The Continuous Logit Dynamic and Price Dispersion

Ratul Lahkar* and Frank Riedel†

August 5, 2014

Abstract

We define the logit dynamic for games with continuous strategy spaces and establish its fundamental properties , i.e. the existence, uniqueness and continuity of solutions. We apply the dynamic to the analysis of the Burdett and Judd (1983) model of price dispersion. Our objective is to assess the stability of the logit equilibrium corresponding to the unique Nash equilibrium of this model. Although a direct analysis of local stability is difficult due to technical difficulties, an appeal to finite approximation techniques suggest that the logit equilibrium is unstable. Price dispersion, instead of being an equilibrium phenomenon, is a cyclical phenomenon. We also establish a result on the Lyapunov stability of logit equilibria in negative definite games.
Keywords: Price dispersion; Evolutionary game theory; Logit dynamic.

JEL classification: C72; C73; L11.

*Department of Economics, Ashoka University, Rajiv Gandhi Education City, Kundli, Haryana 131028, India. email: ratul.lahkar@ashoka.edu.in.

†Institute for Mathematical Economics, Bielefeld University, Postfach 100131, 33501 Bielefeld, Germany. email: friedel@uni-bielefeld.de.

1 Introduction

Evolutionary game theory seeks to provide dynamic foundations to equilibrium behavior. The two most well known evolutionary dynamics are the replicator dynamic (Taylor and Jonker, 1978) and the logit dynamic (Fudenberg and Levine, 1998; Hofbauer and Sandholm, 2002, 2007). The more usual approach in this field of research is to consider games with a finite number of strategies and assess whether dynamics converge to Nash equilibria in such games. Alongside this approach, however, there has also emerged a significant literature that seeks to extend such evolutionary dynamics to games with continuous strategy sets. For example, the continuous strategy version of the replicator dynamic has been studied by Bomze (1990, 1991), Oechssler and Riedel (2001,2002), Cressman (2005), Cressman and Hofbauer (2005), and Cressman, Hofbauer and Riedel (2006). Hofbauer et al. (2009) develop the infinite dimensional extension of the Brown-von Neumann-Nash (BNN) dynamic. There has, however, been a significant gap in this literature on infinite dimensional dynamics. The continuous strategy version of the logit dynamic has remained relatively unexplored. This paper seeks to address this lacuna and provide foundations to the logit dynamic for games with continuous strategy sets.

The interest in the logit dynamic in the evolutionary literature emerges from the fact that it preserves a close approximation of the canonical game theoretic best response behavioral model while still being amenable to analysis using standard ODE techniques. Since the finite dimensional logit dynamic has been extensively studied, we believe it is a worthwhile exercise to establish its continuous strategy version. Our more general interest in infinite dimensional evolutionary dynamics is due to two reasons. First, these dynamics are not straightforward extensions of their finite dimensional counterparts. Instead, their analysis requires us to address certain subtle mathematical and technical issues whose resolution offers further insight into how these dynamics behave. For example, since the state space of continuous strategy evolutionary dynamics is the space of probability measures, defining these dynamics require us to make a non–trivial choice of topology with which to define neighborhoods in this space. As Oechssler and Riedel (2002) show, this has implications on important issues like the choice of stability criterion to assess convergence of solution trajectories to equilibria. Second, economic problems are usually modeled as games with continuous strategy sets. It is to be hoped, therefore, that the development of a theory of infinite dimensional evolutionary dynamics will make evolutionary game theory more amenable for application to economic problems.

In fact, the second major objective of this paper is to apply the infinite dimensional logit dynamic to a problem of economic interest. This is the problem of price dispersion which refers to the phenomenon of different sellers selling the same homogeneous commodity at different prices. Various theoretical models explain price dispersion as a mixed strategy Nash equilibrium, called a *dispersed price equilibrium*. The mixed equilibrium emerges due to the presence of some imperfection in the market; for example, the lack of information or the presence of search costs which prevent consumers from knowing about the lowest prevailing price.[1] Therefore, sellers can earn more by simply selling to the fraction of consumers who are willing to pay a higher price instead of undercutting each other to marginal cost as in Bertrand competition.

In this paper, we consider specifically the Burdett and Judd (1983) model of price dispersion. There are two types of consumers in this model, one more informed and the other less informed. The proportion of each type is exogenously fixed.[2] Sellers choose prices from a continuous set. Since the first type of consumer does not search for a lower price, the model does not admit a Bertrand equilibrium. Instead, there exists a dispersed price equilibrium that specifies the proportion of sellers choosing any specific subset of prices.

The interpretation of price dispersion as a mixed equilibrium is a static explanation requiring instantaneous coordination on that equilibrium. It is, however, reasonable to seek a dynamic explanation of this phenomenon as it is not immediately clear how a large population of sellers will be able to achieve such coordination on a probability distribution. It is precisely with this motivation that Lahkar (2011), building on previous work by Hopkins and Seymour (2002), seeks an evolutionary explanation of price dispersion. Using the logit dynamic, Lahkar (2011) shows that the distribution of prices in the population of sellers in the Burdett and Judd (1983) model converges not to a mixed equilibrium, but to a limit cycle. Hence, while there is persistent price dispersion, that paper suggests that price dispersion is a cyclical rather than an equilibrium phenomenon. This conclusion is consistent with empirical and experimental evidence on price dispersion as discovered by, for example, Lach (2002) and Cason et. al (2005).[3]

However, due to the technical apparatus of the continuous logit dynamic not being available, Lahkar (2012) uses the logit dynamic for games with a finite strategy set. This, in turn, required the analysis in that paper to be based not on the original Burdett and Judd (1983) model but on a finite approximation of that model. This approach does provide some significant insights into the dynamic behaviour of price dispersion. However, in certain respects, this approach is unsatisfactory. First, as Lahkar (2011) shows, in the finite dimensional version of the Burdett and Judd (1983) model, we encounter multiple mixed equilibria and, thereby, lose the elegance of the unique equilibrium prediction of the original model. Second, the finite dimensional evolutionary approach does not actually explain the dynamic nature of price dispersion in the original continuous strategy set Burdett and Judd (1983) model. Intuitively, it is plausible that the instability result of this approach extends to the continuous model. But to establish this conclusion rigorously, one needs to conduct an evolutionary analysis directly in the setting of the original model. We seek to complete this exercise in this paper using the continuous logit dynamic. Not only does it remedy a technical drawback in the analysis in Lahkar (2012), but it also establishes a framework for the

[1] See Lahkar (2011) for a more comprehensive review of the literature on price dispersion.

[2] Burdett and Judd (1983) use this simple case of exogenously fixed consumer types as a means to prove their more general dispersed price equilibrium in which the two types of buyers and their proportions emerge endogenously as an equilibrium condition. We focus on this simpler version for reasons of tractability.

[3] The choice of the logit dynamic in Lahkar (2011) is motivated by the experimental findings of Cason et. al (2005) that this dynamic provides a better explanation of divergence from mixed equilibria than the replicator dynamic, which was used in Hopkins and Seymour (2002).

application of this dynamic to other economic models with continuous strategy sets; for example, the model of congestion studied in Hu (2011), which we elaborate upon further later in this section.

Our first task in this paper, therefore, is to rigorously define the logit dynamic for games with continuous strategy spaces. In such games, a population state is a probability measure that specifies the distribution of the use of the various strategies in the population. The logit dynamic for such games is based on the logit choice measure which is a probability measure that maximizes a perturbed version of the players' payoffs, where the perturbation arises from the entropy of the strategy the agent uses (Mattsson and Weibull, 2002). The logit dynamic moves the current population state towards the logit choice. Rest points of this dynamic are fixed points of the logit choice measure. We call such a fixed point a logit equilibrium. We note that the basic definition of the continuous logit dynamic has already been provided in Lahkar (2007). That work is, however, of a preliminary nature since it did not contain a very rigorous analysis of the properties of the dynamic. This paper rectifies this problem and establishes a set of technical results that provides precise foundations to the dynamic; namely, the existence of a logit equilibrium, its convergence to a Nash equilibrium as the perturbation factor becomes small, and the the existence, uniqueness and continuity of solution trajectories of the logit dynamic. Given the abstract nature of the state space, showing these properties require us to make appropriate choice of topologies in the space of probability measures. Following earlier work such as Oechssler and Riedel (2001, 2002) and Hofbauer et al. (2009), we choose between the strong and weak topologies in the set of signed measures for establishing these results.

We then apply the dynamic to negative definite games with continuous strategy sets. We extend the Lyapunov function used in Hobauer and Sandholm (2007) for finite strategy games to our setting and establish the Lyapunov stability of the unique logit equilibrium in such games. Once again, this result hinges on the appropriate choice between the strong and weak topologies in the state space. Apart from being of theoretical interest, this result is also applicable to some existing work by Hu (2011) who applies the logit dynamic, as defined in Lahkar (2007), to a continuous strategy model of congestion. This model satisfies negative definiteness and anticipates some of our methods since it applies the same Lyapunov function that we use. It concludes that logit equilibrium in the model is asymptotically stable, but does so without exploring the technical issues related to the definition of the dynamic or the notion of stability used. By establishing the technical foundations of the dynamic, this paper puts Hu's (2011) application on a sounder footing.

The Burdett and Judd (1983) model is, however, not a negative definite but a positive definite game. Intuition derived from known results in finite strategy positive definite games suggest that the mixed equilibrium in this game should be evolutionarily unstable. It is, however, an extremely challenging task to prove this rigorously for continuous strategy games, and one that we have been able to complete only partially. In finite dimensional evolutionary analysis, the standard way to establish instability of an equilibrium is to linearize the dynamic at that equilibrium and check if any eigenvalue has a positive real part. Unfortunately, as far as we know, there is no such clearly defined theory of spectral analysis that can be applied to to a dynamic defined in a measure theoretic setting. Nor have we been able to develop such a theory. Instead, we follow an alternative strategy where we show that the solution trajectories of the logit dynamic in the continuous strategy Burdett and Judd (1983) model can be arbitrarily well approximated by the solution trajectories of the dynamic in the finite analogue of this model. We know from Lahkar (2012) that mixed equilibria are unstable in the finite strategy model. The approximation result then suggests that the same is true in the continuous strategy model.

We need to note that our conclusion only suggests instability of the dispersed price equilibrium and is not conclusive proof of its instability. This also raises the question of whether attempts to develop the continuous logit dynamic has been redundant since in our primary application, we have had to rely on a finite approximation to arrive at our conclusions. We argue that this is not so. For the finite approximation technique to be meaningful, it is necessary that a continuous strategy dynamic exists as a limiting case for the finite approximations. Hence, providing rigorous foundations to the continuous strategy version of the logit dynamic is a necessary prelude to applying any finite approximation method. In fact, the very question of whether any logit equilibrium, including the one in the Burdett and Judd (1983) model, is stable or not acquires substance only after we have able to establish the fundamental properties of the continuous logit dynamic. Our approach also suggests that pending the development of a complete theory of the spectral analysis of such measure theoretic dynamics, this is perhaps the only feasible way of establishing instability in continuous strategy models. We leave the spectral analysis question as a topic for further research.

Finally, we note that this paper is most closely related to Perkins and Leslie (2014). The authors, in work done simultaneously but independently from our paper, consider a model of stochastic fictitious play in a finite player game with a continuous strategy set. They show that the evolution of mixed strategies of the players in this framework is determined by the logit dynamic. While the two papers have different motivations, they are also complementary in certain respects. First, they combine our dynamical systems results with stochastic approximation theory to develop the learning variant of the logit dynamic. Hence, their approach provides another way to provide microfoundations to the dynamic. Second, they extend our single population logit dynamic to a multi-player game or, equivalently, to a multiple population game. Third, they extend our result on Lyapunov stability of the logit equilibrium in a negative definite game to global asymptotic stability. Perkins and Leslie (2014), therefore, extend our analysis of the logit dynamic in important directions and establishes it more robustly as a dynamic for the analysis of games with continuous strategy sets.

The rest of the paper is structured as follows. In Section 2, we introduce populations games with continuous strategy spaces and illustrate such games using the Burdett and Judd (1983) model. Section 3 introduces the logit dynamic in a measure theoretic setting. In Section 4, we establish the existence of a logit equilibrium in a population game. Section 5 proves the existence, uniqueness and continuity of the logit dynamic with respect to initial conditions. In Section 6, we discuss the microfoundations of this dynamic using the idea of maximization of the perturbed payoff. Section 7 establishes the Lyapunov stability of logit equilibrium in negative definite games. In Section

8, we apply the finite approximation technique to arrive at our conjecture on instability of logit equilibrium in the Burdett and Judd (1983) model. Some proofs are in the Appendix.

2 Population Games

We consider a strategic situation in which a continuum of agents, called a *population*, play a game with a strategy set S which is a compact subset of $\mathbf{R}$. Let $\mathcal{B}$ be the Borel σ-algebra on S. We identify a population state with a probability measure over the measurable space $(S, \mathcal{B})$. We denote the set of population states by Δ. The population state describes the distribution of agents over S. For example, if P is a population state, then $P(A)$, $A \in \mathcal{B}$, denotes the proportion or mass of agents who play strategies in A. If P is absolutely continuous with respect to the Lebesgue measure, we denote by p its density function. In that case, $P(A) = \int_A p(x)dx$.

We characterize the population game by the bounded measurable payoff function $\pi : S \times \Delta \to \mathbf{R}$ (where, as usual, Δ is taken to be a subset of the Banach space of finite σ–additive measures on the measurable space $(S, \mathcal{B})$, see Section 4 for more details). The scalar $\pi(x, P)$ denotes the payoff of an agent who plays strategy x given the population state P. A population state P^* is a Nash equilibrium if for all $x \in \text{support}\, P^*$ and $y \in S$ $\pi(x, P^*) \geq \pi(y, P^*)$[4].

2.1 The Burdett and Judd Model

We present the Burdett and Judd (1983) model of price dispersion as an example of an economic problem that can be modeled as a population game. Each seller in a population sells one unit of the same good at a given time t. Sellers choose a price independently from the set $[0, 1]$ at which to sell the product. We interpret 0 as the cost of production for the sellers and 1 as the common reservation price of consumers which is known to the sellers. Consumers buy the product after taking a sample of a certain number of prices that sellers are charging. We assume that there are two types of consumers, the proportion of each type being exogenously fixed. A proportion q_1, $0 < q_1 < 1$, samples just one price and buy the product at that price. The remaining proportion, denoted $q_2 = 1 - q_1$, pick out two buyers at random and buy from the one charging the lower price. If both prices are equal, we assume they randomize equally between the two sellers.[5]

Let P represent the population state of sellers and let F_P be its distribution function. Then, given our assumptions, the expected payoff of a seller who charges price x is

$$\pi(x, P) = x \left(q_1 + 2q_2 \left(\frac{P(\{x\})}{2} + (1 - F_P(x)) \right) \right). \tag{1}$$

[4]We use this useful characterization (and slight abuse of terminology) of symmetric Nash equilibria (P^*, P^*) here and in the sequel. Strictly speaking, a Nash equilibrium consists of a *pair* of strategies (P, Q) that satisfy the best reply property, of course.

[5]In this paper, we focus on games with only one population. In this simple version of the Burdett and Judd model, the relevant population is the population of sellers. However, in the more general version with endogenous consumer behaviour (footnote 2), we would need to consider the dynamics of both the sellers' and the buyers' populations. See Lahkar (2011) for the finite dimensional evolutionary analysis of both versions.

The expected mass of consumers who sample the seller's price is $q_1 + 2q_2$. If the consumer samples only the price of this particular seller, he obtains a payoff of x. If the consumer samples two prices, including this particular one, then $\left(\frac{P(\{x\})}{2} + (1 - F_P(x)\right)$ represents the probability that the seller actually manages to sell to the consumer. This gives to the expected payoff in (1). If P is absolutely continuous, (1) simplifies to

$$\pi(x, P) = x\left(q_1 + 2q_2\left(1 - F_P(x)\right)\right). \tag{2}$$

Burdett and Judd (1983) show that this model has a unique mixed strategy equilibrium P^* with distribution function F_{P^*} given by

$$F_{P^*}(x) = \begin{cases} 1 - \frac{q_1(1-x)}{2(1-q_1)x}, & \text{for } x \in [\frac{q_1}{2-q_1}, 1] \\ 0, & \text{for } x \in [0, \frac{q_1}{2-q_1}] \end{cases}. \tag{3}$$

To verify that (3) is a Nash equilibrium, note that $\pi_x(P^*) = q_1$ for $x \in [\frac{q_1}{2-q_1}, 1]$ and $\pi_x(P^*) < q_1$ for $x \in [0, \frac{q_1}{2-q_1}]$. For the proof of the uniqueness of equilibrium, see Burdett and Judd (1983). For later application, we note that P^* is not a fully mixed equilibrium in that it does not have full support on S. This feature, in turn, emerges from the fact that prices lower than $\frac{q_1}{2-q_1}$, the lowest price in the support of the equilibrium, are dominated by 1. To see this, note that the maximum payoff from strategy 1 is q_1, whereas the maximum payoff from any other strategy x is $x(q_1 + 2q_2) = x(2 - q_1)$.[6]

The Burdett and Judd (1983) model is a static model in which the mixed equilibrium (3) provides the explanation of price dispersion. We seek to analyze this model from a dynamic perspective. Let $P(t)$ denote the population state of sellers at time t.[7] We assume that sellers (or agents in a more general game) undertake a strategy revision process (see Section 6) that generates the logit dynamic which changes the population state $P(t)$ over time. We show at the end of our analysis in Sections 7 and 8 that $P(t)$ diverges away from the equilibrium (3) rendering it unstable. Therefore, this equilibrium cannot explain price dispersion in our dynamic model. Instead, we argue that price dispersion manifests itself as a limit cycle.

3 The Logit Dynamic

Our objective is to define an evolutionary dynamic on Δ for the population game π such that from any initial population state $P(0) \in \Delta$, we are able to trace the change in the population

[6] The equilibrium (3) holds for the case $0 < q_1 < 1$. If $q_1 = 0$ so that all consumers sample two prices, the unique equilibrium is the Bertrand equilibrium where all sellers charge the marginal cost zero. If all consumers sample just one price, i.e. $q_1 = 1$, then the unique equilibrium is the monopoly equilibrium where all sellers charge the reservation price of one.

[7] For notational clarity, we note that we typically use t to denote time and x to denote any particular strategy in S. Therefore, when we write $P(t)$, we refer to the population state at time t whereas $P([x])$ refers to the population mass of strategy x under population state P. Throughout the paper, we occasionally refer to the initial population state $P(0)$. From the context, it should be clear that this refers to the population state at time 0 and not the population mass of strategy 0.

state over time. An evolutionary dynamic is a vector field that changes one probability measure in Δ into another probability measure in Δ. Therefore, the dynamic is a component of the linear span of Δ, $\mathcal{M}^e(S,\mathcal{B})$, which is the space of all finite and signed measures on $(S,\mathcal{B})$.[8] As in the finite dimensional case, we can interpret a evolutionary dynamic as a differential equation system $\dot{P}(A) = V(P)(A)$ for $A \in \mathcal{B}$. If the current population state is $P(t)$, then $V(P)(A)$ gives the direction and magnitude of the instantaneous change at time t of the mass of population using strategies in A.[9] If $\dot{P}$ is absolutely continuous with respect to the Lebesgue measure, then we denote by $\dot{p}$ its density function. In this case, $\dot{P}(A) = \int_A \dot{p}(x)dx$.

To be admissible as an evolutionary dynamic, we require that from every initial condition $P(0) \in \Delta$, the differential equation $\dot{P} = V(P)$ admits a solution, preferably unique, $\{P(t)\}_{t\geq 0}$ with $P(t) \in \Delta$, for all $t \geq 0$. In the infinite dimensional context, the existence of unique solutions is a technically demanding condition that we establish in Section 5. However, once we establish this property, the second condition that $P(t) \in \Delta$, for all $t \geq 0$, follows readily if $V(P)(S) \in T\Delta$ where

$$T\Delta = \{\mu_Z \in \mathcal{M}^e(S,\mathcal{B}) : \mu_Z(S) = 0\} \tag{4}$$

is the *tangent space* of Δ. This property that $V(P)(S) \in T\Delta$ is called *forward invariance* and it ensures that P_t is a probability measure at all time. Given $\mu_Z \in T\Delta$, Z denotes its distribution function. If μ_Z is absolutely continuous with respect to the Lebesgue measure, we denote its density function by z.

We now define the logit dynamic for the population game π. Let P be the present population state and let $A \in \mathcal{B}$. Then, under the logit dynamic, the rate of change in the mass of $P(A)$ is given by

$$\dot{P}(A) = L_\eta(P)(A) - P(A), \tag{5}$$

where $L_\eta : \Delta \to \Delta$ is the logit choice measure defined, given parameter $\eta > 0$, by

$$L_\eta(P)(A) = \int_A \frac{\exp\left(\eta^{-1}\pi(x,P)\right)}{\int_S \exp\left(\eta^{-1}\pi(y,P)\right)dy}dx. \tag{6}$$

Since both $L_\eta(P)$ and P are probability measures, $\dot{P} \in T\Delta$.

We denote by $\Delta^{AC} \subset \Delta$ the set of probability measures that are *uniformly equivalent* with respect to the Lebesgue measure. These are measures that have a strictly positive, bounded, and bounded away from zero density function. With our assumptions, $L_\eta(P) \in \Delta^{AC}$ with density function (see also Mattsson and Weibull, 2002)

$$l_\eta(P)(x) = \frac{\exp\left(\eta^{-1}\pi(x,P)\right)}{\int_S \exp\left(\eta^{-1}\pi(y,P)\right)dy}. \tag{7}$$

[8] Recall that μ is a finite signed measure on $(S,\mathcal{B})$ if there are two finite measures ν^1 and ν^2 such that for all $A \in \mathcal{B}$, $\mu(A) = \nu^1(A) - \nu^2(A)$.

[9] Strictly speaking, we should use the notation $V_\pi(P)(A)$ since the dynamic depends upon the payoff function π. However, since the underlying game is usually clear from the context, we dispense with this extra piece of notation.

We note that even if P does not have a density function, $L_\eta(P)$ admits this density function. If P has density function p, then $\dot{P}$ also admits the density function

$$\dot{p}(x) = l_\eta(P)(x) - p(x). \tag{8}$$

Since $l_\eta(P)(x)$ takes positive value over all $x \in S$, all sets of positive Lebesgue measure receive positive probability under the logit choice measure. This has a clear analogy in the finite dimensional case where logit choice function puts positive probability on all strategies. Section 6 elaborates the logit choice measure further by generating it through a process of perturbed optimizing behaviour by individual agents.

We call a fixed point of $L_\eta(P)$ a *logit equilibrium* of the population game π. Clearly, the set of rest points of the dynamic coincide with the set of fixed points of the logit choice measure.

In the finite dimensional case, the logit dynamic has some technical and behavioral features that make it preferable to the best response dynamic. For example, the logit dynamic is a smooth differential equation that makes it amenable to analysis by using standard techniques, whereas the best response dynamic is a differential inclusion that is difficult to analyze. In the infinite dimensional case, this advantage of the logit dynamic is even more striking: the logit choice measure is always well defined whereas the proper best response may not even exist. For example, in the Burdett and Judd model in Section 2.1, suppose $P([x]) = 1$ for some $x > \frac{q_1}{2-q_1}$. These strategies are not dominated by strategy 1. There is no best response to this population state although any strategy sufficiently close to x is a better response. However, the logit choice measure is well defined since the density function exists at all points in S. Moreover, it also manages to capture the intuition of the best response. For η sufficiently small, the logit choice measure assigns most of the probability mass on a small interval immediately to the left of x.

4 Existence of Logit Equilibrium

Before applying the logit dynamic to any problem of interest, we need to address a set of important technical questions. The first is whether a logit equilibrium exists in any population game. The second is the existence of, preferably unique and continuous, solution trajectories for the logit dynamic. We analyze the first question in this section while deferring the second to the next.

The resolution of these questions depends upon the choice of topology in Δ. In the finite dimensional case, the choice of topology is not of much consequence since all norms are equivalent. However, in the infinite dimensional case, this choice is important since the structure of the neighborhood of a probability measure can depend upon the topology we choose.

We use three topologies in this paper. The first is the strong topology or the variational topology. This is the topology induced by the variational norm on $\mathcal{M}^e(S, \mathcal{B})$, given by $\|\mu\| = \sup_g |\int_S g d\mu|$ where g is a measurable function $g : S \to \mathbf{R}$ such that $\sup_{x \in S} |g(s)| \leq 1$. If P, Q are two probability measures, then the distance between them under this norm is (Shiryaev, 1995)

$$\|P-Q\| = 2\sup_{A\in\mathcal{B}} |P(A) - Q(A)|.$$

The strong topology turns the vector space of all signed finite measures into a Banach space. This allows us to consider ordinary differential equations is such a measure space and establish some fundamental properties of such a dynamic, for example the existence and uniqueness of solution trajectories. This topology finds application in Oechssler and Riedel (2001) in the analysis of the replicator dynamic, and in Hofbauer et al. (2009) in the analysis of the BNN dynamic.

The second topology we use is the weak topology or the topology induced by convergence in distribution. Intuitively, under this topology, two probability measures are close to each other if their distribution functions are close to each other. The weak topology has the advantage that certain sets which are not compact under the strong topology become so under this topology. For example, the set Δ itself is not compact under the strong topology but is so under the weak topology. This topology, therefore, is useful in applications that require compactness, for example, the existence of equilibria. The weak topology also induces a richer structure to a neighborhood of a probability measure. States which are not close to one another under the strong topology, for example two Dirac measures δ_x and $\delta_{(x+\varepsilon)}$, $x \in S \subseteq R$ and ε small, are close under the weak topology. It can, therefore, be a more appropriate topology when one considers stability issues of a dynamic; although, as we shall see, we cannot apply it to the stability analysis of the logit dynamic in the Burdett and Judd (1983) model.

On the whole vector space $\mathcal{M}^e(S,\mathcal{B})$, there is no norm that metrizes the weak topology (see, for example, Billingsley, 1999). Therefore, $\mathcal{M}^e(S,\mathcal{B})$ is not a Banach space under this topology. So this topology cannot be applied, for example, to define an evolutionary dynamic on $\mathcal{M}^e(S,\mathcal{B})$. On the other hand, for positive measures, and on Δ in particular, several norms lead to the same topology as the weak one. One such norm is the Levy–Prohorov norm under which the distance between $P, Q \in \Delta$ is

$$\rho_L(P,Q) = \inf\{\varepsilon > 0, Q_F(x-\varepsilon) - \varepsilon \le P_F(x) \le Q_F(x+\varepsilon) + \varepsilon, \text{ for all } x \in S\}.$$

Another norm that metrizes the weak topology on Δ is the Bounded Lipschitz (BL) norm. We refer the reader to, for example, Oechssler and Riedel (2002) and Hofbauer et al. (2009) for applications of this norm.

The final notion of distance in Δ that we use is the Kolmogorov (or uniform) metric for probability measures. Under the Kolmogorov metric, the distance between two probability measures P, Q is

$$\rho_K(P,Q) = \sup_{x\in S} |F_P(x) - F_Q(x)|. \tag{9}$$

The Kolmogorov metric does not metrize the weak topology. However, for certain applications like the Burdett and Judd (1983) model in which distribution functions play a prominent role, this norm is analytically much more convenient than norms that metrize the weak topology on Δ (see Section 8). Moreover, it is known that $\rho_L(P,Q) \leq \rho_K(P,Q)$. Therefore, convergence under the Kolmogorov metric implies convergence under the weak topology.[10]

Using the Schauder fixed point theorem, we now establish the existence of a logit equilibrium in Δ. This theorem is an extension of Brouwer's fixed point theorem to infinite dimensional settings. The Schauder fixed point theorem requires Δ to be a compact set. Therefore, we need to impose the weak topology on Δ to establish this theorem. Let $LE\,(\pi, \eta)$ be the set of logit equilibria of the population game π, given parameter $\eta > 0$.

Theorem 4.1 *Consider a population game π. If the mapping $P \mapsto L_\eta(P)$ is continuous with respect to the weak topology, then the set $LE\,(\pi, \eta)$ is non-empty.*

Proof. Under the weak topology, $\mathcal{M}^e(S, \mathcal{B})$ is a locally convex, Hausdorff, topological vector space (Billingsley, 1999) and Δ is a compact, nonempty, and convex subset of $\mathcal{M}^e(S, \mathcal{B})$. The existence of a logit equilibrium then follows from the Schauder fixed point theorem. ■

We denote a logit equilibrium as $\tilde{P}_\eta$ and note that it has has full support since $l_\eta(P)(x) > 0$ for all $x \in S$. The following theorem shows that if a family of logic equilibria converges, then it converges to a Nash equilibrium. Therefore, the stability analysis of logit equilibria provides a strong assessment of the dynamic properties of the corresponding Nash equilibrium, at least for η sufficiently small. The proof of the theorem is in Appendix A.1.

Theorem 4.2 *Let $\{\tilde{P}_\eta\}_{\eta>0}$ be a family of logit equilibria in the population game π. Let $\pi(x, P)$ be continuous in the Euclidean norm for x and in the weak topology for P.*

The family has accumulation points with respect to the weak topology. Each accumulation point P^* is a Nash equilibrium.

We now apply Theorem 4.1 to the Burdett and Judd (1983) model defined by the payoff function (1). While doing so, we need to note that (1) is not continuous with respect to the weak topology. This is due the presence of the term $P([\{x\}])$ in the expression of $\pi(x, P)$ in case the distribution function F_P is discontinuous at the strategy $x \in [0, 1]$. However, the discontinuity of the payoff function does not affect the continuity of the mapping $P \mapsto L_\eta(P)$ with respect to the weak topology. A distribution function can be discontinuous at at most a countable number of points. Such a set of points is of measure zero and measure zero sets do not affect the logit choice measure. We, therefore, have the following proposition. The proof of the proposition, in Appendix A.1, involves working with an auxiliary payoff function $\phi(x, P) = x(q_1 + 2q_2(1 - F_P(x)))$ which is identical to (1) everywhere except a set of measure zero.

Proposition 4.3 *In the Burdett and Judd model, the mapping $P \mapsto L_\eta(P)$ is continuous with respect to the weak topology. Thus there exists a logit equilibrium in the Burdett and Judd model*

[10]For more on the relation between various metrics on the space of probability measures, see, for example, Gibbs and Su (2002).

with payoff function (1). Moreover, if $\{\tilde{P}_\eta\}_{\eta>0}$ is a sequence of logit equilibria in this model, then $\lim_{\eta\to 0}\tilde{P}_\eta = P^*$, where P^* is its unique Nash equilibrium (3).

We note that the final claim of this proposition about the convergence of the sequence of logit equilibria to the Nash equilibrium is an implication of Theorem 4.2. In principle, a family of logit equilibria need not converge. But if there is a unique Nash equilibrium, as in the Burdett and Judd (1983) model, then one also obtains convergence because there is only one possible accumulation point.

5 Logit Dynamic: Existence, Uniqueness and Continuity of Solution

For the logit dynamic to be a meaningful description of population behavior, it is necessary that a forward invariant solution exists for the dynamic from every initial population state in Δ. To enable a precise prediction of the path of social evolution, each initial point should generate a unique solution trajectory. Since a initial state can be feasibly known only with approximate accuracy, it is also desirable that solution trajectories are continuous with respect to initial conditions.

In this section, we establish these fundamental properties for the logit dynamic. We first extend the logit dynamic to the entire vector space $\mathcal{M}^e(S,\mathcal{B})$ and prove the existence of a unique solution from every initial state in this space. We then appeal to forward invariance to conclude that solution trajectories that start in Δ remain in Δ forever. Our approach, in turn, requires a norm on $\mathcal{M}^e(S,\mathcal{B})$. We, therefore, use the strong topology on $\mathcal{M}^e(S,\mathcal{B})$. We condition this result on the Lipschitz continuity of the payoff function π. We extend π to the set

$$\mathcal{M}_2 = \left\{P \in \mathcal{M}^e\left(S,\mathcal{B}\right) : \|P\| \le 2\right\}, \tag{10}$$

and define Lipschitz continuity on $\mathcal{M}_2$ with respect to the variational norm as follows.

Definition 5.1 *The extended payoff function $\pi : S \times \mathcal{M}_2 \to \mathbf{R}$ is Lipschitz continuous on $\mathcal{M}_2$ with respect to the variational norm if there exists a constant K such that*

$$|\pi(x,P) - \pi(x,Q)| \le K\|P - Q\|, \tag{11}$$

for all $P, Q \in \mathcal{M}_2$, all $x \in S$.

We also define the semiflow of the logit dynamic $\xi : \Delta\times[0,\infty] \to \Delta$, where $\xi(P,t) = P(t)$ denotes the population at time t when logit dynamic starts from $P = P(0)$. The continuity of the semiflow with respect to $P(0)$ implies that solutions to the dynamic are continuous in initial conditions. With this definition, we can now state the following theorem on the fundamental properties of the logit dynamic. The second part of the theorem applies these properties to the Burdett and Judd (1983) model (1). The proof of the theorem is in Appendix A.2.

Theorem 5.2 *Let the extended payoff function* $\pi : S \times \mathcal{M}_2 \to \mathbf{R}$ *be measurable, bounded and Lipschitz continuous on* $\mathcal{M}_2$ *with respect to the variational norm on* $\mathcal{M}^e(S, \mathcal{B})$*. Then, from each initial condition* $P = P(0) \in \Delta$*, there exists a unique solution* $P(t) \in \Delta$ *of the logit dynamic defined by the ordinary differential equation (5) for* $t \in [0, \infty)$*. Further, the semiflow* ξ *is continuous with respect to the variational norm.*

The following corollary applies Theorem 5.2 to the Burdett and Judd (1983) model.

Corollary 5.3 *In the Burdett and Judd (1983) model characterized by payoff function (1), the logit dynamic admits a unique solution from every initial condition* $P(0) \in \Delta$*. Moreover, the semiflow of the dynamic is continuous with respect to the variational norm.*

We have noted earlier that if the population state P has a density function p, then the logit dynamic can be written in its density function form $\dot{p}(x)$ given by (8). Together with Theorem 5.2, this allows us to conclude that the logit dynamic is invariant on Δ^{AC}. Therefore, instead of the sequence of probability measures $\{P(t)\}_{t\geq 0}$, we can equivalently express the solution trajectory in terms of the sequence of density functions $\{p(t)\}_{t\geq 0}$. We formalize this conclusion in the following corollary.

Corollary 5.4 *Let the initial population state* $P(0)$ *be absolutely continuous with density function* $p(0)$*. Let* $\{P(t)\}_{t\geq 0}$ *be the solution trajectory of logit dynamic in the population game* π*. Then, for all* $t > 0$*,* $P(t)$ *admits a density function* $p(t)$*. Further, we can describe the evolution of* $p(t)$ *with the differential equation* $\dot{p}(x) = l_\eta(P)(x) - p(x)$ *where* $l_\eta(P)(x) = \frac{\exp\left(\eta^{-1}\pi(x,P)\right)}{\int_S \exp(\eta^{-1}\pi_y(P))dy}$.

6 Microfoundations of the Logit Dynamic

Cheung (2014) provides a general derivation of evolutionary dynamics in population games with continuous strategy space by using the notion of a *mean dynamic*. Let $B = B(S, \mathcal{B})$ be the vector space of bounded measurable functions. We write $\pi(P)$ for the bounded measurable function from S to the reals with $\pi(P)(x) = \pi(x, P)$. Note that $\pi(P) \in B$. We now define the conditional switch rate ρ as the bounded measurable function

$$\rho : B \times \Delta \times S \times S \to \mathbf{R}_+$$

such that for a given population state P, $\rho(\pi(P), P, x, y)$ is proportional to the probability with which an agent who is currently playing strategy $x \in S$ switches to another strategy in $y \in S$. Let $\mu \in \mathcal{M}^e(S, \mathcal{B})$ be the rate at which a revising agent chooses the various candidate strategies. Together, the conditional switch rate ρ and the measure μ describe the *revision protocol* that players use to revise strategies. Cheung (2013) then defines the mean dynamic as the following differential equation on $\mathcal{M}^e(S, \mathcal{B})$:

$$\dot{P}(A) = \int_S \int_A \rho\left(\pi(P), P, y, x\right) \mu(dx) P(dy) - \int_S \int_A \rho\left(\pi(P), P, x, y\right) P(dx) \mu(dy), \tag{12}$$

for all $A \in \mathcal{B}$. The first expression on the right hand side of this equation rate of inflow of mass into A. The second expression gives the gross outflow from A. The difference represents the change in the mass of A. By specifying the conditional switch rate ρ and the measure μ, we may then obtain different evolutionary dynamics. As examples, we first present the replicator dynamic and then, the logit dynamic.

Example 6.1 *Let* $\rho(\pi(P), P, y, x) = (K - \pi(y, P))$, *where* $K > \max_{y \in S, P \in \Delta} \pi(y, P)$, *and* $\mu(dx) = P(dx)$. *This revision protocol is the extension of Björnerstedt and Weibull's (1996) pure imitation driven by dissatisfaction protocol to games with continuous strategy space. We may interpret* K *as an aspiration parameter, with the difference* $K - \pi(y, P)$ *measuring the agent's dissatisfaction from playing* y *given population state* P. *The agent then randomly chooses another member of the population and imitates that member's strategy. Applying this revision protocol in (12), we obtain*

$$\begin{aligned}
\dot{P}(A) &= \int_S \int_A (K - \pi(y, P))\, P(dx)P(dy) - \int_S \int_A (K - \pi(x, P))\, P(dx)P(dy) \\
&= P(A) \int_S (K - \pi(y, P))\, P(dy) - \int_S \left(KP(A) - \int_A \pi(x, P)P(dx) \right) P(dy) \\
&= KP(A) - P(A) \int_S \pi(y, P)P(dy) - KP(A) + \int_A \pi(x, P)P(dx) \\
&= \int_A (\pi(x, P) - \bar{\pi}(P))\, P(dx), \qquad (13)
\end{aligned}$$

where $\bar{\pi}(P) = \int_S \pi(y, P)P(dy)$ *is the average payoff under population state* P. *Equation (13) is the replicator dynamic for games with continuous strategy spaces analyzed, for example, in Oechssler and Riedel (2001,2002).*

Example 6.2 *To generate the logit dynamic, we specify* $\rho(\pi(P), P, y, x) = \frac{\exp\left(\eta^{-1}\pi(x,P)\right)}{\int_S \exp(\eta^{-1}\pi(w,P))dw}$ *and* $\mu(dx) = dx$. *Therefore, from (12),*

$$\begin{aligned}
\dot{P}(A) &= \int_S \int_A \frac{\exp\left(\eta^{-1}\pi(x, P)\right)}{\int_S \exp\left(\eta^{-1}\pi(w, P)\right) dw} dx P(dy) - \int_S \int_A \frac{\exp\left(\eta^{-1}\pi(y, P)\right)}{\int_S \exp\left(\eta^{-1}\pi(w, P)\right) dw} P(dx)dy \\
&= \int_A \frac{\exp\left(\eta^{-1}\pi(x, P)\right)}{\int_S \exp\left(\eta^{-1}\pi(w, P)\right) dw} dx \int_S P(dy) - \int_S \frac{\exp\left(\eta^{-1}\pi(y, P)\right)}{\int_S \exp\left(\eta^{-1}\pi(w, P)\right) dw} dy \int_A P(dx) \\
&= \int_A \frac{\exp\left(\eta^{-1}\pi(x, P)\right)}{\int_S \exp\left(\eta^{-1}\pi(w, P)\right) dw} dx - P(A),
\end{aligned}$$

since $\int_S \frac{\exp\left(\eta^{-1}\pi(y,P)\right)}{\int_S \exp(\eta^{-1}\pi(w,P))dw} dy = \int_S P(dy) = 1$.

To complete the microfoundations of the logit dynamic, we need to identify a process through which individual behaviour generates the logit choice measure. Mattsson and Weibull (2002) generates the logit choice through a deterministic perturbation of the expected payoff of a strategy by the entropy of that strategy. Let $q(x)$ be the density function of a probability measure $Q \in \Delta^{AC}$, the set of absolutely continuous probability measures. Consider the entropy of $Q \in \Delta^{AC}$ defined as

$$v(Q) = \eta \int_S \log q(x) q(x) dx. \tag{14}$$

Clearly, $v(Q)$ is only defined for such Q whose density $q(x) > 0$ for all $x \in S$. We then define the agent's *perturbed payoff* $\int_S \pi(x, P) Q(dx) - v(Q)$ from playing Q. The entropy function, therefore, represents a perturbation of the agent's true payoff. As $\eta \to 0$, $v(Q) \to 0$ and, therefore, the perturbed payoff approaches the true payoff. Mattsson and Weibull (2002) show that the logit choice measure is the probability measure in Δ^{AC} that maximizes the perturbed payoff. Therefore,

$$L_\eta(P) = \underset{Q \in \Delta^{AC}}{\operatorname{argmax}} \int_S \pi(x, P) Q(dx) - v(Q). \tag{15}$$

This approach restricts our search for the logit choice measure to Δ^{AC}. The obvious reason for this is that the entropy function, and hence the perturbed payoff, is defined only on Δ^{AC}. However, to complete our microfoundations, we need to extend these definitions to all of Δ, and ascertain that the perturbed payoff maximizer is not a member of $\Delta \setminus \Delta^{AC}$. A standard procedure is to equate the entropy function to infinity on $\Delta \setminus \Delta^{AC}$ (see, for example, Definition 3.20, Föllmer and Schied, 2010). Therefore, $v(\hat{Q}) = \infty$ if $\hat{Q} \in \Delta \setminus \Delta^{AC}$. We may provide an intuitive explanation of this definition as follows. Let us approximate $\hat{Q}$ with a sequence of absolutely continuous probability measures $Q_n \to \hat{Q}$ under the weak topology. Let q_n be the density function of Q_n. We define the entropy function for $\hat{Q}$ as

$$v(\hat{Q}) = \eta \lim_{n \to \infty} \int_S \log q_n(x) q_n dx.$$

We note that if $q_n(x) \to \infty$ over a set of positive Q_n-measure, then $\int_S \log q_n(x) q_n dx \to \infty$. Since $\hat{Q}$ is not absolutely continuous, there exists at least one point x where its distribution function is discontinuous. In the neighborhood of that x which has positive measure under Q_n, $q_n \to \infty$ as $n \to \infty$. Hence, $v(\hat{Q}) \to \infty$.

In any case, this definition makes it clear that no $\hat{Q} \in \Delta \setminus \Delta^{AC}$ can be a maximizer of the perturbed payoff. The perturbed payoff is $-\infty$ on $\Delta \setminus \Delta^{AC}$. Therefore, it is justified that we restrict ourselves to Δ^{AC} as in (15) in maximizing the perturbed payoff to find the logit choice measure.

7 Negative and Positive Definite Games

Negative and positive definite games have been extensively analyzed in the literature on finite dimensional evolutionary dynamics. All standard evolutionary dynamics converge to the unique equilibrium in negative definite games (Hofbauer and Sandholm, 2009). In contrast, in positive definite games, dynamics diverge away from mixed equilibria (Benaim et. al, 2009). In this section, we extend these results for the logit dynamic to games with continuous strategy sets. We then apply these results to the Burdett and Judd (1983) model. Hopkins and Seymour (2002) and Lahkar (2011) show that the finite analogue of the Burdett and Judd (1983) model satisfies positive definiteness. They use this result to conclude instability of dispersed price equilibrium in the finite dimensional model. We expect a similar result to hold for the infinite dimensional case.

We first define negative and positive definite games in the infinite dimensional setting.

Definition 7.1 *The population game π is negative definite if*

$$\int_S (\pi(x,Q) - \pi(x,P))(Q-P)(dx) < 0, \tag{16}$$

for all $Q, P \in \Delta$. If the inequality is reversed, then the game is positive definite. If (16) holds weakly, then the game is negative semi-definite.

We use this definition to establish that the Burdett and Judd model (1983) in Section 2.1 is a positive definite game.[11] We state the result formally in the following proposition. The proof is in Appendix A.3.

Proposition 7.2 *The population game π defined by the Burdett and Judd (1983) payoff function (1) is positive definite.*

We can equivalently define negative or positive definiteness of a game by considering the Gateaux derivative of the payoff function $\pi(x,P)$. The Gateaux derivative of $\pi(x,P)$ in the direction $\mu_Z \in T\Delta$ is defined as

$$D\pi\left(x,P\right)\mu_Z = \lim_{\varepsilon\to 0} \frac{\pi\left(x, P+\varepsilon\mu_Z\right) - \pi\left(x,P\right)}{\varepsilon}. \tag{17}$$

The Gateaux derivative extends the notion of the directional derivative in a finite dimensional space to an infinite dimensional setting. It satisfies all the usual rules of calculating derivatives, including the product rule and the chain rule.[12] If this derivative is well defined, then the population game π is negative definite if

$$\int_S D\pi(x,P)\mu_Z\mu_Z(dx) < 0 \tag{18}$$

for all $P \in \Delta$, $\mu_Z \in T\Delta$, $\mu_Z \neq 0$. If the inequality is reversed, then the game is positive definite. To see that (18) is equivalent to (16), we can take $Q = P + \mu_Z$, for μ_Z small, and apply Taylor's approximation. Hence, $\pi(x,Q) \approx \pi(x,P) + D\pi(x,P)\mu_Z$ so that

$$\int_S (\pi(x,Q) - \pi(x,P))(Q-P)(dx) \approx \int_S D\pi(x,P)\mu_Z\mu_Z(dx).$$

It is important, however, to note that the validity of the definition in (18) may depend upon the topology under consideration. For example, in the Burdett and Judd (1983) model, the payoff function (1) is not continuous under the weak topology. Hence, the derivative $D\pi\left(x,P\right)\mu_Z$ is not

[11] For examples of negative definite games like the "War of Attrition", see Hofbauer et. al. (2009). Hu (2011) studies a class of dynamic congestion games called departure–time choice games. These games satisfy negative definiteness not on entire Δ, but on the subset of absolutely continuous probability measures. The details of the these games are somewhat involved. Hence, we refer the interested reader to the original work by Hu.

[12] See, for example, Atkinson and Han (2009) for more a detailed analysis of the Gateaux derivative.

defined under this topology. Under the strong topology, however, this derivative is well-defined, and, therefore, (18) is valid.

Intuitively, μ_Z represents a mutation away from the current population state P. The derivative $D\pi_x(P)\,\mu_z$ then represents the marginal change in the payoff of strategy x due to the mutation, and $\int_S D\pi(x,P)\mu_Z\mu_Z(dx)$ is the weighted average of these changes. If this integral is negative, then the average change in the payoff of the mutant group in negative. This property, termed "self-defeating externalities" by Hofbauer and Sandholm (2009), generalizes Maynard Smith's (1982) notion of an evolutionary stable strategy (ESS). As shown by Hofbauer et al. (2009), negative definite games have a unique Nash equilibrium. Due to the property of "self-defeating externalities", we expect this unique equilibrium to be evolutionarily stable. In contrast, positive definite games are characterized by "self-improving externalities" in that the mutant group's average payoff changes in a positive direction.

Hofbauer and Sandholm (2007) construct a Lyapunov function to prove stability of logit equilibria in negative definite games with finite strategy sets. We extend their methodology to seek corresponding results in the infinite dimensional setting. First, we establish the uniqueness of logit equilibrium in negative definite games with continuous strategy sets. The proof this lemma is in Appendix A.3.

Lemma 7.3 *If π is a negative definite or a negative semi-definite population game, then it has a unique logit equilibrium.*

We now define the measure theoretic analogue of Hofbauer and Sandholm's (2007) Lyapunov function, $G : \Delta^{AC} \to \mathbf{R}$, in the following lemma and characterize some of its obvious properties without proof. The restriction of the domain of G to Δ^{AC} is due to the fact that the entropy function, which forms a part of the definition of G, is defined only on Δ^{AC}.

Lemma 7.4 *Given the population game π, define $G : \Delta^{AC} \to \mathbf{R}$ as*

$$G(P) = \sup_{Q\in\Delta^{AC}} \left[\int_S \pi(x,P)Q(dx) - v(Q)\right] - \left[\int_S \pi(x,P)p(x)dx - v(P)\right] \tag{19}$$

Clearly, $G(P) \geq 0$ for all $P \in \Delta^{AC}$ and $G(P) = 0$ if and only if P is a logit equilibrium.

Since the domain of G is restricted to Δ^{AC} instead of all of Δ, any result that we obtain directly from the analysis of this function can apply only to Δ^{AC}. With this caveat, we establish a relatively weak stability result for the unique logit equilibrium in negative definite games; if the population starts close to the equilibrium in the strong topology, then it stays close to the equilibrium in the weak topology. The proof of this theorem is in Appendix A.3.

Theorem 7.5 *Let π be a negative definite or a negative semi-definite population game that is Gâteaux differentiable with respect to the variational norm and continuous with respect to the weak topology. Let P^η be its unique logit equilibrium. Then, P^η is stable in the following sense: for* *every $\varepsilon > 0$ there exists $\delta > 0$ such that for initial populations $P_0 \in \Delta^{AC}$ with $\|P_0 - P^\eta\| < \delta$, the trajectory $(P(t))$ with $P(0) = P_0$ satisfies $\rho_L(P(t), P^\eta) < \varepsilon$.*

The conclusion in this theorem is similar to a Lyapunov stability result, but involving the use of two different topologies. Our candidate Lyapunov function G is not continuous with respect to the weak topology. This, in turn, is due to the fact that the entropy function v is not continuous with respect to this topology. This requires the use of the variational norm to evaluate $\dot{G}(P)$. The state space Δ, however, is not compact under the strong topology. Therefore, once we establish that $\dot{G}(P) < 0$, we need the weak topology to define a compact $\varepsilon-$ neighborhood around the logit equilibrium. Since the Lyapunov function is lower semicontinuous with respect to the weak topology (this follows from continuity of π with respect to the weak topology), it attains a minimum in this neighborhood. This allows us to conclude that if a trajectory starts close to the logit equilibrium (in terms of the variational norm), it remains in this weak $\varepsilon-$ neighborhood.

We note that Hu (2011) uses the same Lyapunov function in the stability analysis of a dynamic model of congestion, although without clarifying the technical issues related to its application. We also note that Perkins and Leslie (2014) have extended this result to show convergence of the logit equilibrium in two player zero sum games. Their key innovation is to show that although Δ is not compact under the strong topology, a bounded subset of absolutely continuous probability measures (i.e. a subset of Δ^{AC}) is indeed compact. Using the same Lyapunov function that we are using, they then show convergence under the variational norm to the logit equilibria from initial states in this subset. They then extend this result to all of Δ by showing that from any arbitrary initial state, the logit trajectory comes close to this bounded subset of probability density functions. By appealing to continuity of solutions with respect to initial conditions, they conclude global convergence to the logit equilibrium. While their proof is for two player zero sum games, the method, as they note, can be extended to a more general negative definite or semi–definite game.[13]

The Burdett and Judd (1983) model is, however, as shown in Proposition 7.2, a positive definite game. We cannot use the Lyapunov function to arrive at any definite conclusion about whether a logit equilibrium is stable or not in such games. To get a glimpse of why this is difficult, consider (33) in Appendix A.3. This provides an expression for $\dot{G}(P)$ which is negative for negative definite and semi–definite games. However, for positive definite games, the sign of $\dot{G}(P)$ would depend very subtly on the value of η. Hence, using a Lyapunov function argument, it may not be possible to conclude whether solution trajectories stay near or move away from the logit equilibrium in the Burdett and Judd (1983). Unlike in finite dimensional dynamics, it is also not possible to conduct a local stability analysis of the logit equilibria in the absence of a well developed mathematical theory of spectral analysis of measure theoretic dynamics. Therefore, in the absence of any feasible method of rigorous stability analysis, we consider an approximation of the solution trajectories of

[13]We also refer the reader to Anderson et al. (2004) for an application of the Lyapunov function method to the stability analysis of a logit equilibrium in a potential game. Anderson et al. (2004) derive the logit equilibrium as the rest point of an evolutionary process in which, subject to a normal error, agents revise their strategies in the direction of increasing payoffs. They show that the resulting dynamic process follows the Focker–Planck equation, in which the logit equilibrium is a steady state.

the logit dynamic in the Burdett and Judd (1983) with those in a finite approximation of the model. An appeal to the results in Lahkar (2012) then suggests that the logit equilibrium corresponding to the unique Nash equilibrium in this model is unstable.

8 Price Dispersion as a Limit Cycle

Lahkar (2011) approximates of the continuous strategy set $S = [0,1]$ of the Burdett and Judd (1983) model with the discrete set

$$S^n = \left\{0, \frac{1}{n}, \frac{2}{n}, \cdots, \frac{n-1}{n}, 1\right\}.$$

of $(n+1)$ prices. Let $x_i^n = \frac{i}{n}$, $i \in \{0, 1, \cdots, n\}$ be the $(i+1)-$th price in S^n and

$$\Delta^n = \left\{p^n \in \mathbf{R}_+^{n+1} : \sum_{i=0}^{n} p_i^n = 1\right\}$$

be the set of probability distributions on S^n. Then, the payoff function of this discrete model

$$\pi_i^n\left(p^n\right) = x_i^n \left(q_1 + 2q_2\left(\frac{p_i^n}{2} + \sum_{j>i} p_j^n\right)\right) \tag{20}$$

represents the finite dimensional analogue of the Burdett and Judd (1983) payoff function (1).

Using the fact that (20) represents a positive definite game, Lahkar (2011) shows that for all η sufficiently low, logit equilibria corresponding to the dispersed price equilibria in this discrete model are unstable under the logit dynamic. This conclusion of course relates to the instability of logit equilibria and not the dispersed price equilibria (mixed Nash equilibria) in the model. But since a logit equilibrium converges to a Nash equilibrium as $\eta \to 0$, this result suffices, at least qualitatively, to imply that the population state moves away the dispersed price equilibria. Hence, price dispersion in the finite strategy Burdett and Judd (1983) is not an equilibrium phenomenon. Instead, it seems to be a cyclical phenomenon as simulations in Lahkar (2011) show that the trajectories of the logit dynamic converge to a limit cycle in Δ^n. The cyclical steady states persist even in simulations with reasonably large values of n, for example, $n = 100$. This suggests that price dispersion in the discrete model is cyclical for all n, howsoever large. It therefore, seems, reasonable to expect a similar cyclical pattern even in the infinite dimensional model. Solution trajectories starting near the dispersed price equilibria converges to a limit cycle. Formally, in the measure theoretic context, a limit cycle of the logit dynamic is a solution trajectory $\{P(t)\}_{t\geq 0} \in \Delta$ of dynamic with the property that for every t, there exist $\tau > 0$ such that $P(t+\tau) = P(t)$. A trajectory $Q(t)$ converges to the limit cycle $P(t)$ if, as $t \to \infty$, $Q(t) \to P(t)$ with respect to a given norm.

To provide formal support to this hypothesis, we show that solution trajectories of the logit dynamic in Δ may be approximated with the trajectories of the finite dimensional logit dynamic in Δ^n. We use the Kolmogorov norm for this purpose and our approach broadly follows Oechssler and Riedel's (2002) finite approximation of the infinite dimensional replicator dynamic under the strong and weak topologies.

We begin with the approximation of any arbitrary probability measure P with an appropriate finite dimensional probability distribution. A natural approximation is the distribution $p^n \in \Delta^n$ such that

$$\begin{aligned} p_i^n &= P\left(\left[\frac{i}{n}, \frac{i+1}{n}\right)\right) \text{ for } i = 0, 1, \cdots, n-2, \\ p_{n-1}^n &= P\left(\left[\frac{n-1}{n}, 1\right]\right) \text{ and } p_n^n = 0. \end{aligned} \tag{21}$$

We note that under this approximation, as n increases, p^n approaches P under the Kolmogorov norm, i.e. $\rho_K(P, p^n) \to 0$ as $n \to \infty$.

Let $P(0)$ be an initial population state in Δ and let $P(t)$ be the solution trajectory of the logit dynamic from $P(0)$ in the infinite dimensional Burdett and Judd model (1). We also consider an additional trajectory that starts from $p^n(0)$, where $p^n(0)$ is the approximation of $P(0)$ as defined in (21). This trajectory is determined by the finite dimensional logit dynamic

$$\dot{p}_i^n = \mathcal{L}_{\eta,i}(p^n) - p_i^n, \tag{22}$$

where

$$\mathcal{L}_{\eta,i}(p^n) = \frac{\exp(\eta^{-1})\pi_i^n\left(p^n\right)}{\sum_{S^n} \exp(\eta^{-1})\pi_j^n\left(p^n\right)}$$

is the $i-$th component of the finite dimensional logit best response function $\mathcal{L}_\eta(p^n) \in \Delta^n$ defined with respect to the finite dimensional Burdett and Judd payoff function (20).

Strictly speaking, the finite distribution p^n is a component of Δ^n. However, we may also regard it as probability measure on Δ, one that puts positive mass only on the subset S^n of S. Similarly, we may also regard $\mathcal{L}_\eta(p^n)$ as a probability measure on Δ such that $\mathcal{L}_\eta(p^n)(A) = \sum_{i \in A} \mathcal{L}_{\eta,i}(p^n)$. This measure differs from the logit choice measure $L_\eta(p^n)$.[14] However, if p^n is an approximation of P, then, given our definitions of the finite and infinite dimensional logit dynamics, it is clear that as $n \to \infty$, $\mathcal{L}_\eta(p^n)(A) \to L_\eta(P)(A)$. Intuitively, this suggests that if n is sufficiently large and if $p^n(0)$ is an approximation of $P(0)$, then the two trajectories $\{p^n(t)\}_{t \geq 0}$ and $\{P(t)\}_{t \geq 0}$ generated by the finite and infinite dimensional logit dynamics are also close to each other. We formalize this intuition in the following proposition. The proof is in Appendix A.4.

Proposition 8.1 *Let $\{P(t)\}_{t \geq 0}$ be the solution trajectory of the logit dynamic from initial state $P(0)$ in the infinite dimensional Burdett and Judd (1983) model defined by payoff function (1). Let $p^n(0)$ be the finite dimensional approximation of $P(0)$, where the approximation is as defined in*

[14] To see this, note that $L_\eta(p^n)(A) > 0$ for all $A \in \mathcal{B}$ whereas $\mathcal{L}_{\eta,i}(p^n) = 0$ for all $i \notin S^n$.

(21). Let $\{p^n(t)\}_{t\geq 0}$ be the solution trajectory of the finite dimensional logit dynamic (22) defined with respect to the finite dimensional Burdett and Judd (1983) model defined by payoff function (20). Then for any $T > 0$ and any $t \in [0, T]$,

$$\rho_K(P(t), p^n(t)) \to 0$$

as $n \to \infty$.

Proposition 8.1 leads to the following conjecture. For η sufficiently small, let $\tilde{P}_\eta$ be the logit equilibrium corresponding to the unique Nash equilibrium P^* of the continuous strategy Burdett and Judd (1983) model as defined in (1). Consider an initial state $P(0) \in \Delta$ close to P^* and approximate it with $p^n(0) \in \Delta$ under the Komogorov norm. Typically, for n large and η small, $p^n(0)$ will not be a logit equilibrium of (20). Therefore, as the results and simulations in Lahkar (2011) imply, $p^n(t)$ converges to a limit cycle in $\Delta^n \subset \Delta$. This is true for all large n. But Proposition 8.1 shows that the logit trajectory from $P(0)$, $P(t)$, is arbitrarily close to $p^n(t)$ under the Kolmogorov norm, for n sufficiently large. Hence, for low values of η, $P(t)$ should also converge to a limit cycle in Δ.

We should note, however, this is a conjecture. Proposition 8.1 does not rigorously prove that the dispersed price equilibrium in the infinite dimensional Burdett and Judd (1983) model is unstable. It is possible that as n becomes large, the extent to which a trajectory deviates from the equilibrium becomes progressively less so that in the limiting continuous case, there is no deviation at all and the equilibrium is stable. This is, however, unlikely, because simulations in Lahkar (2011) suggest that once n becomes sufficiently large, the probability measure $P(t)$ on Δ induced by the limit cycle on Δ^n remains largely unchanged.[15] Hence, it is reasonable to conclude that even in the original Burdett and Judd (1983) model, price dispersion manifests itself as a cyclical phenomenon.

9 Conclusion

In this paper, we have introduced the logit dynamic for such population games with continuous strategy sets. We have established the fundamental properties of the logit dynamic of existence, uniqueness and continuity of solutions to the dynamic with respect to initial conditions. We have shown that the logit equilibrium in a negative definite game with a continuous strategy set is Lyapunov stable under the logit dynamic.

We have also applied the logit dynamic to the Burdett and Judd (1983) model with exogenous consumer behaviour in its original setting of sellers choosing prices from a continuous set. We have established that the dynamic is well defined in this model. We have shown that this model is a positive definite game which gives rise to the conjecture that the dispersed price equilibrium in the

[15]This is in the following sense. Lahkar (2011) suggests that for n large, all solution trajectories converge to a unique limit cycle in Δ^n. Let $\{p^n(t)\}_{t\geq 0}$ be this limit cycle. Let $P(t)$ be the probability measure induced by this limit cycle in Δ at time t. Hence, for $A \in S$, $P(t)(A) = \sum_{i\in A} p_i^n(t)$. Once n exceeds a certain value, $P(t)$ remains largely unchanged.

model is unstable. We have not been able to rigourously establish this instability result. But we show that the solution trajectories of the logit dynamic in the continuous strategy Burdett and Judd (1983) model can be well approximated with the trajectories in the finite analogue of the model. An appeal to the results in Lahkar (2011) then lends support to our conjecture that the dispersed price equilibria in the continuous strategy Burdett and Judd (1983) model is unstable. Instead, solution trajectories converge to a limit cycle.

The paper immediately suggests an important research question for the future. This is the development of a theory of the spectral analysis of ODE systems defined in a measure space. Once such a theory is developed, it should be possible to apply it to economic problems like the Burdett and Judd (1983) model for a more rigorous assessment of local stability of equilibria in such models.

A Appendix

A.1 Proofs of Section 4

Proof of Theorem 4.2. By compactness, accumulation points exist. Let P^* be such an accumulation point. Take a sequence of positive numbers $\eta_n \downarrow 0$ and let P_n be a logit equilibrium for η_n such that $P_n \to P^*$ in the weak topology. Let $A \in \mathcal{B}$ be the support of P^*.

Suppose P^* is not a Nash equilibrium. Then, there exist $y \in S$ and $x \in A$ such that $\pi(y, P^*) > \pi(x, P^*) =$ for all $x \in A$. By continuity of π in the first variable, there exist open intervals B_x, B_y around x and y and an $\varepsilon > 0$ such that

$$\inf_{z \in B_y} \pi(z, P^*) \geq 2\varepsilon + \sup_{z \in B_y} \pi(z, P^*) .$$

By continuity of π in the second variable, we then have for sufficiently large $n \geq N \in \mathbb{N}$

$$\inf_{z \in B_y} \pi(z, P_n) \geq \varepsilon + \sup_{z \in B_y} \pi(z, P_n) .$$

As P_n are logit equilibria ($P_n = L(P_n)$), it follows that

$$\frac{P_n(B_y)}{P_n(B_x)} = \frac{\int_{B_y} \exp\left(\pi\left(z, P_n\right)/\eta_n\right) dz}{\int_{B_x} \exp\left(\pi\left(z, P_n\right)/\eta_n\right) dz} \geq \frac{\exp\left(\inf_{z \in B_y} \pi\left(z, P_n\right)/\eta_n\right) \lambda(B_y)}{\exp\left(\sup_{z \in B_x} \pi\left(z, P_n\right)/\eta_n\right) \lambda(B_y)} \geq \exp(\varepsilon/\eta_n) \frac{\lambda(B_y)}{\lambda(B_x)} .$$

As open intervals have positive Lebesgue mass and $\eta_n \downarrow 0$, the last term converges to infinity. It follows that $P_n(B_x) \to 0$.

As P_n converges weakly to P^*, the Portemanteau theorem implies that $P^*(B_x) = 0$. But this contradicts $x \in A$. (As x is in the support of P^*, and B_x is open, we must have $P^*(B_x) > 0$.) ■

Proof of Proposition 4.3. For strategy $x \in S = [0, 1]$, define $\phi(x, P) = x(q_1 + 2q_2(1 - F_P(x)))$. If $P([x]) = 0$, $\phi(x, P) = \pi(x, P)$, where $\pi(x, P)$ is given by (1). The set of $x \in S$ with $P(\{x\}) > 0$

is countable, thus of Lebesgue measure zero. Therefore, for $A \in \mathcal{B}$,

$$L_\eta(P)(A) = K_\eta(P)(A) := \int_A \frac{\exp\left(\eta^{-1}\phi(x,P)\right)}{\int_0^1 \exp\left(\eta^{-1}\phi_y(P)\right) dy} dx.$$

Therefore, to establish the continuity of $L_\eta(P)$, it is enough to show that $K_\eta(P)$ is continuous with respect to the weak topology.

Let (P_n) be a sequence of probability measures that converges weakly to a probability measure P. We have to show that the induced measures $K_\eta(P_n)$ converge weakly to the measure $K_\eta(P)$. For every point $x \in S$ of continuity of P (or its distribution function F_P, i.e. with $P(\{x\}) = 0$), we have

$$F_{P_n}(x) \to F_P(x)\,.$$

Again, as the set of discontinuities of F_P is countable, it follows that

$$F_{P_n}(x) \to F_P(x) \quad \text{for almost every } x \in S\,.$$

Now fix $z \in S$. The distribution function of $K_\eta(P_n)$ is the ratio of

$$\int_0^z \exp(\eta^{-1}\phi(x,P_n))dx$$

and

$$\int_0^1 \exp(\eta^{-1}\phi(x,P_n))dx\,.$$

By Lebesgue's dominated convergence theorem, the definition of ϕ, and our above argument, both terms converge for $n \to \infty$ towards

$$\int_0^z \exp(\eta^{-1}\phi(x,P))dx$$

and

$$\int_0^1 \exp(\eta^{-1}\phi(x,P))dx\,.$$

As the last term is strictly positive, the ratio also converges. This establishes the continuity of $K_\eta(P)$, and, therefore, $K_\eta(P)$ with respect to the weak topology. The existence of a logit equilibrium follows from Theorem 4.1.

For the second part, note that if $\{\tilde{P}_\eta\}_{\eta>0}$ is a sequence of logit equilibria, then this sequence has accumulation points (by Theorem 4.2). Moreover, each accumulation point is a Nash equilibrium. Since there is one Nash equilibrium in the Burdett and Judd model, we have convergence. ■

A.2 Proofs of Section 5

Proof of Theorem 5.2. Given $P \in \mathcal{M}^e(S, \mathcal{B})$, we define the auxiliary function

$$\tilde{V}(P) = (2 - \|P\|)_+ V(P),$$

where $\|\cdot\|$ is the variational norm. We note that $\tilde{V}$ coincides with V on Δ. We show that the function $\tilde{V}(P)$ is Lipschitz continuous on $\mathcal{M}^e(S, \mathcal{B})$. Formally, we want to show that for all $P, Q \in \mathcal{M}^e(S, \mathcal{B})$, there exists $K > 0$ such that

$$\|\tilde{V}(P) - \tilde{V}(Q)\| \leq K\|P - Q\|. \tag{23}$$

To show (23), we need to distinguish between three cases. First, if both $\|P\|, \|Q\| \geq 2$, then $\tilde{V}(P) = 0$ and there is nothing to show. Next, we consider $\|P\| \geq 2 > \|Q\|$. Hence, $\tilde{V}(P) = 0$ and the left hand side of (23) is therefore

$$\|\tilde{V}(Q)\| = (2 - \|Q\|)\|V(Q)\|.$$

Given any $Q \in \mathcal{M}^e(S, \mathcal{B})$, the logit best response function to it is a probability measure; i.e. $\|L(Q)\| = 1$. Hence,

$$\|V(Q)\| = 1 + \|Q\| \leq 3.$$

We therefore have

$$\|\tilde{V}(P) - \tilde{V}(Q)\| = \|\tilde{V}(Q)\| \leq 3(2 - \|Q\|) \leq 3[\|P\| - \|Q\|] \leq 3\|P - Q\|.$$

Finally, we consider the case where both $\|P\|, \|Q\| \leq 2$. Continuing from (23), we can write

$$\begin{aligned}
\|\tilde{V}(P) - \tilde{V}(Q)\| &= \|(2 - \|P\|)V(P) - (2 - \|Q\|)V(Q)\| \\
&\leq (2 - \|P\|)\|V(P) - V(Q)\| + \|V(Q)\| \left(|\|P\| - \|Q\||\right) \\
&\leq (2 - \|P\|)\|V(P) - V(Q)\| + 3\|P - Q\| \\
&\leq 2\|V(P) - V(Q)\| + 3\|P - Q\|.
\end{aligned}$$

The proof is complete if we show that $V(P)$ is Lipschitz continuous for both $\|P\|, \|Q\| \leq 2$. To show this, it is sufficient to show that the logit best response function $L(P)$ is Lipschitz continuous on $\mathcal{M}_2$. Since $L(P)$ is a probability measure, $\|L(P) - L(Q)\| = 2\sup_{A \in \mathcal{B}} |P(A) - Q(A)|$. It is, therefore, sufficient to show that for all $x \in S$,

$$\left| \frac{\exp\left(\pi(x, P)\right) dx}{\int_S \exp\left(\pi(x, P)\right) dx} - \frac{\exp \pi(x, Q) dx}{\int_S \exp \pi(x, Q) dx} \right| \leq K_1 \|P - Q\|,$$

for some K_1 independent of P, Q. Here, without loss of generality, we have taken $\eta = 1$.

By the Lipschitz continuity of π_x and the fact that S is bounded, $e^{\pi(x,P)}$ is Lipschitz continuous.

Since payoffs are bounded on the bounded set S, the denominator $\int_S e^{\pi(x,P)}dx$ Lipschitz continuous and bounded away from zero. The ratio of two Lipschitz continuous functions is Lipschitz continuous as long as the denominator stays away from zero. Therefore, the Lipschitz continuity of π_x implies the Lipschitz continuity of $L(P)$.

The Lipschitz continuity of $\tilde{V}$, therefore, implies that the differential equation $\dot{P} = \tilde{V}(P)$ has a unique solution in $\mathcal{M}^e(S, \mathcal{B})$. Since $\tilde{V}$ coincides with V on Δ, there also exists a unique solution $P(t)$ from every initial point $P(0) \in \Delta$. Finally, the facts that $\dot{P}(S) = 0$, $\dot{P}(A) > 0$ if $P(A) = 0$ and $\dot{P}(A) < 0$ if $P(A) = 1$ imply that if $P(0) \in \Delta$, $P(t) \in \Delta$ for all $t > 0$.

Since the logit dynamic is Lipschitz continuous with respect to the variational norm, the semiflow ξ is continuous with respect to initial conditions under the same norm. This follows from the application of Gronwall's lemma (see Zeidler (1986), Propositions 3.10 and 3.11). By Gronwall's lemma, if $P(t)$ and $Q(t)$ are solutions to the logit dynamic from initial points $P(0)$ and $Q(0)$ respectively, then then there exists a constant $K_2 > 0$ such that

$$\|P(t) - Q(t)\| \leq e^{K_2 t}\|P(0) - Q(0)\|.$$

This establishes continuity of solution trajectories $P(t)$ and $Q(t)$ with respect to initial conditions $P(0)$ and $Q(0)$. ■

Proof of Corollary 5.3. Note that (1) is measurable and bounded on $\mathcal{M}_2$. We show the Lipschitz continuity of (1) under the variational norm.

Consider the strategy set $S = [0, 1]$ and let $\mathbf{1}_{(x,1]}$ denote the identity function on $(x, 1]$. We note that $\mathbf{1}_{(x,1]}$ is a measurable function with $\sup_{y \in S} |\mathbf{1}_{(x,1]}(y)| = 1$. Then,

$$\begin{aligned} |\pi(x, P) - \pi(x, Q)| &= 2q_2 x \left| \left(\frac{P([x])}{2} + P(x, 1] \right) - \left(\frac{Q(x)}{2} + Q(x, 1] \right) \right| \\ &\leq 2 \left(|P(x, 1] - Q(x, 1]| + \left| \frac{P([x])}{2} - \frac{Q(x)}{2} \right| \right). \end{aligned}$$

Let $\mathbf{1}_{(x,1]}$ denote the identity function on $(x, 1]$. Then

$$\begin{aligned} |P(x, 1] - Q(x, 1]| &= \left| \int_S \mathbf{1}_{(x,1]} d(P - Q) \right| \\ &\leq \left| \sup \int g d(P - Q) \right|, \text{ such that } \|g\| \leq 1 \\ &\leq \sup \left| \int g d(P - Q) \right| \\ &= \|P - Q\|. \end{aligned}$$

Similarly, $\left|\frac{P([x])}{2} - \frac{Q(x)}{2}\right| \leq \|P - Q\|$. Hence,

$$|\pi(x,P) - \pi(x,Q)| \leq 4\|P - Q\|. \blacksquare$$

A.3 Proofs of Section 7

Proof of Proposition 7.2. Given the Burdett-Judd payoff function $\pi(x,P)$ in (1), we need to show that for all $P, Q \in \Delta$,

$$\int_S (\pi(x,Q) - \pi(x,P))(Q - P)(dx) > 0.$$

Denote $Q - P = \mu_Z$. Hence, the difference of the distribution functions is $F_Q(x) - F_P(x) = Z(x)$. Therefore,

$$\pi(x,Q) - \pi(x,P) = 2xq_2 \left(\frac{\mu_Z(x)}{2} - Z(x)\right).$$

Hence,

$$\begin{aligned}
\int_S (\pi(x,Q) - \pi(x,P))(Q - P)(dx) &= \int_S (\pi(x,Q) - \pi(x,P))dZ(x) \\
&= 2q_2 \int\limits_S x \left(\frac{\mu_Z(x)}{2} - Z(x)\right) dZ(x) \\
&= -q_2 \int\limits_S x dZ^2(x),
\end{aligned}$$

where we apply integration by parts for the Stieltjes integral to obtain

$$Z(x)dZ(x) = \frac{1}{2}dZ^2(x) + \frac{\mu_Z(x)dZ(x)}{2}.$$

Then, by using integration by parts for the Stieltjes integral and the fact that $Z(0) = Z(1) = 0$, we have

$$\int_S (\pi(x,Q) - \pi(x,P))(Q - P)(dx) = q_2 \int\limits_S Z^2(x)dx > 0. \blacksquare$$

Proof of Lemma 7.3: The proof follows the method established by Hofbauer and Sandholm (2007, Theorem 3.1). Clearly, it suffices to show uniqueness on Δ^{AC}. Consider $P, Q \in \Delta^{AC}$ and let p, q be their density functions. Define the *virtual payoff*

$$\hat{\pi}(x,P) = \pi(x,P) - \eta(1 + \log p(x)),$$

Let $z(x) = p(x) - q(x)$ and denote by μ_Z the associated signed measure. Define

$$\varphi(P, \mu_Z, t) = \int_S \hat{\pi}(x, P + t\mu_Z)\mu_Z(dx),$$

where the scalar $t > 0$ is such that $p(x)+tz(x) > 0$ for all $x \in S$. This is possible since z is bounded and bounded away from zero. Using the properties of the Gateaux derivative, we compute

$$\begin{aligned}\frac{\partial \varphi(P,\mu_Z,t)}{\partial t} &= \int_S D\hat{\pi}(x,P+t\mu_Z)\mu_Z\mu_Z(dx) \\ &= \int_S D\pi(x,P+t\mu_Z)\mu_Z\mu_Z(dx) - \int_S \frac{z^2(x)}{p(x)+tz(x)}dx < 0, \end{aligned} \tag{24}$$

by the negative semi-definiteness of π. Hence, for all $P \in \Delta$ and $\mu_Z \in T\Delta$, $\varphi(P,\mu_Z,t)$ is declining in t.

Let $\tilde{P}_\eta$ be a logit equilibrium. Then, $\hat{\pi}(x,\tilde{P}_\eta) = \eta\left(\log\int_S \exp\left(\pi(y,\tilde{P}_\eta)\right)dy - 1\right)$, so that $\varphi(\tilde{P}_\eta,\mu_Z,0) = \int_S \hat{\pi}(x,\tilde{P}_\eta)\mu_Z(dx) = 0$, for all $\mu_Z \in TX$. Now, let $Q = \tilde{P}_\eta + t\mu_Z$ for some non-zero $\mu_Z \in T\Delta$ such that $Q \in \operatorname{int}\Delta$ is distinct from $\tilde{P}_\eta$. By definition, $\int_S \hat{\pi}(x,Q)\mu_Z(dx) = \varphi(\tilde{P}_\eta,\mu_Z,t)$, and if $\varphi(\tilde{P}_\eta,\mu_Z,0) = 0$, then from (24), we conclude that $\varphi(\tilde{P}_\eta,\mu_Z,t) < 0$. Hence, $\int_S \hat{\pi}(x,Q)\mu_Z(dx) < 0$, and so, Q cannot be a logit equilibrium. ■

Proof of Theorem 7.5: Let $P \in \Delta^{AC}$. Consider the candidate Lyapunov function $G(P)$ in (19) and note that it is equivalent to

$$G(P) = \left[\int_S \pi(x,P)l_\eta(P)(x)dx - v(l_\eta(P))\right] - \left[\int_S \pi(x,P)p(x)dx - v(P)\right].$$

Recall from Lemma 7.4 that $G(P) \geq 0$ for all $P \in \Delta^{AC}$, with equality only at a logit equilibrium $\tilde{P}_\eta$. Moreover,

$$\begin{aligned} v\left(L_\eta(P)\right) &= \eta\int_S \log(L_\eta(P)(x))l_\eta(P)(x)dx \\ &= \int_S \pi(x,P)l_\eta(P)(x)dx - \eta\log\left[\int_S \exp\left(\eta^{-1}\pi(x,P)\right)dx\right]. \end{aligned}$$

Therefore, $G(P)$ reduces to

$$G(P) = \eta\log\left[\int_S \exp\left(\eta^{-1}\pi(x,P)\right)dx\right] - \left[\int_S \pi(x,P)p(x)dx - v(P)\right]. \tag{25}$$

G is continuous in the variational norm on Δ^{AC} and lower semicontinuous with respect to the weak topology. The latter fact follows from the variational characterization of entropy as a supremum of continuous linear functionals.

We apply the definition of the Gateaux derivative (17) to compute $\dot{G}(P)$. Therefore,

$$\begin{aligned} \dot{G}(P) &= DG(P)\dot{P} \\ &= \eta D\log\left[\int_S \exp\left(\eta^{-1}\pi(x,P)\right)dx\right]\dot{P} - \left[D\left(\int_S \pi(x,P)p(x)dx\right)\dot{P} - Dv(P)\dot{P}\right]. \end{aligned} \tag{26}$$

We calculate each part of (26) separately. Applying the chain rule to the first part, we have

$$
\begin{aligned}
&\eta D \log \left[\int_S \exp\left(\eta^{-1}\pi(x,P)\right) dx \right] \dot{P} \\
&= \frac{\eta}{\int_S \exp\left(\eta^{-1}\pi(y,P)\right) dy} \left[\int_S \exp\left(\eta^{-1}\pi(x,P)\right) \eta^{-1} D\pi(x,P)\dot{P} dx \right] \\
&= \int_S D\pi(x,P)\dot{P} \frac{\exp\left(\eta^{-1}\pi(x,P)\right)}{\int_S \exp\left(\eta^{-1}\pi(y,P)\right) dy} dx \\
&= \int_S D\pi(x,P)\dot{P} l_\eta(P)(x) dx. \qquad (27)
\end{aligned}
$$

Through the product rule, the second part of (26) is

$$
\begin{aligned}
D\left(\int_S \pi(x,P)p(x)dx \right) \dot{P} &= \int_S D\pi(x,P)\dot{P}p(x)dx + \int_S Dp(x)\dot{P}\pi(x,P)dx \\
&= \int_S D\pi(x,P)\dot{P}p(x)dx + \int_S \pi(x,P)\dot{p}(x)dx, \qquad (28)
\end{aligned}
$$

since $Dp(x)\dot{P} = \dot{p}(x)$. Again applying the product rule, the third part of (26) is

$$
\begin{aligned}
Dv(P)\dot{P} &= \eta D \left(\int_S \log p(x) p(x) dx \right) \dot{P} \\
&= \eta \int_S D \log p(x) \dot{P} p(x) dx + \int_S \log p(x) Dp(x) \dot{P} p(x) dx \\
&= \eta \int_S \frac{1}{p(x)} Dp(x) \dot{P} p(x) dx + \int_S \log p(x) \dot{p}(x) dx \\
&= \eta \int_S \dot{p}(x) dx + \int_S \log p(x) \dot{p}(x) dx \\
&= \eta \int_S \log p(x) \dot{p}(x) dx, \qquad (29)
\end{aligned}
$$

since $Dp(x)\dot{P} = \dot{p}(x)$ and $\int_S \dot{p}(x)dx = 0$. Using (27), (28) and (29) in (26), we obtain

$$
\begin{aligned}
\dot{G}(P) &= \int_S D\pi(x,P)\dot{P} l_\eta(P)(x) dx - \int_S D\pi(x,P)\dot{P}p(x)dx - \int_S \left(\pi(x,P) - \eta \log p(x)\right) \dot{p}(x) dx \\
&= \int_S D\pi(x,P)\dot{P} \left(l_\eta(P)(x) - p(x) \right) dx - \int_S \left(\pi(x,P) - \eta \log p(x)\right) \dot{p}(x) dx \\
&= \int_S D\pi(x,P)\dot{P}\dot{P}(dx) - \int_S \left(\pi(x,P) - \eta \log p(x)\right) \dot{p}(x) dx. \qquad (30)
\end{aligned}
$$

Using the definition of $l_\eta(P)$, we can write

$$
\pi(x,P) = \eta \log l_\eta(P)(x) + \eta \left(\log \left(\int_S \exp(\eta^{-1}\pi(y,P)) dy \right) \right). \qquad (31)
$$

Therefore,

$$
\begin{aligned}
&\int_S (\pi(x,P) - \eta \log p(x))\, \dot{p}(x) dx \\
=&\int_S \left(\eta \log l_\eta(P)(x) + \eta \left(\log \left(\int_S \exp(\eta^{-1}\pi(y,P)) dy \right)\right) - \eta \log p(x) \right) \dot{p}(x) dx \\
=&\eta \int_S (\log l_\eta(P)(x) - \log p(x))\, \dot{p}(x) dx. \qquad (32)
\end{aligned}
$$

Using (32) in (30), we obtain

$$
\begin{aligned}
\dot{G}(P) &= \int_S D\pi(x,P)\dot{P}\dot{P}(dx) - \eta \int_S (\log l_\eta(P)(x) - \log p(x))\, \dot{p}(x) dx \\
&= \int_S D\pi(x,P)\dot{P}\dot{P}(dx) - \eta \int_S (\log l_\eta(P)(x) - \log p(x))\, (l_\eta(P)(x) - p(x))\, dx. \qquad (33)
\end{aligned}
$$

As π is negative-semidefinite, the first term is nonnegative. The second term is zero for a logit equilibrium and positive else since the log function is an increasing function. Hence, for any negative semi-definite game, $\dot{G}(P) \leq 0$, with equality only a rest point of the logit dynamic, or a logit equilibrium.

Now take $\varepsilon > 0$. The function G assumes its minimum on the weakly compact sphere

$$
S = \{P \in \Delta : \rho(P, P^\eta) = \varepsilon\}
$$

because it is lower semicontinuous with respect to the weak topology. Thus,

$$
m = \min_{P \in S} G(P) > 0\,.
$$

As G is continuous in the strong topology and $G(P^\eta) = 0$, there exists $\delta > 0$ such that for all $Q \in \Delta$ with $\|Q - P^\eta\| < \delta$ we have $G(Q) < m$. When we thus start in such a population state Q, we have, by the above, $G(P(t)) < m$ for the whole trajectory $(P(t))$. As the trajectory is continuous in time in the strong, and then in the weak topology, it can never leave the weak ε-ball around P^η. ■

A.4 Proofs of Section 8

In order to prove Proposition 8.1, we need to establish the following lemma.

Lemma A.1 *Let $P, Q \in \Delta$. Then, there exists a constant $K > 0$ such that $\rho_K(L_\eta(P), L_\eta(Q)) \leq K\rho_K(P,Q)$.*

Proof. As in the proof of Proposition 4.3, define $\phi(x,P) = x(q_1 + 2q_2(1 - F_P(x)))$, for all $x \in S$. Note that $\phi(x,P) = \pi(x,P)$ *a.e.* We show that $\phi(x,P)$ is Lipschitz continuous with respect to Kolmogorov norm. We need to show that for every $x \in S$, $P, Q \in \Delta$, $|\phi(x,P) - \phi(x,Q)| <$

$K_1\rho_K(P,Q)$, for some constant K_1. But

$$|\phi(x,P)-\phi(x,Q)| = 2xq_2|F_P(x)-F_Q(x)| < 2\rho_K(P,Q).$$

By an argument akin to the proof of Theorem 5.2, Lipschitz continuity of the payoff function ϕ implies the Lipschitz continuity of the logit choice measure defined for ϕ. But the logit choice measure defined for ϕ is the same as that defined for π. This establishes the result. ■

Proof of Proposition 8.1. Let $F_{P(t)}$ and $F_{p^n(t)}$ be the distribution functions of $P(t)$ and $p^n(t)$ respectively. Interpret $p^n(t)$ as a probability measure in Δ that puts zero mass on $S \setminus S^n$. For any $x \in S$, we then obtain

$$\begin{aligned}
&\left|F_{P(t)}(x) - F_{p^n(t)}(x)\right| \\
&= \left|\int_0^x dP(t) - \int_0^x dp^n(t)\right| \\
&= \left|F_{P(0)}(x) - F_{p^n(0)}(x) + \int_0^t \left(\int_0^x d\dot{P}(s) - \int_0^x d\dot{p}^n(s)\right) ds\right| \\
&= \left|\varepsilon_1(n) + \int_0^t \left(\int_0^x d\left(L_\eta(P(s)) - P(s)\right) - \int_0^x d\left(\mathcal{L}_\eta(p^n(s)) - p^n(s)\right)\right) ds\right| \\
&\le |\varepsilon_1(n)| + \int_0^t \left|\int_0^x dL_\eta(P(s)) - \int_0^x d\mathcal{L}_\eta(p^n(s))\right| ds + \int_0^t \left|\int_0^x dP(s) - \int_0^x dp^n(s)\right| ds \\
&= |\varepsilon_1(n)| + \int_0^t \left|\int_0^x d\left(L_\eta(P(s)) - L_\eta(p^n(s))\right) + \int_0^x d\left(L_\eta(p^n(s)) - \mathcal{L}_\eta(p^n(s))\right)\right| ds \\
&\quad + \int_0^t \left|\int_0^x dP(s) - \int_0^x dp^n(s)\right| ds \\
&\le |\varepsilon_1(n)| + \int_0^t \left|\int_0^x d\left(L_\eta(P(s)) - L_\eta(p^n(s))\right)\right| ds + \int_0^t \left|\int_0^x d\left(L_\eta(p^n(s)) - \mathcal{L}_\eta(p^n(s))\right)\right| ds \\
&\quad + \int_0^t \left|\int_0^x dP(s) - \int_0^x dp^n(s)\right| ds. \qquad (34)
\end{aligned}$$

The first term on the right hand of (34), $|\varepsilon_1(n)| \to 0$ as $n \to \infty$. The second term is of the order of $\rho_K(P(s), p^n(s))$ (by Lemma A.1). The third term is $\rho_K(P(s), p^n(s))$. Let us denote the second term as $\varepsilon_2(n)$. As $n \to \infty$, $\mathcal{L}_\eta(p^n(s))(A) \to L_\eta(p^n(s))$. Hence, $\varepsilon_2(n) \to 0$ as $n \to \infty$.

From (34), we, therefore, obtain

$$\left|F_{P(t)}(x) - F_{p^n(t)}(x)\right| \le |\varepsilon_1(n)| + \int_0^t |\varepsilon_2(n)| + K\int_0^t \left(\rho_K(P(s), p^n(s))\right) ds,$$

for some constant $K > 0$. Let $|\varepsilon_1(n)| + \int_0^T |\varepsilon_2(n)| = \varepsilon(n) \to 0$ as $n \to \infty$. Clearly, $|\varepsilon_1(n)| + \int_0^t |\varepsilon_2(n)| \le \varepsilon(n)$ for all $t \in [0,T]$. Hence, for every x,

$$\left|F_{P(t)}(x) - F_{p^n(t)}(x)\right| \le \varepsilon(n) + K\int_0^t \left(\rho_K(P(s), p^n(s))\right) ds.$$

The same holds true at the particular x at which $\left|F_{P(t)}(x) - F_{p^n(t)}(x)\right|$ is maximized so that

$$\rho_K\left(P(t), p^n(t)\right) \leq \varepsilon(n) + K\int_0^t \left(\rho_K(P(s), p^n(s))\right) ds.$$

The result then follows from from an application of Gronwall's lemma. ■

Acknowledgments

The authors thank David Leslie, Steven Perkins and Bill Sandholm for comments and suggestions. This work was partially done when the first author was employed in IFMR, Chennai. The author acknowledges the generous support and contribution of IFMR in completing this paper.

References

[1] Anderson, S., Goeree, J., Holt, C., 2004, Noisy Directional Learning and the Logit Equilibrium, Scand. J. Econ. 106, 581–602.

[2] Atkinson, K., Han, W., 2009, Theoretical Numerical Analysis: A Functional Analysis Approach, Texts in Applied Mathematics, Springer.

[3] Benaim, M., Hofbauer, J., Hopkins, E., 2009, Learning in Games with Unstable Equilibria, J. Econ. Theory 144, 1694–1709.

[4] Billingsley, P., Convergence of Probability Measures, Wiley Series in Probability and Statistics, 1999.

[5] Björnerstedt, J., Weibull, J., 1996, Nash equilibrium and evolution by imitation. In: K. Arrow et al. (Eds.), The Rational Foundations of Economic Behavior. St. Martins, New York, pp. 155-181.

[6] Bomze, I., 1990, Dynamical aspects of evolutionary stability, Monatsh. Math 110, 189–206.

[7] Bomze, I., 1991, Cross entropy minimization in uninvadable states of complex populations, J. Math. Biol. 30, 73–87.

[8] Burdett, K., Judd, K., 1983, Equilibrium Price Dispersion, Econometrica 51, 955–969.

[9] Cheung, M.W., 2014, Pairwise Comparison Dynamics for Games with Continuous Strategy Space, forthcoming, J. Econ. Theory, `http://dx.doi.org/10.1016/j.jet.2014.07.001`.

[10] Cressman, R., 2005, Stability of the replicator equation with continuous strategy space. Math. Soc. Sci. 50, 127–147.

[11] Cressman, R., Hofbauer, J., 2005, Measure dynamics on a one-dimensional continuous trait space: Theoretical foundations for adaptive dynamics, Theor. Popul. Biol. 67, 47–59.

[12] Cressman, R., Hofbauer, J., Riedel, F., 2006, Stability of the replicator equation for a single-species with a multi-dimensional continuous trait space. J Theor Biol 239, 273–288.

[13] Cason, T., Friedman, D., Wagener, F., 2005, The Dynamics of Price Dispersion or Edgeworth Variations, J. Econ. Dynam. Control 29, 801–822.

[14] Föllmer, H., Schied, A., 2010, Stochastic Finance, De Gruyter, Berlin.

[15] Fudenberg, D., Levine, D., 1998, Theory of Learning in Games. MIT Press, Cambridge, Massachusetts.

[16] Gibbs, A., Su, F., 2002, On Choosing and Bounding Probability Metrics, Available at arXiv:math.

[17] Hofbauer, J., Sandholm, W., 2002, On the Global Convergence of Stochastic Fictitious Play Econometrica 70, 2265–2294.

[18] Hofbauer, J., Sandholm, W., 2007, Evolution in Games with Randomly Disturbed Payoffs, J. Econ. Theory 132, 47–69.

[19] Hofbauer, J., Sandholm, W., 2009, Stable Games and their Dynamics, J. Econ. Theory 144, 1665–1693

[20] Hofbauer, J., Oechssler, J., Riedel, F., 2009, Brown–von Neumann–Nash dynamics: The continuous strategy case, Games Econ Behav. 65, 406–429.

[21] Hopkins, E., Seymour, R., 2002, The Stability of Price Dispersion under Seller and Consumer Learning, Int. Econ. Rev. 43, 1157–1190.

[22] Hu, D., 2011, Essays on Dynamic Congestion Games, PhD Dissertation, University of Wisconsin-Madison, 2011.

[23] Lahkar, R., 2007, Essays in Evolutionary Game Theory, PhD Dissertation, University of Wisconsin-Madison, 2007.

[24] Lahkar, R., 2011, The Dynamic Instability of Dispersed Price Equilibria, J. Econ. Theory 146, 1796–1827.

[25] Lach, S., 2002Existence and Persistence of Price Dispersion: An Empirical Analysis, Rev. Econ. Statist. 84, 433–444.

[26] Mattsson, L., Weibull, J., 2002, Probabilistic choice and procedurally bounded rationality, Games Econ Behav. 41, 61–78.

[27] Maynard Smith, J., 1982, Evolution and the Theory of Games, Cambridge University Press, Cambridge, UK.

[28] Oechssler, J., Riedel, F., 2001, Evolutionary dynamics on infinite strategy spaces, Econ. Theory 17, 141-162.

[29] Oechssler, J., Riedel, F., 2002, On the Dynamics Foundations of Evolutionary Stability in Continuous Models, J. Econ. Theory 146, 223–252.

[30] Perkins, S., Leslie, D., 2014, Stochastic fictitious play with continuous action sets, J. Econ. Theory 152, 179–213.

[31] Shiryaev, A., 1995, Probability, 2nd ed., Springer Verlag, New York.

[32] Taylor, P.D., Jonker, L., 1978. Evolutionarily stable strategies and game dynamics. Math. Biosci. 40, 145-156.

[33] Zeidler, E., 1986, Nonlinear Functional Analysis and its Applications, Vol.I, Springer Verlag, New York.