\eta(\bar{\mu}^{n,N}) - 2\eta n \log 2 - \Delta\varphi(\bar{\mu}^{n,N}) \right| = 0. \tag{4.40}$$

Proof 16 *See Appendix C.6.*

As with Theorem 8, it will be useful to express (4.40) equivalently as

$$\Xi^{n,N}_\eta(\bar{\mu}^{n,N}) \approxeq \exp\left[\eta^{-1} n N \log 2 \left(\frac{\Delta\varphi(\bar{\mu}^{n,N})}{n \log 2} + \mathscr{E}_{n,N,\eta}(\bar{\mu}^{n,N}) \right)\right], \quad \text{uniformly in } \bar{\mathscr{P}}_{n,N}(\mathscr{S}). \tag{4.41}$$

Here $\mathscr{E}_{n,N,\eta}$ denotes a term which converges uniformly in $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ to 0 as N followed by n approaches infinity finally followed by η approaching 0. Again, as in Theorem 8, once we take the double limit $\lim_{n\to\infty} \lim_{N\to\infty}$, the FPFSD game $\mathscr{D}^{n,N}_\pi$ begins to resemble our original game π. When we take the additional limit $\eta \downarrow 0$, we can apply Theorem 9 to the actual potential function φ of π, as can be seen from the $\Delta\varphi$ term in (4.40). By (4.33), $\Delta\varphi(\bar{\mu}^{n,N}) \leq 0$ with equality only at the global maximizer of φ in $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$. Hence, $\Xi^{n,N}_\eta(\bar{\mu}^{n,N}) \approx \exp\left(\eta^{-1} N \Delta\varphi(\bar{\mu}^{n,N})\right) \to 0$ as $N \to \infty$ at all points except the global maximizer on $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$. The entire mass of this stationary distribution then becomes concentrated on the global maximizer of φ on $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$.

But how does the global maximizer of φ on $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ relate to the global maximizer on our original state space $\mathscr{P}(\mathscr{S})$, which is our real object of interest. Recall that as N followed by n approaches infinity, $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ approaches not $\mathscr{P}(\mathscr{S})$ but the subset of absolutely continuous probability measures under the variational norm. Also recall all maximizers of φ on $\mathscr{P}(\mathscr{S})$ are Nash equilibria. Therefore, if the global maximizer of φ on $\mathscr{P}(\mathscr{S})$ is a Nash equilibrium with a density function, then we are done. By Theorem 9, the stationary distribution $\Xi^{n,N}_\eta$ would indeed put all its mass on the global maximizer of φ on $\mathscr{P}(\mathscr{S})$. But the potential maximizing Nash equilibrium may not have a density function. This would be the case with a pure or monomorphic equilibrium. An argument like footnote 24 will then not work. Instead, we need to argue that the set of absolutely continuous step probability measures is dense in $\mathscr{P}(\mathscr{S})$. Hence, $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ approaches $\mathscr{P}(\mathscr{S})$ but under the weak topology. Therefore, in that case, the global maximizer of φ on $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ would approximate arbitrarily well the global maximizer on $\mathscr{P}(\mathscr{S})$ under the weak topology.

Theorem 9 is our main equilibrium selection result for the potential game π. Once we take the triple limit in (4.40), the stationary distribution of the direct exponential stochastic process (4.31) will put all its mass on or near the Nash equilibrium that globally maximizes the potential function φ of π. This conclusion holds irrespective of the initial state of the process. But this selection holds only in the long run. Thus, as $\tau \to \infty$, the proportion of time the stochastic process spends on the globally potential maximizing Nash equilibrium goes to 1.

4.6 Discussion and Conclusion

We now apply our results in Section 4.5 to the examples we had discussed in Section 4.3. Both Examples 6 and 7 are continuous strategy potential games. Thus, Example 6 is an aggregative potential game with potential function (4.10) while Example 7 is a doubly symmetric game with potential function is $\frac{1}{2}\bar{\pi}(\mu)$, where $\bar{\pi}(\mu)$ is the aggregate payoff of the game. We had noted in Section 4.3 that both these examples have multiple Nash equilibria. Further all those Nash equilibria that are maximizers, local or global, are locally stable under deterministic evolutionary dynamics in the medium run. Therefore, one question we had raised in Section 4.3 was whether stochastic evolutionary analysis can provide an exact equilibrium selection result in continuous strategy potential games. Theorem 9 provides the answer in the affirmative. It shows that the direct exponential stochastic process (4.31) will converge globally to the Nash equilibrium that is the global maximizer of the respective potential functions in both these examples.

Our second question in Section 4.3 was whether exact equilibrium selection also implies selection of the Pareto optimal Nash equilibrium. We address this question for Example 6 first. It turns out that the answer is not necessarily so and it follows from Propositions 7 and 8. Part 3 of Proposition 7 shows that in any aggregative potential game as discussed in Example 6, the Pareto optimal Nash equilibrium is δ_1. But Part 2 of Proposition 8 shows that δ_1 is not always the globally potential maximizing Nash equilibrium. Only if $\alpha^* < \frac{1}{2}$ is δ_1 also globally potential maximizing. In contrast, if $\alpha^* > \frac{1}{2}$, it is the Pareto inferior Nash equilibrium δ_0 that maximizes the potential function globally. Since stochastic evolution selects the Nash equilibrium that is the global maximizer of the potential function, it follows that *exact equilibrium selection in Example 6 implies selection of the Pareto optimal equilibrium only if* $\alpha^* < \frac{1}{2}$.

But α^* is also the aggregate strategy level of one of the three Nash equilibria of Example 6 (part 1 of Proposition 7). For convenience, we refer to α^* as the "interior equilibrium" as it lies in $(0,1)$ and is, hence, in the interior of $\mathscr{S}$. Figure 4.1 illustrates this equilibrium as the middle intersection of $BR(\alpha)$ and α. We can provide some intuition for why this interior equilibrium plays a role in the selection of the Pareto optimal Nash equilibrium by considering random matching in a two–player symmetric game with two strategies, 1 and 2. Let z be the proportion of players playing strategy 1 so that $1-z$ is the proportion of players using strategy 2. Further, suppose the game has three Nash equilibria, $z=0$, $z=1$ and $z=z^* \in (0,1)$. One such game is the stag–hunt game an example of which is

$$\begin{pmatrix} 3 & 0 \\ 2 & 2 \end{pmatrix}. \tag{4.42}$$

This matrix depicts the payoffs of the row player. The payoff of strategy 1 is $3z$ while that of strategy 2 is $2z+2(1-z)=2$. Hence, it is easily verified that $z^* = \frac{2}{3}$.

Consider now the best response revision protocol in this stag–hunt game. It is easily verified that under such a protocol, $(0, z^*)$ is the basin of attraction of the $z=0$ equilibrium while $(z^*, 1)$ is the basin of attraction of the $z=1$ equilibrium.[25] Stochastic evolutionary processes like best response with mutation selects the Nash equilibrium with the larger basin of attraction in two–strategy symmetric games like (4.42) (Kandori et al. (1993), Young (1993)). Intuitively, if $z^* > \frac{1}{2}$, then $z=0$ would have a bigger basin of attraction. Hence, best response with mutation would select $z=0$ as it would require a larger number of mutations to escape that equilibrium than $z=1$. It, therefore, becomes easier for process to escape $z=1$ due to which, its stationary distribution puts zero weight on that Nash equilibrium. In two–strategy games like (4.42), an equilibrium with the larger basin of attraction is also the risk dominant equilibrium.

A similar insight applies to Example 6 as well but at the level of the aggregate strategy. Even though there is a continuum of strategies, there are only three levels of aggregate strategy that are relevant for equilibrium selection. They are 0, 1 and the "interior equilibrium" $\alpha^* = \frac{d}{b-2k}$. This is akin to the three Nash equilibria in the two–strategy symmetric game (4.42). Consider now Figure 4.1 and the two intervals $(0, \alpha^*)$ and $(\alpha^*, 1)$. We interpret these two intervals as the basins of attraction of the two equilibrium levels of aggregate strategy 0 and 1 respectively under the best response protocol. Suppose the present population state is μ such that the aggregate strategy is $A(\mu) = \alpha$. The best response protocol would require all agents to respond to μ with the best response $BR(\alpha)$ defined in (4.11). Figure 4.1 shows that if $\alpha > \alpha^*$, $BR(\alpha) > \alpha$ so that if agents follow the best response protocol, the aggregate strategy level would be pushed higher towards 1. Conversely, if $\alpha < \alpha^*$, $BR(\alpha) < \alpha$ so that the best response protocol would further reduce α towards 0. Thus, $(0, \alpha^*)$ and $(\alpha^*, 1)$ are the basins of attraction of 0 and 1 respectively.

A direct exponential protocol like the logit protocol (4.29) may be seen as a perturbation of the best response protocol. Agents put nearly the entire weight on the best response. Due to best response, there is a natural tendency to move towards the two equilibrium levels of $\alpha = 0$ or $\alpha = 1$. But the irreducible nature of the process causes the system to keep escaping these two equilibria. Following the intuition of the two–strategy game (4.42), we can then argue that it is more difficult for the stochastic process to escape the equilibrium with the larger basin of attraction. If $\alpha^* > \frac{1}{2}$, that equilibrium is $\alpha = 0$ and, hence, stochastic evolution would select $\alpha = 0$. This is the case in the left panel of Figure 4.1 where $\frac{d}{b-2k} = \frac{2}{3} > \frac{1}{2}$. On the other hand, if $\alpha^* < \frac{1}{2}$, then stochastic evolution would select $\alpha = 1$. The right panel of Figure 4.1 where $\frac{d}{b-2k} = \frac{2}{5} < \frac{1}{2}$ illustrates this case.

But as noted earlier, this is the exact conclusion we obtain from the analysis of the potential function. Therefore, we now have an intuitive understanding of when our model of stochastic evolution selects the Pareto optimal Nash equilibrium δ_1 in Example 6. It does when that equilibrium has the larger basin of attraction under the best response protocol at the level of the aggregate strategy. This condition, in turn, is satisfied when the aggregate strategy level at the interior equilibrium, $\alpha^* < \frac{1}{2}$.[26]

---

[25] This is because if $z \in (0, z^*)$, then the best response to z is 0 whereas if $z \in (z^*, 1)$, then the best response is 1.

[26] In the non–generic case when $\alpha^* = \frac{1}{2}$, the stationary distribution will put an equal weight of $\frac{1}{2}$ on both 0 and 1.

We now consider the same question for the doubly symmetric game in Example 7. Thus, we ask whether in this game, our stochastic evolutionary model will select the Pareto optimal equilibrium. Here, the answer is a definite yes and it follows from the last paragraph of Section 4.3. The global maximizer of the potential function $\frac{1}{2}\bar{\pi}(\mu)$ is not just a Nash equilibrium, it is the Pareto optimal state in the game. Therefore, in a doubly symmetric game, the selection of the globally potential maximizing Nash equilibrium by stochastic evolution always implies selection of the Pareto optimal equilibrium.

The overall insight we receive is, therefore, mixed. In certain continuous strategy potential games like doubly symmetric games, the direct exponential stochastic process uniquely selects the Pareto optimal Nash equilibrium in the long run. But in other games like aggregative potential games, that may not happen. Therefore, stochastic evolution may well lead to the selection of the Pareto inferior Nash equilibrium in the long run.

We conclude with a few remarks on related research questions for the future. One criticism of stochastic evolutionary models is that the waiting time required for the exact equilibrium selection result to be realized can be astronomically long. This is particularly so when the population size is reasonably large and the noise parameter is reasonably low.[27] The intuition behind this argument is that a large number of agents have to simultaneously play a non–best response in order to escape the basin of attraction of a Nash equilibrium. The likelihood of that happening diminishes exponentially as the population size grows and the noise parameter, which measures the probability of playing a non–best response, becomes small. In view of this complication, the applicability of stochastic evolution to problems of social or economic interest may be limited. Instead, it is the medium run deterministic evolutionary methods that is more appropriate for such applications, even if such methods require us to sacrifice the possibility of exact equilibrium selection. One way to resolve this dilemma is to retain stochastic evolution but assume agents interact with their neighbors in a network instead of the whole population. Existing results on finite strategy stochastic models with such local interaction show that waiting times for equilibrium selection can go down significantly (Ellison (1993)). It would be of interest to see whether similar results can also be obtained with local interaction in a continuous strategy games.

Another research question would be to extend the present results to other strategy revision protocols. Apart from direct exponential protocols, Sandholm (2010) also considers stochastic evolution under imitative exponential revision protocols and establishes equilibrium selection results in finite strategy potential games under such protocols.[28] The main difference is that in direct exponential protocols, players choose a strategy on the basis of the payoffs of different strategies while in imitative exponential protocols, an agent imitates the strategy of a randomly chosen opponent. Our conjecture is that the approximation approach of the present paper can also be applied to such imitative exponential protocols. Details of the analysis would, of course, vary and we leave that as a future exercise.

[27] See Section 11.3.2 in Sandholm (2010) for a discussion of this issue.

[28] See Theorem 11.5.13 in Sandholm (2010).

CHAPTER

5

REGULARIZED BAYESIAN BEST RESPONSE LEARNING IN FINITE GAMES

5.1 Introduction

5.1.1 Background and Motivation of the problem

Evolutionary game theory is primarily concerned with the behavior of large populations of strategically interacting players. In such games, an evolutionary dynamic governs the behavior of the population dynamically as a function of (discrete or continuous) time. The standard literature typically assumes that the players in the underlying population are homogeneous, that is, are of the same type (Hofbauer and Sandholm (2002), Hofbauer and Sandholm (2007), Sandholm (2010)). Such homogeneity is sometimes referred to as anonymity in the game theoretic literature. However, in real-life situations, the population can be heterogeneous in nature, consisting of players with diverse preferences/utilities. For instance, depending on the infrastructure, labor force, and resource contracts, different players participating in a Cournot competition may have different payoff functions. Similarly, in a prisoner's dilemma game, different individuals may incur different costs for the same number of years of imprisonment due to factors such as age, health, and socioeconomic status. In view of this, we attempt to generalize evolutionary games to accommodate heterogeneity in the population. We call such games Bayesian population games.

Best response is one of the most standard behavioral protocols considered for homogeneous population evolutionary games. However, the multiplicity of the best response map makes it difficult to analyze the evolutionary dynamic (Gilboa and Matsui (1991), Hofbauer (1995)). To circumvent the problem, several authors have considered the uniquely defined perturbed (smoothed) version of the best response. Extrapolating this idea, we introduce the notion of perturbed Bayesian best response for heterogeneous population games, which we call regularized Bayesian best response (RBBR) dynamic. In RBBR, the expected payoff of a sub-population (consisting of players of a particular type) is perturbed by a regularizer (for eg. an entropy function), and the sub-population best responds to this perturbed expected payoff.

Regularized best response dynamics are one of the most important classes of dynamics in the literature of evolutionary game theory. The logit dynamic (Fudenberg and Levine (1998)) is a standard example of such an evolutionary dynamic for homogeneous population games. In this dynamic, the expected payoff at a population state is regularized by perturbing via the well-known Shannon-Gibbs entropy. In practice, however, the structure of the perturbation may be unknown (Coucheney et al. (2015), Lahkar et al. (2022b)). To the best of our knowledge, regularized best response have been studied primarily in the context of homogeneous population games. This paper is an attempt to extend the notion of perturbed best response dynamics to the class of heterogeneous (Bayesian) games. (see Section 5.2.2 for some examples of such games). To model such a situation for heterogeneous population games, instead of assuming a particular form of the perturbation function, RBBR learning dynamic is generated via perturbation by an arbitrary lower semicontinuous and strongly convex function. Examples of such regularizers include the Tsallis entropy and the Burg entropy, which have appeared in the literature of deterministically perturbed best response dynamics (Hofbauer and Sandholm (2002)), stochastically perturbed best response dynamics (Hofbauer and Sandholm (2007)), as well as in the context of online non-convex optimization problems (Héliou et al. (2020)).

The brief heuristics of our model is borrowed from Ely and Sandholm (2005). We consider a population where the players are distributed over a complete separable metric space consisting of types, equipped with a pre-specified probability distribution. We call such a probability space the type space of the population. The state space of the RBBR learning dynamic is the space of Bayesian strategies defined as the measurable maps from the type space to the collection of all probability distributions over the strategies. To define the RBBR dynamic, we extend Bayesian strategies to integrable signed Bayesian strategies defined as the measurable maps from the type space into $\mathbb{R}^n$, and consider the normed linear space of these strategies. Such an extension is necessary so that the learning dynamic lives in an underlying vector space. Following Ely and Sandholm (2005), we equip the space of integrable signed Bayesian strategies with the strong norm, which turns it into a Banach space (Diestel and Uhl Jr (1978)). We define the RBBR dynamic as the instantaneous change in the population state occurred due to its regularized Bayesian best response (to the current state). More formally, the RBBR dynamic is defined as the difference between the regularized Bayesian best response to the current population state and the current population state. Since the difference between two Bayesian strategies is a signed Bayesian strategy, we require the RBBR to be defined on the extended space of signed Bayesian strategies. In this paper, we study several important properties of the RBBR dynamic, which we explain in detail in the next section.

5.1.2 Our contribution

An overview of our results is as follows. First, we establish the existence of rest points of the RBBR learning dynamic. These rest points are fixed points of the regularized Bayesian best response map and which we call regularized Bayesian equilibria. Next, we show that these regularized equilibria act as approximations to the Bayesian equilibria of the original game for vanishingly small regularizations. This approximation result is non-trivial and requires one of the two conditions, namely Condition C1 or Condition C2, on the regularizers apart from strong convexity. Condition C1 ensures that the population concentrates on the strategies that provide maximum payoff as the regularization vanishes, and Condition C2 requires the regularizer to satisfy some mild continuity conditions in the sense that small changes in the population state can only lead to a small changes in the value of the regularizer. We show that the Shannon-Gibbs entropy satisfies Condition C1, whereas, the Tsallis and Burg entropy satisfy Condition C2 (See Appendix D.7).

Next, we prove the existence of a unique continuous solution trajectory to the RBBR dynamic. We further show that the solutions to the RBBR dynamic are continuous with respect to the initial states (that is, if we start from two initial states that are very close, then the respective solutions to the RBBR dynamic will also be close enough).

Finally, we introduce the notions of Bayesian potential and Bayesian negative semidefinite games and provide results related to the convergence of the RBBR dynamic to the regularized Bayesian equilibrium for such games. The Bayesian potential games are heterogeneous population extensions of standard homogeneous potential games. The notion of such games have been studied very recently by Zusai (2023) in the context of mean-field revision protocols. To the best of our knowledge, our paper is the first to consider the heterogeneous version of negative semidefinite games.[1]

We provide a brief discussion on the technical challenges in deriving our results. Unlike evolutionary games with finitely many strategies where the state space of the underlying dynamic is finite dimensional, the space of Bayesian strategies on which the RBBR learning dynamic is defined is an infinite dimensional Banach space. As a result, the RBBR dynamic is a non trivial extension of the existing results on homogeneous population games with finitely many strategies. In particular, the choice of topology on the state space plays an important role in deriving the results. For example, the weak topology on the space of integrable signed Bayesian strategies turns out to be the appropriate topology to establish results related to the existence of regularized Bayesian equilibrium. Whereas, the strong topology (induced by the strong norm) is necessary to show the existence of solution trajectories to the RBBR learning dynamic. Establishing convergence results for Bayesian potential and Bayesian negative semidefinite games requires a further exercise of characterizing (strong) norm compact and forward invariant subsets of the space of signed Bayesian strategies. This is subsequently used for constructing a suitable Lyapunov function so as to apply relevant results from the theory of dynamical systems (Bhatia and Szegö (2006)) for proving convergence results related to Bayesian potential and Bayesian negative semidefinite games.

[1] See Monderer and Shapley (1996), Cheung (2014), Lahkar and Riedel (2015), and Lahkar et al. (2022b) for a complete exposition of homogeneous population potential and negative semidefinite games.

5.1.3 Literature review

Evolution in heterogeneous (Bayesian) games has not been explored much in the literature of evolutionary game theory. The theory was first introduced by Ely and Sandholm (2005) to study the Bayesian best response dynamic when players have diverse preferences. The authors, instead of considering regularization of the Bayesian best response, provided sufficient conditions under which the Bayesian best response correspondence is unique (single-valued). In another result, they show the equilibrium behavior and the stability of the Bayesian best response dynamic can be fully characterized by the equilibrium behavior and the stability of the aggregation dynamic. The aggregation dynamic is obtained by taking the expectation of Bayesian strategies over types.

Later, Sandholm et al. (2007) studies the evolutionary stability of purified equilibria of two-player normal form games under the Bayesian best response dynamic. Staudigl (2013) considers a Bayesian model for network evolution where players have diverse preferences and provides a closed-form expression for the invariant distribution of the underlying co-evolutionary process for potential games. Very recently in Zusai (2023), the evolutionary dynamics in heterogeneous populations is obtained via a mean dynamic based on revision protocol functions. Under some regularity conditions on the revision protocol function, they show that the evolutionary dynamic admits a unique solution trajectory. They further establish the equilibrium stability in heterogeneous potential games.

Our approach is different from Zusai (2023) as we generate the best response evolutionary dynamic via regularization of the expected payoff of the population. We do this mainly due to two major reasons. Firstly, a vast amount of literature in evolutionary game theory is devoted specifically to studying the perturbed best response learning (Fudenberg and Kreps (1993), Kaniovski and Young (1995), Benaım and Hirsch (1999), Lahkar et al. (2022b)), and secondly, to address questions related to the existence and approximation of regularized equilibria, and convergence of RBBR dynamic in certain classes of games, namely Bayesian potential and Bayesian negative semidefinite games.

5.1.4 Organization of our paper

The paper is organized as follows. Section 6.2 contains the preliminaries of the model. In Section 5.3, we first introduce the notion of a regularizer, followed by the definition of regularized Bayesian best response function, and the associated RBBR learning dynamic. We then establish the existence of rest points of the RBBR learning dynamic, in particular, the existence of regularized Bayesian equilibrium followed by approximation results concerning regularized Bayesian equilibria. In Section 5.4, we establish some fundamental properties of the RBBR learning dynamic subject to some conditions on the type space. As applications of these results, we prove convergence results pertaining to Bayesian potential games and Bayesian negative semidefinite games in Section 5.5 and Section 5.6, respectively.

5.2 Preliminaries

5.2.1 Basic notations

For $n \geqslant 1$, let $\mathbb{R}^n$ denote the n-dimensional Euclidean space endowed with its Borel sigma algebra $\mathscr{B}_n$. Let (Ω, d_Ω) be a complete separable metric space equipped with its Borel sigma algebra $\mathscr{B}_\Omega$ and a Borel probability measure ξ. Let $S := \{1, \dots, n\}$ and denote by Δ be the collection of all probability distributions on S, identified by the simplex $\Delta := \{\mathbf{x} = (x_1, \dots, x_n) \in [0,1]^n : x_1 + \cdots + x_n = 1\}$. We use the notations Δ° and $\partial\Delta$ to denote the relative interior and boundary of Δ, respectively. Let $\mathscr{C}_n(\Delta)$ be the collection of all continuous maps from Δ into $\mathbb{R}^n$ endowed with the topology of uniform convergence.[2]

5.2.2 Model

Consider a large population of heterogeneous players whose mass is normalized to unity. By a heterogeneous population, we mean that the players are of different types. We assume that the players are distributed over Ω according to the given probability measure ξ. The probability measure ξ thus accounts for the degree of heterogeneity within the population. We call the probability space (Ω, ξ) the **type space**, and the probability measure ξ, the **type measure** of the population. The set S denotes the strategy space of the players in the population. In homogeneous population games, a probability distribution $\mathbf{x} \in \Delta$ denotes the social state/aggregate behavior of the population. When we say that the aggregate behavior of the population is $\mathbf{x}$, we mean that $\mathbf{x}_j$ proportion of players in the population are exercising strategy $j \in S$. However, for heterogeneous populations, we consider a refined version of aggregate behavior called a Bayesian strategy.[3]

A *Bayesian strategy* is a ξ-measurable mapping $\sigma : \Omega \to \Delta$, where $\sigma(\omega)$ denotes the belief (distribution) induced by the players of type ω over the strategy space S.[4] Loosely speaking, a Bayesian strategy can be viewed as a random vector defined on the type space (Ω, ξ), taking values in the simplex Δ. We say that two Bayesian strategies σ and ρ are equivalent and write $\sigma \sim \rho$ if $\sigma(\omega) = \rho(\omega)$, for ξ-a.e $\omega \in \Omega$. Thus, with an abuse of notation, we denote the equivalence classes of all Bayesian strategies by Σ itself. We conclude this section with the definition of Bayesian population games followed by some examples on Bayesian population games. As we shall observe later, Bayesian population games are extensions of homogeneous population games to heterogeneous populations.

Definition 21 *A* ***Bayesian population game*** *is a map $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ such that for every $\sigma \in \Sigma$, the mapping*

$$\omega \mapsto \mathscr{G}(\sigma, \omega) \qquad \textit{is } \mathscr{B}_\Omega / \mathscr{B}_n \text{ measurable.}$$ [5]

The definition of Bayesian population games has the following interpretation. For a Bayesian strategy $\sigma \in \Sigma$, $\mathscr{G}(\sigma, \omega)$ denotes the payoff vector of the players in the population of type $\omega \in \Omega$. More precisely, if

[2]Suppose $(f_m)_{m \geqslant 1}, f \in \mathscr{C}_n(\Delta)$. Then $f_m \to f$ uniformly iff $\|f_m - f\|_\infty := \sup_{x \in \Delta} \|f_m(x) - f(x)\| \to 0$ as $m \to \infty$.

[3]The notion of a Bayesian strategy was first introduced in the context of evolutionary games by Ely and Sandholm (2005).

[4]See Appendix D.1, Definition 50.

[5]Here $\mathscr{B}_\Omega / \mathscr{B}_n$ measurable refers to measurability of $\mathscr{G}(\sigma, \dot{})$ seen as a map from $(\Omega, \mathscr{B}_\Omega)$ into $(\mathscr{R}^n, \mathscr{B}_n)$.

the population is playing Bayesian strategy σ, then a player of type ω exercising strategy i receives a payoff $\mathscr{G}^i(\sigma,\omega)$. Zusai (2023) calls such a game a heterogeneous population game. The following is a special example of a Bayesian population game, where the space of types Ω are identified by the payoff functions (population games) which the players in the population use to calculate their payoff. We call such a game, a Bayesian aggregative population game. Such games have been considered earlier in Ely and Sandholm (2005). We clarify that our definition of a Bayesian population game has a more general structure and all our results in the subsequent sections of the paper hold true for Bayesian populations games with the general set-up.

Example 8 (Bayesian aggregative population games, Ely and Sandholm (2005)) *In Bayesian aggregative population games, the type space of the agents in the population is the collection of continuous maps from Δ to $\mathbb{R}^n$. Mathematically, we have $(\Omega, d_\Omega) := (\mathscr{C}_n(\Delta), \|\cdot\|_\infty)$, where $\|\cdot\|_\infty$ denotes the sup-norm metric, and is defined as $\|\omega-\bar{\omega}\|_\infty := \sup_{x\in\Delta}\|\omega(\boldsymbol{x})-\bar{\omega}(\boldsymbol{x})\|$ for all $\omega,\bar{\omega}\in\Omega$. Since Δ is compact, we have that $\mathscr{C}_n(\Delta)$, and thus, Ω is a complete and separable metric space. Let ξ be a probability measure on Ω. Note that from the standard theory of evolutionary games, the elements of Ω are population games. The expectation operator is a mapping $\mathscr{E}:\Sigma\to\Delta$ such that for every $\sigma\in\Sigma$,*

$$\mathscr{E}^i(\sigma) := \int_\Omega \sigma^i(\omega)\xi(d\omega), \quad \textit{for all } i\in S.^{6}$$

Thus, given $\sigma\in\Sigma$, $\mathscr{E}(\sigma)$ denotes the aggregate belief of the population over the strategies with respect to the probability measure ξ. Following Ely and Sandholm (2005), we define the Bayesian aggregative population game as a mapping $\mathscr{G}:\Sigma\times\Omega\to\mathbb{R}^n$ such that for every $\sigma\in\Sigma$,

$$\mathscr{G}(\sigma,\omega) := \omega(\mathscr{E}(\sigma)), \qquad \textit{for all } \omega\in\Omega.$$

The definition of Bayesian aggregative population game has the following interpretation. Given a Bayesian strategy σ, $\omega^i(\mathscr{E}(\sigma))$ denotes the payoff to a player of type $\omega\in\Omega$ exercising strategy $i\in S$ at aggregate population state $\mathscr{E}(\sigma)$. Thus, the payoff to a player of type ω exercising strategy i depends only on the aggregate belief $\mathscr{E}(\sigma)$ over the strategies. Note that $\mathscr{G}$ above is well-defined since for every $\sigma\in\Sigma$, the mapping $\omega\mapsto\omega(\mathscr{E}(\sigma))$ is continuous, and hence measurable.

Example 9 (Asymmetric War of Attrition Game) *The (symmetric) war of attrition game is a model to mathematically analyze conflicts between different species of the animal population. It was first introduced from a biological perspective by Maynard Smith (1982). In such a game, two animals (players) are engaged in competition for an indivisible resource. A player's strategy is to decide whether to dispute (exit) given that the other player is still involved in the conflict. The player who persists longer is the winner of the game. For such games, evolutionary stable strategies (ESS) were introduced and examined by Maynard Smith (1982). The concept of a "generalized war of attrition game" with a general payoff structure was later presented by Bishop and Cannings (1978), who also obtained mixed evolutionary*

[6] Since the Bayesian strategies take values in the compact set Δ, the expectation operator $\mathscr{E}$ is well–defined.

stable strategies for such games. However, the nature of the war of attrition game can be asymmetric (heterogeneous) in real-world situations. Such asymmetry within the population is denoted by Hammerstein and Parker (1982) as 'roles' and 'situations', respectively. We refer to such heterogeneity as 'types' in this paper. As a result, the roles and situations (and thus, the types) in which the game is being played determine the value and cost of the resources.

Suppose that a player with type ω encounters a player with type $\bar{\omega}$. Then the value and cost of the game corresponding to this environment are $V(\omega,\bar{\omega})$ and $C(\omega,\bar{\omega})$, respectively. As a result, the payoff to the player of type ω persisting time i against the player of type $\bar{\omega}$ persisting time j is given by the following payoff function:

$$\pi^{ij}(\omega,\bar{\omega}) := \begin{cases} V(\omega,\bar{\omega}) - jC(\omega,\bar{\omega}) & \text{if } i > j \\ \frac{V(\omega,\bar{\omega})}{2} - iC(\omega,\bar{\omega}) & \text{if } i = j \\ -iC(\omega,\bar{\omega}) & \text{if } i < j. \end{cases}$$

Thus, if we consider the random matching version of π in the Bayesian setup, we have the following asymmetric war of attrition game $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ such that for all $\sigma \in \Sigma$ and all $\omega \in \Omega$,

$$\mathscr{G}^i(\sigma,\omega) := \sum_{j=1}^{n} \int_{\Omega} \pi^{ij}(\omega,\bar{\omega}) \sigma_j(\bar{\omega}) \xi(d\bar{\omega}) \quad \text{for all } i \in S.$$

Example 10 (Additively Separable Anonymous Game) *Another example of a Bayesian population game that can be found in the literature on discrete choice models (Anderson et al. (1992)) is the additively separable anonymous game, or ASAG (Zusai (2023), Pg. 1221, Example 1). An additive idiosyncratic shock that solely depends on the type of player introduces heterogeneity in the payoff in this class of games. There is an additive relationship between the Bayesian strategy of the population and the type of the player.*

To illustrate the game mathematically, let us suppose that the type space Ω of the population is $\mathbb{R}^n$. More formally, a player of type ω exercising strategy i induces an idiosyncratic shock ω^i to the payoff of the game. Let $\pi : \Delta \to \mathbb{R}^n$ be a given homogeneous population game. Thus, an ASAG $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ is defined as follows: for a Bayesian strategy $\sigma \in \Sigma$, and a type $\omega \in \mathbb{R}^n$, the payoff is given by

$$\mathscr{G}^i(\sigma,\omega) := \pi^i(\mathscr{E}(\sigma)) + \omega^i \quad \text{for all } i \in S,$$

where $\mathscr{E}(\theta)$ denotes the expectation of the Bayesian strategy θ. The payoff at a Bayesian strategy σ and a type ω in an additively separable anonymous game comprises of a common payoff (homogeneous component) and a type dependent idiosyncratic shock (heterogeneous component). As a result, the overall payoff function of the population is heterogeneous.

Example 11 (Incomplete Information Game) *According to Zusai (2023), Pg. 1221, Example 2, an incomplete information game is also an example of a Bayesian population game. Formally, let us consider a finite set of signals $\Theta = \{\theta_1, \ldots, \theta_k\}$. Let μ be a common prior belief on Θ. Let Ω represent the type space of the population. Given a signal $\theta_j \in \Theta$, we denote the conditional distribution over Ω by $\nu(\cdot|\theta_j)$. The* *posterior belief on* Θ *can therefore be obtained by using the Bayes' updation rule via the following relation:*

$$\mu(\theta_j|\omega) \propto \nu(\omega|\theta_j) \times \mu(\theta_j) \quad \textit{for } j = 1, \dots, k. \tag{5.1}$$

Furthermore, let $\mathscr{G}_{\theta_j} : \Sigma \times \Omega \to \mathbb{R}^n$ *be the Bayesian population game corresponding to signal* $\theta_j \in \Theta$*, where* $j = 1, \dots, k$*. The expected payoff at the signal dependent Bayesian population games* $(\mathscr{G}_{\theta_j})_{j=1,\dots,k}$ *with respect to the posterior belief on* Θ *in (6.3) defines the Bayesian game for the population. Thus, mathematically, we define the Bayesian population game as a map* $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ *such that*

$$\mathscr{G}(\sigma, \omega) := \sum_{j=1}^{k} \mathscr{G}_{\theta_j}(\sigma, \omega)\mu(\theta_j|\omega) \quad \textit{for all } \sigma \in \Sigma \textit{ and } \omega \in \Omega.$$

5.3 Regularization of Bayesian best response

In this section we seek to introduce the notion of regularized Bayesian best response learning dynamic. This, in turn, requires us to define the notion of a Bayesian best response map. It is an extension of best response maps to the context of Bayesian (heterogeneous) population games. For a Bayesian strategy $\sigma \in \Sigma$, the expected payoff to a player of type ω against an aggregate population state (aggregate belief) $\mathbf{y} \in \Delta$ is given by $\langle \mathbf{y}, \mathscr{G}(\sigma, \omega) \rangle$, where $\langle \cdot, \cdot \rangle$ denotes the usual inner product on $\mathbb{R}^n$. Quite naturally, the **Bayesian best response correspondence** is a mapping $\beta : \Sigma \twoheadrightarrow \Sigma$ such that for all $\sigma \in \Sigma$,

$$\beta[\sigma](\omega) := \arg\max_{\mathbf{y} \in \Delta} \langle \mathbf{y}, \mathscr{G}(\sigma, \omega) \rangle \quad \text{for all } \omega \in \Omega. \tag{5.2}$$

Note that from the definition of Bayesian population games, $\mathscr{G}(\sigma, \omega)$ denotes the payoff vector a player of type ω faces at Bayesian strategy σ. Thus, the best response to a player of type ω at Bayesian strategy σ is any probability distribution $\mathbf{y}$ on S which maximizes the expected payoff $\langle \mathbf{y}, \mathscr{G}(\sigma, \omega) \rangle$.

A Bayesian strategy $\sigma^\circ \in \Sigma$ is said to be a **Bayesian equilibrium** of the game $\mathscr{G}$ if $\sigma^\circ \in \beta(\sigma^\circ)$. In terms of the game $\mathscr{G}$, the above condition can equivalently be expressed in the following way: σ° is a Bayesian equilibrium if for ξ-a.e $\omega \in \Omega$,

$$\sigma^\circ_j(\omega) > 0 \implies \mathscr{G}_j(\sigma^\circ, \omega) \geqslant \mathscr{G}_i(\sigma^\circ, \omega), \quad \text{for all } i \in S. \tag{5.3}$$

In principle, the above condition suggests that a Bayesian strategy σ° is an equilibrium provided the following condition holds: if a strategy $j \in S$ has positive probability in the induced belief $\sigma^\circ(\omega)$, then the payoff vector $\mathscr{G}(\sigma, \omega)$ should have a component-wise maximum at strategy j. Thus, strategies in the support of the induced belief must have maximum payoff. Let $\mathrm{BE}(\mathscr{G}, \xi)$ denote the collection of Bayesian equilibria corresponding to the game $\mathscr{G}$. In order to state the result on existence of Bayesian equilibrium, we require the notion of an appropriate topology on the space of Bayesian strategies, which essentially is the weak topology in this case (see Definition 52).

Definition 22 *A Bayesian population game* $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ *is said to be* ***weakly continuous*** *if the mapping* $\sigma \mapsto \mathscr{G}(\sigma, \omega)$ *is continuous relative to the weak topology on* Σ*, for* ξ*-a.e* $\omega \in \Omega$*.*

Proposition 11 *Suppose that the population game $\mathscr{G}$ is weakly continuous. Then $BE(\mathscr{G},\xi)$ is a non-empty and compact subset of Σ.*

Proof 17 *Follows from Kakutani-Glicksberg-Fan theorem (Aliprantis and Border (2013)).*

Suppose now that we fix a type $\omega\in\Omega$. We then define the **game sliced at type** ω as a mapping $\mathscr{G}_\omega:\Sigma\to\mathbb{R}^n$ such that $\mathscr{G}_\omega(\sigma):=\mathscr{G}(\sigma,\omega)$ for all $\sigma\in\Sigma$. In words, the sliced game determines how payoffs behave as a function of Bayesian strategies, keeping a type fixed. Thus, we can view the sliced game $\mathscr{G}_\omega$ as the game corresponding to the underlying sub-population consisting of players of exactly one type ω. We then introduce the ω**-sliced best response** for the game $\mathscr{G}_\omega$ as a mapping $\beta_\omega:\Sigma\twoheadrightarrow\Delta$ such that

$$\beta_\omega(\sigma):=\beta[\sigma](\omega),\quad \text{for all } \sigma\in\Sigma. \tag{5.4}$$

We say that a Bayesian strategy σ° is an ω**-sliced Bayesian equilibrium** of the ω-sliced game $\mathscr{G}_\omega$ if $\sigma^\circ(\omega)\in\beta_\omega(\sigma^\circ)$. We then have the following characterization result for Bayesian equilibria of the original game $\mathscr{G}$ and the equilibria of the sliced game counterparts $\mathscr{G}_\omega$.

Proposition 12 *A Bayesian strategy σ° is a Bayesian equilibrium of the game $\mathscr{G}$ iff σ° is an ω-sliced Bayesian equilibrium of the game $\mathscr{G}_\omega$ for ξ-a.e $\omega\in\Omega$.*

Proof 18 *See Appendix D.2.*

As the title of the paper suggests, we now shift our focus towards the regularized version of the Bayesian best response. Recall that the Bayesian best response is a multi-valued mapping, which, under standard evolutionary dynamics, becomes difficult to analyze. To circumvent this problem, we resort to a regularized (smoothed) version of the Bayesian best response mapping in a view to obtain a unique smoothed best response which can then be analyzed using the existing theory of differential equations and dynamical systems. We begin with the definition of a regularizer.

Definition 23 *A mapping $\boldsymbol{v}:\Delta\to\mathbb{R}\cup\{\infty\}$ is called a* ***regularizer*** *on Δ if:*

- *it is finite except possibly on the relative boundary $\partial\Delta$ of Δ;*
- *it is lower semi-continuous on Δ;*
- *it is strongly convex on Δ°, in the sense that there exists $\gamma>0$ such that for all $0\leqslant\alpha\leqslant1$, and all $\boldsymbol{x},\boldsymbol{y}\in\Delta^\circ$,*

$$\boldsymbol{v}(\alpha\boldsymbol{x}+(1-\alpha)\boldsymbol{y})\leqslant\alpha\boldsymbol{v}(\boldsymbol{x})+(1-\alpha)\boldsymbol{v}(\boldsymbol{y})-\frac{\gamma\alpha(1-\alpha)}{2}\|\boldsymbol{x}-\boldsymbol{y}\|^2.$$

It is well-known (see Coucheney et al. (2015), Mertikopoulos and Sandholm (2016)) that regularization ensures the existence of a unique best response. We bring to notice that our definition of a regularizer is weaker as compared to the one used in the references mentioned just prior, in the sense that regularizers in Coucheney et al. (2015) and Mertikopoulos and Sandholm (2016) are smooth in the interior Δ°; whereas in our case, we only require lower semicontinuity on Δ.

Definition 24 (Regularized Bayesian best response) *Given a noise parameter $\epsilon > 0$, the regularized Bayesian best response (RBBR) is a mapping $\beta_\epsilon : \Sigma \to \Sigma$ such that for every $\sigma \in \Sigma$,*

$$\beta_\epsilon[\sigma](\omega) := \arg\max_{\boldsymbol{y} \in \Delta} \Big[\langle \boldsymbol{y}, \mathscr{G}(\sigma, \omega) \rangle - \epsilon \boldsymbol{v}(\boldsymbol{y}) \Big], \quad \textit{for all } \omega \in \Omega.^{7} \tag{5.5}$$

As we mentioned before, the purpose behind defining the regularized version of the Bayesian best response correspondence is to ensure the existence of a unique maximizer in (5.5), when the expected payoff is perturbed via the regularizer $\mathbf{v}$, thereby making the study of the evolutionary learning dynamic amenable to the existing theory of ordinary differential equations and dynamical systems. Suppose further that the regularizer is steep, in the sense that $\|\nabla \mathbf{v}(\mathbf{x}_n)\| \to \infty$ as $\mathbf{x}_n \to \partial \Delta$. It then follows that for every $\sigma \in \Sigma$, the best response $\beta_\epsilon[\sigma](\omega)$ to type $\omega \in \Omega$ lies in the relative interior Δ° (see Rockafellar (1970), Chapter 26).

Also it worthwhile to mention that the notion of deterministic and stochastically perturbed best response in homogeneous populations with finitely many strategies has been studied in several occasions such as Benaım and Hirsch (1999), Hofbauer and Sandholm (2002), Hofbauer and Hopkins (2005), Hofbauer and Sandholm (2007) to name a few. In the recent past, Mertikopoulos and Sandholm (2016) considers reinforcement learning under regularization. In the context of homogeneous population games with a continuum of strategies, the study of perturbed best response dynamics has been generalized to logit best response in Lahkar and Riedel (2015) and Perkins and Leslie (2014). Very recently, a general class of perturbed best response dynamics generated via an arbitrary perturbation function (regularizer in our case) has been considered in Lahkar et al. (2022b). In this paper, we thus focus on extending the notion of perturbed best response to heterogeneous population games with finitely many strategies.

5.3.1 Some examples of regularizers

Below we provide some examples of widely used regularizers which have been extensively used in the literature of evolution and learning in games, both in the context of finite strategy games and infinite strategy homogeneous population games. They have also been used extensively in reinforcement learning and online non-convex optimization problems (see Mertikopoulos and Sandholm (2016), Héliou et al. (2020)). In the case of continuum strategy games, perturbed best response dynamics using general regularizers have been recently considered in Lahkar et al. (2022b).

Example 12 (The Shannon-Gibbs entropy) *The most classic and prototypical example of a regularizer is the well-known Shannon-Gibbs entropy which is defined as*

$$\boldsymbol{v_s}(\boldsymbol{x}) := \sum_{j \in S} \boldsymbol{x}_j \log \boldsymbol{x}_j, \quad \textit{for all } \boldsymbol{x} \in \Delta. \tag{5.6}$$

[7] The fact that the perturbation function $\mathbf{v}$ is strongly convex implies that for every fixed $\sigma \in \Sigma$, the mapping $\mathbf{y} \mapsto \langle \mathbf{y}, \mathscr{G}(\sigma, \omega) \rangle - \epsilon \mathbf{v}(\mathbf{y})$ has a unique maximum.

As is well-known from the literature (see Fudenberg and Levine (1998)), the regularized Bayesian best response map corresponding to $\boldsymbol{v_s}$ is such that for every $\sigma \in \Sigma$,

$$\beta^{\boldsymbol{s}}_{\epsilon,j}[\sigma](\omega) = \frac{\exp(\epsilon^{-1}\mathscr{G}_j(\sigma,\omega))}{\sum_{i\in S}\exp(\epsilon^{-1}\mathscr{G}_i(\sigma,\omega))} \quad \textit{for all } \omega \in \Omega, \textit{ and all } j \in S. \tag{5.7}$$

Example 13 (The Tsallis entropy) *A well-known generalization of the Shannon-Gibbs entropy is the so-called Tsallis entropy which is defined as*

$$\boldsymbol{v_t}(\boldsymbol{x}) := (q-q^2)^{-1}\sum_{j\in S}(\boldsymbol{x}_j - \boldsymbol{x}_j^q), \quad \textit{for } 0<q<1, \textit{ for all } \boldsymbol{x} \in \Delta. \tag{5.8}$$

Notice that the Shannon-Gibbs entropy can be obtained from the Tsallis entropy by letting $q \to 1$. It can be shown (see for example Coucheney et al. (2015), Héliou et al. (2020)) that under perturbation by $\boldsymbol{v_t}$, the regularized Bayesian best response map takes the following form: for every $\sigma \in \Sigma$,

$$\beta^{\boldsymbol{t}}_{\epsilon,j}[\sigma](\omega) := \Big[\frac{\epsilon}{1-q}\Big]^{1/(1-q)} \frac{1}{(\theta^{\boldsymbol{t}}_{\epsilon}(\sigma,\omega)-\mathscr{G}_j(\sigma,\omega))^{1/(1-q)}}, \quad \textit{for all } \omega \in \Omega, \textit{ and all } j \in S, \tag{5.9}$$

where $\theta^{\boldsymbol{t}}(\sigma,\omega) > \max_{j\in S}\mathscr{G}_j(\sigma,\omega)$ is such that $\beta^{\boldsymbol{t}}_{\epsilon}[\sigma](\omega)$ is a valid probability distribution on S for every $\sigma \in \Sigma$ and for every $\omega \in \Omega$.

Example 14 (The log-barrier entropy) *Another important example of a regularizer is the logarithmic barrier entropy, which is sometimes also called the Burg entropy (see Hofbauer and Sandholm (2002)) and is defined as*

$$\boldsymbol{v_b}(\boldsymbol{x}) := -\sum_{j\in S}\log \boldsymbol{x}_j, \quad \textit{for all } \boldsymbol{x} \in \Delta. \tag{5.10}$$

In case of log-barrier regularization, the closed form expression of the regularized Bayesian best response map is such that for every $\sigma \in \Sigma$,

$$\beta^{\boldsymbol{b}}_{\epsilon,j}[\sigma](\omega) := \frac{\epsilon}{\theta^{\boldsymbol{b}}_{\epsilon}(\sigma,\omega)-\mathscr{G}_j(\sigma,\omega)}, \quad \textit{for all } \omega \in \Omega, \textit{ and all } j \in S, \tag{5.11}$$

where $\theta^{\boldsymbol{b}}_{\epsilon}(\sigma,\omega) > \max_{j\in S}\mathscr{G}_j(\sigma,\omega)$ is such that $\beta^{\boldsymbol{b}}_{\epsilon}[\sigma](\omega)$ is a valid probability distribution on S for every $\sigma \in \Sigma$ and for every $\omega \in \Omega$ (see Héliou et al. (2020) for the continuum strategy counterpart).

Before moving ahead with the notion of regularized Bayesian best response, it is necessary to ensure that β_ϵ is well-defined in the sense that $\beta_\epsilon[\sigma]: \Omega \to \Delta$ is a ξ-measurable mapping for every $\sigma \in \Sigma$. To this end, we have the following proposition.

Proposition 13 *For every $\epsilon > 0$ and $\sigma \in \Sigma$, the mapping $\omega \mapsto \beta_\epsilon[\sigma](\omega)$ is ξ-measurable.*

Proof 19 *See Appendix D.3.*

5.3.2 Regularized Bayesian best response dynamic

In this section, we seek to define the notion of the regularized Bayesian best response dynamic on the space of Bayesian strategies. In what follows, we first extend the space of Bayesian strategies to the space of **integrable signed Bayesian strategies** which is defined below as:

$$\widehat{\Sigma} := \left\{ \widehat{\sigma} : \Omega \to \mathbb{R}^n \ \mid \ \widehat{\sigma} \text{ is } \xi\text{−measurable and } |||\widehat{\sigma}||| := \int_{\Omega} \|\widehat{\sigma}(\omega)\| \xi(d\omega) < \infty \right\}.^{8} \tag{5.12}$$

Definition 25 *Let $\epsilon > 0$ be a noise parameter and let the RBBR map be as defined in (5.5). Then the* ***regularized Bayesian best response learning dynamic*** *is defined as*

$$\dot{\sigma} = \beta_\epsilon[\sigma] - \sigma.^{9} \tag{5.13}$$

From Definition 25, it follows that under the RBBR learning dynamic, a Bayesian strategy σ moves towards its regularized Bayesian best response $\beta_\epsilon[\sigma]$. A rest point of the RBBR learning dynamic is a Bayesian strategy σ° satisfying $\dot{\sigma}^\circ \equiv 0$. Thus any rest point of the dynamic (5.13) is a Bayesian strategy σ° which satisfies $\beta_\epsilon[\sigma^\circ] = \sigma^\circ$, and hence becomes a fixed point of the regularized Bayesian best response map. Such a fixed point is called a **regularized Bayesian equilibrium**. Let $\mathrm{RBE}(\mathscr{G}, \xi, \epsilon)$ be the collection of all regularized Bayesian equilibrium of the game $\mathscr{G}$ with noise parameter ϵ under type measure ξ, that is, $\mathrm{RBE}(\mathscr{G}, \xi; \epsilon) := \{\sigma^\circ \in \Sigma : \beta_\epsilon[\sigma^\circ] = \sigma^\circ\}$. We now state the existence theorem for regularized Bayesian equilibria in Bayesian population games.

Theorem 10 *Suppose that the Bayesian population game $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ is weakly continuous. Then for every $\epsilon > 0$, the RBBR mapping $\sigma \mapsto \beta_\epsilon[\sigma]$ admits a fixed point, that is, there exists $\sigma^\circ_\epsilon \in \Sigma$ such that $\beta_\epsilon[\sigma^\circ_\epsilon] = \sigma^\circ_\epsilon$.*

Proof 20 *See Appendix D.4.*

As a corollary to the above theorem, we have the following result which states the existence of regularized equilibrium in Bayesian aggregative population games.

Corollary 3 *Let $\mathscr{G} : \Sigma \times \mathscr{C}_n(\Delta) \to \mathbb{R}^n$ be a Bayesian aggregative population game as defined in Example 8. Then for every noise parameter $\epsilon > 0$, there exists a regularized Bayesian equilibrium of the game $\mathscr{G}$.*

Proof 21 *See Appendix D.5.*

[8]Thus, we identify the space of integrable signed Bayesian strategies by the Böchner space $\mathscr{L}^1(\Omega, \xi, \mathbb{R}^n)$. See Appendix D.1.

[9]We review the notions of continuity and differentiability of trajectories through the space $(\widehat{\Sigma}, |||\cdot|||)$ from Ely and Sandholm (2005). Let $(\widehat{\sigma}_t)_{t \geqslant 0}$ be a trajectory through $\widehat{\Sigma}$. An integrable signed Bayesian strategy $\widehat{\sigma} \in \widehat{\Sigma}$ is a *strong limit* of $(\widehat{\sigma}_t)_{t \geqslant 0}$ as t approaches t_0 if $\lim_{t \to t_0} |||\widehat{\sigma}_t - \widehat{\sigma}||| = 0$. The trajectory $(\widehat{\sigma}_t)_{t \geqslant 0}$ is *continuous* at time t if $\lim_{s \to t} |||\widehat{\sigma}_s - \widehat{\sigma}_t||| = 0$. We say that the trajectory $(\widehat{\sigma}_t)_{t \geqslant 0}$ is continuous if it is continuous at every $t \geqslant 0$. The trajectory $(\widehat{\sigma}_t)_{t \geqslant 0}$ is *differentiable* if for every $t \geqslant 0$, there exists $\dot{\widehat{\sigma}} \in \widehat{\Sigma}$ such that

$$\lim_{\epsilon \longrightarrow 0} \left|\left|\left| \frac{\widehat{\sigma}_{t+\epsilon} - \widehat{\sigma}_t}{\epsilon} - \dot{\widehat{\sigma}} \right|\right|\right| = 0.$$

We now aim to prove the convergence of regularized Bayesian equilibria to the Bayesian equilibria of the underlying game $\mathscr{G}$ when we allow the noise parameter to be arbitrarily small. In such a scenario, we can then argue that when the noise parameter is close to 0, the regularized Bayesian equilibrium acts as a smoothed approximation of the original Bayesian equilibrium relative to the underlying weak topology. Such approximation results have been considered, for example in Hofbauer and Sandholm (2002) for finite strategy games and in Lahkar and Riedel (2015), Lahkar et al. (2022b) for continuum strategy homogeneous population games. Thus, this calls for an extension of such a result in heterogeneous population games, towards which we dedicate the rest of the section. In what follows, we require some technical conditions on the regularizer. We begin with the following definition.

Definition 26 *For an extended real valued function $f:\mathbb{R}^n \to \mathbb{R}\cup\{-\infty,\infty\}$, the convex conjugate $f^*:\mathbb{R}^n \to \mathbb{R}\cup\{-\infty,\infty\}$ is defined by*

$$f^*(\boldsymbol{y}) := \sup_{\boldsymbol{x}\in\mathbb{R}^n}(\langle \boldsymbol{y},\boldsymbol{x}\rangle - f(\boldsymbol{x})), \quad \textit{for all } \boldsymbol{y}\in\mathbb{R}^n.$$

We also require two technical conditions on the regularizer, namely Condition C1 and Condition C2. We provide some heuristics behind them prior to the definition of both the conditions.

Condition C1 requires that as the noise level in the learning dynamic gradually vanishes, the population, while playing their best response at a particular population state must concentrate themselves over that strategy which has the maximum payoff at that state in the limit. For $\epsilon > 0$, let us denote $\epsilon\mathbf{v}$ by $\mathbf{v}_\epsilon$. Then, Condition C1 is defined as follows.

Definition 27 *A regularizer $\boldsymbol{v}:\Delta \to \mathbb{R}\cup\{\infty\}$ satisfies **Condition C1** if for all $i,j\in S$ with $i\neq j$ and $\boldsymbol{u}\in\mathbb{R}^n$ with $\boldsymbol{u}_j > \boldsymbol{u}_i$, there exists a function $\lambda:(0,\infty)\to\mathbb{R}$ with $\lim_{x\to 0}\lambda(x)=\infty$ such that*

$$\nabla \boldsymbol{v}^*_{\epsilon,j}(\boldsymbol{u}) \geqslant \lambda(\epsilon)\nabla \boldsymbol{v}^*_{\epsilon,i}(\boldsymbol{u}), \quad \textit{for sufficiently small } \epsilon > 0, \tag{5.14}$$

where $\nabla \boldsymbol{v}^*_\epsilon(\boldsymbol{u})$ denotes the gradient of the convex conjugate $\boldsymbol{v}^*_\epsilon$ of $\boldsymbol{v}_\epsilon$ at the point $\boldsymbol{u}\in\mathbb{R}^n$.

The following condition (Condition C2) ensures that a small change in the population states can only lead to a small change in the value of the regularizer. Before defining the condition technically, we require the following notion; that of a shift of a Bayesian strategy.

Definition 28 *Let $\sigma\in\Sigma$ be a Bayesian strategy, and let $\bar{\Omega}\subseteq\Omega$ be such that $\xi(\bar{\Omega})>0$. We say that $\widetilde{\sigma}$ is a **shift** of σ over $\bar{\Omega}$ from $\boldsymbol{i}$ to $\boldsymbol{j}$, where $\boldsymbol{i},\boldsymbol{j}:\bar{\Omega}\to S$ with $\boldsymbol{i}(\omega)\neq\boldsymbol{j}(\omega)$ for all $\omega\in\bar{\Omega}$, if for all $\omega\in\bar{\Omega}$, the following three conditions are satisfied:*

- $\widetilde{\sigma}^{\boldsymbol{i}(\omega)}(\omega) \leqslant \sigma^{\boldsymbol{i}(\omega)}(\omega)$,
- $\widetilde{\sigma}^{\boldsymbol{j}(\omega)}(\omega) = \sigma^{\boldsymbol{j}(\omega)}(\omega) + (\sigma^{\boldsymbol{i}(\omega)}(\omega) - \widetilde{\sigma}^{\boldsymbol{i}(\omega)}(\omega))$, *and*
- $\widetilde{\sigma}^k(\omega) = \sigma^k(\omega)$, *for all $k\neq\boldsymbol{i}(\omega),\boldsymbol{j}(\omega)$.*

We now proceed to state the Condition C2 which calls for a 'mild' continuity condition on the regularizer in the sense that small changes in the population state do not lead to abrupt changes in the regularizer. We clarify for the reader that the required notion of continuity imposed on the regularizer is much weaker than the usual notion of continuity for extended real valued functions.

Definition 29 *Let $(\sigma_m)_{m\geqslant 1}$ be a sequence of Bayesian strategies and let $\bar{\Omega}\subseteq\Omega$ be such that $\xi(\bar{\Omega})>0$. We say that a regularizer $\boldsymbol{v}$ satisfies **Condition C2** with respect to the pair $((\sigma_m)_{m\geqslant 1},\bar{\Omega})$ if for all $\omega\in\bar{\Omega}$, all distinct $\boldsymbol{i},\boldsymbol{j}:\bar{\Omega}\to S$, with*

- $\lim_{m\to\infty}\sigma_m^{\boldsymbol{i}(\omega)}(\omega)>0$,
- $\mathscr{G}^{\boldsymbol{j}(\omega)}(\sigma_m,\omega)>\mathscr{G}^{\boldsymbol{i}(\omega)}(\sigma_m,\omega)$ *for sufficiently large m (depending on ω),*

there exists $(\widetilde{\sigma}_m)_{m\geqslant 1}$, where $\widetilde{\sigma}_m$ is a shift of σ_m over $\bar{\Omega}$ from $\boldsymbol{i}$ to $\boldsymbol{j}$, and a constant $\kappa_i(\omega)>0$ (independent of m) such that

$$|\boldsymbol{v}(\sigma_m(\omega))-\boldsymbol{v}(\widetilde{\sigma}_m(\omega))|\leqslant\kappa_{\boldsymbol{i}(\omega)}(\omega)(\sigma_m^{\boldsymbol{i}(\omega)}(\omega)-\widetilde{\sigma}_m^{\boldsymbol{i}(\omega)}(\omega)), \tag{5.15}$$

for sufficiently large m (depending on ω).

*We say that a regularizer $\boldsymbol{v}$ satisfies **Condition C2** with respect to a sequence of Bayesian strategies $(\sigma_m)_{m\geqslant 1}$ if it satisfies Condition C2 with respect to the pair $((\sigma_m)_{m\geqslant 1},\bar{\Omega})$ for every $\bar{\Omega}\subseteq\Omega$ with $\xi(\bar{\Omega})>0$.*

We dedicate Appendix D.7 to show that the Shannon-Gibbs entropy satisfies Condition C1 and that the Tsallis and the Burg entropy satisfy Condition C2. We now state the following approximation theorem.

Theorem 11 *Let $(\epsilon_m)_{m\geqslant 1}$ be a sequence of positive noise parameters converging to 0 and let $(\sigma_m^{\circ})_{m\geqslant 1}$ be the corresponding sequence of ϵ_m-regularized Bayesian equilibria such that $\sigma_m^{\circ}\in RBE(\mathscr{G},\xi,\epsilon_m)$ for all $m\geqslant 1$. Suppose that the regularizer $\boldsymbol{v}$ satisfies either Condition C1, or Condition C2 with respect to the sequence of Bayesian strategies $(\sigma_m^{\circ})_{m\geqslant 1}$. Then $(\sigma_m^{\circ})_{m\geqslant 1}$ has accumulation points with respect to the weak topology and any such accumulation point is a Bayesian equilibrium of the original game $\mathscr{G}$.*

Proof 22 *See Appendix D.6.*

5.4 Fundamental properties of RBBR learning dynamic

In this section, we establish two fundamental properties of the regularized Bayesian best response dynamic, namely, the existence of solutions from arbitrary initial conditions in Σ and continuity of these solutions with respect to the initial states. We resort to the infinite dimensional Picard-Lindelöff theorem (Zeidler (1986), Pg. 79) and the Grönwall's inequality (see Zeidler (1986), Pg. 82, Proposition 3.10 and 3.11) in order to prove the main results of this section. We begin with the following two definitions.

Definition 30 *A Bayesian population game $\mathscr{G}:\Sigma\times\Omega\to\mathbb{R}^n$ is said to be **strongly Lipschitz** if there exists $\kappa>0$ such that for ξ-a.e $\omega\in\Omega$,*

$$\|\mathscr{G}(\sigma,\omega)-\mathscr{G}(\rho,\omega)\|\leqslant\kappa\left|\!\left|\!\left|\sigma-\rho\right|\!\right|\!\right|, \quad \textit{for all } \sigma,\rho\in\Sigma.$$

The strong Lipschitz assumption on Bayesian population games is a standard assumption in the literature of Bayesian evolutionary games (Ely and Sandholm (2005)). In particular, the random matching versions of Bayesian aggregate population games, asymmetric war of attrition games, additively separable anonymous games, incomplete information games satisfy Lipschitz continuity according to the above definition.

Definition 31 *Consider an abstract differential equation*

$$\phi'(t) = \Lambda(\phi(t))$$

on a Banach space $(\mathbb{X}, \|\cdot\|_{\mathbb{X}})$. *Suppose that a unique solution to the above differential equation exists for every initial condition* $\phi(0) = \gamma$ *and denote the solution with initial condition* γ *by* $\phi_\gamma(t)$. *Then the semiflow of the dynamic is the map* $\zeta : \mathbb{X} \times [0, \infty) \to \mathbb{X}$ *defined as* $\zeta(\gamma, t) = \phi_\gamma(t)$ *for all* $\gamma \in \mathbb{X}$ *and* $t \in [0, \infty)$.

We now state the main theorem of this section.

Theorem 12 *Consider the type space* (Ω, ξ). *Suppose that the Bayesian population game* $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ *is strongly Lipschitz. Then the following statements hold true:*

1. *For every initial condition* $\sigma_0 \in \Sigma$, *the RBBR dynamic admits a unique solution* $(\sigma_t)_{t \geqslant 0}$.
2. *The semiflow* $\zeta : \widehat{\Sigma} \times [0, \infty) \to \widehat{\Sigma}$ *of the RBBR dynamic is continuous in the initial conditions with respect to the strong topology on* $\widehat{\Sigma}$.

Proof 23 *See Appendix E.5.*

Following is a corollary, which states that the above result holds true for Bayesian aggregative population games in Example 8. In what follows we require the following definition.

Definition 32 *A population game* $\omega : \Delta \to \mathbb{R}^n$ *is said to be Lipschitz if there exists a real number* $\alpha > 0$ *such that*

$$\|\omega(\mathbf{x}) - \omega(\mathbf{y})\| \leqslant \alpha \|\mathbf{x} - \mathbf{y}\|, \quad \text{for all } \mathbf{x}, \mathbf{y} \in \Delta.$$

Let $\text{Lip}_\alpha(\Delta)$ be the collection of α-Lipschitz population games which map Δ into $\mathbb{R}^n$. Observe that under the uniform topology, $\text{Lip}_\alpha(\Delta)$ is a closed subset of $\mathscr{C}_n(\Delta)$. Thus $\text{Lip}_\alpha(\Delta)$ is a possible candidate for the support of the type measure ξ.

Corollary 4 *Let* $\mathscr{G} : \Sigma \times \mathscr{C}_n(\Delta) \to \mathbb{R}^n$ *be a Bayesian aggregate population game. Suppose that there exists* $\alpha > 0$ *such that* $\xi(Lip_\alpha(\Delta)) = 1$. *Then the conclusions of Theorem 17 hold true.*

Proof 24 *The proof of the corollary follows immediately and is left to the reader.*

We now focus on some applications of the RBBR dynamic in order to obtain convergence results for certain classes of Bayesian population games. As in case of homogeneous population games with finite or continuum strategies, we consider here the class of Bayesian potential and Bayesian negative semi-definite games. To this end, we need to restrict our attention to norm-compact subsets of the space of Bayesian strategies which are forward invariant under the RBBR dynamic. We thus provide a sufficient condition to identify such subsets on which the solution trajectories of the RBBR dynamic can lead to nice convergence results. This calls for some regularity on the Bayesian strategies which we present in the following definition.

Definition 33 *A Bayesian strategy $\sigma:\Omega\to\Delta$ is said to be Lipschitz if there exists $\alpha>0$ such that*

$$\|\sigma(\omega)-\sigma(\bar{\omega})\|\leqslant \alpha d_{\Omega}(\omega,\bar{\omega}), \quad \textit{for all } \omega,\bar{\omega}\in\Omega.$$

Intuitively, the Lipschitz condition ensures that if the types of any two agents are close, then the induced beliefs of the two agents are also close. Formally, if two types ω and $\bar{\omega}$ are close in the d_Ω metric, then the induced mixed strategies $\sigma(\omega)$ and $\sigma(\bar{\omega})$ are also close with respect to the Euclidean metric.

We need a regularity condition (Lipschitz continuity in this case) on the Bayesian strategies to establish the convergence of the dynamic to the perturbed equilibrium of the underlying game. Such a convergence requires certain subsets of Σ to be norm compact and forward invariant under the regularized Bayesian best response dynamic, which is ensured by the Lipschitz continuity (together with other conditions).

Next we define a subset of the space of Bayesian strategies as follows. Let $M\subseteq\Delta^{\circ}$ be a compact set. Define

$$\Sigma_M := \{\sigma\in\Sigma : \sigma(\omega)\in M \text{ for } \xi-\text{a.e } \omega\in\Omega\}. \tag{5.16}$$

In words, Σ_M is the collection of those Bayesian strategies such that for ξ-almost all types $\omega\in\Omega$, the induced belief over the strategies lies in M. In principle this means that the beliefs induced by Bayesian strategies in Σ are almost surely bounded away from the boundary $\partial\Delta$ of the simplex Δ. As a result, the induced belief assigns a positive probability to each strategy in S. To present the main result on convergence in Bayesian potential games, we restrict attention to the following subset of Bayesian strategies. Fix $\alpha>0$, $M\subseteq\Delta^{\circ}$ compact, and $\Xi\subseteq\Omega$. Define

$$\Sigma_M^{\alpha}(\Xi) := \{\sigma\in\Sigma : \sigma \text{ is } \alpha\text{-Lipschitz on } \Xi\}\cap\Sigma_M. \tag{5.17}$$

Thus, $\Sigma_M^{\alpha}(\Xi)$ is the collection of Bayesian strategies which are α-Lipschitz on the subset Ξ with range ξ-a.s contained in the compact set $M\subseteq\Delta^{\circ}$. In the following lemma (Lemma 5), we show that $\Sigma_M^{\alpha}(\Xi)$ is relatively norm compact subject to suitable conditions on the type measure ξ. We also show that $\Sigma_M^{\alpha}(\Xi)$ is forward invariant under the RBBR dynamic under some further conditions on the game $\mathscr{G}$ and on the Lipschitz coefficient in Definition 33. As a result, the lemma provides a candidate space to prove convergence results for Bayesian potential games and Bayesian negative semidefinite games.

Lemma 5 *Suppose that there exists a compact set $\Xi\subseteq\Omega$ such that $\xi(\Xi)=1$. Then the following statements hold true:*

1. *For every $\alpha>0$ and every $M\subseteq\Delta^{\circ}$ compact, the subset $\Sigma_M^{\alpha}(\Xi)$ (as defined in (5.17)) is relatively norm compact in $\widehat{\Sigma}$.*

2. *Fix $\epsilon,\gamma>0$. Let $\boldsymbol{v}$ be steep and γ-strongly convex. Suppose that for every $\sigma\in\Sigma$, the following assumptions are satisfied:*

   - $\|\mathscr{G}(\sigma,\omega)\|\leqslant c$, *for ξ-a.e $\omega\in\Xi$, for some $c>0$.*
   - *the mapping $\omega\mapsto\mathscr{G}(\sigma,\omega)$ is λ-Lipschitz on Ξ, for some $\lambda>0$.*[10]

   Then there exists a compact set $M\subseteq\Delta^{\circ}$ (depending on c) such that $\Sigma_M^{\alpha}(\Xi)$ is forward invariant under the RBBR learning dynamic, provided $\alpha\geqslant\frac{\lambda}{\epsilon\gamma}$.

Proof 25 *See Appendix D.9.*

5.5 Convergence in Bayesian potential games

In this section we define the notion of Bayesian potential games and study the convergence of the RBBR dynamic to the regularized Bayesian equilibria in such games.[11] For this purpose, we first require the notion of Fréchet derivative, which is a generalization of the standard notion of derivatives to Banach spaces. Let $\mathbb{V}$ and $\mathbb{W}$ be two Banach spaces. An operator $\phi:\mathbb{V}\to\mathbb{W}$ is Fréchet differentiable at $x\in\mathbb{V}$ if there exists a bounded linear operator $A_x:\mathbb{V}\to\mathbb{W}$ such that

$$\lim_{\|h\|_{\mathbb{V}}\to 0}\frac{\|\phi(x+h)-\phi(x)-A_x(h)\|_{\mathbb{W}}}{\|h\|_{\mathbb{V}}}=0.$$

Thus, if A is the Fréchet derivative of ϕ, then $\phi(x+h)=\phi(x)+A_x(h)+\mathbf{o}(\|h\|_{\mathbb{V}})$ for all h near zero in $\mathbb{V}$.

In this paper, we consider Fréchet differentiability of functions $\phi:\widehat{\Sigma}\to\mathbb{R}$. The Fréchet derivative of ϕ is denoted by $\mathrm{D}\phi$. Let $\mathbb{R}_0^n$ be the collection of n-dimensional vectors such that the sum of the components is zero, that is, $\mathbb{R}_0^n:=\{x=(x_1,\dots,x_n)\in\mathbb{R}^n: x_1+\cdots+x_n=0\}$. Let $\Sigma_0:=\{\widehat{\sigma}\in\widehat{\Sigma}:\widehat{\sigma}(\omega)\in\mathbb{R}_0^n$ for $\xi-$a.e $\omega\in\Omega\}$. For a Fréchet differentiable function $\phi:\widehat{\Sigma}\to\mathbb{R}$, the *Fréchet gradient* $\nabla^{\mathbf{F}}\phi:\widehat{\Sigma}\to\mathscr{L}^{\infty}(\Omega,\xi,\mathbb{R}^{n*})$ of ϕ is a bounded measurable map defined as follows: for every $\widehat{\sigma}\in\widehat{\Sigma}$ and every $\sigma_0\in\Sigma_0$,

$$\mathrm{D}\phi[\widehat{\sigma}](\sigma_0):=\int_{\Omega}\langle\sigma_0,\nabla^{\mathbf{F}}\phi(\widehat{\sigma})\rangle(\omega)\xi(d\omega)=\int_{\Omega}\langle\nabla^{\mathbf{F}}\phi[\widehat{\sigma}](\omega),\sigma_0(\omega)\rangle\xi(d\omega),\tag{5.18}$$

where the inner product in above integral after the first equality denotes the bilinear pairing as in Definition 52 and the inner product in the integral after the second equality is the usual inner product in $\mathbb{R}^n$.

[10] We bring to notice to the readers that this condition is not very restrictive in the sense that the Bayesian aggregative population games (Example 8) satisfies this condition with $\lambda=1$.

[11] See Monderer and Shapley (1996), Sandholm (2010), and Lahkar and Riedel (2015) for a complete exposition of homogeneous potential games.

Definition 34 (Zusai (2023)) *A mapping $\mathscr{G}:\Sigma\times\Omega\to\mathbb{R}^n$ is called a **Bayesian potential game** if there exists a Fréchet differentiable map $\varphi:\widehat{\Sigma}\to\mathbb{R}$ such that for every $\sigma\in\Sigma$, we have*

$$\nabla^{\boldsymbol{F}}\varphi[\sigma](\omega)=\mathscr{G}(\sigma,\omega),\quad \xi-a.e\,\omega\in\Omega.$$

The map φ is called the Bayesian potential function of the Bayesian potential game $\mathscr{G}$.

We then introduce the notion of entropy adjusted Bayesian potential function. For a compact set $M\subseteq\Delta^\circ$, we then define a map $\widetilde{\mathbf{v}}:\Sigma_M\to\mathbb{R}$ such that

$$\widetilde{\mathbf{v}}(\sigma):=\int_\Omega \mathbf{v}(\sigma(\omega))\xi(d\omega),\quad \text{for all } \sigma\in\Sigma_M.$$

The fact that the regularizer $\mathbf{v}$ is strongly convex on Δ° implies that $\mathbf{v}$ is continuous on M. Thus $\mathbf{v}$ is bounded on M and hence $\widetilde{\mathbf{v}}$ is well defined. This leads us to the following important definition.

Definition 35 *For $M\subseteq\Delta^\circ$ compact, the **entropy adjusted Bayesian potential function** is a mapping $\widetilde{\varphi}:\Sigma_M\to\mathbb{R}$ defined as*

$$\widetilde{\varphi}(\sigma)=\varphi(\sigma)-\widetilde{\boldsymbol{v}}(\sigma),\quad \textit{for all } \sigma\in\Sigma_M.$$

We require an additional result for establishing convergence in Bayesian potential games. We shall say that the RBBR dynamic satisfies **Bayesian positive correlation** with respect to the game $\mathscr{G}:\Sigma\times\Omega\to\mathbb{R}^n$ if $\int_\Omega\langle\mathscr{G}(\sigma,\omega),\dot{\sigma}(\omega)\rangle\xi(d\omega)\geqslant 0$ for all $\sigma\in\Sigma$ with equality only if $\dot{\sigma}=0$. For the rest of our analysis on Bayesian potential and Bayesian negative semidefinite games, we assume that the regularizer is a $\mathscr{C}^1$ function. The reason for such an assumption will be clear from the following definition. Define the virtual payoff vector at a Bayesian strategy $\sigma\in\Sigma$ for a player of type $\omega\in\Omega$ as

$$\widetilde{\mathscr{G}}(\sigma,\omega)=\mathscr{G}(\sigma,\omega)-\nabla^{\mathbf{G}}\widetilde{\mathbf{v}}[\sigma](\omega).\text{[12]}\tag{5.19}$$

The following lemma states that the RBBR dynamic satisfies Bayesian positive correlation with respect to the virtual payoff $\widetilde{\mathscr{G}}$. The lemma imposes a further condition (5.20) on the regularizer. Let $\mathscr{C}^1((0,1))$ denote the collection of all continuously differentiable functions $\theta:(0,1)\to\mathbb{R}$.

Definition 36 *A regularizer $\boldsymbol{v}:\Delta\to\mathbb{R}^n\cup\{\infty\}$ is said to be **decomposable** if there exists a $\mathscr{C}^1$ function θ such that the regularizer $\boldsymbol{v}$ can be expressed as*

$$\mathbf{v}(\mathbf{x}):=\sum_{i=1}^n\theta(\mathbf{x}_i),\quad \text{for all } \mathbf{x}\in\Delta^\circ.\text{[13]}\tag{5.20}$$

The three entropies we consider, namely the Shannon-Gibbs entropy, the Tsallis entropy, and the Burg entropy, all satisfy condition (5.20). Following is an intermediate lemma which we use to prove convergence in Bayesian potential games.

[12]The fact that $\mathbf{v}$ is smooth in the interior Δ° and bounded on M implies that $\widetilde{\mathbf{v}}$ is Gâteaux differentiable.
[13]Such regularizers have been consider earlier, for instance in Coucheney et al. (2015), Mertikopoulos and Sandholm (2016).

Lemma 6 *Let $\mathscr{G}:\Sigma\times\Omega\to\mathbb{R}^n$ be a strongly Lipschitz Bayesian potential game. Assume that the conditions of part (ii) of Lemma 5 are satisfied. Suppose that the regularizer satisfies (5.20). Then the following statements hold:*

1. *Suppose that θ' is strictly increasing on $(0,1)$. Then for every $\epsilon>0$, the RBBR dynamic satisfies Bayesian positive correlation with respect to the virtual payoff $\widetilde{\mathscr{G}}$ as defined in (5.19).*

2. *Consider the norm compact set $\Sigma_M^\alpha(\Xi)$ obtained in Lemma 5. Let $\delta=(\delta_1,\dots,\delta_n)\in(0,\frac{1}{2})^n$ be such that $M\subseteq[\delta_1,1-\delta_1]\times\cdots\times[\delta_n,1-\delta_n]$. Suppose that θ' is strictly increasing on the interval $[\min_{i=1,\dots,n}\delta_i, 1-\min_{i=1,\dots,n}\delta_i]$. Then the RBBR dynamic satisfies positive correlation with respect to the virtual payoff $\widetilde{\mathscr{G}}$ defined in (5.19).*

Proof 26 *See Appendix D.10.*

We shall see in Theorem 18 below that $\widetilde{\varphi}$ restricted to $\Sigma_M^\alpha(\Xi)$ will act as a Lyapunov function for the RBBR dynamic when the underlying game is a Bayesian potential game. Also, the fact that $\Sigma_M^\alpha(\Xi)$ is relatively norm compact and forward invariant under the RBBR learning dynamic will allow us to use results from the theory of dynamical systems to prove the desired convergence of the solutions to the regularized Bayesian equilibrium. In what follows, we present the main theorem of this section.

Theorem 13 *Suppose that the assumptions of Theorem 5, part (ii) of Lemma 5 and Lemma 6 are satisfied. Let $\mathscr{G}:\Sigma\times\Omega\to\mathbb{R}^n$ be a Bayesian potential game with Bayesian potential function $\varphi:\Sigma\to\mathbb{R}$ and let $\widetilde{\varphi}_M^\alpha:\Sigma_M^\alpha(\Xi)\to\mathbb{R}$ be the entropy adjusted Bayesian potential function restricted to $\Sigma_M^\alpha(\Xi)$. Then the following statements hold true:*

1. *$\widetilde{\varphi}_M^\alpha$ increases weakly along every solution trajectory to RBBR learning dynamic that originates in $\Sigma_M^\alpha(\Xi)$ and increases strictly across every non-stationary solution trajectory.*

2. *The set of omega–limit points (in the strong topology) of any trajectory to the RBBR learning dynamic is a non-empty connected compact set of perturbed Bayesian equilibria. Moreover, such limit points are local maximizers of the entropy adjusted Bayesian potential function $\widetilde{\varphi}_M^\alpha$ on $\Sigma_M^\alpha(\Xi)$.*

Proof 27 *See Appendix D.11.*

5.6 Convergence in Bayesian negative semidefinite games

In this section we seek to define the notion of Bayesian negative semidefinite games, which is a generalization of negative semidefinite population games on finite strategy spaces to heterogeneous populations. We also generalize the notion of self defeating externalities to Bayesian self defeating externalities. Following Cheung (2014), we show the equivalence between Bayesian negative semidefinite games and Bayesian self defeating externalities. Similar to Hofbauer and Sandholm (2007) in finite strategy negative semidefinite games and Lahkar and Riedel (2015) in continuum strategy negative semidefinite games under homogeneous set-up, we show uniqueness of regularized equilibrium. Finally, we conclude the section by proving a convergence result which shows that the solution to the RBBR learning dynamic with initial condition in norm compact forward invariant subsets of Σ converges to the unique regularized Bayesian equilibrium under the strong topology on Σ, when the underlying game is a Bayesian negative semidefinite game. We begin with the following definition.

Definition 37 *A mapping $\mathscr{G}:\Sigma\times\Omega\to\mathbb{R}^n$ is called a* ***Bayesian negative semidefinite*** *game if*

$$\int_{\Omega}\langle\mathscr{G}(\sigma,\omega)-\mathscr{G}(\rho,\omega),\sigma(\omega)-\rho(\omega)\rangle\xi(d\omega)\leqslant 0\quad \textit{for all } \sigma,\rho\in\Sigma. \tag{5.21}$$

Hofbauer and Sandholm (2009) calls such games stable games. The above definition is a generalization of stable games to heterogeneous populations. Below we provide some examples of games which are Bayesian negative semi-definite.

Example 15 (Bayesian aggregative RPS game) *In this example, we define the notion of a Bayesian aggregative Rock-Paper-Scissor game. We will use the abbreviation BA-RPS for the same. In standard RPS games, the players are allowed three actions, namely, 'rock (R)', 'paper (P)' and 'scissors (S)'. The P covers the R, the R smashes the S, and the S cuts to the P. Now consider an agent of type ω. Suppose that the Bayesian strategy exercised by the underlying population is Σ. Then the payoff vector of the agent is given by $A_\omega\mathscr{E}(\sigma)$, where A_ω is the payoff matrix specific to agents of type ω and is defined as*

$$A_\omega=\begin{pmatrix}0 & -a(\omega) & b(\omega)\\ b(\omega) & 0 & -a(\omega)\\ -a(\omega) & b(\omega) & 0\end{pmatrix} \tag{5.22}$$

where $a,b:\Omega\to[0,\infty)$ are measurable functions. Formally, a Bayesian aggregate RPS is a mapping $\mathscr{G}_{BA\text{-}RPS}:\Sigma\times\Omega\to\mathbb{R}^n$ such that for all $\sigma\in\Sigma$,

$$\mathscr{G}_{BA-RPS}(\sigma,\omega):=A_\omega\mathscr{E}(\sigma),\quad \textit{for all } \omega\in\Omega. \tag{5.23}$$

A special case of the above game is called a standard BA-RPS game if $a(\omega)=b(\omega)$ for ξ-a.e $\omega\in\Omega$. Now suppose that $b(\omega)-a(\omega)\geqslant 0$ for ξ-a.e $\omega\in\Omega$, then from Definition 37 it follows that $\mathscr{G}_{BA\text{-}RPS}$ is a Bayesian aggregative RPS game.[14]

Example 16 (Bayesian symmetric zero-sum game) *For every $\omega\in\Omega$, let A_ω be a skew symmetric matrix, that is, $A_\omega^{ij}=-A_\omega^{ji}$ for all $i,j\in S$. We then define the Bayesian symmetric zero-sum game as a mapping $\mathscr{G}_{sym}:\Sigma\times\Omega\to\mathbb{R}$ such that for every $\sigma\in\Sigma$,*

$$\mathscr{G}_{sym}(\sigma,\omega):=A_\omega\mathscr{E}(\sigma),\quad \textit{for all } \omega\in\Omega. \tag{5.24}$$

Thus, in such games, the inequality (5.21) is satisfied and hence $\mathscr{G}_{sym}$ is a Bayesian negative semidefinite game.

[14] This follows computations similar to Hofbauer and Sandholm (2009).

We now turn to investigate the equivalence between Bayesian negative semidefiniteness and Bayesian self-defeating externalities. Recall from Section 5.3 that for a fixed $\omega \in \Omega$, the slice game $\mathscr{G}_{\omega} : \Sigma \to \mathbb{R}^n$ is defined as $\mathscr{G}_{\omega}(\sigma) := \mathscr{G}(\sigma, \omega)$, for all $\sigma \in \Sigma$. We require this in the definition of Bayesian self defeating externalities.

Definition 38 *A Bayesian population game $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ satisfies **Bayesian self defeating externalities** if the following two conditions are satisfied:*

[(i)]

1. For every $\omega \in \Omega$, the slice game $\mathscr{G}_{\omega} : \widehat{\Sigma} \to \mathbb{R}^n$ is $\mathscr{C}^1$-Fréchet differentiable for ξ-a.e $\omega \in \Omega$, and
2. $\int_{\Omega} \langle \mathrm{D}\mathscr{G}_{\omega}[\sigma](\sigma_0), \sigma_0(\omega) \rangle \xi(d\omega) \leqslant 0$, for all $\sigma \in \Sigma$, and all $\sigma_0 \in \Sigma_0$.

Hofbauer and Sandholm (2009) first proved that the notion of negative semidefiniteness and self-defeating externalities are equivalent in the case of finite strategy homogeneous population games, followed by Cheung (2014) who proved the equivalence in the context of continuum strategy homogeneous population games. Our next lemma states that indeed, the same holds true in the case of finite strategy Bayesian population games also.

Proposition 14 *A Bayesian population game $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ is Bayesian negative semidefinite if and only if it satisfies Bayesian self defeating externalities.*

Proof 28 *See Appendix D.12.*

Following is the result on the uniqueness of regularized Bayesian equilibrium for Bayesian negative semidefinite games. This result extends the uniqueness result of Hofbauer and Sandholm (2009) in the heterogeneous set-up for finitely many strategies.

Proposition 15 *Fix a noise parameter $\epsilon > 0$. Suppose that $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ is a Bayesian negative semidefinite game which is $\mathscr{C}^1$ in the sense of Fréchet differentiability. Suppose that the regularizer $\boldsymbol{v}$ satisfies $\langle d\nabla \boldsymbol{v}_{\epsilon}[\boldsymbol{x} + t\boldsymbol{z}_0](\boldsymbol{z}_0), \boldsymbol{z}_0 \rangle \geqslant 0$, for all $t \in \mathbb{R}$, $\boldsymbol{x} \in \Delta$ and $\boldsymbol{z}_0 \in \mathbb{R}^n_0$ such that $\boldsymbol{x} + t\boldsymbol{z}_0 \in \Delta$. Then the game $\mathscr{G}$ has a unique regularized Bayesian equilibrium.*

Proof 29 *See Appendix D.13.*

Now, we proceed to study the convergence of Bayesian negative semidefinite games under the RBBR dynamic. To do so we again restrict the dynamic to the norm compact forward invariant subset $\Sigma^{\alpha}_M(\Xi)$ as defined in Section 5.5. On this restricted space we define the appropriate Lyapunov function as a mapping $\psi^{\alpha}_{M,\epsilon} : \Sigma^{\alpha}_M(\Xi) \to \mathbb{R}$ such that for all $\sigma \in \Sigma^{\alpha}_M(\Xi)$,

$$\psi^{\alpha}_{M,\epsilon}(\sigma) := \Big[\int_{\Omega} \langle \mathscr{G}(\sigma,\omega), \beta_{\epsilon}[\sigma](\omega) \rangle \xi(d\omega) - \widetilde{\mathbf{v}}_{\epsilon}(\beta_{\epsilon}(\sigma)) \Big] - \Big[\int_{\Omega} \langle \mathscr{G}(\sigma,\omega), \sigma(\omega) \rangle \xi(d\omega) - \widetilde{\mathbf{v}}_{\epsilon}(\sigma) \Big]. \tag{5.25}$$

Theorem 14 *Let $\mathscr{G} : \Sigma \times \Omega \to \mathbb{R}^n$ be a Bayesian negative semidefinite game such that the slice map $\mathscr{G}_\omega$ is $\mathscr{C}^1$-Fréchet differentiable in the norm topology for ξ-a.e $\omega \in \Omega$. Then, the Lyapunov function $\psi^\alpha_{M,\epsilon} : \Sigma^\alpha_M(\Xi) \to \mathbb{R}$ as defined in (E.19) decreases monotonically to zero along every solution trajectory to the RBBR learning dynamic that originates in $\Sigma^\alpha_M(\Xi)$, and decreases strictly along every non stationary solution. Hence, every solution trajectory to the RBBR learning dynamic in $\Sigma^\alpha_M(\Xi)$ converges with respect to the norm topology on $\Sigma^\alpha_M(\Xi)$ to the unique regularized Bayesian equilibrium.*

Proof 30 *See Appendix D.14.*

CHAPTER

6

PERTURBED BAYESIAN BEST RESPONSE DYNAMICS IN CONTINUUM GAMES

6.1 Introduction

6.1.1 Background of the Problem

Evolutionary game theory, from the perspective of an economist, is concerned with the behavior of large populations consisting of strategically interacting 'small' players who recurrently revise their strategies based on the availability of payoff opportunities. From the point of view of a biologist, the theory can also be explained as a tool to model the interactions between various animal populations whose genetic features lead to uncertain levels of success in their reproductivity. Most standard models in the literature of evolutionary games are built upon two objects: population games and revision protocols. A population game provides a complete description of the payoff structure for every distribution of the players over the common strategy space of the population. On the other hand, a revision protocol (perturbed best response in our case) captures how a player switches to a profitable strategy upon receiving an opportunity.

The literature on evolutionary game theory is majorly concerned with homogeneous populations where players have the same payoff structure. However, there can arise real-life situations where the homogeneity assumption does not hold. For instance, different firms in a Cournot competition may have different payoff functions based on their infrastructure, manpower, and availability/contracts over resources. Even in a prisoner's dilemma game, the same number of years of imprisonment may incur different costs to different players based on their socioeconomic backgrounds, age, health, etc.

The objective of this paper is to investigate Bayesian evolutionary games with a continuum of strategies when players employ perturbed Bayesian best response as the revision protocol. In a Bayesian population game, each player has a 'type', and a population state is no longer a distribution of the players in the population over strategies, rather, it is a function from the space of types to the space of distributions across strategies.[1] As we shall observe in the various examples presented in the paper, the type of an agent is typically characterized by the utility function of the player. The relevance of continuum strategy evolutionary games, although it makes the analysis technically involved, is omnipresent in various domains of economics (Lahkar (2017)), biology (Hammerstein and Parker (1982)), and computer science (Hsieh et al. (2021)). We also provide some examples of Bayesian population games with a continuum of strategies, namely asymmetric war of attrition games, additively separable anonymous games, games with incomplete information, Bayesian aggregative population games in Section 6.2. As in the case of homogeneous population games, the perturbed Bayesian best response for Bayesian games is obtained when the payoff corresponding to players of each type is subject to perturbation by the Shannon entropy and players best respond to these perturbed payoffs. As a result, the logit Bayesian best response dynamic is generated.[2] We consider the perturbed best response approach to avoid multiplicity of best responses as indicated by Gilboa and Matsui (1991) and Hofbauer (1995).

6.1.2 Our Contribution

We now provide a detailed description of the results in the paper. We assume that there is a complete separable metric space of types and the population is distributed over the type space according to some given probability measure. The state of the population is described by a Bayesian strategy which is a map from the space of types to the space of distribution over the strategies. Thus, the state space of the population is the collection of all Bayesian strategies. We consider three different topologies, namely; the usual 'weak topology', a variation of the Bounded Lipschitz topology, and a variation of the total variation topology on the space of Bayesian strategies to prove the results in this paper.[3] In Theorem 15, we prove the existence of logit Bayesian equilibrium in the weak topology under some mild regularity conditions on the Bayesian population game. Theorem 16 establishes that the logit Bayesian equilibrium approximates the original Nash equilibrium of the game for vanishingly small perturbations under a modification of the Bounded Lipschitz topology. Under a version of the total variation norm on the space of Bayesian strategies, we prove the existence of solutions of the perturbed Bayesian best response dynamic (Theorem 17) and provide sufficient conditions for subsets of Bayesian strategies to be norm-compact and forward invariant under the perturbed Bayesian best response dynamic (Lemma 7). Finally, Theorem 18 and Theorem 19, respectively, establish results related to

[1]A population state is referred to as a 'Bayesian strategy' of the population in Ely and Sandholm (2005).

[2]The logit best response dynamic, first introduced by Fudenberg and Levine (1998) in the context of homogeneous population finite strategy games is one of the most well–studied and prototypical form of perturbed best response dynamics in the literature of evolutionary game theory.

[3]As mentioned earlier, the extension to perturbed best response dynamics to continuum strategy heterogeneous/Bayesian population games is technically demanding and requires us to use techniques from probability theory, functional analysis, differential equations, and dynamical systems. In particular, since the state space (Bayesian strategies in this case) is infinite-dimensional, the choice of the topology plays a very crucial role in almost all the results in the paper.

convergence in Bayesian potential games and Bayesian negative semidefinite games.

The paper is organized as follows. In Section 6.2, we provide the preliminaries and model of continuum strategy Bayesian population games followed by some examples of such games. Section 6.3 establishes the existence of perturbed Bayesian equilibrium of the Bayesian population game and exhibits some fundamental properties of the dynamic. In Section 6.4 and Section 6.5, we establish the convergence of solutions of the perturbed best response dynamic to the logit Bayesian equilibria in Bayesian potential games and Bayesian negative semidefinite games. Finally, Section 6.6 concludes the paper.

6.1.3 Literature Review

One can broadly classify the literature of evolutionary games into four categories based on the type of the population (homogeneous or heterogeneous) and the type of the strategy space (finite or continuum).

For homogeneous populations having finite strategy spaces, Benaım and Hirsch (1999) studies fictitious play in infinitely repeated randomly perturbed games where they establish that the sample path process of the repeated game converges almost surely for 2×2 games with countably many Nash equilibria. Hofbauer and Sandholm (2002) establish global convergence results for stochastic fictitious play in four classes of games: games with an interior ESS, zero sum games, potential games, and supermodular games. Hofbauer and Hopkins (2005) investigates stability of mixed strategy equilibria in two-person (bimatrix) games under the perturbed best response dynamics. Later, Hofbauer and Sandholm (2007) considers a stochastic evolutionary game dynamic obtained through random disturbances in payoffs. They show that the resulting evolutionary process converges to an approximate Nash equilibrium in both the medium run and the long run for three general classes of population games: stable games, potential games, and supermodular games. Mertikopoulos and Sandholm (2016) considers a class of reinforcement learning dynamics in finite games obtained via regularization. They study elimination of dominated strategies, asymptotic stability of strict Nash equilibria, and convergence of time-averaged trajectories in zero-sum games with interior Nash equilibrium under such dynamics.[4]

The extension of homogeneous population games for a continuum of strategies was led by Oechssler and Riedel (2001). They study the stability and instability of rest points of the replicator dynamics on continuum strategy spaces with applications in Nash demand game, war of attrition, linear-quadratic game, and the harvest preemtion game. In a companion paper, Oechssler and Riedel (2002) provides robust static stability concepts of the Nash equilibrium in doubly-symmetric games which are sufficient to guarantee evolutionary stability under the replicator dynamics. Hofbauer and Sorin (2006) studies best response dynamics in continuous time for concave-convex zero-sum games with finitely many strategies and provides a dynamical proof of the minmax theorem by establishing the convergence of the evolutionary trajectories to the set of saddle points. Hofbauer et al. (2009) studies Nash stationarity for continuous payoff functions. In addition, they also show asymptotic stability of evolutionarily robust Nash equilibria along with stability analysis for doubly symmetric games. Cheung (2014) considers global convergence and local stability under general deterministic evolutionary dynamics in

[4] We also refer the reader to Sandholm (2010) for a complete exposition of such classes of games.

potential games, and global asymptotic stability under pairwise comparison dynamics in contractive games. The replicator dynamics for continuum strategy asymmetric games has been considered in Mendoza-Palacios and Hernández-Lerma (2015) where they study the stability of critical points of the system. Cheung (2016) studies local stability and global convergence in continuum strategy potential games under the imitative dynamics. Lahkar and Riedel (2015) was dedicated to studying convergence in potential and negative semidefinite games under the continuum strategy logit dynamic. Very recently, convergence results for the same classes of games under the perturbed best response dynamics generated out of arbitrary perturbation functions was considered in Lahkar et al. (2022b).

The model considered in our paper is more aligned towards Ely and Sandholm (2005) who were the first to introduce the notion of heterogeneity in finite strategy population games and coined the name Bayesian population games for the corresponding heterogeneous population. Ely and Sandholm (2005) study the conditions necessary for the Bayesian best response correspondence to be unique. Under such a scenario, they explore the disequilibrium behavior of the Bayesian best response dynamic and also provide results under which the dynamic converges to the equilibrium in some classes of games. In a companion paper, Sandholm et al. (2007) studies the evolutionary stability of purified equilibrium in two-player normal form games under Bayesian evolution. Later Zusai (2023) deals with heterogeneous population finite strategy games generated out of mean-field revision protocols. Until very recently, the study of a perturbed version of the Bayesian best response dynamic has been considered in Mukherjee and Roy (2021) for games with finitely many strategies. To the best our knowledge, ours is the first paper to study evolutionary games with heterogeneous populations and a continuum of strategies.

6.2 Preliminaries and Model

Let $\mathscr{S}$ be a compact convex subset of $\mathbb{R}$ endowed with the Borel sigma algebra $\mathscr{B}$. Consider a large population of heterogeneous players with unit mass exercising strategies from $\mathscr{S}$. By heterogeneity, we mean that the players in the population are diverse and are characterized by types which are distributed over a complete separable metric space (Ω, d_Ω) according to some given Borel probability measure ξ.[5] The probability measure ξ thus accounts for the heterogeneity within the population. We call the probability space (Ω,ξ) the **type space** of the population. Let $\mathbf{Lip}(\mathscr{S})$ denote the collection of all Lipschitz functions on $\mathscr{S}$. The bounded Lipschitz norm is defined as

$$\|g\|_{\mathbf{BL}} := \sup_{x\in\mathscr{S}} |g(x)| + \sup_{x\neq y} \frac{|g(x)-g(y)|}{|x-y|}, \quad \text{for all } g\in \mathbf{Lip}(\mathscr{S}).$$

Let $\mathscr{P}$ denote the collection of all probability measures and let $\mathscr{M}$ denote the collection of all finite signed measures on $\mathscr{S}$. The dual $\mathbf{BL}^*$-norm on $\mathscr{M}$ is then defined as

$$\|\mu-\nu\|_{\mathbf{BL}^*} := \sup_{g\in\mathbf{Lip}(\mathscr{S}):\|g\|_{\mathbf{BL}}\leq 1} \Big(\int_{\mathscr{S}} f\,d\mu - \int_{\mathscr{S}} f\,d\nu\Big), \quad \text{for all } \mu,\nu\in\mathscr{M}. \tag{6.1}$$

[5]The probability measure ξ is defined on the Borel subsets $\mathscr{B}_\Omega$ of Ω induced by the underlying metric d_Ω.

It is well-known (Dudley (1966); Perkins and Leslie (2014), Pg. 185) that the space of finite signed measures $\mathscr{M}$ endowed with the $\mathbf{BL}^*$-norm is a Banach space. Moreover, the weak topology on $\mathscr{P}$ coincides with the topology generated by the $\mathbf{BL}^*$-norm on $\mathscr{M}$. Therefore, weak convergence of probability measures on $\mathscr{S}$ is equivalent to convergence with respect to the $\mathbf{BL}^*$-norm. We now proceed with the definition of a Bayesian strategy.

Definition 39 *A **Bayesian strategy** is a ξ-measurable map $\theta:\Omega\to\mathscr{P}$.*[6]

Thus, a Bayesian strategy θ maps every type $\omega\in\Omega$ to a distribution $\theta(\omega)\in\mathscr{P}$ over the strategies. We say two Bayesian strategies θ and ρ are equivalent if $\theta(\omega)=\rho(\omega)$ for ξ-a.e $\omega\in\Omega$. The equivalence class of Bayesian strategies is denoted by Θ. We conclude this section with the definition of Bayesian population games followed by some real-life examples. Let $\mathscr{C}$ denote the collection of continuous maps on $\mathscr{S}$ endowed with the sup-norm topology, that is, $\|g\|_\infty := \sup_{x\in\mathscr{S}}|g(x)|$ for $g\in\mathscr{C}$.

Definition 40 *A **Bayesian population game** is a map $\mathscr{G}:\Theta\times\Omega\to\mathscr{C}$ such that for every $\theta\in\Theta$, the mapping*

$$\Omega\ni\omega\mapsto\mathscr{G}(\theta,\omega)\in\mathscr{C}\quad \text{is } \mathscr{B}_\Omega/\mathscr{B}(\mathscr{C}) \text{ measurable.} \tag{6.2}$$

The above definition can be interpreted in the following way. For a Bayesian strategy $\theta\in\Theta$, the payoff to a player of type $\omega\in\Omega$ exercising strategy $x\in\mathscr{S}$ is given by $\mathscr{G}_x(\theta,\omega)$ such that the map in (6.2) is measurable. We shall assume throughout the paper that the game is bounded, that is, there exists $M>0$ such that for every $\theta\in\Theta$, we have $\|\mathscr{G}(\theta,\omega)\|_\infty\le M$ for ξ-a.e $\omega\in\Omega$. Below, we provide some examples of a Bayesian population games arising from the fields of biology and economics.

Example 17 (Bayesian aggregative population games) *In the literature of evolutionary games with a continuum of strategies, homogeneous aggregative population games are a special class of population games where the payoff at a strategy depends on the state (distribution of players over the strategies) of the population only through the expected value of the state. Some examples of such games include the class of large population aggregative potential games (Lahkar (2017), Pg. 447); which, in principle, includes large population Cournot games. This idea can similarly be extrapolated to heterogeneous populations. For example, Ely and Sandholm (2005) studied the Bayesian best response dynamics in such classes of games with finitely many strategies, though they did not call them Bayesian aggregative population games.*

Below, we mathematically introduce the notion of such games when the strategy space is a continuum and the population is heterogeneous. As we shall eventually see, the payoff in such a game depends on the state of the population (Bayesian strategy) θ only through the expectation $\mathscr{E}(\theta)$ and hence the name Bayesian aggregative population games.

[6] A function $\theta:\Omega\to\mathscr{P}$ is ξ-measurable if it is the pointwise ξ-almost everywhere (a.e) limit of a sequence of simple functions $\theta_n:\Omega\to\mathscr{P}$. Whereas, the usual notion measurability would require that the inverse image of any Borel subset (generated by open subsets w.r.t the weak topology) of $\mathscr{P}$ under the map θ should lie in $\mathscr{B}_\Omega$. Thus, the notion of ξ-measurability is a stronger notion of measurability. We require such a condition on the Bayesian strategies to enable us to view them as elements of an appropriate Böchner space. We refer the reader to Diestel and Uhl Jr (1978) for a complete exposition on the theory of Böchner spaces.

Recall that $\mathscr{P}$ and $\mathscr{C}$ respectively denote the collection of probability distributions and continuous functions on the strategy space $\mathscr{S}$. For a homogeneous population game $\varphi : \mathscr{P} \to \mathscr{C}$, the payoff to strategy x at an aggregate population state μ is $\varphi_x(\mu)$. A population game φ is said to be weakly continuous if

$$\mu_n \to^{\boldsymbol{BL}^*} \mu \quad \text{implies that} \quad \varphi(\mu_n) \to \varphi(\mu) \text{ uniformly.}[7]$$

Let Φ_{wc} denote the space of weakly continuous population games. Since $\mathscr{P}$ is compact under the $\mathbf{BL}^*$-norm and that $\mathscr{C}$ is a Polish space under the sup-norm topology, it follows that the space Φ_{wc} is a Polish space (Munkres (2014)). Consequently, Φ_{wc} is a possible candidate for the type space of a population. Suppose now, that the type space Ω of the population is Φ_{wc}. As usual, let Θ be the collection of Bayesian strategies. Then the expectation operator is a mapping $\mathscr{E} : \Theta \to \mathscr{P}$ such that for every $\theta \in \Theta$,

$$\mathscr{E}[\theta](A) := \int_{\Omega} \theta[\omega](A)\xi(d\omega) \quad \text{for all } A \in \mathscr{B},$$

where the above integral is a Böchner integral.[8] Finally, a mapping $\mathscr{G} : \Theta \times \Phi_{wc} \to \mathscr{C}$ defined as

$$\mathscr{G}(\theta, \varphi) := \varphi(\mathscr{E}(\theta)) \quad \text{for all } \theta \in \Theta, \text{ and all } \varphi \in \Phi_{wc}$$

is called a Bayesian aggregative population game. Since types are identified by population games, the payoff function at a Bayesian strategy θ for a population game φ is the payoff of φ at the expectation $\mathscr{E}(\theta)$ of the Bayesian strategy θ. We note that $\mathscr{G}$ is well-defined since for every $\theta \in \Theta$, the mapping $\varphi \mapsto \varphi(\mathscr{E}(\theta))$ is continuous and hence measurable.

Example 18 (Asymmetric War of Attrition Game) *The (symmetric) war of attrition game, first introduced from a biological perspective by Smith (1974), is a model to study conflicts between various species of the animal population. The game involves two animals (players) competing over an indivisible resource. In such a game, the strategy of a player is to select the time to dispute (exit) given that the opponent is still continuing to last in the conflict. The player who persists longer is the winner of the game. Thus, the game induces a trade-off between strategic gains in outlasting the competitor and loss incurred in enduring the conflict over the passage of time. Smith (1974) formulated and studied evolutionary stable strategies (ESS) for such games. Later, Bishop and Cannings (1978) introduced the notion of 'generalized war of attrition game' and obtained mixed evolutionary stable strategies for such games. However, in real life scenarios, the war of attrition game can be asymmetric (heterogeneous) in nature. In particular, the specie of the animal as well as the environmental conditions involved in the conflict might pose undue advantages or disadvantage to the competitors. Hammerstein and Parker (1982) denotes such asymmetry within the population as 'roles' and 'situations' respectively. In this paper, we call such asymmetry as types. Thus, the value and cost of the resources depends on the roles and situations (and therefore, types) in which the game is being played.*

[7]See (B.14) for the definition of $\mathbf{BL}^*$-topology on $\mathscr{P}$.

[8]We refer the reader to Chap II of Diestel and Uhl Jr (1978) for an exposition on Böchner integrals.

For a formal description of such games, suppose that a player with type ω encounters a player with type $\bar{\omega}$. Then the value and cost of the game corresponding to this environment are $V_{\omega,\bar{\omega}}$ and $C_{\omega,\bar{\omega}}$, respectively. As a result, the payoff to the player of type ω persisting time x against the player of type $\bar{\omega}$ persisting time y is given by the following payoff function:

$$g_{\omega,\bar{\omega}}(x,y) := \begin{cases} V_{\omega,\bar{\omega}} - C_{\omega,\bar{\omega}}y & \text{if } x > y \\ \frac{V_{\omega,\bar{\omega}}}{2} - C_{\omega,\bar{\omega}}x & \text{if } x = y \\ -C_{\omega,\bar{\omega}}x & \text{if } x < y. \end{cases}$$

Thus, if we consider the random matching version of g in the Bayesian setup, we have the following asymmetric war of attrition game $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ such that for all $\theta \in \Theta$ and all $\omega \in \Omega$,

$$\mathscr{G}_x(\theta,\omega) := \int_\Omega \int_{\mathscr{S}} g_{\omega,\bar{\omega}}(x,y)\theta(\bar{\omega})(dy)\xi(d\bar{\omega}) \quad \text{for all } x \in \mathscr{S}.$$

Example 19 (Additively Separable Anonymous Game) *An additively separable anonymous game (ASAG) (Zusai (2023), Pg. 1221, Example 1) is an example of a Bayesian population game which appears in the literature of discrete choice models (Anderson et al. (1992)), where the heterogeneity in the payoff is introduced via an additive idiosyncratic shock which depends only on the type of the player for which the payoff is calculated. In such games, the interaction between the Bayesian strategy of the population and the type of a player is additive in nature.*

To illustrate the example mathematically, let us assume that the strategy space of the population is given by $\mathscr{S} := [0,1]$. Suppose that the type space Ω of the population is the collection of continuous functions $\mathscr{C}$ on $\mathscr{S}$ endowed with the sup-norm.[9] In other words, a player of type ω exercising strategy x incorporates an idiosyncratic shock $\omega(x)$ to the payoff of the game. Again, as in Example 17, let $\varphi : \mathscr{P} \to \mathscr{C}$ be a fixed homogeneous population game. Since we shall use this fixed population game φ to define additively separable anonymous games, we shall refer to φ as the common payoff function. This is because ASAG will be defined as a sum of a common payoff function and a type dependent idiosyncratic shock. More formally, an ASAG $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ has the following additively separable structure: for a Bayesian strategy $\theta \in \Theta$, and a type $\omega \in \Omega$, the payoff is given by

$$\mathscr{G}_x(\theta,\omega) := \underbrace{\varphi_x(\mathscr{E}(\theta))}_{\text{common payoff at Bayesian strategy } \theta} + \underbrace{\omega(x)}_{\text{idiosyncratic shock at type } \omega} \quad \text{for all } x \in \mathscr{S},$$

where $\mathscr{E}(\theta)$ denotes the expectation of the Bayesian strategy θ. Thus, according to the definition above, the payoff in an additively separable anonymous game comprises of two components, namely; a common payoff at every Bayesian strategy (homogeneous component) and a type dependent idiosyncratic shock (heterogeneous component). As a result, the overall payoff function of the population is heterogeneous.

[9]Notice that $\mathscr{C}$ is a complete and separable metric space and hence a possible candidate for the type space according to our definition of a Bayesian population game.

Example 20 (Incomplete Information Game) *An incomplete information game is another example of a Bayesian population game. We borrow this example from Zusai (2023) (Pg. 1221, Example 2). To begin with, let us consider a finite set of signals $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$, and let μ be a common prior belief on Σ. Let Ω denote the type space of the population. We denote by $\nu(\cdot|\sigma)$ the conditional distribution on the type space Ω given a signal $\sigma \in \Sigma$. Then, using the Bayes' rule, the posterior belief on Σ can be obtained via the following relation:*

$$\mu(\sigma|\omega) \propto \nu(\omega|\sigma) \times \mu(\sigma) \quad \textit{for } \sigma \in \Sigma. \tag{6.3}$$

Further for every signal σ, let $\mathscr{G}_\sigma(\theta, \omega)$ denote the payoff function of the population at Bayesian strategy θ for the players of type ω. Then the payoff corresponding to the Bayesian population game $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ can be obtained as the expected payoff at the signal dependent Bayesian population game $\mathscr{G}_\sigma$ with respect to the posterior belief on Σ in (6.3) as

$$\mathscr{G}(\theta, \omega) := \sum_\sigma \mathscr{G}_\sigma(\theta, \omega)\mu(\sigma|\omega) \quad \textit{for all } \theta \in \Theta \textit{ and } \omega \in \Omega.$$

6.3 Perturbed Bayesian Best Response Dynamic

In this paper, we aim to introduce the notion of logit Bayesian best response dynamic for Bayesian (heterogeneous) population games with a continuum of strategies. In order to do so, we first require the notion of Bayesian best response correspondence. Let $\langle g, \mu \rangle := \int_{\mathscr{S}} g(x)\mu(dx)$ for $g \in \mathscr{C}$ and $\mu \in \mathscr{P}$. For a Bayesian strategy $\theta \in \Theta$, the expected payoff to a player of type ω against a probability measure $\mu \in \mathscr{P}$ is given by $\langle \mathscr{G}(\theta, \omega), \mu \rangle$. Thus, the **Bayesian best response correspondence** is a mapping $\mathbf{B} : \Theta \rightrightarrows \Theta$ such that for all $\theta \in \Theta$,

$$\mathbf{B}[\theta](\omega) := \operatorname*{argmax}_{\mu \in \mathscr{P}} \langle \mathscr{G}(\theta, \omega), \mu \rangle, \quad \text{for all } \omega \in \Omega. \tag{6.4}$$

The following proposition ensures the measurability of the Bayesian best response correspondence.

Proposition 16 *Let $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ be a Bayesian population game. Then for every $\theta \in \Theta$, the correspondence $\boldsymbol{B}[\theta] : \Omega \rightrightarrows \mathscr{P}$ is ξ-measurable.*

Proof 31 *See Appendix E.1.*

A Bayesian strategy $\theta^\circ \in \Theta$ is called a **Bayesian equilibrium** of the game $\mathscr{G}$ if $\theta^\circ \in \mathbf{B}(\theta^\circ)$. Suppose now that we fix a type $\omega \in \Omega$. Then the **game sliced at type** ω is defined as a mapping $\mathscr{G}_\omega : \Theta \to \mathscr{C}$ such that $\mathscr{G}_\omega(\theta) := \mathscr{G}(\theta, \omega)$ for all $\theta \in \Theta$. Thus, we can view the sliced game $\mathscr{G}_\omega$ as the game corresponding to the underlying subpopulation consisting of players of exactly one type ω. As a result, the ω**-sliced best response** for the game $\mathscr{G}_\omega$ is a mapping $\mathbf{B}_\omega : \Theta \rightrightarrows \mathscr{P}$ such that

$$\mathbf{B}_\omega(\theta) := \mathbf{B}[\theta](\omega), \quad \text{for all } \theta \in \Theta. \tag{6.5}$$

A Bayesian strategy θ° is an ω**-sliced Bayesian equilibrium** of the ω-sliced game $\mathscr{G}_\omega$ if $\theta^\circ(\omega) \in \mathbf{B}_\omega(\theta^\circ)$. The following result establishes a relation between Bayesian equilibria of the original game $\mathscr{G}$ and equilibria of the sliced game counterparts $\mathscr{G}_\omega$. It states that a Bayesian strategy is a Bayesian equilibrium of the game if and only if it is a sliced Bayesian equilibrium for the corresponding sliced counterparts of the game.

Proposition 17 *Let $\mathscr{G}:\Theta\times\Omega\to\mathscr{C}$ be a Bayesian population game. A Bayesian strategy θ° is a Bayesian equilibrium of $\mathscr{G}$ iff θ° is an ω-sliced Bayesian equilibrium of the game $\mathscr{G}_\omega$ for ξ-a.e $\omega\in\Omega$.*

Proof 32 *The proof is similar to the proof of Proposition 3.2 in Mukherjee and Roy (2021).*

To avoid multiplicity of the Bayesian best response correspondence and to make the analysis of best response tractable, we consider the perturbed version of **B**. In what follows, let λ denote the Lebesgue measure on $\mathscr{S}$ and let $\mathscr{P}_{ac}$ denote the collection of all probability measures on $\mathscr{S}$ which admits density with respect to λ.

Definition 41 *Let $\eta>0$. Then the **logit Bayesian best response** (LBBR) under perturbation level η is a mapping $\boldsymbol{L}_\eta:\Theta\to\Theta$ such that for every $\theta\in\Theta$,*

$$\boldsymbol{L}_\eta[\theta](\omega):=\operatorname*{argmax}_{\mu\in\mathscr{P}}\Big[\int_{\mathscr{S}}\mathscr{G}_x(\theta,\omega)\mu(dx)-\eta\boldsymbol{v}(\mu)\Big],\quad \textit{for all } \omega\in\Omega, \tag{6.6}$$

where $\boldsymbol{v}(\mu)$ is the negative Shannon entropy of a probability measure μ and is defined as $\boldsymbol{v}(\mu)=\int_{\mathscr{S}}\mu_x\log\mu_x\,dx$ if $\mu\in\mathscr{P}_{ac}$, and $\boldsymbol{v}(\mu)=\infty$ otherwise.[10]

We refer to $\mathbf{L}_\eta(\theta)$ as the logit Bayesian best response (LBBR) to the Bayesian strategy θ with respect to the perturbation $\mathbf{v}$ at pertubation level η. The fact that $\mathbf{v}(\mu)=\infty$ for all $\mu\notin\mathscr{P}_{ac}$ implies that for every $\theta\in\Theta$, we have $\mathbf{L}_\eta[\theta](\omega)\in\mathscr{P}_{ac}$ for all $\omega\in\Omega$. It is also well-known (Lahkar and Riedel (2015)) that for every $\theta\in\Theta$, and every $\omega\in\Omega$,

$$\mathbf{L}_\eta[\theta][\omega](A)=\int_A\frac{\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))dy}{\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_z(\theta,\omega))dz},\quad \text{for all } A\in\mathscr{B}. \tag{6.7}$$

Before we introduce the notion of logit Bayesian best response dynamic, we argue that the logit Bayesian best response map is well-defined, that is, $\mathbf{L}_\eta(\theta)$ can be expressed as a limit of (simple) Bayesian strategies for every $\theta\in\Theta$. Thus, we have the following proposition.

Proposition 18 *For every $\theta\in\Theta$, the mapping $\theta\mapsto\boldsymbol{L}_\eta(\theta)$ is well-defined.*

Proof 33 *The proof of the proposition follows from Pettis Measurability Theorem (Diestel and Uhl Jr (1978), Pg. 42, Theorem 2).*

In order to define the logit Bayesian best response dynamic, we require the notion of strong topology on the space of Bayesian strategies. In what follows, we first extend the space of Bayesian strategies to the space of *integrable signed Bayesian strategies* $\hat{\Theta}$. Thus, we identify $\hat{\Theta}$ with the collection $\mathscr{L}^1(\Omega,\xi,\mathscr{M})$

[10]Whenever a probability measure μ admits a probability density function, we shall denote the same by μ_x.

of Böchner integrable functions from Ω into the Banach space $\mathscr{M}$. The *strong topology* on $\hat{\Theta}$ is then induced by the following norm:

$$||\hat{\theta}-\hat{\rho}||_{\hat{\Theta}\mathbf{TV}}=\int_{\Omega}||\hat{\theta}(\omega)-\hat{\rho}(\omega)||_{\mathbf{TV}}\xi(d\omega), \quad \text{for all } \hat{\theta},\hat{\rho}\in\hat{\Theta}.$$ [11]

Definition 42 *For a perturbation level $\eta>0$, the* ***logit Bayesian best response dynamic*** *is defined as*

$$\dot{\theta}=\boldsymbol{L}_{\eta}(\theta)-\theta. \tag{6.8}$$ [12]

A rest point of the dynamic (6.8) is a Bayesian strategy θ° that satisfies $\mathbf{L}_{\eta}(\theta^{\circ})=\theta^{\circ}$. In other words, it is a fixed point of the logit Bayesian best response map. We call such a fixed point, a **logit Bayesian equilibrium** of the game $\mathscr{G}$. It follows from Proposition 18 that the logit Bayesian best response map $\mathbf{L}_{\eta}(\theta)$ is well-defined. Also, for every fixed $\theta\in\Theta$, the logit Bayesian best response $\mathbf{L}_{\eta}[\theta](\omega)\in\mathscr{P}$ for every $\omega\in\Omega$. Therefore, by Lahkar and Riedel (2015), we can conclude that Θ is positively invariant under the dynamic (6.8). As a result, we obtain that the logit Bayesian best response dynamic (6.8) is well-defined. Let $\text{LBE}(\mathscr{G},\xi,\eta)$ be the collection of logit Bayesian equilibria at perturbation level η. Thus, notationally we have $\text{LBE}(\mathscr{G},\xi,\eta):=\{\theta^{\circ}\in\Theta:\mathbf{L}_{\eta}(\theta^{\circ})=\theta^{\circ}\}$. We now aim to prove the existence of logit Bayesian equilibria in our setup. In what follows, we first define the following subset of $\mathscr{P}$: for $\beta>1$ let

$$\mathscr{P}_{\beta}:=\Big\{\mu\in\mathscr{P}_{\text{ac}}:\frac{1}{\beta}\leq\mu_{x}\leq\beta \text{ for all } x\in\mathscr{S} \text{ and } |\mu_{x}-\mu_{y}|\leq\beta|x-y| \text{ for all } x,y\in\mathscr{S}\Big\}.$$

Again for $\beta>1$, let us also define the following subset of the space of Bayesian strategies:

$$\Theta_{\beta}:=\Big\{\theta\in\Theta:\theta(\omega)\in\mathscr{P}_{\beta} \text{ for } \xi\text{-a.e } \omega\in\Omega\Big\}.$$

To prove the existence of logit Bayesian equilibrium, we shall first fix a perturbation level $\eta>0$. As we shall observe in the proof of Theorem 15 that for a given $\eta>0$, there exists $\beta>1$ (depending on η) such that $\mathbf{L}_{\eta}(\theta)\in\Theta_{\beta}$ for all $\theta\in\Theta$.[13] Thus, it is enough to prove the existence of a logit Bayesian equilibrium in the subset Θ_{β}. Let $\mathscr{L}^{2}(\lambda):=\{f:\mathscr{S}\to\mathbb{R}:||f||_{\mathscr{L}^{2}}:=\int_{\mathscr{S}}|f(x)|^{2}\lambda(dx)<\infty\}$ denote the collection of all square integrable functions on $\mathscr{S}$. We now make the following observation. Note that since $\text{Range}(\mathscr{G})$ is contained in the space of continuous maps $\mathscr{C}$, the Radon-Nikodym derivative of $\mathbf{L}_{\eta}[\theta](\omega)$ is a continuous probability density function on $\mathscr{S}$ for every $\theta\in\Theta$ and every $\omega\in\Omega$. Furthermore, since the strategy space $\mathscr{S}$ is compact, the convergence of probability density functions with respect to the $||\cdot||_{\hat{\Theta}\mathbf{TV}}$-norm and $||\cdot||_{\hat{\Theta}\mathscr{L}^{2}}$-norm are equivalent.[14] Thus, we may identify Θ_{β} as a subset of $\mathscr{L}^{2}(\lambda)$. We now introduce the notion of weak topology on $\mathscr{L}^{1}(\Omega,\xi,\mathscr{L}^{2})$. For $\hat{\theta}\in\mathscr{L}^{1}(\Omega,\xi,\mathscr{L}^{2})$ and $\hat{\rho}\in\mathscr{L}^{\infty}(\Omega,\xi,\mathscr{L}^{2})$, let $\langle\hat{\theta},\hat{\rho}\rangle(\omega):=\hat{\rho}(\omega)(\hat{\theta}(\omega))=\int_{\mathscr{S}}\hat{\rho}_{x}(\omega)\hat{\theta}_{x}(\omega)\lambda(dx)$ for all $\omega\in\Omega$.[15] Thus, weak topology on $\mathscr{L}^{1}(\Omega,\xi,\mathscr{L}^{2})$

[11] For two probability measures μ and ν, the total variation distance is defined as $||\mu-\nu||_{\mathbf{TV}}:=2\sup_{A\in\mathscr{B}}|\mu(A)-\nu(A)|$.

[12] The derivative in the dynamic is with respect to the $||\cdot||_{\hat{\Theta}\mathbf{TV}}$-topology defined above. See Ely and Sandholm (2005) for more details on how the dynamic is defined for finite strategy games.

[13] We shall remove the dependence of η on β to keep the notations simple.

[14] Here, $||\hat{\theta}-\hat{\rho}||_{\hat{\Theta}\mathscr{L}^{2}}:=\int_{\Omega}||\hat{\theta}(\omega)-\hat{\rho}(\omega)||_{\mathscr{L}^{2}}\xi(d\omega)$.

[15] It follows that $\mathscr{L}^{1}(\Omega,\xi,\mathscr{L}^{2})^{*}\cong\mathscr{L}^{\infty}(\Omega,\xi,\mathscr{L}^{2})$ (Diestel and Uhl Jr (1978), Pg. 98, Theorem 1).

is the topology induced by the following convergence:

$$\hat{\theta}_n \to^w \hat{\theta} \iff \int_\Omega \langle \hat{\theta}_n, \hat{\rho} \rangle(\omega)\xi(d\omega) \to \int_\Omega \langle \hat{\theta}, \hat{\rho} \rangle(\omega)\xi(d\omega), \quad \text{for all } \hat{\rho} \in \mathscr{L}^\infty(\Omega, \xi, \mathscr{L}^2).$$

Definition 43 *A Bayesian population game $\mathscr{G} : \theta \times \Omega \to \mathscr{C}$ is said to be a **weakly continuous**, if for ξ-a.e $\omega \in \Omega$, the mapping $\theta \mapsto \mathscr{G}(\theta, \omega)$ is continuous relative to the weak topology on Θ.*

We now state the main result of this section. It guarantees the existence of a logit Bayesian equilibrium subject to a mild continuity condition on the Bayesian population game.

Theorem 15 *Suppose that the Bayesian population game $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ is bounded and weakly continuous. Then for every $\eta > 0$, the perturbed best response mapping $\theta \mapsto \boldsymbol{L}_\eta(\theta)$ admits a fixed point, that is, there exists $\theta^\circ_\eta \in \Theta$ such that $\boldsymbol{L}_\eta(\theta^\circ_\eta) = \theta^\circ_\eta$.*

Proof 34 *See Appendix E.2.*

Corollary 5 *Consider the type space $(\Phi_{\boldsymbol{wc}}, \xi)$. Let $\mathscr{G} : \Theta \times \Phi_{wc} \to \mathscr{C}$ be a bounded aggregative Bayesian population game as defined in Example 17. Then for every perturbation parameter $\eta > 0$, there exists a logit Bayesian equilibrium.*

Proof 35 *See Appendix E.3.*

Note that it is natural to think of logit Bayesian equilibria as smooth approximations of the original Bayesian equilibria for small enough perturbations η. This requires an appropriate theoretical justification, and towards which we dedicate the next result. Such approximations have been considered in finite strategy homogeneous population games by Hofbauer and Sandholm (2009), continuum strategy homogeneous population games by Lahkar and Riedel (2015) and finite strategy heterogeneous population games by Mukherjee and Roy (2021). To state the desired theorem, we require the notion of $\|\cdot\|_{\hat{\Theta}\mathbf{BL}^*}$-*topology* on $\hat{\Theta}$ which is induced by the following norm:

$$\|\hat{\theta} - \hat{\rho}\|_{\hat{\Theta}\mathbf{BL}^*} = \int_\Omega \|\hat{\theta}(\omega) - \hat{\rho}(\omega)\|_{\mathbf{BL}^*} \xi(d\omega), \quad \text{for all } \hat{\theta}, \hat{\rho} \in \hat{\Theta}.$$

Theorem 16 *Let $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ be a bounded and weakly continuous Bayesian population game. Suppose further that $\mathscr{G}$ is continuous in the $\|\cdot\|_{\hat{\Theta}\boldsymbol{BL}^*}$-topology. Consider a sequence of positive perturbation parameters $(\eta_n)_{n\geq 1}$ such that $\eta_n \downarrow 0$. For $n \geq 1$, let θ°_n be an η_n-logit Bayesian equilibrium of $\mathscr{G}$. If θ°_n converges to some θ° in the $\|\cdot\|_{\hat{\Theta}\boldsymbol{BL}^*}$-topology, then θ° is a Bayesian equilibrium of the original game $\mathscr{G}$.*

Proof 36 *See Appendix E.4.*

Theorem 16 guarantees that any convergent sequence of perturbed Bayesian equilibria along minuscule perturbations approximate a Bayesian equilibrium of the game $\mathscr{G}$. However, one may wonder whether the converse of Theorem 16 also holds. Formally, whether it is true that for every Bayesian equilibrium θ° of $\mathscr{G}$, there exists a sequence of perturbed equilibria that approximates θ° for vanishingly small perturbations. This problem is not addressed in this paper and we leave it for future research.

Corollary 6 *Consider the type space $(\Phi_{\boldsymbol{wc}}, \xi)$. Let $\mathscr{G} : \Theta \times \Phi_{wc} \to \mathscr{C}$ be a bounded aggregative Bayesian population game as defined in Example 17. Then the conclusion of Theorem 16 holds for $\mathscr{G}$.*

Proof 37 *The proof of Corollary 6 is straightforward and is left as an exercise.*

We now establish two fundamental properties of the logit Bayesian best response dynamic, namely, the existence of solutions from arbitrary initial conditions and the continuity of the solutions thus obtained with respect to the initial conditions. We shall resort to the infinite dimensional Picard-Lindelöff theorem (Zeidler (1986), Pg. 79) and the Grönwall's lemma (Zeidler (1986), Pg. 82, Proposition 3.10 and 3.11) in order to prove the result.

Definition 44 *A Bayesian population game $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ is said to be **strongly Lipschitz** if there exists $\kappa > 0$ such that for ξ-a.e $\omega \in \Omega$,*

$$\|\mathscr{G}(\theta, \omega) - \mathscr{G}(\rho, \omega)\|_\infty \le \kappa \|\theta - \rho\|_{\hat{\Theta}\boldsymbol{TV}}, \quad \textit{for all } \theta, \rho \in \Theta. \tag{6.9}$$

In other words, $\mathscr{G}$ is strongly Lipschitz if it is Lipschitz with respect to the $\|\cdot\|_{\hat{\Theta}\mathbf{TV}}$-topology for ξ-a.e $\omega \in \Omega$. As mentioned in Chapter 5, the Lipschitz condition on Bayesian population games is a standard assumption in the literature of Bayesian evolutionary games. The following theorem shows that the strong Lipschitz property of the Bayesian population game is sufficient to ensure the existence of solutions to the logit Bayesian best response dynamic from any given initial condition and to ensure continuity of the corresponding semiflow with respect to the initial conditions.

Theorem 17 *Suppose that the Bayesian population game $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ is strongly Lipschitz. Then the following statements hold true:*

1. *For every initial condition $\theta_0 \in \Theta$, the LBBR dynamic admits a unique solution $(\theta_t)_{t \ge 0}$.*
2. *The semiflow $\zeta : \Theta \times [0, \infty) \to \Theta$ of the LBBR dynamic is continuous in the initial conditions with respect to the strong topology on $\hat{\Theta}$.*

Proof 38 *See Appendix E.5.*

The following corollary is dedicated to the existence of solutions to the logit Bayesian best response dynamic for the class of Bayesian aggregative population games. For $\kappa > 0$, let

$$\mathbf{Lip}_\kappa(\mathscr{P}) := \{\varphi \in \Phi_{wc} : \left\|\varphi(\mu) - \varphi(\nu)\right\|_\infty \le \kappa \|\mu - \nu\|_{\mathbf{TV}} \text{ for all } \mu, \nu \in \mathscr{P}\}. \tag{6.10}$$

Corollary 7 *Let $\mathscr{G} : \Theta \times \Phi_{wc} \to \mathscr{C}$ be Bayesian aggregative population game. Suppose that there exists $\kappa > 0$ such that $\xi(\boldsymbol{Lip}_\kappa(\mathscr{P})) = 1$. Then the conclusions of Theorem 17 hold for $\mathscr{G}$.*

Proof 39 *The proof of the corollary follows immediately and is left to the reader.*

We now aim to prove convergence of solutions to the logit Bayesian equilibria in two classes of games, namely, the class of Bayesian potential games and the class of Bayesian negative semidefinite games. This requires us to find an appropriate norm compact forward invariant subset of the space of Bayesian strategies on which the dynamic can lead to nice convergence results. For $\gamma > 0$ and $\Xi \subseteq \Omega$, let $\Theta_\gamma(\Xi)$ be the collection of Bayesian strategies which are γ-Lipschitz on Ξ.[16] For $\beta > 1, \gamma > 0$ and $\Xi \subseteq \Omega$, let us consider the following subset of Θ:

$$\Theta_{\beta,\gamma}(\Xi) := \Theta_\beta \cap \Theta_\gamma(\Xi).$$

The following lemma establishes relative compactness and forward invariance of the subcollection $\Theta_{\beta,\gamma}(\Xi)$ of Bayesian strategies under some mild conditions on the type space of the population and the Bayesian population game.

Lemma 7 *Fix a perturbation level $\eta > 0$. Suppose that $\Xi \subseteq \Omega$ be compact such that $\xi(\Xi) = 1$. Then the following statements hold true:*

1. *There exist $\beta > 1$ and $\gamma > 0$ (both depending on η) such that $\Theta_{\beta,\gamma}(\Xi)$ is relatively compact with respect to the $\|\cdot\|_{\hat{\Theta}\boldsymbol{TV}}$-topology.*

2. *Suppose further that there exists $\hat{\kappa} > 0$ such that the mapping $\omega \mapsto \mathscr{G}(\theta, \omega)$ is $\hat{\kappa}$-Lipschitz on Ξ for every $\theta \in \Theta$. Then there exist $\beta > 1$ and $\gamma > 0$ (both depending on η and $\hat{\kappa}$) such that $\Theta_{\beta,\gamma}(\Xi)$ is forward invariant with respect to the LBBR dynamic.*

Proof 40 *See Appendix E.6.*

6.4 Convergence in Bayesian Potential Games

In this section we define the notion of Bayesian potential games with a continuum of strategies and prove related convergence results. The notion of potential games for continuum strategy homogeneous population games was introduced by Cheung (2014), and later studied in the context of perturbed best response dynamics by Lahkar and Riedel (2015). On the other hand, finite strategy heterogeneous potential games was introduced very recently by Zusai (2023). In what follows, we first require the notion of the Fréchet derivative of maps between Banach spaces. Let $\mathscr{V}$ and $\mathscr{W}$ be Banach spaces. An operator $\phi : \mathscr{V} \to \mathscr{W}$ is Fréchet differentiable at u if there exists a bounded linear operator $A_u : \mathscr{V} \to \mathscr{W}$ such that

$$\lim_{\|h\|_{\mathscr{V}} \to 0} \frac{\|\phi(u+h) - \phi(u) - A_u(h)\|_{\mathscr{W}}}{\|h\|_{\mathscr{V}}} = 0. \tag{6.11}$$

[16] Let $\gamma > 0$ and $\Xi \subseteq \Omega$. A Bayesian strategy $\theta : \Omega \to \mathscr{P}$ is said to be γ-Lipschitz on Ξ if

$$\|\theta(\omega) - \theta(\bar{\omega})\|_{\mathbf{TV}} \le \gamma d_\Omega(\omega, \bar{\omega}), \quad \text{for all } \omega, \bar{\omega} \in \Xi.$$

Definition 45 *A Bayesian population game $\mathscr{G} : \hat{\Theta} \times \Omega \to \mathscr{C}$ is called an* ***extended Bayesian potential game*** *if there exists a $\mathscr{C}^1$-Fréchet differentiable map $\varphi : \hat{\Theta} \to \mathbb{R}$ such that for every $\theta \in \Theta$,*

$$\nabla^{\boldsymbol{F}} \varphi[\theta](\omega) = \mathscr{G}(\theta, \omega), \quad \textit{for } \xi\textit{-a.e } \omega \in \Omega.^{17} \tag{6.12}$$

The map φ is called the *Bayesian potential function* of the Bayesian potential game $\mathscr{G}$. Note that in homogeneous population games both in the context of finite strategies, its continuum counterpart as well as in heterogeneous population games with finitely many strategies, construction of an appropriate Lyapunov function serves as a necessary tool to prove convergence results in the class of potential games. We apply the same technique here. Thus, we have the following definition.

Definition 46 *For $\beta > 1$, $\gamma > 0$, and $\Xi \subseteq \Omega$ compact, the* ***entropy adjusted Bayesian potential function*** *is a mapping $\hat{\varphi} : \Theta_{\beta,\gamma}(\Xi) \to \mathbb{R}$ defined by*

$$\hat{\varphi}(\theta) = \varphi(\theta) - \underbrace{\int_{\Omega} \int_{\mathscr{S}} \theta_x(\omega) \log \theta_x(\omega) \lambda(dx) \xi(d\omega)}_{=:\hat{\boldsymbol{v}}(\theta)}, \quad \textit{for all } \theta \in \Theta_{\beta,\gamma}(\Xi). \tag{6.13}$$

Observe that the mapping $\hat{\varphi}$ is well-defined since for every $\theta \in \Theta_{\beta,\gamma}(\Xi)$, and for ξ-a.e $\omega \in \Omega$, the quantity $0 \le \theta_x(\omega) \le \beta$ for every $x \in \mathscr{S}$. Further, observe that the Fréchet differentiability of the potential function φ and the Bayesian version $\hat{\mathbf{v}}$ of the Shannon entropy on $\Theta_{\beta,\gamma}(\Xi)$ with respect to the strong topology together imply differentiability (and hence continuity) of the entropy adjusted potential function (Lyapunov function) $\hat{\varphi}$ on $\Theta_{\beta,\gamma}(\Xi)$ for the Bayesian potential game $\mathscr{G}$. We are now ready to state the main convergence result pertaining to Bayesian potential games.

Theorem 18 *Fix a perturbation level $\eta > 0$. Let $\Xi \subseteq \Omega$ be compact with $\xi(\Xi) = 1$. Let $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$ be a Bayesian potential game such that the map $\omega \to \mathscr{G}(\theta, \omega)$ is $\hat{\kappa}$-Lipschitz on Ξ for every $\theta \in \Theta$. Then, there exist $\beta > 1$ and $\gamma > 0$ (both depending on η and $\hat{\kappa}$) such that the following statements hold true for the entropy adjusted Bayesian potential function $\hat{\varphi} : \Theta_{\beta,\gamma}(\Xi) \to \mathbb{R}$ of the game $\mathscr{G}$:*

1. *$\hat{\varphi}$ increases weakly along every solution trajectory to the LBBR dynamic that originates in $\Theta_{\beta,\gamma}(\Xi)$, and increases strictly across every non-stationary solution trajectory.*

2. *The set of omega-limit points (in the $\|\cdot\|_{\hat{\Theta}_{TV}}$-topology) of any trajectory to LBBR dynamic is a non-empty connected compact set of perturbed Bayesian equilibria. Moreover, such limit points are local maximizers of the entropy adjusted Bayesian potential function $\hat{\varphi}$ on $\Theta_{\beta,\gamma}(\Xi)$.*

Proof 41 *See Appendix E.7.*

[17] Here $\nabla^{\mathbf{F}} \varphi$ denotes the gradient of the Fréchet differentiable potential function φ. Note that since φ is a $\mathscr{C}^1$-Fréchet differentiable map, $\nabla^{\mathbf{F}} \varphi[\theta](\omega)$ is a continuous function on $\mathscr{S}$ for every $\theta \in \Theta$ and ξ-a.e $\omega \in \Omega$.

6.5 Convergence in Bayesian Negative Semidefinite Games

In this section we introduce the notion of Bayesian negative semidefinite games for heterogeneous population games with a continuum of strategies. These classes of games have been considered in homogeneous population finite strategy setup by Hofbauer and Sandholm (2009), homogeneous population continuum strategy setup by Lahkar and Riedel (2015), and heterogeneous population finite strategy setup Mukherjee and Roy (2021). We also generalize the notion of self defeating externalities to Bayesian self defeating externalities. Following Cheung (2014), we show the equivalence between Bayesian negative semidefinite games and Bayesian self defeating externalities. Finally, we conclude the section by proving a convergence result related to Bayesian negative semidefinite games.

Definition 47 *An extended Bayesian population game $\mathscr{G}:\hat{\Theta}\times\Omega\to\mathscr{C}$ is called an **extended Bayesian negative semidefinite game** if*

$$\int_{\Omega}\langle\mathscr{G}(\theta,\omega)-\mathscr{G}(\rho,\omega),\theta(\omega)-\rho(\omega)\rangle\xi(d\omega)\leq 0 \quad \textit{for all } \theta,\rho\in\Theta. \tag{6.14}$$

Since $\mathscr{G}$ is bounded, the integral in (6.14) is well-defined. Following Cheung (2014), we now characterize the notion of Bayesian semidefinite games using the notion of Bayesian self defeating externalities. For a fixed $\omega\in\Omega$, recall the ω-sliced game $\mathscr{G}_\omega$ associated to type $\omega\in\Omega$ defined in Section 6.3. Let $\mathscr{M}_0$ be the collection of all finite signed measures on $\mathscr{S}$ whose measure on the full space $\mathscr{S}$ equals 0. We define

$$\Theta_0:=\{\theta_0\in\hat{\Theta}:\theta_0(\omega)\in\mathscr{M}_0 \text{ for } \xi-\text{a.e } \omega\in\Omega\}.$$

In other words, Θ_0 can be seen as the collection of integrable signed Bayesian strategies in $\hat{\Theta}$ that can be expressed as the difference between two Bayesian strategies.

Definition 48 *A Bayesian population game $\mathscr{G}:\Theta\times\Omega\to\mathscr{C}$ satisfies **Bayesian self defeating externalities** if the following two conditions hold:*

1. *The slice map $\theta\mapsto\mathscr{G}_\omega(\theta)$ is $\mathscr{C}^1$-Fréchet differentiable in the norm topology for ξ-a.e $\omega\in\Omega$.*
2. *$\int_\Omega\langle D\mathscr{G}_\omega[\theta](\theta_0),\theta_0(\omega)\rangle\xi(d\omega)\leq 0$, for all $\theta\in\Theta$ and all $\theta_0\in\Theta_0$.*[18]

Cheung (2014) proved that negative semidefiniteness and self defeating externalities are equivalent in the context of homogeneous population games with a continuum of strategies. Our next lemma states that similar result holds in the context of Bayesian negative semidefinite games.

Lemma 8 *Let $\mathscr{G}:\Theta\times\Omega\to\mathscr{C}$ be a Bayesian population game. Suppose that the slice map $\theta\mapsto\mathscr{G}_\omega(\theta)$ is $\mathscr{C}^1$-Fréchet differentiable in the norm topology for ξ-a.e $\omega\in\Omega$. Then, $\mathscr{G}$ is Bayesian negative semidefinite if and only if it satisfies Bayesian self defeating externalities.*

Proof 42 *See Appendix E.8.*

[18]Here $D\mathscr{G}_\omega$ denotes the Fréchet derivative of the sliced game $\mathscr{G}_\omega$.

We now proceed to state the main convergence theorem of this section pertaining to Bayesian negative semidefinite games. In order to do so we need to first define the appropriate Lyapunov function. Recall the norm-compact forward invariant subset $\Theta_{\beta,\gamma}(\Xi)$ for $\gamma>0$, $\beta>1$, and $\Xi\subseteq\Omega$ compact. We then define the Lyapunov function as a map $\hat{\psi}_\eta:\Theta_{\beta,\gamma}(\Xi)\to\mathbb{R}$ such that for all $\theta\in\Theta_{\beta,\gamma}(\Xi)$,

$$\hat{\psi}_\eta(\theta):=\Big[\int_\Omega\langle\mathscr{G}(\theta,\omega),\mathbf{L}_\eta[\theta](\omega)\rangle\xi(d\omega)-\hat{\mathbf{v}}_\eta(\mathbf{L}_\eta(\theta))\Big]$$
$$-\Big[\int_\Omega\langle\mathscr{G}(\theta,\omega),\theta(\omega)\rangle\xi(d\omega)-\hat{\mathbf{v}}_\eta(\theta)\Big]. \tag{6.15}$$

The following theorem ensures the convergence of every solution trajectory of the logit Bayesian best response dynamic under the strong topology to the unique logit Bayesian equilibrium of the Bayesian negative semidefinite game. Notice that the assumption of Fréchet differentiability of the game $\mathscr{G}$, the differentiability of the Bayesian version $\hat{\mathbf{v}}$ of the Shannon entropy, and the forward invariance property of the perturbed Bayesian best response dynamics on $\Theta_{\beta,\gamma}(\Xi)$ together imply differentiability (and hence continuity) of the Lyapunov function $\hat{\psi}_\eta$ on $\Theta_{\beta,\gamma}(\Xi)$ for the Bayesian negative semidefinite game $\mathscr{G}$.

Theorem 19 *Fix a perturbation level $\eta>0$. Let $\Xi\subseteq\Omega$ be compact with $\xi(\Xi)=1$. Let $\mathscr{G}:\Theta\times\Omega\to\mathscr{C}$ be a Bayesian negative semidefinite game such that the slice map $\mathscr{G}_\omega$ is $\mathscr{C}^1$-Fréchet differentiable in the $\|\cdot\|_{\hat{\Theta}\boldsymbol{TV}}$-topology for ξ-a.e $\omega\in\Omega$ and the map $\omega\mapsto\mathscr{G}(\theta,\omega)$ is $\hat{\kappa}$-Lipschitz on Ξ for every $\theta\in\Theta$. Then, there exists $\beta>1$ and $\gamma>0$ (both depending on η and $\hat{\kappa}$) such that the Lyapunov function $\hat{\psi}_\eta:\Theta_{\beta,\gamma}(\Xi)\to\mathbb{R}$ defined in (6.15) decreases monotonically to zero along every solution trajectory to the logit Bayesian best response dynamic that originates in $\Theta_{\beta,\gamma}(\Xi)$, and decreases strictly along every non stationary solutions. Hence, every solution trajectory to the LBBR dynamic in $\Theta_{\beta,\gamma}(\Xi)$ converges with respect to the $\|\cdot\|_{\hat{\Theta}\boldsymbol{TV}}$-topology to the unique logit Bayesian equilibrium.*

Proof 43 *See Appendix E.9.*

6.6 Conclusion

This paper provides an approach for modeling several real-life evolutionary situations by allowing the population to be heterogeneous. More precisely, it analyzes the perturbed best response dynamics for continuum strategy Bayesian (heterogeneous) population games. It characterizes the rest points of the logit Bayesian best response dynamic and thereby establishes the existence of perturbed (logit) equilibria for a Bayesian population game. These equilibria are shown to approximate the original equilibrium of the game for small levels of perturbation. It further establishes the existence of solutions to the logit Bayesian best response dynamic from any initial condition and the continuity of these solutions with respect to the initial conditions. Sufficient conditions for a subset of Bayesian strategies to be norm-compact and forward invariant under the perturbed Bayesian best response dynamic are provided. Finally, given any initial condition in some forward invariant subset, this paper establishes the convergence of the logit Bayesian best response dynamic to the logit equilibrium for two classes of games, namely; Bayesian potential and Bayesian negative semidefinite games.[19]

[19]See Lahkar and Riedel (2015) for the existing results on continuum strategy homogeneous population games.

REFERENCES

Charalambos D. Aliprantis and Kim C Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag Berlin and Heidelberg GmbH & Company KG, 2013. [15, 87, 125, 177, 199]

Carlos Alós-Ferrer and Nick Netzer. The logit-response dynamics. *Games and Economic Behavior*, 68(2): 413–427, 2010. [52]

Simon P Anderson, Andre De Palma, and Jacques-Francois Thisse. *Discrete choice theory of product differentiation*. MIT press, 1992. [85, 107]

Erik J Balder, Maria Girardi, and Vincent Jalby. From weak to strong types of l_e^1-convergence by the bocce-criterion. *arXiv preprint math/9402210*, 1994. [188, 205]

Dieter Balkenborg, Josef Hofbauer, and Christoph Kuzmics. Refined best reply correspondence and dynamics. *Theoretical Economics*, 8(1):165–192, 2013. [9]

Sarvesh Bandhu and Ratul Lahkar. Evolutionary robustness of dominant strategy implementation. *Economic Theory*, pages 1–37, 2022. [53]

Michel Benaım and Morris W Hirsch. Mixed equilibria and dynamical systems arising from fictitious play in perturbed games. *Games and Economic Behavior*, 29(1-2):36–72, 1999. [8, 82, 88, 103]

Michel Benaïm and Jörgen W Weibull. Deterministic approximation of stochastic evolution in games. *Econometrica*, 71(3):873–903, 2003. [34, 35, 36, 44, 51, 145, 149, 150]

Nam P Bhatia and George P Szegö. *Dynamical systems: stability theory and applications*, volume 35. Springer, 2006. [81, 138, 140, 193, 197, 209, 213]

Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2008. [159]

Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013. [14, 126, 152, 153, 202, 203]

Ken Binmore and Larry Samuelson. Muddling through: Noisy equilibrium selection. *journal of economic theory*, 74(2):235–265, 1997. [52]

Kenneth G Binmore, Larry Samuelson, and Richard Vaughan. Musical chairs: Modeling noisy evolution. *Games and economic behavior*, 11(1):1–35, 1995. [34, 52]

DT Bishop and Chris Cannings. A generalized war of attrition. *Journal of theoretical biology*, 70(1): 85–124, 1978. [84, 106]

Lawrence E Blume. The statistical mechanics of strategic interaction. *Games and economic behavior*, 5 (3):387–424, 1993. [52, 54]

Lawrence E Blume. How noise matters. *Games and Economic Behavior*, 44(2):251–271, 2003. [52, 54]

Lawrence E Blume. Population games. In *The economy as an evolving complex system II*, pages 425–460. CRC Press, 2018. [52]

Tilman Börgers and Rajiv Sarin. Learning through reinforcement and replicator dynamics. *Journal of economic theory*, 77(1):1–14, 1997. [34]

Richard T Boylan. Continuous approximation of dynamical systems with randomly matched individuals. *Journal of Economic Theory*, 66(2):615–625, 1995. [34]

Haim Brezis. *Functional analysis, Sobolev spaces and partial differential equations*, volume 2. Springer, 2011. [202]

Charles Castaing and Michel Valadier. *Convex Analysis and Measurable multifunctions*. Springer, 1977. [198]

Man-Wah Cheung. Pairwise comparison dynamics for games with continuous strategy space. *Journal of Economic Theory*, 153:344–375, 2014. [9, 10, 22, 25, 34, 35, 52, 56, 58, 81, 97, 99, 103, 113, 115, 193, 209]

Man-Wah Cheung. Imitative dynamics for games with continuous strategy space. *Games and Economic Behavior*, 99:206–223, 2016. [9, 35, 52, 104]

Man-Wah Cheung and Ratul Lahkar. Nonatomic potential games: the continuous strategy case. *Games and Economic Behavior*, 108:341–362, 2018. [9, 10, 27, 28, 30, 33, 52, 56, 57, 62, 74]

Luis C Corchon. Comparative statics for aggregative games the strong concavity case. *Mathematical Social Sciences*, 28(3):151–165, 1994. [27, 57]

Pierre Coucheney, Bruno Gaujal, and Panayotis Mertikopoulos. Penalty-regulated dynamics and robust learning procedures in games. *Mathematics of Operations Research*, 40(3):611–633, 2015. [80, 87, 89, 96]

J Diestel and JJ Uhl Jr. Vector measures. *Bull. Am. Math. Soc*, 84(4):681–685, 1978. [80, 105, 106, 109, 110, 175, 176, 177, 178, 199]

Rl Dudley. Convergence of baire measures. *Studia Mathematica*, 3(27):251–268, 1966. [105]

Glenn Ellison. Learning, local interaction, and coordination. *Econometrica: Journal of the Econometric Society*, pages 1047–1071, 1993. [78]

Jeffrey C Ely and William H Sandholm. Evolution in bayesian games i: theory. *Games and Economic Behavior*, 53(1):83–109, 2005. [5, 80, 82, 83, 84, 90, 93, 102, 104, 105, 110]

Dietmar Ferger. A continuous mapping theorem for the argmax-functional in the non-unique case. *Statistica Neerlandica*, 58(1):83–96, 2004. [15]

Dean Foster and Peyton Young. Stochastic evolutionary game dynamics. *Theoretical population biology*, 38(2):219–232, 1990. [51]

Drew Fudenberg and David M Kreps. Learning mixed equilibria. *Games and economic behavior*, 5(3): 320–367, 1993. [82]

Drew Fudenberg and David K. Levine. *The theory of learning in games*, volume 2. MIT press, 1998. [8, 19, 35, 42, 80, 89, 102]

Itzhak Gilboa and Akihiko Matsui. Social stability and equilibrium. *Econometrica: Journal of the Econometric Society*, pages 859–867, 1991. [8, 79, 102]

Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. *Quantal Response Equilibrium: A Stochastic Theory of Games*. Princeton University Press, 2016. [8]

Peter Hammerstein and Geoffrey A Parker. The asymmetric war of attrition. *Journal of Theoretical Biology*, 96(4):647–682, 1982. [85, 102, 106]

Harald Hanche-Olsen and Helge Holden. The kolmogorov–riesz compactness theorem. *Expositiones Mathematicae*, 28(4):385–394, 2010. [21, 199]

John C Harsanyi. Oddness of the number of equilibrium points: a new proof. *International Journal of Game Theory*, 2(1):235–250, 1973. [20]

Amélie Héliou, Matthieu Martin, Panayotis Mertikopoulos, and Thibaud Rahier. Online non-convex optimization with imperfect feedback. *Proceedings of the 34th International Conference on Neural Information Processing Systems (Accepted)*, 2020. [13, 19, 20, 80, 88, 89, 128, 131, 133]

Einar Hille and Ralph Saul Phillips. *Functional analysis and semi-groups*, volume 31. American Mathematical Soc., 1996. [175]

Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of convex analysis*. Springer Science & Business Media, 2004. [135]

Josef Hofbauer. Stability for the best response dynamics. *preprint*, 1995. [5, 8, 79, 102]

Josef Hofbauer and Ed Hopkins. Learning in perturbed asymmetric games. *Games and Economic Behavior*, 52(1):133–152, 2005. [5, 8, 26, 35, 88, 103]

Josef Hofbauer and William H Sandholm. On the global convergence of stochastic fictitious play. *Econometrica*, 70(6):2265–2294, 2002. [5, 8, 9, 19, 20, 31, 35, 79, 80, 88, 89, 91, 103]

Josef Hofbauer and William H Sandholm. Evolution in games with randomly disturbed payoffs. *Journal of economic theory*, 132(1):47–69, 2007. [5, 8, 10, 25, 27, 34, 35, 37, 45, 46, 47, 48, 49, 54, 71, 72, 79, 80, 88, 97, 103, 156, 158, 159, 194]

Josef Hofbauer and William H Sandholm. Stable games and their dynamics. *Journal of Economic theory*, 144(4):1665–1693, 2009. [10, 26, 98, 99, 111, 115]

Josef Hofbauer and Karl Sigmund. Adaptive dynamics and evolutionary stability. *Applied Mathematics Letters*, 3(4):75–79, 1990. [24]

Josef Hofbauer and Sylvain Sorin. Best response dynamics for continuous zero-sum games. *Discrete and Continuous Dynamical Systems Series B*, 6(1):215, 2006. [5, 103]

Josef Hofbauer, Jörg Oechssler, and Frank Riedel. Brown–von neumann–nash dynamics: the continuous strategy case. *Games and Economic Behavior*, 65(2):406–429, 2009. [9, 10, 24, 25, 30, 35, 52, 53, 62, 63, 103]

Yu-Guan Hsieh, Kimon Antonakopoulos, and Panayotis Mertikopoulos. Adaptive learning in continuous games: Optimal regret bounds and convergence to nash equilibrium. In *Conference on Learning Theory*, pages 2388–2422. PMLR, 2021. [102]

Michihiro Kandori, George J Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica: Journal of the Econometric Society*, pages 29–56, 1993. [52, 77]

Yuri M Kaniovski and H Peyton Young. Learning dynamics in games with stochastic perturbations. *Games and economic behavior*, 11(2):330–363, 1995. [82]

Srinivasan Kesavan. *Functional analysis*, volume 52. Springer, 2009. [177, 202]

Thomas G Kurtz. Solutions of ordinary differential equations as limits of pure jump markov processes. *Journal of applied Probability*, 7(1):49–58, 1970. [34]

Ratul Lahkar. Large population aggregative potential games. *Dynamic Games and Applications*, 7: 443–467, 2017. [53, 102, 105]

Ratul Lahkar. Elimination of non-individualistic preferences in large population aggregative games. *Journal of Mathematical Economics*, 84:150–165, 2019. [59]

Ratul Lahkar. Convergence to walrasian equilibrium with minimal information. *Journal of Economic Interaction and Coordination*, 15(3):553–578, 2020. [9]

Ratul Lahkar and Saptarshi Mukherjee. Evolutionary implementation in a public goods game. *Journal of Economic Theory*, 181:423–460, 2019. [9, 10, 31, 35]

Ratul Lahkar and Saptarshi Mukherjee. Evolutionary implementation in aggregative games. *Mathematical Social Sciences*, 109:137 – 151, 2021. ISSN 0165-4896. doi: https://doi.org/10.1016/j.mathsocsci.2020.11.004. [9, 35, 52]

Ratul Lahkar and Vinay Ramani. An evolutionary approach to pollution control in competitive markets. *Dynamic Games and Applications*, pages 1–25, 2021. [35]

Ratul Lahkar and Frank Riedel. The logit dynamic for games with continuous strategy sets. *Games and Economic Behavior*, 91:268–282, 2015. [9, 10, 13, 14, 15, 18, 22, 23, 25, 26, 27, 34, 35, 39, 40, 52, 56, 74, 81, 88, 91, 95, 97, 104, 109, 110, 111, 113, 115, 117, 138, 139, 143, 159, 194]

Ratul Lahkar and William H Sandholm. The projection dynamic and the geometry of population games. *Games and Economic Behavior*, 64(2):565–590, 2008. [34]

Ratul Lahkar, Sayan Mukherjee, and Souvik Roy. Generalized perturbed best response dynamics with a continuum of strategies. *Journal of Economic Theory*, 200(105398), 2022a. [35, 52, 56, 62, 63]

Ratul Lahkar, Sayan Mukherjee, and Souvik Roy. Generalized perturbed best response dynamics with a continuum of strategies. *Journal of Economic Theory*, 200(105398), 2022b. [80, 81, 82, 88, 91, 104, 194]

Ratul Lahkar, Sayan Mukherjee, and Souvik Roy. The logit dynamic in supermodular games with a continuum of strategies: A deterministic approximation approach. *Games and Economic Behavior*, 139:133–160, 2023. [52]

Lars-Göran Mattsson and Jörgen W Weibull. Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior*, 41(1):61–78, 2002. [13]

John Maynard Smith. *Evolution and the Theory of Games*. Cambridge university press, 1982. [10, 84]

Saul Mendoza-Palacios and Onésimo Hernández-Lerma. Evolutionary dynamics on measurable strategy spaces: Asymmetric games. *Journal of Differential Equations*, 259(11):5709–5733, 2015. [104]

Panayotis Mertikopoulos and William H Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research*, 41(4):1297–1324, 2016. [87, 88, 96, 103]

Paul Milgrom and John Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica: Journal of the Econometric Society*, pages 1255–1277, 1990. [37]

Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996. [10, 22, 52, 81, 95]

Sayan Mukherjee and Souvik Roy. Regularized bayesian best response learning in finite games. *Games and Economic Behavior*, 149:1–31, 2021. [104, 109, 111, 115, 206]

James Munkres. Topology, 2014. [106]

Roger B Myerson and Philip J Reny. Perfect conditional ε-equilibria of multi-stage games with infinite sets of signals and actions. *Econometrica*, 88(2):495–531, 2020. [31]

Jörg Oechssler and Frank Riedel. Evolutionary dynamics on infinite strategy spaces. *Economic Theory*, 17(1):141–162, 2001. [9, 10, 35, 52, 103]

Jörg Oechssler and Frank Riedel. On the dynamic foundation of evolutionary stability in continuous models. *Journal of Economic Theory*, 107(2):223–252, 2002. [9, 10, 35, 39, 52, 103]

Daisuke Oyama, William H Sandholm, and Olivier Tercieux. Sampling best response dynamics and deterministic equilibrium selection. *Theoretical Economics*, 10(1):243–281, 2015. [52]

Steven Perkins and David S Leslie. Stochastic fictitious play with continuous action sets. *Journal of Economic Theory*, 152:179–213, 2014. [9, 14, 21, 34, 35, 36, 37, 39, 52, 88, 105, 152, 199, 206]

R Tyrrell Rockafellar. *Convex analysis*. Number 28. Princeton university press, 1970. [88, 136, 177, 191, 195]

Walter Rudin. Functional analysis 2nd ed. *International Series in Pure and Applied Mathematics. McGraw-Hill, Inc., New York*, 1991. [202]

Walter Rudin. *Real and complex analysis*. Tata McGraw-hill education, 2006. [133, 168, 200, 201]

William H Sandholm. Potential games with continuous player sets. *Journal of Economic theory*, 97(1): 81–108, 2001. [10, 28, 53, 62, 65]

William H Sandholm. Evolution and equilibrium under inexact information. *Games and economic behavior*, 44(2):343–378, 2003. [34]

William H Sandholm. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132(1):367–382, 2007. [52]

William H Sandholm. *Population games and evolutionary dynamics*. MIT press, 2010. [9, 33, 34, 47, 51, 52, 53, 54, 63, 65, 66, 68, 69, 70, 72, 78, 79, 95, 103, 167]

William H Sandholm. Stochastic imitative game dynamics with committed agents. *Journal of Economic Theory*, 147(5):2056–2071, 2012. [52]

William H Sandholm and Mathias Staudigl. Large deviations and stochastic stability in the small noise double limit. *Theoretical Economics*, 11(1):279–355, 2016. [52]

William H Sandholm et al. Evolution in bayesian games ii: Stability of purified equilibria. *Journal of Economic Theory*, 136(1):641–667, 2007. [82, 104]

Karl H Schlag. Why imitate, and if so, how?: A boundedly rational approach to multi-armed bandits. *Journal of economic theory*, 78(1):130–156, 1998. [34]

A. N. Shiryaev. *Probability, 2nd edition.* Springer Verlag, New York, 1995. [14]

Hal L Smith. *Monotone dynamical systems: an introduction to the theory of competitive and cooperative systems: an introduction to the theory of competitive and cooperative systems.* Number 41. American Mathematical Soc., 1995. [37]

J Maynard Smith. The theory of games and the evolution of animal conflicts. *Journal of theoretical biology*, 47(1):209–221, 1974. [106]

Mathias Staudigl. Stochastic stability in asymmetric binary choice coordination games. *Games and Economic Behavior*, 75(1):372–401, 2012. [52]

Mathias Staudigl. Co-evolutionary dynamics and bayesian interaction games. *International Journal of Game Theory*, 42(1):179–210, 2013. [82]

Donald M Topkis. *Supermodularity and complementarity.* Princeton university press, 1998. [37, 46]

Constantino Tsallis. Possible generalization of boltzmann-gibbs statistics. *Journal of statistical physics*, 52(1-2):479–487, 1988. [19]

Aad W Van Der Vaart and Jon A Wellner. Weak convergence. In *Weak convergence and empirical processes*, pages 16–28. Springer, 1996. [15, 125]

Xavier Vives. Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19(3):305–321, 1990. [37]

H Peyton Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84, 1993. [52, 77]

E Zeidler. Nonlinear functional analysis and its applications, i springer. *New York, NY*, 1986. [92, 112, 134, 135, 186, 187, 204, 205]

Dai Zusai. Evolutionary dynamics in heterogeneous populations: a general framework for an arbitrary type distribution. *International Journal of Game Theory*, 52(4):1215–1260, 2023. [5, 81, 82, 84, 85, 96, 104, 107, 108, 113]

APPENDICES

APPENDIX

A SUPPLEMENT TO CHAPTER 1

A.1 Proof of Theorem 1

The proof of Theorem 1 requires the application of the Brouwer-Schauder-Tychonoff fixed point theorem (see Aliprantis and Border (2013), Pg. 583, Corollary 17.56), which we state below.

Theorem 20 (Brouwer-Schauder-Tychonoff FPT) *Let V be a non-empty compact convex subset of a locally convex Hausdorff space, and let $f: V \to V$ be a continuous function. Then the set of fixed points of f is compact and nonempty.*

We apply Theorem 24 with $V = \mathscr{P}(S)$ and $f(\mathbb{P}) = \mathscr{G}_\eta(\mathbb{P})$. Recall the Prohorov metric ρ. Since $(\mathscr{P}(S), \rho)$ is a metric space, it is Hausdorff. Furthermore, it is convex by definition. To show that f is continuous in the weak topology, it is sufficient to show that $\mathscr{G}_\eta$ is continuous in the weak topology. We show this in the following lemma, which is the key result in the proof of Theorem 1.

Lemma 9 *Suppose that the conditions in Theorem 1 are satisfied. Then the mapping $\mathscr{G}_\eta : \mathscr{P}(S) \to \mathscr{P}(S)$ is continuous in the weak topology.*

The proof of Lemma 9 will, in turn, require the following theorem.

Theorem 21 (Argmax Continuous Mapping Theorem, Van Der Vaart and Wellner (1996)) *Let $(\mathbb{M}_n)_{n\geq 1}$ and $\mathbb{M}$ be stochastic processes indexed by a metric space $\mathscr{H}$ and defined on some common probability space $(\Omega, \mathscr{F}, \mathbb{P})$. Consider a sequence $(\hat{h}_n)_{n\geq 1}$, satisfying*

$$\mathbb{M}_n(\hat{h}_n) \geq \sup_{h\in\mathscr{H}} \mathbb{M}_n(h) - \alpha_n$$

where $\alpha_n = o_{\mathbb{P}}(1)$. Suppose that the following assumptions are fulfilled :

1. $\mathbb{M}_n \rightsquigarrow \mathbb{M}$ *in* $\mathcal{L}^\infty(G)$ *for every compact* $G \subseteq \mathcal{H}$.

2. *The trajectories of* $\mathbb{M}$ *are almost surely upper semicontinuous and posses a unique maximizer* $\hat{h}$, *which is Borel measurable.*

3. *The sequence* $(\hat{h}_n)_{n\geq 1}$ *is uniformly asymptotically tight, i.e. for every* $\epsilon > 0$, *there is a compact set* $G_\epsilon \subseteq \mathcal{H}$ *such that*
$$\limsup_{n\to\infty} \mathbb{P}^*(\hat{h}_n \notin G_\epsilon) \leq \epsilon,$$
where $\mathbb{P}^*$ *denotes the outer probability of* $\mathbb{P}$.

Then
$$\hat{h}_n \rightsquigarrow \hat{h} \;\; in \;\; \mathcal{H}.$$

Proof 44 (Proof of Lemma 9) In order to apply Theorem 21 to Lemma 9, we firstly show that the perturbation function v is weakly lower semicontinuous on $\mathcal{P}_0(S)$. Since $\|\mathbb{Q}_f - \mathbb{Q}_g\|_{\mathbf{TV}} = \frac{1}{2}\|f-g\|_{\mathcal{L}^1(S)}$, it is easy to see that v is strongly convex with respect to the total variation norm. Now we show that the perturbation function v is strongly convex with respect to the Prohorov metric. Fix two probability measures $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(S)$ and $A \in \mathcal{B}(S)$. Let $\epsilon := \|\mathbb{P} - \mathbb{Q}\|_{\mathbf{TV}}$. Then

$$\begin{aligned}
\mathbb{P}(A) &\leq \mathbb{P}(A) - \mathbb{Q}(A) + \mathbb{Q}(A) \\
&\leq \|\mathbb{P} - \mathbb{Q}\|_{\mathbf{TV}} + \mathbb{Q}(A^\epsilon) \\
&= \mathbb{Q}(A^\epsilon) + \epsilon.
\end{aligned}$$

Therefore by the definition of Prohorov metric, we have that $\rho(\mathbb{P}, \mathbb{Q}) \leq \|\mathbb{P} - \mathbb{Q}\|_{\mathbf{TV}}$ for all $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(S)$. This along with the fact that v is strongly convex with respect to the total variation norm implies that v is strongly convex with respect to the Prohorov metric. It then follows from elementary properties of strongly convex functions that

$$v(\tilde{\mathbb{Q}}) \geq v(\mathbb{Q}) + \int_S \nabla^{\mathbf{G}} v(\mathbb{Q}) d(\tilde{\mathbb{Q}} - \mathbb{Q}) + \frac{L}{2}\rho(\tilde{\mathbb{Q}}, \mathbb{Q})^2 \tag{A.1}$$

for all $\mathbb{Q}, \tilde{\mathbb{Q}} \in \mathcal{P}_0(S)$. Now consider a sequence $(\mathbb{Q}_n)_{n\geq 1}, \mathbb{Q} \subset \mathcal{P}_0^c(S)$ such that $\rho(\mathbb{Q}_n, \mathbb{Q}) \to 0$. Setting $\tilde{\mathbb{Q}} = \mathbb{Q}_n$ for $n \geq 1$ in Equation (A.1) we have

$$v(\mathbb{Q}_n) \geq v(\mathbb{Q}) + \int_S \nabla^{\mathbf{G}} v(\mathbb{Q}) d(\mathbb{Q}_n - \mathbb{Q}) + \frac{L}{2}\rho(\mathbb{Q}_n, \mathbb{Q})^2$$

for all $n \geq 1$. Since $\nabla^{\mathbf{G}} v(\mathbb{Q}) \in \mathcal{C}(S)$ for all $\mathbb{Q} \in \mathcal{P}_0^c(S)$, we have that

$$\liminf_{n\to\infty} v(\mathbb{Q}_n) \geq v(\mathbb{Q}).$$

This proves that v is weakly lower semicontinuous on $\mathcal{P}_0^c(S)$.

Consider the metric space $(\mathcal{P}(S), \rho)$. Since S is compact, weak convergence of probability measures and convergence with respect to ρ on the space $\mathcal{P}(S)$ are equivalent (see Billingsley (2013), Pg. 73). Let

$\mathcal{O}(\mathcal{P}(S))$ denote the collection of all open sets in $\mathcal{P}(S)$ under ρ. For $D > 1$, let

$$\mathrm{U}(\mathcal{P}_D(S)) := \left\{ \psi : \mathcal{P}_D(S) \to \mathbb{R} \;\middle|\; \psi \text{ is weakly upper semicontinuous} \right\}.$$

Recall that $\mathcal{P}_D(S)$ is a compact subset of $\mathcal{P}(S)$ under the strong topology. Since $\rho(\mathbb{P},\mathbb{Q}) \le \|\mathbb{P}-\mathbb{Q}\|_{\mathbf{TV}}$ for all $\mathbb{P},\mathbb{Q} \in \mathcal{P}(S)$, it follows that the identity map $\iota : (\mathcal{P}_D(S), \|\cdot\|_{\mathbf{TV}}) \to (\mathcal{P}_D(S), \rho)$ is continuous. Since the image of a compact set under a continuous map is also compact, we have that $\mathcal{P}_D(S)$ is compact under the weak topology. The fact that upper semicontinuous functions on a compact set attains its maximum implies that $\mathrm{U}(\mathcal{P}_D(S)) \subseteq \mathcal{L}^\infty(\mathcal{P}_D(S))$. We consider the sup-norm topology on $\mathrm{U}(\mathcal{P}_D(S))$. Let $\mathcal{B}(\mathrm{U}(\mathcal{P}_D(S)))$ be the Borel sigma algebra on $\mathrm{U}(\mathcal{P}_D(S))$, that is, the smallest sigma algebra containing all open subsets of $\mathrm{U}(\mathcal{P}_D(S))$ with respect to the sup-norm topology. For the rest of the proof, we assume that $\mathrm{U}(\mathcal{P}_D(S))$ is endowed with the Borel sigma algebra $\mathcal{B}(\mathrm{U}(\mathcal{P}_D(S)))$, and $\mathcal{P}(S)$ is endowed with the Borel sigma algebra $\mathcal{B}(\mathcal{P}(S))$ generated by open sets of $\mathcal{P}(S)$ under the weak topology.

Consider a sequence $(\mathbb{P}_n)_{n\ge1}$ in $\mathcal{P}(S)$ such that $\rho(\mathbb{P}_n,\mathbb{P}) \to 0$ as $n \to \infty$. We need to show that $\rho(\mathcal{G}_\eta(\mathbb{P}_n), \mathcal{G}_\eta(\mathbb{P})) \to 0$. Fix a probability space $(\Omega, \mathcal{F}, \hat{\lambda})$.

For every $n \ge 1$, let $\mathbb{M}_n : \Omega \to \mathrm{U}(\mathcal{P}_D(S))$ and $\hat{h}_n : \Omega \to \mathcal{P}_D(S)$ be degenerate random variables, where $\mathbb{M}_n$ is indexed by $\mathcal{P}_D(S)$ and is defined as $\mathbb{M}_n(\mathbb{Q}) = \int_S \pi(\mathbb{P}_n) d\mathbb{Q} - \eta v(\mathbb{Q})$ for all $\mathbb{Q} \in \mathcal{P}_D(S)$, and $\hat{h}_n = \mathcal{G}_\eta(\mathbb{P}_n)$. Similarly, let $\mathbb{M} : \Omega \to \mathrm{U}(\mathcal{P}_D(S))$ and $\hat{h} : \Omega \to \mathcal{P}_D(S)$ be degenerate random variables indexed by $\mathcal{P}_D(S)$ defined as $\mathbb{M}(\mathbb{Q}) = \int_S \pi(\mathbb{P}) d\mathbb{Q} - \eta v(\mathbb{Q})$ for all $\mathbb{Q} \in \mathcal{P}_D(S)$ and $\hat{h} = \mathcal{G}_\eta(\mathbb{P})$. By our formulation, $(\mathbb{M}_n)_{n\ge1}$ and $\mathbb{M}$ are stochastic processes indexed by $\mathcal{P}_D(S)$ and are defined on the common probability space $(\Omega, \mathcal{F}, \hat{\lambda})$. Moreover, by the definition of $\mathcal{G}_\eta(\mathbb{P}_n)$, we have that $\mathbb{M}_n(\mathcal{G}_\eta(\mathbb{P}_n)) \ge \sup_{\mathbb{Q}\in\mathcal{P}_D(S)} \mathbb{M}_n(\mathbb{Q})$ for all $n \ge 1$. We proceed to check Conditions (i), (ii), and (iii) in Theorem 21.

Condition (i): Let $\epsilon > 0$. Define the open ball of radius ϵ around $\mathbb{Q}$ as $\mathrm{B}_\epsilon(\mathbb{Q}) := \{\mathbb{P} \in \mathcal{P}(S) : \rho(\mathbb{Q},\mathbb{P}) < \epsilon\}$. We have

$$\begin{aligned} |\mathbb{M}_n(\mathbb{Q}) - \mathbb{M}(\mathbb{Q})| &= \left| \int_S [\pi_x(\mathbb{P}_n) - \pi_x(\mathbb{P})] \mathbb{Q}(dx) \right| \\ &\le \int_S |\pi_x(\mathbb{P}_n) - \pi_x(\mathbb{P})| \mathbb{Q}(dx). \end{aligned}$$

Now fix $\epsilon > 0$. Then $(\mathrm{B}_\epsilon(\mathbb{Q}))_{\mathbb{Q}\in\mathcal{P}_D(S)}$ is an open cover of $\mathcal{P}_D(S)$. Note that since $(\mathcal{P}_D(S), \rho)$ is compact, there exists a $K_\epsilon \in \mathbb{N}$ and a finite subcollection $\{\mathbb{Q}_1^\epsilon, \dots, \mathbb{Q}_{K_\epsilon}^\epsilon\}$ such that $\mathcal{P}_D(S) \subseteq \cup_{1\le i\le K_\epsilon} \mathrm{B}_\epsilon(\mathbb{Q}_i^\epsilon)$. Let $\mathbb{Q} \in \mathcal{P}_D(S)$. Let $j \in \{1, \dots, K_\epsilon\}$ be such that $\rho(\mathbb{Q}, \mathbb{Q}_j^\epsilon) \le \epsilon$.

Putting these observations together, we have

$$\begin{aligned} \|\mathbb{M}_n - \mathbb{M}\|_\infty &= \sup_{\mathbb{Q}\in\mathcal{P}_D(S)} |\mathbb{M}_n(\mathbb{Q}) - \mathbb{M}(\mathbb{Q})| \\ &\le \max_{1\le i\le K_\epsilon} \int_S |\pi_x(\mathbb{P}_n) - \pi_x(\mathbb{P})| \mathbb{Q}_i^\epsilon(dx) + \epsilon. \end{aligned}$$

By the assumption of the theorem, π is a bounded and weakly continuous map. Therefore by Dominated Convergence Theorem, we have $\int_S |\pi_x(\mathbb{P}_n) - \pi_x(\mathbb{P})| \mathbb{Q}_i^\epsilon(dx) \to 0$ as $n \to \infty$ for all $1 \le i \le K_\epsilon$.

So we have

$$\lim_{n\to\infty} \|\mathbb{M}_n - \mathbb{M}\|_\infty \le \lim_{n\to\infty} \max_{1\le i\le K_\epsilon} \int_S |\pi_x(\mathbb{P}_n) - \pi_x(\mathbb{P})| \mathbb{Q}_i^\epsilon(dx) + \epsilon$$
$$\to \epsilon.$$

Since ϵ is arbitrary, $\mathbb{M}_n$ converges uniformly over $\mathscr{P}_D(S)$ (and therefore over every compact $G \subseteq \mathscr{P}_D(S)$) to $\mathbb{M}$. Since for all $\mathbb{Q} \in \mathscr{P}_D(S)$, $\mathbb{M}_n(\mathbb{Q})$ is degenerate at $\int_S \pi(\mathbb{P}_n) d(\mathbb{Q}) - \eta v(\mathbb{Q})$ for all $n \ge 1$, and $\mathbb{M}(\mathbb{Q})$ is degenerate at $\int_S \pi(\mathbb{P}_n) d(\mathbb{Q}) - \eta v(\mathbb{Q})$, this implies that $\mathbb{M}_n \rightsquigarrow \mathbb{M}$ in $\mathscr{L}^\infty(G)$ for every compact $G \subseteq \mathscr{P}_D(S)$. This verifies Condition (i) of Theorem 21.

Condition (ii): Since the perturbation function v is lower semicontinuous with respect to ρ (and appears with a minus sign in $\mathbb{M}$) and the mapping $\mathbb{Q} \mapsto \int_S \pi(\mathbb{P}) d\mathbb{Q}$ is continuous with respect to ρ, we have that $\mathbb{M}$ has almost surely upper semicontinuous paths on $\mathscr{P}_D(S)$ and a unique maximizer $\mathscr{G}_\eta(\mathbb{P}) \in \mathscr{P}_D(S)$. This verifies Condition (ii) of Theorem 21 since Range$(\pi) \subseteq \mathscr{C}(S)$ and v is admissible.

Condition (iii): Note that the space $(\mathscr{P}_D(S), \rho)$ is compact. Therefore, Condition (iii) holds if we choose $G_\epsilon = \mathscr{P}_D(S)$ for each $\epsilon > 0$. This verifies Condition (iii) of Theorem 21.

Now, we are ready to complete the proof of Lemma 9. By Theorem 21, we have $\hat{h}_n \rightsquigarrow \hat{h}$ in $\mathscr{P}(S)$. Since $\hat{h}_n$ is degenerate at $\mathscr{G}_\eta(\mathbb{P}_n)$ for all $n \ge 1$, and $\hat{h}$ is degenerate at $\mathscr{G}_\eta(\mathbb{P})$, it follows that $\rho(\mathscr{G}_\eta(\mathbb{P}_n), \mathscr{G}_\eta(\mathbb{P})) \to 0$. This completes the proof of Lemma 9.

The proof of Theorem 1 is now complete by Theorem 21 and Lemma 9.

A.2 Proof of Theorem 2

Let $(\eta_n)_{n\ge1}$ be a sequence of positive perturbation parameters such that $\eta_n \downarrow 0$. Let $\mathbb{P}_n^\circ$ be a sequence of η_n-perturbed equilibrium for every $n \ge 1$. By compactness of $\mathscr{P}(S)$ under the weak topology, there exists a subsequence such that $\mathbb{P}_{n_k}^\circ \to^w \mathbb{P}^\circ$. We need to show that $\mathbb{P}^\circ$ is a Nash equilibrium of the original game. We prove this by contradiction. Suppose that $\mathbb{P}^\circ$ is not a Nash equilibrium. Then there exist $y \in S$ and $x \in \text{supp}(\mathbb{P}^\circ)$, disjoint open balls B_y and B_x, an $\epsilon > 0$, and a large enough N_ϵ such that for all $n_k \ge N_\epsilon$, we have

$$\inf_{z\in B_y} \pi_z(\mathbb{P}_{n_k}^\circ) \ge 2\epsilon + \sup_{z\in B_x} \pi_z(\mathbb{P}_{n_k}^\circ). \tag{A.2}$$

Weak convergence of $\mathbb{P}_{n_k}^\circ$ ensures the convergence of $\mathbb{P}_{n_k}^\circ(B_x)$. We show that $\mathbb{P}_{n_k}^*(B_x)$ converges to 0. First, we establish the result under Condition A1. Set $E = B_x$, $F = B_y$, and $g_{n_k} = \pi(\mathbb{P}_{n_k}^\circ)$. For $n_k > N_\epsilon$, it follows from (A.2) that $\inf_{s\in F} g_{n_k}(s) > \sup_{t\in E} g_{n_k}(t)$. Since v satisfies Condition A1, it follows that for all $n_k > N_\epsilon$,

$$\mathrm{D}\tilde{v}_{\eta_{n_k}}^*(\pi(\mathbb{P}_{n_k}^\circ))(\mathbb{I}_{B_y}) \ge u(\eta_{n_k}) \mathrm{D}\tilde{v}_{\eta_{n_k}}^*(\pi(\mathbb{P}_{n_k}^\circ))(\mathbb{I}_{B_x}). \tag{A.3}$$

By Danskin's theorem (Héliou et al. (2020), Pg. 13, Proof of Lemma B.1), we have

$$\mathrm{D}\tilde{v}_{\eta_{n_k}}^*(\pi(\mathbb{P}_{n_k}^\circ))(\mathbb{I}_{B_y}) = \langle \nabla^{\mathbf{F}} \tilde{v}_{\eta_{n_k}}^*(\pi(\mathbb{P}_{n_k}^\circ)), \mathbb{I}_{B_y} \rangle.$$

Here, the dual pairing $\langle\cdot,\cdot\rangle:\mathscr{L}^1(S)\times\mathscr{L}^\infty(S)\to\mathbb{R}$ is defined as $\langle f,g\rangle:=\int_S f(x)g(x)dx$ for all $f\in\mathscr{L}^1(S)$ and all $g\in\mathscr{L}^\infty(S)$. Moreover,

$$\begin{aligned}\langle\nabla^{\mathbf{F}}\tilde{v}^*_{\eta_{n_k}}(\pi(\mathbb{P}^\circ_{n_k})),\mathbb{I}_{B_y}\rangle &= \int_S \nabla^{\mathbf{F}}\tilde{v}^*_{\eta_{n_k}}(\pi(\mathbb{P}^\circ_{n_k}))(t)\mathbb{I}_{B_y}(t)dt\\ &= \int_{B_y} \nabla^{\mathbf{F}}\tilde{v}^*_{\eta_{n_k}}(\pi(\mathbb{P}^\circ_{n_k}))(t)dt\\ &= \xi(\nabla^{\mathbf{F}}\tilde{v}^*_{\eta_{n_k}}(\pi(\mathbb{P}^\circ_{n_k})))(B_y)\\ &= \mathbb{P}^\circ_{n_k}(B_y).\end{aligned}$$

Combining all these observations, we have $\mathrm{D}\tilde{v}^*_{\eta_{n_k}}(\pi(\mathbb{P}^\circ_{n_k}))(\mathbb{I}_{B_y})=\mathbb{P}^\circ_{n_k}(B_y)$. Similarly, we can obtain $\mathrm{D}\tilde{v}^*_{\eta_{n_k}}(\pi(\mathbb{P}^\circ_{n_k}))(\mathbb{I}_{B_x})=\mathbb{P}^\circ_{n_k}(B_x)$. Therefore, (A.3) implies $\mathbb{P}^\circ_{n_k}(B_y)\geq u(\eta_{n_k})\mathbb{P}^\circ_{n_k}(B_x)$ for $n_k>N_\epsilon$. Since $u(\eta_{n_k})\to\infty$ as $k\to\infty$, this implies that $\mathbb{P}^\circ_{\eta_{n_k}}(B_x)\to 0$. Moreover, since $\mathbb{P}^\circ_{n_k}\to^w\mathbb{P}^\circ$, we have by Portmanteau's theorem that $\mathbb{P}^\circ(B_x)=0$. However, this leads to a contradiction since $x\in\mathrm{supp}(\mathbb{P}^\circ)$. Therefore $\mathbb{P}^\circ$ must be a Nash equilibrium of the original game.

Next, we establish the result under Condition A2. Continuing from (A.2), we once again show by contradiction that $\mathbb{P}^\circ_{n_k}(B_x)\to 0$. Suppose that $\mathbb{P}^\circ_{n_k}(B_x)\to c$, where $c>0$. This implies that $\mathbb{P}^\circ_{n_k}(B_x)\geq\frac{c}{2}$ for large k such that $n_k\geq N'$ (say). Since $\mathbb{P}^\circ_{n_k}(B_x)$ is bounded away from 0 for sufficiently large k, there exists a sequence $(A_{n_k})_{k\geq1}\subseteq B_x$ and a sequence of shifted probability measures $(\tilde{\mathbb{P}}^\circ_{n_k})_{k\geq1}$ where $\tilde{\mathbb{P}}^\circ_{n_k}$ is a shift of $\mathbb{P}^\circ_{n_k}$ from A_{n_k} to B_y for sufficiently large k. Then by some algebra we have that

$$\int_{A_{n_k}\cup B_y}\pi_z(\mathbb{P}^\circ_{n_k})\tilde{\mathbb{P}}^\circ_{n_k}(dz)\geq 2\epsilon(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))+\int_{A_{n_k}\cup B_y}\pi_z(\mathbb{P}^\circ_{n_k})\mathbb{P}^\circ_{n_k}(dz)$$

for all $n_k\geq\max\{N_\epsilon,N'\}$. Since $\eta_n\downarrow 0$ and v satisfies Condition A2 with respect to the sequence $(\mathbb{P}^\circ_n)_{n\geq1}$, there exists N'' such that for all $n_k\geq N''$, we have

$$\begin{aligned}|v_{\eta_{n_k}}(\tilde{\mathbb{P}}^\circ_{n_k})-v_{\eta_{n_k}}(\mathbb{P}^\circ_{n_k})| &= |\eta_{n_k}||v(\tilde{\mathbb{P}}^\circ_{n_k})-v(\mathbb{P}^\circ_{n_k})|\\ &\leq\epsilon(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))\end{aligned}$$

for $n_k\geq\max\{N_\epsilon,N',N''\}$. This in turn implies that

$$v_{\eta_{n_k}}(\tilde{\mathbb{P}}^\circ_{n_k})\leq v_{\eta_{n_k}}(\mathbb{P}^\circ_{n_k})+\epsilon(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))$$

for $n_k\geq\max\{N_\epsilon,N',N''\}$. Note that

$$\int_{S\setminus(A_{n_k}\cup B_y)}\pi_z(\mathbb{P}^\circ_{n_k})\mathbb{P}^\circ_{n_k}(dz)=\int_{S\setminus(A_{n_k}\cup B_y)}\pi_z(\mathbb{P}^\circ_{n_k})\tilde{\mathbb{P}}^\circ_{n_k}(dz)$$

for all $k\geq1$. Therefore we have that

$$\int_S\pi_z(\mathbb{P}^\circ_{n_k})\tilde{\mathbb{P}}^\circ_{n_k}(dz)-\eta_{n_k}v(\tilde{\mathbb{P}}^\circ_{n_k})\geq\epsilon(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))+\int_S\pi_z(\mathbb{P}^\circ_{n_k})\mathbb{P}^\circ_{n_k}(dz)-\eta_{n_k}v(\mathbb{P}^\circ_{n_k})$$

for large enough k, contradicting the fact that $\mathbb{P}^\circ_{n_k}$ is a η_{n_k}-perturbed equilibrium for every $n \geq 1$. Therefore $\mathbb{P}^\circ_{n_k}(B_x) \to 0$. Since $\mathbb{P}^\circ_{n_k} \to \mathbb{P}^\circ$, we must have that $\mathbb{P}^\circ(B_x) = 0$, which is a contradiction, since x is in the support of $\mathbb{P}^\circ$. Therefore $\mathbb{P}^\circ$ is a Nash equilibrium of the original game. This concludes the proof of Theorem 2.

A.3 Tsallis Entropy and Burg Entropy

We show that both the Tsallis(γ) entropy, $0 < \gamma < 1$ (Example 13), and the Burg entropy (Example 5.3.1) satisfy Condition A2 with respect to the sequence $(\mathbb{P}^\circ_{n_k})_{k\geq 1}$, where $(\mathbb{P}^\circ_{n_k})_{k\geq 1}$ is as defined in the proof of Theorem 1. To show this, consider two disjoint subsets $E, F \subseteq S$ with $\mathbb{P}^\circ_{n_k}(E) \nrightarrow 0$ with $\inf_{z\in F} \pi_z(\mathbb{P}^\circ_{n_k}) > \sup_{w\in E} \pi_w(\mathbb{P}^\circ_{n_k})$ for sufficiently large k. Since $(\mathbb{P}^\circ_{n_k})_{k\geq 1} \subseteq \mathscr{P}_0^c(S)$, it follows that there exists $(A_{n_k})_{k\geq 1} \subseteq E$ such that $p^\circ_{n_k}$ is bounded away from 0 uniformly of A_{n_k} for sufficiently large k, in which case there exist a sequence of probability measures $(\tilde{\mathbb{P}}^\circ_{n_k})_{k\geq 1}$, where $\tilde{\mathbb{P}}^\circ_{n_k}$ is a shift of $\mathbb{P}^\circ_{n_k}$ from A_{n_k} to F for all $k \geq 1$ (Definition 4).

1. **Tsallis(γ) entropy satisfies Condition A2:** By the definition of shifted probability measures (Definition 4), for every $k \geq 1$, the probability density function of $\tilde{\mathbb{P}}^\circ_{n_k}$ is the same as the density of $\mathbb{P}^\circ_{n_k}$ on the set $S \setminus (A_{n_k} \cup F)$. Therefore it is enough to bound the difference in entropy on $A_{n_k} \cup F$. Note that

$$|v^{\mathbf{t}}(\tilde{\mathbb{P}}^\circ_{n_k}) - v^{\mathbf{t}}(\mathbb{P}^\circ_{n_k})| = \left| \int_{A_{n_k}\cup F} \frac{\tilde{p}^\circ_{n_k}(x) - \tilde{p}^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)} dx - \int_{A_{n_k}\cup F} \frac{p^\circ_{n_k}(x) - p^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)} dx \right|,$$

where $\tilde{p}^\circ_{n_k}$ and $p^\circ_{n_k}$ denote the probability densities of $\tilde{\mathbb{P}}^\circ_{n_k}$ and $\mathbb{P}^\circ_{n_k}$, respectively, for every $k \geq 1$. By Definition 4, we have that

$$\tilde{p}^\circ_{n_k}(x) = \alpha_{n_k} p^\circ_{n_k}(x) \qquad \text{for all } x \in F,$$

where $\alpha_{n_k} = 1 + \dfrac{\mathbb{P}^\circ_{n_k}(A_{n_k}) - \tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k})}{\mathbb{P}^\circ_{n_k}(F)}$ for all $k \geq 1$. Since $\alpha_{n_k} > 1$ for all $k \geq 1$ and $0 < \gamma < 1$, we have that $\alpha^\gamma_{n_k} - 1 < \alpha_{n_k} - 1$ for all $k \geq 1$. Therefore, we have in particular

$$\alpha^\gamma_{n_k} - 1 < \frac{\mathbb{P}^\circ_{n_k}(A_{n_k}) - \tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k})}{\mathbb{P}^\circ_{n_k}(F)} \qquad \text{for all } k \geq 1.$$

First we bound the difference in the entropy on the ball F. We have

$$\begin{aligned}
&\int_F \left| \frac{\tilde{p}^\circ_{n_k}(x) - \tilde{p}^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)} - \frac{p^\circ_{n_k}(x) - p^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)} \right| dx \\
&\leq \int_F \left| \frac{\tilde{p}^\circ_{n_k}(x) - p^\circ_{n_k}(x)}{\gamma(1-\gamma)} \right| dx + \int_F \left| \frac{\tilde{p}^\circ_{n_k}(x)^\gamma - p^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)} \right| dx \\
&\leq \frac{1}{\gamma(1-\gamma)} \left[(\alpha_{n_k} - 1) \int_F p^\circ_{n_k}(x) dx + (\alpha^\gamma_{n_k} - 1) \int_S p^\circ_{n_k}(x)^\gamma dx \right]
\end{aligned}$$

$$\leq \frac{1}{\gamma(1-\gamma)}\Bigg[(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))+\frac{(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))}{\mathbb{P}^\circ_{n_k}(F)}\left(\int_S p^\circ_{n_k}(x)dx\right)^\gamma\Bigg]$$

$$\leq \frac{(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))}{\gamma(1-\gamma)}\Big[1+\frac{1}{\mathbb{P}^\circ_{n_k}(F)}\Big].$$

Note that the second last inequality is due to Jensen's inequaity since the mapping $t\mapsto t^\gamma$ is concave for $0<\gamma<1$. We now show that $\mathbb{P}^\circ_{n_k}(F)$ is bounded away from 0. Since $\mathbb{P}^\circ_{n_k}$ is a Nash equilibrium for every $k\geq 1$, it follows from Héliou et al. (2020) that

$$\mathbb{P}^\circ_{n_k}(A)=\left[\frac{\eta_{n_k}}{1-\gamma}\right]^{1/(1-\gamma)}\int_A \frac{1}{(\theta_{n_k}-\pi_z(\mathbb{P}^\circ_{n_k}))^{1/(1-\gamma)}}dz$$

for all $A\in\mathscr{B}(S)$, where for every $k\geq 1$, we choose $\theta_{n_k}>\|\pi(\mathbb{P}^\circ_{n_k})\|_\infty$ such that

$$\left[\frac{\eta_{n_k}}{1-\gamma}\right]^{1/(1-\gamma)}\int_S \frac{1}{(\theta_{n_k}-\pi_z(\mathbb{P}^\circ_{n_k}))^{1/(1-\gamma)}}dz=1.$$

Letting $\epsilon\to 0$ in (A.2), we have that

$$\int_F \frac{1}{(\theta_{n_k}-\pi_z(\mathbb{P}^\circ_{n_k}))^{1/(1-\gamma)}}dz\geq \int_E \frac{1}{(\theta_{n_k}-\pi_z(\mathbb{P}^\circ_{n_k}))^{1/(1-\gamma)}}dz$$

for all $k\geq 1$. Therefore we have that $\mathbb{P}^\circ_{n_k}(F)\geq \mathbb{P}^\circ_{n_k}(E)$ for all $k\geq 1$. Since $\mathbb{P}^\circ_{n_k}(B_x)$ is bounded away from 0, we have that $\mathbb{P}^\circ_{n_k}(F)$ is also bounded away from 0. That is, there exists $c>0$ such that $\mathbb{P}^\circ_{n_k}(F)\geq c$ for all $k\geq 1$. Therefore in particular we have

$$\int_F \left|\frac{\tilde{p}^\circ_{n_k}(x)-\tilde{p}^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)}-\frac{p^\circ_{n_k}(x)-p^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)}\right|dx\leq \frac{(c+1)(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))}{c\gamma(1-\gamma)}.$$

Now, it follows that

$$\int_{A_{n_k}} \left|\frac{\tilde{p}^\circ_{n_k}(x)-\tilde{p}^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)}-\frac{p^\circ_{n_k}(x)-p^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)}\right|dx\leq \frac{1}{\gamma(1-\gamma)}\Bigg[\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}) + \int_{A_{n_k}} |p^\circ_{n_k}(x)^\gamma-\tilde{p}^\circ_{n_k}(x)^\gamma|dt\Bigg].$$

It now follows that there exists a constant $\tilde{c}>0$ such that for sufficiently large k, we have $p^\circ_{n_k}(x)\geq \tilde{c}$ for all $x\in A_{n_k}$. Therefore, for sufficiently large k we have that the mapping $t\mapsto t^\gamma$ is Lipschitz on A_{n_k}, in which case we have that

$$\int_{A_{n_k}} |p^\circ_{n_k}(x)^\gamma-\tilde{p}^\circ_{n_k}(x)^\gamma|dt\leq \frac{\gamma}{\tilde{c}^{1-\gamma}}(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))$$

for sufficiently large k. Therefore in particular we have that

$$\int_{A_{n_k}} \left|\frac{\tilde{p}^\circ_{n_k}(x)-\tilde{p}^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)} - \frac{p^\circ_{n_k}(x)-p^\circ_{n_k}(x)^\gamma}{\gamma(1-\gamma)}\right| dx \le \frac{\left(1+\frac{\gamma}{\bar{c}^{1-\gamma}}\right)(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))}{\gamma(1-\gamma)}.$$

Finally, for sufficiently large k, we have that

$$|v^{\mathbf{t}}_\eta(\tilde{\mathbb{P}}^\circ_{n_k}) - v^{\mathbf{t}}_\eta(\mathbb{P}^\circ_{n_k})| \le \frac{\eta\left(2c+\frac{\gamma}{\bar{c}^{1-\gamma}}+1\right)(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))}{c\gamma(1-\gamma)}.$$

Since $E, F \subseteq S$ are arbitrary, it follows that Tsallis entropy satisfies Condition A2.

2. **Burg entropy satisfies Condition A2:** As in the previous case, we have that the densities of $\mathbb{P}^\circ_{n_k}$ and $\tilde{\mathbb{P}}^\circ_{n_k}$ remain unchanged on $S\setminus(A_{n_k}\cup F)$ for every $k \ge 1$. Therefore it is enough to bound the difference in entropy on $A_{n_k}\cup F$. Note that,

$$|v^{\mathbf{b}}(\tilde{\mathbb{P}}^\circ_{n_k}) - v^{\mathbf{b}}(\mathbb{P}^\circ_{n_k})| = \left|\int_{A_{n_k}\cup F} \log \tilde{p}^\circ_{n_k}(x)\,dx - \int_{A_{n_k}\cup F} \log p^\circ_{n_k}(x)\,dx\right|.$$

Observe that

$$\tilde{p}^\circ_{n_k}(x) = \alpha_{n_k} p^\circ_{n_k}(x) \qquad \text{for all } x \in F,$$

where $\alpha_{n_k} = 1 + \dfrac{\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k})}{\mathbb{P}^\circ_{n_k}(F)}$ for all $k \ge 1$. Since $\log(1+x) < x$ for all $x > 0$, we have that $\log \alpha_{n_k} < \dfrac{\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k})}{\mathbb{P}^\circ_{n_k}(F)}$ for all $k \ge 1$. Therefore, we have

$$\begin{aligned}
\int_F |\log \tilde{p}^\circ_{n_k}(x) - \log p^\circ_{n_k}(x)|\,dx &= \int_F |\log \alpha_{n_k} p^\circ_{n_k}(x) - \log p^\circ_{n_k}(x)|\,dx \\
&= \int_F |\log \alpha_{n_k}|\,dx \\
&= \int_F \left|\log\left(1 + \frac{\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k})}{\mathbb{P}^\circ_{n_k}(F)}\right)\right| dx \\
&\le \frac{\lambda(F)}{\mathbb{P}^\circ_{n_k}(F)}(\mathbb{P}^\circ_{n_k}(A_{n_k})-\tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k})).
\end{aligned}$$

Since $\mathbb{P}^\circ_{n_k}(E)$ is bounded away from 0 for sufficiently large k, we have by similar arguments that $\mathbb{P}^\circ_{n_k}(F)$ is bounded away from 0 for sufficiently large k. In other words, $\mathbb{P}^\circ_{n_k}(F) \ge c$ for sufficiently large k. Therefore we have

$$\int_F |\log \tilde{p}^\circ_{n_k}(x) - \log p^\circ_{n_k}(x)|\,dx \le \frac{\lambda(F)}{c}(\mathbb{P}^\circ_{n_k} A_{n_k}) - \tilde{\mathbb{P}}^\circ_{n_k}(A_{n_k}))$$

for sufficiently large k. By similar arguments, as in the case of Tsallis entropy, there exists $\bar{c} > 0$

such that that $p^{\circ}_{n_k}(x) \geq \bar{c}$ for all $t \in A_{n_k}$ for sufficiently large k. Therefore, for sufficiently large k, the mapping $t \mapsto \log t$ is Lipschitz of A_{n_k}. This implies that

$$\int_{A_{n_k}} |\log \tilde{p}^{\circ}_{n_k}(x) - \log p^{\circ}_{n_k}(x)| \leq \frac{1}{\bar{c}}(\mathbb{P}^{\circ}_{n_k}(A_{n_k}) - \tilde{\mathbb{P}}^{\circ}_{n_k}(A_{n_k}))$$

for sufficiently large k. So we have

$$|v^{\mathbf{b}}_{\eta}(\tilde{\mathbb{P}}^{\circ}_{n_k}) - v^{\mathbf{b}}_{\eta}(\mathbb{P}^{\circ}_{n_k})| \leq \eta\left(\frac{\lambda(F)}{c} + \frac{1}{\bar{c}}\right)(\mathbb{P}^{\circ}_{n_k}(A_{n_k}) - \tilde{\mathbb{P}}^{\circ}_{n_k}(A_{n_k}))$$

for sufficiently large k. Since $E, F \subseteq S$ are arbitrary, it follows that Burg entropy satisfies Condition A2.

A.4 Appendix to Section 2.4

Proof of Lemma 1

Recall the definitions of $\mathscr{L}^{1,1}_{+}(S)$ in (2.1) and $\mathscr{L}^{1,1}_{+,D}(S)$ in (2.15). Let ξ be the bijective map $\xi : \mathscr{L}^{1,1}_{+}(S) \to \mathscr{P}_0(S)$ defined by $\xi(f) = \mathbb{Q}_f$, where $\mathbb{Q}_f$ is as defined immediately following (2.1). Note that the restriction of ξ to $\mathscr{L}^{1,1}_{+,D}(S)$ is a bijection between $\mathscr{L}^{1,1}_{+,D}(S)$ and $\mathscr{P}_D(S)$ as defined in (2.14).

Without loss of generality, it is enough to prove the lemma for $\eta = 1$. Let $\pi : \mathscr{P}(S) \to \mathscr{L}^{\infty}(S)$ and let $\nabla^{\mathbf{F}} \tilde{v}^*(g) \in \mathscr{L}^{1,1}_{+,D}(S)$ for all $g \in \mathscr{L}^{\infty}(S)$. We show that Δ_D is forward invariant under the generalized dynamic, that is, $\mathscr{G}_1(\mathbb{P}) \in \mathscr{P}_D(S)$ for all $\mathbb{P} \in \mathscr{P}_D(S)$. Since $\mathscr{L}^1(S)^* \cong \mathscr{L}^{\infty}(S)$ (see Rudin (2006), Pg. 127, Theorem 6.16), we have $\tilde{v}^*(g) = \sup_{f \in \mathscr{L}^1(S)}\left[\langle f, g\rangle - \tilde{v}(f)\right]$ for all $g \in \mathscr{L}^{\infty}(S)$, where $\langle f, g\rangle := \int_S f(x)g(x)dx$. Since $\tilde{v}$ is closed and L-strong convex, we have for all $g \in \mathscr{L}^{\infty}(S)$ (see Héliou et al. (2020), Pg. 13, Proof of Lemma B1),

$$\nabla^{\mathbf{F}} \tilde{v}^*(g) = \arg\max_{f \in \mathscr{L}^1(S)} \left[\langle f, g\rangle - \tilde{v}(f)\right]. \tag{A.4}$$

Since $\tilde{v}(f) = \infty$ for all $f \in \mathscr{L}^1(S) \setminus \mathscr{L}^{1,1}_{+}(S)$ we have

$$\arg\max_{f \in \mathscr{L}^1(S)} \left[\langle f, g\rangle - \tilde{v}(f)\right] = \arg\max_{f \in \mathscr{L}^{1,1}_{+}(S)} \left[\langle f, g\rangle - \tilde{v}(f)\right]. \tag{A.5}$$

Therefore,

$$\begin{aligned}
\arg\max_{f \in \mathscr{L}^1(S)} \left[\langle f, g\rangle - \tilde{v}(f)\right] &= \arg\max_{f \in \mathscr{L}^{1,1}_{+}(S)} \left[\langle f, g\rangle - \tilde{v}(f)\right] \\
&= \arg\max_{f \in \mathscr{L}^{1,1}_{+}(S)} \left(\int_S f(x)g(x)\mathrm{d}x - \tilde{v}(f)\right) \\
&= \arg\max_{f \in \mathscr{L}^{1,1}_{+}(S)} \left(\int_S g(x)f(x)\mathrm{d}x - \tilde{v}(f)\right) \\
&= \arg\max_{f \in \mathscr{L}^{1,1}_{+}(S)} \left(\int_S g\,\mathrm{d}\mathbb{Q}_f - v(\mathbb{Q}_f)\right).
\end{aligned}$$

This implies

$$\xi\Big(\arg\max_{f\in\mathscr{L}_+^{1,1}(S)}\big[\langle f,g\rangle-\tilde{v}(f)\big]\Big)=\arg\max_{\mathbb{Q}_f\in\mathscr{P}_0(S)}\Big[\int_S g\,\mathrm{d}\mathbb{Q}_f-v(\mathbb{Q}_f)\Big]. \tag{A.6}$$

By the assumption of the lemma, $\nabla^{\mathbf{F}}\tilde{v}^*(g)\in\mathscr{L}_{+,D}^{1,1}(S)$ for all $g\in\mathscr{L}^\infty(S)$. Fix $\mathbb{P}\in\mathscr{P}_D(S)$. Taking $g=\pi(\mathbb{P})$ and using (A.6) we have

$$\begin{aligned}\mathscr{G}_1(\mathbb{P})&=\arg\max_{\mathbb{Q}g\in\mathscr{P}_0(S)}\Big[\int_S\pi(\mathbb{P})\mathrm{d}\mathbb{Q}_g-v(\mathbb{Q}_g)\Big]\\&=\xi\Big(\arg\max_{f\in\mathscr{L}_+^{1,1}(S)}\big[\langle f,\pi(\mathbb{P})\rangle-\tilde{v}(f)\big]\Big)\\&=\xi(\nabla^{\mathbf{F}}\tilde{v}^*(\pi(\mathbb{P})))\in\mathscr{P}_D(S).\end{aligned}$$

This concludes our proof.

Proof of Theorem 3

We use the generalized Picard-Lindelöf Theorem to prove Theorem 3.

Theorem 22 (The Generalized Picard-Lindelöf Theorem, Zeidler (1986), Pg. 79) *Consider the initial value problem*

$$x'(t)=f(t,x(t));\qquad x(t_0)=y_0$$

Let Y be a Banach space and let $f:\mathbb{R}\times Y\to Y$ be a mapping with the following properties:

1. *f is continuous on $\mathbb{R}\times Y$ and locally Lipschitz continuous with respect to Y; i.e for every $(t_0,y_0)\in\mathbb{R}\times Y$ there exists real numbers $a,b>0$ and $L\geq 0$ such that*

   $$\|f(t,x)-f(t,y)\|\leq L\|x-y\|,$$

   for all $t\in\mathbb{R}$, $x,y\in Y$ with $|t-t_0|\leq a$, $\|x-y_0\|\leq b$, and $\|y-y_0\|\leq b$.

2. *There exists a constant $K>0$ such that $\|f(t,x(t))\|\leq K$ for all t for which $x(\cdot)$ is a solution to the above initial value problem exists.*

Then the above initial value problem has exactly one continuously differentiable solution $x(\cdot)$ on $\mathbb{R}$ for each initial value $y_0\in Y$.

By Theorem 26 it is enough to show that the mapping $\mathbb{P}\mapsto\mathscr{G}_\eta(\mathbb{P})$ is Lipschitz with respect to the total variation norm. We show this in the following lemma.

Lemma 10 *The mapping $\mathbb{P}\mapsto\mathscr{G}_\eta(\mathbb{P})-\mathbb{P}$ is Lipschitz continuous with respect to the total variation norm.*

Proof 45 By Definition 6, we need to show that

$$||\mathscr{G}_\eta(\mathbb{P})-\mathscr{G}_\eta(\tilde{\mathbb{P}})||_{\mathbf{TV}}\leq||\mathbb{P}-\tilde{\mathbb{P}}||_{\mathbf{TV}} \tag{A.7}$$

for all $\mathbb{P},\tilde{\mathbb{P}} \in \mathscr{P}(S)$. Consider the mapping $\xi : (\mathscr{L}^{1,1}_+(S), \|\,.\,\|_{\mathscr{L}^1(S)}) \to (\mathscr{P}_0(S), \|\,.\,\|_{\mathbf{TV}})$. We show that ξ is 1-Lipschitz. For $f, g \in \mathscr{L}^{1,1}_+(S)$, we have

$$\begin{aligned}\|\xi(f)-\xi(g)\|_{\mathbf{TV}} &= \|\mathbb{Q}_f - \mathbb{Q}_g\|_{\mathbf{TV}} \\ &= \frac{1}{2}\int_S |f(x)-g(x)|\,dx \\ &\le \|f-g\|_{\mathscr{L}^1(S)}.\end{aligned}$$

We complete the proof in two steps: first we prove that the mapping $\mathbb{P} \mapsto \mathscr{G}_\eta(\mathbb{P})$ is Lipschitz with respect to the total variation norm, and second, we use this to prove the same for $\mathbb{P} \mapsto \mathscr{G}_\eta(\mathbb{P}) - \mathbb{P}$. Since π is β-Lipschitz, we have $\|\pi(\mathbb{P})-\pi(\tilde{\mathbb{P}})\|_\infty \le \beta\|\mathbb{P}-\tilde{\mathbb{P}}\|_{\mathbf{TV}}$ for all $\mathbb{P},\tilde{\mathbb{P}} \in \mathscr{P}(S)$. Let $\tilde{v}_\eta(f) := \eta\tilde{v}(f)$ for all $f \in \mathscr{L}^1(S)$. Since $\tilde{v}$ is closed and strongly convex, $\eta\tilde{v}$ is closed and ηL-strongly convex. Therefore, $\|\nabla^{\mathbf{F}}\tilde{v}^*_\eta(g) - \nabla^{\mathbf{F}}\tilde{v}^*_\eta(\tilde{g})\|_\infty \le \dfrac{1}{\eta L}\|g-\tilde{g}\|_\infty$ for all $g, \tilde{g} \in \mathscr{L}^\infty(S)$ (see Hiriart-Urruty and Lemaréchal (2004), Pg. 240 Theorem 4.2.1). Therefore it follows from Lemma 1 that

$$\begin{aligned}\|\mathscr{G}_\eta(\mathbb{P})-\mathscr{G}_\eta(\tilde{\mathbb{P}})\|_{\mathbf{TV}} &= \|\xi(\nabla^{\mathbf{F}}\tilde{v}^*_\eta(\pi(\mathbb{P}))) - \xi(\nabla^{\mathbf{F}}\tilde{v}^*_\eta(\pi(\tilde{\mathbb{P}})))\|_{\mathbf{TV}} \\ &\le \|\nabla^{\mathbf{F}}\tilde{v}^*_\eta(\pi(\mathbb{P})) - \nabla^{\mathbf{F}}\tilde{v}^*_\eta(\pi(\tilde{\mathbb{P}}))\|_{\mathscr{L}^1(S)} \\ &\le \frac{1}{\eta L}\|\pi(\mathbb{P})-\pi(\tilde{\mathbb{P}})\|_\infty \\ &\le \frac{\beta}{\eta L}\|\mathbb{P}-\tilde{\mathbb{P}}\|_{\mathbf{TV}}.\end{aligned}$$

Finally,

$$\begin{aligned}\|(\mathscr{G}_\eta(\mathbb{P})-\mathbb{P})-(\mathscr{G}_\eta(\tilde{\mathbb{P}})-\tilde{\mathbb{P}})\|_{\mathbf{TV}} &\le \|\mathscr{G}_\eta(\mathbb{P})-\mathscr{G}_\eta(\tilde{\mathbb{P}})\|_{\mathbf{TV}} + \|\mathbb{P}-\tilde{\mathbb{P}}\|_{\mathbf{TV}} \\ &\le \left(\frac{\beta}{\eta L}+1\right)\|\mathbb{P}-\tilde{\mathbb{P}}\|_{\mathbf{TV}}.\end{aligned}$$

This completes the proof of Lemma 10.

The proof of Lemma 10 establishes one part of Theorem 3, the existence of a unique solution from every initial state. To prove the continuity of the semiflow (Definition 7) of the dynamic, we apply Gronwall's Lemma which we state below.

Lemma 11 (Zeidler (1986), Pg. 82, Proposition 3.10) *Let $u, v : [t_0-a, t_0+a] \to V$ be maps into the subspace V of a Banach space Y. Suppose $f : [t_0-a, t_0+a]\times V \to Y$ is continuous and Lipschitz continuous with respect to u, i.e.,*

$$\|f(t,u)-f(t,v)\| \le L\|u-v\| \text{ for all } t \in [t_0-a, t_0+a] \text{ and } u, v \in V \text{ and fixed } L > 0.$$

Then $\|u(t)-v(t)\| \le e^{L|t-t_0|}\|u(0)-v(0)\|$ for all $t \in [t_0-a, t_0+a]$.

We apply Lemma 29 with $Y = \mathscr{M}(S), V = \mathscr{P}(S)$ and $f = \mathscr{G}_\eta$. By Lemma 10, the conditions of Lemma 29 are satisfied. Therefore, there exists $k > 0$ such that for any two solutions $(\mathbb{P}_t)_{t\ge0}$ and $(\mathbb{Q}_t)_{t\ge0}$ to the Generalized Dynamic with initial conditions $\mathbb{P}_0$ and $\mathbb{Q}_0$, respectively, we have $\|\mathbb{P}_t - \mathbb{Q}_t\|_{\mathbf{TV}} \le e^{kt}\|\mathbb{P}_0 -$

$\mathbb{Q}_0\|_{\mathbf{TV}}$. This establishes continuity of the semiflow of the Generalized Dynamic. This concludes the proof of Theorem 3.

A.5 Appendix to Section 6.4

We first prove Lemma 2. For that, we require another lemma.

Lemma 12 *Suppose $\Phi : \mathcal{M}(S) \to \mathbb{R}$ is the perturbed expected payoff in a population game π. That is $\Phi(\mu) = \int_S \pi(\mathbb{P}) d\mu - v(\mu)$ where v satisfies the properties as stated before. Then there exists a unique maximizer $\hat{\mu}$ such that*

$$dv(\hat{\mu})(\nu) = \int_S \pi(\mathbb{P}) d\nu.$$

Proof 46 The fact v is strongly convex and lower semicontinuous imply that there exists an unique maximizer of Φ, say $\hat{\mu}$. Then,

$$\begin{aligned}
0 &= \mathrm{d}\Phi(\hat{\mu})(\nu) \\
&= \lim_{\epsilon \to 0} \frac{\Phi(\hat{\mu} + \epsilon\nu) - \Phi(\hat{\mu})}{\epsilon} \\
&= \lim_{\epsilon \to 0} \frac{\epsilon \int_S \pi(\mathbb{P}) d\nu - v(\hat{\mu} + \epsilon\nu) + v(\hat{\mu})}{\epsilon} \\
&= \int_S \pi(\mathbb{P}) d\nu - \mathrm{d}v(\hat{\mu})\nu.
\end{aligned}$$

Therefore we must have $\mathrm{d}v(\hat{\mu})(\nu) = \int_S \pi(\mathbb{P}) d\nu$.

Proof 47 (Proof of Lemma 2) We need to show that $\int_S \tilde{\pi}_x(\mathbb{P})\dot{\mathbb{P}}(dx) \geq 0$ for all $\mathbb{P} \in \mathscr{P}_0(S)$. In other words, we need to show that $\int_S \tilde{\pi}_x(\mathbb{P})(\mathscr{G}_\eta(\mathbb{P}) - \mathbb{P})(dx) \geq 0$. By Lemma 12 we have $\int_S \pi(\mathbb{P}) d\mathbb{P} = \int_S \nabla^{\mathbf{G}} v_\eta(\mathscr{G}_\eta(\mathbb{P})) d\mathbb{P}$ for all $\mathbb{P} \in \mathscr{P}_0(S)$. This yields

$$\int_S \tilde{\pi}_x(\mathbb{P})(\mathscr{G}_\eta(\mathbb{P}) - \mathbb{P})(dx) = \int_S [\nabla^{\mathbf{G}} v_\eta(\mathscr{G}_\eta(\mathbb{P})) - \nabla^{\mathbf{G}} v_\eta(\mathbb{P})] d(\mathscr{G}_\eta(\mathbb{P}) - \mathbb{P}).$$

Since v_η is strongly convex (and hence convex), it follows (see Rockafellar (1970), Pg. 242, Theorem 25.1) that $v_\eta(\tilde{\mathbb{Q}}) \geq v_\eta(\mathbb{Q}) + \int_S \nabla^{\mathbf{G}} v_\eta(\mathbb{Q}) d(\tilde{\mathbb{Q}} - \mathbb{Q})$ for all $\tilde{\mathbb{Q}}, \mathbb{Q} \in \mathscr{P}(S)$. Setting $\tilde{\mathbb{Q}} = \mathscr{G}_\eta(\mathbb{P})$ and $\mathbb{Q} = \mathbb{P}$, and interchanging the roles of $\tilde{\mathbb{Q}}$ and $\mathbb{Q}$, we obtain

$$\int_S [\nabla^{\mathbf{G}} v_\eta(\mathscr{G}_\eta(\mathbb{P})) - \nabla^{\mathbf{G}} v_\eta(\mathbb{P})] d(\mathscr{G}_\eta(\mathbb{P}) - \mathbb{P}) \geq 0. \tag{A.8}$$

It remains to show that equality holds in (A.8) only if $\mathscr{G}_\eta(\mathbb{P}) = \mathbb{P}$, (i.e, $\dot{\mathbb{P}} = 0$). We first show that

$$\mathrm{d}v_\eta(\mathbb{P})\mu = \int_S \Theta' \circ p(x)\mu(dx) \tag{A.9}$$

for all $\mathbb{P}\in\mathscr{P}_0(S)$ and all $\mu\in\mathscr{M}_0(S)$. Fix $\mathbb{P}\in\mathscr{P}_0(S)$ and $\mu\in\mathscr{M}_0(S)$ with density functions p and μ_0 respectively. Since the perturbation function satisfies (2.17), it follows that for every $\epsilon>0$

$$v_\eta(\mathbb{P}+\epsilon\mu)=\int_S \Theta\circ(p+\mu_0)(x)dx=\int_S \Theta\circ p(x)dx+\epsilon\mu_0(x)\Theta'\circ p(x)dx+\mathbf{o}(\epsilon^2). \tag{A.10}$$

Therefore (A.9) follows that Equation (A.10) and the notion of the Gateaux derivative. It now follows that $\nabla^{\mathbf{G}}v_\eta(\mathbb{P})=\Theta'\circ p$ for all $\mathbb{P}\in\mathscr{P}_0(S)$.

We now proceed to prove part (i) of Lemma 2. Let p_η denote the density function of $\mathscr{G}_\eta(\mathbb{P})$. Therefore, by Equation (A.9), we have $\nabla^{\mathbf{G}}v_\eta(\mathscr{G}_\eta(\mathbb{P}))(x)=\Theta'\circ p_\eta(x)$ and $\nabla^{\mathbf{G}}v_\eta(\mathbb{P})(x)=\Theta'\circ p(x)$ for all $x\in S$. This in particular means that

$$\int_S[\nabla^{\mathbf{G}}v_\eta(\mathscr{G}_\eta(\mathbb{P}))-\nabla^{\mathbf{G}}v_\eta(\mathbb{P})]d(\mathscr{G}_\eta(\mathbb{P})-\mathbb{P})=\int_S(\Theta'\circ p_\eta(x)-\Theta'\circ p(x))(p_\eta(x)-p(x))dx. \tag{A.11}$$

By assumption of part (i) the lemma, Θ is either strictly increasing or strictly decreasing on $[0,\infty)$. Therefore the integrand in the RHS Equation (A.11) is either non negative or non positive for all $x\in S$. This implies that equality holds in Equation (A.11) if and only if $p_\eta(x)=p(x)$ for all $x\in S$. This proves part (i) of Lemmma 2.

We now prove part (ii) of Lemma 2. It follows from Lemma 1 that for every $\eta>0$, $\mathscr{G}_\eta(\mathbb{P})\in\mathscr{P}_{D_\eta}(S)$ for every $\mathbb{P}\in\mathscr{P}(S)$. This in particular means that $D_\eta^{-1}\le p_\eta(x)\le D_\eta$ for all $x\in S$. By assumption of part (ii) of the lemma, there exists $0\le\alpha<1<\beta$ such that Θ is either strictly increasing or strictly decreasing on $[\alpha,\beta]$.

Let $\eta>0$ be such that $[D_\eta^{-1},D_\eta]\subset[\alpha,\beta]$. It then follows that

$$(\Theta'\circ p_\eta(x)-\Theta'\circ p(x))(p_\eta(x)-p(x))$$

is either non negative or non positive for all $x\in S$ and all $\mathbb{P}\in\mathscr{P}_{D_\eta}(S)$ according as Θ is strictly increasing or decreasing on $[\alpha,\beta]$. This implies that

$$\int_S[\nabla^{\mathbf{G}}v_\eta(\mathscr{G}_\eta(\mathbb{P}))-\nabla^{\mathbf{G}}v_\eta(\mathbb{P})]d(\mathscr{G}_\eta(\mathbb{P})-\mathbb{P})=\int_S(\Theta'\circ p_\eta(x)-\Theta'\circ p(x))(p_\eta(x)-p(x))dx \tag{A.12}$$

is either non negative or non positive according as Θ is strictly increasing or decreasing on $[\alpha,\beta]$, with equality if and only if $p_\eta(x)=p(x)$ for all $x\in S$. This proves part (ii) of Lemma 2.

This concludes the proof of Lemma 2.

Proof 48 (Proof of Theorem 4) Let $\{\mathbb{P}_t\}_{t\ge0}$ be the solution trajectory with initial condition $\mathbb{P}_0\in\mathscr{P}_D(S)$. We have already shown that $\mathscr{P}_D(S)$ is forward invariant under the generalized dynamic. This implies that $\mathbb{P}_t \in \mathscr{P}_D(S)$ for all $t > 0$ whenever $\mathbb{P}_0 \in \mathscr{P}_D(S)$. Note that

$$\begin{aligned}
\dot{\tilde{\varphi}}(\mathbb{P}) &= \dot{\varphi}(\mathbb{P}) - \dot{v}_\eta(\mathbb{P}) \\
&= \mathrm{D}\varphi(\mathbb{P})\dot{\mathbb{P}} - \mathrm{d}v_\eta(\mathbb{P})\dot{\mathbb{P}} \\
&= \langle \nabla^{\mathbf{G}}\varphi(\mathbb{P}), \dot{\mathbb{P}}\rangle - \langle \nabla^{\mathbf{G}} v_\eta(\mathbb{P}), \dot{\mathbb{P}}\rangle \\
&= \langle \pi(\mathbb{P}), \dot{\mathbb{P}}\rangle - \langle \nabla^{\mathbf{G}} v_\eta(\mathbb{P}), \dot{\mathbb{P}}\rangle \\
&= \int_S \pi(\mathbb{P})\mathrm{d}\dot{\mathbb{P}} - \int_S \nabla^{\mathbf{G}} v_\eta(\mathbb{P})\mathrm{d}\dot{\mathbb{P}} \\
&= \int_S [\pi(\mathbb{P}) - \nabla^{\mathbf{G}} v_\eta(\mathbb{P})]\mathrm{d}\dot{\mathbb{P}} \\
&= \int_S \tilde{\pi}(\mathbb{P})\mathrm{d}\dot{\mathbb{P}} \geq 0.
\end{aligned}$$

The last inequality follows from Lemma 2. Hence $\Lambda(\mathbb{P}_t)$ increases strictly whenever $\mathbb{P}_0$ is not the rest point of the generalized dynamic.

We know that the space $\mathscr{P}_D(S)$ is compact with respect to the strong (total variation) topology. The fact that the vector field in the generalized dynamic is Lipschitz, a continuous and unique solution exists starting from $\mathbb{P}_0 \in \mathscr{P}_D(S)$. Form standard results of dynamical systems theory (see Bhatia and Szegö (2006), Pg. 117, Theorem 2.2.9) it follows that the ω-limit set of the trajectory $\{\mathbb{P}_t\}_{t\geq 0}$ is non-empty compact and connected in the strong topology. Moreover this set is contained in the set consisting of the rest points of the dynamic. Hence any element in the ω-limit set is a rest point of the generalized dynamic and hence a generalized equilibrium. Also since the entropy adjusted potential function is increasing along every solution trajectory, any point in the ω-limit set will be a maximizer of the function. This concludes the proof of Theorem 4.

A.6 Appendix to Section 6.5

Proof 49 (Proof of Lemma 3) We prove the lemma along the lines of Lahkar and Riedel (2015). Since the generalized dynamic is forward invariant on $\mathscr{P}_D(S)$, it suffices to prove uniqueness on $\mathscr{P}_D(S)$. Fix $\mathbb{P} \in \mathscr{P}_D(S)$ and consider the virtual payoff

$$\tilde{\pi}(\mathbb{P}) = \pi(\mathbb{P}) - \nabla^{\mathbf{G}} v_\eta(\mathbb{P})$$

for every $\mathbb{P} \in \mathscr{P}(S)$. Then for every $\mu \in \mathscr{M}_0(S)$ and $t \in \mathbb{R}$, such that $\mathbb{P} + t\mu \in \mathscr{P}_D(S)$, define

$$\varrho(\mathbb{P},\mu,t) := \int_S \tilde{\pi}(\mathbb{P} + t\mu)\mu(dx).$$

We now compute the Gâteaux derivative of $\varrho(\mathbb{P},\mu,.)$ with respect to t. By similar arguments as in Lahkar and Riedel (2015), we obtain

$$\frac{\partial \varrho(\mathbb{P},\mu,t)}{\partial t} = \int_S \mathrm{d}\tilde{\pi}_x(\mathbb{P} + t\mu)\mu(dx).$$

We now compute integrand on the right hand side of the above equation. Note that by definition of $\tilde{\pi}$, we have

$$(\mathrm{d}\tilde{\pi}(\mathbb{P}+t\mu)\mu)(x) = (\mathrm{d}\pi(\mathbb{P}+t\mu)\mu)(x) - \mathrm{d}\nabla^{\mathbf{G}} v_\eta(\mathbb{P}+t\mu)\mu(x).$$

Therefore integrating both sides of the above equation we have,

$$\int_S \mathrm{d}\tilde{\pi}_x(\mathbb{P}+t\mu)\mu(dx) = \int_S \mathrm{d}\pi_x(\mathbb{P}+t\mu)\mu(dx) - \int_S \mathrm{d}\nabla^{\mathbf{G}} v_\eta(\mathbb{P}+t\mu)\mu(dx).$$

Since π is a Fréchet differentiable, the Gâteaux derivative of π coincides with the Fréchet derivative of π. Also since negative semidefinite game, we have

$$\int_S \mathrm{d}\pi_x(\mathbb{P}+t\mu)\mu(dx) = \int_S \mathrm{D}\pi_x(\mathbb{P}+t\mu)\mu(dx) \leq 0.$$

By assumptions of the lemma, we have

$$\int_S \mathrm{d}\nabla^{\mathbf{G}} v_\eta(\mathbb{P}+\epsilon\mu)\mu(dx) \geq 0.$$

This in turn implies that

$$\frac{\partial \varrho(\mathbb{P},\mu,t)}{\partial t} < 0.$$

The proof now follows along the lines of Lahkar and Riedel (2015). This completes the proof of Theorem 3.

Proof 50 (Proof of Theorem 5) We first define some notations which we require in the proof. Let $\Lambda_D^1(\mathbb{P}) := \int_S \pi_x(\mathbb{P})\mathscr{G}_\eta(\mathbb{P})(dx)$, $\Lambda_D^2(\mathbb{P}) := v_\eta(\mathscr{G}_\eta(\mathbb{P}))$, $\Lambda_D^3(\mathbb{P}) := \int_S \pi_x(\mathbb{P})\mathbb{P}(dx)$ and $\Lambda_D^4(\mathbb{P}) := v_\eta(\mathbb{P})$ for all $\mathbb{P} \in \mathscr{P}_0(S)$. Since Gâteaux derivative is linear, it is enough to compute the Gâteaux derivative of $\Lambda_D^i(\mathbb{P})$ for $i = 1, \ldots, 4$ separately. Differentiating Λ_D^1 we have

$$\begin{aligned}
\mathrm{D}\Lambda_D^1(\mathbb{P})\dot{\mathbb{P}} &= \int_S \mathrm{D}\pi(\mathbb{P})\dot{\mathbb{P}}(x)\mathscr{G}_\eta(\mathbb{P})(dx) + \int_S \pi_x(\mathbb{P})\dot{\mathscr{G}}_\eta(\mathbb{P})(dx) \\
&= \langle \mathrm{D}\pi(\mathbb{P})\dot{\mathbb{P}}, \mathscr{G}_\eta(\mathbb{P})\rangle + \langle \pi(\mathbb{P}), \dot{\mathscr{G}}_\eta(\mathbb{P})\rangle.
\end{aligned}$$

Since the perturbation function v is Gâteaux differentiable, we have

$$\begin{aligned}
\mathrm{d}\Lambda_D^2(\mathbb{P})\dot{\mathbb{P}} &= \mathrm{d}v_\eta(\mathscr{G}_\eta(\mathbb{P}))\dot{\mathscr{G}}_\eta(\mathbb{P}) \\
&= \langle \nabla^{\mathbf{G}} v_\eta(\mathscr{G}_\eta(\mathbb{P})), \dot{\mathscr{G}}_\eta(\mathbb{P})\rangle.
\end{aligned}$$

Now, differentiating Λ_D^3 we have

$$\begin{aligned}
\mathrm{D}\Lambda_D^3(\mathbb{P})\dot{\mathbb{P}} &= \int_S \pi_x(\mathbb{P})\dot{\mathbb{P}}\mathbb{P}(dx) + \int_S \pi_x(\mathbb{P})\dot{\mathbb{P}}(dx) \\
&= \langle \mathrm{D}\pi(\mathbb{P})\dot{\mathbb{P}}, \mathbb{P}\rangle + \langle \pi(\mathbb{P}), \dot{\mathbb{P}}\rangle.
\end{aligned}$$

Finally differentiating we have Λ_D^4 we have

$$\begin{aligned}\mathrm{d}\Lambda_D^4(\mathbb{P})\dot{\mathbb{P}} &= \mathrm{d}v_\eta(\mathbb{P})\dot{\mathbb{P}}\\ &= \langle \nabla^{\mathbf{G}} v_\eta(\mathbb{P}), \dot{\mathbb{P}}\rangle.\end{aligned}$$

Combining all the observations from above we have

$$\begin{aligned}\dot{\Lambda}_D(\mathbb{P}) &= \langle \mathrm{D}\pi(\mathbb{P})\dot{\mathbb{P}}, \mathscr{G}_\eta(\mathbb{P})\rangle + \langle \pi(\mathbb{P}), \dot{\mathscr{G}}_\eta(\mathbb{P})\rangle - \langle \nabla^{\mathbf{G}} v_\eta(\mathscr{G}_\eta(\mathbb{P})), \dot{\mathscr{G}}_\eta(\mathbb{P})\rangle\\ &\qquad - \langle \mathrm{D}\pi(\mathbb{P})\dot{\mathbb{P}}, \mathbb{P}\rangle - \langle \pi(\mathbb{P}), \dot{\mathbb{P}}\rangle + \langle \nabla^{\mathbf{G}} v_\eta(\mathbb{P}), \dot{\mathbb{P}}\rangle.\\ &= \langle \mathrm{D}\pi(\mathbb{P})\dot{\mathbb{P}}, \mathscr{G}_\eta(\mathbb{P}) - \mathbb{P}\rangle + \langle \pi(\mathbb{P}), \dot{\mathscr{G}}_\eta(\mathbb{P}) - \dot{\mathbb{P}}\rangle + \langle \nabla^{\mathbf{G}} v_\eta(\mathbb{P}), \dot{\mathbb{P}}\rangle - \langle \nabla^{\mathbf{G}} v_\eta(\mathscr{G}_\eta(\mathbb{P})), \dot{\mathscr{G}}_\eta(\mathbb{P})\rangle.\end{aligned}$$

Since π is a negative semidefinite game and $\dot{\mathbb{P}} = \mathscr{G}_\eta(\mathbb{P}) - \mathbb{P} \in \mathscr{M}_0(S)$, we have from the definition of negative semidefinite game that the first term in the RHS on the above equation, namely $\langle \mathrm{D}\pi(\mathbb{P})\dot{\mathbb{P}}, \mathscr{G}_\eta(\mathbb{P}) - \mathbb{P}\rangle \le 0$. Using this inequality in the RHS of the above expression and rearranging the terms we have

$$\dot{\Lambda}_D(\mathbb{P}) \le \langle \pi(\mathbb{P}) - \nabla^{\mathbf{G}} v_\eta(\mathscr{G}_\eta(\mathbb{P})), \dot{\mathscr{G}}_\eta(\mathbb{P})\rangle - \langle \pi(\mathbb{P}) - \nabla^{\mathbf{G}} v_\eta(\mathbb{P}), \dot{\mathbb{P}}\rangle.$$

Observe that the first term on the right is 0 due to Lemma 12 and the second term on the right is non negative due to Lemma 2. Therefore $\dot{\Lambda}_D(\mathbb{P}) \le 0$ along every solution which originated in $\mathscr{P}_D(\mathrm{S})$. This in particular means the Lyapunov function as defined in (2.22) decrease along every solution trajectory which originates in $\mathscr{P}_D(S)$. Since $\mathscr{P}_D(S)$ is compact under the strong topology and $\dot{\Lambda}_D(\mathbb{P}) \le 0$ with equality only from the rest points, it follows from standard results of dynamical systems theory (Bhatia and Szegö (2006), Pg. 117, Theorem 2.2.9) implies that ω-limit of every solution trajectory satisfies $\dot{\Lambda}_D(\mathbb{P}) = 0$, which is the set of all rest points of the GPBR dynamic. In particular, we have from Lemma 3 every non stationary solution originating in $\mathscr{P}_D(S)$ converges to the unique generalized perturbed equilibrium.

A.7 Appendix to Section 2.7

In this Appendix we calculate the value of the constants $\theta_\eta^{\mathbf{t}}$ in (2.12) and $\theta_\eta^{\mathbf{b}}$ in (2.13) in the following manner.

Tsallis Entropy

Recall from (2.12) that for Tsallis(γ) entropy, the generalized perturbed best response is

$$\mathscr{G}_\eta(\mathbb{P})(A) = \left[\frac{\eta}{1-\gamma}\right]^{1/(1-\gamma)} \int_A \frac{\mathrm{d}x}{(\theta_\eta^{\mathbf{t}} - \pi_x(\mathbb{P}))^{1/(1-\gamma)}},$$

where $\theta_\eta^{\mathbf{t}}$ is such that the integral if the above density is 1. Here we have taken $\gamma = 0.5$. Therefore the generalized best response becomes

$$\mathscr{G}_\eta(\mathbb{P})(A) = 4\eta^2 \int_A \frac{\mathrm{d}x}{(\theta_\eta^{\mathbf{t}} - \pi_x(\mathbb{P}))^2}.$$

Since we are looking at the Cournot competition model (2.23), we have $\pi_x(\mathbb{P}) = (1-\alpha)x - 0.2x$, where $\alpha = A(\mathbb{P})$ is the aggregate output at $\mathbb{P}$ or, equivalently, the mean of the distribution $\mathbb{P}$. Therefore,

$$\mathscr{G}_\eta(\mathbb{P})(A) = \int_A \frac{4\eta^2 \mathrm{d}x}{(\theta_\eta^{\mathbf{t}} - (1-\alpha)x + 0.2x)^2}, \tag{A.13}$$

where $\theta_\eta^{\mathbf{t}}$ is such that the integral of the above density is 1. After computing the above integral we see that the value of $\theta_\eta^{\mathbf{t}}$ for which the integral in (A.13) is 1 is given by

$$\theta_\eta^{\mathbf{t}} = \frac{(0.8-\alpha) + \sqrt{(0.8-\alpha)^2 + 16\eta^2}}{2}.$$

Burg Entropy

Similarly, for the Burg entropy, we have from (2.13) that the perturbed best response is

$$\mathscr{G}_\eta(\mathbb{P})(A) = \eta \int_A \frac{\mathrm{d}x}{(\theta_\eta^{\mathbf{b}} - \pi_x(\mathbb{P}))},$$

where $\theta_\eta^{\mathbf{b}}$ is such that the integral of the above density is 1. For the Cournot competition model (2.23), the generalized perturbed best response becomes

$$\mathscr{G}_\eta(\mathbb{P})(A) = \int_A \frac{\eta \mathrm{d}x}{(\theta_\eta^{\mathbf{b}} - (1-\alpha)x + 0.2x)}. \tag{A.14}$$

Similar to the Tsallis entropy, the value of $\theta_\eta^{\mathbf{b}}$ for which the integral in (A.14) is 1 is

$$\theta_\eta^{\mathbf{b}} = \frac{(0.8-\alpha)\exp\left(\frac{0.8-\alpha}{\eta}\right)}{\exp\left(\frac{0.8-\alpha}{\eta}\right) - 1}.$$

APPENDIX

B

SUPPLEMENT TO CHAPTER 2

B.1 Proof of Proposition 2

To keep the notations simple, we shall, without of loss of generality, fix η to be equal to 1. Now fix $A \in \mathscr{B}(\mathscr{S})$. We then observe that,

$$
\begin{aligned}
|\mu[t](A)-\mu^n[t](A)| &= \left|\int_0^t \dot{\mu}[s](A)\mathrm{d}s - \int_0^t \dot{\mu}^n[s](A)ds\right| \\
&\leq \left|\int_0^t \mathfrak{L}[\mu(s)](A)\mathrm{d}s - \int_0^t \mathfrak{L}^n[\mu^n(s)](A)ds\right| + \left|\int_0^t \mu[s](A)-\mu^n[s](A)\right|ds \\
&\leq \left|\int_0^t \mathfrak{L}[\mu(s)](A)ds - \int_0^t \mathfrak{L}[\mu^n(s)](A)ds\right| + \left|\int_0^t \mu[s](A)-\mu^n[s](A)\right|ds \\
&\qquad + \left|\int_0^t \mathfrak{L}(\mu^n(s))(A)\mathrm{d}s - \int_0^t \mathfrak{L}^n(\mu^n(s))(A)ds\right| \\
&\leq \int_0^t \underbrace{(\|\mathfrak{L}(\mu(s))-\mathfrak{L}(\mu^n(s))\|_{\mathbf{TV}}+\|\mu(s)-\mu^n(s)\|_{\mathbf{TV}})}_{=:\mathscr{T}_1^n(s)}ds \\
&\qquad + \int_0^t \int_A \underbrace{\left|\frac{\exp(\varphi_x(\mu^n(s)))}{\int_{\mathscr{S}}\exp(\varphi_y(\mu^n(s)))dy} - \frac{\exp(\varphi_x^n(\mu^n(s)))}{\int_{\mathscr{S}}\exp(\varphi_y^n(\mu^n(s)))dy}\right|}_{=:\mathscr{T}_2^n(s,x)} dx\,ds
\end{aligned}
$$

It follows from Lahkar and Riedel (2015) that there exists $\kappa > 0$ such that

$$
\begin{aligned}
\mathscr{T}_1^n(s) &:= \|\mathfrak{L}(\mu(s))-\mathfrak{L}(\mu^n(s))\|_{\mathbf{TV}}+\|\mu(s)-\mu^n(s)\|_{\mathbf{TV}} \\
&\leq (1+\kappa)\|\mu(s)-\mu^n(s)\|_{\mathbf{TV}},
\end{aligned} \tag{B.1}
$$

for all $s \geq 0$. We now state and prove a lemma which provides an appropriate upper bound on the second term $\mathscr{T}_2^n(s,x)$ in order for us to apply Grönwall's lemma in the subsequent stages of the proof.

Lemma 13 *Fix $n \geq 1$. Then for every $x \in \mathscr{S}$ and every $s \geq 0$, it holds that:*

$$
\mathscr{T}_2^n(s,x) \leq 2\|f_n - f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty). \tag{B.2}
$$

Proof 51 Fix $x \in \mathscr{S}$. Since the two-player games $(f_n)_{n\geq 1}$ and f are bounded below by definition, it then follows from the definition of $\mathscr{T}_2^n(s,x)$ that,

$$\begin{aligned}
\mathscr{T}_2^n(s,x) &= \left|\frac{\int_S \exp(\varphi_x(\mu^n(s)))\exp(\varphi_y^n(\mu^n(s)))dy - \int_S \exp(\varphi_x^n(\mu^n(s)))\exp(\varphi_y(\mu^n(s)))dy}{\left(\int_S \exp(\varphi_y(\mu^n(s)))dy\right)\left(\int_S \exp(\varphi_y^n(\mu^n(s)))dy\right)}\right| \\
&\le \exp(\|f_n\|_\infty + \|f\|_\infty)\left|\int_S \exp(\varphi_x(\mu^n(s)))\exp(\varphi_y^n(\mu^n(s)))dy - \int_S \exp(\varphi_x^n(\mu^n(s)))\exp(\varphi_y(\mu^n(s)))dy\right| \\
&= \exp(\|f_n\|_\infty + \|f\|_\infty)\left|\int_S \exp[\varphi_x(\mu^n(s)) + \varphi_y^n(\mu^n(s))]dy - \int_S \exp[\varphi_x^n(\mu^n(s)) + \varphi_y(\mu^n(s))]dy\right| \\
&\le \exp(\|f_n\|_\infty + \|f\|_\infty)\Big[\int_S \underbrace{|\exp[\varphi_x(\mu^n(s)) + \varphi_y^n(\mu^n(s))] - \exp[\varphi_x^n(\mu^n(s)) + \varphi_y^n(\mu^n(s))]|}_{=:\mathscr{T}_{21}^n(s,x)} dy \\
&\qquad + \int_S \underbrace{|\exp[\varphi_x^n(\mu^n(s)) + \varphi_y^n(\mu^n(s))] - \exp[\varphi_x^n(\mu^n(s)) + \varphi_y(\mu^n(s))]|}_{=:\mathscr{T}_{22}^n(s,x)} dy\Big].
\end{aligned}$$

The first inequality in the above string of inequalities occurs since the payoff functions $(f_n)_{n\ge1}$ and f are bounded below by $(-\|f_n\|_\infty)_{n\ge1}$ and $-\|f\|_\infty$, in which case the integrals $\int_S \exp(\varphi_y(\mu^n(s)))\mathrm{d}y$ and $\int_S \exp(\varphi_y^n(\mu^n(s)))\mathrm{d}y$ are bounded below by $\exp(-\|f_n\|_\infty)_{n\ge1}$ and $\exp(-\|f\|_\infty)$. Now we bound the terms $\mathscr{T}_{21}^n(s,x)$ and $\mathscr{T}_{22}^n(s,x)$ separately. For any two bounded functions $g_1, g_2 : \mathscr{S} \to \mathbb{R}$, it follows by an application of mean-value theorem that

$$|\exp(g_1(x)) - \exp(g_2(x))| \le \exp(\|g_1\|_\infty \vee \|g_2\|_\infty)|g_1(x) - g_2(x)|, \quad \text{for all } x \in \mathscr{S}. \tag{B.3}$$

Therefore by an application of (B.3), we have for all $s \ge 0$, and $x \in \mathscr{S}$, that

$$\begin{aligned}
\mathscr{T}_{21}^n(s,x) &= \int_{\mathscr{S}} |\exp(\varphi_y^n(\mu^n(s)))||\exp(\varphi_x(\mu^n(s))) - \exp(\varphi_x^n(\mu^n(s)))|dy \\
&\le \exp\|f_n\|_\infty \exp(\|f\|_\infty \vee \|f_n\|_\infty)\int_{\mathscr{S}} |\varphi_x(\mu^n(s)) - \varphi_x^n(\mu^n(s))|dy \\
&\le \exp(\|f_n\|_\infty + \|f\|_\infty \vee \|f_n\|_\infty)\|f - f_n\|_\infty \\
&\le \exp(2\|f\|_\infty \vee \|f_n\|_\infty)\|f - f_n\|_\infty.
\end{aligned}$$

Proceeding similarly as above we observe that

$$\mathscr{T}_{22}^n(s,x) \le \exp(2\|f\|_\infty \vee \|f_n\|_\infty)\|f - f_n\|_\infty, \quad \text{for all } x \in \mathscr{S}.$$

Combining the above inequalities we have,

$$\begin{aligned}
\mathscr{T}_2^n(s,x) &\le \exp(\|f_n\|_\infty + \|f\|_\infty)\big[\mathscr{T}_{21}^n(s,x) + \mathscr{T}_{22}^n(s,x)\big] \\
&\le 2\|f_n - f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty),
\end{aligned}$$

for all $x \in \mathscr{S}$. This completes the proof of Lemma 13.

We now proceed with the proof of the proposition. Using Lemma 13 in conjunction with (B.1), we now arrive at

$$
\begin{aligned}
|\mu[t](A)-\mu^n[t](A)| &\le \int_0^t \mathcal{T}_1^n(s)ds + \int_0^t \mathcal{T}_2^n(s,x)ds \\
&\le (1+\kappa)\int_0^t \|\mu(s)-\mu^n(s)\|_{\mathbf{TV}}ds + 2\int_0^t \|f_n-f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty)ds \\
&\le (\kappa+3)\int_0^t \Big[\|f_n-f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty) + \|\mu(s)-\mu^n(s)\|_{\mathbf{TV}}\Big]ds.
\end{aligned}
$$

Taking supremum over all $A \in \mathcal{B}(\mathcal{S})$, we have

$$
\|\mu(t)-\mu^n(t)\|_{\mathbf{TV}} \le 2(\kappa+3)\int_0^t \Big[\exp(4\|f\|_\infty \vee \|f_n\|_\infty)\|f_n-f\|_\infty + \|\mu(s)-\mu^n(s)\|_{\mathbf{TV}}\Big]ds. \tag{B.4}
$$

We now define a mapping $\Lambda_n : [0,T] \to \mathbb{R}$ such that

$$
\Lambda_n(t) := \|f_n-f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty) + \|\mu(t)-\mu^n(t)\|_{\mathbf{TV}}, \quad \text{for all } 0 \le t \le T.
$$

From the above inequality (B.4), it then follows that

$$
\Lambda_n(t) \le 2\|f_n-f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty) + 2(\kappa+3)\int_0^t \Lambda_n(s)ds, \text{ for all } 0 \le t \le T.
$$

As a result, by Grönwall's inequality, we have for all $0 \le t \le T$ that

$$
\begin{aligned}
\Lambda_n(t) &\le 2\|f_n-f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty) \times \exp(2t\kappa+6t) \\
&\le 2\|f_n-f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty) \times \exp(2T\kappa+6T).
\end{aligned}
$$

Since $f_n \to f$ uniformly, $\|f_n\|_\infty$ is uniformly bounded, in which case it follows using the definition of Λ_n that

$$
\sup_{0 \le t \le T} \|\mu(t)-\mu^n(t)\|_{\mathbf{TV}} \le 2\|f_n-f\|_\infty \exp(4\|f\|_\infty \vee \|f_n\|_\infty) \times \exp(2(\kappa+3)T-1) \to 0, \quad \text{as } n \to \infty.
$$

This completes the proof of Proposition 2.

B.2 Proof of Proposition 3

As in the proof of Proposition 2, we assume without loss of generality that η is equal to 1. We begin the proof of the proposition along the lines of Benaïm and Weibull (2003). To get rid of cumbersome notations in the proof, we shall, for every $n \ge 1$, denote the vector field $\mathcal{P}(\mathcal{S}) \ni \mu \mapsto \mathfrak{L}^n(\mu) - \mu \in \mathcal{M}_0(\mathcal{S})$ by $\mu \mapsto \mathfrak{L}^n_{\mathsf{vect}}(\mu)$ for all $\mu \in \mathcal{P}(\mathcal{S})$. For $N, n \ge 1$, consider the stochastic evolutionary process $(\mathbb{X}_n^N(k))_{k\ge1}$ as defined in (3.16). We also use the notation δ_N to denote the quantity $1/N$ for every $N \ge 1$. Now, let

$\mathscr{U}_n^N : [0,\infty) \to \mathscr{M}(\mathscr{S})$ be a mapping defined as

$$\mathscr{U}_n^N(k\delta_N) = \frac{1}{\delta_N}[\mathbb{X}_n^N((k+1)\delta_N) - \mathbb{X}_n^N(k\delta_N)] - \mathfrak{L}_{\mathsf{vect}}^n(\mathbb{X}_n^N(k\delta_N)), \text{ for all } k \geq 1. \tag{B.5}$$

We now define the continuous time process as $\mathscr{U}_n^N(t) := \mathscr{U}_n^N(k\delta_N)$ for all $k\delta_N \leq t < (k+1)\delta_N$. For $k \geq 1$, let $\mathscr{F}_k$ denote the σ-algebra generated by $\{\mathbb{X}_n^N(0), \ldots, \mathbb{X}_n^N(k\delta_N)\}$. To proceed with the proof, we also require the following notation. Let $\bar{\mathbb{X}}_n^N(t) := \mathbb{X}_n^N(k\delta_N)$, for all $k\delta_N \leq t < (k+1)\delta_N$, be the step process corresponding to $(\mathbb{X}_n^N(k))_{k\geq 1}$. It now follows that

$$\begin{aligned}
\widehat{\mathbb{X}}_n^N(t) - \mu_0^n &= \int_0^t [\mathscr{L}_{\mathsf{vect}}^n(\bar{\mathbb{X}}_n^N(s)) + \mathscr{U}_n^N(s)]ds \\
&= \int_0^t [\mathscr{L}_{\mathsf{vect}}^n(\widehat{\mathbb{X}}_n^N(s)) + \mathscr{L}_{\mathsf{vect}}^n(\bar{\mathbb{X}}_n^N(s)) - \mathscr{L}_{\mathsf{vect}}^n(\widehat{\mathbb{X}}_n^N(s)) + \mathscr{U}_n^N(s)]ds.
\end{aligned}$$

We now make the following important observations about the smoothing operator $\mathfrak{s}_n$ before we proceed with the proof of the proof of Proposition 3. Firstly, we note that the smoothing operator is linear in its arguments. Secondly, we can extend the domain of definition of the smoothing operator in such a way so that it acts as an identity operator on the subset of probability measures which already admits a probability density function. Thus, by definition of the smoothing map $\mathfrak{s}_n$ (Definition 12), it follows that

$$\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t)) - \mathfrak{s}_n(\mu_{\mathbf{p}^n}^N) = \int_0^t [\mathfrak{L}_{\mathsf{vect}}^n(\widehat{\mathbb{X}}_n^N(s)) + \mathfrak{L}_{\mathsf{vect}}^n(\bar{\mathbb{X}}_n^N(s)) - \mathfrak{L}_{\mathsf{vect}}^n(\widehat{\mathbb{X}}_n^N(s)) + \mathfrak{s}_n(\mathscr{U}_n^N(s))]ds.$$

Also since we have $\mu^n(t) - \mu_0^n = \int_0^t \mathfrak{L}_{\mathsf{vect}}^n(\mu^n(s))ds$, we have that

$$\begin{aligned}
\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t)) - \mu^n(t)\|_{\mathbf{TV}} \leq &\int_0^t \|\mathfrak{L}_{\mathsf{vect}}^n(\widehat{\mathbb{X}}_n^N(s)) - \mathfrak{L}_{\mathsf{vect}}^n(\mu^n(s))\|_{\mathbf{TV}} ds \\
&\int_0^t \|\mathfrak{L}_{\mathsf{vect}}^n(\bar{\mathbb{X}}_n^N(s)) - \mathfrak{L}_{\mathsf{vect}}^n(\widehat{\mathbb{X}}_n^N(s))\|_{\mathbf{TV}} ds + \Psi(T),
\end{aligned}$$

where $\Psi(T) = \sup_{0\leq t\leq T} \|\int_0^T \mathfrak{s}_n(\mathscr{U}_n^N(s))ds\|_{\mathbf{TV}}$. We now require the following two lemmas to complete the proof of Proposition 3.

Lemma 14 *Fix $n \geq 1$. Let $\mu, \nu \in \mathscr{P}(\mathscr{S}_n)$ be two arbitrary purely atomic probability measures. Then*

$$\|\mathfrak{s}_n(\mu) - \mathfrak{s}_n(\nu)\|_{\boldsymbol{TV}} \leq \frac{1}{4}\|\mu - \nu\|_{\boldsymbol{TV}}. \tag{B.6}$$

Proof 52 By definition of smoothing of a probability measure, we have that $\mathfrak{s}_n(\mu), \mathfrak{s}_n(\mu)$ are both absolutely continuous with respect to the Lebesgue measure. Now by definition of total variation distance, we have

$$
\begin{aligned}
\|\mathfrak{s}_n(\mu)-\mathfrak{s}_n(\nu)\|_{\mathbf{TV}} &= \frac{1}{2}\int_S \left|\frac{d\mathfrak{s}_n(\mu)}{d\lambda}(x)-\frac{d\mathfrak{s}_n(\nu)}{d\lambda}(x)\right|dx \\
&= \frac{1}{2}\int_S \left|\sum_{k=0}^{2^n-1} 2^n\mu(\{\alpha_{n,k}\})\mathbb{I}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)-\sum_{k=0}^{2^n-1} 2^n\nu(\{\alpha_{n,k}\})\mathbb{I}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)\right|dx \\
&\le \frac{1}{2}\sum_{k=0}^{2^n-1}|\mu(\{\alpha_{n,k}\})-\nu(\{\alpha_{n,k}\})| \\
&= \frac{1}{4}\|\mu-\nu\|_{\mathbf{TV}}.
\end{aligned}
$$

This completes the proof of Lemma 14.

Lemma 15 *Fix $\epsilon>0$. Then for every $\mu,\nu\in\mathscr{P}(\mathscr{S})$, we have*

$$
\|\mathfrak{L}^n_{vect}(\mu)-\mathfrak{L}^n_{vect}(\nu)\|_{TV}\le 4\epsilon\exp(4\|f_n\|_\infty)+(1+2\|f\|_\infty\exp(4\|f_n\|_\infty))\|\mu-\nu\|_{TV}, \tag{B.7}
$$

for sufficiently large n.

Proof 53 For every $n\ge 1$, let us define a sequence $\epsilon_n:=\|f_n-f\|_\infty$. We shall make use of this sequence in the proof of the lemma as well as in the later part of the proof of Proposition 3. We now proceed with the proof of the lemma. Fix $\mu,\nu\in\mathscr{P}(\mathscr{S})$. Now note that by definition of total-variation norm, and again exploiting the fact that the two-player games $(f_n)_{n\ge1}$ are bounded below, we have

$$
\begin{aligned}
\|\mathfrak{L}^n(\mu)-\mathfrak{L}^n(\nu)\|_{\mathbf{TV}} &= \int_{\mathscr{S}}\left|\frac{\exp(\varphi^n_x(\mu))}{\int_{\mathscr{S}}\exp(\varphi^n_y(\mu))dy}-\frac{\exp(\varphi^n_x(\nu))}{\int_{\mathscr{S}}\exp(\varphi^n_y(\nu))dy}\right|dx \\
&= \int_{\mathscr{S}}\left|\frac{\int_{\mathscr{S}}\exp(\varphi^n_x(\mu))\exp(\varphi^n_y(\nu))dy-\int_{\mathscr{S}}\exp(\varphi^n_x(\nu))\exp(\varphi^n_y(\mu))dy}{(\int_{\mathscr{S}}\exp(\varphi^n_y(\mu))dy)(\int_{\mathscr{S}}\exp(\varphi^n_y(\nu))dy)}\right|dx \\
&\le \exp(2\|f_n\|_\infty)\int_{\mathscr{S}}\Big|\int_{\mathscr{S}}\exp(\varphi^n_x(\mu)+\varphi^n_y(\nu))dy-\int_{\mathscr{S}}\exp(\varphi^n_x(\mu)+\varphi^n_y(\nu))dy \\
&\qquad +\int_{\mathscr{S}}\exp(\varphi^n_x(\mu)+\varphi^n_y(\nu))dy-\int_{\mathscr{S}}\exp(\varphi^n_x(\nu)+\varphi^n_y(\mu))dy\Big| \\
&\le \exp(2\|f_n\|_\infty)\int_{\mathscr{S}}|\exp(\varphi^n_x(\mu))(\exp(\varphi^n_y(\nu))-\exp(\varphi^n_y(\mu)))|dy \\
&\qquad \int_{\mathscr{S}}|\exp(\varphi^n_y(\mu))(\exp(\varphi^n_x(\mu))-\exp(\varphi^n_x(\nu)))|dy \\
&\le 2\exp(2\|f_n\|_\infty)\exp(\|\varphi^n(\mu)\|_\infty+\|\varphi^n(\mu)\|_\infty\vee\|\varphi^n(\nu)\|_\infty)\|\varphi^n(\mu)-\varphi^n(\nu)\|_\infty \\
&\le 2\exp(4\|f_n\|_\infty)\|\varphi^n(\mu)-\varphi^n(\nu)\|_\infty \\
&\le 2\exp(4\|f_n\|_\infty)(\|\varphi^n(\mu)-\varphi(\mu)\|_\infty+\|\varphi(\mu)-\varphi(\nu)\|_\infty+\|\varphi^n(\nu)-\varphi(\nu)\|_\infty) \\
&\le 2\exp(4\|f_n\|_\infty)(2\epsilon_n+\|f\|_\infty\|\mu-\nu\|_{\mathbf{TV}}) \\
&= 4\epsilon_n\exp(4\|f_n\|_\infty)+2\|f\|_\infty\exp(4\|f_n\|_\infty)\|\mu-\nu\|_{\mathbf{TV}}.
\end{aligned}
$$

As a result, we have by definition of $\mathfrak{L}^n_{\text{vect}}$ that

$$\begin{aligned}
\|\mathfrak{L}^n_{\text{vect}}(\mu)-\mathfrak{L}^n_{\text{vect}}(\nu)\|_{\mathbf{TV}} &= \|(\mathfrak{L}^n(\mu)-\mu)-(\mathfrak{L}^n(\nu)-\nu)\|_{\mathbf{TV}} \\
&\leq \|\mathfrak{L}^n(\mu)-\mathfrak{L}^n(\nu)\|_{\mathbf{TV}}+\|\mu-\nu\|_{\mathbf{TV}} \\
&\leq 4\epsilon_n \exp(4\|f_n\|_\infty)+2\|f\|_\infty \exp(4\|f_n\|_\infty)\|\mu-\nu\|_{\mathbf{TV}}+\|\mu-\nu\|_{\mathbf{TV}} \\
&= 4\epsilon_n \exp(4\|f_n\|_\infty)+(1+2\|f\|_\infty \exp(4\|f_n\|_\infty))\|\mu-\nu\|_{\mathbf{TV}} \qquad \text{(B.8)}
\end{aligned}$$

for all $\mu,\nu\in\mathscr{P}(\mathscr{S})$. Finally fix $\epsilon>0$. We first observe that since $f_n\to f$ uniformly, we have $\epsilon_n\leq\epsilon$ for sufficiently large n. Thus, we have that

$$\|\mathfrak{L}^n_{\text{vect}}(\mu)-\mathfrak{L}^n_{\text{vect}}(\nu)\|_{\mathbf{TV}}\leq 4\epsilon \exp(4\|f_n\|_\infty)+(1+2\|f\|_\infty \exp(4\|f_n\|_\infty))\|\mu-\nu\|_{\mathbf{TV}},$$

for sufficiently large n. This concludes the proof of Lemma 15.

We now proceed with the proof of the proposition. Notice that using (B.8), we can thus say that

$$\begin{aligned}
\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\mu^n(t)\|_{\mathbf{TV}} \leq{}& 8\epsilon_n t\exp(4\|f_n\|_\infty)+(1+2\|f\|_\infty \exp(4\|f_n\|_\infty))\delta_N T \\
&+(1+2\|f\|_\infty \exp(4\|f_n\|_\infty))\int_0^t \|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(s))-\mu^n(s)\|_{\mathbf{TV}}\,ds+\Psi(T).
\end{aligned}$$

Since $f_n\to f$ uniformly, there exists $B>0$ such that $\|f_n\|_\infty\leq B$ for all $n\geq 1$. Therefore, by taking supremum on both sides of the above inequality over the time horizon 0 to T, we have

$$\begin{aligned}
\sup_{0\leq t\leq T}\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\mu^n(t)\|_{\mathbf{TV}} \leq{}& \Big[8\epsilon_n T\exp(4B)+(1+2\|f\|_\infty \exp(4B))\delta_N T \\
&+\Psi(T)\Big]+(1+2\|f\|_\infty \exp(4B))\int_0^t \|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(s))-\mu^n(s)\|_{\mathbf{TV}}\,ds.
\end{aligned}$$

Therefore by Grönwall's inequality we have the following bound:

$$\begin{aligned}
\sup_{0\leq t\leq T}\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\mu^n(t)\|_{\mathbf{TV}} \leq{}& \Big[8\epsilon_n T\exp(4B)+(1+2\|f\|_\infty \exp(4B))\delta_N T \\
&+\Psi(T)\Big]\exp((1+2\|f\|_\infty \exp(4B))T). \qquad \text{(B.9)}
\end{aligned}$$

Recall the definition of ϵ_n from the proof of Lemma 15. Since both ϵ_n and δ_N converges to 0 as $n\to\infty$ and $N\to\infty$, we can choose n and N sufficiently large such that

$$\epsilon_n\vee\delta_N\leq\Big(\frac{\epsilon}{32T\exp(4B)}\wedge\frac{\epsilon}{4T(1+2\|f\|_\infty\exp(4\|f\|_\infty))}\Big)\times\exp(-(1+2\|f\|_\infty\exp(4B))T), \qquad \text{(B.10)}$$

in which case we have

$$\Big\{\sup_{0\leq t\leq T}\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\mu^n(t)\|_{\mathbf{TV}}\geq\epsilon\Big\}\implies\Big\{\Psi(T)\geq\frac{\epsilon}{2}\exp(-(1+2\|f\|_\infty\exp(4B))T)\Big\} \qquad \text{(B.11)}$$

As a result, we have

$$\mathbb{P}_{\mu_0^n}\left(\sup_{0\le t\le T}\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))-\mu^n(t)\|_{\mathbf{TV}}\ge\epsilon\right)\le\mathbb{P}_{\mu_0^n}\left(\Psi(T)\ge\frac{\epsilon}{2}\exp(-(1+2\|f\|_\infty\exp(4B))T)\right).$$

We now proceed to bound the probability on the right hand side of the above inequality. Let us define the following collection of functions:

$$\bar{\mathcal{L}}^\infty(\mathcal{S}):=\left\{\bar{g}\in\mathcal{L}^\infty(\mathcal{S}):\exists\,(\beta_k)_{k=0}^{2^n-1}\subseteq\mathbb{R},\text{ such that }\bar{g}(x):=\sum_{k=0}^{2^n-1}\beta_k\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x),\text{ for all }x\in\mathcal{S}\right\}.$$

Thus, in principle, $\bar{\mathcal{L}}^\infty(\mathcal{S})$ is the sub-collection of bounded step functions on the strategy space $\mathcal{S}$.

Lemma 16 *Let $\gamma:=(\sqrt{2}+\sup_{\mu\in\mathcal{P}(\mathcal{S})}\|\mathfrak{L}^n_{\text{vect}}(\mu)\|_{\mathbf{TV}})^2$. For $k\ge1$, let $\mathbb{Z}_k:\bar{\mathcal{L}}^\infty(\mathcal{S})\to\mathbb{R}$ be a mapping defined as*

$$\mathbb{Z}_k(\bar{g}):=\exp\Big(\sum_{i=0}^{k}\langle\bar{g},\delta_N\int^n(\mathcal{U}^N_{n,i})\rangle-\frac{\gamma}{2}k\delta_N^2\|\bar{g}\|_\infty^2\Big),\quad\textit{for all }\bar{g}\in\bar{\mathcal{L}}^\infty(\mathcal{S}).$$

Then, for every $\bar{g}\in\bar{\mathcal{L}}^\infty(S)$, the process $(\mathbb{Z}_k(\bar{g}))_{k\ge1}$ is a supermartingale with respect to the filtration $(\mathcal{F}_k)_{k\ge1}$.

Proof 54 From definition of $(\mathcal{U}^N_{n,k})$, we observe that

$$\begin{aligned}\|\mathfrak{s}_n(\mathcal{U}^N_{n,k})\|_{\mathbf{TV}}&\le\frac{1}{\delta_N}\|\mathbb{X}_n^N((k+1)\delta_N)-\mathbb{X}_n^N(k\delta_N)\|_{\mathbf{TV}}+\|\mathfrak{L}^n_{\text{vect}}(\mathbb{X}_n^N(k\delta_N))\|_{\mathbf{TV}}\\&\le\sqrt{2}+\sup_{\mu\in\mathcal{P}(\mathcal{S})}\|\mathfrak{L}^n_{\text{vect}}(\mu)\|_{\mathbf{TV}}=\sqrt{\gamma}.\end{aligned}$$

Now it follows along the lines of Benaïm and Weibull (2003), that $\mathbb{E}(\exp(\langle\bar{g},\mathcal{U}^N_{n,k}\rangle))\le\exp(\frac{\gamma\|\bar{g}\|_\infty^2}{2})$, and thus completes the proof of Lemma 16.

Thus for any $\theta>0$, using Lemma 16, we have that

$$\begin{aligned}\mathbb{P}_{\mu_0^n}\Big(\max_{0\le k\le m}\Big\langle\bar{g},\sum_{i=0}^{k-1}\delta_N\mathfrak{s}_n(\mathcal{U}^N_{n,i})\Big\rangle\ge\theta\Big)&\le\mathbb{P}_{\mu_0^n}\Big(\max_{0\le k\le m}\mathbb{Z}_k(\bar{g})\ge\exp(\theta-\frac{\gamma}{2}\|\bar{g}\|_\infty^2m\delta_N^2)\Big)\\&\le\exp\Big(\frac{\gamma}{2}\|\bar{g}\|_\infty^2m\delta_N^2-\theta\Big)\end{aligned}$$

For $k\ge1$, let $h_k:S\to\mathbb{R}$ be defined as $h_k(x):=\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)$, for $x\in S$. It then follows that $(h_k)_{0\le k\le 2^n-1}$ forms a basis of $\bar{\mathcal{L}}^\infty(S)$. Let $h=h_k$ or $h=-h_k$ for some k. Set $\theta=\frac{\epsilon^2}{\gamma m\delta^2}$. Then there exists $\mathbf{h}=(h(\alpha_{n,0}),\ldots,h(\alpha_{n,2^n-1}))\in\mathbb{R}^{2^n}$ and $\mathbf{U}_n^N\in\mathbb{R}^{2^n}$ such that $\Big\langle h,\sum_{i=0}^{k-1}\delta_N\mathfrak{s}_n(\mathcal{U}^N_{n,i})\Big\rangle=\Big\langle\mathbf{h},\sum_{i=0}^{k-1}\delta_N\mathbf{U}^N_{n,i}\Big\rangle$. Now set $g=$

$\left(\frac{\theta}{\epsilon}\right)h$ and define $\mathbf{g} := \left(\frac{\theta}{\epsilon}\right)\mathbf{h}$. It then follows from Benaïm and Weibull (2003) that

$$\begin{aligned}\mathbb{P}_{\mu_0^n}\Big(\max_{0\le k\le m}\Big\langle h,\sum_{i=0}^{k-1}\delta_N\mathfrak{s}_n(U_{n,i}^N)\Big\rangle\ge\epsilon\Big) &= \mathbb{P}_{\mu_0^n}\Big(\max_{0\le k\le m}\Big\langle \mathbf{h},\sum_{i=0}^{k-1}\delta_N\mathbf{U}_{n,i}^N\Big\rangle\ge\epsilon\Big)\\ &= \mathbb{P}_{\mu_0^n}\Big(\max_{0\le k\le m}\Big\langle \mathbf{g},\sum_{i=0}^{k-1}\delta_N\mathbf{U}_{n,i}^N\Big\rangle\ge\theta\Big)\\ &\le \exp\Big(\frac{-\epsilon^2}{2m\gamma\delta_N^2}\Big).\end{aligned}$$

It now follows from Lemma 14 that

$$\mathbb{P}_{\mu_0^n}(\Psi(T)\ge\epsilon)\le 2(2^n+1)\exp\Big(\frac{-\epsilon^2}{8\delta_N\gamma T}\Big).$$

Define $C := \exp(-(1+2\|f\|_\infty\exp(4B))T) > 0$. Note that by definition of C, it is independent of μ_0^n. As a result, we then arrive the inequality:

$$\begin{aligned}\mathbb{P}_{\mu_0^n}\Big(\sup_{0\le t\le T}\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))-\mu^n(t)\|_{\mathbf{TV}}\ge\epsilon\Big) &\le \mathbb{P}_{\mu_0^n}\Big(\Psi(T)\ge\frac{C\epsilon}{2}\Big)\\ &\le 2(2^n+1)\exp\left(\frac{-C^2\epsilon^2}{32\delta_N\gamma T}\right).\end{aligned}$$

Let $\bar{C} := C^2/32\gamma T$. As $\delta_N = 1/N$, we have that

$$\mathbb{P}_{\mu_0^n}\Big(\sup_{0\le t\le T}\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))-\mu^n(t)\|_{\mathbf{TV}}\ge\epsilon\Big)\le 2(2^n+1)\exp(-N\bar{C}\epsilon^2).$$

This concludes the proof of Proposition 3.

B.3 Proof of Theorem 6

The proof Theorem 6 follows from Propositions 2 and 3. To make the proof mathematically rigorous, consider the fixed probability space $(\Omega,\mathscr{F},\mathbb{P})$ on which the stochastic evolutionary process $(\widehat{\mathbb{X}}_n^N(t))_{t\ge0}$ is defined. For $n\ge1$, let $\widehat{\mathbb{X}}_n:\Omega\to\mathscr{P}(\mathscr{S})$ be the degenerate stochastic process such that for al $\omega\in\Omega$, we have $\widehat{\mathbb{X}}_n(\omega)(t) := \mu^n(t)$, for all $t\ge0$. Also define the process $\widehat{\mathbb{X}}:\Omega\to\mathscr{P}(\mathscr{S})$ such that for all $\omega\in\Omega$, $\widehat{\mathbb{X}}(\omega)(t) := \mu(t)$, for all $t\ge0$. For $\epsilon>0$, we have that

$$\begin{aligned}\Big\{\sup_{0\le t\le T}\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))-\widehat{\mathbb{X}}(t)\|_{\mathbf{TV}}\ge\epsilon\Big\} &\subseteq \Big\{\sup_{0\le t\le T}\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))-\widehat{\mathbb{X}}_n(t)\|_{\mathbf{TV}}\ge\epsilon/2\Big\}\\ &\cup\Big\{\sup_{0\le t\le T}\|\widehat{\mathbb{X}}_n(t)-\widehat{\mathbb{X}}(t)\|_{\mathbf{TV}}\ge\epsilon/2\Big\}.\end{aligned}$$

As a result we have that

$$\mathbb{P}_{\mu^N_{\mathbf{p}^n}}\left(\sup_{0\le t\le T}\left\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\widehat{\mathbb{X}}(t)\right\|_{\mathbf{TV}}\ge\epsilon\right)\le\mathbb{P}_{\mu^N_{\mathbf{p}^n}}\left(\sup_{0\le t\le T}\left\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\widehat{\mathbb{X}}_n(t)\right\|_{\mathbf{TV}}\ge\epsilon/2\right)$$
$$+\mathbb{P}_{\mu^N_{\mathbf{p}^n}}\left(\sup_{0\le t\le T}\left\|\widehat{\mathbb{X}}_n(t)-\widehat{\mathbb{X}}(t)\right\|_{\mathbf{TV}}\ge\epsilon/2\right).$$

By assumption of the theorem, we have that the initial conditions $\mu^N_{\mathbf{p}^n}=\mu_0$, for all $N,n\ge 1$, implies that

$$\mathbb{P}_{\mu_0}\left(\sup_{0\le t\le T}\left\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\widehat{\mathbb{X}}(t)\right\|_{\mathbf{TV}}\ge\epsilon\right)\le\mathbb{P}_{\mu_0}\left(\sup_{0\le t\le T}\left\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\widehat{\mathbb{X}}_n(t)\right\|_{\mathbf{TV}}\ge\epsilon/2\right)$$
$$+\mathbb{P}_{\mu_0}\left(\sup_{0\le t\le T}\left\|\widehat{\mathbb{X}}_n(t)-\widehat{\mathbb{X}}(t)\right\|_{\mathbf{TV}}\ge\epsilon/2\right).$$

Fix $n\ge 1$. Then by taking limits as $N\to\infty$ we have by Proposition 3 that

$$\lim_{N\to\infty}\mathbb{P}_{\mu_0}\left(\sup_{0\le t\le T}\left\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\widehat{\mathbb{X}}(t)\right\|_{\mathbf{TV}}\ge\epsilon\right)\le\mathbb{P}_{\mu_0}\left(\sup_{0\le t\le T}\left\|\widehat{\mathbb{X}}_n(t)-\widehat{\mathbb{X}}(t)\right\|_{\mathbf{TV}}\ge\epsilon/2\right).\tag{B.12}$$

Now, allowing $n\to\infty$, in (B.12) we have by an application of Proposition 2 that

$$\lim_{n\to\infty}\lim_{N\to\infty}\mathbb{P}_{\mu_0}\left(\sup_{0\le t\le T}\left\|\mathfrak{s}_n(\widehat{\mathbb{X}}^N_n(t))-\widehat{\mathbb{X}}(t)\right\|_{\mathbf{TV}}\ge\epsilon\right)=0.$$

Finally, using the fact that $\widehat{\mathbb{X}}(t)=\mu(t)$, for all $t\ge 0$, $\mathbb{P}_{\mu_0}$-a.s concludes the proof of Theorem 6.

B.4 Proof of Proposition 4

Only-if part: Suppose that φ is a supermodular game as per Definition 13. We need to show that the underlying two-player game f is supermodular. That is, we need to show that $f(y,z)-f(x,z)$ is increasing in z if $x<y$. Let $z_1\le z_2$. Define $\mu_1:=\delta_{z_1}$ and $\mu_2:=\delta_{z_2}$. It then follows that $\mu_1\preceq\mu_2$. Since φ is a supermodular game, we have that

$$\begin{aligned}
f(y,z_1)-f(x,z_1)&=\int_{\mathscr{S}}f(y,z)\mu_1(dz)-\int_{\mathscr{S}}f(x,z)\mu_1(dx)\\
&=\varphi_y(\mu_1)-\varphi_x(\mu_1)\\
&\le\varphi_y(\mu_2)-\varphi_x(\mu_2)\\
&=\int_{\mathscr{S}}f(y,z)\mu_2(dz)-\int_{\mathscr{S}}f(x,z)\mu_2(dx)\\
&=f(y,z_2)-f(x,z_2).
\end{aligned}$$

This proves that $f(y,z)-f(x,z)$ is non–decreasing in z if $x<y$. Thus f is supermodular on $\mathscr{S}$ by Definition 14.

If part: We now show that if f is supermodular in $\mathscr{S}$ then the population game φ is a supermodular game. To this end, fix $x < y$. Since f is supermodular on $\mathscr{S}$, we have that $f(y,z)-f(x,z)$ is non–decreasing in z by Definition 14. Now fix $\mu \preceq \nu$. By (3.1), we have

$$\begin{aligned}\varphi_y(\mu)-\varphi_x(\mu) &= \int_{\mathscr{S}} (f(y,z)-f(x,z))\mu(dz) \\ &\leq \int_{\mathscr{S}} (f(y,z)-f(x,z))\nu(dz) \\ &= \varphi_y(\nu)-\varphi_x(\nu).\end{aligned}$$

This proves that φ is supermodular according to Definition 13. This concludes the proof of Proposition 4.

B.5 Proof of Proposition 5

For this proof, we topologize the space of probability measures $\mathscr{P}(\mathscr{S})$ with the $\mathbf{BL}^*$-norm. In order to define such a norm, we first require the notion of a $\mathbf{BL}$-norm. The bounded Lipschitz norm is defined as

$$\|g\|_{\mathbf{BL}} := \sup_{x\in\mathscr{S}} |g(x)| + \sup_{x\neq y} \frac{|g(x)-g(y)|}{|x-y|}, \quad \text{for all } g \in \mathsf{Lip}(\mathscr{S}), \tag{B.13}$$

where $\mathsf{Lip}(\mathscr{S})$ denotes the collection of all Lipschitz functions on $\mathscr{S}$. Now let

$$\mathsf{Bounded\text{-}Lip}(\mathscr{S}) := \{g : \mathscr{S} \to \mathbb{R} \mid g \text{ is bounded and Lipschitz with } \|g\|_{\mathbf{BL}} \leq 1\}$$

be the collection of bounded Lipschitz continuous functions with $\mathbf{BL}$-norm bounded by 1. Denote by $\mathscr{M}(\mathscr{S})$, the space of all signed measures on $\mathscr{S}$. Then the dual $\mathbf{BL}^*$-norm on $\mathscr{M}(\mathscr{S})$ is defined as

$$\|\mu-\nu\|_{\mathbf{BL}^*} := \sup_{g\in\mathsf{Bounded\text{-}Lip}(\mathscr{S})} \left(\int_{\mathscr{S}} f\,d\mu - \int_{\mathscr{S}} f\,d\nu \right), \quad \text{for all } \mu,\nu \in \mathscr{M}(\mathscr{S}). \tag{B.14}$$

It is well-known (see Billingsley (2013), Perkins and Leslie (2014)) that the weak topology on $\mathscr{P}(\mathscr{S})$ coincides with the topology generated by the $\mathbf{BL}^*$-norm on $\mathscr{P}(\mathscr{S})$. Therefore, weak convergence is equivalent to convergence with respect to the $\mathbf{BL}^*$-norm.

We now proceed with the proof of the proposition. Suppose that φ is supermodular and let $n \geq 1$. We show that the level-n discretization $\mathscr{D}_{\varphi}^{n}$ of φ as defined in Definition 11 is a supermodular game. Fix $\mathbf{x},\mathbf{y} \in \Delta(S_n)$ with $\mathbf{x} \succeq \mathbf{y}$. It then follows that $\mu_{\mathbf{x}} := \sum_{0\leq j\leq 2^n} \mathbf{x}_j \delta_{\alpha_{n,j}} \succeq \sum_{0\leq j\leq 2^n} \mathbf{y}_j \delta_{\alpha_{n,j}} =: \mu_{\mathbf{y}}$. Now fix $i < j$. It then follows from the definition of $\mathscr{D}^n_\varphi$ that,

$$\begin{aligned}
\mathscr{D}^n_\varphi[\mathbf{x}](\alpha_{n,j})-\mathscr{D}^n_\varphi[\mathbf{x}](\alpha_{n,i}) &= \langle f_n(\alpha_{n,j},\cdot),\mathbf{x}\rangle-\langle f_n(\alpha_{n,i},\cdot),\mathbf{x}\rangle\\
&= \sum_{0\le k\le 2^n} f_n(\alpha_{n,j},\alpha_{n,k})\mathbf{x}_k-\sum_{0\le k\le 2^n} f_n(\alpha_{n,i},\alpha_{n,k})\mathbf{x}_k\\
&= \int_S f_n(\alpha_{n,j},t)\mu_{\mathbf{x}}(dt)-\int_S f_n(\alpha_{n,i},t)\mu_{\mathbf{x}}(dt)\\
&= \varphi_{\alpha_{n,j}}(\mu_{\mathbf{x}})-\varphi_{\alpha_{n,i}}(\mu_{\mathbf{x}})\\
&\ge \varphi_{\alpha_{n,j}}(\mu_{\mathbf{y}})-\varphi_{\alpha_{n,i}}(\mu_{\mathbf{y}})\\
&= \int_S f_n(\alpha_{n,j},t)\mu_{\mathbf{y}}(dt)-\int_S f_n(\alpha_{n,i},t)\mu_{\mathbf{y}}(dt)\\
&= \sum_{0\le k\le 2^n} f_n(\alpha_{n,j},\alpha_{n,k})\mathbf{y}_k-\sum_{0\le k\le 2^n} f_n(\alpha_{n,i},\alpha_{n,k})\mathbf{y}_k\\
&= \langle f_n(\alpha_{n,j},\cdot),\mathbf{y}\rangle-\langle f_n(\alpha_{n,i},\cdot),\mathbf{y}\rangle\\
&= \mathscr{D}^n_\varphi[\mathbf{y}](\alpha_{n,j})-\mathscr{D}^n_\varphi[\mathbf{y}](\alpha_{n,i}).
\end{aligned}$$

This completes the proof that $\mathscr{D}^n_\varphi$ is a supermodular game. Since $n\ge 1$ is arbitrary, this concludes the proof of “only-if part" of Proposition 5.

We now proceed to prove the “if-part" of the theorem. Now suppose that $\mathscr{D}^n_\varphi$ is a supermodular game for every $n\ge 1$. We need to show that φ is a supermodular game. To this end, first $x>y$. Let $S_\infty := \cup_{n\ge 1} S_n$. Note that S_∞ consists of all dyadic rationals in $\mathscr{S}$ and hence is dense in $\mathscr{S}$. Thus, there exists two subsequences $\{n_k\}_{k\ge 1},\{m_k\}_{k\ge 1}\subseteq\mathbb{N}$ and sequences $\{\alpha^x_{n_k,j_{n_k}}\}\subseteq S_{n_k}$ and $\{\alpha^y_{m_k,j_{m_k}}\}\subseteq S_{m_k}$ such that $\alpha^x_{n_k,j_{n_k}}\downarrow x$ and $\alpha^y_{m_k,j_{m_k}}\uparrow y$ as $k\to\infty$. Again, since S_∞ is dense in $\mathscr{S}$, it follows (see Billingsley (2013)) that that space of atomic probability measures $\mathscr{P}(S_\infty)$ on S_∞ is dense in $\mathscr{P}(\mathscr{S})$ relative to the topology of weak convergence. Since weak topology in equivalent to the $\mathbf{BL}^*$-norm on $\mathscr{P}(\mathscr{S})$, we therefore have that for every $\mu\in\mathscr{P}(\mathscr{S})$, there exists $(\mu_n)_{n\ge 1}\subseteq\mathscr{P}(S_\infty)$ such that $\mu_n\xrightarrow{\mathbf{BL}^*}\mu$. By definition since φ is a weakly continuous population game, we have $\varphi_x(\mu_n)$ converges to $\varphi_x(\mu)$ for all $x\in\mathscr{S}$. As a result we have for all $x\in\mathscr{S}$ that,

$$\begin{aligned}
\varphi_x(\mu) &= \lim_{n\to\infty}\varphi_x(\mu_n)\\
&= \lim_{n\to\infty}\int_S f(x,t)\mu_n(dt).
\end{aligned}$$

By the definition in (4.16), the sequence $(f_n)_{n\ge 1}$ converges uniformly to f. This, in conjuction with Bounded Convergence Theorem (Billingsley (2013)) implies that

$$\begin{aligned}
\varphi_x(\mu) &= \lim_{n\to\infty}\int_S f(x,t)\mu_n(dt).\\
&= \lim_{n\to\infty}\lim_{k\to\infty}\int_S f_k(\alpha^x_{n_k,j_{n_k}},t)\mu_n(dt)\\
&= \lim_{n\to\infty}\lim_{k\to\infty}\mathscr{D}^n_\varphi[\mathbf{p}_{\mu_n}](\alpha^x_{n_k,j_{n_k}}).
\end{aligned}$$

We are now ready to conclude the proof of the theorem. Fix $\mu, \nu \in \mathscr{P}(\mathscr{S})$ with $\mu \succeq \nu$. Then there exists two sequences of atomic probability measures $(\mu_n)_{n\geq 1} \subseteq \mathscr{P}(S_\infty)$ and $(\nu_n)_{n\geq 1} \subseteq \mathscr{P}(S_\infty)$ such that $\mu_n \succeq \nu_n$ for all $n \geq 1$, such that $\mu_n \xrightarrow{\mathbf{BL}^*} \mu$ and $\nu_n \xrightarrow{\mathbf{BL}^*} \nu$. By hypothesis of the theorem, the level-n discretization $\mathscr{D}_\varphi^n$ is a supermodular game. Also by construction, we have that $\alpha^x_{n_k, j_{n_k}} > \alpha^y_{m_k, j_{m_k}}$ for every $k \geq 1$. Therefore we have

$$\begin{aligned}
\varphi_x(\mu) - \varphi_y(\mu) &= \lim_{n\to\infty} \lim_{k\to\infty} \mathscr{D}_\varphi^n[\mathbf{p}_{\mu_n}](\alpha^x_{n_k, j_{n_k}}) - \lim_{n\to\infty} \lim_{k\to\infty} \mathscr{D}_\varphi^n[\mathbf{p}_{\mu_n}](\alpha^y_{m_k, j_{m_k}}) \\
&= \lim_{n\to\infty} \lim_{k\to\infty} \Big[\mathscr{D}_\varphi^n[\mathbf{p}_{\mu_n}](\alpha^x_{n_k, j_{n_k}}) - \mathscr{D}_\varphi^n[\mathbf{p}_{\mu_n}](\alpha^y_{m_k, j_{m_k}})\Big] \\
&\geq \lim_{n\to\infty} \lim_{k\to\infty} \Big[\mathscr{D}_\varphi^n[\mathbf{p}_{\nu_n}](\alpha^x_{n_k, j_{n_k}}) - \mathscr{D}_\varphi^n[\mathbf{p}_{\nu_n}](\alpha^y_{m_k, j_{m_k}})\Big] \\
&= \lim_{n\to\infty} \lim_{k\to\infty} \mathscr{D}_\varphi^n[\mathbf{p}_{\nu_n}](\alpha^x_{n_k, j_{n_k}}) - \lim_{n\to\infty} \lim_{k\to\infty} \mathscr{D}_\varphi^n[\mathbf{p}_{\nu_n}](\alpha^y_{m_k, j_{m_k}}) \\
&= \varphi_x(\nu) - \varphi_y(\nu).
\end{aligned}$$

Since $\mu \succeq \nu$ is arbitrary, this proves that φ is a supermodular game and hence concludes the proof of Proposition 5.

B.6 Proof of Proposition 6

Recall that the two-player game f which induces φ in Example 5 is defined as $f(x,y) = m(x)y - c(x)$, for all $x, y \in \mathscr{S}$, where m is an increasing function of x and c is an arbitrary function of x. Using (4.16), we define the step–wise approximation of f as

$$f_n(x,y) = \begin{cases} m(\alpha_{n,j})\alpha_{n,k} - c(\alpha_{n,j}), \text{ if } (x,y) \in [\alpha_{n,j}, \alpha_{n,j+1}) \times [\alpha_{n,k}, \alpha_{n.k+1}) \text{ and } j,k = 0,\ldots,2^n-2 \\ m(\alpha_{n,2^n-1})\alpha_{n,2^n-1} - c(\alpha_{n,2^n-1}) \text{ if } (x,y) \in [\alpha_{n,2^n-1}, 1] \times [\alpha_{n,2^n-1}, 1]. \end{cases} \tag{B.15}$$

This implies that for every $n \geq 1$, we have from (3.13),

$$\mathscr{D}_\varphi^n[\mathbf{p}^n](\alpha_{n,j}) = \sum_{0\leq k\leq 2^n-1} f_n(\alpha_{n,j}, \alpha_{n,k})\mathbf{p}^n_k, \quad \text{for all } j = 0,1,\ldots,2^n-1.$$

We now proceed to show that $\mathscr{D}_\varphi^n$ is irreducible. Fix $\mathbf{p}^n \in \Delta(S_n)$ and $K \subset \widehat{S}_n$ arbitrary as in Definition 16. Now let $j \in K$ and $i \in \widehat{S}_n \setminus K$. Since m is an increasing function of x, we have

$$\begin{aligned}
\frac{\partial \mathscr{D}_\varphi^n[\mathbf{p}^n](\alpha_{n,j+1})}{\partial \mathbf{p}^n_{i+1}} - \frac{\partial \mathscr{D}_\varphi^n[\mathbf{p}^n](\alpha_{n,j+1})}{\partial \mathbf{p}^n_{i}} &= f_n(\alpha_{n,j+1}, \alpha_{n,i+1}) - f_n(\alpha_{n,j+1}, \alpha_{n,i}) \\
&= m(\alpha_{n,j+1})(\alpha_{n,i+1} - \alpha_{n,i}) \\
&> m(\alpha_{n,j})(\alpha_{n,i+1} - \alpha_{n,i}) \\
&= f_n(\alpha_{n,j}, \alpha_{n,i+1}) - f_n(\alpha_{n,j}, \alpha_{n,i}) \\
&= \frac{\partial \mathscr{D}_\varphi^n[\mathbf{p}^n](\alpha_{n,j})}{\partial \mathbf{p}^n_{i+1}} - \frac{\partial \mathscr{D}_\varphi^n[\mathbf{p}^n](\alpha_{n,j})}{\partial \mathbf{p}^n_{i}}
\end{aligned}$$

This proves that $\mathscr{D}^n_\varphi$ is irreducible and hence concludes the proof of Proposition 6.

B.7 Proof of Lemma 4

To prove the lemma we shall first exploit the properties of $\mathbf{BL}^*$-norm, or equivalently weak topology on the space $\mathscr{P}(\mathscr{S})$. Observe that the sequence of level-n step logit equilibria $\int^{n}(\mu_{\mathbf{p}^*_n}) \in \mathscr{P}(\mathscr{S})$ for every $n \geq 1$. Since the weak topology, and hence the $\mathbf{BL}^*$-norm renders the space $\mathscr{P}(\mathscr{S})$ compact, there exists a subsequence $(n_k)_{k\geq 1}$ such that $\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})$ converges in the $\mathbf{BL}^*$-norm to some $\mu^* \in \mathscr{P}(\mathscr{S})$. We now show that μ^* is a logit equilibrium of the original game φ. To this end, we need to show that

$$\mu^*(A) = \mathscr{L}_\eta[\mu^*](A), \quad \text{for all } A \in \mathscr{B}(\mathscr{S}).$$

Note that since $\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})$ is a level-n_k step logit equilibrium of the game φ^{n_k} for every $k \geq 1$, we have that

$$\int^{n_k}[\mu_{\mathbf{p}^*_{n_k}}](A) = \mathscr{L}^{n_k}_\eta[\int^{n_k}(\mu_{\mathbf{p}^*_{n_k}})](A), \quad \text{for all } A \in \mathscr{B}(\mathscr{S}) \text{ such that } \mu^*(\partial A) = 0. \tag{B.16}$$

Now, we show that $\varphi^{n_k}(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})) \to \varphi(\mu^*)$ as $k \to \infty$. Fix $x \in \mathscr{S}$. Since f is Lipschitz continuous, we have

$$\begin{aligned}
|\varphi^{n_k}_x(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})) - \varphi_x(\mu^*)| &= \left| \int_{\mathscr{S}} f_{n_k}(x,y)\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})(dy) - \int_{\mathscr{S}} f(x,y)\mu^*(dy) \right| \\
&\leq \int_{\mathscr{S}} |f_{n_k}(x,y) - f(x,y)|\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})(dy) + \left| \int_{\mathscr{S}} f(x,y)(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}}) - \mu^*)(dy) \right| \\
&\leq 2\|f_{n_k} - f\|_\infty + \|f\|_\infty \|\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}}) - \mu^*\|_{\mathbf{BL}^*}.
\end{aligned}$$

As a result, using the fact that $f_n \to f$ uniformly, and that $\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})$ converges in the $\mathbf{BL}^*$-norm to μ^*, we have

$$\begin{aligned}
\|\varphi^{n_k}(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})) - \varphi(\mu^*)\|_\infty &= \sup_{x\in\mathscr{S}} |\varphi^{n_k}_x(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})) - \varphi_x(\mu^*)| \\
&\leq 2\|f_{n_k} - f\|_\infty + \|f\|_\infty \|\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}}) - \mu^*\|_{\mathbf{BL}^*} \\
&\to 0, \quad \text{as } k \to \infty.
\end{aligned}$$

Also since

$$\exp(\varphi^{n_k}_x(\int^{n_k}(\mu_{\mathbf{p}^*_{n_k}}))) \leq \exp(2(1 + \|f\|_\infty)), \quad \text{for all } x \in \mathscr{S},$$

we have using Dominated Convergence Theorem that $\int_{\mathscr{S}} \exp(\varphi^{n_k}_y(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})))dy \to \int_{\mathscr{S}} \exp(\varphi_y(\mu_*))dy$. Now, in order to prove that $\mathfrak{L}^{n_k}_\eta(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}}))$ converges in the variational norm to $\mathfrak{L}_\eta(\mu_*)$, it is enough to show that

$$\frac{d\mathscr{L}^{n_k}_\eta(\int^{n_k}(\mu_{\mathbf{p}^*_{n_k}}))}{d\lambda} \to \frac{d\mathscr{L}_\eta(\mu_*)}{d\lambda} \quad \text{in } \mathscr{L}^1(\mathscr{S}) \text{ as } k \to \infty.[1]$$

[1] The space $\mathscr{L}^1(\mathscr{S})$ is defined as the collection of all measurable functions $g : \mathscr{S} \to \mathbb{R}$ such that $\int_{\mathscr{S}} |f(x)|dx < \infty$. We say that $g_n \to g$ in $\mathscr{L}^1(\mathscr{S})$ if $\int_{\mathscr{S}} |g_n(x) - g(x)|dx \to 0$ as $n \to \infty$.

Note that the density function of $\mathfrak{L}_\eta^{n_k}(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}}))$ satisfies

$$\frac{d\mathscr{L}_\eta^{n_k}(\int^{n_k}(\mu_{\mathbf{p}^*_{n_k}}))}{d\lambda}(x) \leq \exp(2(1+\|f\|_\infty)), \quad \text{for all } x\in\mathscr{S}.$$

We also have that for all $x\in\mathscr{S}$,

$$\frac{d\mathfrak{L}_\eta^{n_k}(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}}))}{d\lambda}(x) \to \frac{d\mathfrak{L}_\eta(\mu_*)}{d\lambda}(x) \quad \text{as } k\to\infty.$$

Again, by an application of Dominated Convergence Theorem, we conclude that $\mathfrak{L}_\eta^{n_k}(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}}))$ converges in the variational norm to $\mathfrak{L}_\eta(\mu_*)$. Therefore, for every $A\in\mathscr{B}(\mathscr{S})$, we have that $\mathfrak{L}_\eta^{n_k}[\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})](A)\to\mathfrak{L}_\eta[\mu^*](A)$ as $k\to\infty$. Using this observation in (B.16), we have that μ^* is a fixed point of $\mathfrak{L}_\eta$ and hence a logit equilibrium of φ. In fact, we proved that that $\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}})=\mathfrak{L}_\eta^{n_k}(\mathfrak{s}_{n_k}(\mu_{\mathbf{p}^*_{n_k}}))\to\mathfrak{L}_\eta(\mu^*)=\mu^*$ as $k\to\infty$ with respect to the total variation norm. This concludes the proof of Lemma 4.

B.8 Proof of Theorem 7

The proof of Theorem 7 is subdivided into several steps.

Step 1: In this step, we show that the discretized version of the stochastic evolutionary process converges in the medium run to an element of $\mathrm{LE}_\eta(\mathscr{D}_\varphi^n)$ as $N\to\infty$. Fix $\epsilon>0$. Pick $N,n\geq 1$ and consider the interpolated stochastic evolutionary process $(\widehat{\mathbb{X}}_n^N(t))_{t\geq 0}$. We observe that though $(\widehat{\mathbb{X}}_n^N(t))_{t\geq 0}$ is a stochastic process in taking values in $\mathscr{P}(\mathscr{S})$, however in principle, $(\widehat{\mathbb{X}}_n^N(t))_{t\geq 0}$ takes values in the space of atomic measures on $\mathscr{S}$ with support S_n. Note that any probability measure $\mu\in\mathscr{P}(S_n)$ can be expressed as $\mu=\sum_{0\leq k\leq 2^n-1}\mathbf{p}_j^n\delta_{\alpha_{n,j}}$, for some $\mathbf{p}^n\in\Delta(S_n)$ and vice-versa. We therefore define the bijective mapping $\mathfrak{d}_n:\mathscr{P}(S_n)\to\Delta(S_n)$ such that $\mathfrak{d}_n(\sum_{0\leq k\leq 2^n-1}\mathbf{p}_j^n\delta_{\alpha_{n,j}})=\mathbf{p}^n$ for all $\mathbf{p}^n\in\Delta(S_n)$. It then follows that $(\mathfrak{d}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\geq 0}$ is the interpolated stochastic evolutionary process on $\Delta(S_n)$. Let $\mathrm{LE}_\eta(\mathscr{D}_\varphi^n)$ denote the collection of discretized logit equilibria corresponding to the level-n discretized game $\mathscr{D}_\varphi^n$. More precisely, $\mathrm{LE}_\eta(\mathscr{D}_\varphi^n):=\{\mathbf{p}_\eta^{n,*}: L_\eta^n[\mathbf{p}_\eta^{n,*}]=\mathbf{p}_\eta^{n,*}\}$. By assumption of the theorem, we have that the original population game φ is supermodular. Therefore by Proposition 5, we have that the level-n discretization $\mathscr{D}_\varphi^n$ is supermodular. Also by hypothesis of the theorem, we have that $\mathscr{D}_\varphi^n$ is irreducible. Thus, by Hofbauer and Sandholm (2007), the process $(\mathfrak{d}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\geq 0}\subseteq\Delta(S_n)$ converges in the medium run to an element of $\mathrm{LE}_\eta(\mathscr{D}_\varphi^n)$ as $N\to\infty$. Let us denote by $\mathbf{p}_n^*$ the element in $\mathrm{LE}_\eta(\mathscr{D}_\varphi^n)$ to which $(\mathfrak{d}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\geq 0}$ converges in the medium run as $N\to\infty$.[2] That is, there exists $T_\epsilon>0$ such that for all $T\geq T_\epsilon$, such that

$$\lim_{N\to\infty}\mathbb{P}_{\mu_0}\Big(\mathfrak{d}_n(\widehat{\mathbb{X}}_n^N(t))\in\mathscr{O}^\epsilon(\mathbf{p}_n^*) \text{ for all } T_\epsilon\leq t\leq T\Big)=1.$$

We now proceed to show that the smoothed version of the stochastic evolutionary process $(\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\geq 0}$ converges in the medium run to an element of $\mathrm{LE}_\eta(\varphi^n)$ as $N\to\infty$. In particular, we show that $(\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\geq 0}$ converges in the medium run to $\mathfrak{s}_n(\mu_{\mathbf{p}_n^*})$. The fact that the processes $(\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\geq 0}$ and $(\mathfrak{d}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\geq 0}$ are in one-to-one correspondence with each other implies that $\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))_{t\geq 0}$ converges in the medium run to an element of $\int^n(\mathrm{LE}_\eta(\mathscr{D}_\varphi^n))$.

[2] To keep the notations simple, we remove the dependence of η in $\mathbf{p}_n^*$.

Step 2: As a next step in the proof, we proceed to show that $\int^n(\mathrm{LE}_\eta(\mathscr{D}_\varphi^n))$ coincides with $\mathrm{LE}_\eta(\varphi^n)$, thereby proving that $\mathfrak{s}_n(\mu_{\mathbf{p}_n^*})\in\mathrm{LE}_\eta(\varphi^n)$. We prove this in the following lemma.

Lemma 17 *For $n\geq 1$, let $LE_\eta(\mathscr{D}_\varphi^n)$ and $LE_\eta(\varphi^n)$ be the collection of level-n discretized and level-n step logit equilibria corresponding to $\mathscr{D}_\varphi^n$ and φ^n respectively. Then $\boldsymbol{p}_\eta^{n,*}\in LE_\eta(\mathscr{D}_\varphi^n)$ if and only if $\mathfrak{s}_n(\mu_{\boldsymbol{p}_\eta^{n,*}})\in LE_\eta(\varphi^n)$.*

Proof 55 Again, as before, we shall use the notation $\mathbf{p}^*$ to denote a discretized level-n logit equilibrium to simplify the presentation of the proof. First observe that for any $\mathbf{p}^n\in\Delta(S_n)$, we have by definition of $\mathfrak{s}_n$ that

$$\mu_{\mathbf{p}^n}[\alpha_{n,k},\alpha_{n,k+1})=\mathfrak{s}_n(\mu_{\mathbf{p}^n})[\alpha_{n,k},\alpha_{n,k+1})=\mathbf{p}_k^n,\quad \text{for all } k=0,1,\ldots,2^n-1.$$

Suppose that $\mathbf{p}^*\in\mathrm{LE}_\eta(\mathscr{D}_\varphi^n)$. Then we have $\mathbf{p}^*=L_\eta^n[\mathbf{p}^*]$. This, in particular implies that

$$\mu_{\mathbf{p}^*}=\sum_{0\leq k\leq 2^n-1}L_{\eta,k}^n[\mathbf{p}^*]\delta_{\alpha_{n,k}}.$$

Thus by definition of $\mathfrak{s}_n$ (Definition 12), above equality implies that for all $x\in\mathscr{S}$,

$$\begin{aligned}
\frac{\mathfrak{s}_n(\mu_{\mathbf{p}^*})}{d\lambda}(x)&=\sum_{0\leq k\leq 2^n-1}2^nL_{\eta,k}^n[\mathbf{p}^*]\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)\\
&=\sum_{0\leq k\leq 2^n-1}2^n\frac{\exp(\eta^{-1}\langle f_n(\alpha_{n,k},\cdot),\mathbf{p}^*\rangle)}{\sum_{j=0}^{2^n-1}\exp(\eta^{-1}\langle f_n(\alpha_{n,j},\cdot),\mathbf{p}^*\rangle)}\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)\\
&=\sum_{0\leq k\leq 2^n-1}\frac{\exp(\eta^{-1}\sum_{0\leq j\leq 2^n-1}f_n(\alpha_{n,k},\alpha_{n,j})\mathbf{p}_j^*)}{2^{-n}\sum_{j=0}^{2^n-1}\exp(\eta^{-1}\sum_{0\leq i\leq 2^n-1}f_n(\alpha_{n,j},\alpha_{n,i})\mathbf{p}_i^*))}\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)\\
&=\sum_{0\leq k\leq 2^n-1}\frac{\exp(\eta^{-1}\sum_{0\leq j\leq 2^n-1}f_n(\alpha_{n,k},\alpha_{n,j})\mu_{\mathbf{p}^*}[\alpha_{n,j},\alpha_{n,j+1}))}{2^{-n}\sum_{j=0}^{2^n-1}\exp(\eta^{-1}\sum_{0\leq i\leq 2^n-1}f_n(\alpha_{n,j},\alpha_{n,i})\mu_{\mathbf{p}^*}[\alpha_{n,i},\alpha_{n,i+1}))}\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)\\
&=\sum_{0\leq k\leq 2^n-1}\frac{\exp(\eta^{-1}\int_{\mathscr{S}}f_n(\alpha_{n,k},y)\mu_{\mathbf{p}^*}(dy))}{2^{-n}\sum_{j=0}^{2^n-1}\exp(\eta^{-1}\int_{\mathscr{S}}f_n(\alpha_{n,j},y)\mu_{\mathbf{p}^*}(dy))}\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)\\
&=\sum_{0\leq k\leq 2^n-1}\frac{\exp(\eta^{-1}\varphi_{\alpha_{n,k}}^n(\mu_{\mathbf{p}^*}))}{2^{-n}\sum_{j=0}^{2^n-1}\exp(\eta^{-1}\varphi_{\alpha_{n,j}}^n(\mu_{\mathbf{p}^*}))}\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)\\
&=\sum_{0\leq k\leq 2^n-1}\frac{\exp(\eta^{-1}\varphi_{\alpha_{n,k}}^n(\mu_{\mathbf{p}^*}))}{\sum_{j=0}^{2^n-1}\exp(\eta^{-1}\varphi_{\alpha_{n,j}}^n(\mu_{\mathbf{p}^*}))\lambda([\alpha_{n,j},\alpha_{n,j+1})}\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x)\\
&=\sum_{0\leq k\leq 2^n-1}\frac{\exp(\eta^{-1}\varphi_{\alpha_{n,k}}^n(\mu_{\mathbf{p}^*}))}{\int_{\mathscr{S}}\exp(\eta^{-1}\varphi_y^n(\mu_{\mathbf{p}^*}))\lambda(dy)}\mathbb{1}_{[\alpha_{n,k},\alpha_{n,k+1})}(x).
\end{aligned}\tag{B.17}$$

The proof of Lemma 17 concludes if we are able to show that $\varphi_x^n(\mathfrak{s}_n(\mu_{\mathbf{p}^*}))=\varphi_x^n(\mu_{\mathbf{p}^*})$ for all $x\in\mathscr{S}$. In fact, a more general equality holds: for all $x\in\mathscr{S}$, we have

$$\varphi_x^n(\mathfrak{s}_n(\mu_{\mathbf{p}^n}))=\varphi_x^n(\mu_{\mathbf{p}^n}),\quad \text{for all } \mathbf{p}^n\in\Delta(S_n).$$

To this end, fix $x \in \mathscr{S}$. Then, using the definition of φ^n, we have that

$$\begin{aligned}
\varphi_x^n(\mathfrak{s}_n(\mu_{\mathbf{p}^n})) &= \int_{\mathscr{S}} f_n(x,y)\mathfrak{s}_n(\mu_{\mathbf{p}^n})(dy) \\
&= \sum_{0 \le k \le 2^n - 1} f_n(x, \alpha_{n,k})\mathfrak{s}_n(\mu_{\mathbf{p}^n})[\alpha_{n,k}, \alpha_{n,k+1}) \\
&= \sum_{0 \le k \le 2^n - 1} f_n(x, \alpha_{n,k})\mathbf{p}_k^n \\
&= \int_{\mathscr{S}} f_n(x,y)\mu_{\mathbf{p}^n}(dy) \\
&= \varphi_x^n(\mu_{\mathbf{p}^n}).
\end{aligned}$$

Plugging in the above observation in (B.17), we have the following equality:

$$\frac{\int^n(\mu_{\mathbf{p}^*})}{d\lambda}(x) = \sum_{0 \le k \le 2^n - 1} \frac{\exp(\eta^{-1}\varphi^n_{\alpha_{n,k}}(\int^n(\mu_{\mathbf{p}^*})))}{\int_{\mathscr{S}} \exp(\eta^{-1}\varphi_y^n(\int^n(\mu_{\mathbf{p}^*})))\lambda(dy)} \mathbb{1}_{[\alpha_{n,k}, \alpha_{n,k+1})}(x), \quad \text{for all } x \in \mathscr{S}.$$

Thus, using the above equality, we have by definition of level-n step logit equilibrium that $\mathfrak{s}_n(\mu_{\mathbf{p}^*}) = \mathfrak{L}_\eta^n(\mathfrak{s}_n(\mu_{\mathbf{p}^*}))$. As a result, we have that $\mathfrak{s}_n(\mu_{\mathbf{p}^*}) \in \mathrm{LE}_\eta(\varphi^n)$. Similarly, suppose $\mathbf{p}^*$ be such that $\mathfrak{s}_n(\mu_{\mathbf{p}^*}) \in \mathrm{LE}_\eta(\varphi^n)$. It then follows using similar arguments that $\mathbf{p}^* \in \mathrm{LE}_\eta(\mathscr{D}_\varphi^n)$. This concludes the proof of Lemma 17.

Therefore by applying Lemma 17 along with Hofbauer and Sandholm (2007), we then have that for every $\epsilon > 0$, interpolated smoothed version $\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t))_{t \ge 0}$ converges in the medium run to $\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))$.

Step 3: In this step, we use the deterministic approximation result (Proposition 3) to obtain that for sufficiently large $T \ge T_\epsilon$, the solution to the level-n step logit dynamic $(\mu^n(t))_{t \ge 0}$ is close to an element of $\mathrm{LE}_\eta(\varphi^n)$ in the medium run. We require the following lemma in this regard.

Lemma 18 *For every $\epsilon > 0$, there exits $T_\epsilon > 0$ such that for all $T \ge T_\epsilon$, we have $\mu^n(t) \in \mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))$ for all $T_\epsilon \le t \le T$.*

Proof 56 Consider the probability space $(\Omega, \mathscr{F}, \mathbb{P})$ on which the interpolated stochastic evolutionary process $(\widehat{\mathbb{X}}_n^N(t))_{t \ge 0}$ is defined. Let $(\widehat{\mathbb{X}}_n(t))_{t \ge 0}$ be a process such that $\widehat{\mathbb{X}}_n(t): \Omega \to \mathscr{P}(\mathscr{S})$ is defined as $\widehat{\mathbb{X}}_n(t,\omega) = \mu^n(t)$ for all $\omega \in \Omega$. The fact that $(\widehat{\mathbb{X}}_n(t))_{t \ge 0}$ is a deterministic process implies that it is measurable. To prove the lemma, we need to show that

$$\mathbb{P}_{\mu_0}\Big(\widehat{\mathbb{X}}_n(t) \in \mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*})) \text{ for all } T_\epsilon \le t \le T\Big) = 1.$$

To this end, we first show that the solution trajectory $\mu^n(t)$ is continuous in t. Note that by definition of φ^n, we have that for all $\mu, \nu \in \mathscr{P}(\mathscr{S})$ and all $x \in \mathscr{S}$,

$$\begin{aligned}
|\varphi_x^n(\mu) - \varphi_n^n(\nu)| &= \Big| \int_{\mathscr{S}} f_n(x,y)\mu(dy) - \int_{\mathscr{S}} f_n(x,y)\nu(dy) \Big| \\
&\le \|f_n\|_\infty \|\mu - \nu\|_{\mathbf{TV}}.
\end{aligned}$$

As a result, we have $\|\varphi^n(\mu)-\varphi^n(\nu)\|_\infty \le \|f_n\|_\infty\|\mu-\nu\|_{\mathbf{TV}}$ for all $\mu,\nu\in\mathscr{P}(\mathscr{S})$. Thus, φ^n is a Lipschitz population game, and hence by Lahkar and Riedel (2015), we have that the step-wise logit dynamic admits a continuous solution $(\mu^n(t))_{t\ge0}$. Therefore $\widehat{\mathbb{X}}_n$ has continuous paths $\mathbb{P}$-a.s. This implies that

$$\mathbb{P}_{\mu_0}\Big(\widehat{\mathbb{X}}_n(t)\in\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\text{ for all }T_\epsilon\le t\le T\Big)=\mathbb{P}_{\mu_0}\Big(\widehat{\mathbb{X}}_n(t)\in\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\text{ for all }t\in[T_\epsilon,T]\cap\mathbb{Q}^+\Big),$$

where $\mathbb{Q}^+$ denotes the collection of all positive rationals. Let $(r_k)_{k\ge1}$ be an enumeration of rationals in $[T_\epsilon,T]$. Thus, in order to prove the lemma, we need to show that $\mathbb{P}_{\mu_0}\big(\cap_{k=1}^\infty\widehat{\mathbb{X}}_n(r_k)\in\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\big)=1$. To this end, we show that $\mathbb{P}_{\mu_0}\big(\widehat{\mathbb{X}}_n(r_k)\in\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\big)=1$ for all $k\ge1$. Fix $\bar\epsilon>0$. Then, for all $N\ge1$, it holds that

$$\begin{aligned}\Big\{\|\widehat{\mathbb{X}}_n(r_k)-\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\|_{\mathbf{TV}}>\bar\epsilon\Big\}\subseteq&\Big\{\|\widehat{\mathbb{X}}_n(r_k)-\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(r_k))\|_{\mathbf{TV}}>\bar\epsilon/2\Big\}\\&\cup\Big\{\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(r_k)-\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\|_{\mathbf{TV}}>\bar\epsilon/2\Big\}.\end{aligned}$$

This implies that

$$\begin{aligned}\mathbb{P}_{\mu_0}\Big(\|\widehat{\mathbb{X}}_n(r_k)-\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\|_{\mathbf{TV}}>\bar\epsilon\Big)\le&\,\mathbb{P}_{\mu_0}\Big(\|\widehat{\mathbb{X}}_n(r_k)-\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(r_k))\|_{\mathbf{TV}}>\bar\epsilon/2\Big)\\&+\mathbb{P}_{\mu_0}\Big(\|\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(r_k)-\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\|_{\mathbf{TV}}>\bar\epsilon/2\Big).\end{aligned}$$

We now apply Theorem 3 and Theorem 4.1 of Hofbauer and Sandholm (2007) to conclude that there exists $N_{\bar\epsilon}\ge1$ sufficiently large such that for all $N\ge N_{\bar\epsilon}$, we have

$$\begin{aligned}\mathbb{P}_{\mu_0}\Big(\|\widehat{\mathbb{X}}_n(r_k)-\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))\|_{\mathbf{TV}}>\bar\epsilon\Big)&\le2(2^n+1)e^{-N\bar C\bar\epsilon^2}+\bar\epsilon\\&\le2\bar\epsilon\end{aligned}$$

Since $\bar\epsilon$ is arbitrary, we conclude that $\mathbb{P}_{\mu_0}(\widehat{\mathbb{X}}_n(r_k)\in\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*})))=1$. Also, since $k\ge1$ is arbitrary, we have that $\mathbb{P}(\widehat{\mathbb{X}}_n(r_k)\in\mathscr{O}^\epsilon(\mathrm{LE}_\eta(\varphi^n)))=1$ for all $k\ge1$. Now, by an application of Bonferroni's inequality (see Billingsley (2008)), we have that $\mathbb{P}(\cap_{k=1}^\infty\widehat{\mathbb{X}}_n(r_k)\in\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*})))=1$. This concludes the proof of Step 3. Since $(\widehat{\mathbb{X}}_n(t))_{t\ge0}$ is a degenerate stochastic process, we have that $\mu^n(t)\in\mathscr{O}^\epsilon(\mathfrak{s}_n(\mu_{\mathbf{p}_n^*}))$ for all $T_\epsilon\le t\le T$.

Step 4: As a final step of the proof, we now proceed to show that the solution to the continuum strategy logit dynamic $(\mu(t))_{t\ge0}$ is close to the continuum strategy logit equilibria $\mathrm{LE}_\eta(\varphi)$ in the medium run. We again show this in the following lemma.

Lemma 19 *For $\epsilon>0$, there exists $T_\epsilon>0$ such that for all $T\ge T_\epsilon$, we have $\mu(t)\in\mathscr{O}^{3\epsilon}(LE_\eta(\varphi))$ for all $T_\epsilon\le t\le T$.*

Proof 57 First note that since the original supermodular game φ is supermodular, we have by Proposition 5 that the level-n discretization $\mathscr{D}_\varphi^n$ is supermodular for every $n\ge1$. By assumption of the theorem, $\mathscr{D}_\varphi^n$ is irreducible for every $n\ge1$. Thus, it follows from Hofbauer and Sandholm (2007) that the discretized stochastic evolutionary process $(\mathfrak{d}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\ge0}$ converges in the medium run to an element of $\mathrm{LE}_\eta(\mathscr{D}_\varphi^n)$. Let us denote by $\mathbf{p}_*^n$ the medium run limit. Again, it follows from Lemma 17 that $(\mathfrak{s}_n(\widehat{\mathbb{X}}_n^N(t)))_{t\ge0}$ converges in the medium run to $\mathfrak{s}_n(\mu_{\mathbf{p}_*^n})$, which is an element of $\mathrm{LE}_\eta(\varphi^n)$. By applying Lemma 18, we have that $\mu^n(t)$ stays in the medium run near $\mathfrak{s}_n(\mu_{\mathbf{p}_*^n})$. By Lemma 4, there exists a subsequence of $\mathfrak{s}_n(\mu_{\mathbf{p}_*^n})$ which converges to μ^* under the variational norm for some $\mu^*\in\mathscr{P}(\mathscr{S})$. Again by Lemma 4, the limit μ^* is a continuum strategy logit equilibrium corresponding to the game φ, that is, $\mu^*\in\mathrm{LE}_\eta(\varphi)$. We now proceed show that the solution to the continuum strategy logit dynamic $(\mu(t))_{t\geq 0}$ stays near μ^* in the medium run. To this end we make use of the approximation result (Proposition 2). The fact that $\sup_{0\leq t\leq T}\|\mu^n(t)-\mu(t)\|_{\mathbf{TV}}\to 0$ as $n\to\infty$, implies that for every $\epsilon>0$, there exists $N_\epsilon\geq 1$ such that $\sup_{0\leq t\leq T}\|\mu^n(t)-\mu(t)\|_{\mathbf{TV}}\leq\epsilon$ for all $n\geq N_\epsilon$. For $T\geq T_\epsilon$, we have by Lemma 18 that $\mu^n(t)$ stays in the medium run near $\mathfrak{s}_n(\mu_{\mathbf{p}_*^n})$. Again by Lemma 4, we have that $\mathfrak{s}_n(\mu_{\mathbf{p}_*^n})\to\mu^*$ in the total variation norm, with μ^* being a continuum strategy logit equilibrium. Combining all these observations we have that $\mu(t)$ stays in the 3ϵ neighborhood of μ^* in the medium run. This completes the proof of Lemma 19.

This concludes the proof of Theorem 7.

APPENDIX

C

SUPPLEMENT TO CHAPTER 3

C.1 Proof of Proposition 7

Part 1 follows from the observation preceding the statement of the proposition that Nash equilibria of an aggregative game like (4.9) can be characterized by solutions to (4.12). Part 2 is immediate.

For part 3, note that all three Nash equilibria in (4.13) are in monomorphic states. At any monomorphic state $\mu = \delta_\alpha$, the aggregate strategy is $A(\mu) = \alpha$. Hence, applying (4.9), we can write the payoff of any player at such a monomorphic state as

$$G(\alpha) = b\alpha^2 - (d\alpha + k\alpha^2). \tag{C.1}$$

Clearly, $G(0) = 0$ and $G(1) = b - (d+k) > 0$ given our assumptions. Hence, the Nash equilibrium δ_1 generates a strictly higher payoff than the Nash equilibrium δ_0. We now show that δ_1 also generates a strictly higher payoff than the third Nash equilibrium, $\delta_{\frac{d}{b-2k}}$, or $G(1) > G\left(\frac{b}{d-2k}\right)$. That would be sufficient to show that δ_1 is the Pareto optimal Nash equilibrium.

For this, note that $G'(\alpha) = 2\alpha(b-k) - d$. Hence, $G'(\alpha) < 0$ if $\alpha < \frac{d}{2(b-k)}$ and $G'(\alpha) > 0$ if $\alpha > \frac{d}{2(b-k)}$. It is straightforward to check that $\frac{d}{2b-2k} < \frac{d}{b-2k} < 1$. Hence, $G(\alpha)$ is strictly increasing at both $\frac{b}{d-2k}$ and 1. Hence, it must be that $G(1) > G\left(\frac{b}{d-2k}\right)$.

C.2 Proof of Proposition 8

Define the function $g : [0,1] \to \mathbb{R}$ as

$$g(\alpha) = \frac{b}{2}\alpha^2 - (d\alpha + k\alpha^2) = \left(\frac{b}{2} - k\right)\alpha^2 - d\alpha. \tag{C.2}$$

1. Notice from (4.10) and (C.2) that at a monomorphic state $\mu=\delta_\alpha$, $\varphi(\delta_\alpha)=g(\alpha)$. On the other hand, if μ is a polymorphic state with $A(\mu)=\alpha$, then $g(\alpha)>\varphi(\mu)$. This is due to the strict convexity of the cost function $c(x)=dx+kx^2$. Therefore, if μ is such that $A(\mu)=\alpha$, then $g(\alpha)\geqslant\varphi(\mu)$ with equality only if $\mu=\delta_\alpha$.

   It is easy to verify from (C.2) that due to our assumption $b>2k+d$ (see Proposition 7), $g(\alpha)$ is strictly convex and is minimized at $\frac{d}{b-2k}\in(0,1)$. Hence, (C.2) must be locally maximized at $\alpha=0$ and $\alpha=1$. Further, $g(0)=\varphi(\delta_0)$. We now argue that the potential function (6) is also locally maximized at δ_0. Consider $\tilde{\mu}$ close to δ_0 under the weak topology so that $A(\tilde{\mu})=0$. First, suppose $\tilde{\mu}$ is monomorphic. Hence, $\tilde{\mu}=\delta_\varepsilon$, for $\varepsilon\approx0$. Then, $\varphi(\delta_0)=g(0)>g(\varepsilon)=\varphi(\delta_\varepsilon)$. Now consider a polymorphic $\tilde{\mu}$ such that $A(\tilde{\mu})=\varepsilon$. Then, $\varphi(\delta_0)=g(0)>g(\varepsilon)>\varphi(\tilde{\mu})$. Therefore, for all $\tilde{\mu}$ close to δ_0 under the weak topology, we obtain $\varphi(\delta_0)>\varphi(\tilde{\mu})$.

   An analogous argument holds for δ_1 as well.

2. From (4.10), it is easily verified that $\varphi(0)=0$ and $\varphi(1)=\frac{b}{2}-(d+k)$. Hence, $\varphi(1)>\varphi(0)$ if and only if $b>2(d+k)$. The identity of the global maximizer of (4.10) then follows immediately. Recall from part 2 of Proposition 7 that $b>2(d+k)$ is equivalent to $\frac{d}{b-2k}<\frac{1}{2}$. Hence, the remainder of the statement follows.

3. Consider $\delta_{\frac{d}{b-2k}}$. As noted earlier, $\frac{d}{b-2k}$ is the unique minimizer of (C.2). Hence, $g\left(\frac{d}{b-2k}\right)=\varphi\left(\delta_{\frac{d}{b-2k}}\right)<\varphi(\delta_\alpha)=g(\alpha)$ for all $\alpha\in[0,1]\setminus\frac{2}{3}$. That suffices to establish that this Nash equilibrium is not a local maximizer of the potential function (4.10) under the weak topology.

C.3 Proof of Proposition 9

Proof 58 (Proof of part (i)) *In order to prove part (i) of Proposition 9, we require the following two lemmas.*

Lemma 20 *Consider the population game π as defined in (4.1). Then the following statement holds true:*

$$\pi_{\alpha_{n,j}}(\beta_n(\boldsymbol{p}^n))=\mathscr{D}_\pi^n[\boldsymbol{p}^n](\alpha_{n,j}),\quad \textit{for } \boldsymbol{p}^n\in\Delta(S_n) \textit{ and } j\in\{0,1,\dots,2^n-1\}. \tag{C.3}$$

Proof 59 Since the population game π is assumed to be a pairwise interaction game, we observe using (4.1) that for $\mathbf{p}^n\in\Delta(S_n)$ and $j\in\{0,1,\dots,2^n-1\}$,

$$\begin{aligned}
\pi_{\alpha_{n,j}}(\beta_n(\mathbf{p}^n))&=\int_{\mathscr{S}} f(\alpha_{n,j},y)\beta_n(\mathbf{p}^n)(dy)\\
&=\int_{\mathscr{S}} f(\alpha_{n,j},y)\sum_{k=0}^{2^n-1}\mathbf{p}_k^n\delta_{\alpha_{n,k}}(dy)\\
&=\sum_{k=0}^{2^n-1} f(\alpha_{n,j},\alpha_{n_k})\mathbf{p}_k^n\\
&=\langle f(\alpha_{n,j},\cdot),\mathbf{p}^n\rangle\\
&\overset{(4.16)}{=}\langle f_n(\alpha_{n,j},\cdot),\mathbf{p}^n\rangle\\
&=\mathscr{D}_\pi^n[\mathbf{p}^n](\alpha_{n,j}).
\end{aligned}$$

This completes the proof of Lemma 20.

Lemma 21 *One has,*

$$\nabla \mathscr{D}^n_{\varphi}[\boldsymbol{p}^n](\alpha_{n,j}) = \nabla_{\mathscr{F}} \varphi_{\alpha_{n,j}}(\beta_n(\boldsymbol{p}^n)), \quad \textit{for } \boldsymbol{p}^n \in \Delta(S_n) \textit{ and } j \in \{0,1,\dots,2^n-1\}. \tag{C.4}$$

Proof 60 We first note that the map $\mu \mapsto \varphi(\mu)$ is Fréchet differentiable in the variational norm. Therefore by Definition 9, we have that the map $\mathbf{p}^n \mapsto \mathscr{D}^n_{\varphi}(\mathbf{p}^n)$ is differentiable in the usual Euclidean topology. Fix $\mathbf{p}^n \in \Delta(S_n)$, $\mathbf{z}^n \in \mathbb{R}^{2^n}$ and denote the gradient of $\mathscr{D}^n_{\varphi}$ at the point $\mathbf{p}^n$ by $\nabla \mathscr{D}^n_{\varphi}(\mathbf{p}^n)$. Then by (4.25), and the relationship between derivative and its gradient, we have

$$\begin{aligned}
\sum_{j=0}^{2^n-1} \nabla \mathscr{D}^{n,j}_{\varphi}[\mathbf{p}^n]\mathbf{z}^n_j &= \langle \nabla \mathscr{D}^n_{\varphi}(\mathbf{p}^n), \mathbf{z}^n \rangle \\
&= \mathbf{D}\mathscr{D}^n_{\varphi}[\mathbf{p}^n](\mathbf{z}^n) \\
&= \lim_{\epsilon \downarrow 0} \Big[\epsilon^{-1} \big(\mathscr{D}^n_{\varphi}(\mathbf{p}^n + \epsilon \mathbf{z}^n) - \mathscr{D}^n_{\varphi}(\mathbf{p}^n) \big) \Big] \\
&= \lim_{\epsilon \downarrow 0} \Big[\epsilon^{-1} \big(\varphi(\beta_n(\mathbf{p}^n + \epsilon \mathbf{z}^n)) - \varphi(\beta_n(\mathbf{p}^n)) \big) \Big] \\
&= \lim_{\epsilon \downarrow 0} \Big[\epsilon^{-1} \big(\varphi(\beta_n(\mathbf{p}^n) + \epsilon \beta_n(\mathbf{z}^n)) - \varphi(\beta_n(\mathbf{p}^n)) \big) \Big] \\
&= \mathbf{D}_{\mathscr{F}} \varphi[\beta_n(\mathbf{p}^n)](\beta_n(\mathbf{z}^n)) \\
&= \langle \nabla_{\mathscr{F}} \varphi(\beta_n(\mathbf{p}^n)), \beta_n(\mathbf{z}^n) \rangle \\
&= \sum_{j=0}^{2^n-1} \nabla_{\mathscr{F}} \varphi_{\alpha_{n,j}}(\beta_n(\mathbf{p}^n)) \mathbf{z}^n_j.
\end{aligned}$$

Since the above equality holds true for arbitrary $\mathbf{z}^n \in \mathbb{R}^{2^n}$, one concludes that

$$\nabla \mathscr{D}^{n,j}_{\varphi}(\mathbf{p}^n) = \nabla_{\mathscr{F}} \varphi_{\alpha_{n,j}}(\beta_n(\mathbf{p}^n)), \quad \text{for all } j \in \{0,1,\dots,2^n\}. \tag{C.5}$$

This completes the proof of Lemma 21.

We now complete the proof of the proposition using Lemma 20 and Lemma 21. Since π is a potential game, we can conclude that,

$$\nabla \mathscr{D}^{n,j}_{\varphi}[\mathbf{p}^n](\alpha_{n,j}) = \nabla_{\mathscr{F}} \varphi_{\alpha_{n,j}}(\beta_n(\mathbf{p}^n)) = \pi_{\alpha_{n,j}}(\beta_n(\mathbf{p}^n)) = \mathscr{D}^n_{\pi}[\mathbf{p}^n](\alpha_{n,j}). \tag{C.6}$$

Thus, $\mathscr{D}^n_{\pi}$ is a potential game with potential function $\mathscr{D}^n_{\varphi}$. This concludes the proof of part (i) of Proposition 9.

Proof 61 (Proof of part (ii)) *Follows from (4.22) and part (i) of Proposition 9.*

Proof 62 (Proof of part (iii)) Fix $n \geqslant 1$. We need to show that π^n (Definition 18) is a potential game with potential function φ^n (see (4.27)). We prove this in a few steps. First, we show that φ^n is well-defined, in the sense that the definition of φ^n is independent of the choice of the absolutely continuous curve γ. However, the definition of φ^n might depend on the choice of the initial condition of γ. But, this does not affect the proof of the proposition since changing the initial condition of γ only leads to a uniform translation in the value of φ^n, the effect of which diminishes while computing its Fréchet gradient. Then we prove that φ^n is a Fréchet differentiable map with respect to the strong topology. After showing that φ^n is well-defined, we show that the Fréchet gradient $\nabla_{\mathscr{F}}\varphi^n$ of φ^n is identified by π^n (see (4.3) and the paragraph therein). Proceeding, fix $\mu \in \mathscr{P}(\mathscr{S})$ and consider a $\gamma \in \mathbf{AC}((0,1);\mathscr{P}(\mathscr{S}))$ with $\gamma_0 = \delta_0$ and $\gamma_1 = \mu$. It then follows after some computations that,

$$\begin{aligned}\langle \pi^n(\gamma_t), \dot{\gamma}_t\rangle &= \int_{\mathscr{S}} \pi^n_x(\gamma_t)\dot{\gamma}_t(dx) \\ &= \int_{\mathscr{S}}\int_{\mathscr{S}} f_n(x,y)\gamma_t(dy)\dot{\gamma}_t(dx) \\ &= \int_{\mathscr{S}} \sum_{k=0}^{2^n-1} f_n(x,\alpha_{n,k})\gamma_t(\mathscr{S}_{n,k})\dot{\gamma}_t(dx) \\ &= \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\gamma_t(\mathscr{S}_{n,k})\dot{\gamma}_t(\mathscr{S}_{n,j}).\end{aligned}$$

Thus, by an application of the definition (4.27), we have,

$$\varphi^n(\mu) = \varphi(\delta_0) + \int_0^1 \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\gamma_t(\mathscr{S}_{n,k})\dot{\gamma}_t(\mathscr{S}_{n,j})dt. \tag{C.7}$$

Given the curve γ, and $n \geqslant 1$ (fixed before hand), we construct a new curve $\gamma^n : (0,1) \to \Delta(S_n)$ such that for all $0 < t < 1$,

$$\gamma^{n,k}_t = \gamma_t(\mathscr{S}_{n,k}), \quad \text{for all } k = 0,1,\ldots,2^n-1. \tag{C.8}$$

In words, for every $0 < t < 1$, γ^n_t is an element of $\Delta(S_n)$ such that the k-th component $\gamma^{n,k}_t$ of γ^n_t is the probability that the original curve γ assigns to the sub-interval $\mathscr{S}_{n,k}$ at time t. Since $\gamma \in \mathbf{AC}((0,1);\mathscr{P}(\mathscr{S}))$, and the fact that the variational norm on $\Delta(S_n)$ is equivalent to the Euclidean norm on $\Delta(S_n)$, implies that $\gamma^n \in \mathbf{AC}(0,1;\Delta(S_n))$. Also notice that the initial conditions of the curve γ^n is given by $\gamma^n_0 = \mathbf{e}^n_0$ and $\gamma^n_1 = (\mu(\mathscr{S}_{n,0}),\mu(\mathscr{S}_{n,1}),\ldots,\mu(\mathscr{S}_{n,2^n-1}))$. Thus, using definition of $\mathscr{D}^n_\varphi$ in (4.25), we can rewrite (C.7) as

$$\begin{aligned}\varphi^n(\mu) &= \varphi(\delta_0) + \int_0^1 \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\gamma^{n,k}_t\dot{\gamma}^{n,k}_t dt \\ &= \mathscr{D}^n_\varphi(\mathbf{e}^n_0) + \int_0^1 \sum_{j=0}^{2^n-1} \mathscr{D}^n_\pi(\gamma^n_t)(\alpha_{n,j})\dot{\gamma}^{n,j}_t dt \\ &= \mathscr{D}^n_\varphi(\mathbf{e}^n_0) + \int_0^1 \langle \mathscr{D}^n_\pi(\gamma^n_t)(\cdot), \dot{\gamma}^{n,\cdot}_t\rangle dt.\end{aligned} \tag{C.9}$$

However, by part (i) of Proposition 9, $\mathscr{D}^n_\pi$ is a potential game with potential function $\mathscr{D}^n_\varphi$. As a result, the vector field $\mathbf{p}^n \mapsto \mathscr{D}^n_\pi(\mathbf{p}^n)$ is a conservative. Thus the integral in the RHS of Equation (C.9) is in fact independent of the curve γ^n. Since $\gamma \in \mathbf{AC}((0,1);\mathscr{P}(\mathscr{S}))$ is arbitrary, we conclude that the definition of

φ^n is Equation (4.27) is independent of the curve γ.

We now proceed to show that φ^n is Fréchet differentiable and that its Fréchet gradient $\nabla_{\mathscr{F}}\varphi^n$ is indeed identified by the population game π^n. Recall from the proof just prior that the definition of φ^n is does not depend on the choice of the curve γ. Thus, in order to complete the proof of part (ii) of Proposition 9, we shall, in particular, fix a curve γ according to our convenience which we define below. Let $\gamma : (0,1) \to \mathscr{P}(\mathscr{S})$ be defined as

$$\gamma_t^\epsilon := \begin{cases} (1-t)\delta_0 + t\mu, & \text{if } 0 \leqslant t \leqslant 1 \\ \left(\frac{\epsilon - t + 1}{\epsilon}\right)\mu + \left(\frac{t-1}{\epsilon}\right)(\mu + \epsilon\,\nu), & \text{if } 1 < t \leqslant 1+\epsilon. \end{cases} \tag{C.10}$$

In words, in between time points 0 and 1, the curve γ is a line segment joining δ_0 and μ followed by a line segment joining μ and $\mu + \epsilon\,\nu$ between time points 1 and $1+\epsilon$ respectively. Thus, in principle, γ is a piece-wise linear curve essentially joining δ_0, μ, and $\mu+\epsilon\,\nu$ and therefore belongs to $\mathbf{AC}((0,1);\mathscr{P}(\mathscr{S}))$. As a consequence, we have the following two observations:

- $\lim_{\epsilon\downarrow 0}\gamma_1^\epsilon = \mu$.
- $\dot{\gamma}_t^\epsilon = \nu$ for all $1 < t \leqslant 1+\epsilon$.

Since φ^n is Fréchet differentiable, we have by (4.27) that for $\mu \in \mathscr{P}(\mathscr{S})$ and $\nu \in \mathscr{M}(\mathscr{S})$,

$$\begin{aligned} \mathbf{D}_{\mathscr{F}}\varphi^n[\mu](\nu) &= \lim_{\epsilon\downarrow 0}\frac{\varphi^n(\mu+\epsilon\,\nu) - \varphi^n(\nu)}{\epsilon} \\ &= \lim_{\epsilon\downarrow 0}\frac{1}{\epsilon}\Big[\int_0^{1+\epsilon}\langle \pi^n(\gamma_t), \dot{\gamma}_t\rangle\,dt - \int_0^1 \langle \pi^n(\gamma_t), \dot{\gamma}_t\rangle\,dt\Big] \\ &= \lim_{\epsilon\downarrow 0}\frac{1}{\epsilon}\int_1^{1+\epsilon}\langle \pi^n(\gamma_t), \dot{\gamma}_t\rangle\,dt. \end{aligned}$$

Thus, in order to complete the proof of the proposition, we need to show that the limit in the RHS of the above string of equalities is $\langle \pi^n(\mu), \nu\rangle$. Therefore,

$$\begin{aligned}
\int_1^{1+\epsilon} \langle \pi^n(\gamma_t), \dot{\gamma}_t\rangle dt &= \int_1^{1+\epsilon} \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\gamma_t(\mathscr{S}_{n,k})\dot{\gamma}_t(\mathscr{S}_{n,j}) \\
&= \int_1^{1+\epsilon} \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\Big\{\frac{\epsilon-t+1}{\epsilon}\mu(\mathscr{S}_{n,k}) \\
&\qquad\qquad + \frac{t-1}{\epsilon}(\mu(\mathscr{S}_{n,k})+\epsilon\,\nu(\mathscr{S}_{n,k}))\Big\}\nu(\mathscr{S}_{n,j})dt \\
&= \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\Big[\int_1^{1+\epsilon}\Big\{\frac{\epsilon-t+1}{\epsilon}\mu(\mathscr{S}_{n,k}) \\
&\qquad\qquad + \frac{t-1}{\epsilon}(\mu(\mathscr{S}_{n,k})+\epsilon\,\nu(\mathscr{S}_{n,k}))\Big\}dt\Big]\nu(\mathscr{S}_{n,j}) \\
&= \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\Big[\frac{\mu(\mathscr{S}_{n,k})}{\epsilon}(\epsilon t-\frac{t^2}{2}+t)_1^{1+\epsilon} \\
&\qquad\qquad + \frac{\mu(\mathscr{S}_{n,k})+\epsilon\,\nu(\mathscr{S}_{n,k})}{\epsilon}(\frac{t^2}{2}-t)_1^{1+\epsilon}\Big]\nu(\mathscr{S}_{n,j}) \\
&= \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\Big[\frac{\epsilon^2\mu(\mathscr{S}_{n,k})}{2\epsilon}+\frac{\epsilon^2\mu(\mathscr{S}_{n,k})+\epsilon^3\,\nu(\mathscr{S}_{n,k})}{2\epsilon}\Big]\nu(\mathscr{S}_{n,j}) \\
&= \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\Big[\frac{\epsilon\mu(\mathscr{S}_{n,k})}{2}+\frac{\epsilon\mu(\mathscr{S}_{n,k})+\epsilon^2\,\nu(\mathscr{S}_{n,k})}{2}\Big]\nu(\mathscr{S}_{n,j}) \\
&= \sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1} f_n(\alpha_{n,j},\alpha_{n,k})\Big[\epsilon\mu(\mathscr{S}_{n,k})+\frac{\epsilon^2}{2}\,\nu(\mathscr{S}_{n,k})\Big]\nu(\mathscr{S}_{n,j}).
\end{aligned}$$

As a result, using the fact that the step-approximations $(f_n)_{n\geqslant 1}$ are uniformly bounded, we have

$$\begin{aligned}
\lim_{\epsilon\downarrow 0}\Big|\frac{1}{\epsilon}\int_1^{1+\epsilon}\langle \pi^n(\gamma_t),\dot{\gamma}_t\rangle dt - \langle \pi^n(\mu), \nu\rangle\Big| &\leqslant \lim_{\epsilon\downarrow 0}\frac{\epsilon}{2}\sum_{j=0}^{2^n-1}\sum_{k=0}^{2^n-1}|f_n(\alpha_{n,j},\alpha_{n,k})|\,\nu(\mathscr{S}_{n,k})\,\nu(\mathscr{S}_{n,j}) \\
&\leqslant \lim_{\epsilon\downarrow 0}\frac{\epsilon}{2}\|f_n\|_\infty \\
&= 0.
\end{aligned}$$

Thus, we arrive at the following equality: for all $\mu\in\mathscr{P}(\mathscr{S})$,

$$\mathbf{D}_{\mathscr{F}}\varphi^n[\mu](\nu) = \langle \pi^n(\mu), \nu\rangle, \quad \text{for all } \nu\in\mathscr{M}(\mathscr{S}). \tag{C.11}$$

This implies that $\nabla_{\mathscr{F}}\varphi^n(\mu)=\pi^n(\mu)$ for all $\mu\in\mathscr{P}(\mathscr{S})$, verifying the fact that π^n is a potential game with potential function φ^n. This concludes the proof of Proposition 9.

C.4 Proof of Proposition 10

We prove the proposition in two steps. We first verify that the sequence of rescaled potential functions $(N^{-1}\mathscr{D}_{\varphi}^{n,N})_{N\geqslant 1}$ converge uniformly to $\mathscr{D}_{\varphi}^{n}$. Then we show that $\mathscr{D}_{\varphi}^{n}$ is a $\mathscr{C}^1$-potential function.

Step 1: Observe that the potential function φ is Fréchet differentiable on $\mathscr{P}(\mathscr{S})$ under the variational norm, and hence continuous on $\mathscr{P}(\mathscr{S})$. Also notice that $\Delta(S_n)$ is a compact under the variational norm when seen as an embedded subset of $\mathscr{P}(\mathscr{S})$. Thus, one can say that there exists $M>0$ such that $|\varphi(\beta_n(\mathbf{p}^n))|\leqslant M$ for all $\mathbf{p}^n\in\Delta(S_n)$. Recall that by definition, we have for all $\mathbf{p}^{n,N}\in\Delta_N(S_n)$,

$$
\begin{aligned}
\max_{\mathbf{p}^{n,N}\in\Delta_N(S_n)}|\epsilon_N\mathscr{D}_{\varphi}^{n,N}(\mathbf{p}^{n,N})-\mathscr{D}^{n}(\mathbf{p}^{n})| &= \max_{\mathbf{p}^{n,N}\in\Delta_N(S_n)}\Big|\frac{1}{N-1}\Big(N\mathscr{D}_{\varphi}^{n}(\mathbf{p}^{n,N})\\
&\qquad -\sum_{k\in S_n} f_n(\alpha_{n,k},\alpha_{n,k})\mathbf{p}_k^{n,N}\Big)-\mathscr{D}_{\varphi}^{n}(\mathbf{p}^{n,N})\Big|\\
&\leqslant \frac{1}{N-1}\sup_{\mathbf{p}^n\in\Delta(S_n)}|\mathscr{D}_{\varphi}^{n}(\mathbf{p}^{n})|+\frac{1}{N-1}\|f_n\|_\infty\\
&= \frac{1}{N-1}\sup_{\mathbf{p}^n\in\Delta(S_n)}|\varphi(\beta_n(\mathbf{p}^{n}))|+\frac{1}{N-1}\|f_n\|_\infty\\
&\leqslant \frac{M}{N-1}+\frac{1}{N-1}\|f_n\|_\infty\\
&\to 0,\quad \text{as } N\to\infty.
\end{aligned}
$$

Step 2: In this step, we need to verify that $\mathscr{D}_{\varphi}^{n}$ is a $\mathscr{C}^1$-potential function. To this end, we show that φ is a $\mathscr{C}^1$-potential function. Let $\mathbb{B}(\mathscr{M}(\mathscr{S}),\mathbb{R})$ denote the space of bounded linear operators from $\mathscr{M}(\mathscr{S})$ to $\mathbb{R}$ equipped with the operator norm. To prove that $\mathscr{D}_{\varphi}^{n}$ is a $\mathscr{C}^1$-potential function, it is enough to show that the mapping $\mathbf{D}_{\mathscr{F}}\varphi:\mathscr{M}(\mathscr{S})\to\mathbb{B}(\mathscr{M}(\mathscr{S}),\mathbb{R})$ is continuous. Let $(\mu_m)_{n\geqslant 1},\mu\subseteq\mathscr{M}(\mathscr{S})$ be such that $\mu_m\to\mu$ in the variational norm as $m\to\infty$. It then follows that,

$$
\begin{aligned}
\|\mathbf{D}_{\mathscr{F}}\varphi(\mu_m)-\mathbf{D}_{\mathscr{F}}\varphi(\mu)\|_{\mathbf{op}} &= \sup_{\|\nu\|_{\mathbf{TV}}\leqslant 1}|\mathbf{D}_{\mathscr{F}}\varphi(\mu_m)(\nu)-\mathbf{D}_{\mathscr{F}}\varphi(\mu)(\nu)|\\
&= \sup_{\|\nu\|_{\mathbf{TV}}\leqslant 1}|\langle\pi(\mu_m),\nu\rangle-\langle\pi(\mu),\nu\rangle|\\
&\leqslant 2\|f\|_\infty\|\mu_m-\mu\|_{\mathbf{TV}}\\
&\to 0,\quad \text{as } m\to\infty.
\end{aligned}
$$

Thus, the conditions of Theorem 12.2.3 in Sandholm (2010) are verified and hence concludes the proof of Proposition 10.

C.5 Proof of Theorem 8

In order to prove the theorem, we shall make use of Proposition 10. In order to proceed with the proof of the theorem, we require the following observation, which we present in the form of a lemma.

Lemma 22 *One has*

$$
\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n,N})=\Delta\mathscr{D}_{\varphi}^{n,\eta}(\bar{\beta}_{n,N}^{-1}(\bar{\mu}^{n,N})),\quad \textit{for all } \bar{\mu}^{n,N}\in\bar{\mathscr{P}}^{n,N}(\mathscr{S}),\tag{C.12}
$$

Proof 63 Fix $\bar{\mu}^{n,N}(\mathscr{S})\in\bar{\mathscr{P}}^{n,N}(\mathscr{S})$. We then observe that

$$\begin{aligned}
\varphi^n_\eta(\bar{\mu}^{n,N}) &= \varphi^n(\bar{\mu}^{n,N})-\eta\int_{\mathscr{S}}\bar{\mu}^{n,N}_x\log\bar{\mu}^{n,N}_x\,dx\\
&=\varphi(\delta_0)+\int_0^1\langle\pi^n(\gamma_t),\dot{\gamma}_t\rangle\,dt-\eta\sum_{j=0}^{2^n-1}2^n\mathbf{p}^{n,N}_j\log(2^n\mathbf{p}^{n,N}_j)\lambda(\mathscr{S}_{n,j})\\
&=\varphi(\delta_0)+\int_0^1\langle\pi^n(\gamma_t),\dot{\gamma}_t\rangle\,dt-\eta\sum_{j=0}^{2^n-1}\mathbf{p}^{n,N}_j\log\mathbf{p}^{n,N}_j-\eta n\log 2\\
&=\varphi(\delta_0)+\int_0^1\langle\mathscr{D}^n_\pi[\gamma^n_t],\dot{\gamma}^n_t\rangle\,dt-\eta\sum_{j=0}^{2^n-1}\mathbf{p}^{n,N}_j\log\mathbf{p}^{n,N}_j-\eta n\log 2,
\end{aligned}$$

where $\gamma^n:(0,1)\to\Delta(S_n)$ is the mapping such that for all $0<t<1$, we have $\gamma^{n,j}_t:=\gamma_t(\mathscr{S}_{n,j})$, for all $j=0,1,\ldots,2^n-1$, with $\gamma^n_0=\mathbf{e}^n_0$ and $\gamma^n_1=(\bar{\mu}^{n,N}(\mathscr{S}_{n,0}),\bar{\mu}^{n,N}(\mathscr{S}_{n,1}),\ldots,\bar{\mu}^{n,N}(\mathscr{S}_{n,2^n-1}))=\bar{\beta}^{-1}_{n,N}(\bar{\mu}^{n,N})$. We now state a representation result which we shall use in the course of the proof of Lemma 22.

Lemma 23 *Suppose π is a potential game with potential function φ. For every $n\geqslant 1$, let $\mathscr{D}^n_\pi$ be the corresponding LPFSD game induced by π. Then the following representation holds modulo translation:*

$$\varphi(\beta_n(\boldsymbol{p}^n))=\varphi(\beta_n(\boldsymbol{e}^n_0))+\int_0^1\langle\mathscr{D}^n_\pi[\gamma^n_t],\dot{\gamma}^n_t\rangle\,dt,\quad \textit{for all } \boldsymbol{p}^n\in\Delta(S_n),\tag{C.13}$$

where $\gamma^n\in\boldsymbol{AC}(0,1;\Delta(S_n))$ satisfying $\gamma^n_0=\boldsymbol{e}^n_0$ and $\gamma^n_1=\boldsymbol{p}^n$.

Proof 64 Fix $\mu,\nu\in\mathscr{P}(\mathscr{S})$ and consider an $\mathscr{C}^1$-regular curve γ with $\gamma_0=\mu$ and $\gamma_1=\nu$. Observe that by chain rule of differentiation, we have

$$\frac{d}{dt}\varphi^n\circ\gamma_t=\langle\nabla^{\mathbf{F}}\varphi^n(\gamma_t)),\dot{\gamma}_t\rangle,\quad\text{for all } 0<t<1.\tag{C.14}$$

Thus, using Fundamental theorem of calculus (Rudin (2006)), (C.14), and the fact that $\mathscr{D}^n_\pi$ is a potential game with potential function $\mathscr{D}^n_\varphi$, one has

$$\begin{aligned}
\mathscr{D}^n_\varphi(\gamma^n_t)-\mathscr{D}^n_\varphi(\mathbf{e}^n_0) &= \mathscr{D}^n_\varphi(\gamma^n_t)-\mathscr{D}^n_\varphi(\gamma^n_0)\\
&=\int_0^t\frac{d}{du}\mathscr{D}^n_\varphi\circ\gamma^n_u\,du\\
&=\int_0^t\langle\nabla\mathscr{D}^n_\varphi(\gamma^n_u),\dot{\gamma}^n_u\rangle\,du\\
&=\int_0^t\langle\mathscr{D}^n_\pi(\gamma^n_u),\dot{\gamma}^n_u\rangle\,du.
\end{aligned}\tag{C.15}$$

We make the following observations. Firstly, by construction, f_n is bounded. Since $\mathscr{D}^n_\pi$ is generated via f_n according to (4.17) (see Definition 18), it is uniformly bounded. As a result, the RHS of (C.15) is uniformly bounded. Also, notice that γ^n is a $\mathscr{C}^1$-regular on $\Delta(S_n)$ and $\mathscr{D}^n_\varphi$ is differentiable. Hence, $t \mapsto \gamma^n_t$ and $\mathbf{p}^n \mapsto \mathscr{D}^n_\varphi(\mathbf{p}^n)$ is continuous. Thus, by an application of Bounded Convergence Theorem,

$$\begin{aligned}\mathscr{D}^n_\varphi(\mathbf{p}^n) - \mathscr{D}^n_\varphi(\mathbf{e}^n_0) &= \lim_{t\uparrow 1} \mathscr{D}^n_\varphi(\gamma^n_t) - \mathscr{D}^n_\varphi(\mathbf{e}^n_0)\\ &= \lim_{t\uparrow 1}\int_0^t \langle \mathscr{D}^n_\pi(\gamma^n_u), \dot{\gamma}^n_u\rangle du\\ &= \int_0^1 \langle \mathscr{D}^n_\pi(\gamma^n_u), \dot{\gamma}^n_u\rangle du.\end{aligned}$$

Therefore, by definition of $\mathscr{D}^n_\varphi$, we can conclude the proof of Lemma 23 by observing that

$$\begin{aligned}\varphi(\beta_n(\mathbf{p}^n)) - \varphi(\beta_n(\mathbf{e}^n_0)) &= \mathscr{D}^n_\varphi(\mathbf{p}^n) - \mathscr{D}^n_\varphi(\mathbf{e}^n_0)\\ &= \int_0^1 \langle \mathscr{D}^n_\pi(\gamma^n_u), \dot{\gamma}^n_u\rangle du.\end{aligned}$$

This concludes the proof of Lemma 23.

We now proceed with the proof of the lemma. Thus, by an application of Lemma 23, we observe that

$$\begin{aligned}\Delta\varphi^n_\eta(\bar{\mu}^{n,N}) &= \int_0^1 \langle \mathscr{D}^n_\pi(\gamma^n_t), \dot{\gamma}^n_t\rangle dt - \eta\sum_{j=0}^{2^n-1}\bar{\beta}^{-1}_{n,N,j}(\bar{\mu}^{n,N})\log\bar{\beta}^{-1}_{n,N,j}(\bar{\mu}^{n,N})\\ &\qquad - \max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}^{n,N}(\mathscr{S})}\Big(\int_0^1 \langle \mathscr{D}^n_\pi(\gamma^n_t), \dot{\gamma}^n_t\rangle dt - \eta\sum_{j=0}^{2^n-1}\bar{\beta}^{-1}_{n,N,j}(\bar{\mu}^{n,N})\log\bar{\beta}^{-1}_{n,N,j}(\bar{\mu}^{n,N})\Big)\\ &= \mathscr{D}^n_\varphi(\bar{\beta}^{-1}_n(\bar{\mu}^{n,N})) - \eta\sum_{j=0}^{2^n-1}\bar{\beta}^{-1}_{n,N,j}(\bar{\mu}^{n,N})\log\bar{\beta}^{-1}_{n,N,j}(\bar{\mu}^{n,N})\\ &\qquad - \max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}^{n,N}(\mathscr{S})}\Big(\mathscr{D}^n_\varphi(\bar{\beta}^{-1}_{n,N}(\bar{\mu}^{n,N})) - \eta\sum_{j=0}^{2^n-1}\bar{\beta}^{-1}_{n,N,j}(\bar{\mu}^{n,N})\log\beta^{-1}_{n,j}(\bar{\mu}^{n,N})\Big)\\ &= \Delta\mathscr{D}^n_\varphi(\bar{\beta}^{-1}_{n,N}(\bar{\mu}^{n,N})).\end{aligned}$$

This concludes the proof of Lemma 22.

For $n \geqslant 1$, let us introduce the following subset of $\mathscr{P}(\mathscr{S})$:

$$\bar{\mathscr{P}}_n(\mathscr{S}) := \left\{\mu\in\mathscr{P}(\mathscr{S}) : \mu(A) = \sum_{j=0}^{2^n-1} 2^n\mathbf{p}^n_j\lambda(A\cap\mathscr{S}_{n,j}) \text{ for all } A\in\mathscr{B}(\mathscr{S}) \text{ for some } \mathbf{p}^n\in\Delta(S_n)\right\}. \tag{C.16}$$

Thus, $\bar{\mathscr{P}}_n(\mathscr{S})$ is the collection of all absolutely continuous probability measures on $\mathscr{S}$ with step probability density functions, such that the probability in the sub-intervals $\{\mathscr{S}_{n,j} : j = 0,1,\dots 2^n-1\}$ is identified by some element in $\Delta(S_n)$.

We now proceed with the proof of the concerned theorem. We do so by invoking the result from

Lemma 22 in the second line of the string of inequalities which follows below:

$$\begin{aligned}\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}(\bar{\mu}^{n,N})-\Delta\varphi_{\eta}(\bar{\mu}^{n,N})\right| &\leqslant \sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}(\bar{\mu}^{n,N})-\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n,N})\right| \\ &\quad + \sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}|\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n,N})-\Delta\varphi_{\eta}(\bar{\mu}^{n,N})| \\ &= \sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}(\bar{\mu}^{n,N})-\Delta\mathscr{D}_{\varphi}^{n,\eta}(\bar{\beta}_{n,N}^{-1}(\bar{\mu}^{n,N}))\right| \\ &\quad + \sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}|\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n,N})-\Delta\varphi_{\eta}(\bar{\mu}^{n,N})|.\end{aligned}$$

Thus, allowing N to approach infinity, we have

$$\begin{aligned}&\lim_{N\to\infty}\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}(\bar{\mu}^{n,N})-\Delta\varphi_{\eta}(\bar{\mu}^{n,N})\right| \\ &\leqslant \lim_{N\to\infty}\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\left|\frac{\eta}{N}\log\mu_{\eta}^{n,N}\circ\bar{\beta}_{n,N}^{-1}(\bar{\mu}^{n,N})-\Delta\mathscr{D}_{\varphi}^{n,\eta}(\bar{\beta}_{n,N}^{-1}(\bar{\mu}^{n,N}))\right| \\ &\quad + \sup_{\bar{\mu}^{n}\in\bar{\mathscr{P}}_{n}(\mathscr{S})}|\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n})-\Delta\varphi_{\eta}(\bar{\mu}^{n})| \\ &\leqslant \lim_{N\to\infty}\sup_{\mathbf{p}^{n,N}\in\Delta_{N}(S_n)}\left|\frac{\eta}{N}\log\mu_{\eta}^{n,N}(\mathbf{p}^{n,N})-\Delta\mathscr{D}_{\varphi}^{n,\eta}(\mathbf{p}^{n,N})\right| \\ &\quad + \sup_{\bar{\mu}^{n}\in\bar{\mathscr{P}}_{n}(\mathscr{S})}|\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n})-\Delta\varphi_{\eta}(\bar{\mu}^{n})|.\end{aligned}$$

By applying Proposition 10, we have that the first term in the RHS of the above inequality

$$\lim_{N\to\infty}\sup_{\mathbf{p}^{n,N}\in\Delta_{N}(S_n)}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}\circ\bar{\beta}_{n,N}^{-1}(\mathbf{p}^{n,N})-\Delta\mathscr{D}_{\varphi}^{n,\eta}(\mathbf{p}^{n,N})\right|=0.$$

As a result, we arrive at the following inequality:

$$\lim_{N\to\infty}\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}(\bar{\mu}^{n,N})-\Delta\varphi_{\eta}(\bar{\mu}^{n,N})\right|\leqslant\sup_{\bar{\mu}^{n}\in\bar{\mathscr{P}}_{n}(\mathscr{S})}\left|\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n})-\Delta\varphi_{\eta}(\bar{\mu}^{n})\right|.$$

Finally, allowing n to approach infinity, we have

$$\lim_{n\to\infty}\lim_{N\to\infty}\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}(\bar{\mu}^{n,N})-\Delta\varphi_{\eta}(\bar{\mu}^{n,N})\right|\leqslant\lim_{n\to\infty}\sup_{\bar{\mu}^{n}\in\bar{\mathscr{P}}_{n}(\mathscr{S})}|\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n})-\Delta\varphi_{\eta}(\bar{\mu}^{n})|.$$

Thus, in order to complete the proof of the theorem, we need to show that the RHS of the above inequality is 0. Thus, proceeding with the right hand side of the above inequality we have

$$\begin{aligned}\lim_{n\to\infty}\sup_{\bar{\mu}^{n}\in\bar{\mathscr{P}}_{n}(\mathscr{S})}|\Delta\varphi_{\eta}^{n}(\bar{\mu}^{n})-\Delta\varphi_{\eta}(\bar{\mu}^{n})| &= \lim_{n\to\infty}\sup_{\bar{\mu}^{n}\in\bar{\mathscr{P}}_{n}(\mathscr{S})}|\Delta\varphi^{n}(\bar{\mu}^{n})-\Delta\varphi(\bar{\mu}^{n})| \\ &\leqslant 2\lim_{n\to\infty}\sup_{\bar{\mu}^{n}\in\bar{\mathscr{P}}_{n}(\mathscr{S})}|\varphi^{n}(\bar{\mu}^{n})-\varphi(\bar{\mu}^{n})|.\end{aligned}$$

The above computation thus implies that it is enough to show that the potential function φ^n converges uniformly to φ over $\bar{\mathscr{P}}_n(\mathscr{S})$. Therefore, we proceed as follows:

$$\begin{aligned}
\sup_{\bar{\mu}^n\in\bar{\mathscr{P}}_n(\mathscr{S})}|\varphi^n(\bar{\mu}^n)-\varphi(\bar{\mu}^n)| &= \sup_{\bar{\mu}^n\in\bar{\mathscr{P}}_n(\mathscr{S})}|(\varphi^n(\bar{\mu}^n)-\varphi(\delta_0))-(\varphi(\bar{\mu}^n)-\varphi(\delta_0))| \\
&\leqslant \sup_{\bar{\mu}^n\in\bar{\mathscr{P}}_n(\mathscr{S})}\sup_{\{\gamma^n\in\mathbf{AC}^1((0,1);\bar{\mathscr{P}}_n(\mathscr{S})):\gamma_0^n=\delta_0,\gamma_1^n=\bar{\mu}^n\}}\left(\left|\int_0^1\langle\pi^n(\gamma_t^n),\dot{\gamma}_t^n\rangle dt\right.\right. \\
&\qquad\left.\left.-\int_0^1\langle\pi(\gamma_t^n),\dot{\gamma}_t^n\rangle dt\right|\right) \\
&\leqslant \sup_{\bar{\mu}^n\in\bar{\mathscr{P}}_n(\mathscr{S})}\sup_{\{\gamma^n\in\mathbf{AC}^1((0,1);\bar{\mathscr{P}}_n(\mathscr{S})):\gamma_0^n=\delta_0,\gamma_1^n=\bar{\mu}^n\}}\int_0^1|\langle\pi^n(\gamma_t^n)-\pi(\gamma_t^n),\dot{\gamma}_t^n\rangle|dt \\
&\leqslant \sup_{\bar{\mu}^n\in\bar{\mathscr{P}}_n(\mathscr{S})}\|f_n-f\|_\infty\|\bar{\mu}^n\|_{\mathbf{TV}} \\
&\leqslant 2\|f_n-f\|_\infty.
\end{aligned}$$

Thus, we arrive at

$$\begin{aligned}
\lim_{n\to\infty}\sup_{\bar{\mu}^n\in\bar{\mathscr{P}}_n(\mathscr{S})}|\varphi^n(\bar{\mu}^n)-\varphi(\bar{\mu}^n)| &\leqslant \lim_{n\to\infty}\|f_n-f\|_\infty \\
&=0.
\end{aligned}$$

This concludes the proof of Theorem 8.

C.6 Proof of Theorem 9

We shall make use of Theorem 8 in order to complete the proof of this theorem. To this end we first obtain the following upper bound. Notice that since $n\log 2\geqslant 1$ for every $n\geqslant 4$, we have via an application of triangle inequality that:

$$\begin{aligned}
&\frac{1}{n\log 2}\left|\frac{\eta}{N}\log\Xi_\eta^{n,N}(\bar{\mu}^{n,N})-2\eta n\log 2-\Delta\varphi(\bar{\mu}^{n,N})\right| \\
&\quad\leqslant\frac{1}{n\log 2}\left|\frac{\eta}{N}\log\Xi_\eta^{n,N}(\bar{\mu}^{n,N})-\Delta\varphi_\eta(\bar{\mu}^{n,N})\right|+\frac{1}{n\log 2}\left|\Delta\varphi_\eta(\bar{\mu}^{n,N})-\Delta\varphi(\bar{\mu}^{n,N})-2\eta n\log 2\right| \\
&\quad\leqslant\left|\frac{\eta}{N}\log\Xi_\eta^{n,N}(\bar{\mu}^{n,N})-\Delta\varphi_\eta(\bar{\mu}^{n,N})\right|+\frac{1}{n\log 2}\left|\Delta\varphi_\eta(\bar{\mu}^{n,N})-\Delta\varphi(\bar{\mu}^{n,N})-2\eta n\log 2\right|.
\end{aligned}$$

Note that as compared to the proof of Theorem 8, this proof requires an additional 'outer-limit' as the noise parameter $\eta\downarrow 0$. Therefore, taking supremum over $\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ followed by the required limits on both sides of the above inequality we have:

$$\lim_{\eta\downarrow 0}\lim_{n\to\infty}\lim_{N\to\infty}\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\frac{1}{n\log 2}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}(\bar{\mu}^{n,N})-2\eta n\log 2-\Delta\varphi(\bar{\mu}^{n,N})\right|$$
$$\leqslant\lim_{\eta\downarrow 0}\lim_{n\to\infty}\lim_{N\to\infty}\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\left|\frac{\eta}{N}\log\Xi_{\eta}^{n,N}(\bar{\mu}^{n,N})-\Delta\varphi_{\eta}(\bar{\mu}^{n,N})\right|$$
$$+\lim_{\eta\downarrow 0}\lim_{n\to\infty}\lim_{N\to\infty}\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\frac{1}{n\log 2}\left|\Delta\varphi_{\eta}(\bar{\mu}^{n,N})-\Delta\varphi(\bar{\mu}^{n,N})-2\eta n\log 2\right|.$$

Thus, in order to complete the proof of Theorem 9, we need to show that both the limits in the right hand side of the above inequality is 0. Notice that by Theorem 8, we have for every $\eta>0$, the first term in the right hand side of the above inequality is 0. As a result, it is only enough to show that the second term in the right hand side of the above inequality is 0. To this end, we require the following lemma.

Lemma 24 *Fix $n,N\geqslant 1$ and consider the map $\bar{\beta}_{n,N}:\Delta_N(S_n)\to\bar{\mathscr{P}}_{n,N}(\mathscr{S})$ as defined in (4.36). Then,*

$$\int_{\mathscr{S}}\bar{\mu}_x^{n,N}\log\bar{\mu}_x^{n,N}(dx)=\sum_{j=0}^{2^n-1}\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})\log\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})+n\log 2,\quad \textit{for all } \bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S}). \tag{C.17}$$

Proof 65 Fix $n,N\geqslant 1$ and consider an arbitrary step probability measure $\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})$. Then, using the definition of β_n one can express the Shannon entropy associated to $\bar{\mu}^{n,N}$ in the following manner:

$$\begin{aligned}\int_{\mathscr{S}}\bar{\mu}_x^{n,N}\log\bar{\mu}_x^{n,N}(dx)&=\int_{\mathscr{S}}\sum_{j=0}^{2^n-1}2^n\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})\log(2^n\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N}))\mathbb{1}_{\mathscr{S}_{n,j}}(x)dx\\&=\sum_{j=0}^{2^n-1}2^n\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})[n\log 2+\log\beta_{n,j}^{-1}(\bar{\mu}^{n,N})]\lambda(\mathscr{S}_{n,j})\\&=\sum_{j=0}^{2^n-1}\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})[n\log 2+\log\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})]\\&=\sum_{j=0}^{2^n-1}\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})\log\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})+n\log 2.\end{aligned}$$

This concludes the proof of Lemma 24.

We now continue with the proof of the theorem. In other words, we proceed to show that the second term in the inequality just prior equals 0. Using the definition of $\Delta\varphi_\eta$ and $\Delta\varphi$ we have by an application of triangle inequality that:

$$\begin{aligned}
&\frac{1}{n\log 2}|\Delta\varphi_\eta(\bar{\mu}^{n,N})-\Delta\varphi(\bar{\mu}^{n,N})-2\eta n\log 2| \\
&\qquad =\frac{1}{n\log 2}\Big|\Big(\varphi_\eta(\bar{\mu}^{n,N})-\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\varphi_\eta(\bar{\mu}^{n,N})\Big) \\
&\qquad\qquad -\Big(\varphi(\bar{\mu}^{n,N})-\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\varphi(\bar{\mu}^{n,N})\Big)-2\eta n\log 2\Big| \\
&\qquad =\frac{1}{n\log 2}\Big|\Big(\varphi_\eta(\bar{\mu}^{n,N})-\varphi(\bar{\mu}^{n,N})-\eta n\log 2\Big) \\
&\qquad\qquad -\Big(\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\varphi_\eta(\bar{\mu}^{n,N})-\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\varphi(\bar{\mu}^{n,N})-\eta n\log 2\Big)\Big| \\
&\qquad \leqslant\frac{1}{n\log 2}|\varphi_\eta(\bar{\mu}^{n,N})-\varphi(\bar{\mu}^{n,N})-\eta n\log 2| \\
&\qquad\qquad +\frac{1}{n\log 2}\Big|\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\varphi_\eta(\bar{\mu}^{n,N})-\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\varphi(\bar{\mu}^{n,N})-\eta n\log 2\Big| \\
&\qquad \leqslant\frac{1}{n\log 2}|\varphi_\eta(\bar{\mu}^{n,N})-\varphi(\bar{\mu}^{n,N})-\eta n\log 2| \\
&\qquad\qquad +\frac{1}{n\log 2}\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}|\varphi_\eta(\bar{\mu}^{n,N})-\varphi(\bar{\mu}^{n,N})-\eta n\log 2| \\
&\qquad \leqslant\frac{2}{n\log 2}\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}|\varphi_\eta(\bar{\mu}^{n,N})-\varphi(\bar{\mu}^{n,N})-\eta n\log 2| \\
&\qquad \leqslant\frac{2\eta}{n\log 2}\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\Big|\int_{\mathscr{S}}\bar{\mu}_x^{n,N}\log\bar{\mu}_x^{n,N}(dx)-n\log 2\Big|.
\end{aligned}$$

We now apply Lemma 24 to obtain the following inequality:

$$\begin{aligned}
&\sup_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\frac{1}{n\log 2}\Big|\Delta\varphi_\eta(\bar{\mu}^{n,N})-\Delta\varphi(\bar{\mu}^{n,N})-2\eta n\log 2\Big| \\
&\qquad \leqslant\frac{2\eta}{n\log 2}\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\Big|\int_{\mathscr{S}}\bar{\mu}_x^{n,N}\log\bar{\mu}_x^{n,N}(dx)-n\log 2\Big| \\
&\qquad =\frac{2\eta}{n\log 2}\max_{\bar{\mu}^{n,N}\in\bar{\mathscr{P}}_{n,N}(\mathscr{S})}\Big|\sum_{j=0}^{2^n-1}\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})\log\bar{\beta}_{n,N,j}^{-1}(\bar{\mu}^{n,N})\Big| \\
&\qquad =\frac{2\eta}{n\log 2}\max_{\mathbf{p}^{n,N}\in\Delta_N(S_n)}\Big|\sum_{j=0}^{2^n-1}\mathbf{p}_j^{n,N}\log\mathbf{p}_j^{n,N}\Big| \\
&\qquad \leqslant\frac{2\eta}{n\log 2}\sup_{\mathbf{p}^{n}\in\Delta(S_n)}\Big|\sum_{j=0}^{2^n-1}\mathbf{p}_j^{n}\log\mathbf{p}_j^{n}\Big|
\end{aligned}$$

Now we observe that the term in the right hand side of the above inequality is in particular the Shannon entropy itself. We know that the supremum of the quantity is attained at the uniform distribution. More precisely, if all the strategies in S_n is allocated the same probabilities, then the term is maximized. Note that there are 2^n-many strategies in S_n for every $n\geqslant 1$. Therefore, the absolute entropy is maximized at

$\mathbf{p}_j^{n,N} = \frac{1}{2^n}$ for all $j = 0, 1, \ldots, 2^n - 1$. In that case, we have

$$\begin{aligned}\frac{2\eta}{n\log 2} \sup_{\mathbf{p}^n \in \Delta(S_n)} \left| \sum_{j=0}^{2^n-1} \mathbf{p}_j^n \log \mathbf{p}_j^n \right| &\leqslant \frac{2\eta}{n\log 2} \left| \sum_{j=0}^{2^n-1} \frac{1}{2^n} \log \frac{1}{2^n} \right| \\ &\leqslant 2\eta.\end{aligned}$$

This leads to the fact that

$$\begin{aligned}\lim_{\eta\downarrow 0} \lim_{n\to\infty} \lim_{N\to\infty} \sup_{\bar{\mu}^{n,N} \in \bar{\mathscr{P}}_{n,N}(\mathscr{S})} \frac{1}{n\log 2} \left| \Delta\varphi_\eta(\bar{\mu}^{n,N}) - \Delta\varphi(\bar{\mu}^{n,N}) - 2\eta n \log 2 \right| \\ \leqslant \lim_{\eta\downarrow 0} \lim_{n\to\infty} \lim_{N\to\infty} \sup_{\mathbf{p}^n \in \Delta(S_n)} \frac{2\eta}{n\log 2} \left| \sum_{j=0}^{2^n-1} \mathbf{p}_j^n \log \mathbf{p}_j^n \right| \\ = 0.\end{aligned}$$

This completes the proof of Theorem 9.

APPENDIX

D

SUPPLEMENT TO CHAPTER 4

D.1 Preliminaries of Böchner Spaces

In this section, we provide a brief exposition on the preliminary theory of Böchner spaces which we use in our proofs.

Definition 49 *Let $(\Omega, \mathscr{A}, \mu)$ be a probability space and let $\mathbb{X}$ be a Banach space. A function $f : \Omega \to \mathbb{X}$ is simple if there exists $x_1, x_2, \ldots, x_k \in \mathbb{X}$ and $E_1, E_2, \ldots, E_k \in \mathscr{A}$ such that $f(\omega) = \sum_{1 \leqslant i \leqslant k} x_i \chi_{E_i}(\omega)$ for all $\omega \in \Omega$, where χ_{E_i} denotes the indicator function on E_i for all $i = 1, 2, \ldots, k$.*

Definition 50 *Let $(\Omega, \mathscr{A}, \mu)$ be a probability space and let $\mathbb{X}$ be a Banach space. A function $f : \Omega \to \mathbb{X}$ is μ-measurable if there exists a sequence of simple functions $(f_n)_{n \geqslant 1}$ with $\lim_{n \to \infty} \|f_n - f\|_{\mathbb{X}} = 0$, μ-almost everywhere.*

Definition 51 *Let $(\Omega, \mathscr{A}, \mu)$ be a probability space and let $\mathbb{X}$ be a Banach space. A function $f : \Omega \to \mathbb{X}$ is weakly μ-measurable if for every $x^* \in \mathbb{X}^*$, the numerical function $x^* f$ is μ-measurable.*

Let $\mathscr{L}^p(\Omega, \mu, \mathbb{X})$ denote the space of all (μ-a.s.) equivalence classes of Böchner integrable functions $f : \Omega \to \mathbb{X}$ with norm defined as

$$\|f\| := \int_{\Omega} \|f(\omega)\|_{\mathbb{X}}^p \mu(d\omega) < \infty.$$

The Böchner space $\mathscr{L}^p(\Omega, \mu, \mathbb{X})$ is a generalization of standard $\mathscr{L}^p(\mu)$ spaces.[1] Note that the space of integrable signed Bayesian strategies as defined in (5.12) can be viewed as the Böchner space $\mathscr{L}^1(\Omega, \mu, \mathbb{X})$, where $\Omega = \Omega$, $\mathscr{A} = \mathscr{B}_\Omega$, $\mu = \xi$, and $\mathbb{X} = \mathbb{R}^n$.

[1] See Diestel and Uhl Jr (1978) and Hille and Phillips (1996) for further details on Böchner spaces.

We now introduce the notion of weak topology on $\mathscr{L}^p(\Omega,\mu,\mathbb{X})$ for $1 \leqslant p < \infty$. Let $f \in \mathscr{L}^p(\Omega,\mu,\mathbb{X})$ and $g \in \mathscr{L}^q(\Omega,\mu,\mathbb{X}^*)$, where $p^{-1}+q^{-1}=1$.[2] Define $\langle f,g\rangle(\omega) := g(\omega)(f(\omega))$ for all $\omega \in \Omega$.[3]

Definition 52 *Let $(\Omega,\mathscr{A},\mu)$ be a probability space and let $\mathbb{X}$ be a Banach space. The weak topology on $\mathscr{L}^p(\Omega,\mu,\mathbb{X})$ is the topology induced by the convergence:*

$$f_n \to^w f \iff \int_\Omega \langle f_n,g\rangle(\omega)\mu(d\omega) \to \int_\Omega \langle f,g\rangle(\omega)\mu(d\omega), \quad \textit{for all } g \in \mathscr{L}^q(\Omega,\mu,\mathbb{X}^*).$$

The following three definitions are required in the context of Theorem 5.

Definition 53 *Let $f \in \mathscr{L}^1(\Omega,\mu,\mathbb{X})$ and $A \in \mathscr{A}$. Then the average value of f over A is defined as*

$$avg(f;A) := \frac{1}{\mu(A)}\int_A f(\omega)\mu(d\omega).$$

Definition 54 *Let $f \in \mathscr{L}^1(\Omega,\mu,\mathbb{X})$ and $A \in \mathscr{A}$. Then the Bocce oscillation of f over A is defined as*

$$Bocce\text{-}osc(f;A) := \frac{1}{\mu(A)}\int_A \|f(\omega) - avg(f;A)\|_{\mathbb{X}}\mu(d\omega).$$

Definition 55 *A subset $\mathbb{K}$ of $\mathscr{L}^1(\Omega,\mu,\mathbb{X})$ is said to satisfy small Bocce oscillation if for every $\epsilon > 0$, there exists a finite measurable partition $\{A_1,\ldots,A_p\} \subseteq \Omega$ such that for all $f \in \mathbb{K}$, we have*

$$\sum_{1\leqslant i\leqslant p} \mu(A_i) Bocce\text{-}osc(f;A_i) < \epsilon.$$

D.2 Proof of proposition 17

We first prove the 'only-if' part. Note that the definition of Bayesian best response (6.4) can be equivalently written as

$$\beta(\sigma) := \Big\{\rho \in \Sigma : \rho(\omega) \in \arg\max_{\mathbf{y}\in\Delta}\langle \mathbf{y}, \mathscr{G}(\sigma,\omega)\rangle \text{ for } \xi-\text{a.e } \omega \in \Omega\Big\}.$$

Now suppose that σ° is a Bayesian equilibrium of $\mathscr{G}$. Then we must have $\sigma^\circ \in \beta(\sigma^\circ)$ which implies by the definition of Bayesian equilibrium that

$$\sigma^\circ(\omega) \in \arg\max_{\mathbf{y}\in\Delta}\langle \mathbf{y}, \mathscr{G}(\sigma^\circ,\omega)\rangle \quad \text{for } \xi-\text{a.e } \omega \in \Omega.$$

As a result, by the original definition of Bayesian best response correspondence, we have $\sigma^\circ(\omega) \in \beta[\sigma^\circ](\omega)$ for ξ-a.e $\omega \in \Omega$. Finally, by the definition of ω-sliced Bayesian best response (6.5), we have that $\sigma^\circ(\omega) \in \beta_\omega(\sigma^\circ)$ for ξ-a.e $\omega \in \Omega$. This completes the proof of the only-if part.

We now proceed with the proof of 'if-part'. So let us suppose that σ° is an ω-sliced Bayesian equilibrium of the game $\mathscr{G}_\omega$ for ξ-a.e $\omega \in \Omega$. Similarly, back tracing as in the proof of only if part, we have σ° is a Bayesian equilibrium of $\mathscr{G}$.

[2] Here $\mathbb{X}^*$ denotes the dual space of $\mathbb{X}$.

[3] Suppose that the Banach space $\mathbb{X}$ is reflexive. Then for $1 \leqslant p < \infty$, the dual of $\mathscr{L}^p(\tilde{\mu},\mathbb{X})$ is identified by $\mathscr{L}^q(\tilde{\mu},\mathbb{X}^*)$, where $p^{-1}+q^{-1}=1$. That is, $\mathscr{L}^p(\tilde{\mu},\mathbb{X})^* \cong \mathscr{L}^q(\tilde{\mu},\mathbb{X}^*)$, for $1 \leqslant p < \infty$. See Diestel and Uhl Jr (1978) for a complete exposition on weak topology on Böchner spaces.

D.3 Proof of Proposition 13

To prove the proposition, we resort to the Pettis Measurability Theorem theorem. We provide the statement of the theorem for completeness.

Theorem 23 (Pettis Measurability Theorem (Diestel and Uhl Jr (1978))) *Let $(\widehat{\Omega}, \mathscr{A}, \mu)$ be a probability space and let $\mathbb{X}$ be a Banach space. A function $f : \widehat{\Omega} \to \mathbb{X}$ is μ-measurable if and only if*

1. *f is μ-essentially separable valued, that is, there exists $E \in \mathscr{A}$ with $\mu(E) = 0$ and such that $f(\widehat{\Omega} \setminus E)$ is (norm) separable subset of $\mathbb{X}$, and*
2. *f is weakly μ-measurable.*

We now proceed with the proof of the proposition. We apply Theorem 23 with $\widehat{\Omega} = \Omega$, $\mu = \xi$, and $\mathbb{X} = \mathbb{R}^n$. To show that the mapping $\sigma \mapsto \beta_\epsilon(\sigma)$ is well-defined, we need to show that for every $\sigma \in \Sigma$, $\beta_\epsilon(\sigma) \in \Sigma$. In other words, we need to show that $\beta_\epsilon[\sigma]$ is ξ-measurable. We thereby verify the conditions of Theorem 23. Note that by definition, $\beta[\sigma] : \Omega \to \Delta$. Since Δ is compact, it is separable. This verifies condition (i). To verify condition (ii), we need to show that $\beta_\epsilon(\sigma) : \Omega \to \Delta$ is weakly ξ-measurable in the sense of Definition 51. Thus, fix $\mathbf{u} \in \mathbb{R}^n$. We need to show that the map $\mathbf{u} \circ \beta_\epsilon[\sigma] : \Omega \to \mathbb{R}$ defined by $\mathbf{u} \circ \beta_\epsilon[\sigma](\omega) := \langle \mathbf{u}, \beta_\epsilon[\sigma](\omega) \rangle$ for all $\omega \in \Omega$ is weakly ξ-measurable. Since $\mathbf{v}$ is a strongly convex function, it follows that $\beta_\epsilon[\sigma](\omega) = \nabla \mathbf{v}_\epsilon^*(\mathscr{G}(\sigma, \omega))$, in which case we have

$$\langle \mathbf{u}, \beta_\epsilon[\sigma](\omega) \rangle = \langle \mathbf{u}, \nabla \mathbf{v}_\epsilon^*(\mathscr{G}(\sigma, \omega)) \rangle \quad \text{for all } \omega \in \Omega.$$

Also since $\mathscr{G}$ is a Bayesian population game, it follows from the definition that the map $\omega \mapsto \mathscr{G}(\sigma, \omega)$ is measurable. Also, it follows from Rockafellar (1970) that $\nabla \mathbf{v}_\epsilon^*$ is $1/\gamma\epsilon$-Lipschitz (hence continuous). Therefore the map $\omega \mapsto \nabla \mathbf{v}_\epsilon^*(\mathscr{G}(\sigma, \omega))$ is measurable. Thus, $\omega \mapsto \langle \mathbf{u}, \nabla \mathbf{v}_\epsilon^*(\mathscr{G}(\sigma, \omega)) \rangle$ is measurable. As a result, by simple function approximation theorem, it follows that the map $\omega \mapsto \mathbf{u} \circ \beta_\epsilon[\sigma](\omega)$ is weakly ξ-measurable. Since $\mathbf{u} \in \mathbb{R}^n$ is arbitrary, condition (ii) is verified. This concludes the proof of Proposition 23.

D.4 Proof of Theorem 10

The proof of Theorem 10 is an application of Brouwer-Schauder-Tychonoff fixed point theorem (Aliprantis and Border (2013))). We again provide the statement of the theorem for completeness.

Theorem 24 (Brouwer-Schauder-Tychonoff FPT) *Let $\mathbb{V}$ be a non-empty compact convex subset of a locally convex Hausdorff space, and let $f : \mathbb{V} \to \mathbb{V}$ be a continuous function. Then the set of fixed points of f is compact and nonempty.*

In order to apply Theorem 24, recall from (5.12) that the space of integrable signed Bayesian strategies $\widehat{\Sigma}$ is identified by the Böchner space $\mathscr{L}^1(\Omega, \xi, \mathbb{R}^n)$. Since $\widehat{\Sigma}$ is endowed with the weak topology, it follows from Kesavan (2009) that it is a Hausdorff topological space. It also follows from the definition that $\widehat{\Sigma}$ is locally convex. Therefore, under the weak topology, the space of integrable signed Bayesian strategies is a locally convex Hausdorff topological space. We now proceed to show that the space of Bayesian strategies Σ is compact under the weak topology. The fact that Σ is convex is obvious.

Lemma 25 *The space of Bayesian strategies Σ is compact under the weak topology.*

Proof 66 The proof of Lemma 33 is an application of Dunford's Theorem (Diestel and Uhl Jr (1978)) which characterizes relatively weakly compact subsets of $\mathscr{L}^1(\Omega, \xi, \mathbb{R}^n)$.

Theorem 25 (Dunford's Theorem (Diestel and Uhl Jr (1978))) *Let $(\widehat{\Omega}, \mathscr{A}, \mu)$ be a finite measure space and $\mathbb{X}$ be a Banach space such that both $\mathbb{X}$ and $\mathbb{X}^*$ satisfy the Radon-Nikodym property.*[4] *A subset $\mathbb{K}$ of $\mathscr{L}^1(\mu, \mathbb{X})$ is relatively weakly compact if*

1. *the subset $\mathbb{K}$ is bounded,*
2. *the subset $\mathbb{K}$ is uniformly integrable, and*
3. *for each $E \in \mathscr{A}$, the set $\left\{ \int_E f(\omega) d\mu(d\omega) : f \in \mathbb{K} \right\}$ is relatively weakly compact.*

We apply Theorem 25 with $\widehat{\Omega} = \Omega$, $\mu = \xi$, $\mathbb{X} = \mathbb{R}^n$, and $\mathbb{K} = \Sigma$. We need to show that the space of Bayesian strategies Σ satisfies the conditions stated in Theorem 25. Note that $\|\sigma(\omega)\| \leqslant \sqrt{n}$, for all $\omega \in \Omega$. Therefore, we have

$$\sup_{\sigma \in \Sigma} |||\sigma||| = \sup_{\sigma \in \Sigma} \int_{\Omega} \|\sigma(\omega)\| \xi(d\omega) \leqslant \sqrt{n}.$$

Therefore the subset Σ is bounded. This verifies condition (i) of Theorem 25.

To verify condition (ii), note that Σ is uniformly integrable, since every Bayesian strategy $\sigma \in \Sigma$ is uniformly bounded by $\sqrt{n}$, for every $\omega \in \Omega$, in which case we have

$$\lim_{\epsilon \to \infty} \sup_{\sigma \in \Sigma} \int_{\|\sigma\| > \epsilon} \|\sigma(\omega)\| \xi(d\omega) \leqslant \sqrt{n} \lim_{\epsilon \to \infty} \sup_{\sigma \in \Sigma} \xi(\|\sigma\| > \epsilon) = 0.$$

Finally we verify condition (iii) of Theorem 25. To this end, let $E \in \mathscr{B}_{\Omega}$. Consider the set

$$\mathscr{I}_E := \left\{ \int_E \sigma(\omega) \xi(d\omega) : \sigma \in \Sigma \right\}.$$

The fact that Δ is a compact convex subset of $\mathbb{R}^n$ and that Range$(\sigma) \subseteq \Delta$ for all $\sigma \in \Sigma$, together imply that $\mathscr{I}_E$ is relatively compact. Since E is arbitrary, condition (iii) of Theorem 25 is satisfied.

This proves that Σ is relatively compact under the weak topology. However since Σ is itself closed under the weak topology, it is compact. This concludes the proof of lemma 33.

To conclude the proof of Theorem 10, we need to show that the mapping $\sigma \mapsto \beta_\epsilon(\sigma)$ is continuous in the weak topology. We show this in the following lemma.

Lemma 26 *The mapping $\sigma \mapsto \beta_\epsilon(\sigma)$ is continuous in the weak topology.*

Proof 67 To prove the lemma, consider a sequence $(\sigma_n)_{n \geqslant 1} \subseteq \Sigma$, such that $\sigma_n \to^w \sigma$. We need to show that $\beta_\epsilon(\sigma_n) \to^w \beta_\epsilon(\sigma)$. Since the regularizer is a γ-strongly convex function and $\mathscr{G}$ is a weakly continuous Bayesian population game, there exists $\bar{\Omega} \subseteq \Omega$ with $\xi(\bar{\Omega}) = 1$ such that for every fixed $\omega \in \bar{\Omega}$, we have that

$$\begin{aligned} \|\beta_\epsilon[\sigma_n](\omega) - \beta_\epsilon[\sigma](\omega)\| &= \|\nabla \mathbf{v}_\epsilon^*(\mathscr{G}(\sigma_n, \omega)) - \nabla \mathbf{v}_\epsilon^*(\mathscr{G}(\sigma, \omega))\| \\ &\leqslant \frac{1}{\epsilon\gamma} \|\mathscr{G}(\sigma_n, \omega) - \mathscr{G}(\sigma, \omega)\| \to 0. \end{aligned}$$

[4]See Diestel and Uhl Jr (1978), Chap III for the definition of Radon-Nikodym property. It is known that if $\mathbb{X}$ is a reflexive Banach space, then $\mathbb{X}$ has the Radon-Nikodym property. In our case, $\mathbb{X} = \mathbb{R}^n$. Thus $\mathbb{X}$ follows the Radon Nikodym property.

This proves the fact that $\beta_\epsilon[\sigma_n](\omega) \to \beta_\epsilon[\sigma](\omega)$ for every $\omega \in \bar{\Omega}$. Now let $g \in \mathscr{L}^\infty(\Omega, \xi, \mathbb{R}^{n*})$. This implies that $g(\omega)$ is a bounded linear functional from $\mathbb{R}^n$ to $\mathbb{R}$ for every $\omega \in \Omega$. Therefore, for every $\omega \in \Omega$, we have $g[\omega](\mathbf{x}_n) \to g[\omega](\mathbf{x})$, whenever $\mathbf{x}_n \to \mathbf{x}$. Since $\beta_\epsilon[\sigma_n](\omega) \to \beta_\epsilon[\sigma](\omega)$ for every $\omega \in \bar{\Omega}$, we have that

$$\begin{aligned}
\int_\Omega \langle \beta_\epsilon(\sigma_n), g\rangle(\omega)\xi(d\omega) &= \int_\Omega g(\omega)(\beta_\epsilon[\sigma_n](\omega))\xi(d\omega) \\
&= \int_{\bar{\Omega}} \langle g(\omega), \beta_\epsilon[\sigma_n](\omega)\rangle \xi(d\omega) \\
&\to \int_{\bar{\Omega}} \langle g(\omega), \beta_\epsilon[\sigma](\omega)\rangle \xi(d\omega) \\
&= \int_\Omega g(\omega)(\beta_\epsilon[\sigma](\omega))\xi(d\omega) \\
&= \int_\Omega \langle \beta_\epsilon(\sigma), g\rangle(\omega)\xi(d\omega).
\end{aligned}$$

Since $g \in \mathscr{L}^\infty(\Omega, \xi, \mathbb{R}^{n*})$ is arbitrary, this proves that the mapping $\sigma \mapsto \beta_\epsilon(\sigma)$ is continuous in the weak topology on Σ which completes the proof of Lemma 35.

Since Σ is a non empty compact convex subset of $\widehat{\Sigma}$ and β_ϵ is a continuous map, the conditions of Theorem 24 are satisfied and the existence of a fixed point is established. This concludes the proof of Theorem 10.

D.5 Proof of Corollary 3

In order to prove the corollary, it is enough to show that the aggregate Bayesian population game $\mathscr{G}$ is weakly continuous. Note that from the definition of $\mathscr{G}$ it is enough to show that $\mathscr{E}$ is continuous. Let $\sigma_n \to \sigma$. We show that $\mathscr{E}(\sigma_n) \to \mathscr{E}(\sigma)$. To this end, fix $i \in S$. Define the projection map $\pi_i : \mathscr{C}_n(\Delta) \to \mathbb{R}^{n*}$, such that for every $\omega \in \mathscr{C}_n(\Delta)$,

$$\pi_i[\omega](\mathbf{x}) := \mathbf{x}_i, \quad \text{for all } \mathbf{x} \in \mathbb{R}^n.$$

Note that $\pi_i \in \mathscr{L}^\infty(\mathscr{C}_n(\Delta), \xi, \mathbb{R}^{n*})$ for all $i \in S$. Since $\sigma_n \to^w \sigma$, we have from the definition of weak topology that

$$\begin{aligned}
\mathscr{E}^i(\sigma_n) &= \int_\Omega \sigma_n^i(\omega)\xi(d\omega) \\
&= \int_\Omega \pi_i(\omega)(\sigma_n(\omega))\xi(d\omega) \\
&= \int_\Omega \langle \sigma_n, \pi_i \rangle(\omega)\xi(d\omega) \\
&\to \int_\Omega \langle \sigma, \pi_i \rangle(\omega)\xi(d\omega) \\
&= \int_\Omega \pi_i(\omega)(\sigma(\omega))\xi(d\omega) \\
&= \int_\Omega \sigma^i(\omega)\xi(d\omega) \\
&= \mathscr{E}^i(\sigma).
\end{aligned}$$

This proves that $\mathscr{E}(\sigma_n) \to \mathscr{E}(\sigma)$. Thus $\sigma_n \to^w \sigma$ implies that $\mathscr{G}(\sigma_n, \omega) = \omega(\mathscr{E}(\sigma_n)) \to \omega(\mathscr{E}(\sigma)) = \mathscr{G}(\sigma, \omega)$ for all $\omega \in \mathscr{C}_n(\Delta)$. Hence $\mathscr{G}$ is weakly continuous. This concludes the proof of Corollary 3.

D.6 Proof of Theorem 11

Let $(\epsilon_m)_{m\geqslant 1}$ be an arbitrary sequence of positive noise parameters such that $\epsilon_m \downarrow 0$. For every $m \geqslant 1$, let σ_m° denote the corresponding ϵ_m-regularized Bayesian equilibrium, the existence of which is guaranteed by Theorem 10. Note that by Lemma 33, the space of Bayesian strategies Σ is compact relative to the weak topology. Thus there exists a subsequence $(\sigma_{m_k})_{k\geqslant 1} \subseteq \Sigma$ such that $\sigma_{m_k}^\circ \to^w \sigma^\circ$, for some $\sigma^\circ \in \Sigma$. We need to show that σ° is a Bayesian equilibrium of the game $\mathscr{G}$. We first prove the following lemma.

Lemma 27 *Fix a noise parameter $\epsilon > 0$. Then σ_ϵ° is a regularized Bayesian equilibrium of the game $\mathscr{G}$ iff $\sigma_\epsilon^\circ(\omega)$ is a regularized equilibrium of the ω-sliced game $\mathscr{G}_\omega$ for ξ-a.e $\omega \in \Omega$.*

Proof 68 The proof of the lemma follows using the definitions of β_ϵ and β_ω^ϵ and is left to the reader.

We now proceed with the proof of the theorem. We shall prove this by the method of contradiction. Suppose on the contrary that σ° is not a Bayesian equilibrium. Then by definition, there exists $\bar{\Omega} \subseteq \Omega$ with $\xi(\bar{\Omega}) > 0$ such that $\sigma^\circ(\omega) \notin \beta[\sigma^\circ](\omega)$ for all $\omega \in \bar{\Omega}$, which in principle implies that $\sigma^\circ(\omega) \notin \operatorname{argmax}_{\mathbf{y}\in\Delta} \langle \mathbf{y}, \mathscr{G}(\sigma^\circ, \omega)\rangle$ for all $\omega \in \bar{\Omega}$. Thus, there exists $\mathbf{i}, \mathbf{j} : \bar{\Omega} \to S$ such that for every $\omega \in \bar{\Omega}$, there exists $\mathbf{j}(\omega) \in S$, $\mathbf{i}(\omega) \in \operatorname{supp}(\sigma^\circ(\omega))$ and an $\eta(\omega) > 0$ satisfying the following inequality:

$$\mathscr{G}^{\mathbf{j}(\omega)}(\sigma^\circ, \omega) > 2\eta(\omega) + \mathscr{G}^{\mathbf{i}(\omega)}(\sigma^\circ, \omega).$$

Since $\mathscr{G}$ is a weakly continuous population game, there exists for every $\omega \in \bar{\Omega}$, a $K(\omega) \geqslant 1$ sufficiently large enough such that

$$\mathscr{G}^{\mathbf{j}(\omega)}(\sigma_{m_k}^\circ, \omega) > \eta(\omega) + \mathscr{G}^{\mathbf{i}(\omega)}(\sigma_{m_k}^\circ, \omega), \quad \text{for sufficiently large } k \geqslant K(\omega). \tag{D.1}$$

We now make the following observation. By our assumptions, we have that $\sigma^\circ_{m_k} \to^w \sigma^\circ$. We show that $\sigma^\circ_{m_k}(\omega) \to \sigma^\circ(\omega)$ as $k \to \infty$ for every $\omega \in \Omega$. To prove the claim, fix an $\omega \in \Omega$ and $i \in S$. Define a function $g^i_\omega : \Omega \to \mathbb{R}^n$ as $g^i_\omega(\bar{\omega}) := \mathbf{e}_i \mathbb{1}_\omega(\bar{\omega})$, for all $\bar{\omega} \in \Omega$. It then follows that $\|g^i_\omega\|_\infty = 1$ and that $g^i_\omega \in \mathscr{L}^\infty(\Omega, \xi, \mathbb{R}^{n*})$, for every $\omega \in \Omega$ and every $i \in S$. Since $\sigma^\circ_{m_k} \to \sigma^\circ$, we have by definition of weak topology that,

$$
\begin{aligned}
\sigma^{\circ,i}_{m_k}(\omega) &= \langle \sigma^\circ_{m_k}(\omega), \mathbf{e}_i \rangle \\
&= \int_\Omega \langle \sigma^\circ_{m_k}, g^i_\omega \rangle(\bar{\omega}) \xi(d\bar{\omega}) \\
&\to \int_\Omega \langle \sigma^\circ, g^i_\omega \rangle(\bar{\omega}) \xi(d\bar{\omega}) \\
&= \langle \sigma^\circ(\omega), \mathbf{e}_i \rangle \\
&= \sigma^{\circ,i}(\omega).
\end{aligned}
$$

Since $\mathbf{i}(\omega) \in \operatorname{supp}(\sigma^\circ(\omega))$ for all $\omega \in \bar{\Omega}$, we must have that $\sigma^{\circ,\mathbf{i}(\omega)}(\omega) > 0$ for all $\omega \in \bar{\Omega}$. However, we show that this cannot happen. Thus, in order to arrive at a contradiction, we show that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \to 0$ for some $\omega \in \bar{\Omega}$. We show this separately for regularizers satisfying Condition C1 followed by regularizers satisfying Condition C2.

Proof 69 (Proof for regularizers satisfying Condition C1) Proceeding, first let us suppose that the regularizer satisfies Condition C1 (Definition 27). We prove a much stronger statement, that is, we show that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \to 0$ for all $\omega \in \bar{\Omega}$. To this end, fix $\omega \in \bar{\Omega}$. Then by (D.1), there exists $\mathbf{i}(\omega), \mathbf{j}(\omega) \in S$ such that $\mathbf{u}^{\mathbf{j}(\omega)}_{m_k} := \mathscr{G}^{\mathbf{j}(\omega)}(\sigma^\circ_{m_k}, \omega) > \mathscr{G}^{\mathbf{i}(\omega)}(\sigma^\circ_{m_k}, \omega) =: \mathbf{u}^{\mathbf{i}(\omega)}_{m_k}$ for $k \geqslant K(\omega)$. Since $\mathbf{v}$ satisfies Condition C1, we have

$$
\nabla \mathbf{v}^*_{\epsilon_{m_k}, \mathbf{j}(\omega)}(\mathscr{G}(\sigma^\circ_{m_k}, \omega)) \geqslant \lambda(\epsilon_{m_k}) \mathbf{v}^*_{\epsilon_{m_k}, \mathbf{i}(\omega)}(\mathscr{G}(\sigma^\circ_{m_k}, \omega)), \quad \text{for all } k \geqslant K(\omega).
$$

Next, note that $\nabla \mathbf{v}^*_{\epsilon_{m_k}, \mathbf{j}(\omega)}(\mathscr{G}(\sigma^\circ_{m_k}, \omega)) = \beta_{\epsilon_{m_k}, \mathbf{j}(\omega)}[\sigma^\circ_{m_k}](\omega)$ and $\nabla \mathbf{v}^*_{\epsilon_{n_k}, \mathbf{i}(\omega)}(\mathscr{G}(\sigma^\circ_{m_k}, \omega)) = \beta_{\epsilon_{m_k}, \mathbf{i}(\omega)}[\sigma^\circ_{m_k}](\omega)$. Also since $\sigma^\circ_{m_k}$ is a ϵ_{m_k}-regularized Bayesian equilibrium of the game $\mathscr{G}$, we have that $\nabla \mathbf{v}^*_{\epsilon_{m_k}, \mathbf{j}(\omega)}(\mathscr{G}(\sigma^\circ_{m_k}, \omega)) = \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ and $\nabla \mathbf{v}^*_{\epsilon_{m_k}, \mathbf{i}(\omega)}(\mathscr{G}(\sigma^\circ_{m_k}, \omega)) = \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$. In what follows, we have

$$
\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) \geqslant \lambda(\epsilon_{m_k}) \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega), \quad \text{for all } k \geqslant K(\omega).
$$

Since ϵ_{m_k} converges to 0 as $k \to \infty$ and that $\mathbf{v}$ satisfies Condition C1, we have $\lambda(\epsilon_{m_k}) \to \infty$ as $k \to \infty$. Thus we can conclude that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \to 0$. However, this leads to a contradiction since $\mathbf{i}(\omega) \in \operatorname{supp}(\sigma^\circ(\omega))$. Since $\omega \in \bar{\Omega}$ is arbitrary, we have that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \to 0$ for all $\omega \in \bar{\Omega}$. Thus, the required claim is proved. Therefore σ° must be a Bayesian equilibrium of the original game $\mathscr{G}$.

Proof 70 (Proof for regularizers satisfying Condition C2) Next we establish the same under Condition C2 (Definition 29). We again prove this via the method of contradiction. However, contrary to Case 1, we show that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \to 0$ for some $\omega \in \bar{\Omega}$, which, as mentioned in the previous arguments, is enough to conclude that σ° is a Bayesian equilibrium. Suppose not. Then we must have that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ converges to $c(\omega)$, where $c(\omega) > 0$ for all $\omega \in \bar{\Omega}$.[5] As a result, $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ is bounded away from 0 for all $\omega \in \bar{\Omega}$. Thus,

[5] We note that the convergence of $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ is guaranteed by weak convergence of $(\sigma^\circ_{m_k})_{k \geqslant 1}$.

we have via Condition C2 that there exists a sequence of shifted Bayesian strategies $(\widetilde{\sigma}^{\circ}_{m_k})_{k\geqslant 1}$, where $(\widetilde{\sigma}^{\circ}_{m_k})_{k\geqslant 1}$ is a shift of $(\sigma^{\circ}_{m_k})_{k\geqslant 1}$ over $\bar{\Omega}$, such that, after some algebra, the following condition holds true for all $\omega\in\bar{\Omega}$:

$$\sum_{\alpha\in\{\mathbf{i}(\omega),\mathbf{j}(\omega)\}}\mathscr{G}^{\alpha}(\sigma^{\circ}_{m_k},\omega)\widetilde{\sigma}^{\circ,\alpha}_{m_k}(\omega)\geqslant 2\eta(\omega)(\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)-\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega))\\ +\sum_{\alpha\in\{\mathbf{i}(\omega),\mathbf{j}(\omega)\}}\mathscr{G}^{\alpha}(\sigma^{\circ}_{m_k},\omega)\sigma^{\circ,\alpha}_{m_k}(\omega)$$

for sufficiently large $k\geqslant K_2(\omega)$. Now fix an $\epsilon>0$. Note that since $\epsilon_{m_k}\downarrow 0$, there exists $N\geqslant 1$ such that $\epsilon_{m_k}<\epsilon$ for all $k\geqslant N$. Also since the regularizer $\mathbf{v}$ satisfies Condition C2. Thus, in particular, it satisfies Condition C2 with respect to the sequence of Bayesian strategies $(\sigma^{\circ}_{m_k})_{k\geqslant 1}$ over the subset $\bar{\Omega}$, we have for every $\omega\in\bar{\Omega}$,

$$|\mathbf{v}_{\epsilon_{m_k}}(\widetilde{\sigma}^{\circ}_{m_k}(\omega))-\mathbf{v}_{\epsilon_{m_k}}(\sigma^{\circ}_{m_k}(\omega))|=|\epsilon_{m_k}||\mathbf{v}_{\epsilon_{m_k}}(\widetilde{\sigma}^{\circ}_{m_k}(\omega))-\mathbf{v}_{\epsilon_{m_k}}(\sigma^{\circ}_{m_k}(\omega))|\\ \leqslant\epsilon(\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)-\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega))$$

for sufficiently large $k\geqslant\max\{N,K_2(\omega)\}$. As a result, we have that

$$\mathbf{v}_{\epsilon_{m_k}}(\widetilde{\sigma}^{\circ}_{m_k}(\omega))\leqslant\mathbf{v}_{\epsilon_{m_k}}(\sigma^{\circ}_{m_k}(\omega))+\epsilon(\sigma^{\circ,i(\omega)}_{m_k}(\omega)-\widetilde{\sigma}^{\circ,i(\omega)}_{m_k}(\omega))$$

for $k\geqslant\max\{N,K_2(\omega)\}$. Now note that for all $\omega\in\bar{\Omega}$ and $\alpha\in S\setminus\{\mathbf{i}(\omega),\mathbf{j}(\omega)\}$, we have

$$\sum_{\alpha\in S\setminus\{\mathbf{i}(\omega),\mathbf{j}(\omega)\}}\mathscr{G}^{\alpha}(\sigma^{\circ}_{m_k},\omega)\widetilde{\sigma}^{\circ,\alpha}_{m_k}(\omega)=\sum_{\alpha\in S\setminus\{\mathbf{i}(\omega),\mathbf{j}(\omega)\}}\mathscr{G}^{\alpha}(\sigma^{\circ}_{m_k},\omega)\sigma^{\circ,\alpha}_{m_k}(\omega).$$

Combining all the above observations, we arrive at the following inequality:

$$\langle\mathscr{G}(\sigma^{\circ}_{m_k},\omega),\widetilde{\sigma}^{\circ}_{m_k}(\omega)\rangle-\epsilon_{m_k}\mathbf{v}(\widetilde{\sigma}^{\circ}_{m_k}(\omega))\geqslant\eta(\omega)(\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)-\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega))\\ +\langle\mathscr{G}(\sigma^{\circ}_{m_k},\omega),\sigma^{\circ}_{m_k}(\omega)\rangle-\epsilon_{m_k}\mathbf{v}(\sigma^{\circ}_{m_k}(\omega)) \tag{D.2}$$

for large $k\geqslant\max\{N,K_2(\omega)\}$. Thus in terms of the sliced game, the above inequality can equivalently be written as

$$\langle\mathscr{G}_{\omega}(\sigma^{\circ}_{m_k}),\widetilde{\sigma}^{\circ}_{m_k}(\omega)\rangle-\epsilon_{m_k}\mathbf{v}(\widetilde{\sigma}^{\circ}_{m_k}(\omega))\geqslant\eta(\omega)(\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)-\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega))\\ +\langle\mathscr{G}_{\omega}(\sigma^{\circ}_{m_k}),\sigma^{\circ}_{m_k}(\omega)\rangle-\epsilon_{m_k}\mathbf{v}(\sigma^{\circ}_{m_k}(\omega)) \tag{D.3}$$

for large $k\geqslant\max\{N,K_2(\omega)\}$. However note that since $\sigma^{\circ}_{m_k}$ is an ϵ_{m_k}-regularized Bayesian equilibrium, we have by Lemma 27 that $\sigma^{\circ}_{m_k}(\omega)$ is an ϵ_{m_k}-perturbed equilibrium of the sliced game $\mathscr{G}_{\omega}$ for every $k\geqslant 1$. Thus we have by definition that $\sigma^{\circ}_{m_k}(\omega)=\beta_{\epsilon_{m_k}}[\sigma^{\circ}_{m_k}](\omega)$ for ξ-a.e $\omega\in\Omega$. As a result, it should be the case that

$$\langle\mathscr{G}_{\omega}(\sigma^{\circ}_{m_k}),\sigma^{\circ}_{m_k}(\omega)\rangle-\epsilon_{m_k}\mathbf{v}(\sigma^{\circ}_{m_k}(\omega))\geqslant\langle\mathscr{G}_{\omega}(\sigma^{\circ}_{m_k}),\mathbf{y}\rangle-\epsilon_{m_k}\mathbf{v}(\mathbf{y}),\quad\text{for all }\mathbf{y}\in\Delta.$$

But we observe from (D.3) that for all $\omega\in\bar{\Omega}$ the above inequality does not hold. Therefore, we have constructed a sequence of shifted Bayesian strategies $(\tilde{\sigma}^{\circ}_{m_k})_{k\geqslant 1}$ which have a greater expected payoff under the ω-sliced game $\mathscr{G}_\omega$ for every $\omega\in\bar{\Omega}$. We also note that the shifted sequence of Bayesian strategies differ from the original sequence of regularized Bayesian equilibria on the set $\bar{\Omega}$ of positive measure. Hence, arrive at a contradiction. Thus, $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)\to 0$ for some $\omega\in\bar{\Omega}$. Also since we have that $\sigma^{\circ}_{m_k}\to\sigma^{\circ}$ relative to the weak topology on Σ, we have that $\sigma^{\circ,\mathbf{i}(\omega)}(\omega)=0$. But again, we arrive at a contradiction since by our hypothesis, $\mathbf{i}(\omega)\in\operatorname{supp}(\sigma^{\circ}(\omega))$ for all $\omega\in\bar{\Omega}$. Therefore, σ° is a Nash equilibrium of the original game $\mathscr{G}$.

D.7 Regularizers satisfy Condition C1 or Condition C2

In this section, we shall show that the Shannon-Gibbs entropy satisfies Condition C1, whereas the Tsallis entropy and the Burg entropy satisfy Condition C2.

Shannon-Gibbs entropy satisfies condition C1

In this subsection, we show that the Shannon-Gibbs entropy (Example 12) satisfies Condition C1 (Definition 27). Fix $\mathbf{u}\in\mathbb{R}^n$ and $i,j\in S$ with $i\neq j$ such that $\mathbf{u}_j>\mathbf{u}_i$. In that case, it follows from (5.7) that

$$\frac{\nabla\mathbf{v}^*_{\mathbf{s},\epsilon,j}(\mathbf{u})}{\nabla\mathbf{v}^*_{\mathbf{s},\epsilon,i}(\mathbf{u})}=\frac{\exp(\epsilon^{-1}\mathbf{u}_j)}{\exp(\epsilon^{-1}\mathbf{u}_i)}=\exp(\epsilon^{-1}(\mathbf{u}_j-\mathbf{u}_i)).$$

Thus, Condition C1 is satisfied with $\lambda(\epsilon):=\exp(\epsilon^{-1}(\mathbf{u}_j-\mathbf{u}_i))$. Note that since $\mathbf{u}_j>\mathbf{u}_i$, this implies that $\lambda(\epsilon)\to\infty$ as $\epsilon\to 0$.

Tsallis entropy satisfies Condition C2

In this subsection, we show that Tsallis entropy (Example 13) satisfies Condition C2 (Definition 29) with respect to the sequence $(\sigma^{\circ}_{m_k})_{k\geqslant 1}$, where $(\sigma^{\circ}_{m_k})_{k\geqslant 1}$ is sequence of Bayesian strategies obtained in Theorem 10. Fix an arbitrary $\bar{\Omega}\subseteq\Omega$ such that $\xi(\bar{\Omega})>0$. Consider $\omega\in\bar{\Omega}$, distinct $\mathbf{i},\mathbf{j}:\bar{\Omega}\to S$ with $\lim_{k\to\infty}\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)>0$ and $\mathscr{G}^{\mathbf{j}(\omega)}(\sigma^{\circ}_{m_k},\omega)>\mathscr{G}^{\mathbf{i}(\omega)}(\sigma^{\circ}_{m_k},\omega)$ for sufficiently large k (depending on ω). The fact that $\lim_{k\to\infty}\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)>0$ implies that $\lim_{k\to\infty}\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ is bounded away from 0, in which case we can construct a corresponding sequence of Bayesian strategies $(\tilde{\sigma}^{\circ}_{m_k})_{k\geqslant 1}$, where $(\tilde{\sigma}^{\circ}_{m_k})_{k\geqslant 1}$ is a shift of $(\sigma^{\circ}_{m_k})_{k\geqslant 1}$ over $\bar{\Omega}$ from $\mathbf{i}$ to $\mathbf{j}$. First we observe that for every $\omega\in\bar{\Omega}$, the induced belief under the shifted sequence $(\tilde{\sigma}^{\circ}_{m_k})_{k\geqslant 1}$ differs from the original sequence $(\sigma^{\circ}_{m_k})_{k\geqslant 1}$ only at the co-ordinates $\mathbf{i}$ and $\mathbf{j}$ in the sense that for all $\omega\in\bar{\Omega}$, we have $\tilde{\sigma}^{\circ,k}_{m_k}(\omega)=\sigma^{\circ,k}_{m_k}(\omega)$ for all $k\in S\setminus\{\mathbf{i}(\omega),\mathbf{j}(\omega)\}$. Therefore, in order to verify Condition C2, it is enough to bound the difference in the entropy function over the strategies $\mathbf{i}(\omega)$ and $\mathbf{j}(\omega)$. Thus, by definition of Tsallis entropy, we have

$$\begin{aligned}|\mathbf{v}^{\mathbf{t}}(\tilde{\sigma}^{\circ}_{m_k}(\omega))-\mathbf{v}^{\mathbf{t}}(\sigma^{\circ}_{m_k}(\omega))|\leqslant\ &\frac{1}{q(1-q)}\Big[\underbrace{|\tilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)-\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)|}_{=:T_1^k(\omega)}+\underbrace{|\tilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^q-\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^q|}_{=:T_2^k(\omega)}\Big]\\ &+\frac{1}{q(1-q)}\Big[|\tilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)-\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)|+\underbrace{|\tilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)^q-\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)^q|}_{=:T_3^k(\omega)}\Big].\end{aligned}\tag{D.4}$$

We bound the terms $T_1^k(\omega)$, $T_2^k(\omega)$ and $T_3^k(\omega)$ separately. Note that by the very definition of $\widetilde{\sigma}^{\circ}_{m_k}$ (via Definition 28) that,

$$T_1^k(\omega) = |\widetilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) - \sigma^{\circ,\mathbf{j}}_{m_k}(\omega)| = |\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)|. \tag{D.5}$$

In order to bound the term $T_2^k(\omega)$, we first observe that $\widetilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)$ can be expressed as

$$\widetilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) = \delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega)\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega), \tag{D.6}$$

where $\delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega) := 1 + \frac{\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)}{\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)}$. Since by definition, $\delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega) > 1$, and $0 < q < 1$, we must have $\delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega)^q - 1 < \delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega) - 1$. Thus we arrive at

$$\begin{aligned}
T_2^k(\omega) &= |\widetilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^q - \sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^q| \\
&= |(\delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega)^q - 1)\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^q| \\
&\leqslant |(\delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega) - 1)\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^q| \\
&= \frac{\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)}{\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^{1-q}}.
\end{aligned}$$

Thus, in order to bond the term $T_2^k(\omega)$, we need to lower bound the term $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^{1-q}$. Note that since $\sigma^{\circ}_{m_k}(\omega)$ is a regularized equilibrium of the sliced game $\mathscr{G}_\omega$, we have that

$$\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) = \left[\frac{\epsilon_{m_k}}{1-q}\right]^{1/(1-q)} \frac{1}{(\theta^{\mathbf{t}}_{\epsilon}(\sigma^{\circ}_{m_k}, \omega) - \mathscr{G}^{\mathbf{j}(\omega)}_{\omega}(\sigma^{\circ}_{m_k}))^{1/(1-q)}} \tag{D.7}$$

Thus, letting $\eta(\omega) \downarrow 0$ in (D.1), we have that $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) \geqslant \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ for all k sufficiently large (depending on ω). As a result, we have that $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)^{q-1} \leqslant \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)^{q-1}$ for sufficiently large k (depending on ω). Note that since $\lim_{k\to\infty} \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) > 0$, there exists $c_{\mathbf{i}(\omega)}(\omega) > 0$ such that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \geqslant \frac{c_{\mathbf{i}(\omega)}(\omega)}{2}$ for sufficiently large k (depending on ω). This, in conjunction with the fact that $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) \geqslant \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ for all k sufficiently large (depending on ω) implies that $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) \geqslant \frac{c_{\mathbf{i}(\omega)}(\omega)}{2}$ for all k sufficiently large (depending on ω). Combining all the above observations, we arrive at

$$T_2^k(\omega) \leqslant \frac{2^{1-q}}{c_{\mathbf{i}(\omega)}(\omega)^{1-q}}(\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)). \tag{D.8}$$

Finally, we proceed to bound the term T_3. Recall that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \geqslant \frac{c_{\mathbf{i}(\omega)}(\omega)}{2}$ for sufficiently large k (depending on ω). Now, by an application of the mean-value theorem, we have

$$\begin{aligned}
T_3^k(\omega) &= |\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)^q - \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)^q| \\
&\leqslant \frac{q2^{1-q}}{c_{\mathbf{i}(\omega)}(\omega)^{1-q}}|\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)|.
\end{aligned}$$

Combining all the inequalities obtained above in (D.5), (D.8) and (D.9) we have for sufficiently large

k that,

$$
\begin{aligned}
|\mathbf{v}^{\mathbf{t}}(\widetilde{\sigma}^{\circ}_{m_k}(\omega)) - \mathbf{v}^{\mathbf{t}}(\sigma^{\circ}_{m_k}(\omega))| &\leqslant \Big(2 + \frac{2^{1-q}}{c_{\mathbf{i}(\omega)}(\omega)^{1-q}} + \frac{q2^{1-q}}{c_{\mathbf{i}(\omega)}(\omega)^{1-q}}\Big)|\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)| \\
&= \Big(\frac{2c_{\mathbf{i}(\omega)}(\omega)^{1-q} + 2^{1-q}(1+q)}{q(1-q)c_{\mathbf{i}(\omega)}(\omega)^{1-q}}\Big)|\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)|.
\end{aligned}
$$

Thus setting $\kappa_{\mathbf{i}(\omega)}(\omega) := \frac{2c_{\mathbf{i}(\omega)}(\omega)^{1-q}+2^{1-q}(1+q)}{q(1-q)c_{\mathbf{i}(\omega)}(\omega)^{1-q}}$, and observing the fact that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) > \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$, proves the fact that the Tsallis entropy satisfies Condition C2.

Burg entropy satisfies Condition C2

In this subsection, we show that Burg entropy (Example 5.3.1) satisfies Condition C2. Fix $\omega \in \Omega$, distinct $\mathbf{i}, \mathbf{j} : \bar{\Omega} \to S$ with $\lim_{k\to\infty} \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) > 0$ and $\mathscr{G}^{\mathbf{j}(\omega)}(\sigma^{\circ}_{m_k}, \omega) > \mathscr{G}^{\mathbf{i}(\omega)}(\sigma^{\circ}_{m_k}, \omega)$ for sufficiently large k (depending on ω). As in the case of the proof for Tsallis entropy, it is enough to bound the Burg entropy at co-ordinates $\mathbf{i}(\omega)$ and $\mathbf{j}(\omega)$.

We proceed similar to the case of Tsallis entropy. Note that by definition of Burg entropy, we have

$$
|\mathbf{v}^{\mathbf{b}}(\widetilde{\sigma}^{\circ}_{m_k}(\omega)) - \mathbf{v}^{\mathbf{b}}(\sigma^{\circ}_{m_k}(\omega))| \leqslant \underbrace{|\log \widetilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) - \log \sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)|}_{=:U_1^k(\omega)} + \underbrace{|\log \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \log \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)|}_{=:U_2^k(\omega)}. \tag{D.9}
$$

We bound the terms $U_1^k(\omega), U_2^k(\omega)$ separately. We first observe that $\widetilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)$ can be expressed as

$$
\widetilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) = \delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega)\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega), \tag{D.10}
$$

where $\delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega) := 1 + \frac{\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)}{\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)}$. Since $\log(1+x) < x$ for all $x > 0$, we have that $\log \delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega) < \frac{\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)}{\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)}$. Now, by definition of Burg entropy, we have

$$
\begin{aligned}
U_1^k(\omega) &= |\log \widetilde{\sigma}^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) - \log \sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)| \\
&= |\log\big[\delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega)\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)\big] - \log \sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)| \\
&= |\log \delta^{\mathbf{i}(\omega)\mathbf{j}(\omega)}_{m_k}(\omega)| \\
&< \frac{\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)}{\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)}
\end{aligned}
$$

Thus, similar to Tsallis entropy, we need to lower bound the term $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega)$. Note that since $\sigma^{\circ}_{m_k}(\omega)$ is a regularized equilibrium of the sliced game $\mathscr{G}_{\omega}$, we have that

$$
\sigma^{\circ,\mathbf{j}}_{m_k}(\omega) = \frac{\epsilon_{m_k}}{(\theta^{\mathbf{b}}_{\epsilon}(\sigma^{\circ}_{m_k}, \omega) - \mathscr{G}^{\mathbf{j}(\omega)}_{\omega}(\sigma^{\circ}_{m_k}))}. \tag{D.11}
$$

Thus, letting $\eta(\omega) \downarrow 0$ in (D.1), we have that $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) \geqslant \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ for all k sufficiently large (depending on ω). Note that since $\lim_{k\to\infty} \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) > 0$, there exists $\widehat{c}_{\mathbf{i}(\omega)}(\omega) > 0$ such that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \geqslant \frac{\widehat{c}_{\mathbf{i}(\omega)}(\omega)}{2}$ for sufficiently large k (depending on ω). This, in conjunction with the fact that $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) \geqslant \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ for all k sufficiently large (depending on ω) implies that $\sigma^{\circ,\mathbf{j}(\omega)}_{m_k}(\omega) \geqslant \frac{\widehat{c}_{\mathbf{i}(\omega)}(\omega)}{2}$ for all k sufficiently large (depending on ω). Combining the above observations, we arrive at

$$U^k_1(\omega) \leqslant \frac{2}{\widehat{c}_{\mathbf{i}(\omega)}(\omega)} (\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)). \tag{D.12}$$

We now proceed to bound the term $U^k_2(\omega)$. Notice that the mapping $x \mapsto \log x$ is Lipschitz on $[a,\infty)$, where $a > 0$. The fact that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ is bounded away from 0, implies that the shifted sequence of Bayesian strategies can be chosen in such a way that $\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$ is bounded away from 0. More precisely, we can construct a shifted sequence of Bayesian strategies such that $\widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) \geqslant \frac{\widehat{c}_{\mathbf{i}(\omega)}(\omega)}{4}$. Thus, by an application of mean value theorem, we have

$$\begin{aligned} U^k_2(\omega) &= |\log \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \log \sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)| \\ &\leqslant \frac{4}{\widehat{c}_{\mathbf{i}(\omega)}(\omega)} (\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)). \end{aligned} \tag{D.13}$$

Therefore, combining (D.12) and (D.13), we arrive at

$$\begin{aligned} |\mathbf{v}^{\mathbf{b}}(\widetilde{\sigma}^{\circ}_{m_k}(\omega)) - \mathbf{v}^{\mathbf{b}}(\sigma^{\circ}_{m_k}(\omega))| &\leqslant U^k_1(\omega) + U^k_2(\omega) \\ &\quad \Big(\frac{2}{\widehat{c}_{\mathbf{i}(\omega)}(\omega)} + \frac{4}{\widehat{c}_i(\omega)}\Big)(\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)) \\ &\leqslant \frac{6}{\widehat{c}_{\mathbf{i}(\omega)}(\omega)} (\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) - \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)). \end{aligned}$$

Thus, setting $\widehat{\kappa}_{\mathbf{i}(\omega)}(\omega) := \frac{6}{\widehat{c}_{\mathbf{i}(\omega)}(\omega)}$ and observing the fact that $\sigma^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega) > \widetilde{\sigma}^{\circ,\mathbf{i}(\omega)}_{m_k}(\omega)$, proves the fact that the Burg entropy satisfies Condition C2.

D.8 Proof of Theorem 17

We use the infinite dimensional Picard-Lindeloff Theorem (see Zeidler (1986) for details) to prove Theorem 17.

Theorem 26 (Picard-Lindelöff Theorem, Zeidler (1986)) *Let $\mathbb{Y}$ be a Banach space. Let $y_0 \in \mathbb{Y}$, $t_0 \in \mathbb{R}$, and*

$$\mathscr{Q}_b := \{(t,y) \in \mathbb{R} \times \mathbb{Y} \mid |t - t_0| \leqslant a \text{ and } \|y - y_0\| \leqslant b\}$$

for fixed $a, b > 0$. Suppose $f : \mathscr{Q}_b \to \mathbb{Y}$ is continuous and that

$$\|f(t,x) - f(t,y)\| \leqslant L\|x - y\| \text{ for all } (t,x),(t,y) \in \mathscr{Q}_b, \text{ and}$$

$$\|f(t,x)\| < M \text{ for all } (t,x) \in \mathscr{Q}_b$$

where $L \geqslant 0$ and $M > 0$ are fixed. Choose c such that $0 < c < a$ and $Mc < b$. Then the initial value problem

$$\dot{x}(t) = f(t, x(t)), \qquad x(t_0) = y_0,$$

where $x : [t_0 - c, t_0 + c] \to \mathbb{Y}$, has exactly one continuously differentiable solution on the interval $[t_0 - c, t_0 + c]$.

Now we proceed with the proof of Theorem 17. To prove Theorem 17, it is by Theorem 26 enough to show that the mapping $\sigma \mapsto \beta_\epsilon(\sigma)$ is Lipschitz continuous with respect to the strong topology on Σ. We show this in the following lemma.

Lemma 28 *The mapping $\sigma \mapsto \beta_\epsilon(\sigma)$ is Lipschitz continuous with respect to the strong norm on Σ.*

Proof 71 To prove the lemma, we need to show that there exists a constant $\kappa > 0$, such that

$$\left\|\left|\beta_\epsilon(\sigma) - \beta_\epsilon(\rho)\right|\right\| \leqslant \bar{\kappa} \left\|\left|\sigma - \rho\right|\right\|, \quad \text{for all } \sigma, \rho \in \Sigma. \tag{D.14}$$

Fix $\sigma, \rho \in \Sigma$. By (5.5) and the definition of strong norm, we have that

$$\begin{aligned}
\left\|\left|\beta_\epsilon(\sigma) - \beta_\epsilon(\rho)\right|\right\| &= \int_\Omega \|\beta_\epsilon[\sigma](\omega) - \beta_\epsilon[\rho](\omega)\| \xi(d\omega) \\
&= \int_\Omega \|\nabla \mathbf{v}_\epsilon^*(\mathscr{G}(\sigma, \omega)) - \nabla \mathbf{v}_\epsilon^*(\mathscr{G}(\rho, \omega))\| \xi(d\omega) \\
&\leqslant \frac{1}{\epsilon\gamma} \int_\Omega \|\mathscr{G}(\sigma, \omega) - \mathscr{G}(\rho, \omega)\| \xi(d\omega) \\
&\leqslant \frac{\kappa}{\epsilon\gamma} \left\|\left|\sigma - \rho\right|\right\|.
\end{aligned}$$

The second last inequality occurs since $\nabla \mathbf{v}_\epsilon^*$ is $(\epsilon\gamma)^{-1}$-Lipschitz, and the last inequality occurs since the Bayesian population game $\mathscr{G}$ is κ-strongly Lipschitz. Finally we have that

$$\left\|\left|\beta_\epsilon(\sigma) - \beta_\epsilon(\rho)\right|\right\| \leqslant \bar{\kappa} \left\|\left|\sigma - \rho\right|\right\|, \quad \text{for all } \sigma, \rho \in \Sigma.$$

Therefore (D.14) is satisfied with $\bar{\kappa} = \frac{\kappa}{\epsilon\gamma}$. This completes the proof of Lemma 36.

We now apply Grönwall's Lemma (see Zeidler (1986) for details) to prove the continuity of the semiflow, that is, part (ii) of Theorem 17. We state the lemma below.

Lemma 29 (Grönwalls's Lemma, Zeidler (1986)) *Let $u, v : [t_0 - a, t_0 + a] \to \mathbb{V}$ be maps into the subspace $\mathbb{V}$ of a Banach space $\mathbb{Y}$. Suppose $f : [t_0 - a, t_0 + a] \times \mathbb{V} \to \mathbb{Y}$ is continuous and Lipschitz continuous with respect to u, that is,*

$$\|f(t, u) - f(t, v)\|_{\mathbb{Y}} \leqslant L\|u - v\|_{\mathbb{V}} \text{ for all } t \in [t_0 - a, t_0 + a] \text{ and } u, v \in \mathbb{V} \text{ and fixed } L > 0.$$

Then $\|u(t) - v(t)\|_{\mathbb{V}} \leqslant e^{L|t - t_0|}\|u(0) - v(0)\|$ for all $t \in [t_0 - a, t_0 + a]$.

We apply Lemma 29 with $\mathbb{Y} = \widehat{\Sigma}, \mathbb{V} = \Sigma$, and $f = \beta_\epsilon$. By Lemma 36, the conditions of Lemma 29 are satisfied. Therefore, there exists $k > 0$ such that for any two solutions $(\sigma_t)_{t \geqslant 0}$ and $(\widetilde{\sigma}_t)_{t \geqslant 0}$ to the Generalized Dynamic with initial conditions σ_0 and $\widetilde{\sigma}_0$, respectively, we have $\|\sigma_t - \widetilde{\sigma}_t\| \leqslant e^{kt} \|\sigma_0 - \widetilde{\sigma}_0\|$. This establishes continuity of the semiflow of the Generalized Dynamic. This concludes the proof of Theorem 17.

D.9 Proof of Lemma 5

To prove part (i) of the theorem, we require the following lemma.

Lemma 30 (Balder et al. (1994)) *A subset $\mathbb{K}$ of $\mathscr{L}^1(\mu,\mathbb{X})$ is relatively norm compact if and only if the following conditions are satisfied:*

1. *The subset $\mathbb{K}$ is relatively weakly compact.*
2. *The subset $\mathbb{K}$ satisfies small Bocce oscillation.*
3. *The subset $\{avg(f;A): f\in\mathbb{K}\}$ is relatively norm compact in $\mathbb{X}$, for every $A\in\mathscr{B}_\Omega$, such that $\mu(A)>0$.*

Throughout the proof, we fix a compact set $\Xi\subseteq\Omega$ with $\xi(\Xi)=1$. To verify condition (i) of Lemma 37, note that for every $\alpha>0$ and $M\subseteq\Delta^\circ$, we have $\Sigma_M^\alpha(\Xi)$ is a subset of Σ. Therefore, the fact that Σ is relatively weakly compact implies that $\Sigma_M^\alpha(\Xi)$ is also relatively weakly compact.

We now proceed to verify condition (ii) of Lemma 37, that is, we need to show that $\Sigma_M^\alpha(\Xi)$ satisfies small Bocce oscillation. To this end, fix $\eta>0$. Since $\Sigma_M^\alpha(\Xi)$ is a family of Bayesian strategies which are α-Lipschitz on Ξ, it follows that $\Sigma_M^\alpha(\Xi)$ is equicontinuous on Ξ. Therefore, by the definition, there exists $\delta>0$, such that $\omega,\bar\omega\in\Xi$ and $d_\Omega(\omega,\bar\omega)<\delta$ imply that

$$\|\sigma(\omega)-\sigma(\bar\omega)\|<\eta,\quad \text{for all } \sigma\in\Sigma_M^\alpha(\Xi). \tag{D.15}$$

Fix $\omega\in\Xi$, and let $B(\omega;\delta)$ denote the open ball of radius δ about ω. Then, for every $\sigma\in\Sigma_M^\alpha(\Xi)$, we have that $\|\sigma(\omega)-\sigma(\bar\omega)\|<\eta$, for all $\bar\omega\in B(\omega;\delta)$. We now proceed to show that Bocce-osc$(\sigma,B(\omega;\delta))<\eta$, for every $\sigma\in\Sigma_M^\alpha(\Xi)$. Fix $\sigma\in\Sigma_M^\alpha(\Xi)$. Note that for every $\bar\omega\in B(\omega;\delta)$, we have $\|\sigma(\omega)-\sigma(\bar\omega)\|<\eta$. This implies that

$$\begin{aligned}\|\sigma(\omega)-\text{avg}(\sigma,B(\omega;\delta))\| &= \left\|\sigma(\omega)-\frac{1}{\xi(B(\omega;\delta))}\int_{B(\omega;\delta)}\sigma(\bar\omega)\xi(d\bar\omega)\right\|\\ &= \frac{1}{\xi(B(\omega;\delta))}\left\|\int_{B(\omega;\delta)}\sigma(\omega)\xi(d\bar\omega)-\int_{B(\omega;\delta)}\sigma(\bar\omega)\xi(d\bar\omega)\right\|\\ &\leqslant \frac{1}{\xi(B(\omega;\delta))}\int_{B(\omega;\delta)}\|\sigma(\omega)-\sigma(\bar\omega)\|\xi(d\bar\omega)\\ &<\eta.\end{aligned}$$

This implies Bocc-osc$(\sigma,B(\omega,\delta))<\eta$. Since $\Sigma_M^\alpha(\Xi)$ is equicontinuous on Ξ, we have from (D.15), that

$$\text{Bocce-osc}(\sigma;B(\omega;\delta))<\eta,\quad \text{for all } \sigma\in\Sigma_M^\alpha(\Xi).$$

Note that $(B(\omega;\delta))_{\omega\in\Omega}$ is an open cover of Ω, and hence an open cover of Ξ. Since Ξ is compact, it follows that there exist $\omega_1,\ldots,\omega_p\in\Xi$, such that $\Xi\subseteq\cup_{1\leqslant i\leqslant p}B(\omega_i;\delta)$. Suppose that the collection $(B(\omega_i;\delta))_{1\leqslant i\leqslant p}$ is disjoint. Then,

$$\{B(\omega_i;\delta): 1\leqslant i\leqslant p\}\cup(\Omega\setminus\Xi)$$

forms a finite measurable partition of Ω. However, if it happens to be the case that the collection $(B(\omega_i;\delta))_{1\leqslant i\leqslant p}$ is not disjoint, we disjointify the collection and extract a collection $(G(\omega_i;\delta))_{1\leqslant i\leqslant p}$ such that $\cup_{1\leqslant i\leqslant p}G(\omega;\delta)=\cup_{1\leqslant i\leqslant p}B(\omega_i;\delta)$ and that $G(\omega_i;\delta)\cap G(\omega_j;\delta)=\emptyset$ for all $i\neq j$. In such a case, we have that

$$\{G(\omega_i;\delta): 1\leqslant i\leqslant p\}\cup(\Omega\setminus\Xi)$$

forms a measurable partition of Ω. The fact that the elements of the new collection $(G(\omega_i;\delta))_{1\leqslant i\leqslant p}$ are measurable subsets follows since both open sets and closed sets are measurable, and the Borel sigma-algebra is closed under intersection of finitely many open and closed sets. Therefore we finally have that

$$\sum_{1\leqslant i\leqslant p} \xi(G(\omega_i;\delta))\text{Bocc-osc}(\sigma;G(\omega_i;\delta))<\eta.$$

This proves that $\Sigma_M^\alpha(\Xi)$ has small Bocce oscillation thereby verifying condition (ii) of Lemma 37.

Finally we verify condition (iii) of Lemma 37. Fix $A\in\mathscr{B}_\Omega$, such that $\xi(A)>0$. By the definition of Bayesian strategies, we have that $\{\text{avg}(\sigma;A):\sigma\in\Sigma_M^\alpha(\Xi)\}\subseteq\Delta$. Since Δ is compact in the Euclidean topology, we have that $\{\text{avg}(\sigma;A):\sigma\in\Sigma_M^\alpha(\Xi)\}$ is relatively norm compact in $\mathbb{R}^n$. Since $A\in\mathscr{B}_\Omega$ is arbitrary, Condition (iii) of Lemma 37 is satisfied. This concludes the proof of part (i) of Theorem 5.

We now proceed to prove part (ii) of Lemma 5. We first observe that the game $\mathscr{G}$ is uniformly bounded in the sense that there exists $c>0$ such that for every $\sigma\in\Sigma$, $\|\mathscr{G}(\sigma,\omega)\|\leqslant c$ for ξ-a.e $\omega\in\Omega$. Let $B_c:=\{\mathbf{u}\in\mathbb{R}^n:\|\mathbf{u}\|\leqslant c\}$. Since the mapping $\mathbf{u}\mapsto\nabla\mathbf{v}_\epsilon^*(\mathbf{u})$ is continuous and B_c is a compact set, we have that the image of B_c is a compact set. Also since the regularized best response maps payoff-vectors to the interior of Δ, we have $\nabla\mathbf{v}_\epsilon^*(\mathbf{u})\in\Delta^\circ$ for every $\mathbf{u}\in\mathbb{R}^n$, we have that the image of B_c is a compact subset of Δ°. Thus, for a given $c>0$, consider the set B_c and define M_c to be the image of B_c under the map $\mathbf{u}\mapsto\nabla\mathbf{v}_\epsilon^*(\mathbf{u})$. Therefore, we have that $\beta_\epsilon[\sigma]\in\Sigma_{M_c}$ whenever $\sigma\in\Sigma$. This implies, in particular, that $\beta_\epsilon[\sigma]\in\Sigma_{M_c}$ whenever $\sigma\in\Sigma_{M_c}$. Also, by assumption of lemma, we have for every $\sigma\in\Sigma$, the map $\omega\mapsto\mathscr{G}(\sigma,\omega)$ is λ-Lipschitz for ξ-a.e $\omega\in\Xi$. As a result, fixing $\omega,\bar{\omega}\in\Xi$, we have

$$\begin{aligned}\|\beta_\epsilon[\sigma](\omega)-\beta_\epsilon[\sigma](\bar{\omega})\|&=\|\nabla\mathbf{v}_\epsilon^*(\mathscr{G}(\sigma,\omega))-\nabla\mathbf{v}_\epsilon^*(\mathscr{G}(\sigma,\bar{\omega}))\|\\&\leqslant\frac{1}{\epsilon\gamma}\|\mathscr{G}(\sigma,\omega)-\mathscr{G}(\sigma,\bar{\omega})\|\\&\leqslant\frac{\lambda}{\epsilon\gamma}d_\Omega(\omega,\bar{\omega})\\&\leqslant\alpha d_\Omega(\omega,\bar{\omega}).\end{aligned}$$

Combining the above observations, we have that $\beta_\epsilon[\sigma]\in\Sigma_{M_c}^\alpha(\Xi)$ whenever $\sigma\in\Sigma$. This again, in particular, implies that $\beta_\epsilon[\sigma]\in\Sigma_{M_c}^\alpha(\Xi)$ whenever $\sigma\in\Sigma_{M_c}^\alpha(\Xi)$. Thus, we have proved that given $c>0$, there exists a compact set $M_c\subseteq\Delta^\circ$ such that $\beta_\epsilon[\sigma]\in\Sigma_{M_c}^\alpha(\Xi)$ whenever $\sigma\in\Sigma_{M_c}^\alpha(\Xi)$, provided the Lipschitz coefficient of the Bayesian strategy $\alpha\geqslant\frac{\lambda}{\epsilon\gamma}$. Hence $\Sigma_{M_c}^\alpha(\Xi)$ is forward invariant under the RBBR learning dynamic. This completes the proof of part (ii) of Lemma 5.

D.10 Proof of Lemma 6

In order to prove the lemma, we need to show that

$$\int_\Omega\langle\widetilde{\mathscr{G}}(\sigma,\omega),\dot{\sigma}(\omega)\rangle\xi(d\omega)\geqslant 0\quad\text{for all }\sigma\in\Sigma,$$

with equality only if $\dot{\sigma}=0$. We first prove the inequality. We require the following lemma in this regard.

Lemma 31 *For every $\sigma \in \Sigma_M^\alpha$, the following equality holds:*

$$\int_\Omega \langle \nabla^{\mathbf{G}} \widetilde{\boldsymbol{v}}_\epsilon[\beta_\epsilon(\sigma)](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) = \int_\Omega \langle \mathscr{G}(\sigma,\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega).$$

Proof 72 To prove the lemma, fix $\sigma \in \Sigma_M^\alpha$ and $\omega \in \Omega$ and consider a mapping $\Phi_{\sigma,\omega}^\epsilon : \Delta^\circ \to \mathbb{R}$ defined as

$$\Phi_{\sigma,\omega}^\epsilon(\mathbf{y}) := \langle \mathbf{y}, \mathscr{G}(\sigma,\omega)\rangle - \epsilon \mathbf{v}(\mathbf{y}), \quad \text{for all } \mathbf{y} \in \Delta^\circ. \tag{D.16}$$

Now note that by the definition of regularized Bayesian best response map and the fact that $\mathbf{v}$ is strongly convex, we have that $\beta_\epsilon[\sigma](\omega)$ is the unique maximizer of $\Phi_{\sigma,\omega}^\epsilon$. Therefore, we must have $d\Phi_{\sigma,\omega}^\epsilon[\beta_\epsilon[\sigma](\omega)](\mathbf{z}_0) = 0$ for all $\mathbf{z}_0 \in \mathbb{R}_0^n$. As a result, for $\mathbf{z}_0 \in \mathbb{R}_0^n$, we have

$$\begin{aligned}
0 &= d\Phi_{\sigma,\omega}^\epsilon[\beta_\epsilon[\sigma](\omega)](\mathbf{z}_0) \\
&= \lim_{\eta\to 0} \frac{\Phi_{\sigma,\omega}^\epsilon(\beta_\epsilon[\sigma](\omega) + \eta \mathbf{z}_0) - \Phi_{\sigma,\omega}^\epsilon(\beta_\epsilon[\sigma](\omega))}{\eta} \\
&= \langle \mathbf{z}_0, \mathscr{G}(\sigma,\omega)\rangle - \lim_{\eta\to 0} \frac{\mathbf{v}_\epsilon(\beta_\epsilon[\sigma](\omega) + \eta \mathbf{z}_0) - \mathbf{v}_\epsilon(\beta_\epsilon[\sigma](\omega))}{\eta} \\
&= \langle \mathbf{z}_0, \mathscr{G}(\sigma,\omega)\rangle - d\mathbf{v}_\epsilon[\beta_\epsilon[\sigma](\omega)](\mathbf{z}_0).
\end{aligned} \tag{D.17}$$

Thus, we have $d\mathbf{v}_\epsilon[\beta_\epsilon[\sigma](\omega)](\mathbf{z}_0) = \langle \mathscr{G}(\sigma,\omega), \mathbf{z}_0\rangle$ for all $\mathbf{z}_0 \in \mathbb{R}_0^n$. Fix $\sigma \in \Sigma_M^\alpha$ and $\sigma_0 \in \Sigma_0$. Recall from the definition of $\widetilde{\mathbf{v}}$ that for fixed $\eta > 0$, we have

$$\widetilde{\mathbf{v}}_\epsilon(\beta_\epsilon[\sigma] + \eta\sigma_0) = \int_\Omega \mathbf{v}_\epsilon(\beta_\epsilon[\sigma](\omega) + \eta\sigma_0(\omega))\xi(d\omega),$$

and that

$$\widetilde{\mathbf{v}}_\epsilon(\beta_\epsilon[\sigma]) = \int_\Omega \mathbf{v}_\epsilon(\beta_\epsilon[\sigma](\omega))\xi(d\omega).$$

Since $\sigma \in \Sigma_M^\alpha$, we have that $\mathbf{v}(\sigma(\omega))$ is uniformly bounded for ξ-a.e $\omega \in \Omega$, in which case, we have

$$d\widetilde{\mathbf{v}}_\epsilon[\beta_\epsilon(\sigma)](\sigma_0) = \int_\Omega d\mathbf{v}_\epsilon(\beta_\epsilon[\sigma](\omega))(\sigma_0(\omega))\xi(d\omega). \tag{D.18}$$

However using (D.17) and the definition of Gâteaux derivative, we have

$$\begin{aligned}
\int_\Omega \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_\epsilon[\beta_\epsilon(\sigma)](\omega)(\sigma_0(\omega))\xi(d\omega) &= \int_\Omega \langle \sigma_0(\omega), \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_\epsilon[\beta_\epsilon(\sigma)](\omega)\rangle \xi(d\omega) \\
&= \int_\Omega d\mathbf{v}_\epsilon(\beta_\epsilon[\sigma](\omega))(\sigma_0(\omega))\xi(d\omega) \\
&= \int_\Omega \langle \mathscr{G}(\sigma,\omega), \sigma_0(\omega)\rangle \xi(d\omega).
\end{aligned}$$

Note that the above equality holds true for arbitrary $\sigma_0 \in \Sigma_0$. Since, $\dot{\sigma} = \beta_\epsilon(\sigma) - \sigma \in \Sigma_0$, we have that

$$\int_\Omega \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_\epsilon[\beta_\epsilon(\sigma)](\omega)(\dot{\sigma}(\omega))\xi(d\omega) = \int_\Omega \langle \mathscr{G}(\sigma,\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega). \tag{D.19}$$

This completes the proof of the Lemma 31.

We now proceed with the proof of Lemma 6. Note that, as a consequence of Lemma 31, we have the following equality:

$$\begin{aligned}\int_\Omega \langle \widetilde{\mathscr{G}}(\sigma,\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) &= \int_\Omega \langle \mathscr{G}(\sigma,\omega) - \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}[\sigma](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) \\ &= \int_\Omega \langle \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}[\beta_\epsilon(\sigma)](\omega) - \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}[\sigma](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega).\end{aligned}$$

Therefore, in order to prove Theorem 18, firstly we need to show that

$$\int_\Omega \langle \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}[\beta_\epsilon(\sigma)](\omega) - \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}[\sigma](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) \geqslant 0. \tag{D.20}$$

Since $\mathbf{v}$ is strongly convex (and hence convex), we have that $\widetilde{\mathbf{v}}$ is also convex. As a result, we have from Rockafellar (1970), that for every $\sigma, \rho \in \Sigma$,

$$\begin{aligned}\widetilde{\mathbf{v}}(\rho) &\geqslant \widetilde{\mathbf{v}}(\sigma) + \int_\Omega \langle \rho - \sigma, \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}(\sigma)\rangle(\omega)\xi(d\omega) \\ &= \widetilde{\mathbf{v}}(\sigma) + \int_\Omega \langle \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}[\sigma](\omega), \rho(\omega) - \sigma(\omega)\rangle \xi(d\omega).\end{aligned}$$

Putting $\rho = \beta_\epsilon[\sigma]$ in the above inequality we have that

$$\widetilde{\mathbf{v}}(\beta_\epsilon[\sigma]) \geqslant \widetilde{\mathbf{v}}(\sigma) + \int_\Omega \langle \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}[\sigma](\omega), \beta_\epsilon[\sigma](\omega) - \sigma(\omega)\rangle \xi(d\omega). \tag{D.21}$$

Now interchanging ρ and σ, we have that

$$\widetilde{\mathbf{v}}(\sigma) \geqslant \widetilde{\mathbf{v}}(\beta_\epsilon[\sigma]) + \int_\Omega \langle \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}(\beta_\epsilon[\sigma])(\omega), \sigma(\omega) - \beta_\epsilon[\sigma](\omega)\rangle \xi(d\omega). \tag{D.22}$$

Combining inequalities (D.21) and (D.22) above, we arrive at the required inequality (D.20). Hence this completes the proof of Lemma 6.

It now remains to show that equality in (D.20) holds only if $\beta_\epsilon(\sigma) = \sigma$, that is, $\dot{\sigma} = 0$. Since $\mathbf{v}$ is a decomposable regularizer, there exists a $\mathscr{C}^1$-function $\theta : (0,1) \to \mathbb{R}$ such that $\mathbf{v}(\mathbf{x}) = \sum_{1 \leqslant i \leqslant n} \theta(\mathbf{x}_i)$, for all $\mathbf{x} \in \Delta^\circ$. We first require the following lemma.

Lemma 32 *Suppose that $\boldsymbol{v}$ is a regularizer decomposable via the function $\theta : (0,1) \to \mathbb{R}$. Then for every $\sigma \in \Sigma_M^\alpha$, we have*

$$d\widetilde{\boldsymbol{v}}_\epsilon[\sigma](\sigma_0) := \epsilon \sum_{i=1}^n \int_\Omega \theta'(\sigma^i(\omega))\sigma_0^i(\omega)\xi(d\omega), \quad \textit{for all } \sigma_0 \in \Sigma_0.$$

Proof 73 For $\mathbf{x} \in \Delta^\circ$ and $\mathbf{z}_0 \in \mathbb{R}_0^n$, we have

$$\begin{aligned} d\mathbf{v}_\epsilon[\mathbf{x}](\mathbf{z}_0) &= \lim_{\eta\to 0} \frac{\mathbf{v}_\epsilon(\mathbf{x}+\eta\mathbf{z}_0) - \mathbf{v}_\epsilon(\mathbf{x})}{\eta} \\ &= \epsilon \lim_{\eta\to 0} \frac{\sum_{1\leqslant i\leqslant n}\theta(\mathbf{x}_i+\eta\mathbf{z}_i^0) - \sum_{1\leqslant i\leqslant n}\theta(\mathbf{x}_i)}{\eta} \\ &= \epsilon\sum_{i=1}^n \mathbf{z}_0^i\theta'(\mathbf{x}_i). \end{aligned}$$

It now follows from (D.18) and the previous arguments that for $\sigma \in \Sigma_M^\alpha$ and $\sigma_0 \in \Sigma_0$, we have

$$\begin{aligned} d\widetilde{\mathbf{v}}_\epsilon[\sigma](\sigma_0) &= \int_\Omega d\mathbf{v}_\epsilon(\sigma(\omega))(\sigma_0(\omega))\xi(d\omega) \\ &= \epsilon\sum_{i=1}^n\int_\Omega \theta'(\sigma^i(\omega))\sigma_0^i(\omega)\xi(d\omega). \end{aligned}$$

This completes the proof of Lemma 32.

We now proceed with the proof of part (i) of Lemma 6 in the context of decomposable regularizers. By Lemma 32, we have that the Gâteaux gradient $\nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon(\beta_\epsilon[\sigma])$ at the type ω is identified by the vector $\epsilon(\theta'(\beta_\epsilon^1[\sigma](\omega)),\ldots,\theta'(\beta_\epsilon^n[\sigma](\omega)))$. Similarly, we have that $\nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon[\sigma](\omega) = \epsilon(\theta'(\sigma^1(\omega)),\ldots,\theta'(\sigma^n(\omega)))$ for all $\omega\in\Omega$. This in particular implies that

$$\begin{aligned} &\int_\Omega \langle \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}[\beta_\epsilon(\sigma)](\omega) - \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}[\sigma](\omega), \dot{\sigma}(\omega)\rangle\xi(d\omega) \\ &\qquad = \epsilon\sum_{i=1}^n\int_\Omega [\theta'(\beta_\epsilon^i[\sigma](\omega)) - \theta'(\sigma^i(\omega))][\beta_\epsilon^i[\sigma](\omega) - \sigma^i(\omega)]\xi(d\omega). \end{aligned} \tag{D.23}$$

By assumption of the lemma, we have that the function θ' is strictly increasing on $(0,1)$. Thus, we have that the RHS of the equality (D.23) is positive for ξ-a.e $\omega\in\Omega$. This implies that the RHS of (D.23) equals 0 iff $\beta_\epsilon[\sigma]=\sigma$, that is equality holds iff $\dot{\sigma}=0$.

Finally we proceed with the proof of part (ii) of Lemma 6. We note by Lemma 5, that there exists a compact set $M\subseteq\Delta^\circ$ such that $\Sigma_M^\alpha(\Xi)$ is forward invariant under the RBBR dynamic. Thus, we can find $\delta\in(0,\frac{1}{2})^n$ such that $M\subseteq[\delta^1,1-\delta^1]\times\cdots\times[\delta^n,1-\delta^n]$.[6] Now suppose that $\sigma\in\Sigma_M^\alpha(\Xi)$. By forward invariance, $\beta_\epsilon[\sigma]\in\Sigma_M^\alpha(\Xi)$. Also by assumption in part (ii) of the lemma, we have that θ' is strictly increasing on the interval $[\min_{1\leqslant i\leqslant n}\delta^i, 1-\min_{1\leqslant i\leqslant n}\delta^i]$. Thus, in such a scenario, the RHS of the equality (D.23) is positive for ξ-a.e $\omega\in\Omega$. This implies that the RHS of (D.23) equals 0 iff $\beta_\epsilon[\sigma]=\sigma$, that is equality holds iff $\dot{\sigma}=0$. Thus, the RBBR learning dynamic satisfies positive correlation with respect to the virtual payoff $\widetilde{\mathscr{G}}$.

[6]We clarify the reader that the compact set M and $\delta\in(0,\frac{1}{2})^n$ depends on the underlying parameters stated precisely in Lemma 5. We remove the dependence of the parameters for a clean presentation of the proof.

D.11 Proof of Theorem 18

We first observe that by Theorem 5, the subset $\Sigma^{\alpha}_{M}(\Xi)$ is norm compact and forward invariant under the RBBR learning dynamic. Thus, in order to prove the theorem, we need to show the Bayesian entropy adjusted potential function $\widetilde{\varphi}^{\alpha}_{M}$ increases along every solution trajectory to the regularized Bayesian best response learning dynamic which originates in $\Sigma^{\alpha}_{M}(\Xi)$. Note that

$$\begin{aligned}
\dot{\widetilde{\varphi}}(\sigma) &= \dot{\varphi}(\sigma) - \dot{\widetilde{\mathbf{v}}}(\sigma) \\
&= \mathrm{D}\varphi[\sigma](\dot{\sigma}) - d\widetilde{\mathbf{v}}[\sigma](\dot{\sigma}) \\
&= \int_{\Omega} \langle \dot{\sigma}, \nabla^{\mathbf{F}}\varphi(\sigma)\rangle(\omega)\xi(d\omega) - \int_{\Omega} \langle \dot{\sigma}, \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}(\sigma)\rangle(\omega)\xi(d\omega) \\
&= \int_{\Omega} \nabla^{\mathbf{F}}\varphi[\sigma](\omega)(\dot{\sigma}(\omega))\xi(d\omega) - \int_{\Omega} \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}[\sigma](\omega)(\dot{\sigma}(\omega))\xi(d\omega) \\
&= \int_{\Omega} \langle \nabla^{\mathbf{F}}\varphi[\sigma](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) - \int_{\Omega} \langle \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}[\sigma](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) \\
&= \int_{\Omega} \langle \mathscr{G}(\sigma,\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) - \int_{\Omega} \langle \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}[\sigma](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) \\
&= \int_{\Omega} \langle \mathscr{G}(\sigma,\omega) - \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}[\sigma](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega) \\
&= \int_{\Omega} \langle \widetilde{\mathscr{G}}(\sigma,\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega)
\end{aligned}$$

Now we complete the proof by showing that the solution trajectory to the RBBR dynamic converges to the set of regularized Bayesian equilibria under the strong topology. By assumption of the theorem, it follows from Theorem 5 that $\Sigma^{\alpha}_{M}(\Xi)$ is norm compact and forward invariant. Also, since the vector field is Lipschitz under the norm $|||\,.\,|||$, it follows from Theorem 17 that a continuous and unique solution exists, given any initial condition $\sigma_0 \in \Sigma$. It then follows from standard results of dynamical systems theory (see Bhatia and Szegö (2006)) that the omega–limit set of the trajectory $(\sigma_t)_{t\geqslant 0}$ is non–empty, compact, and connected in the norm topology. Moreover this set is contained in the set of rest points of the RBBR learning dynamic. Hence any element in the omega–limit set is a rest point of the RBBR learning dynamic and hence a regularized Bayesian equilibrium. Also since the entropy adjusted Bayesian potential function is increasing along every solution trajectory, any point in the omega–limit set will be a maximizer of the function. This concludes the proof of Theorem 18.

D.12 Proof of Proposition 8

The proof of the proposition follows along the lines of Cheung (2014). Fix $\omega \in \Omega$, $\sigma \in \Sigma$, and $\sigma_0 \in \Sigma_0$. Since the slice map $\mathscr{G}_{\omega} : \widehat{\Sigma} \to \mathbb{R}^n$ is $\mathscr{C}^1$-Fréchet differentiable for ξ-a.e $\omega \in \Omega$, we have

$$\mathscr{G}_{\omega}(\sigma + \epsilon\sigma_0) = \mathscr{G}_{\omega}(\sigma) + \mathrm{D}\mathscr{G}_{\omega}[\sigma](\epsilon\sigma_0) + \mathbf{o}(\epsilon\|\sigma_0\|).$$

This implies that

$$\mathscr{G}_{\omega}(\sigma + \epsilon\sigma_0) - \mathscr{G}_{\omega}(\sigma) = \mathrm{D}\mathscr{G}_{\omega}[\sigma](\epsilon\sigma_0) + \mathbf{o}(\epsilon\|\sigma_0\|).$$

Taking inner product with respect to $\epsilon\sigma_0(\omega)$ on both sides of the above equation and integrating with respect to ξ, we have

$$\int_\Omega \langle \mathscr{G}_\omega(\sigma+\epsilon\sigma_0)-\mathscr{G}_\omega(\sigma),\epsilon\sigma_0(\omega)\rangle\xi(d\omega)=\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\sigma](\epsilon\sigma_0),\sigma_0(\omega)\rangle\xi(d\omega)+\mathbf{o}(\epsilon^2\|\sigma_0\|^2).$$

Since $\mathscr{G}$ is a negative semidefinite game, we have

$$\epsilon^2\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega(\sigma)(\sigma_0),\sigma_0(\omega)\rangle\xi(d\omega)+\mathbf{o}(\epsilon^2\|\sigma_0\|^2)\leqslant 0.$$

Since ϵ is arbitrary, we have by dividing both sides of the above inequality by ϵ^2 and allowing $\epsilon\to 0$, that

$$\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\sigma](\sigma_0),\sigma_0(\omega)\rangle\xi(d\omega)\leqslant 0.$$

This proves that $\mathscr{G}$ satisfies self defeating externalities if and only if $\mathscr{G}$ is a negative semidefinite game, which concludes the proof of Proposition 8.

D.13 Proof of Proposition 15

Proof 74 Hofbauer and Sandholm (2007) show uniqueness of logit equilibrium for finite strategy negative semidefinite games under homogeneous population set-up. We extend this result to finite strategy negative semidefinite games under heterogeneous set-up, which we call Bayesian negative semidefinite games. We prove the proposition along the lines of Lahkar and Riedel (2015), Lahkar et al. (2022b). Note that since the space $\Sigma_M^\alpha(\Xi)$ is forward invariant under the RBBR dynamic, it is enough to prove the uniqueness result on $\Sigma_M^\alpha(\Xi)$. Recall from Section 5.5 that the virtual payoff vector is such that for all $\sigma\in\Sigma_M^\alpha(\Xi)$,

$$\widetilde{\mathscr{G}}(\sigma,\omega)=\mathscr{G}(\sigma,\omega)-\nabla^{\mathbf{G}}\tilde{\mathbf{v}}[\sigma](\omega),\quad \text{for all } \omega\in\Omega. \tag{D.24}$$

Now, for every $\widehat{\sigma}_0\in\widehat{\Sigma}_0$ and $t\in\mathbb{R}$ such that $\sigma+t\widehat{\sigma}_0\in\Sigma_M^\alpha(\Xi)$, we define the following quantity:

$$\lambda(\sigma,\widehat{\sigma}_0,t):=\int_\Omega \langle\widetilde{\mathscr{G}}(\sigma+t\widehat{\sigma}_0,\omega),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega). \tag{D.25}$$

We would now like to calculate the Gâteaux derivative of the function $\lambda(\sigma,\widehat{\sigma}_0,\cdot)$ with respect to t. Notice that

$$\begin{aligned}\lambda(\sigma,\widehat{\sigma}_0,t+\eta)&=\int_\Omega\langle\widetilde{\mathscr{G}}(\sigma+(t+\eta)\widehat{\sigma}_0,\omega),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)\\&=\int_\Omega[\langle\widetilde{\mathscr{G}}(\sigma+t\widehat{\sigma}_0,\omega)+d\widetilde{\mathscr{G}}_\omega[\sigma+t\widehat{\sigma}_0](\eta\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\rangle+\mathbf{o}(|\eta|)]\xi(d\omega)\\&=\lambda(\sigma,\widehat{\sigma}_0,t)+\int_\Omega\langle d\widetilde{\mathscr{G}}_\omega[\sigma+t\widehat{\sigma}_0](\eta\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)+\mathbf{o}(|\eta|).\end{aligned}$$

Thus, we have that
$$\frac{\partial \lambda(\sigma,\widehat{\sigma}_0,t)}{\partial t}=\int_\Omega \langle d\widetilde{\mathcal{G}}_\omega[\sigma+t\widehat{\sigma}_0](\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega).$$

In order for us to complete the proof of the proposition, we need to compute the right hand side of the above equation. Note that from the definition of $\widetilde{\mathcal{G}}$, we have

$$\int_\Omega \langle d\widetilde{\mathcal{G}}_\omega[\sigma+t\widehat{\sigma}_0](\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)=\int_\Omega \langle d\mathcal{G}_\omega[\sigma+t\widehat{\sigma}_0](\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)$$
$$-\int_\Omega \langle d\nabla^{\mathbf{G}}\widetilde{\mathbf{v}}[\sigma+t\widehat{\sigma}_0](\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega). \tag{D.26}$$

Since the game $\mathcal{G}$ is Fréchet differentiable, $\mathcal{G}$ is Gâteaux differentiable too in which case, using the fact that $\mathcal{G}$ satisfies Bayesian self-defeating externalities, we have the following inequality:

$$\int_\Omega \langle d\mathcal{G}_\omega(\sigma+t\widehat{\sigma}_0)(\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)=\int_\Omega \langle D\mathcal{G}_\omega(\sigma+t\widehat{\sigma}_0)(\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)\leqslant 0.$$

We now show that the second term in the RHS of (E.18) is positive. For this, we note that for every $\sigma\in\Sigma_M^\alpha$, the following equality holds:

$$\nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon[\sigma](\omega)=\nabla\mathbf{v}_\epsilon(\sigma(\omega)),\quad \text{for all } \omega\in\Omega. \tag{D.27}$$

With the above equality in hand, we have that

$$\int_\Omega \langle d\nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon[\sigma+t\widehat{\sigma}_0](\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\xi(d\omega)=\int_\Omega \langle d\nabla\mathbf{v}_\epsilon[\sigma(\omega)+t\widehat{\sigma}_0(\omega)](\widehat{\sigma}_0(\omega)),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega). \tag{D.28}$$

By assumption of the lemma, $\langle d\nabla\mathbf{v}_\epsilon[\mathbf{x}+t\mathbf{z}_0](\mathbf{z}_0),\mathbf{z}_0\rangle\geqslant 0$, for all $t\in\mathbb{R}$, $\mathbf{x}\in\Delta$ and $\mathbf{z}_0\in\mathbb{R}_0^n$ such that $\mathbf{x}+t\mathbf{z}_0\in\Delta$. Thus, noticing the fact that $\sigma(\omega)+t\widehat{\sigma}_0(\omega)\in\Delta$, $\widehat{\sigma}_0(\omega)\in\mathbb{R}_0^n$ for all $\omega\in\Omega$ and using the equality relation in (D.28), we have that

$$\int_\Omega \langle d\nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon[\sigma+t\widehat{\sigma}_0](\widehat{\sigma}_0),\widehat{\sigma}_0(\omega)\xi(d\omega)=\int_\Omega \langle d\nabla\mathbf{v}_\epsilon[\sigma(\omega)+t\widehat{\sigma}_0(\omega)](\widehat{\sigma}_0(\omega)),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)\geqslant 0. \tag{D.29}$$

Therefore, the LHS of (E.18), and thus $\frac{\partial\lambda(\sigma,\widehat{\sigma}_0,t)}{\partial t}<0$. Now, let us fix a regularized Bayesian equilibrium σ°. It then follows from standard theory of gradient of convex conjugate (see for instance Rockafellar (1970)) that $\widetilde{\mathcal{G}}^i(\sigma^\circ,\omega)$ is independent of i for every $\omega\in\Omega$. In other words, given the equilibrium σ°, we have for every $\omega\in\Omega$, a $c(\sigma^\circ,\omega)$ such that $\widetilde{\mathcal{G}}^i(\sigma^\circ,\omega)=c(\sigma^\circ,\omega)$ for all $\omega\in\Omega$ and $i\in S$. As a result, we have that $\lambda(\sigma^\circ,\widehat{\sigma}_0,0)=\int_\Omega\langle\mathcal{G}(\sigma^\circ,\omega),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)=0$ for all $\widehat{\sigma}_0\in\widehat{\Sigma}$. Now define $\bar{\sigma}:=\sigma^\circ+t\widehat{\sigma}_0$ for some non-zero $\widehat{\sigma}_0\in\widehat{\Sigma}$ such that $\bar{\sigma}$ is distinct from σ°. By definition, we have $\int_\Omega\langle\widetilde{\mathcal{G}}(\bar{\sigma},\omega),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)=\lambda(\sigma^\circ,\widehat{\sigma}_0,t)$.Also, from the previous arguments, we have that $\lambda(\sigma^\circ,\widehat{\sigma}_0,0)=0$. This, in conjunction with the fact that $\frac{\partial\lambda(\bar{\sigma},\widehat{\sigma}_0,t)}{\partial t}<0$, implies that $\lambda(\bar{\sigma},\widehat{\sigma}_0,t)<0$. Therefore, we have that $\int_\Omega\langle\widetilde{\mathcal{G}}(\bar{\sigma},\omega),\widehat{\sigma}_0(\omega)\rangle\xi(d\omega)<0$, and hence $\bar{\sigma}$ cannot be a regularized equilibrium. This completes the proof of Proposition 15.

D.14 Proof of Theorem 19

Similar to the proof of Theorem 18, we need to show that the Lyapunov function $\psi^{\alpha}_{M,\epsilon}$ decreases along every solution trajectory to the RBBR learning dynamic which originates in $\Sigma^{\alpha}_{M}(\Xi)$. To this end let us define the following quantities:

- $\psi^1_\epsilon(\sigma) := \int_\Omega \langle \mathscr{G}(\sigma,\omega), \beta_\epsilon[\sigma](\omega)\rangle \xi(d\omega)$, for all $\sigma \in \Sigma^{\alpha}_{M}(\Xi)$,
- $\psi^2_\epsilon(\sigma) := \widetilde{\mathbf{v}}_\epsilon(\beta_\epsilon(\sigma))$, for all $\Sigma^{\alpha}_{M}(\Xi)$,
- $\psi^3_\epsilon(\sigma) := \int_\Omega \langle \mathscr{G}(\sigma,\omega), \sigma(\omega)\rangle \xi(d\omega)$, for all $\Sigma^{\alpha}_{M}(\Xi)$, and
- $\psi^4_\epsilon(\sigma) := \widetilde{\mathbf{v}}_\epsilon(\sigma)$, for all $\Sigma^{\alpha}_{M}(\Xi)$.

To prove the theorem, it is enough to show that

$$\dot{\psi}_\epsilon(\sigma) = \dot{\psi}^1_\epsilon(\sigma) + \dot{\psi}^2_\epsilon(\sigma) + \dot{\psi}^3_\epsilon(\sigma) + \dot{\psi}^4_\epsilon(\sigma) \leqslant 0.$$

We now compute the derivatives of ψ^i_ϵ, $i = 1, \ldots, 4$ separately. We have

$$\mathrm{D}\psi^1_\epsilon[\sigma](\dot{\sigma}) = \int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\sigma](\dot{\sigma}), \beta_\epsilon[\sigma](\omega)\rangle \xi(d\omega) + \int_\Omega \langle \mathscr{G}_\omega(\sigma), \dot{\beta}_\epsilon[\sigma](\omega)\rangle \xi(d\omega).$$

Since the regularizer $\widetilde{\mathbf{v}}_\epsilon$ is Gâteaux differentiable, we have that

$$\begin{aligned}
\mathrm{D}\psi^2_\epsilon[\sigma](\dot{\sigma}) &= \mathrm{d}\widetilde{\mathbf{v}}_\epsilon(\beta_\epsilon(\sigma))(\dot{\beta}_\epsilon(\sigma)) \\
&= \int_\Omega \langle \dot{\beta}_\epsilon(\sigma), \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon(\beta_\epsilon(\sigma))\rangle \xi(d\omega) \\
&= \int_\Omega \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon[\beta_\epsilon(\sigma)](\omega)(\dot{\beta}_\epsilon[\sigma](\omega))\xi(d\omega).
\end{aligned}$$

For the third term, we have

$$\mathrm{D}\psi^3_\epsilon[\sigma](\dot{\sigma}) = \int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\sigma](\dot{\sigma}), \sigma(\omega)\rangle \xi(d\omega) + \int_\Omega \langle \mathscr{G}_\omega(\sigma), \dot{\sigma}(\omega)\rangle \xi(d\omega).$$

Finally for the fourth term we have

$$\begin{aligned}
\mathrm{D}\psi^4_\epsilon[\sigma](\dot{\sigma}) &= \mathrm{d}\widetilde{\mathbf{v}}_\epsilon[\sigma](\dot{\sigma}) \\
&= \int_\Omega \langle \dot{\sigma}, \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon(\sigma)\rangle \xi(d\omega) \\
&= \int_\Omega \nabla^{\mathbf{G}}\widetilde{\mathbf{v}}_\epsilon[\sigma](\omega)(\dot{\sigma}(\omega))\xi(d\omega).
\end{aligned}$$

Combining all the four equations above, we have after simplification that

$$\begin{aligned}
\dot{\psi}^{\alpha}_{M,\epsilon}(\sigma) = & \int_{\Omega} \langle \mathrm{D}\mathscr{G}_{\omega}[\sigma](\dot{\sigma}), \beta_{\epsilon}(\omega)(\omega)\rangle \xi(d\omega) + \int_{\Omega} \langle \mathscr{G}_{\omega}(\sigma), \dot{\beta}_{\epsilon}[\sigma](\omega)\rangle \xi(d\omega) \\
& - \int_{\Omega} \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_{\epsilon}[\beta_{\epsilon}(\sigma)](\omega)(\dot{\beta}_{\epsilon}[\sigma](\omega))\xi(d\omega) - \int_{\Omega} \langle \mathrm{D}\mathscr{G}_{\omega}[\sigma](\dot{\sigma}), \sigma(\omega)\rangle \xi(d\omega) \\
& - \int_{\Omega} \langle \mathscr{G}_{\omega}(\sigma), \dot{\sigma}(\omega)\rangle \xi(d\omega) + \int_{\Omega} \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_{\epsilon}[\sigma](\omega)(\dot{\sigma}(\omega))\xi(d\omega) \\
= & \int_{\Omega} \langle \mathrm{D}\mathscr{G}_{\omega}[\sigma](\dot{\sigma}), \dot{\sigma}(\omega)\rangle \xi(d\omega) + \int_{\Omega} \langle \mathscr{G}_{\omega}(\sigma), \dot{\beta}_{\epsilon}[\sigma](\omega) - \dot{\sigma}(\omega)\rangle \xi(d\omega) \\
& - \int_{\Omega} \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_{\epsilon}[\beta_{\epsilon}(\sigma)](\omega)(\dot{\beta}_{\epsilon}[\sigma](\omega))\xi(d\omega) + \int_{\Omega} \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_{\epsilon}[\sigma](\omega)(\dot{\sigma}(\omega))\xi(d\omega).
\end{aligned}$$

Since $\mathscr{G}$ is a Bayesian negative semidefinite game and $\dot{\sigma} \in \Sigma_0$, it follows from Lemma 8 that the first term in the RHS of the above equation is less than or equal to zero, that is,

$$\int_{\Omega} \langle \mathrm{D}\mathscr{G}_{\omega}[\sigma](\dot{\sigma}), \dot{\sigma}(\omega)\rangle \xi(d\omega) \leqslant 0.$$

Using this fact and rearranging the other terms in the RHS of the above equation, we have

$$\begin{aligned}
\dot{\psi}^{\alpha}_{M,\epsilon}(\sigma) \leqslant & \int_{\Omega} \langle \mathscr{G}_{\omega}(\sigma) - \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_{\epsilon}[\beta_{\epsilon}(\sigma)](\omega), \dot{\beta}_{\epsilon}[\sigma](\omega)\rangle \xi(d\omega) \\
& \qquad - \int_{\Omega} \langle \mathscr{G}_{\omega}(\sigma) - \nabla^{\mathbf{G}} \widetilde{\mathbf{v}}_{\epsilon}[\sigma](\omega), \dot{\sigma}(\omega)\rangle \xi(d\omega).
\end{aligned}$$

It now follows from Lemma (31) that the first term in the RHS of the above inequality is 0. It also follows that the second term in the RHS of the above inequality is non-negative. As a result, we have $\dot{\psi}^{\alpha}_{M,\epsilon} \leqslant 0$ along every solution which originates in $\Sigma^{\alpha}_{M}(\Xi)$. In particular, this implies that the Lyapunov function defined in (E.19) decreases along every non–stationary solution which originates in $\Sigma^{\alpha}_{M}(\Xi)$. By part (i) of Theorem 5, we have that $\Sigma^{\alpha}_{M}(\Xi)$ is compact under the norm topology on $\widehat{\Sigma}$ and $\dot{\psi}^{\alpha}_{M,\epsilon} \leqslant 0$ with equality only at the rest points, it follows from the standard theory of dynamical systems (Bhatia and Szegö (2006)) that the omega–limit of every solution trajectory satisfies $\dot{\psi}^{\alpha}_{M,\epsilon} = 0$, which is the set of all rest points of the RBBR dynamic. We therefore have that every non–stationary solution arising in $\Sigma^{\alpha}_{M}(\Xi)$ converges to the unique regularized Bayesian equilibrium.

APPENDIX

E

SUPPLEMENT TO CHAPTER 5

E.1 Proof of Proposition 16

We first present the formal definition of Böchner measurability for correspondences, and then state a useful result for our proof that relates Böchner measurability to a relatively simpler notion called scalar measurability.

A correspondence $\Gamma : \Theta \rightrightarrows \mathscr{M}$ is Böchner measurable if there exists a sequence of simple correspondences $(\Gamma_n)_{n\geq 1}$ such that for ξ-a.e $\omega \in \Omega$,

$$\mathscr{H}(\Gamma_n(\omega), \Gamma(\omega)) \to 0 \text{ as } n \to \infty,$$

where $\mathscr{H}$ denotes the Hausdorff distance on $\mathscr{P}$.[1]

Let $\mathscr{M}^*$ denote the dual space of $\mathscr{M}$. For $E \subseteq \mathscr{M}$, the support function of E is a mapping $\sigma(\cdot, E) : \mathscr{M}^* \to \mathbb{R}$ such that

$$\sigma(\mu^*, E) := \sup\{\langle \mu^*, \mu \rangle : \mu \in E\} \quad \text{for all } \mu^* \in \mathscr{M}^*.$$

A correspondence $\Gamma : \Theta \rightrightarrows \mathscr{M}$ is said to be scalarly measurable if for all $\mu^* \in \mathscr{M}^*$ the map

$$\omega \mapsto \sigma(\mu^*, \Gamma(\omega)) \text{ is measurable.}$$

From Theorem III.2 and III.15 in Castaing and Valadier (1977), it follows that for compact and convex valued correspondences, the Böchner measurability coincides with scaler measurability.

We now show that $\mathbf{B}[\theta] : \Omega \rightrightarrows \mathscr{P}$ is Böchner measurable. First, we show that $\mathbf{B}[\theta](\omega)$ is compact

[1] For $E, F \subseteq \mathscr{M}$, the Hausdorff distance $\mathscr{H}(E, F)$ between E and F is defined as

$$\mathscr{H}(E, F) := \max\left\{ \sup_{\mu \in E}\left(\inf_{\nu \in F} \|\mu - \nu\|_{\mathbf{BL}^*} \right),\ \sup_{\nu \in F}\left(\inf_{\mu \in E} \|\mu - \nu\|_{\mathbf{BL}^*} \right) \right\}.$$

and convex valued for every $\omega \in \Omega$, and then we conclude the proof by showing that $\mathbf{B}[\theta]$ is scalarly measurable, that is, the function $\omega \mapsto \sigma(\mu^*, \mathbf{B}[\theta](\omega))$ is measurable.

The fact that $\mathbf{B}[\theta](\omega)$ is convex valued follows from the fact that any convex combination of two probability measures maximizing the expected payoff of $\mathscr{G}(\theta,\omega)$ also does so due to the linearity of the expectation operation. We show that $\mathbf{B}[\theta](\omega)$ is compact-valued for every $\omega \in \Omega$. Note that $\mathscr{P}$ is compact since it is endowed with the weak topology. Thus, it is enough to show that $\mathbf{B}[\theta](\omega)$ is a closed subset of $\mathscr{P}$ for all $\omega \in \Omega$. Let $(\mu_n)_{n\geq 1}$ be a sequence in $\mathbf{B}[\theta](\omega)$ such that $\mu_n \to^{\mathbf{BL}^*} \mu$. This implies

$$\langle \mathscr{G}(\theta,\omega), \mu_n \rangle \geq \langle \mathscr{G}(\theta,\omega), \nu \rangle \quad \text{for all } \nu \in \mathscr{P}.$$

Since $\mathscr{G}(\theta,\omega) \in \mathscr{C}$ for every $\omega \in \Omega$, weak convergence of the sequence $(\mu_n)_{n\geq 1}$ to μ implies that

$$\langle \mathscr{G}(\theta,\omega), \mu \rangle \geq \langle \mathscr{G}(\theta,\omega), \nu \rangle \quad \text{for all } \nu \in \mathscr{P},$$

and hence $\mu \in \mathbf{B}[\theta](\omega)$. This shows $\mathbf{B}[\theta](\omega)$ is closed (and hence compact) for every $\omega \in \Omega$.

Finally, we show that the map $\omega \mapsto \sigma(\mu^*, \mathbf{B}[\theta](\omega))$ is measurable. First, note that since $\langle \mu^*, \cdot \rangle : \mathscr{M} \to \mathbb{R}$ is a bounded linear functional, $\mu \mapsto \langle \mu^*, \mu \rangle$ is measurable (under the topology induced by the $\mathbf{BL}^*$-norm). Furthermore, $\mathbf{B}[\theta](\omega)$ is separable since it is a compact subset of $\mathscr{M}$. Since, the map $\omega \mapsto \mathscr{G}(\theta,\omega)$ is measurable and $\mathscr{G}(\theta,\omega) \in \mathscr{C}$ for every $\omega \in \Omega$, we have by the application of Measurable Maximum Theorem (Aliprantis and Border (2013), Pg. 508, Theorem 14.91) that

$$\omega \mapsto \sigma(\mu^*, \mathbf{B}[\theta](\omega)) \text{ is measurable,}$$

and hence the correspondence Γ is scalarly measurable. This completes the proof that $\mathbf{B}[\theta]$ is Böchner measurable.

E.2 Proof of Theorem 15

The proof of Theorem 15 is an application of Brouwer-Schauder-Tychonoff fixed point theorem (Aliprantis and Border (2013), Pg. 583, Corr. 17.56). We first make the following observation. Suppose that for a given $\eta > 0$, a logit Bayesian equilibrium θ°_η exists. Then, we must have that $\theta^\circ_\eta \in \Theta_\beta$ for some $\beta > 1$ (depending on the noise parameter η).[2] Thus, we complete the proof of Theorem 15 in two steps. First, we show that Θ_β is compact under the weak topology. Then, we show that the mapping $\theta \mapsto \mathbf{L}_\eta(\theta)$ is continuous on Θ_β with respect to the same topology.

Lemma 33 *The subset Θ_β is compact under the weak topology.*

Proof 75 We shall invoke Dunford's Theorem (Diestel and Uhl Jr (1978), Pg. 101, Theorem 1) which characterizes relatively weakly compact subsets of Böchner spaces. In particular, we first need to show that the subset Θ_β is bounded and uniformly integrable. This follows immediately since Θ_β consists of Bayesian strategies, the range of which is contained in the set $\mathscr{P}_\beta$, which in turn, contains probability measures with densities uniformly bounded by β. Thus to conclude the proof of the lemma, we finally need to show that for every $\Xi \subseteq \Omega$, the set $\{\langle \theta, \xi \rangle_\Xi : \theta \in \Theta_\beta\}$ has compact closure as a subset of $\mathscr{P}$ under the $\|\cdot\|_{\mathscr{L}^2}$-topology.[3] Note that by Kolmogorov-Reisz compactness theorem (Hanche-Olsen and Holden (2010), Perkins and Leslie (2014)), the subset $\mathscr{P}_\beta$ is compact under the $\|\cdot\|_{\mathbf{TV}}$-topology. Since

[2] See the proof of Theorem 15 for validity of the statement.

[3] Here, $\langle \theta, \xi \rangle_\Xi := \int_\Xi \theta(\omega)\xi(d\omega)$.

$\{\langle\theta,\xi\rangle_{\Xi}:\theta\in\Theta_\beta\}$ can be viewed as an embedded subset of $\mathscr{P}_\beta$ for every $\Xi\subseteq\Theta$, the proof of the lemma follows. This concludes the proof of lemma 33.

To proceed with the proof of the theorem, we require another lemma which we shall also use in this as well as in the subsequent proofs of the paper.

Lemma 34 *The following inequality holds true: for all $\theta,\rho\in\Theta$,*

$$\|\boldsymbol{L}_\eta[\theta](\omega)-\boldsymbol{L}_\eta[\rho](\omega)\|_{\mathscr{L}^2}\leq\eta^{-1}\exp(6\eta^{-1}M)\|\mathscr{G}(\theta,\omega)-\mathscr{G}(\rho,\omega)\|_\infty,\quad \textit{for all } \omega\in\Omega. \tag{E.1}$$

Proof 76 Consider $\theta,\rho\in\Theta$ and $\omega\in\Omega$. Using the definition of $\|\cdot\|_{\mathscr{L}^2}$, we have the following:

$$\begin{aligned}
&\|\mathbf{L}_\eta[\theta](\omega)-\mathbf{L}_\eta[\rho](\omega)\|_{\mathscr{L}^2}\\
&=\frac{1}{2}\int_{\mathscr{S}}|\mathbf{L}_\eta^x[\theta](\omega)-\mathbf{L}_\eta^x[\rho](\omega)|^2dx\\
&=\frac{1}{2}\int_{\mathscr{S}}\left|\frac{\exp(\eta^{-1}\mathscr{G}_x(\theta,\omega))}{\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))dy}-\frac{\exp(\eta^{-1}\mathscr{G}_x(\rho,\omega))}{\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_y(\rho,\omega))dy}\right|^2dx\\
&\leq\exp(2\eta^{-1}M)\int_{\mathscr{S}}\left|\frac{\exp(\eta^{-1}\mathscr{G}_x(\theta,\omega))}{\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))dy}-\frac{\exp(\eta^{-1}\mathscr{G}_x(\rho,\omega))}{\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_y(\rho,\omega))dy}\right|dx\\
&\leq\exp(4\eta^{-1}M)\int_{\mathscr{S}}\left|\int_{\mathscr{S}}\Big[\exp(\eta^{-1}(\mathscr{G}_x(\theta,\omega)+\mathscr{G}_y(\rho,\omega)))-\exp(\eta^{-1}(\mathscr{G}_x(\rho,\omega)+\mathscr{G}_y(\theta,\omega)))\Big]dy\right|dx\\
&\leq\exp(4\eta^{-1}M)\int_{\mathscr{S}}\left|\int_{\mathscr{S}}\Big[\exp(\eta^{-1}(\mathscr{G}_x(\theta,\omega)+\mathscr{G}_y(\rho,\omega)))-\exp(\eta^{-1}(\mathscr{G}_x(\theta,\omega)+\mathscr{G}_y(\theta,\omega)))\Big]dy\right|dx\\
&\qquad+\left|\int_{\mathscr{S}}\Big[\exp(\eta^{-1}(\mathscr{G}_x(\theta,\omega)+\mathscr{G}_y(\theta,\omega)))-\exp(\eta^{-1}(\mathscr{G}_x(\rho,\omega)+\mathscr{G}_y(\theta,\omega)))\Big]dy\right|dx\\
&=\exp(4\eta^{-1}M)\Big[\int_{\mathscr{S}}\left|\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_x(\theta,\omega))\Big[\exp(\eta^{-1}\mathscr{G}_y(\rho,\omega))-\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))\Big]dy\right|dx\\
&\qquad\qquad+\int_{\mathscr{S}}\left|\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))\Big[\exp(\eta^{-1}\mathscr{G}_x(\theta,\omega))-\exp(\eta^{-1}\mathscr{G}_x(\rho,\omega))\Big]dy\right|dx\Big].
\end{aligned}$$

To proceed with the proof of the lemma, let us define the map $g:[-M,M]\to\mathbb{R}$ as $g(u):=\exp(\eta^{-1}u)$ for all $-M\leq u\leq M$. Thus, by an application of the standard mean value theorem (Rudin (2006)) for real-valued functions we have the following inequality:

$$|g(u)-g(v)|\leq\eta^{-1}\exp(\eta^{-1}M)|u-v|,\quad \text{for all } -M\leq u,v\leq M.$$

Now set $u=\mathscr{G}_y(\rho,\omega),v=\mathscr{G}_y(\theta,\omega)$ and $u=\mathscr{G}_x(\theta,\omega),v=\mathscr{G}_x(\rho,\omega)$ subsequently and make use of the above inequality in order to obtain the fact that there exists a constant $\gamma(\eta)>0$ (depending on η) such that the mapping $\omega\mapsto\mathbf{L}_\eta[\theta](\omega)$ is $\gamma(\eta)$-Lipschitz. Continuing from the above string of inequalities we have:

$$
\begin{aligned}
&\|\mathbf{L}_\eta[\theta](\omega)-\mathbf{L}_\eta[\rho](\omega)\|_{\mathscr{L}^2}\\
&\le \exp(4\eta^{-1}M)\Big[\int_{\mathscr{S}}\Big|\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_x(\theta,\omega))\Big[\exp(\eta^{-1}\mathscr{G}_y(\rho,\omega))-\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))\Big]dy\Big|dx\\
&\qquad\qquad+\int_{\mathscr{S}}\Big|\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))\Big[\exp(\eta^{-1}\mathscr{G}_x(\theta,\omega))-\exp(\eta^{-1}\mathscr{G}_x(\rho,\omega))\Big]dy\Big|dx\Big]\\
&\le \exp(5\eta^{-1}M)\Big[\int_{\mathscr{S}}\Big|\exp(\eta^{-1}\mathscr{G}_y(\rho,\omega))-\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))\Big|dy\\
&\qquad\qquad\qquad\qquad+\int_{\mathscr{S}}\Big|\exp(\eta^{-1}\mathscr{G}_y(\rho,\omega))-\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))\Big|dy\Big]\\
&\le \eta^{-1}\exp(6\eta^{-1}M)\|\mathscr{G}(\theta,\omega)-\mathscr{G}(\rho,\omega)\|_\infty.
\end{aligned}
$$

This concludes the proof of Lemma 34.

Finally, to conclude the proof of Theorem 15, we require the following lemma.

Lemma 35 *The mapping $\theta\mapsto\mathbf{L}_\eta(\theta)$ is continuous on Θ_β in the weak topology.*

Proof 77 Fix $\theta\in\Theta_\beta$, and consider a sequence $(\theta_n)_{n\ge1}\subseteq\Theta_\beta$, such that $\theta_n\to^w\theta$. We need to show that $\mathbf{L}_\eta(\theta_n)\to^w\mathbf{L}_\eta(\theta)$. Note that by assumption of the lemma, the game $\mathscr{G}$ is uniformly bounded for ξ-a.e $\omega\in\Omega$. In other words, there exists $M>0$ such that for all $\theta\in\Theta$, we have

$$
\|\mathscr{G}(\theta,\omega)\|_\infty\le M,\quad \text{for } \xi-\text{a.e } \omega\in\Omega.
$$

As a result, one obtains for ξ-a.e $\omega\in\Omega$, the following bounds on the probability density function of the logit Bayesian best response $\mathbf{L}_\eta(\theta)$ at Bayesian strategy $\theta\in\Theta$:

$$
\exp(-2\eta^{-1}M)\le\mathbf{L}^x_\eta[\theta](\omega):=\frac{\exp(\eta^{-1}\mathscr{G}_x(\theta,\omega))}{\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))dy}\le\exp(2\eta^{-1}M),\quad\text{for all } x\in\mathscr{S}. \tag{E.2}
$$

Note that, by the assumption of the theorem, the game $\mathscr{G}$ is continuous relative to the weak topology on Θ_β. As a result, $\mathscr{G}(\theta_n,\omega)$ converges to $\mathscr{G}(\theta,\omega)$ in the sup-norm for ξ-a.e $\omega\in\Omega$. Thus, by an application of Lemma 34, we can conclude that $\mathbf{L}_\eta[\theta_n](\omega)\to\mathbf{L}_\eta[\theta](\omega)$ under the $\|\cdot\|_{\mathscr{L}^2}$-topology for ξ-a.e $\omega\in\Omega$. Now let us fix $\rho\in\mathscr{L}^\infty(\Omega,\xi,\mathscr{L}^2)$. Then, by an application of Hölder's inequality (Rudin (2006)), we have the following observation:

$$\begin{aligned}
\Big|\int_\Omega \langle \mathbf{L}_\eta[\theta_n]-\mathbf{L}_\eta[\theta],\rho\rangle(\omega)\xi(d\omega)\Big| &= \Big|\int_\Omega \langle \mathbf{L}_\eta[\theta_n](\omega)-\mathbf{L}_\eta[\theta](\omega),\rho(\omega)\rangle\xi(d\omega)\Big| \\
&= \Big|\int_\Omega\int_{\mathscr{S}} (\mathbf{L}^x_\eta[\theta_n](\omega)-\mathbf{L}^x_\eta[\theta](\omega))\rho_x(\omega)\lambda(dx)\xi(d\omega)\Big| \\
&\le \int_\Omega\Big|\int_{\mathscr{S}} (\mathbf{L}^x_\eta[\theta_n](\omega)-\mathbf{L}^x_\eta[\theta](\omega))\rho_x(\omega)\lambda(dx)\Big|\xi(d\omega) \\
&\le \int_\Omega\int_{\mathscr{S}} \|\mathbf{L}_\eta[\theta_n](\omega)-\mathbf{L}_\eta[\theta](\omega)\|^{1/2}_{\mathscr{L}^2}\|\rho(\omega)\|^{1/2}_{\mathscr{L}^2}\lambda(dx)\xi(d\omega) \\
&\le \operatorname*{ess\,sup}_{\omega\in\Omega}\|\rho(\omega)\|^{1/2}_{\mathscr{L}^2}\int_\Omega \|\mathbf{L}_\eta[\theta_n](\omega)-\mathbf{L}_\eta[\theta](\omega)\|^{1/2}_{\mathscr{L}^2}\xi(d\omega).
\end{aligned}$$

We need to show that the term in the RHS of the above inequality converges to 0. First, observe that $\rho\in\mathscr{L}^\infty(\Omega,\xi,\mathscr{L}^2)$. As a result, $\operatorname{ess\,sup}_{\omega\in\Omega}\|\rho(\omega)\|^{1/2}_{\mathscr{L}^2}<\infty$. Since the Bayesian population game $\mathscr{G}$ is bounded, we have from (E.2) that the probability density $\mathbf{L}_\eta[\theta](\omega)$ is bounded above by $\exp(2\eta^{-1}M)$. As a result, the integrand in the last line of the above string of inequalities is bounded above by $2\exp(2\eta^{-1}M)$. Since $\mathbf{L}_\eta[\theta_n](\omega)\to\mathbf{L}_\eta[\theta](\omega)$ under the $\|\cdot\|_{\mathscr{L}^2}$-topology for ξ-a.e $\omega\in\Omega$, we have the desired convergence by an application of Dominated Convergence theorem. This proves that the mapping $\theta\mapsto\mathbf{L}_\eta(\theta)$ is continuous in the weak topology on Θ_β which completes the proof of Lemma 35.

Since $\hat\Theta$ is endowed with the weak topology, it is a Hausdorff topological space (Kesavan (2009); Brezis (2011), Pg. 57, Proposition 3.3). Therefore, it follows from standard results in functional analysis (Rudin (1991), Pg. 64, Theorem 3.10) that $\hat\Theta$ is a locally convex Hausdorff topological space under the weak topology. By Lemma 33 and Lemma 35, it follows that Θ_β is a non-empty compact convex subset of $\hat\Theta$ and $\mathbf{L}_\eta$ is a continuous map. Thus the conditions of the Brouwer-Schauder-Tychonoff Theorem are satisfied and the existence of a fixed point is established. This concludes the proof of Theorem 15.

E.3 Proof of Corollary 5

To prove the corollary, we need to show that the Bayesian aggregative population game $\mathscr{G}(\theta,\varphi):=\varphi(\mathscr{E}(\theta))$ is weakly continuous. As stated in the proof of Theorem 15, it is therefore enough to show that the mapping $\theta\mapsto\mathscr{E}(\theta)$ is continuous with respect to the weak topology on Θ_β. Again, let $\theta\in\Theta_\beta$ and consider a sequence $(\theta_n)_{n\ge1}$ such that $\theta_n\to^w\theta$. We show that $\mathscr{E}(\theta_n)\to^w\mathscr{E}(\theta)$. Notice that for every $n\ge1$, the existence of a probability density function for the probability measure $\theta_n(\omega)$ is guaranteed for ξ-a.e $\omega\in\Omega$ since $\theta_n\in\Theta_\beta$. Therefore, for $\varphi\in\Phi_{wc}$ and $x\in\mathscr{S}$, let us define a bounded linear functional $\rho_x[\varphi]:\mathscr{L}^2\to\mathbb{R}$ such that

$$\rho_x[\varphi](g):=g_x,\quad \text{for all } g\in\mathscr{L}^2. \tag{E.3}$$

In other words, ρ_x is the evaluation of any map $g\in\mathscr{L}^2$ at the point x. We now proceed to conclude the proof of the corollary. Thus, let us fix $A\in\mathscr{B}$. Applying Fubini's theorem (Billingsley (2013)), we observe that,

$$
\begin{aligned}
\mathscr{E}[\theta_n](A) &= \int_{\Phi_{wc}} \theta_n[\varphi](A)\xi(d\varphi) \\
&= \int_{\Phi_{wc}} \int_A \theta_{n,x}(\varphi)\lambda(dx)\xi(d\varphi) \\
&= \int_{\Phi_{wc}} \int_A \rho_x[\varphi](\theta_n(\varphi))\lambda(dx)\xi(d\varphi) \\
&= \int_{\Phi_{wc}} \int_A \langle \theta_n, \rho_x\rangle(\varphi)\lambda(dx)\xi(d\varphi) \\
&= \int_A \int_{\Phi_{wc}} \langle \theta_n, \rho_x\rangle(\varphi)\xi(d\varphi)\lambda(dx) \\
&\to \int_A \int_{\Phi_{wc}} \langle \theta, \rho_x\rangle(\varphi)\xi(d\varphi)\lambda(dx) \\
&= \int_{\Phi_{wc}} \int_A \langle \theta, \rho_x\rangle(\varphi)\lambda(dx)\xi(d\varphi) \\
&= \int_{\Phi_{wc}} \int_A \rho_x[\varphi](\theta(\varphi))\lambda(dx)\xi(d\varphi) \\
&= \int_{\Phi_{wc}} \int_A \theta_x(\varphi)\lambda(dx)\xi(d\varphi) \\
&= \int_{\Phi_{wc}} \theta[\varphi](A)\xi(d\varphi) \\
&= \mathscr{E}[\theta](A).
\end{aligned}
$$

Therefore by Portmanteau's theorem (Billingsley (2013), Pg. 11, Theorem 2.1), we have that $\mathscr{E}(\theta_n) \to^w \mathscr{E}(\theta)$. Since the type space Φ_{wc} consists of weakly continuous population games, we have that $\theta_n \to^w \theta$, implies $\mathscr{G}(\theta_n, \varphi) = \varphi(\mathscr{E}(\theta_n)) \to \varphi(\mathscr{E}(\theta)) = \mathscr{G}(\theta, \varphi)$, for all $\varphi \in \Phi_{wc}$. Thus, $\mathscr{G}$ is a weakly continuous Bayesian population game and hence concludes the proof of Corollary 5.

E.4 Proof of Theorem 16

Consider a sequence of noise levels $(\eta_n)_{n\geq 1}$ such that $\eta_n \downarrow 0$. Since $\mathscr{G}$ is continuous in the weak topology, we have by Theorem 15 the existence of η_n-logit Bayesian equilibrium for every $n \geq 1$. Let θ_n° denote the logit Bayesian equilibrium corresponding to noise parameter η_n such that $\theta_n^\circ \to^w \theta^\circ$. We need to show that θ° is a Bayesian equilibrium of the original game $\mathscr{G}$. Assume on the contrary that θ° is not a Bayesian equilibrium. Thus, by definition, there exists $\bar{\Omega} \subseteq \Omega$ with $\xi(\bar{\Omega}) > 0$ such that $\theta^\circ(\omega) \notin \mathbf{B}[\theta^\circ](\omega)$ for all $\omega \in \bar{\Omega}$. This, in particular, implies that $\theta^\circ(\omega) \notin \operatorname{argmax}_{\mu \in \mathscr{P}} \langle \mathscr{G}(\theta^\circ, \omega), \mu\rangle$ for $\omega \in \bar{\Omega}$. Therefore, for every $\omega \in \bar{\Omega}$, there exists $y_\omega \in \mathscr{S}$, $x_\omega \in \operatorname{supp}(\theta^\circ(\omega))$ and $\epsilon_\omega > 0$ satisfying the following inequality:

$$
\mathscr{G}_{y_\omega}(\theta^\circ, \omega) > 2\epsilon_\omega + \mathscr{G}_{x_\omega}(\theta^\circ, \omega). \tag{E.4}
$$

The fact that $\mathscr{G}(\theta^\circ,\omega)\in\mathscr{C}$ implies that there exists open balls $B_{x_\omega}, B_{y_\omega}$ around x_ω and y_ω such that

$$\inf_{z\in B_{y_\omega}} \mathscr{G}_z(\theta^\circ,\omega)\geq 2\epsilon_\omega \sup_{z\in B_{x_\omega}} \mathscr{G}_z(\theta^\circ,\omega). \tag{E.5}$$

Also by assumption of the theorem, we have that the Bayesian population game $\mathscr{G}$ is continuous with respect to the $\|\cdot\|_{\hat{\Theta}\mathbf{BL}^*}$-topology, there exists $K_\omega\geq 1$ sufficiently large such that

$$\inf_{z\in B_{y_\omega}} \mathscr{G}_z(\theta_n^\circ,\omega) > \epsilon_\omega + \sup_{z\in B_{x_\omega}} \mathscr{G}_z(\theta_n^\circ,\omega), \quad \text{for sufficiently large } n\geq K_\omega. \tag{E.6}$$

Notice that θ_n° is a logit Bayesian equilibrium of the game $\mathscr{G}$ corresponding to noise parameter η_n for every $n\geq 1$. Therefore, $\theta_n^\circ = \mathbf{L}_{\eta_n}(\theta_n^\circ)$ for every $n\geq 1$. As a result, we have

$$\theta_n^\circ(\omega) = \mathbf{L}_{\eta_n}[\theta_n^\circ](\omega), \quad \text{for } \xi\text{-a.e } \omega\in\Omega. \tag{E.7}$$

We are now ready to conclude the proof of the theorem. Let us fix $\omega\in\bar{\Omega}$. Then consider the following computation:

$$\begin{aligned}
\frac{\theta_n^\circ[\omega](B_{y_\omega})}{\theta_n^\circ[\omega](B_{x_\omega})} &= \frac{\mathbf{L}_{\eta_n}[\theta_n^\circ][\omega](B_{y_\omega})}{\mathbf{L}_{\eta_n}[\theta_n^\circ][\omega](B_{x_\omega})} \\
&= \frac{\int_{B_{y_\omega}} \exp(\eta_n^{-1}\mathscr{G}_z(\theta_n^\circ,\omega))dz}{\int_{B_{x_\omega}} \exp(\eta_n^{-1}\mathscr{G}_z(\theta_n^\circ,\omega))dz} \\
&\geq \frac{\exp(\eta_n^{-1}\inf_{z\in B_{y_\omega}}\mathscr{G}_z(\theta_n^\circ,\omega))\lambda(B_{y_\omega})}{\exp(\eta_n^{-1}\sup_{z\in B_{x_\omega}}\mathscr{G}_z(\theta_n^\circ,\omega))\lambda(B_{x_\omega})} \\
&= \exp(\epsilon_\omega/\eta_n)\frac{\lambda(B_{y_\omega})}{\lambda(B_{x_\omega})}.
\end{aligned}$$

First observe that B_{x_ω} and B_{y_ω} are open balls and thus have positive Lebesgue measure. Thus, allowing n to approach infinity in the above inequality, we obtain that $\theta_n^\circ[\omega](B_{x_\omega})\to 0$ as $n\to\infty$. In other words, for all $\omega\in\bar{\Omega}$, there exists $B_{x_\omega}\in\mathscr{B}$ such that $\theta_n^\circ[\omega](B_{x_\omega})\to 0$. However, since $x_\omega\in\operatorname{supp}(\theta^\circ(\omega))$, we must have $\theta^\circ[\omega](B_{x_\omega})>0$. Thus, there exists a set $\bar{\Omega}$ with positive ξ-measure such that $\liminf_{n\to\infty}|\theta_n^\circ[\omega](B_{x_\omega}) - \theta^\circ[\omega](B_{x_\omega})|>0$ for all $\omega\in\bar{\Omega}$. This contradicts the fact that $\theta_n^\circ\to\theta^\circ$ in the $\|\cdot\|_{\hat{\Theta}\mathbf{BL}^*}$-topology. Hence θ° is a Bayesian equilibrium of the game $\mathscr{G}$. This concludes the proof of Theorem 16.

E.5 Proof of Theorem 17

The proof of the theorem is an application of the infinite dimensional Picard-Lindelöff Theorem (Zeidler (1986), Pg. 79). The existence of solutions to the LBBR dynamic is guaranteed once we show that the mapping $\theta\mapsto\mathbf{L}_\eta(\theta)$ is Lipschitz continuous with respect to the $\|\cdot\|_{\hat{\Theta}\mathbf{TV}}$-topology.

Lemma 36 *The mapping $\theta\mapsto\boldsymbol{L}_\eta(\theta)$ is Lipschitz continuous with respect to $\|\cdot\|_{\hat{\Theta}\boldsymbol{TV}}$-topology on $\hat{\Theta}$.*

Proof 78 Consider $\theta,\rho\in\Theta$ and $\omega\in\Omega$. Using the conclusion of Lemma 34 in conjunction with the fact that the Bayesian population game $\mathscr{G}$ is strongly κ-Lipschitz with respect to the $\|\cdot\|_{\hat{\Theta}\mathbf{TV}}$-topology, we have that

$$\|\mathbf{L}_\eta[\theta](\omega)-\mathbf{L}_\eta[\rho](\omega)\|_{\mathbf{TV}} \le \kappa\eta^{-1}\exp(4\eta^{-1}M)\|\theta-\rho\|_{\hat{\Theta}\mathbf{TV}}.$$[4]

We now integrate both sides of the above inequality with respect to the probability measure ξ to prove that the mapping $\theta \mapsto \mathbf{L}_\eta(\theta)$ is continuous with respect to the strong topology on $\hat{\Theta}$. This concludes the proof of Theorem 17.

We now apply the infinite dimensional Grönwall's Lemma (Zeidler (1986), Propositions 3.10 and 3.11) to prove the continuity of the semiflow part 2 of Theorem 17. Since the conditions of Grönwall's Lemma is satisfied, there exists $\alpha > 0$ such that for any two solutions $(\theta_t)_{t\ge 0}$ and $(\tilde{\theta}_t)_{t\ge 0}$ to the LBBR dynamic with initial conditions θ_0 and $\tilde{\theta}_0$, respectively, we have $\|\theta_t - \tilde{\theta}_t\|_{\hat{\Theta}\mathbf{TV}^*} \le e^{\alpha t}\|\theta_0 - \tilde{\theta}_0\|_{\hat{\Theta}\mathbf{TV}}$ for all $t > 0$. This establishes continuity of the semiflow of the LBBR dynamic and hence concludes the proof of Theorem 17.

E.6 Proof of Lemma 7

To prove the lemma, we require the following result (Balder et al. (1994), Pg. 3, Theorem 2.4). We provide the result below for completeness.

Lemma 37 *A subset $\mathscr{K}$ of $\mathscr{L}^1(\mu, \boldsymbol{X})$ is relatively norm compact if and only if the following conditions are satisfied.*

1. *The subset $\mathscr{K}$ is relatively weakly compact.*
2. *The subset $\mathscr{K}$ satisfies small Bocce oscillation.*
3. *The subset $\{avg(f;A) : f \in \mathscr{K}\}$ is relatively norm compact in $\boldsymbol{X}$, for every $A \in \mathscr{B}_\Omega$, such that $\mu(A) > 0$.*[5]

To apply the above lemma, we take $\mathbf{X} = \mathscr{L}^2(\xi)$. The fact that $\Theta_{\beta,\gamma}(\Xi)$ is bounded and uniformly integrable follows since $\Theta_{\beta,\gamma}(\Xi)$ consists of only those Bayesian strategies which induces probability density functions that are bounded by β. Fix $\rho^* \in \mathscr{L}^2(\xi)$. For $\theta \in \Theta_{\beta,\gamma}(\Xi)$, let us define the mapping $\langle \rho^*, \theta\rangle : \Omega \to \mathbb{R}$ such that

$$\langle \rho^*, \theta\rangle(\omega) = \rho^*(\theta(\omega)) := \int_{\mathscr{S}} \rho^*(x)\theta_x(\omega)\lambda(dx), \quad \text{for all } \omega \in \Omega.$$

We need to show that the collection of maps $\Lambda_{\rho^*} := \{\langle \rho^*, \theta\rangle : \theta \in \Theta_{\beta,\gamma}(\Xi)\rangle\}$ is relatively norm compact in $\mathscr{L}^1(\xi)$ for all $\rho^* \in \mathscr{L}^2(\xi)$. By Lemma 37, it is enough to show that the above collection is relatively weakly compact and satisfies small Bocce-oscillation. The relative weak compactness follows from the following inequality:

$$\begin{aligned} |\langle \rho^*, \theta\rangle| &\le \|\rho^*\|^{1/2}\|\theta\|^{1/2} \\ &\le \beta^{1/2}\|\rho^*\|^{1/2}. \end{aligned}$$

The proof of the lemma concludes if we are able to show that the collection Λ_{ρ^*} satisfies small Bocce-oscillation for every $\rho^* \in \mathscr{L}^2(\xi)$. Thus fix $\rho^* \in \mathscr{L}^2(\xi)$ and $\theta \in \Theta_{\beta,\gamma}(\Xi)$. Then we have the following

[4]The proof of the inequality is similar to the proof of Lemma 34 replacing $\|\cdot\|_{\mathscr{L}^2}$ by $\|\cdot\|_{\mathbf{TV}}$.

[5]See Balder et al. (1994) for the definitions of 'avg' and 'small Bocce-oscillation'.

computation: for $\omega \in \Omega$, we have

$$
\begin{aligned}
|\langle \rho^*, \theta\rangle(\omega) - \operatorname{avg}(\langle \rho^*, \theta\rangle, B_\omega^\delta)| &= \left| \langle \rho^*, \theta\rangle(\omega) - \frac{1}{\xi(B_\omega^\delta)} \int_{B_\omega^\delta} |\langle \rho^*, \theta\rangle(\bar{\omega}) \xi(d\bar{\omega}) \right| \\
&= \frac{1}{\xi(B_\omega^\delta)} \left| \int_{B_\omega^\delta} \left[\langle \rho^*, \theta\rangle(\omega) - \langle \rho^*, \theta\rangle(\bar{\omega}) \right] \xi(d\omega) \right| \\
&\le \frac{1}{\xi(B_\omega^\delta)} \int_{B_\omega^\delta} |\langle \rho^*, \theta\rangle(\omega) - \langle \rho^*, \theta\rangle(\bar{\omega})| \xi(d\omega) \\
&\le \frac{1}{\xi(B_\omega^\delta)} \int_{B_\omega^\delta} \int_{\mathscr{S}} |\rho^*(x)\theta_x(\omega) - \rho^*(x)\theta_x(\bar{\omega})| \lambda(dx) \xi(d\omega) \\
&= \frac{1}{\xi(B_\omega^\delta)} \int_{B_\omega^\delta} \int_{\mathscr{S}} |\rho^*(x)||\theta_x(\omega) - \theta_x(\bar{\omega})| \lambda(dx) \xi(d\omega) \\
&\le \frac{\|\rho^*\|_{\mathscr{L}^2}}{\xi(B_\omega^\delta)} \int_{B_\omega^\delta} \|\theta(\omega) - \theta(\bar{\omega})\|_{\mathbf{TV}} \xi(d\omega) \\
&\le \|\rho^*\|_{\mathscr{L}^2} \epsilon.
\end{aligned}
$$

Thus the Bocce-osc$(\langle \rho^*, \theta\rangle, B_\omega^\delta) \le \|\rho^*\|_{\mathscr{L}^2} \epsilon$. The rest of the proof now follows along the lines of Mukherjee and Roy (2021).

We now proceed to show that the subset $\Theta_{\beta,\gamma}(\Xi)$ is forward invariant under the LBBR dynamic. Thus, we need to show that $\theta \in \Theta_{\beta,\gamma}(\Xi)$ implies $\mathbf{L}_\eta(\theta) \in \Theta_{\beta,\gamma}(\Xi)$. In other words, we need to show that $\mathbf{L}_\eta(\theta)$ belongs to both Θ_β and $\Theta_\gamma(\Xi)$. First, note that by the assumption of the lemma, the game $\mathscr{G}$ is uniformly bounded for ξ-a.e $\omega \in \Omega$, that is there exists $M > 0$ such that for all $\theta \in \Theta$,

$$
\|\mathscr{G}(\theta, \omega)\|_\infty \le M, \quad \text{for } \xi-\text{a.e } \omega \in \Omega.
$$

Using the above inequality, the following bound holds for the logit Bayesian best response density function $\mathbf{L}_\eta^x[\theta](\omega)$ for all $\theta \in \Theta$ and ξ-a.e $\omega \in \Omega$:

$$
\exp(-2\eta^{-1}M) \le \mathbf{L}_\eta^x[\theta](\omega) \le \exp(2\eta^{-1}M) \quad \text{for all } x \in \mathscr{S}. \tag{E.8}
$$

Secondly, by assumption of the lemma, we also have that the game $\mathscr{G}$ is κ-strongly Lipschitz for ξ-a.e $\omega \in \Omega$. Therefore, it follows from Perkins and Leslie (2014) (Pg. 201, Lemma C.1) that

$$
|\mathbf{L}_\eta^x[\theta](\omega) - \mathbf{L}_\eta^y[\theta](\omega)| \le \frac{2\kappa \exp(2\eta^{-1}M)}{\eta} |x - y|, \quad \text{for all } x, y \in \mathscr{S}.
$$

Notice that the perturbation parameter $\eta > 0$ is fixed. Thus, given η, we define

$$
\beta := \max\{\exp(2\eta^{-1}M), 2\eta^{-1}\kappa \exp(2\eta^{-1}M)\}.
$$

Therefore for ξ-a.e $\omega \in \Omega$, we have the following two inequalities:

- $\frac{1}{\beta} \le \mathbf{L}_\eta^x[\theta](\omega) \le \beta$ for all $x \in \mathscr{S}$.
- $|\mathbf{L}_\eta^x[\theta](\omega) - \mathbf{L}_\eta^y[\theta](\omega)| \le \beta |x - y|$ for all $x, y \in \mathscr{S}$.

As a result, we can conclude that $\mathbf{L}_\eta(\theta) \in \Theta_\beta$. Now, in order to complete the proof of the lemma, we need to show that $\mathbf{L}_\eta(\theta)$ belongs to $\Theta_\gamma(\Xi)$. In other words, we need to show that there exists a constant $\gamma > 0$ (depending on η) such that the logit Bayesian best response $\mathbf{L}_\eta(\theta)$ at the Bayesian strategy $\theta \in \Theta$ viewed as a map $\Omega \ni \omega \mapsto \mathbf{L}_\eta[\theta](\omega) \in \mathscr{P}$ is γ-Lipschitz. More precisely, we need to show that

$$\|\mathbf{L}_\eta[\theta](\omega) - \mathbf{L}_\eta[\theta](\bar{\omega})\|_{\mathbf{TV}} \leq \gamma d_\Omega(\omega, \bar{\omega}), \quad \text{for } \xi\text{-a.e } \omega, \bar{\omega} \in \Omega.$$

The proof of the above statement follows using similar arguments as in the proof of Lemma 34 in conjunction with the fact that for every $\theta \in \Theta$, the Bayesian population game $\mathscr{G}$ is $\hat{\kappa}$-Lipschitz for ξ-a.e $\omega, \bar{\omega} \in \Omega$. We ommit the proof to avoid redundancy. This proves that the map $\Omega \ni \omega \mapsto \mathbf{L}_\eta[\theta](\omega) \in \mathscr{P}$ is γ-Lipschitz and hence $\mathbf{L}_\theta(\theta) \in \Theta_\gamma(\Xi)$. As a result, we have that $\mathbf{L}_\eta(\theta) \in \Theta_{\beta,\gamma}(\Xi)$. This shows that $\Theta_{\beta,\gamma}(\Xi)$ is forward invariant under the LBBR dynamic and hence the proof of Lemma 7 concludes.

E.7 Proof of Theorem 18

Since the conditions of Lemma 7 are satisfied, we first conclude that the subset $\Theta_{\beta,\gamma}(\Xi)$ is relatively norm compact and forward invariant under the logit Bayesian best response dynamic. Thus, in order to prove the theorem, we need to show the entropy adjusted Bayesian potential function $\hat{\varphi}$ is an appropriate Lyapunov function corresponding to the underlying evolutionary dynamic. In particular, we need to show that $\hat{\varphi}$ increases weakly along every solution trajectory to the Bayesian best response dynamic and increases strictly across every non-stationary solution trajectory that originates in $\Theta_{\beta,\gamma}(\Xi)$. To this end, we require the following lemma.

Lemma 38 *Consider the Bayesian population game $\mathscr{G} : \Theta \times \Omega \to \mathscr{C}$. Then the following statements are true:*

1. *For every $\theta \in \Theta$ and $\omega \in \Omega$, one has*

$$\mathscr{G}_x(\theta, \omega) = \eta \log \boldsymbol{L}_\eta^x[\theta](\omega) + \eta \log \int_{\mathscr{S}} \exp(\eta^{-1} \mathscr{G}_y(\theta, \omega)) dy, \quad \textit{for all } x \in \mathscr{S}. \tag{E.9}$$

2. *Consider the virtual payoff function $\hat{\mathscr{G}}$. Then, the logit Bayesian best response dynamic satisfies virtual positive correlation at θ. That is,*

$$\int_\Omega \langle \hat{\mathscr{G}}(\theta, \omega), \dot{\theta}(\omega) \rangle \xi(d\omega) \geq 0 \tag{E.10}$$

with equality only at rest points of the dynamic.

Proof 79 To prove part 1 of the lemma, we observe that for every $\theta \in \Theta$ and every $\omega \in \Omega$,

$$\mathbf{L}_\eta^x[\theta](\omega) = \frac{\exp(\eta^{-1} \mathscr{G}_x(\theta, \omega))}{\int_{\mathscr{S}} \exp(\eta^{-1} \mathscr{G}_y(\theta, \omega)) dy}, \quad \text{for all } x \in \mathscr{S}. \tag{E.11}$$

Taking log followed by multiplication by η on both sides of the above equality yields

$$\mathscr{G}_x(\theta, \omega) = \eta \log \mathbf{L}_\eta^x[\theta](\omega) + \eta \log \int_{\mathscr{S}} \exp(\eta^{-1} \mathscr{G}_y(\theta, \omega)) dy, \quad \text{for all } x \in \mathscr{S}. \tag{E.12}$$

This concludes the proof of part 1 of the lemma.

We now proceed with the proof of part 2 of the lemma. By using the definition of $\hat{\mathscr{G}}$ and an application of part 1 of the lemma, we have the following string of inequalities.

$$\begin{aligned}
\int_{\Omega}\langle\hat{\mathscr{G}}(\theta,\omega),\dot{\theta}(\omega)\rangle\xi(d\omega) &= \int_{\Omega}\langle\mathscr{G}(\theta,\omega)-\eta\log\theta_{\cdot}(\omega),\dot{\theta}(\omega)\rangle\xi(d\omega)\\
&= \int_{\Omega}\int_{\mathscr{S}}(\mathscr{G}_x(\theta,\omega)-\eta\log\theta_x(\omega))\dot{\theta}_x(\omega)(dx)\xi(d\omega)\\
&= \int_{\Omega}\int_{\mathscr{S}}(\eta\log\mathbf{L}^x_\eta[\theta](\omega)+\eta\log\int_{\mathscr{S}}\exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))dy\\
&\qquad\qquad -\eta\log\theta_x(\omega))(\mathbf{L}^x_\eta[\theta](\omega)-\theta_x(\omega))(dx)\xi(d\omega)\\
&= \int_{\Omega}\int_{\mathscr{S}}(\log\mathbf{L}^x_\eta[\theta](\omega)-\log\theta_x(\omega))(\mathbf{L}^x_\eta[\theta](\omega)-\theta_x(\omega))(dx)\xi(d\omega).
\end{aligned}$$

Since the mapping $x\mapsto\log x$ is increasing, we note that for every $\omega\in\Omega$ and every $x\in\mathscr{S}$,

$$\log\mathbf{L}^x_\eta[\theta](\omega)\geq\log\theta_x(\omega)\quad\text{if and only if}\quad\mathbf{L}^x_\eta[\theta](\omega)\geq\theta_x(\omega).$$

Thus, the above integral is non-negative. Also, we notice that equality holds in the above inequality only if $\dot{\theta}=0$. This concludes the proof of part 2 of the lemma and hence, the proof of Lemma 38.

We now proceed with the proof of Theorem 18. Let $(\theta_t)_{t\geq 0}$ denote the solution trajectory of the logit Bayesian best response dynamic with initial condition $\theta_0\in\Theta_{\beta,\gamma}(\Xi)$. Since $\Theta_{\beta,\gamma}(\Xi)$ is forward invariant under the best response dynamic (Lemma 7), it follows that $\theta_t\in\Theta_{\beta,\gamma}(\Xi)$ for all $t>0$. To conclude that the entropy adjusted Bayesian potential function increases along this solution trajectory, we compute the derivative of $\hat{\varphi}$ using the chain rule of differentiation. Note that $\hat{\varphi}$ consists of two terms, namely, the original potential function and the associated Shannon entropy. Since the potential function is Fréchet differentiable, we can apply chain rule directly. However, the term associated to Shannon entropy, though not Fréchet differentiable, is continuously Gâteaux differentiable since, in $\Theta_{\beta,\gamma}(\Xi)$, The Gâteaux derivative of $\hat{\mathbf{v}}$ defines a continuous bilinear mapping

$$\langle\nabla^{\mathbf{G}}\hat{\mathbf{v}}(\theta),\theta_0\rangle=\eta\int_{\Omega}\int_{\mathscr{S}}\log\theta_x(\omega)\theta_0^x(\omega)(dx)\xi(d\omega),\tag{E.13}$$

for $\theta_0\in\Theta_0$ such that $\theta_0(\omega)$ admits a density function for ξ-a.e $\omega\in\Omega$. Thus, using the chain rule of differentiation, we obtain:

$$
\begin{aligned}
\dot{\hat{\varphi}}(\theta) &= \dot{\varphi}(\theta) - \dot{\hat{\mathbf{v}}}(\theta) \\
&= \mathrm{D}\varphi[\theta](\dot{\theta}) - \mathrm{d}\hat{\mathbf{v}}[\theta](\dot{\theta}) \\
&= \int_\Omega \langle \dot{\theta}, \nabla^{\mathbf{F}}\varphi(\theta)\rangle(\omega)\xi(d\omega) - \int_\Omega \langle \dot{\theta}, \nabla^{\mathbf{G}}\hat{\mathbf{v}}(\theta)\rangle(\omega)\xi(d\omega) \\
&= \int_\Omega \nabla^{\mathbf{F}}\varphi[\theta](\omega)(\dot{\theta}(\omega))\xi(d\omega) - \int_\Omega \nabla^{\mathbf{G}}\hat{\mathbf{v}}[\theta](\omega)(\dot{\theta}(\omega))\xi(d\omega) \\
&= \int_\Omega \langle \nabla^{\mathbf{F}}\varphi[\theta](\omega), \dot{\theta}(\omega)\rangle\xi(d\omega) - \int_\Omega \langle \nabla^{\mathbf{G}}\hat{\mathbf{v}}[\theta](\omega), \dot{\theta}(\omega)\rangle\xi(d\omega) \\
&= \int_\Omega \langle \mathscr{G}(\theta,\omega), \dot{\theta}(\omega)\rangle\xi(d\omega) - \int_\Omega \langle \eta\log\theta_{\cdot}(\omega), \dot{\theta}(\omega)\rangle\xi(d\omega) && \text{(E.14)} \\
&= \int_\Omega \langle \mathscr{G}(\theta,\omega) - \eta\log\theta_{\cdot}(\omega), \dot{\theta}(\omega)\rangle\xi(d\omega) \\
&= \int_\Omega \langle \hat{\mathscr{G}}(\theta,\omega), \dot{\theta}(\omega)\rangle\xi(d\omega) \\
&\geq 0, && \text{(E.15)}
\end{aligned}
$$

with equality only at the rest points of the logit Bayesian best response dynamic. We note that (E.14) follows from (E.13) and (E.15) follows as a consequence of Lemma 38.

We are now ready to complete the proof of the theorem. Notice that the space $\Theta_{\beta,\gamma}(\Xi)$ is relatively norm compact under the $\|\cdot\|_{\hat{\Theta}\mathbf{TV}}$-topology and the logit Bayesian best response dynamic generates a continuous semiflow on this forward invariant space with respect to the same topology (Lemma 7). Therefore, by standard results from the theory of dynamical systems (see Bhatia and Szegö (2006), pg. 117, Theorem 2.2.9), we obtain that the omega-limit sets of the solution trajectory θ_t is non-empty, compact and connected. Moreover, this set is contained in the set of Bayesian strategies which satisfy $\dot{\hat{\varphi}}(\theta)=0$. Therefore, we can conclude that the omega-limit points of the trajectory are rest points of the dynamic and is, hence, a logit Bayesian equilibrium. Furthermore, since the entropy adjusted Bayesian potential function is shown to increase along the solution trajectory $(\theta_t)_{t\geq 0}$, we must have that any such limit point is a local maximizer of $\hat{\varphi}$. This concludes the proof of Theorem 18.

E.8 Proof of Lemma 8

The proof of the lemma follows along the lines of Cheung (2014). Fix $\omega\in\Omega$, $\theta\in\Theta$, and $\theta_0\in\Theta_0$. Since the slice map $\mathscr{G}_\omega:\hat{\Theta}\to\mathscr{C}$ is $\mathscr{C}^1$-Fréchet differentiable for ξ-a.e $\omega\in\Omega$, we have

$$\mathscr{G}_\omega(\theta+\eta\theta_0) = \mathscr{G}_\omega(\theta) + \mathrm{D}\mathscr{G}_\omega[\theta](\eta\theta_0) + \mathbf{o}(\eta\|\theta_0\|).$$

As a result, we have that

$$\mathscr{G}_\omega(\theta+\eta\theta_0) - \mathscr{G}_\omega(\theta) = \mathrm{D}\mathscr{G}_\omega[\theta](\eta\theta_0) + \mathbf{o}(\eta\|\theta_0\|).$$

Thus, taking inner product with respect to $\eta\theta_0(\omega)$ on both sides of the above equation and integrating with respect to ξ, we have

$$\int_\Omega \langle \mathscr{G}_\omega(\theta+\eta\theta_0)-\mathscr{G}_\omega(\theta), \eta\theta_0(\omega)\rangle \xi(d\omega) = \int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\eta\theta_0), \eta\theta_0(\omega)\rangle \xi(d\omega) + \mathbf{o}(\eta^2\|\theta_0\|^2).$$

Since $\mathscr{G}$ is a negative semidefinite game, we have by Definition 47 that

$$\eta^2 \int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\theta_0), \theta_0(\omega)\rangle \xi(d\omega) + \mathbf{o}(\eta^2\|\theta_0\|^2) \leq 0.$$

Since η is arbitrary, we have by dividing both sides of the above inequality by η^2 and allowing $\eta \to 0$, that

$$\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\theta_0), \theta_0(\omega)\rangle \xi(d\omega) \leq 0.$$

This proves that $\mathscr{G}$ satisfies self defeating externalities if and only if $\mathscr{G}$ is a negative semidefinite game, which concludes the proof of Lemma 8.

E.9 Proof of Theorem 19

We first show that the game $\mathscr{G}$ admits a unique logit Bayesian equilibrium.

Lemma 39 *The game $\mathscr{G}$ admits a unique logit Bayesian equilibrium.*

Proof 80 We first notice that to prove the lemma, it enough to show uniqueness of equilibrium on $\Theta_{\beta,\gamma}(\Xi)$. Fix $\theta \in \Theta_{\beta,\gamma}(\Xi)$ and $\omega \in \Omega$ and consider the Bayesian virtual payoff

$$\hat{\mathscr{G}}_x(\theta,\omega) := \mathscr{G}_x(\theta,\omega) - \eta \log \theta_x(\omega), \quad \text{for all } x \in \mathscr{S}. \tag{E.16}$$

For every $\hat{\theta}_0 \in \hat{\Theta}$ and $t \in \mathbb{R}$ such that $\theta + t\hat{\theta}_0 \in \Theta_{\beta,\gamma}(\Xi)$, we define a new quantity as

$$\begin{aligned}
\Pi(\theta,\hat{\theta}_0,t+\epsilon) &= \int_\Omega \langle \hat{\mathscr{G}}(\theta+(t+\epsilon)\hat{\theta}_0,\omega), \hat{\theta}_0(\omega)\rangle \xi(d\omega) \\
&= \int_\Omega [\langle \hat{\mathscr{G}}(\theta+t\hat{\theta}_0,\omega) + d\hat{\mathscr{G}}_\omega[\theta+t\hat{\theta}_0](\epsilon\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle + \mathbf{o}(|\epsilon|)]\xi(d\omega) \\
&= \Pi(\theta,\hat{\theta}_0,t) + \int_\Omega \langle d\hat{\mathscr{G}}_\omega[\theta+t\hat{\theta}_0](\epsilon\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle \xi(d\omega) + \mathbf{o}(|\epsilon|)
\end{aligned}$$

Computing the Gâteaux derivative of $\Pi(\theta,\hat{\theta}_0,\cdot)$ with respect to t, one obtains the following

$$\frac{\partial \Pi(\theta,\hat{\theta}_0,t)}{\partial t} = \int_\Omega \langle d\hat{\mathscr{G}}_\omega[\theta+t\hat{\theta}_0](\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle \xi(d\omega). \tag{E.17}$$

We now compute the right hand side of the above equation. Note that from the definition of $\hat{\mathscr{G}}$, we have

$$\int_\Omega \langle d\hat{\mathscr{G}}_\omega[\theta + t\hat{\theta}_0](\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle \xi(d\omega) = \int_\Omega \langle d\mathscr{G}_\omega[\theta + t\hat{\theta}_0](\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle \xi(d\omega) \\ -\eta \int_\Omega \langle d\log[\theta + t\hat{\theta}_0](\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle \xi(d\omega). \tag{E.18}$$

Since the game $\mathscr{G}$ is Fréchet differentiable, $\mathscr{G}$ is Gâteaux differentiable too in which case, using the fact that $\mathscr{G}$ satisfies Bayesian self-defeating externalities, we have the following inequality:

$$\int_\Omega \langle d\mathscr{G}_\omega[\theta + t\hat{\theta}_0](\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle \xi(d\omega) = \int_\Omega \langle D\mathscr{G}_\omega[\theta + t\hat{\theta}_0](\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle \xi(d\omega) \leq 0.$$

In order to complete the proof of the lemma, we need to show that the second term in the RHS of (E.18) is positive. Thus, we compute $d\log[\theta + t\hat{\theta}_0](\hat{\theta}_0)$ explicitly. Fix $\omega \in \Omega$ and $x \in \mathscr{S}$. We then have

$$d\log[\theta_x(\omega) + t\hat{\theta}_0^x(\omega)](\hat{\theta}_0^x(\omega)) = \lim_{\epsilon \to 0} \frac{\log(\theta_x(\omega) + t\hat{\theta}_0^x(\omega) + \epsilon\hat{\theta}_0^x(\omega)) - \log(\theta_x(\omega) + t\hat{\theta}_0^x(\omega))}{\epsilon} \\ = \frac{\hat{\theta}_0^x(\omega)}{\theta_x(\omega) + t\hat{\theta}_0^x(\omega)}.$$

First observe that $t \in \mathbb{R}$ and $\hat{\theta}_0 \in \hat{\Theta}$ is chosen in such a way that $\theta + t\hat{\theta}_0 \in \Theta_{\beta,\gamma}(\Xi)$. Hence the denominator in the above ratio in non-negative. Therefore, the second term in the LHS of (E.18) boils down to

$$\eta \int_\Omega \langle d\log[\theta + t\hat{\theta}_0](\hat{\theta}_0), \hat{\theta}_0(\omega)\rangle \xi(d\omega) = \eta \int_\Omega \int_{\mathscr{S}} \frac{[\hat{\theta}_0^x(\omega)]^2}{\theta_x(\omega) + t\hat{\theta}_0^x(\omega)} (dx)\xi(d\omega) \\ \geq 0,$$

as required. Thus, we have $\frac{\partial \Pi(\theta, \hat{\theta}_0, t)}{\partial t} < 0$. Now, let us fix a regularized Bayesian equilibrium θ°. Notice that $\hat{\mathscr{G}}_x(\theta^\circ, \omega) = \eta \int_{\mathscr{S}} \exp(\eta^{-1}\mathscr{G}_y(\theta^\circ, \omega))dy$, which is a quantity independent of x for every $\omega \in \Omega$. As a result, we have $\Pi(\theta^\circ, \hat{\theta}_0, 0) = \eta \int_\Omega \langle \hat{\mathscr{G}}(\theta^\circ, \omega), \hat{\theta}_0(\omega)\rangle \xi(d\omega) = 0$, for all $\hat{\theta}_0 \in \hat{\Theta}$ which admits a density function. Now let us define a Bayesian strategy ρ as $\rho := \theta^\circ + t\hat{\theta}_0 \in \Theta_{\beta,\gamma}(\Xi)$ such that $\rho \neq \theta^\circ$. By definition, we have $\int_\Omega \langle \hat{\mathscr{G}}(\rho, \omega), \hat{\theta}_0(\omega)\rangle \xi(d\omega) = \Pi(\theta^\circ, \hat{\theta}_0, t)$. Since $\Pi(\theta^\circ, \hat{\theta}_0, 0) = 0$, we have that $\Pi(\theta^\circ, \hat{\theta}_0, t) < 0$. Therefore, we must have $\int_\Omega \langle \hat{\mathscr{G}}(\rho, \omega), \hat{\theta}_0(\omega)\rangle \xi(d\omega) < 0$. Thus, ρ cannot be logit Bayesian equilibrium. This completes the proof of Lemma 39.

We now proceed with the proof of Theorem 19. Similar to the proof of Theorem 18, we show that in case of Bayesian negative semidefinite games, the Lyapunov function $\hat{\psi}_\eta$ in (E.19) decreases along every solution trajectory to the logit Bayesian best response dynamic which originates in $\Theta_{\beta,\gamma}(\Xi)$. Recall that $\hat{\psi}_\eta$ is defined as

$$\hat{\psi}_\eta(\theta) := \Big[\int_\Omega \langle \mathscr{G}(\theta,\omega), \mathbf{L}_\eta[\theta](\omega)\rangle \xi(d\omega) - \hat{\mathbf{v}}_\eta(\mathbf{L}_\eta(\theta))\Big] \\ - \Big[\int_\Omega \langle \mathscr{G}(\theta,\omega), \theta(\omega)\rangle \xi(d\omega) - \hat{\mathbf{v}}_\eta(\theta)\Big], \quad \text{for all } \theta \in \Theta_{\beta,\gamma}(\Xi). \tag{E.19}$$

Note that

$$\begin{aligned}\hat{\mathbf{v}}_\eta(\mathbf{L}_\eta(\theta)) &= \int_\Omega \mathbf{v}_\eta(\mathbf{L}_\eta[\theta](\omega))\xi(d\omega) \\ &= \eta \int_\Omega \int_{\mathscr{S}} \mathbf{L}^x_\eta[\theta](\omega) \log \mathbf{L}^x_\eta[\theta](\omega)(dx)\xi(d\omega) \\ &= \int_\Omega \int_{\mathscr{S}} \mathscr{G}_x(\theta,\omega)\mathbf{L}^x_\eta[\theta](\omega)(dx)\xi(d\omega) - \eta \int_\Omega \Big[\log \int_{\mathscr{S}} \exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))dy\Big]\xi(d\omega).\end{aligned}$$

Therefore plugging in the value of $\mathbf{v}_\eta(\mathbf{L}_\eta(\theta))$ obtained above in (E.19), we observe that the expression of $\hat{\psi}_\eta(\theta)$ reduces to

$$\hat{\psi}_\eta(\theta) = \eta \int_\Omega \Big[\log \int_{\mathscr{S}} \exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))dy\Big]\xi(d\omega) \\ - \Big[\int_\Omega \langle \mathscr{G}(\theta,\omega), \theta(\omega)\rangle \xi(d\omega) - \hat{\mathbf{v}}_\eta(\theta)\Big], \quad \text{for all } \theta \in \Theta_{\beta,\gamma}(\Xi).$$

We now denote the three separate quantities in the right hand side of the above equality by $\hat{\psi}^1_\eta(\theta), \hat{\psi}^2_\eta(\theta)$, and $\hat{\psi}^3_\eta(\theta)$ respectively. More precisely:

- First term: $\hat{\psi}^1_\eta(\theta) := \eta \int_\Omega \Big[\log \int_{\mathscr{S}} \exp(\eta^{-1}\mathscr{G}_y(\theta,\omega))dy\Big]\xi(d\omega)$ for all $\theta \in \Theta_{\beta,\gamma}(\Xi)$,
- Second term: $\hat{\psi}^2_\eta(\theta) := \int_\Omega \langle \mathscr{G}(\theta,\omega), \theta(\omega)\rangle \xi(d\omega)$ for all $\theta \in \Theta_{\beta,\gamma}(\Xi)$, and
- Third term: $\hat{\psi}^3_\eta(\theta) := \hat{\mathbf{v}}_\eta(\theta)$ for all $\theta \in \Theta_{\beta,\gamma}(\Xi)$.

We claim that the map $\hat{\psi}_\eta$ acts as a Lyapunov function for the underlying logit Bayesian best response dynamic subject to the condition that the game is Bayesian negative semidefinite. Thus, in order to prove the theorem, it is enough to show that

$$\dot{\hat{\psi}}_\eta(\theta) = \dot{\hat{\psi}}^1_\eta(\theta) + \dot{\hat{\psi}}^2_\eta(\theta) + \dot{\hat{\psi}}^3_\eta(\theta) \le 0, \tag{E.20}$$

along the solution of the logit Bayesian best response dynamic arising in $\Theta_{\beta,\gamma}(\Xi)$. Thus, we now compute the derivatives of $\hat{\psi}^i_\eta$, $i = 1,2,3$ separately. Firstly, notice that that term $\hat{\psi}^1_\eta$ consists of a smooth transformation of the Bayesian population game $\mathscr{G}$, which is Fréchet differentiable. Therefore, we can conclude that $\hat{\psi}^1_\eta$ is also Fréchet differentiable. We thus compute the Fréchet derivative of $\hat{\psi}^1_\eta$ as follows:

$$\mathrm{D}\hat{\psi}^1_\eta[\theta](\dot{\theta}) = \int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\dot{\theta}), \mathbf{L}_\eta[\theta](\omega)\rangle \xi(d\omega). \tag{E.21}$$

To compute $\mathrm{D}\hat{\psi}^2_\eta$, note that $\hat{\psi}^2_\eta(\theta) = \int_\Omega \langle \mathscr{G}(\theta,\omega), \theta(\omega)\rangle \xi(d\omega)$. Therefore, we have

$$\mathrm{D}\hat{\psi}^2_\eta[\theta](\dot{\theta}) = \int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\dot{\theta}), \theta(\omega)\rangle \xi(d\omega) + \int_\Omega \langle \mathscr{G}_\omega(\theta), \dot{\theta}(\omega)\rangle \xi(d\omega). \tag{E.22}$$

Finally for the third term $\hat{\psi}^3_\eta$, we have

$$\begin{aligned}
\mathrm{D}\hat{\psi}^3_\eta[\theta](\dot{\theta}) &= \mathrm{d}\hat{\mathbf{v}}_\eta[\theta](\dot{\theta}) \\
&= \mathrm{d}\hat{\mathbf{v}}_\eta[\theta](\dot{\theta}) \\
&= \int_\Omega \nabla^{\mathbf{G}}\hat{\mathbf{v}}_\eta[\theta](\omega)(\dot{\theta}(\omega))\xi(d\omega) \\
&= \int_\Omega \langle \mathbf{L}^{\cdot}_\eta[\theta](\omega), \dot{\theta}(\omega)\rangle \xi(d\omega),
\end{aligned} \tag{E.23}$$

where $\mathbf{L}^{\cdot}_\eta[\theta](\omega)$ denotes the logit best response density function corresponding to Bayesian strategy θ for the sub-population of players of type ω. Combining all the three equations above, namely (E.21), (E.22) and (E.23), we have after simplification that

$$\begin{aligned}
\dot{\hat{\psi}}_\eta(\theta) = &\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\dot{\theta}), \mathbf{L}_\eta[\theta](\omega)\rangle \xi(d\omega) - \Big[\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\dot{\theta}), \theta(\omega)\rangle \xi(d\omega) \\
&\qquad - \int_\Omega \langle \mathscr{G}_\omega(\theta), \dot{\theta}(\omega)\rangle \xi(d\omega) - \int_\Omega \langle \mathbf{L}^{\cdot}_\eta[\theta](\omega), \dot{\theta}(\omega)\rangle \xi(d\omega)\Big] \\
= &\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\dot{\theta}), \mathbf{L}_\eta[\theta](\omega) - \theta(\omega)\rangle \xi(d\omega) - \int_\Omega \langle \mathscr{G}_\omega(\theta) - \mathbf{L}^{\cdot}_\eta[\theta](\omega), \dot{\theta}(\omega)\rangle \xi(d\omega) \\
= &\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\dot{\theta}), \dot{\theta}(\omega)\rangle \xi(d\omega) - \int_\Omega \langle \mathscr{G}_\omega(\theta) - \mathbf{L}^{\cdot}_\eta[\theta](\omega), \dot{\theta}(\omega)\rangle \xi(d\omega) \\
\leq &\, 0,
\end{aligned}$$

with equality only at the rest points. The above quantity is non-positive because of the following reason. First notice that since $\mathscr{G}$ is a Bayesian negative semidefinite game and $\dot{\theta} \in \Theta_0$, it follows from part 2 of Definition 48 and Lemma 8 that the first term in the right hand side of the above equation is less than or equal to zero, that is,

$$\int_\Omega \langle \mathrm{D}\mathscr{G}_\omega[\theta](\dot{\theta}), \dot{\theta}(\omega)\rangle \xi(d\omega) \leq 0.$$

The second term in the expression of $\dot{\hat{\psi}}_\eta$ is zero if $\dot{\theta} = 0$ and negative if $\dot{\theta} \neq 0$. This is due to the fact the logit Bayesian best response dynamic satisfies virtual positive correlation with respect to the virtual population game $\hat{\mathscr{G}}$ (see part (ii) of Lemma 38).

As a result, we observe that $\dot{\hat{\psi}}_\eta(\theta) \leq 0$ for all $\theta \in \Theta_{\beta,\gamma}(\Xi)$, with equality only at rest points of the dynamic. Thus, $\hat{\psi}_\eta$ decreases monotonically along every solution to the logit Bayesian best response dynamic which originates in $\Theta_{\beta,\gamma}(\Xi)$, and decreases strictly along every non-stationary solution. We now complete the proof of the theorem. Notice that the subset $\Theta_{\beta,\gamma}(\Xi)$ is a compact subset of the space of Bayesian strategies under the $\|\cdot\|_{\hat{\Theta}\mathbf{TV}}$-topology. Since $\dot{\hat{\psi}}_\eta(\theta) \leq 0$ with equality only at the rest points, we can conclude using standard results of dynamical systems (Bhatia and Szegö (2006), Pg. 117, Theorem

2.2.9) that the omega-limit set corresponding to every solution trajectory satisfies $\dot{\hat{\psi}}_\eta = 0$, which, in particular is the set of logit Bayesian equilibria. Therefore, every solution trajectory converges to the set of logit Bayesian equilibria of the game $\mathscr{G}$. However, we also know from Lemma 39 that a Bayesian negative semidefinite game has a unique Bayesian logit equilibrium. Thus, in particular, every solution trajectory of the logit Bayesian best response dynamic converges to the unique equilibrium of the Bayesian negative semidefinite game $\mathscr{G}$.