APPROXIMATION OF LARGE DYNAMIC GAMES

AARON L. BODOH-CREED - CORNELL UNIVERSITY

Abstract. We provide a framework for approximating the equilibrium set of dynamic games with many players. We show that that the equilibria of a nonatomic dynamic limit game approximation are $\varepsilon-$Bayesian Nash Equilibria of the dynamic game with many players if the limit game is continuous. We also show that the Bayesian-Nash equilibrium correspondence of a large dynamic game is upper hemicontinuous in the number of agents and converges to the set of dynamic competitive equilibria of a nonatomic limit game under stronger continuity conditions. Our techniques provide a framework for simplifying the analysis and estimation of structural models with many agents by using nonatomic approximations. We also use our results to show that repeated static Nash equilibria are the only equilibria of continuous repeated games of perfect or imperfect information, public or private monitoring in the limit as $N \to \infty$.

Date: March 29th, 2011.
Key words and phrases. Approximate Equilibrium, Dynamic Games, Repeated Games, Rational Expectations Equilibrium. JEL Codes: C73, D4, L1.

I would like to thank Ilya Segal, Paul Milgrom, Doug Bernheim, Jon Levin, C. Lanier Benkard, Eric Budish, Matthew Jackson, Ramesh Johari, Philip Reny, Robert Wilson, Marcello Miccoli, Juuso Toikka, Alex Hirsch, Carlos Lever, Ned Augenblick, Isabelle Sin, John Lazarev, and Bumin Yenmez for invaluable discussions regarding this topic. In additional, valuable comments were received during seminars at Stanford, the Stanford Graduate School of Business, Yale School of Management, Kellogg MEDS, Boston University, the University of Rochester, Cornell University, the Ohio State University, and the 2010 IEEE Conference on Decision and Control. Bodoh-Creed gratefully aknowledges the support of the Jacob K. Javits Fellowship program and the Bradley Research Fellowship.
E-mail: acreed@cornell.edu.

1. Introduction

Games with a large but finite number players are important in a variety of fields of economics, although particularly in the areas of industrial organization, mechanism and market design, and macroeconomics. Because equilibrium conditions take into account the incentive effects of off-the-equilibrium-path events, the state space of the model can grow exponentially with the number of agents yielding a *curse of dimensionality*. In applications with more than a handful of economic agents, solving these games (computationally or analytically) or using these games as a foundation for structural analysis is not feasible. The difficulty of discovering and analyzing the exact equilibria for large games motivates research into characterizations of approximate equilibria and the relationship between the sets of approximate and exact equilibria (Benkard et al. [7], Adlakha et al. [1]).

Intuition suggests that when there are many agents and none of the agents are economically large relative to one another, the actions of individual agents in continuous models ought to have little effect on either outcomes in the present period or the future dynamics of the economy. For example, the price taking behavior assumed of a single consumer in a dynamic general equilibrium model captures the intuition that a single person's decision has negligible effect on current or future prices, and this behavioral assumption is precisely correct for an agent in a nonatomic model. Dynamic competitive economy models feature a continuum of nonatomic, utility maximizing agents who behave as if market aggregates, such as price, are exogenous to their own actions.[1] In equilibrium the evolution of market aggregates, as well as the agents' beliefs about the evolution, are consistent with the distribution of actions taken by all agents. By reducing the dimension of the policy space, these assumptions provide crucial tractability in applications ranging from the study of market frictions (e.g. Wolinsky [34]) to macroeconomic dynamic stochastic general equilibrium models. On the other hand, it is usually uncertain how well the continuum of agents in a dynamic competitive economy model captures the behavior of a large finite number of strategic agents.

[1]Throughout this work we use the term *strategic* to mean that the agents choose actions assuming that each agent's individual action can affect market aggregates. Price taking behavior in general equilibrium models is an example of *nonstrategic* behavior.

The primary goal of our study is to provide sufficient conditions under which a tractable nonatomic limit game approximates an underlying large finite game. Our principal theoretical contribution is captured in two approximation theorems. The first theorem provides continuity conditions under which a dynamic competitive equilibria is an $\varepsilon-$Bayesian-Nash equilibria of an N-agent game for sufficiently large N. However, this result leaves open the possibility that there are exact Bayesian-Nash equilibrium strategies of the N agent game that are not well approximated by any equilibrium of the nonatomic limit game.

Our second approximation theorem addresses this question by providing stronger continuity conditions on economic primitives of the nonatomic limit game that are sufficient for the set of Bayesian-Nash equilibria to be upper hemicontinuous in the limit as the number of agents approaches infinity. Therefore, any convergence sequence of equilibrium strategies of the N-agent game approaches a dynamic competitive equilibrium strategy of the nonatomic limit game as $N \to \infty$. This result places bounds on the strategies employed in *any* equilibrium of the large-finite game.[2] We interpret our first approximation result as a proof of approximation in terms of marginal incentives, while our second theorem provides the stronger result showing approximation in terms of observable actions and economic outcomes.

In our second approximation result we provide weaker continuity conditions under which a dynamic competitive equilibria is an $\varepsilon-$Bayesian-Nash equilibria of an N-agent game for sufficiently large N. While this theorem holds under weaker continuity conditions, our second result admits the possibility that there exists an equilibrium of the large finite game that is significantly different than any dynamic competitive equilibrium.

Approximation results in the literature usually focus on complete information games. Our techniques extend the range of models for which approximate solutions are valid to include:

- Dynamic games of incomplete information
- Games with persistent private information

[2]We do not show lower hemicontinuity of the strategy correspondence. Therefore it is possible that some of the equilibria of the nonatomic limit game are distant (in strategy space) from any exact equilibria of the N agent game even for arbitrarily large N.

- Imperfect public or private monitoring structures
- Stochastic games with aggregate shocks
- Dynamic games with large players and a fringe of small agents

Our approximation theorems also intersect the literature on anti-folk theorems (e.g. Green [20], Sabourian [30]). Anti-folk theorems for large games, such as the one we provide for large game of imperfect public and/or private monitoring, prove that the set of equilibrium payoffs in the large game converges to the repeated static Nash payoffs as the number of players approaches infinity. Our results are distinguished by their applicability to a broad class of games and highlight the important role of discontinuities in fostering agent pivotality in large economies. Furthermore, our sufficient conditions provide a microfoundation for nonatomic dynamic general equilibrium models employed by macroeconomists.

Finally, we provide a number of extensions of our results including: games with large players in the limit as $N \to \infty$; Markov perfect equilibria of complete information stochastic games; asymmetric games with multiple player roles (e.g. buyers and sellers); games with asynchronous actions; coalition proof equilibria; and games with (discontinuous) entry and exit decisions.

1.1. **Outline of Paper.** Section 2 outlines the model framework we analyze, describes the stochastic evolution of the games, and defines the equilibrium concepts we employ. Section 3 states the approximation theorems, and section 4 provides applications of our approximation theorems. Sections 4.1 and 4.2 use our results to provide anti-folk theorems for large repeated games of imperfect public or private monitoring. Section 4.3 extends our approximation theorems to games with a few large and many small players, Section 4.4 extends our result to perfect equilibria, and Section 4.5 provides approximations to large games with finite coalitions. Section 4.6 discusses games with entry, exit, and convex investment costs that admit discontinuous equilibrium strategies. Section 4.7 extends our methods to games with asymmetric roles and asynchronous actions. Section 5 relates our study to the existing literature, and section 6 concludes. All proofs are relegated to the appendix.

2. Model Framework

The model framework described in section 2.1 is required to be sufficiently abstract to capture the breadth of games that we approximate. Readers who desire examples are urged to read section 2.2, which provides applications of our framework to a variety of standard classes of games such as games of complete information, games of imperfect monitoring, and stochastic games. Sections 2.3 describes the evolution and equilibrium notions we use in the large finite models we approximate. Section 2.4 defines the evolution and equilibria of the nonatomic games we use as limit approximations of large finite games.

2.1. **Model.** Agent actions are drawn from $\mathcal{A} \subset \mathbb{R}^d$ where $d < \infty$, and we endow the space of lotteries over $\mathcal{A}$, denoted $\Delta(\mathcal{A})$, with the weak-* topology. The agent type space is $\Theta \subset \mathbb{R}^d$ with the Euclidean metric d_Θ, and these types are privately known to the agents. Probability measures over the space Θ are denoted $\Delta(\Theta)$ and endowed with the weak-* topology. When required, we refer to the Lévy-Prokhorov metric over a space of probability measures $\Delta(X)$ by d^X_{LP}.[3] Let those empirical probability measures that can be generated by N realizations from $\mathcal{A}$ and Θ be denoted $\Delta_N(\mathcal{A})$ and $\Delta_N(\Theta)$ respectively. We denote generic elements from these spaces by $\pi^\Theta \in \Delta(\Theta)$ and $\pi^\mathcal{A} \in \Delta(\mathcal{A})$, and let π^Θ_t and $\pi^\mathcal{A}_t$ be the realized distribution of types and actions in period t. All of the measures we employ are with respect to the sigma field generated by the Borel sets over the relevant space, and we use the notation $\mathcal{B}(X)$ to refer to the Borel sets over X. We use the notation $\pi^X[U]$ to refer to the measure of a set $U \in \mathcal{B}(X)$. If $\pi^X : Y \to \Delta(X)$, then we use the notation $\pi^X(y)[U]$ for $y \in Y$.[4]

Let the metric space (Ω, d_Ω) denote the aggregate states of the model. Throughout we assume that Ω is a completely regular Hausdorff space and denote a generic

[3]The Lévy-Prokhorov norm metricizes the weak-* topology over the space.

[4]For example, $\pi^\Theta_{t+1}(\omega_{t+1})$ refers to a measure describing the distribution of types in period $t+1$ conditional on a period $t+1$ aggregate shock, ω_{t+1}. $\pi^\Theta_{t+1}(\omega_{t+1})[U]$, $U \in \mathcal{B}(\Theta)$, refers to the measure of the set U with respect to measure $\pi^\Theta_{t+1}(\omega_{t+1})$.

element realized at time t as $\omega_t \in \Omega$.[5] The aggregate state could include information about aggregate payoff shocks, data available to all of the agents, a variable (unobservable to the agents) that mediates interagent and intertemporal correlation of the agents' private information, or the history of market aggregates in prior periods, $(\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A})$.[6] The initial aggregate state ω_0 is distributed according to some measure ν over Ω.

We assume that the agent types and the state of the economy evolve in the probability space $(\Psi, \mathcal{B}(\Psi), P)$ where $\Psi = \prod_{t=0}^\infty \Omega \times \Theta^\mathbb{N}$ reflects both the aggregate (Ω) and idiosyncratic ($\Theta^\mathbb{N}$) uncertainty in the economy.[7] Types for the agents are drawn at $t = 0$ according to the tight probability measure distribution $\pi_0^\Theta(\omega_0)$.[8] By letting π_0^Θ be conditioned on the $t = 0$ aggregate state, the agent types can incorporate private information about the initial aggregate state.

Throughout this work we will assume that the primitives of the model are anonymous. Anonymity is a symmetry assumption that requires agent utility, strategies, and the evolution of an agent's type depend only on the actions and types of other agents only through their empirical distribution. This structure allows us to treat these distributions as random variables and leverage asymptotic convergence results when analyzing the relationship between large finite games and the nonatomic limit games. In addition, anonymity allows us to elegantly embed the N agent games in the nonatomic limit game as $N \to \infty$.

We assume that the evolution of the agents' types are conditionally independent Markov processes with transition probability function from θ_t to θ_{t+1} denoted

$$T(\circ|\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}) : \mathcal{B}(\Theta) \to [0, 1]$$

which is conditional on the next period's aggregate state (ω_{t+1}), the distribution of actions and types in period t ($\pi_t^\mathcal{A}$ and π_t^Θ), and the agent's current type and action (θ_t

[5]The author is not aware of a space of interest to economists that is not a completely regular Hausdorff spaces. For example, the space $\mathbb{R}^\mathbb{R}$ is a completely regular Hausdorff space under either the standard product or uniform topologies.
[6]$(\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A})$ denotes an aggregate shock (ω_t) and the distribution of types and actions (π_t^Θ and $\pi_t^\mathcal{A}$ respectively) realized in period t.
[7]Throughout this work we use the notation $\prod_{t=M}^N X_i$ to refer to the product space $X_M \times X_{M+1} \times ... \times X_N$ where $N = \infty$ is allowed.
[8]For the definition of a tight measure, see Billingsley [11].

and a_t). By allowing conditioning on $(\omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})$, agent types in period $t+1$ can reflect private information regarding these quantities. Conditioning on (θ_t, a_t) allows the agent's type to reflect the influence of past and present actions. Throughout this work, we omit arguments from the type evolution operator that are not salient for the model of interest. For example, if the model does not have aggregate states, then we employ the notation

$$T(\circ|\theta_t, a_t, \pi_t^\Theta, \pi_t^\mathcal{A}) : \mathcal{B}(\Theta) \to [0,1]$$

The aggregate state evolves according to a Markov process with the associated transition probability function from ω_t to ω_{t+1} denoted

$$G(\circ|\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}) : \mathcal{B}(\Omega) \to [0,1]$$

Although we require restrictions on the space of agent types, we can employ a rich aggregate state space to record a public history. Even though the type evolution of the agents is conditionally independent, the aggregate state space allows us to have many forms of unconditional correlation of agent types both between agents and across periods mediated by ω_t.

If T and G are continuous in the weak-* topology, given continuous functions $r : \Theta \to \mathbb{R}$ and $s : \Omega \to \mathbb{R}$ we have that the expectations of these functions

$$E_t[r(\theta_{t+1})|\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}]$$
$$E_t[s(\omega_{t+1})|\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}]$$

are continuous in the conditioning variables. When we make the stronger assumption that T and G are uniformly continuous, then we have that for any uniformly continuous functions $r_{UC} : \Theta \to \mathbb{R}$ and $s_{UC} : \Omega \to \mathbb{R}$ we have that the expectations of these functions

$$E_t[r_{UC}(\theta_{t+1})|\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}]$$
$$E_t[s_{UC}(\omega_{t+1})|\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}]$$

are uniformly continuous in the conditioning variables.[9]

[9] Uniform continuity of (for example) s is not sufficient to assure that $E_t[s(\omega_{t+1})|\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}]$ is continuous. To see this, suppose that ω_{t+1} is a deterministic continuous function of ω_t, so $\omega_{t+1} =$

Assumption 1. *T and G are uniformly continuous in the weak-* topology.*

We will require that the type transition probability function T be tight. Furthermore, we will require a condition that is intermediate between tightness and uniform tightness of the family $\{T(\circ|\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})\}_{(\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}) \in \Theta \times \mathcal{A} \times \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})}$. In effect we require that the family of measures be tight with respect to transformations of the same set $U \subset \Theta$, a notion that we formalize in our next assumption. This provides the necessary leverage to make our asymptotics claims uniform over the family of type transition probability function. We note that compactness of Θ is obviously sufficient for Assumption 2, but stronger than required.

Assumption 2. *For any $\gamma > 0$ and each $(\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})$ we can choose a compact set U such that $T(U|\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}) > 1 - \gamma$ and U has radius at most $R(\gamma)$.*[10]

The utility function of each agent in each period of the N-agent game is

$$w_N : \Theta \times \mathcal{A} \times \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A}) \to \mathbb{R}$$

where $w_N(\theta, a, \omega, \pi^\Theta, \pi^\mathcal{A})$ is the payoff for an agent taking action a given his own type θ, aggregate state ω, and aggregate type and action distributions in the period equal to π^Θ and $\pi^\mathcal{A}$. The utility function of the nonatomic limit game is denoted

$$w : \Theta \times \mathcal{A} \times \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A}) \to \mathbb{R}$$

Assumption 3. *w is uniformly continuous*[11] *and bounded.*[12]

We assume that agent preferences can be represented in expected utility form when considering stochastic outcomes, Bayesian-Nash equilibria, or equilibria in mixed strategies. Intertemporal utility at time t is the exponentially discounted sum of

$h(\omega_t)$. If $h(\circ)$ is not uniformly continuous, then the composition $s(h(\omega_t)) = E_t[s(\omega_{t+1})|\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}]$ need not be uniformly continuous.

[10]The radius of a set U is $\sup_{\theta, \theta' \in U} d_\Theta(\theta, \theta')$.

[11]We require no continuity conditions on w_N.

[12]Our uniformity conditions can be verified for arbitrary continuous functions over compact domains using the Heine-Borel theorem.

utility in each stage of the dynamic game where $\delta \in (0,1)$ denotes the time discount factor. The following assumption can be viewed as either a regularity assumption on the relationship between the large finite games and the nonatomic limit game or as a definition of the nonatomic limit game approximation.

Assumption 4. ***(Uniform Pointwise Convergence)*** *For all $\varepsilon > 0$, there exists N^* such that for all $N > N^*$, all $\pi^\Theta \in \Delta_N(\Theta)$, $\pi^\mathcal{A} \in \Delta_N(\mathcal{A})$ and $(\theta, a, \omega) \in \Theta \times \mathcal{A} \times \Omega$*

$$|w_N(\theta, a, \omega, \pi^\Theta, \pi^\mathcal{A}) - w(\theta, a, \omega_t, \pi^\Theta, \pi^\mathcal{A})| < \varepsilon$$

The strategy space, a metric space (Σ, d_Σ), is a set of measurable maps from $\Theta \times \Omega \times \Delta(\Theta)$ into $\Delta(\mathcal{A})$ with the the sup-norm over $\sigma, \sigma' \in \Sigma$

$$d_\Sigma(\sigma, \sigma') = \sup_{(\theta,\omega,\pi^\Theta) \in \Theta \times \Omega \times \Delta(\Theta)} d^\mathcal{A}_{LP}(\sigma(\theta, \omega, \pi^\Theta), \sigma'(\theta, \omega, \pi^\Theta))$$

By placing measurability restrictions on the members of Σ, we can refine the set of equilibria. For example, if we wish to restrict our analysis to public equilibria, then we demand that the strategies in Σ be measurable with respect to the space $(\Omega, \mathcal{B}(\Omega))$. Information sets are defined in the model through measurability restrictions on T and Σ.[13] For example, private information about the aggregate state can be reflected by restricting w, w_N and σ to be independent of Ω (which captures the unobservability of ω_t) and allowing the type evolution operator, T, to depend on Ω (reflecting private information about ω_t embedded in the agent types).

Throughout this work we assume that all of the model primitives are measurable, and we are more explicit regarding measurability requirements in sections 2.3 and 2.4 when we define the stochastic evolution operators and equilibrium concepts in the large finite and nonatomic games.[14]

2.2. **Examples.** We now provide examples to illustrate how our framework can be applied.

Example 1. ***(Games of Complete Information)*** *In the simplest example agent types are constant over time, which entails transition probability function $T(\theta, U) =$*

[13]See further examples in section 2.2.

[14]For brevity we will not dwell on measurability requirements and assume throughout that required measurability restrictions are satisfied. The interested reader is urged to consult the excellent books by Pollard [28] or van der Vaart [33] for a full development of these technical issues.

$1\{\theta \in U\}$.[15] *The aggregate states,* $\Omega = \bigcup_{t=0}^{\infty} \left(\prod_{\tau=0}^{t-1} \Delta(\Theta \times \mathcal{A}) \right)$ *, reflect the history of all actions and types realized in the game. Agent utility is a function of the agent's own payoff type, his own action, and the distribution of actions taken by the other agents in the economy. In this setting, each possible (anonymous) history of play is a unique information set.*

Example 2. ***(Complete Information Stochastic Games Without Aggregate Shocks)*** *Let* θ *represent firm capital stock,* $a \in \mathcal{A}$ *be an investment decision, and the type evolution operator reflect stochastic capital investment outcomes. In this example, the type evolution operator is*

$$T(\circ|\theta_t, a_t) : \mathcal{B}(\Theta) \to [0, 1]$$

Firm profit depends on the distribution of types and actions of competitors, so we have $w_N(\theta, a, \pi^\Theta, \pi^\mathcal{A})$*. Finally we let* $\omega_t \in \Omega$ *reflect the complete history of action and type distributions.*

Example 3. ***(Repeated R&D Races)*** *Let* $a \in \mathcal{A}$ *denote an R&D investment decision and* θ *denote a stock of intellectual capital. The type evolution operator*

$$T(\circ|\theta_t, a_t, \pi_t^\Theta, \pi_t^\mathcal{A}) : \mathcal{B}(\Theta) \to [0, 1]$$

depends on π_t^Θ *and* $\pi_t^\mathcal{A}$*, which reflects the influence of investments by competitors on the value of R&D investments. If there are positive spillovers, then a state of the economy reflecting high investment makes it more likely that firm R&D investments succeed. In economies with negative externalities, firms aggressively pursuing research projects make it less likely that another firm's research project wins the race to a novel discovery. Firm profits are defined as in example* 4*, and* $\omega_t \in \Omega$ *reflects the complete history of actions and type distributions.*

Example 4. ***(Games of Imperfect Public Monitoring)*** *Let an agent's type be a fixed private payoff type. The aggregate state* ω_t *is a history of imperfect public signals regarding the actions of the agents in all past periods. The state evolution operator is*

$$G(\circ|\omega_t, \pi_t^\mathcal{A}) : \mathcal{B}(\Omega) \to [0, 1]$$

[15] The notation $1\{E\}$ is an indicator for the event E.

which reflects the distribution of public signals generated in period $t+1$ *as a result of the distribution of actions taken by the agents in period* t, $\pi_t^{\mathcal{A}}$. *The agent utilities are*

$$w_N : \Theta \times \mathcal{A} \times \Omega \to \mathbb{R}$$

so that agent utility depends on the agent's payoff type, the public signals, and the agent's own action. The information sets are represented by histories of public signals, ω_t, *and the private type of the agent,* θ.

Example 5. ***(Games of Imperfect Private Monitoring)*** *An agent's type* $\theta = (\theta_1, ..., \theta_{d+1})$ *is decomposed into a payoff type,* θ_1, *and a length* d *private monitoring history,* $(\theta_2, ..., \theta_{d+1})$, *that may or may not be correlated with the other agents' private information. The type evolution operator is*

$$T(\circ|\theta_t, \omega_t, \pi_t^{\mathcal{A}}) : \mathcal{B}(\Theta) \to [0,1]$$

which reflects the distribution of private signals generated in period $t+1$ *as a result of the distribution of actions taken by the agents in period* t, $\pi_t^{\mathcal{A}}$. *The agent utility*

$$w_N : \Theta \times \mathcal{A} \to \mathbb{R}$$

depends only on the privately observed signals, payoff type, and the agent's own action.

Suppose that $\omega_t \in \{A, B\}$, $\Pr\{\omega_t = A\} \in (0,1)$, *and* ω_t *is drawn independently across periods. Further, suppose that* $T(\circ|\theta_t, A, \pi_t^{\mathcal{A}}) \neq T(\circ|\theta_t, B, \pi_t^{\mathcal{A}})$, *which implies that agent type evolution conditional on* $(\theta_t, \pi_t^{\mathcal{A}})$ *is correlated across agents. By demanding that* G *allow for persistence, we can also induce intertemporal correlation of agent evolution. In the imperfect private monitoring example, information sets are described by private histories of signals observed by each player.*

2.3. **Evolution and Equilibrium - Large Finite Dynamic Games.** In these games a set of $N < \infty$ agents plays an infinite horizon dynamic game. The state space is the projection of Ψ onto $\Psi_N = \prod_{t=1}^{\infty} \Omega \times \Theta^N$ with the associated marginal probability measure P_N.[16] We assume that all random variables in the N-agent

[16]For any $N < \infty$, P_N defines a measure over the space $\Delta_N(\Theta)$ and, when joined with a strategy $\sigma \in \Sigma$, a measure over $\Delta_N(\mathcal{A})$.

games are measurable with respect to a filtration $\{\mathcal{F}_t^N\}_{t=0}^{\infty}$ where $\mathcal{F}_t^N \subset \mathcal{F}_{t+1}^N$ and $\mathcal{F}_t^N \subset \mathcal{F}_t^{N+1}$.[17]

The initial distribution of types in the N agent game is generated by N independent and identically distributed draws from $\pi_0^{\Theta}(\omega_0)$, which is a function of the initial aggregate state ω_0 as drawn from measure ν. The following lemma is immediate from the uniform law of large numbers stated in Corollary 8 in the Appendix.

Lemma 1. *The empirical distribution of types in the period* 0 *of the* N *agent game converges weakly to* $\pi_0^{\Theta}(\omega_0)$ *with a convergence rate of* $O(N^{-0.5})$ *uniformly over* $\pi_0^{\Theta}(\omega_0)$.

The evolution of the aggregates $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ is defined by the combination of $G(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ to describe the conditional distribution of ω_{t+1}, $T(\circ|\theta_t, a_t, \omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ describing the conditional distribution of $\{\theta_t^1, ..., \theta_t^N\}$, and the value of $\pi_{t+1}^{\mathcal{A}} \in \Delta_N(\mathcal{A})$ generated by the strategy $\sigma(\theta_{t+1}^i, \omega_{t+1}, \pi_{t+1}^{\Theta})$. The transition probability function for the large finite game is defined for any $U \subset \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$ and $W \subset \Omega$

$$P_N^A(U \times W|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}) = \int_W \Pr\{(\pi_{t+1}^{\Theta}, \pi_{t+1}^{\mathcal{A}}) \in U|\omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}\} * G(d\omega_{t+1}|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$$

The $\tau < \infty$ fold iteration of this operator is denoted $(P_N^A)^{\tau}(U \times W|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$. For a measurable function $f : \prod_{\tau=t+1}^{\infty} \Omega \times \Theta^N \to \mathcal{X}$, we use the notation $E_t^{\Psi}[f]$ to denote an expectation with respect to $\prod_{\tau=t+1}^{\infty} \Omega \times \Theta^N$ for $N < \infty$.

Having described the evolution of the large finite game conditional on an initial state $(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}}) \in \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$ and a symmetric strategy $\sigma \in \Sigma$, we define intertemporal utility at time t as

$$U_N(\mathbf{a}) = (1-\delta) * E_t^{\Psi}\left[\sum_{\tau=0}^{\infty} \delta^{\tau} w_N(\theta_{t+\tau}^i, a_{t+\tau}, \omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})\right]$$

Agent i's discounted expected utility in the N-agent game can be written in value function form

$$V_N(\theta_t^i, \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}|\sigma) = (1-\delta) * \{E_t^{\Psi}[w_N(\theta_t^i, \sigma(\theta_t^i, \omega_t, \pi_t^{\Theta}), \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})] + \delta E_t^{\Psi}\left[V_N(\theta_{t+1}^i, \omega_{t+1}, \pi_{t+1}^{\Theta}, \pi_{t+1}^{\mathcal{A}}|\sigma)\right]\}$$

[17]The requirement that $\mathcal{F}_t^N \subset \mathcal{F}_t^{N+1}$ reflects the embedding of the N agent game in the $N+1$ agent game.

Let $V_N(\theta_t^i, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma_i', \sigma_{-i})$ denote the utility of agent i in the N-agent game when he follows strategy σ_i' and all other agents follow strategy σ. We use the following notion of an approximate equilibrium in the large finite game.

Definition 1. *A symmetric $\varepsilon-$ Bayesian-Nash Equilibrium ($\varepsilon-$BNE) is a strategy and state $(\sigma^{BNE}, \omega_0, \pi_0^\Theta, \pi_0^\mathcal{A}) \in \Sigma \times \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$ of the N-agent game such that for all $\theta_0^i \in supp[\pi_0^\Theta]$*[18]

$$V_N(\theta_0^i, \omega_0, \pi_0^\Theta, \pi_0^\mathcal{A}|\sigma^{BNE}) + \varepsilon \geq \sup_{\sigma_i' \in \Sigma} V_N(\theta_0^i, \omega_0, \pi_0^\Theta, \pi_0^\mathcal{A}|\sigma_i', \sigma_{-i}^{BNE})$$

2.4. **Evolution and Equilibrium - Nonatomic Dynamic Games.** The initial distribution of types in the nonatomic limit game is $\pi_0^\Theta(\omega_0)$, which is a function of the initial aggregate state ω_0 as drawn from measure ν. In the nonatomic limit game there is no aggregate uncertainty regarding the evolution of the economy once we condition on the evolution of the aggregate state. To see this, note that each agent's type in period $t+1$ given their type and action in period t, the aggregate state ω_{t+1}, and the measures $(\pi_t^\Theta, \pi_t^\mathcal{A})$ is determined by $T(\circ|\theta_t, a_t, \pi_t^\Theta, \pi_t^\mathcal{A})$. Using intuitions based on the law of large numbers for a countable set of random variables, idiosyncratic shocks in the evolution of a single agent's type are eliminated in the aggregate.[19] For all $\omega_{t+1} \in \Omega, \pi_t^\Theta \in \Delta(\Theta)$, $\pi_t^\mathcal{A} \in \Delta(\mathcal{A})$ we define the transition operator for $(\pi_t^\Theta, \pi_t^\mathcal{A})$ as $P^C(\omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma^{DCE}) = (\pi_{t+1}^\Theta, \pi_{t+1}^\mathcal{A})$ where for any $V \in \mathcal{B}(\Theta), U \in \mathcal{B}(\mathcal{A})$[20]

$$\begin{aligned}
\pi_{t+1}^\Theta[V] &= \int_{\mathcal{A}\times\Theta} T(V|\theta, a, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}) * \sigma^{DCE}(\theta, \omega_t, \pi_t^\Theta)[da] * \pi_t^\Theta[d\theta] \\
\pi_{t+1}^\mathcal{A}[U] &= \int_{U\subset\mathcal{A}\times\Theta} \sigma^{DCE}(\theta, \omega_{t+1}, \pi_{t+1}^\Theta)[da] * \pi_{t+1}^\Theta[d\theta]
\end{aligned}$$

Note that given $(\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A})$, the distribution of types and actions realized at time $t+1$ is exogenous to a single agent's decision. The crucial conceptual difference between the large finite game and the nonatomic limit game is that the market

[18]The notation supp$[\pi]$ refers to the support of the probability measure π.

[19]Since the nonatomic dynamic game is founded on a continuum of random variables, $\Theta^{[0,1]}$, we cannot use the traditional law of large numbers and view these dynamics as a definition.

[20]The mapping P^C is not a transition probability function. Our notation differentiates between the deterministic and stochastic components of our nonatomic model.

aggregates are allowed to change in response to a single agent's action in a finite game. In the nonatomic limit game the market aggregates are exogenous to any single agent's action, and an agent's deviation can only affect his present period payoff and his future distribution of types.

We prove in Theorem 1 that the dynamics of the large finite game approach the nonatomic limit dynamics as $N \to \infty$. The state space for the nonatomic limit game is the probability space $(\Omega \times \Theta, \mathcal{B}(\Omega \times \Theta), P_{\Omega\times\Theta})$ where $P_{\Omega\times\Theta}$ denotes the marginal of P on $\Omega \times \Theta$.[21] We assume that all random variables in the nonatomic limit game are measurable with respect to a filtration $\{\mathcal{F}_t^{\infty}\}_{t=0}^{\infty}$. For any measurable function $f : \prod_{\tau=t+1}^{\infty} \Omega \times \Theta \to \mathcal{X}$ the notation $E_t^{\Omega}[f]$ refers to an expectation over the space of aggregate state and individual type evolution paths $\prod_{\tau=t+1}^{\infty} \Omega \times \Theta$. The transition probability function of the nonatomic game has the form for $U \subset \Delta(\Theta) \times \Delta(\mathcal{A}), W \in \mathcal{B}(\Omega)$,

$$P_{\infty}^{\mathcal{A}}(U \times W|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}) = \int_W 1\{P^C(\omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}|\sigma^{DCE}) \in U\} * G(d\omega_{t+1}|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$$

The $\tau < \infty$ fold iteration of this operator is denoted $(P_{\infty}^{\mathcal{A}})^{\tau}(U \times W|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$.

The intertemporal utility in the nonatomic dynamic game is

$$U(\mathbf{a}) = (1-\delta) * E_t^{\Omega}\left[\sum_{\tau=0}^{\infty} \delta^{\tau} w(\theta_{t+\tau}^i, a_{t+\tau}, \omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})\right]$$

where $\delta \in (0,1)$ is the time discount factor, a_t the action of the agent at time t, θ_t^i is agent i's type at time t, and $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ are market aggregates at time t. The value function form of utility in the nonatomic limit game when all agents play the symmetric strategy σ is

$$V_{\infty}(\theta_t^i, \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}|\sigma) = (1-\delta) * \{E_t^{\Omega}[w(\theta_t^i, \sigma(\theta_t^i, \omega_t, \pi_t^{\Theta}), \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})] + \delta E_t^{\Omega}\left[V(\theta_{t+1}^i, \omega_{t+1}, \pi_{t+1}^{\Theta}, \pi_{t+1}^{\mathcal{A}}|\sigma)\right]\}$$

[21]The agent's own type, $\theta_t \in \Theta$, remains a random variable at the individual level.

If agent i deviates from σ to σ', we refer to the strategy vector (σ'_i, σ_{-i}) and employ the value function notation

$$V_\infty(\theta^i_t, \omega_t, \pi^\Theta_t, \pi^\mathcal{A}_t|\sigma) = (1-\delta) * \{E^\Omega_t[w(\theta^i_t, \sigma'_i(\theta^i_t, \omega_t, \pi^\Theta_t), \omega_t, \pi^\Theta_t, \pi^\mathcal{A}_t)] + \delta E^\Omega_t \left[V(\theta^i_{t+1}, \omega_{t+1}, \pi^\Theta_{t+1}, \pi^\mathcal{A}_{t+1}|(\sigma'_i, \sigma_{-i}))\right]\}$$

Note that agent i's deviation affects only his own payoff and the evolution of his own type, while the evolution of the market aggregates $(\omega_t, \pi^\Theta_t, \pi^\mathcal{A}_t)$ is not affected. We define the following equivalent of an ε-BNE for nonatomic dynamic games.

Definition 2. *A symmetric $\varepsilon-$Dynamic Competitive Equilibrium (ε-DCE) consists of a strategy and state $(\sigma^{DCE}, \omega_0, \pi^\Theta_0, \pi^\mathcal{A}_0) \in \Sigma \times \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$ such that for all $\theta^i_0 \in \Theta$*

$$V_\infty(\theta^i_0, \omega_0, \pi^\Theta_0, \pi^\mathcal{A}_0|\sigma^{DCE}) + \varepsilon \geq \sup_{\sigma'_i \in \Sigma} V_\infty(\theta^i_0, \omega_0, \pi^\Theta_0, \pi^\mathcal{A}_0|\sigma'_i, \sigma^{DCE}_{-i})$$

A stationary equilibria is a nonatomic equilibrium concept wherein the economy is assumed to have no aggregate uncertainty and market aggregates remain constant over time.

Definition 3. *Assume $\Omega = \{\omega\}$ is a singleton. A symmetric $\varepsilon-$Stationary Equilibrium (ε-SE) consists of a strategy $\sigma^{SE} : \Theta \to \Delta(\mathcal{A})$ and type distribution $\pi^\Theta_\infty \in \Delta(\Theta)$ such that*

(1) *For all $V \in \mathcal{B}(\Theta), U \in \mathcal{B}(\mathcal{A})$*

$$\begin{aligned} \pi^\Theta_\infty[V] &= \int_{\mathcal{A}\times\Theta} T(V|\theta, a, \pi^\Theta_\infty, \pi^\mathcal{A}_\infty) * \sigma^{SE}(\theta)[da] * \pi^\Theta_\infty[d\theta] \\ \pi^\mathcal{A}_\infty[U] &= \int_U \sigma^{SE}(\theta)[da] * \pi^\Theta_\infty[d\theta] \end{aligned}$$

(2) *For all $\theta \in \Theta$ we have*

$$V_\infty(\theta^i_0, \omega, \pi^\Theta_\infty, \pi^\mathcal{A}_\infty|\sigma^{SE}) + \varepsilon \geq \sup_{\sigma'_i \in \Sigma} V_\infty(\theta^i_0, \omega, \pi^\Theta_\infty, \pi^\mathcal{A}_\infty|\sigma'_i, \sigma^{DCE}_{-i})$$

Conditions (1) implies that endogenous quantities are stationary given the equilibrium strategies. Condition (2) implies that for each type, the action dictated by the strategy is approximately optimal given $(\omega, \pi^\Theta_\infty, \pi^\mathcal{A}_\infty)$.[22] The usefulness of this

[22]For generic games the approximate optimality condition do not hold off the equilibrium path.

equilibrium concept lies in the fact the equilibrium strategy is finite dimensional and hence computationally tractable.[23]

If an analyst attempts to estimate the agent policy function in the context of a complete information stochastic game, the state space of the estimate is $\Theta \times \Omega \times \Delta_N(\Theta)$ and the dimension of the policy space diverges exponentially as $N \to \infty$. In the nonatomic limit game, the state of the economy at time $t+\tau$ is determined uniquely by the sequence of aggregate states $(\omega_{t+1}, ..., \omega_{t+\tau})$ and $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$. Therefore, strategies in a DCE can be written as functions of the initial state, $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$; the history of aggregate states between $t+1$ and $t+\tau$, $(\omega_{t+1}, ..., \omega_{t+\tau})$; and the agent's own type, $\theta_{t+\tau}$. Let $\Omega^t = \prod_{\tau=0}^{t} \Omega$ be the set of aggregate states realized in periods 0 through t with a generic element $\omega^t = (\omega_0, ..., \omega_t) \in \Omega^t$, and $\Omega^\infty = \cup_{t=0}^{\infty} \Omega^t$ denote the set of possible histories of aggregate states. Endow the space Ω^t with the sup-norm.[24] Any equilibrium strategy $\sigma^{DCE} : \Theta \times \Omega \times \Delta(\Theta) \to \Delta(\mathcal{A})$ can be written using an *indirect description* of the form $\sigma_{ID}^{DCE} : \Theta \times \Omega^\infty \to \Delta(\mathcal{A})$.

Given a DCE $(\sigma_{ID}^{DCE}, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$, we denote the associated family of sequences $\{((\omega_t, \pi_t^{\Theta}(\omega^t), \pi_t^{\mathcal{A}}(\omega^t))\}_{\omega^t \in \Omega^\infty}$ where ω_t is the final element of ω^t and

$$(P^C)^t(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}} | \sigma) = (\pi_t^{\Theta}(\omega^t), \pi_t^{\mathcal{A}}(\omega^t))$$

Definition 4. σ ***generates tight outcomes*** *if for all $\omega^t \in \Omega^t$, $t < \infty$, we have that $\pi_t^{\Theta}(\omega^t)$ is a tight probability measure given T, G, σ, and tight measures $\pi_0^{\Theta}(\omega_0)$.*[25]

The assumption that $\pi_t^{\Theta}(\omega^t)$ is a tight measure provides sufficient regularity to prove that the idiosyncratic uncertainty in our model disappears in the limit as $N \to \infty$. If the set Θ is compact, then this immediately implies that any strategy generates tight outcomes with respect to any choice of T, G, and $\pi_0^{\Theta}(\omega_0)$.[26] Although the assumption that σ generates tight outcomes is an assumption on the endogenous outcomes of our model, since we require tightness only over finite horizons (and not

[23]For existence results on exact DCE or SE in nonatomic games, see Bergin and Bernhardt [10].

[24]Therefore $d(\omega^t, \widetilde{\omega}^t) = \sup\{d_\Omega(\omega_t, \widetilde{\omega}_t)\}_{\tau \le t}$ where $\omega^t = (\omega_1, ..., \omega_t)$ and $\widetilde{\omega}^t = (\widetilde{\omega}_1, ..., \widetilde{\omega}_t)$. As a convention, let $d(\omega^r, \omega^s) = \infty$ if $r \neq s$

[25]Note that our definition implies that $\pi_0^{\Theta}(\omega_0)$ is a tight probability measure.

[26]Usually we make the assumption that σ generates tight outcomes without explicitly stating the initial measure $\pi_0^{\Theta}(\omega_0)$.

the stronger condition of uniform tightness) we conjecture that most models in which equilibria exist will satisfy this property.[27]

3. Results

The proofs of the approximation results of Section 3.2, Theorems 4 and 3, require both the convergence of the large finite game to the nonatomic limit game as $N \to \infty$ and that the nonatomic limit game be uniformly continuous in $\Theta \times \mathcal{A} \times \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$. We have assumed that w_N converges to a uniformly continuous and bounded nonatomic limit game utility function w. In Section 3.1 we state Theorem 1 and its corollaries, which provide conditions on the economic primitives sufficient for the dynamics of the market aggregates $(\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}})$ to be continuous in the limit as $N \to \infty$. These mean field results, combined with the continuity of w and a continuous strategy σ, are sufficient to prove that V_N uniformly converges to V_∞ and that V_∞ is uniformly continuous as required.

3.1. Dynamics.

Our analysis focuses on the family of transition probability functions $\{P_N^{\mathcal{A}}(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}})\}_{N=1}^\infty$ induced by the type evolution operator T, the aggregate state transition operator G, and a strategy σ played by all agents in the N agent game. We prove that for any finite τ^*, the behavior of the N-agent economy in periods $t+1$ through $t+\tau^*$ can be approximated by the nonatomic economy transition probability function for sufficiently large N. Intuitively, as the number of agents in the economy becomes large, the idiosyncratic shocks experienced by each agent are smoothed out. The proof of the theorem below uses continuity and stochastic convergence results to make this intuition precise.

Theorem 1. *Fix $\tau^* < \infty$, $\gamma > 0$ and $\rho \in [0,1)$. Assume $\sigma \in \Sigma$ is uniformly continuous. Then there exists $N^*, \overline{\gamma} > 0$ such that for any $(\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})$ where $d_\Omega(\omega_t, \widetilde{\omega}_t) + d_{LP}^\Theta(\pi_t^\Theta, \widetilde{\pi}_t^\Theta) + d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}}, \widetilde{\pi}_t^{\mathcal{A}}) < \overline{\gamma}$, any $N > N^*$, and all $\tau \in \{1, .., \tau^*\}$ we have that*

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^\tau(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})) < \gamma$$

[27] Adlakha et al. [1] discusses a closely related point that a light tailed condition is sufficient for both equilibrium existence and approximation (in a sense difference from our notion) in a family of models related to the model of Ericson and Pakes [17]. It is not clear whether we can replace our tightness assumption with a light-tailed condition as in Benkard et al. [7] or Adlakha et al. [1].

with probability at least $1-\rho$. *Furthermore the convergence rate is* $O(N^{-0.5})$ *and uniform over* $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$.

Convergence rates can be bounded using asymptotics results from empirical process theory. When we claim the convergence rate is $O(N^{-0.5})$ we are formally making the claim that

$$\Pr\{\sqrt{N} * d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^{\tau}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), (P_{\infty}^{\mathcal{A}})^{\tau}(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}})) > \gamma\} < Ce^{-2\gamma^2}$$

where C is a function of the dimension d and τ^*. The two equivalent ways of stating the convergence rate are:

- Holding $\rho > 0$ fixed, δ can be chosen to converge to 0 at a rate $O(N^{-0.5})$
- Holding $\delta > 0$ fixed, ρ can be chosen to go to 0 at a rate $O(e^{-n})$

Both δ and ρ determine the value of epsilon in our approximation theorems, and we cannot in general determine which of δ or ρ (or both) are limiting the approximation precision. We choose the statement emphasizing the convergence of δ as it is likely to be more familiar to the reader. Note that all of the corollaries in this section share the $O(N^{-0.5})$ convergence rate demonstrated in the proof of Theorem 1.

Applications such as games of repeated private monitoring (see Section 4.1) represent cases where the set of types admissible at any time $t < \infty$ is a finite dimensional Euclidean space, but the type space becomes infinite dimensional as $t \to \infty$. For example, an agent's type could reflect a history of signals drawn from $\Phi \subset \mathbb{R}$, where each signal reflects information about the distribution actions in the prior period. Denote the set of types that can be realized in period t as $\Theta^t = \prod_{\tau=0}^{t} \Phi$. We let the distribution of period $t+1$ types be conditionally independent across agents and distributed according to the operator $T_{t+1}(\circ|\theta_t, a_t, \omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$. In this setting we have the following theorem

Corollary 1. *Fix* $\tau^* < \infty$, $\gamma > 0$ *and* $\rho \in [0,1)$. *Assume:*

- $\{T_{t+\tau}(\circ|\theta_t, a_t, \omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})\}_{\tau=1}^{\tau^*}$ *are uniformly continuous*
- $\sigma \in \Sigma$ *is uniformly continuous over* $\Theta^{\tau^*} \times \Omega \times \Delta(\Theta^{\tau^*})$

There exists $N^*, \overline{\gamma} > 0$ *such that for any* $(\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})$ *where* $d_\Omega(\omega_t, \widetilde{\omega}_t) + d_{LP}^\Theta(\pi_t^\Theta, \widetilde{\pi}_t^\Theta) + d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}}, \widetilde{\pi}_t^{\mathcal{A}}) < \overline{\gamma}$*, any* $N > N^*$*, and all* $\tau \in \{1, .., \tau^*\}$ *we have that*

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^A)^\tau(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (P_\infty^A)^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})) < \gamma$$

with probability at least $1 - \rho$*.*

Given a strategy $\sigma \in \Sigma$, consider a deviation by agent i to a strategy $\sigma' \in \Sigma$ and denote the strategy vector in the N agent game as (σ_i', σ_{-i}). Let P_N^A denote the transition probability function given Markov operators T and G and symmetric strategy σ, and let $\widetilde{P}_N^A$ be the corresponding transition probability function given strategy profile (σ_i', σ_{-i}). We now prove that since the evolution operators are continuous, a large deviation in strategy by a single agent has a negligible effect on the market dynamics as $N \to \infty$.

Corollary 2. *Fix* $\tau^* < \infty$, $\gamma > 0$ *and* $\rho \in [0, 1)$*. Assume* $\sigma \in \Sigma$ *is uniformly continuous. Suppose agent* i *deviates from* σ *to* $\sigma' \in \Sigma$ *in period* t*.*[28] *Then there exists* $N^*, \overline{\gamma} > 0$ *such that for any* $(\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})$ *where* $d_\Omega(\omega_t, \widetilde{\omega}_t) + d_{LP}^\Theta(\pi_t^\Theta, \widetilde{\pi}_t^\Theta) + d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}}, \widetilde{\pi}_t^{\mathcal{A}}) < \overline{\gamma}$*, any* $N > N^*$*, and all* $\tau \in \{1, .., \tau^*\}$ *we have that*

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((\widetilde{P}_N^A)^\tau(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (P_\infty^A)^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})) < \gamma$$

with probability at least $1 - \rho$*.*

Our final corollary extends Theorem 1 to prove that for large N the evolution of an economy wherein all of the agent adopt σ is approximated by the evolution of an economy in which all agents choose σ' if σ and σ' are sufficiently close in strategy space. Let P_∞^A denote the evolution of the nonatomic game generated by σ in conjunction with T and G, and let $\widehat{P}_N^A$ reflect the market evolution generated by σ', T, and G in the N agent game.

Corollary 3. *Fix* $\tau^* < \infty$, $\gamma > 0$ *and* $\rho \in [0, 1)$*. Assume* $\sigma, \sigma' \in \Sigma$ *are uniformly continuous. Then there exists* $N^*, \overline{\gamma} > 0$ *such that for any* $(\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})$ *where* $d_\Omega(\omega_t, \widetilde{\omega}_t) + d_{LP}^\Theta(\pi_t^\Theta, \widetilde{\pi}_t^\Theta) + d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}}, \widetilde{\pi}_t^{\mathcal{A}}) + d_\Sigma(\sigma, \sigma') \leq \overline{\gamma}$*, any* $N > N^*$*, and all*

[28] We do not require σ' to be continuous.

$\tau \in \{1, .., \tau^*\}$ *we have that*

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((\widehat{P}_N^{\mathcal{A}})^\tau(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})) < \gamma$$

with probability at least $1 - \rho$.

We note that at the cost of much added notational complexity, we could extend the theorems above to handle cases wherein the stochastic processes in the limit game, T and G, differ from those in the N-agent game, T_N and G_N. It is straightforward to see from our proofs that we would require the following regularity assumption.

Assumption 5. *(**Uniform Pointwise Convergence**) For all* $\gamma > 0$*, there exists* N^* *such that for all* $N > N^*$*, all* $\pi^\Theta \in \Delta_N(\Theta)$*,* $\pi^{\mathcal{A}} \in \Delta_N(\mathcal{A})$ *and* $(\theta, a, \omega) \in \Theta\times\mathcal{A}\times\Omega$

$$\begin{aligned} d_{LP}^\Theta(T_N(\circ|\theta, a, \omega, \pi^\Theta, \pi^{\mathcal{A}}), T(\circ|\theta, a, \omega_t, \pi^\Theta, \pi^{\mathcal{A}})) &< \gamma \\ d_{LP}^\Omega(G_N(\circ|\omega, \pi^\Theta, \pi^{\mathcal{A}}), G(\circ|\omega_t, \pi^\Theta, \pi^{\mathcal{A}})) &< \gamma \end{aligned}$$

3.2. **Approximation Theorems.** The core tools for the proofs of our equilibrium approximation results are Theorem 1 and its corollaries, which show that for any finite time horizon τ the evolution of the large finite game with a sufficiently large number of agents will with high probability resemble the outcomes realized in a nonatomic limit game. Given that the agent utility is bounded and continuous, this implies that the incentives of an agent in the large finite game will be close to those facing an agent of the same type in the nonatomic limit game. It follows that the optimal action for the agent in the nonatomic limit game is close to the optimal action of the agent in the large finite game and strategic convergence follows.

We now use our approximation techniques to show that any DCE represents an $\varepsilon-$BNE when our continuity assumptions hold.[29]

Theorem 2. *Fix* $\varepsilon > 0$*. Assume that* σ_{ID}^{DCE} *is uniformly continuous*[30]*. Then we can choose* N^**such that the DCE* $(\sigma_{ID}^{DCE}, \omega_0, \pi_0^\Theta, \pi_0^{\mathcal{A}})$ *is an* ε*-Bayesian-Nash Equilibrium of the large finite stochastic game for* $N > N^*$ *players. Furthermore,* N^* *can be chosen uniformly across* $(\sigma_{ID}^{DCE}, \omega_0, \pi_0^\Theta, \pi_0^{\mathcal{A}})$.

[29]Our results are similar in character to the nonstationary oblivious equilibria studied by Benkard et al. [8].

[30]The minimal requirement is continuity over Ω^t for sufficiently large t.

Theorem 2 proves that when an econometrician computes the equilibria of a structural model using a nonatomic limit approximation, the equilibrium strategies are an $\varepsilon-$BNE of the corresponding large finite game. Furthermore we can take $\varepsilon \to 0$ as $N \to \infty$, which implies that the dynamic competitive equilibrium strategy exactly satisfies the equilibrium conditions in the asymptotic limit. The DCE strategy requires a dimension[31] of $dim(\Theta) * \tau * \dim(\Omega)$, independent of the number of agents, to describe the equilibrium behavior for the future τ periods.[32] If an analyst attempts to estimate the policy function of a complete information stochastic game, the state space of the estimate would be $\Theta \times \Omega \times \Delta_N(\Theta)$ and the dimension of the policy space diverges exponentially as $N \to \infty$. Therefore, estimating the exact BNE policy function in complete information large games is, in practice, computationally impossible.[33] Furthermore, since the agents gain little by optimizing over this larger state space, the DCE may be an appealing behavioral model if we assume small costs of optimizing with respect to the out-of-equilibrium states.

The convergence rate of ε is important for the practical application of our theorems. If the transition functions (T and G) and the utility function (w) are Lipschitz continuous, then it is straightforward to show that ε shrinks at a rate of $O(N^{-1/2})$.[34] In practice once a DCE strategy has been computed, ε can be found for a particular N-agent game by calculating the optimal deviation for an agent given that all of the other agents follow the DCE strategy. Computing such an optimal deviation is (relatively) tractable since this is the result of a single person decision problem. Prior work suggests that the convergence rate can be much faster - for example

[31]The initial point $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ is formally an infinite dimensional datum. In estimation tasks this datum is embedded into the function estimated, which is a map from $\Theta \times \Omega$ into $\mathcal{A}$. When we describe the function as finite dimensional, we refer to the dimension of $\Theta \times \Omega$ alone.

[32]The dimension of the space Ω^t does diverge to infinity as $t \to \infty$, but some model assumptions must be made to regulate behavior for large t since any estimation must be based on histories of finite length.

[33]One potential technique for reducing the dimension of the policy function of an exact BNE of a game of incomplete information is to make assumptions about the information observed by each agent. For example, public equilibrium policy functions will not suffer a curse of dimensionality so long as the dimension of the public signal remains fixed as N grows large. However, describing the beliefs of the agents (without strong symmetry assumptions) may result in a curse of dimensionality in large games of incomplete information.

[34]*** INSERT REF TO OPTMIZATION BOOK***

static private values double auctions have been shown to converge at a rate equal to $O(N^{-2+\alpha})$ for any $\alpha > 0$ (Cripps and Swinkels [15]).

As an immediate corollary to Theorem 2, we have that for any $\varepsilon > 0$, any stationary equilibrium of the nonatomic limit game is an $\varepsilon-$BNE of a sufficiently large finite game.[35]

Corollary 4. *Consider a Stationary Equilibrium of the nonatomic limit game, $(\sigma^{SE}, \pi_\infty^\Theta)$. Assume that:*

- *There is no aggregate uncertainty ($\Omega = \{\omega\}$)*
- *w is uniformly continuous and bounded in a relatively open set containing $\Theta \times \{\pi_\infty^\Theta\} \times \Delta(\mathcal{A})$*
- *σ^{SE} is uniformly continuous*

For any $\varepsilon > 0$ we can choose $N^* < \infty$ and $\overline{\gamma} > 0$ such that σ^{SE} is an ε-Bayesian-Nash Equilibrium of the large finite dynamic game for $N > N^*$ starting at any $\pi_0^\Theta \in \Delta(\Theta)$ and $\pi_0^\mathcal{A} \in \Delta(\mathcal{A})$ such that $d_{LP}^\Theta(\pi_0^\Theta, \pi_\infty^\Theta) + d_{LP}^\mathcal{A}(\pi_0^\mathcal{A}, \pi_\infty^\mathcal{A}) < \overline{\gamma}$.

Theorem 3 strengthens our approximation results by providing sufficient conditions under which convergent sequences of exact BNE have a limit in the set of DCE. Let the set of exact BNE of the N-agent game be denoted by the correspondence $\mathcal{E} : \mathbb{N} \rightrightarrows \Sigma$, and denote the set of exact DCE of the nonatomic limit game as $\mathcal{E}^{NA}$. The following theorem assures us that a DCE is not merely an ε-BNE, but that any convergent sequence of BNE approaches some DCE strategy. Therefore, the approximate equilibrium strategies that result from computing the set of DCE policy functions delimits the set of possible equilibrium strategies of any exact BNE of the large finite game.

Theorem 3. *Assume that*

- *Θ and Ω are compact*
- *There exists an N^* such that $\cup_{N=N^*}^\infty \mathcal{E}(N)$ is uniformly equicontinuous[36] in $\Theta \times \Omega \times \Delta(\Theta)$*

[35] This result is similar to the claim that Stationary equilibria possess the Asymptotic Markov property of Benkard et al. [7]. See Section 4.6 for a comparison of our conditions to those in Benkard et al.

[36] We assume only that the family of functions $\cup_{N=N^*}^\infty \mathcal{E}(N)$ satisfy the same modulus of continuity over their respective domains. So if the functions are uniformly equicontinuous with modulus of

Then the correspondence $\mathcal{E}$ is upper hemicontinuous with

$$\underset{N\to\infty}{Lim}\mathcal{E}(N) = \mathcal{E}^{\infty} \subset \mathcal{E}^{NA}$$

We interpret Theorem 2 as providing conditions under which the DCE are approximate BNE in terms of marginal incentives, whereas Theorem 3 provides the stronger result that the BNE policy functions may be approximated by (a subset of) the DCE as $N \to \infty$.[37] Theorem 2 does not require the uniform equicontinuity restriction of Theorem 3, an assumption that can be difficult to verify in practice. On the other hand, by neglecting these equicontinuity restrictions we do not have Theorem 3's guarantee of upper hemicontinuity of the equilibrium correspondence, which weakens the link between the set of exact BNE and the approximate BNE described by the DCE of the limit game.

The principal roadblock to directly applying our theorem is the necessity of verifying the continuity of the limit of arbitrary convergent sequences of equilibrium strategies of the large finite games.[38] We focus attention in our examples on cases where these restrictions are weak. For example, a game wherein $\Theta \times \mathcal{A} \times \Omega$ is discrete, $\Delta(\Theta)$ is unobservable, and T and G do not depend on $\Delta(\Theta) \times \Delta(\mathcal{A})$ has only uniformly continuous strategies in the strategy space.[39] Applications with natural discontinuities have to be handled by careful manipulation of topologies to establish the continuity properties required for our results (see Section 4.6).

---

continuity $\delta(\varepsilon)$, then for any $\sigma \in \mathcal{E}(N)$, $N \geq N^*$, and any $(\widetilde{\theta}, \widetilde{\omega}, \widetilde{\pi}^{\Theta}), (\theta, \omega, \pi^{\Theta}) \in \Theta \times \Omega \times \Delta_N(\Theta)$ such that

$$d_{\Theta}(\widetilde{\theta}, \theta) + d_{\Omega}(\widetilde{\omega}, \omega) + d_{LP}^{\Theta}(\widetilde{\pi}^{\Theta}, \pi^{\Theta}) < \delta(\varepsilon)$$

then we have

$$d_{LP}^{\mathcal{A}}(\sigma(\widetilde{\theta}, \widetilde{\omega}, \widetilde{\pi}^{\Theta}), \sigma(\theta, \omega, \pi^{\Theta})) < \varepsilon$$

[37] Examples of cases in which epsilon equilibrium policy functions are not close in strategy space to any exact equilibrium policy function exist in many branches of the game theory literature. In the case of large market games, early examples (and sufficient conditions under which these examples are eliminated) are provided by Roberts and Postlewaite [29].

[38] The Arzelà-Ascioli theorem, which we use to prove continuity of the limit strategy σ^{∞}, necessitates that we assume Θ and Ω are compact and that the equilibrium strategy correspondence becomes equicontinuous for sufficiently large N.

[39] An additional tool is the Heine-Borel theorem, which states that continuity can be strengthened to uniform continuity if the primitives are drawn from compact spaces.

If we do not assume the equicontinuity of $\cup_{N=N^*}^{\infty}\mathcal{E}(N)$, it is not clear that the limit $\sigma^N \to \sigma^\infty$ will be continuous in $\Theta \times \Omega \times \Delta(\Theta)$.[40] We can alternatively *assume* that the limit $\sigma^N \to \sigma^\infty$ exists and σ^∞ is uniformly continuous. We can then strengthen our result by requiring continuity of T, G, and w only in a neighborhood of the support of the aggregate variables in the nonatomic limit game as opposed to the whole space $\Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$. These alternate assumptions lead to the following corollary.

Corollary 5. *Consider a convergent sequence of uniformly continuous Bayesian-Nash equilibrium strategies, $\{\sigma^N\}_{N=1}^{\infty}$ and $\sigma^N \to \sigma^\infty$. Assume σ^∞ is uniformly continuous. Let $\Lambda^\tau = \prod_{t=0}^{\tau} supp[(\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A})]$ be the support of the stochastic process P_∞^A. For N sufficiently large, if there exists an open set U in $\prod_{t=0}^{\infty} \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$ containing Λ^τ such that...*

- *T is uniformly continuous over $\Theta \times \mathcal{A} \times U$*
- *G is uniformly continuous over U*

Then σ^∞ is a Dynamic Competitive Equilibrium of the limit game.

One conclusion we can draw from our corollary is that if we wish to study when a convergent sequence of equilibria $\{\sigma^N\}_{N=1}^{\infty}$ will not yield a DCE in the limit, discontinuities are crucial. Economically these discontinuities may allow the agents to remain pivotal over future outcomes even in cases where the population is large.

4. Applications

4.1. Large Games of Imperfect Private Monitoring.

Games of imperfect private monitoring are remarkable both for the promise of wide application and the difficulty of computing the set of equilibria. For example, games with perfect public information structures are idealizations of a reality wherein each agent obtains a private, noisy observation of a public signal with this noise resulting from the observation technology, mistakes, or agent observations that are not simultaneous. This problem is exacerbated by the fact that communicating data regarding these real world observations is, at best, complex and imperfect (but see McLean et al.

[40] One can easily generate examples wherein Theorem 1 fails to apply and the stochastic processes described by P_N^A and P_∞^A are not similar even in the short run if σ_∞ is not continuous.

[27]). Describing the set of equilibria in classic repeated games such as the prisoner's dilemma remains daunting in the context of private monitoring.

In this section we provide an anti-folk theorem result for large games of imperfect private monitoring. Our results turn on the fact that continuous large games sharply limit the capacity for agents to detect defections from tacit collusion. Most folk theorem results (for example Fudenberg et al. [18]) are inapplicable this condition fails. Our result suggests that relying on (relatively) tractable analysis and estimation techniques designed for competitive economies when studying these games can economize on computational and analytical difficulty as well as the assumptions required to generate predictions. Further, our anti-folk theorem suggests that attempts to find noncompetitive outcomes in large dynamic economies in the industrial organization, labor economics, or macroeconomics literature will be quite difficult. At a minimum, we highlight the roles that perfect monitoring and discontinuity play in theoretical predictions of imperfectly competitive outcomes in large markets.

Assume that in period t, the agent observes a signal ϕ from the discrete set Φ that is informative about the distribution of actions in period $t-1$.[41] An agent's type $\theta_t = (\phi_1, \phi_2, ..., \phi_t)$ reflects a history of these private signals. Note that the type space at period $t < \infty$, denoted Θ^t, is a subset of a finite dimensional Euclidean space with the full type space denoted $\Theta = \cup_{t=1}^{\infty} \Theta^t$. The stochastic process generating each agent's private monitoring signal in period t is

$$\rho(\circ|\pi_{t-1}^{\mathcal{A}}, \omega_t) : \mathcal{B}(\Phi) \to [0,1]$$

which we assume is continuous in $\pi_{t-1}^{\mathcal{A}}$. Therefore for any $\tau < \infty$, the type evolution operators $\{T_{t+\tau}(\circ|\theta_t, \omega_{t+1}, \pi_t^{\mathcal{A}})\}_{t=1}^{\tau}$ are uniformly continuous. The aggregate variable $\omega_t \in \Omega$, unobservable to the agents, mediates correlation across the agents in each period and across time. The evolution of ω_t, $G(\circ|\omega_t)$, is independent of agent actions. Since we have placed no restrictions on Ω, our private monitoring structure allows for rich signal correlations between agents and persistence across time. Our only restriction is that the signals, conditional on ω_t, must be continuous in $\pi_{t-1}^{\mathcal{A}}$.[42]

[41] We can allow for continuous signal spaces if we assume that the equilibria of the large finite game admit only uniformly continuous strategies.

[42] It is straightforward to incorporate privately known payoff types for the agents that evolve in a Markovian fashion or aggregate signals and payoff shocks.

We assume that the agent utilities are of the form

$$w : \Theta \times \mathcal{A} \to \mathbb{R}$$

and that agents employ strategies $\sigma : \Theta \to \Delta(\mathcal{A})$. Since an individual agent cannot affect $\pi_t^{\mathcal{A}}$ in the limit game (and hence has no effect on the information of the other agents in the economy), for all $a^* \in$supp$[\sigma^{DCE}(\theta)]$

$$a^* \in \arg\max_{a \in \mathcal{A}} w(\theta, a)$$

In the case of the repeated prisoner's dilemma, the only equilibrium is for all agents to defect each period.

Since Θ^t is a discrete set, any convergent sequence of strategies $\{\sigma^N : \Theta \to \Delta(\mathcal{A})\}_{N=1}^{\infty}$ is uniformly equicontinuous and hence has a uniformly continuous limit $\sigma^{\infty} : \Theta \to \Delta(\mathcal{A})$. Theorem 3 implies that the equilibria of the nonatomic limit game capture the full range of behavior of the equilibria of sufficiently large finite games. In our prisoner's dilemma example, any sequence of equilibrium strategies converges to the static Nash equilibrium strategy of defection in every period.

Theorem 4. *For any $\delta < 1$, as $N \to \infty$ the set of exact BNE strategies approach the set of repeated static Nash equilibria.*

Proof. Note that uniform continuity of T, σ and G are assumed. Therefore Theorem 3 implies that the equilibrium strategy correspondence is upper hemicontinuous in N. Since agents treat the signal generation process as exogenous to their own action in the nonatomic game, repeated static Nash play is the only equilibria of the nonatomic limit game equilibrium. Our theorem immediately follows. □

Our result is driven by the infinitesimal effect of the agents' individual actions on the signals observed by the other agents. Folk theorems typically generate result for the case wherein N is held fixed and we let $\delta \to 1$. In this setting, even small influences on $\rho(\circ|\pi_{t-1}^{\mathcal{A}}, \omega_t)$ have strong incentive effects on individual agents for large enough δ. However, we hold $\delta < 1$ fixed and let $N \to \infty$, allowing the effect of an agent's action to fall below the threshold where the monitoring and punishment possibilities have any power to provide incentivizes.

4.2. **Large Games of Imperfect Public Monitoring.** We focus on public equilibria[43] in which agent actions are conditioned only on a public history of aggregate signals. At the beginning of period t the agents observe a public signal $y_t \in Y$ of the actions taken in period $t-1$ that is distributed according to a measure $\rho(y_t|\pi^{\mathcal{A}}_{t-1})$. We assume that Y is finite and the aggregate state space is $\Omega = \prod_{t=0}^{\infty}(Y \cup \{\varnothing\})$. The null signal $\varnothing$ is a notational convention for denoting public signals in future periods. A valid aggregate state at time t would be $\omega_t = (y_0, .., y_t, \varnothing, ..., \varnothing, ...)$. To the extent that agents have private information regarding payoffs, such information is reflected by a fixed type θ drawn from a finite set Θ at $t = 0$.

We restrict ourselves to strategies within the set $\Sigma = \{\sigma : \Theta \times \Omega \to \Delta(\mathcal{A})\}$. These strategies are uniformly equicontinuous from the finiteness of $\Theta \times \Omega$. Agent utility is

$$w : \Theta \times \Omega \times \mathcal{A} \to \mathbb{R}$$

Using our approximation theorems we can prove the following anti-folk theorem

Theorem 5. *Assume that ρ is continuous. As $N \to \infty$, the set of exact public equilibrium strategies approaches the set of repeated static Nash equilibria.*

Proof. Note that continuity of G follows from continuity of ρ, and continuity of T is immaterial since each agent's type is fixed. Uniform equicontinuity of Σ is assumed. From Theorem 3 we have that the equilibrium strategy correspondence is upper hemicontinuous. In the nonatomic limit game, the agents treat the signalling process as exogenous, which implies that the only admissible equilibrium strategies involve repeated static Nash play. Therefore, the exact public equilibrium strategies approach repeated static Nash play as $N \to \infty$. □

Now we consider an example wherein our theorems are inapplicable directly, but we can gain insight into the limits of equilibrium outcomes by employing Corollary 5. Consider an N-agent Prisoner's Dilemma game. Let the action space be $\mathcal{A} = \{c, d\}$ and the public signal space be $Y = \{0, 1\}$. Ω denotes an infinite, aggregate history of public signals, while Θ denotes a privately observed randomization device drawn at $t = 0$ independently across agents. Let $f_t = \pi^{\mathcal{A}}_t[\{c\}]$ denote the fraction of agents

[43] We do not impose perfection at this point, but our results carry over to perfect public equilibria (see Section 4.4).

playing c at time t, and assume the signal generation process, $\rho(y_t = 1|f_t)$, is strictly increasing in f_t. Agent utility, $w(a_t, y_t)$, is defined by the following payoff matrix

	$y_t = 1$	$y_t = 0$
$a_t = c$	1	$-\epsilon$
$a_t = d$	$1+\epsilon$	0

for some $\epsilon > 0$.

Since agents in the nonatomic limit game cannot affect f_t, the sole DCE equilibrium strategy, denoted σ^d, is to defect after every history. Theorem 5 implies that for a convergent sequence of exact Bayesian-Nash equilibria of the N-agent games, $\{\sigma^N\}_{N=1}^{\infty}$, $\sigma^N \to \sigma^d$ if $\rho(y_t = 1|f_t)$ is continuous in f_t.

Suppose $\rho(y_t = 1|f_t)$ has a discontinuity at $f_t = \frac{1}{2}$. Corollary 5 implies that σ^N will converge to a strategy σ^∞ such that on the equilibrium path we have $f_t \in \{0, \frac{1}{2}\}$. For example, $\{\sigma^N\}_{N=1}^{\infty}$ could converge to a grim trigger strategy of the form

$$\sigma^\infty(\omega_t)[c] = \{ \begin{array}{l} \frac{1}{2} \text{ if } \omega_t = (y_0 = 1, y_1 = 1, ..., y_{t-1} = 1) \\ 0 \text{ otherwise} \end{array}$$

Alternatively, agent actions could be persistent across periods and asymmetric across agents with these features mediated by the private information, Θ. The agents could also periodically oscillate between the repeated static Nash equilibrium outcome ($f_t = 0$) and collusive behavior ($f_t = \frac{1}{2}$).

Corollary 5 rules out the existence of collusive equilibria except at points of discontinuity, which proves the necessity of these discontinuities for the presence of collusive behavior in equilibrium - the existence of these equilibria also depends on payoffs, discount factors, the information structure, and other primitives of the model. The discontinuity allows agents to have an influence over the evolution of aggregate variables, and hence the actions of other agents, even as $N \to \infty$. The presence of discontinuities as an empirical matter is difficult to test. Even in designed markets, the presence of unmodeled aggregate noise or idiosyncratic errors can smooth out discontinuities deliberately created by the designer. Accepting this fact, if a market designer wishes to support outcomes outside of repeated static Nash play in a large game, the designer is well advised to introduce discontinuous incentives (w_N);

informative, noise-free private signals (T); and to tightly control the aggregate noise present in the environment (G).

4.3. **Games With a Few Large Players.** There are many applications wherein a few economically significant agents participate in a game with a large number of small players. For example, a monopolist supplier could work against the efforts of a cartel of small buyers to reduce the market price. In the context of a corporate takeover a few large shareholders may attempt to reap a profit by making a takeover bid and installing new management, but a fringe of small shareholders may prevent this by free-riding on the efforts of the raiders. Macroeconomists model taxpayers, consumers, financial asset holders, or firms as price takers whereas governments, market makers, and other institutions are game theoretic actors.

Leader-follower games, such as the corporate takeover example above, generate sharply different results if one assumes that the small players take the large players' actions as given instead of treating the small players as strategic actors who act as if the large player could respond to the choices of individual small players (e.g. Fudenberg et al. [19], Levine and Pesendorfer [25]). Our approximation theorems provide conditions under which it is formally sound to assume a large set of small agents can be modeled as a continuum of nonatomic agents, and when such an assumption (often made for tractability) might preclude equilibria that are admissible predictions in a more complex, game theoretic model.

The essence of a large player is one whose actions can have significant effects on market aggregates even in the limit as the number of agents approaches infinity. Our previous results rule out the presence of large players by assuming uniform continuity across the spaces $\Delta(\Theta)$ and $\Delta(\mathcal{A})$. In this section we segment the agents into two groups: a set of M large players and a set of N small players. We provide assumptions on the model that imply that the N small agents can be treated as a measure 1 continuum in the limit as $N \to \infty$ and M remains fixed. In contrast, each of the large agents is treated as a measure 1 atom.[44]

To simplify our discussion, we focus on a model wherein the actions of the large players are observable. Large players choose actions from space $\mathcal{A}_L$, and we denote

[44]Representing the N small agents as a measure 1 continuum is an innocuous normalization.

large player m's choice of action in period t as $a_{m,t}$. A history of actions by the M large players through period τ is an element of $\mathcal{A}_L^{\tau \times M}$. We encode the actions of the large players into the aggregate variable by letting $\Omega = \cup_{\tau=1}^{\infty} \mathcal{A}_L^{\tau \times M}$. Let $\omega_t = \omega_{t-1} \times (a_{1,t}, ..., a_{M,t})$ denote the concatenation of period t action vector $(a_{1,t}, ..., a_{M,t})$ onto history ω_{t-1} and let $\omega_0 = \varnothing$.[45] Let the utility functions of the large players in the game with N small agents be $\{u_{m,N} : \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})\}_{m=1}^{M}$ and the utility of the M large players in the nonatomic limit game be denoted $\{u_m : \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})\}_{m=1}^{M}$. We assume that the utility functions of the M large players are well-defined in the limit as $N \to \infty$.

Assumption 6. ***(Uniform Pointwise Convergence)*** *For all $\varepsilon > 0$, there exists N^* such that for all $N > N^*$, all $\pi^{\Theta} \in \Delta_N(\Theta)$, $\pi^{\mathcal{A}} \in \Delta_N(\mathcal{A})$ and $(\theta, a, \omega) \in \Theta \times \mathcal{A} \times \Omega$*

$$|u_{m,N}(\omega, \pi^{\Theta}, \pi^{\mathcal{A}}) - u_m(\omega_t, \pi^{\Theta}, \pi^{\mathcal{A}})| < \varepsilon$$

Given a fixed vector of strategies for the M large players, $\overline{\sigma} = (\overline{\sigma}_m : \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A}) \to \mathcal{A}_L)_{m=1}^{M}$, the distribution of $(\omega_1, \omega_2, \omega_3, ...)$ is common knowledge. Denote the equilibrium correspondence for the N small players as $\mathcal{E}(N|\overline{\sigma})$. Denote the equilibria of the nonatomic limit game as $\mathcal{E}^{NA}(\overline{\sigma})$.

Theorem 6. *Consider a fixed set of strategies $\overline{\sigma} = (\overline{\sigma}_m : \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A}) \to \mathcal{A}^L)_{m=1}^{M}$ for the M large players. Assume*

- *Θ and Ω are compact*
- *There exists N^* such that $\cup_{N=N^*}^{\infty} \mathcal{E}(N)$ is uniformly equicontinuous in $\Theta \times \Omega \times \Delta(\Theta)$*
- *$\{\sigma_m : \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A}) \to \mathcal{A}_L\}_{m=1}^{M}$ are uniformly equicontinuous*

Then the small player equilibrium strategy correspondence $\mathcal{E}$ is upper hemicontinuous with

$$\underset{N \to \infty}{Lim} \mathcal{E}(N|\overline{\sigma}) \subset \mathcal{E}^{NA}(\overline{\sigma})$$

Proof. Note that assuming that $\{\sigma_m : \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A}) \to \mathcal{A}_L\}_{m=1}^{M}$ is uniformly equicontinuous implies that G is uniformly continuous in $\Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$. Theorem 3 then yields our result. □

[45]Extensions to include aggregate shocks or to use Ω to capture correlations between the type evolution of the N small players are straightforward.

It is straightforward to derive the following extension of Theorem 2.

Theorem 7. *Fix* $\varepsilon > 0$. *Assume that:*

- σ_{ID}^{DCE} *is uniformly continuous*
- $\{\sigma_m : \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A}) \to \mathcal{A}_L\}_{m=1}^{M}$ *are uniformly equicontinuous*

Then we can choose N^**such that the DCE* $(\sigma_{ID}^{DCE}, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ *is an* ε*-Bayesian-Nash Equilibrium for the* N *small players of the large finite stochastic game for* $N > N^*$.

Since our theorem holds for any fixed vector $\overline{\sigma}$, it continues to hold for a vector $\overline{\sigma}$ that satisfies a game-theoretic equilibrium definition for the strategies of the large players.

4.4. **Perfection.** The challenge of using a refinement of BNE in Theorem 3 is that we must show that for any convergent sequence of refined equilibrium strategies, $\sigma^N \to \sigma^\infty$, the refinement holds at σ^∞. For games of complete information, strengthening our equilibrium notion to perfection is straightforward using the following notions of a perfect equilibrium in the large finite and nonatomic limit games.

Definition 5. *A symmetric* $\varepsilon-$ *Perfect Equilibrium* ($\varepsilon-$*PE*) *is a strategy and state* $\sigma^{PE} \in \Sigma$ *of the* N*-agent game such that for all agents* $i \in \{1, ..., N\}$ *we have for all* $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}) \in \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$ *and* $\theta_t^i \in supp[\pi_t^{\Theta}]$

$$V_N(\theta_t^i, \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}} | \sigma^{PE}) + \varepsilon \geq \sup_{\sigma_i' \in \Sigma} V_N(\theta_t^i, \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}} | \sigma_i', \sigma_{-i}^{PE})$$

Definition 6. *A symmetric* $\varepsilon-$*Perfect Competitive Equilibrium* (ε*-PCE*) *consists of a strategy* $\sigma^{PCE} \in \Sigma$ *such that for all* $(\theta_t^i, \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}) \in \Theta \times \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$

$$V_\infty(\theta_t^i, \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}} | \sigma^{PCE}) + \varepsilon \geq \sup_{\sigma_i' \in \Sigma} V_\infty(\theta_t^i, \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}} | \sigma_i', \sigma_{-i}^{PCE})$$

We now generalize both Theorem 2 and Theorem 3 to the case where perfection is required. For both theorems we employ the indirect description of the PCE strategy, which we denote σ_{ID}^{PCE}.

Theorem 8. *Fix $\varepsilon > 0$. Assume that σ_{ID}^{PCE} is uniformly continuous*[46] *and generates tight outcomes. Then we can choose N^* such that the PCE strategy σ^{PCE} is an ε-Perfect Equilibrium of the large finite stochastic game for $N > N^*$ players.*

Proof. A PCE can be represented as a collection of DCE as follows:[47]

$$\left\{(\sigma^{PCE}, \omega, \pi^{\Theta}, \pi^{\mathcal{A}})\right\}_{(\omega,\pi^{\Theta},\pi^{\mathcal{A}})\in\Omega\times\Delta_N(\Theta)\times\Delta_N(\mathcal{A})}$$

For each element of these collection, we consider the indirect description and note that Theorem 2 applies uniformly over this collection, which yields Theorem 8. □

Denote the set of exact perfect equilibrium strategies for the N agent game as $\mathcal{E}^P(N)$ and the set of perfect competitive equilibrium strategies as $\mathcal{E}^{PCE}$.

Theorem 9. *Assume that*

- *Θ and Ω are compact*
- *There exists an N^* such that $\cup_{N=N^*}^{\infty}\mathcal{E}^P(N)$ is uniformly equicontinuous in $\Theta \times \Omega \times \Delta(\Theta)$*

Then the correspondence $\mathcal{E}^P$ is upper hemicontinuous with

$$\underset{N\to\infty}{Lim}\,\mathcal{E}^P(N) \subset \mathcal{E}^{PCE}$$

Proof. Since a PE is a type of BNE, Theorem 3 implies that

$$\underset{N\to\infty}{Lim}\,\mathcal{E}^P(N) \subset \mathcal{E}^{DCE}$$

where $\mathcal{E}^{PCE} \subset \mathcal{E}^{DCE}$. The proof of Theorem 3 argues that if we consider a convergent sequence of PE strategies, $\{\sigma^N\}_{N=N^*}^{\infty}$[48], then the limit strategy, σ^{∞}, is continuous. If there is a perfect competitive equilibrium that is a 2ε-improvement over σ^{∞} at some information set, then we can use Theorem 2 of Khurana [**?**] to generate a uniformly continuous approximation, denoted σ^C, of the profitable deviation that is an $\varepsilon-$improvement over σ^{∞} at the same information set. Paraphrasing our techniques for proving Theorem 3, σ^C then yields an $\frac{\varepsilon}{2}$ improvement over σ^N (for large enough

[46]The minimal requirement is continuity over Ω^t for sufficiently large t.

[47]Note that the equilibrium strategy is identical for each element of the collection.

[48]More formally if we consider a sequence of uniformly continuous extensions of continuous PE strategies to the space $\Theta \times \Omega \times \Delta(\Theta)$.

N) at this information set. But then this yields a contradiction of our assumption that $\sigma^N \in \mathcal{E}^P(N)$ and we are done. □

For games of imperfect information, we conjecture that our results would continue to hold if we insisted that beliefs off the path survive an equilibrium refinement (e.g. the consistency criterion for sequential equilibria). In our formulation these refinements define the admissible conditional probabilities following null events. Define *symmetric consistency* as consistency with respect to limits of symmetric, completely mixed strategies.[49] We conjecture that one can redefine our equilibrium notions to require *symmetric consistency* of off-path beliefs and attain upper hemicontinuity of equilibrium strategy correspondences under the refined equilibrium definition.

4.5. **Coalitions.** Given an equilibrium strategy σ for the N agent game, denote a deviation by a coalition $\mathcal{I} \subset \{1, ..., N\}$ to a strategy vector $\sigma' \in \Sigma^{|\mathcal{I}|}$ by $(\sigma'_{\mathcal{I}}, \sigma_{-\mathcal{I}})$.[50] Consider the following definition of K-coalition proof equilibria in large finite games.

Definition 7. *A symmetric K-Coalition Proof $\varepsilon-$Bayesian-Nash Equilibrium is a strategy and state $(\sigma^{KCP}, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}}) \in \Sigma \times \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$ of the N-agent game such that for any coalition of agents, $\mathcal{I} = \{i_1, .., i_M\} \subset \{1, ..., N\}$ where $|\mathcal{I}| = M \leq K$ and for all $i \in \mathcal{I}, j \in \{1, ..., N\} \backslash \mathcal{I}$*

$$\begin{aligned} V_N(\theta_0^i, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}} | \sigma^{KCP}) + \varepsilon &\geq \sup_{\sigma'_{\mathcal{I}} \in \Sigma^M} V_N(\theta_0^i, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}} | \sigma'_{\mathcal{I}}, \sigma_{-\mathcal{I}}^{KCP}) \\ V_N(\theta_0^j, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}} | \sigma^{KCP}) + \varepsilon &\geq \sup_{\sigma'_j \in \Sigma} V_N(\theta_0^j, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}} | \sigma'_j, \sigma_{-j}^{KCP}) \end{aligned}$$

Let $\widehat{P}_N^A$ denote the stochastic process generated by T and G in conjunction with $(\sigma'_{\mathcal{I}}, \sigma_{-\mathcal{I}}^{KCP})$. We require the following extension of Corollary 6.

Corollary 6. *Fix $\tau^* < \infty, \gamma > 0$ and $\rho \in [0, 1)$. Assume $\sigma \in \Sigma$ is uniformly continuous. Suppose a fixed coalition $\mathcal{I}$, $|\mathcal{I}| = M < N$, deviates from $\sigma \in \Sigma$ to $(\sigma'_1, .., \sigma'_M) \in \Sigma^M$. Then there exists $N^*, \overline{\gamma} > 0$ such that for any $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), (\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}})$ where*

[49]The degree of generality with which a consistency refinement can be formally stated in our nonatomic framework is unclear and remains an interesting avenue for future research.

[50]We do not require the members of the coalition deviate to the same strategy or that $\sigma'_{\mathcal{I}}$ be continuous.

$d_\Omega(\omega_t, \widetilde{\omega}_t) + d_{LP}^{\Theta}(\pi_t^{\Theta}, \widetilde{\pi}_t^{\Theta}) + d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}}, \widetilde{\pi}_t^{\mathcal{A}}) < \overline{\gamma}$, any $N > N^*$, and all $\tau \in \{1, .., \tau^*\}$ we have that

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((\widehat{P}_N^{\mathcal{A}})^{\tau}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^{\tau}(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}})) < \gamma$$

with probability at least $1 - \rho$.

Proof. By treating the deviation by the M agents as a perturbation of the evolution of π_t^{Θ} and $\pi_t^{\mathcal{A}}$ as in the proof of Lemma 8 we get the following result from an essentially identical argument

Lemma 2. *(Modified Lemma 8) Assume all of the agents play symmetric, continuous strategies σ save a coalition $\mathcal{I}$, $|\mathcal{I}| = M < N$. Assume play commences at $(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$. For any $\gamma, \rho > 0$ we have uniformly over $\omega^t \in \Omega^t$ that for large enough N with probability at least $1 - \rho$ that conditional on ω^t*

$$\begin{aligned} d_{LP}^{\Theta}(\pi_t^{\Theta}, \pi_t^{\Theta}(\omega^t)) &< \delta \\ d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}}, \pi_t^{\mathcal{A}}(\omega^t)) &< \delta \end{aligned}$$

The proof of our Corollary then follows an (almost) verbatim use of the proof of Theorem 1. □

It is straightforward to adapt our approximation theorems to show that any DCE is a K-Coalition Proof $\varepsilon-$Bayesian-Nash Equilibrium for sufficiently large N using Corollary 6. Since finite coalitions cannot collude to significantly alter the evolution of large dynamic games and the DCE strategies are best responses to the evolution of $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ given σ^{DCE}, the DCE are approximate best responses to the evolution following a deviation by a K agent coalition.[51]

Theorem 10. *Fix $\varepsilon > 0$ and $K < \infty$. Assume that σ^{DCE} is uniformly continuous. Then we can choose N^*such that the DCE $(\sigma^{DCE}, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ is a K-Coalition Proof ε-Bayesian-Nash Equilibrium of the large finite stochastic game for $N > N^*$.*

Proof. Corollary 6 implies that the stochastic process following any K agent deviation from σ^{DCE}, denoted $\widehat{P}_N^{\mathcal{A}}$, obeys $\widehat{P}_N^{\mathcal{A}} \to P_\infty^{\mathcal{A}}$. The argument of Theorem 2 applies

[51]The extension of Theorem 3 to coalition Proof equilibrium is trivial. Since coalition proof equilibria are a subset of the Nash equilibria and we have already shown that $\mathcal{E}(N)$ is upper hemicontinuous as $N \to \infty$, any convergent sequence of K-Coalition Proof 0-Bayesian-Nash Equilibria converges to a DCE.

directly once we replace references to the stochastic process following a single agent's deviation, $\widetilde{P}_N^A$, with the stochastic process $\widehat{P}_N^A$. □

Our theorem provides a formal statement of the intuitive idea that as markets grow, the power of coalitions to influence the utility of participants and the evolution of market aggregates fades. Formally, this theorem says that for any coalition of maximal size K, as $N \to \infty$ the benefit accrued by firms in any such coalition shrinks to 0. If there are any significant organizational or communication costs for running a coalition, then these coalitions will disappear for sufficiently large N.

4.6. **Entry and Exit Games, Discontinuities and Refined Topologies .** The approximation results of our work are founded on continuity assumptions placed on the model, and in some instances these continuity assumptions may be unappealing. To focus ideas, let us consider the model of industry competition with investment, entry and exit studied by Benkard et al. [7].[52] In this setting the entry and exit actions imply discontinuities in agent strategies. We focus our discussion on applying Theorem 4 in this setting as this is the natural extension of Theorem 5.4 of Benkard et al.

In the model of Benkard et al., continuity of a stationary equilibrium strategy follows from the discreteness of the firm i's state at time t, $x_{it} \in \mathbb{Z}$. Continuity and boundedness of w is assumed in Assumption 3.1. *Uniform* continuity of w is assured by the bounded derivatives of w in Assumption 5.1. Continuity of T and Assumption 2 are obtained in assumption 3.2. Benkard et al. [7] assume $\mathcal{A}$ is bounded, but if this is strengthened to compactness we have uniform continuity of T by the Heine-Borel theorem. There are no aggregate shocks, so G is uniformly continuous.

Entry is regulated by a common entry cost, where firms enter if the cost of entry is lower than the expected net present value of profits upon entering. We define the actions "Enter" and "Stay Out" of potential entrants as discrete elements of $\mathcal{A}$. Since the firms are symmetric prior to entry, we interpret the entry decision as a mixed strategy. In this mixed strategy, the firms mix between "Enter" and "Stay

[52]Convex investment costs as in Caballero and Engel [13] provide another important context with natural discontinuities.

Out" with a probability that is conditional on the state of the economy, π_t^Θ.[53] Since the expected profits upon entry adjust continuously with π_t^Θ, the mixture probability also adjusts continuously.

The exit decision of a firm is determined by whether the net present value of future profits is greater or lower than a scrap value, $\phi_{it} \in \mathbb{R}_+$, which is specific to firm i in period t and independently and identically drawn each period. We include the scrap value in the agent type space, so $(x_{it}, \phi_{it}) \in \Theta = \mathbb{Z}\times\mathbb{R}_+$. The choice to exit is modeled as a discrete action, "Exit," in the action space, and the exit strategy of the firms takes the form of a cutoff strategy. Since the scrap values are continuous random variables, the exit cutoff strategy is a map from a continuous set to a discrete space and is necessarily discontinuous. A crude technique to remove this discontinuity is to discretize the support of the scrap value and hold this support fixed as $N \to \infty$. In this case the exit decision is a uniformly continuous map between two discrete spaces. In this case the assumptions of Theorem 4 hold and the oblivious equilibria Benkard et al. study are $\varepsilon-$Bayesian Nash equilibria.

Alternatively we can modify our asymptotic approximation results to allow for exit strategies that have discontinuities that are well behaved. Intuitively, if we could assure that the equilibrium exit strategy was discontinuous for only a measure 0 of agents in the nonatomic limit game, then we would be able to establish that the equilibrium dynamics of the nonatomic limit game are continuous. This continuity implies that Theorems 1 continues to hold, which is the key result underpinning Theorem 2 that proves DCE are $\varepsilon-$Bayesian Nash equilibria.

In the nonatomic limit game, the equilibrium exit strategies have the form

$$\sigma^{Exit}(x_{it}, \phi_{it}) = \text{"}Exit\text{"}$$

if and only if

$$V_\infty(x_{it}, \omega, \pi_\infty^\Theta, \pi_\infty^\mathcal{A} | \sigma^{SE}) \leq \phi_{it}$$

But note that a discontinuity exists only for an agent of type (x_{it}, ϕ_{it}) such that this inequality holds exactly. If we assume the distribution of ϕ_{it} is nonatomic, then

[53]The potential entrants are indifferent between "Enter" and "Stay Out" since the equilibrium rate of entry is determined by firm indifference between staying out and earning nothing and entering by paying the fixed entry cost that is equal to by the net present value of future profits.

for each x_{it} the discontinuity is nongeneric. Denote $\pi^{\Theta} \in \Delta(\Theta)$ as admissible for a strategy $\sigma \in \Sigma$ if there is an admissible distribution $\widetilde{\pi}^{\Theta} \in \Delta(\Theta)$ such that for any $U \in \mathcal{B}(\Theta)$[54]

$$\pi^{\Theta}(U) = \int_{\Theta} T(U|a,\theta) * \sigma(\theta, \widetilde{\pi}^{\Theta})[da] * \widetilde{\pi}^{\Theta}[ds]$$

Assumption 7. *(Condition A) Fix $\varepsilon > 0$. For each π^{Θ} admissible under strategy σ_{ID}^{DCE} and for any $\omega^t \in \Omega^t$, $\sigma(\circ, \omega^t)$ is uniformly continuous for all but a measure ε of θ with respect to π^{Θ}.*

We only consider pairs $(\pi^{\Theta}, \pi^{\mathcal{A}})$ such that the strategy σ_{ID}^{DCE} in conjunction with π^{Θ} generate $\pi^{\mathcal{A}}$. Written formally, for all $V \in \mathcal{B}(\mathcal{A})$ we have

$$\pi^{\mathcal{A}}[V] = \int_{\Theta} \sigma_{ID}^{DCE}(\theta, \omega^t)[V] * \pi^{\Theta}[d\theta]$$

Denote this dynamic operator, in analogy with the P^C defined in section 3, as

$$P^{ID} : \Delta(\Theta) \times \Omega^{\infty} \to \Delta(\Theta)$$

which defines the evolution of $\pi^{\mathcal{A}} \in \Delta(\mathcal{A})$ implicitly by

$$\begin{aligned} \pi_{t+1}^{\Theta} &= P^{ID}(\pi_t^{\Theta}, \omega^t) \\ \pi_{t+1}^{\mathcal{A}}[\circ] &= \int_{\Theta} \sigma_{ID}^{DCE}(\theta, \omega^t)[\circ] * \pi_{t+1}^{\Theta}[d\theta] \end{aligned}$$

We then have the following continuity theorem for the nonatomic dynamic limit game. It is straightforward to extend our result to uniform continuity of $(P^{ID})^{\tau}$ for any $\tau \in \mathbb{N}$.

Lemma 3. *Assume T is uniformly continuous, hold σ_{ID}^{DCE} fixed, and assume Condition A holds. For any $\varepsilon > 0$, we can find a $\gamma > 0$ such that $d_{LP}^{\Theta}(\pi^{\Theta}, \widetilde{\pi}^{\Theta}) < \gamma$ implies $d_{LP}^{\Theta}(P^{ID}(\circ|\pi^{\Theta}, \omega^t), P^{ID}(\circ|\widetilde{\pi}^{\Theta}, \omega^t)) < \varepsilon$ uniformly over $\pi^{\Theta} \in \Delta(\Theta)$.*

Our continuity results are required to establish that the equilibria of the nonatomic limit game admit uniformly continuous value functions. Lemma 3 proves that well behaved discontinuities of the equilibrium strategy, strategies that satisfy Condition A, do not cause the evolution of the nonatomic equilibrium (and hence the value

[54]Admissiblity is a requirement placed on π^{Θ} realized in the nonatomic limit game.

function for the equilibrium) to be discontinuous. Using this result, we can replace the use of Lemma 4 with Lemma 3 in our proofs to generate analogs to Theorem 1, Corollary 2, and Theorem 2, which allows us to state the following approximation theorem

Theorem 11. *Fix* $\varepsilon > 0$. *Assume that:*

- σ_{ID}^{DCE} *obeys condition A*
- σ_{ID}^{DCE} *is uniformly continuous*

Then we can choose N^**such that the DCE* $(\sigma_{ID}^{DCE}, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ *is an* ε*-Bayesian-Nash Equilibrium of the large finite stochastic game for* $N > N^*$.

In the equilibria considered by Benkard et al. [7], the exit strategy clearly obeys Condition A since only a measure 0 of the agents are indifferent between choosing "Exit" and remaining in the industry. Theorem 11 implies that oblivious equilibria are ε-Bayesian-Nash equilibria. By representing a Perfect Competitive Equilibrium (see Section 4.4) as a collection of DCE $\{(\sigma_{ID}^{DCE}, \omega, \pi^{\Theta}, \pi^{\mathcal{A}})\}_{(\omega, \pi^{\Theta}, \pi^{\mathcal{A}}) \in \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})}$, we can use Theorem 8 to show that oblivious equilibria are $\varepsilon-$Markov Perfect Equilibria. We can also incorporate aggregate shocks so long as the space of aggregate shocks is discrete or the DCE strategies are continuous in the aggregate shocks. We can include the presence of large agents such as monopolists, regulators, or other government bodies (see Section 4.3) and assert that the DCE are proof against finite coalitions (see Section 4.5) without loss of generality.

4.7. **Asymmetric Roles and Alternating Actions.** The agents are ex ante identical in our framework. However, our model can be structured to accommodate different agent roles (e.g. buyers and sellers) without breaking the formal symmetry of our structure. Suppose at $t = 0$ the agent types are independently, but not identically, drawn from a finite set of distributions $\{\pi_i^{\Theta}(\omega)\}_{i=1}^m$ over spaces $\{\Theta_i\}_{i=1}^m$ where Θ_i denotes the types for agents in role i. The full type space is defined as $\Theta = \prod_{i=1}^m \Theta_i$. Denote the number of agents drawn from the i^{th} space as N_i and assume $\frac{N_i}{N} \to \beta_i > 0$ as $N \to \infty$. This is, from the perspective of the agents, a process where the agent characteristics are drawn in two stages: (1) randomly assign the agent's role $i \in \{1, ..., m\}$ and (2) draw the characteristics from the distribution associated with role i. Note that step (1) could be done deterministically. For example, the analyst may be concerned with models wherein the number of buyers and sellers are equal. In this case, buyers and sellers are defined using distinct type distributions and enter the economy in (buyer, seller) pairs with types for each agent in the pair determined independently. Our approximation theorems continue to hold in this setting by applying them to each role's symmetric equilibrium strategy holding the the strategies of the other agents' fixed.

Our structure also assumes players act simultaneously in each period. Our approximation techniques only require that the fraction of players choosing actions in any given period diverges to a positive fraction as $N \to \infty$, which allows us to represent the players acting in each period as a continuum of positive measure. If subsets of the players act in each period of the large finite game (e.g. sellers offer prices in periods $t \in \{0, 2, 4, ...\}$ and buyers search for products in $t \in \{1, 3, 5, ...\}$), we can include this in our model by letting $\omega_t \in \Omega$ encode the agents who act in period t. For example, $\omega_t = "S"$ implies the sellers take actions in period t, whereas $\omega_t = "B"$ implies the buyers take actions. This is formalized by making agents indifferent to the actions of buyers (sellers) when $\omega_t = "S"$ ($\omega_t = "B"$) and defining any distribution of buyer (seller) actions when $\omega_t = "S"$ ($\omega_t = "B"$) as a single information set.[55] This notational artifice allows us to preserve the formal symmetry of the model while allowing agents to act asynchronously.

5. Previous Literature

Our work emphasizes the use of nonatomic limit games as a technique for approximating the full set of exact equilibria of large finite games. This is not the first work to outline conditions under which game-theoretic equilibria approach nonatomic equilibria as the number of agents increases. Green [20], and later Sabourian [30], provide sufficient conditions for the set of nonatomic subgame perfect equilibria (SPE) and the SPE of the analogous large finite repeated game to converge. Our contribution to this vein of literature is two-fold. First, Green [20] and Sabourian [30] study repeated games without private information. We analyze a broad class of games with

[55]This is a restriction on the measurability of the strategy space. We do not let the strategies distinguish histories where buyers (sellers) take different distributions of actions in periods where $\omega_t = "S"$ ($\omega_t = "B"$).

a particular focus on games of interest to applied and empirical practitioners including games of persistent private information, games with imperfect public or private monitoring, and complete information stochastic games. Second, our conditions are limited wherever possible to the economic primitives rather than endogenous quantities, and our theorems rely on assumptions on the model primitives sufficient to establish convergence conditions assumed in earlier works.

A growing literature in econometrics and applied theory develops novel equilibrium concepts that emphasize computational tractability. These works exactly characterize a notion of approximate equilibrium, and a significant challenge is arguing that the approximate equilibrium notions have empirical relevance. Benkard, Van Roy, and Weintraub ([6], [7],[8]) provide an analysis of the stationary oblivious equilibria of a model of industry competition formulated by Ericson and Pakes [17]. Benkard et al. prove that these equilibria asymptotically satisfy the exact equilibrium conditions for a Markov Perfect Equilibrium and provide techniques for bounding the error of the approximation.

Adlakha, Johari and Weintraub [1] extend Bekard et al. [7] by proving existence and asymptotic approximation results for stationary oblivious equilibria in a broad class of complete information stochastic market games. In particular, Adlakha et al. provide sufficient conditions on economic primitives to insure that the *light-tailed property* introduced by Benkard et al. ([7]) holds and find relations between the light-tailed property and dynamics that exhibit decreasing returns to scale. Finally, Adlakha et al. apply their results to a number of dynamic market games to analyze the structure of large industries in stationary equilibria.

Our work is complementary to the research on approximate equilibria in three ways. First, our results apply to a broader class of games including those with public and private information structures and type spaces that are arbitrary subsets of finite dimensional Euclidean spaces. Second, it is well known that approximate equilibria need not resemble any exact equilibria, even when the equilibria become exact asymptotically, except under suitable continuity restrictions. One goal of our work is to fill this gap and provide stronger continuity conditions on the model primitives sufficient for the approximate equilibrium strategies to be close to exact equilibrium strategies. Third, even in the case where the approximate equilibria do faithfully approximate some exact equilibrium, the question remains as to whether there are other equilibria that have been ignored. Our theorems provide an approximate characterization of the equilibrium correspondence, which allows us to circumscribe the limits of equilibrium behavior.[56]

A related literature on large games studies whether individual agents remain pivotal as the total number of agents grows (Fudenberg, Levine and Pesendorfer [19]; Al-Najjar and Smordinsky [4]; Al-Najjar [3], [2]). The theme of this literature is that with minimal conditions, a vanishing fraction (although a potentially infinite number) of players remains pivotal as $N \to \infty$. Fudenberg et al. use this feature to prove that equilibria of a limit game are approximate equilibria of the large finite game, although their notion of approximation is weaker than ours and cannot in general be used to prove upper hemicontinuity of the equilibrium correspondence in the sup-norm.[57]

Our work also relates to the work on convergence (or lack thereof) of markets to perfect competition as sources of friction disappear (Wolinsky [34], Jovanovic and Rosenthal [23], Hopenhayn [21], Serrano [32], Satterthwaite and Shneyerov [31]). These analyses are tractable because the individual agents assume economic aggregates are stationary and exogenous. Our results provide a micro-foundation for these behavioral assumptions in the context of continuous games.

Kubler and Schmedders [24] provide an analysis of the impact of errors in equilibrium calculations in a general equilibrium model. The basic question of their study, the relationship between the computed (and hence approximate) equilibria and exact equilibria, is similar to ours, although the source of the approximation is entirely different.

[56] Whether or not we are able to rule out a significant range of equilibrium strategies will depend on the application. For example, games with equilibria in weakly dominated strategies (e.g. voting games) have a wide array of equilibrium outcomes. Since we are using a limit approximation, in cases such as this it may be unclear whether a failure to tightly circumscribe equilibrium behavior in the large finite game is due to a rich set of exact equilibrium outcomes or a failure of lower hemicontinuity in the limit as $N \to \infty$.

[57] We conjecture, but have not proven, that pivotality results can be used to show upper hemicontinuity of the equilibrium correspondence under the L^0 norm under weaker continuity conditions than we require. However it is not apparent that equilibrium outcomes will remain close even if the strategies converge in the L^0 norm without assuming continuity conditions akin to ours.

Macroeconomics papers use models and equilibrium concepts that are close kin to techniques used in the study of large games. Chari and Kehoe [14] provide an early example that relates the Ramsey allocations of an infinite horizon taxation and investment problem with a continuum of consumers to the perfect Bayesian equilibria of an analogous nonatomic game. In the majority of these studies, the use of a nonatomic model is an untested approximation of an underlying large finite game. Although it is beyond the scope of this paper to investigate particular macroeconomic models, Theorems 2, 3 and 6 provide conditions under which the use of a nonatomic approximation of small players (such as consumers or competitive firms) is valid in a macroeconomic setting. If the continuity properties are satisfied, then the model is a close approximation of the underlying behavior of the agents in the economy. In the event that the continuity properties are not satisfied, care ought to be taken when interpreting the results of the analysis.

6. Conclusion

Our paper illustrates sufficient conditions under which the equilibria of a large finite dynamic game can be approximated by the equilibria of a continuous nonatomic dynamic game. Our principal goal is to provide theoretical tools for applied theorists and empirical researchers interested in studying large economies. Prior research has focused on finding a particular equilibrium of the nonatomic limit game that is an approximate equilibrium in the large finite game (Adlakha et al. [1], Benkard et al. [7]), whereas our approach is to approximate the exact equilibrium correspondence in the limit as $N \to \infty$. Our work applies to a broad array of game theoretic structures and allows for perfect or imperfect monitoring; persistent private information; and continuous state, action and type spaces.

Our approximation theorems provide a framework for simplifying the estimation of structural model based on large finite games by employing tractable nonatomic games as limit approximations. Since agents in the nonatomic limit game take the equilibrium path of market aggregates as given, the nonatomic model's equilibria can be described as a function of aggregate states and an individual agent's type. Therefore the estimated policy functions do not suffer from the curse of dimensionality. We also provide conditions under which the equilibria of the nonatomic limit game approximately bounds the set of equilibrium policy functions of the large finite game, which implies that our nonatomic limit game provides a complete description of equilibrium outcomes in the nonatomic limit game. Finally, our framework is general enough to include stochastic games, games of complete or incomplete information, public or private monitoring, and games with persistent private information.

We use our results to provide a general anti-folk theorem for repeated games in the limit as $N \to \infty$. This anti-folk theorem complements prior results in the literature such as Green [20], Sabourian [30], and Fudenberg, Levine and Pesendorfer [19] amongst others. Our results provide conditions under which game theoretic predictions that turn on pivotality of the agents are robust as the economy grows large. In addition, our sufficient conditions reveal when models that presume a nonatomic continuum of agents, such as dynamic stochastic general equilibrium models used by macroeconomists, can ignore game theoretic interactions without a loss of generality.

There are two paths forward for this research agenda. From a theoretical perspective, it would be interesting to study when we can approximate refinements of the Bayesian-Nash equilibrium correspondence. In games of complete information, it is straightforward to extend our results to the subgame and Markov perfect equilibrium correspondence (Section 4.4). Extending our results to refined equilibria of games with incomplete information requires formally defining refinements such as consistency of beliefs in the nonatomic setting. Chari and Kehoe [14] provide one example of such an extension by placing strong symmetry restrictions on the off-path beliefs of the agents, and we conjecture that such an approach can be employed in our setting to approximate the sequential or perfect Bayesian equilibrium correspondence.

The other route for advancing this research agenda is to apply our results to large games of interest. For example, it would be useful to apply our analysis framework to describe and estimate dynamic auctions or matching models. Furthermore, we can apply our model of large games with a few atomic players (Section 4.3) to analyze structural models with a small number of dominant firms and a competitive fringe of small firms. Examining such hybrid models, both theoretically and empirically, remains an interesting topic for future work.

References

[1] Adlakha, S.; R. Johari; and G. Weintraub. (2010) "Equilibria of Dynamic Games with Many Players: Existence, Approximation, and Market Structure," *mimeo.*

[2] Al-Najjar, N. (2004) "Aggregation and the Law of Large Numbers in Large Economies," *Games and Economic Behavior*, 47, pp. 1-35.

[3] Al-Najjar, N. (2008) "Large Games and the Law of Large Numbers," *Games and Economic Behavior*, 64, pp. 1-34.

[4] Al-Najjar, N. and R. Smordinsky (2000) "Pivotal Players and the Characterization of Influence," *Journal of Economic Theory*, 92, pp. 318-342.

[5] Aliprantis, C. and K. Border (2006) *Infinite Dimensional Analysis: A Hitchhiker's Guide, Third edition*, Springer-Verlag Inc.: New York.

[6] Benkard, C.L.; B. Van Roy; and G. Weintraub (2009) "Industry Dynamics: Foundations for Models with an Infinite Number of Firms," *mimeo.*

[7] Benkard, C.L.; B. Van Roy; and G. Weintraub (2009) "Markov Perfect Industry Dynamics With Many Firms," *Econometrica*, 76, pp. 1375 - 1411.

[8] Benkard, C.L.; B. Van Roy; and G. Weintraub (2009) "Computational Methods for Oblivious Equilibrium," *mimeo.*

[9] Berge, C. (1959) *Topological Spaces* (English Translation by E. Patterson, 1963) Oliver and Boyd: Edinburgh and London.

[10] Bergin, J. and D. Bernhardt (1995) "Anonymous Sequential Games: Existence and Characterization of Equilibrium," *Economic Theory*, 5, pp 461 - 489.

[11] Billingsley, P. (1968) *Convergence of Probability Measures*, John Wiley & Sons, Inc.: New York.

[12] Bodoh-Creed, A. (2010) "Approximation of Large Games," *mimeo.*

[13] Caballero, R. and E. Engel (1999) "Explaining Investment Dynamics in U.S. Manufacturing: A Generalizes (S,s) Approach," *Econometrica*, 67, pp. 783 - 826.

[14] Chari, V. and P. Kehoe (1990) "Sustainable Plans," *Journal of Political Economy*, 98, pp. 783 - 802.

[15] Cripps, M. and J. Swinkels (2006) "Efficiency of Large Double Auctions," *Econometrica*, 74, pp. 47-92.

[16] Dunford, N. and J. Schwartz (1967) *Linear Operators: Part I*, Interscience Publishers, Inc.: New York.

[17] Ericson, R. and A. Pakes (1995) "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62, pp. 53-82.

[18] Fudenberg, D.; D. Levine; and E. Maskin (1994) "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, pp. 997 - 1039.

[19] Fudenberg, D.; D. Levine; and W. Pesendorfer (1998) "When are Nonanonymous Players Negligible?" *Journal of Economic Theory*, 79, pp. 46-71.

[20] Green, E. (1980) "Noncooperative Price Taking in Large Dynamic Markets," *Journal of Economic Theory*, 22, pp. 155-181.

[21] Hopenhayn, H. (1992) "Entry, Exit and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60, pp. 1127 - 1150.

[22] Isbell, J. (1961) "Uniform Neighborhood Retracts," *Pacific Journal of Mathematics*, 11, pp. 609-648.

[23] Jovanovic, B. and R. Rosenthal (1988) "Anonymous Sequential Games," *Journal of Mathematical Economics*, 17, pp. 77 - 87.

[24] Kubler, F. and K. Schmedders (2005) "Approximate Versus Exact Equilibria in Dynamic Economies," *Econometrica*, 73, pp. 1205-1235.

[25] Levine, D. and W. Pesendorfer (1995) "When are Agents Negligible," *American Economic Review*, 85, pp. 1160-1170.

[26] Lindenstrauss, J. (1964) "On nonlinear projections in Banach spaces," *Michigan Mathematics Journal*, 11, pp. 263 - 287.

[27] McLean, R.; I. Obara; and A. Postlewaite (2010) "Informational Smallness and Private Monitoring," *mimeo.*

[28] Pollard, D. (1984) *Convergence of Stochastic Processes*, Springer-Verlag: New York.

[29] Roberts, D. and A. Postlewaite (1976) "The Incentives for Price-Taking Behavior in Large Exchange Economies," *Econometrica*, 44, pp. 115 - 127.

[30] Sabourian, H. (1990) "Anonymous Repeated Games with a Large Number of Players and Random Outcomes," *Econometrica*, 61, pp. 325-351.

[31] Satterthwaite, M. and A. Shneyerov (2007) "Dynamic Matching, Two-Sided Incomplete Information, and Participation Costs: Existence and Convergence to Perfect Competition," *Econometrica,* 75, pp. 155 - 200.

[32] Serrano, R. (2002) "Decentralized information and the Walrasian outcome: a pairwise meetings market with private values," *Journal of Mathematical Economics*, 38, pp. 65-89.

[33] van der Vaart, A. and J. Wellner (1996) *Weak Convergence and Empirical Processes*, Springer-Verlag: New York.

[34] Wolinsky, A. (1988) "Dynamic Markets with Competitive Bidding." *Review of Economic Studies*, 55, pp. 71-84.

Appendix A. Proofs

We begin this section with some useful results from the theory of the weak convergence of empirical processes. Let $d_K : \Delta(\mathcal{S}) \to \Delta(\mathcal{S})$ denote the Kolmogorov metric and reserve the notation $d_K^{\mathcal{S}}$ for when clarity requires us to denote the spaces to which the metric applies.

Theorem 12. *Consider a random variable $X : \Omega \to \mathbb{R}^d$, $d < \infty$, with measure π_0 and associated CDF $F(y) = \int_\Omega 1\{x \leq y\} * \pi_0(dx)$. For N i.i.d. realizations, $\{X_1, ..., X_N\}$ drawn from π_0, denote the N realization empirical CDF as $F_N(y)$. Then we have*

$$d_K(\pi_N, \pi_0) = \sup_{y \in \mathbb{R}^d} |F_N(y) - F(y)| \to 0 \text{ almost surely as } N \to \infty$$

Proof. (Proof of Theorem 12) Follows from van der Vaart et al. [33], p. 135 and noting that the sets of the form $\{x : x \leq y\}$ for $y \in \mathbb{R}^d$ are lower contours and form a VC Class. □

Corollary 7. *Define the empirical measure generated by the counting measure over $\{X_1, ..., X_N\}$ as π_N. Then $\pi_N \to \pi_0$ almost surely in the weak-* topology over $\Delta(\mathbb{R}^d)$*

Proof. (Proof of Corollary 7) From Billingsley (p. 18, [11]) we have that $F_N(y) \to F(y)$ at continuity points of F implies $\pi_N \to \pi_0$ almost surely in the weak-* topology. Since we have uniform convergence $F_N(y) \to F(y)$ for all y almost surely, we have $\pi_N \to \pi_0$ in the weak-* topology. □

Corollary 8. *Consider a random variable $X : \Omega \to \mathbb{R}^d$, $d < \infty$, and associated CDF $F(y)$. Denote the N realization empirical CDF as $F_N(y)$. Then*

$$\Pr\{\sqrt{N} \sup_{y \in \mathbb{R}^d} |F_N(y) - F(y)| > t\} \leq C * e^{-2t^2}$$

where the constant $C > 0$ depends only on the dimension d.

Proof. (Proof of Corollary 8) This result follows directly from Theorems 2.6.7 and 2.14.9 of van der Vaart and Wellner [33]. □

A.1. **Proofs from Section 3.** We begin this section by proving that the nonatomic model transition operator

$$P^C(\omega_{t+1}, \pi_t^\Theta, \pi_t^{\mathcal{A}} | \sigma^{DCE}) = (\pi_{t+1}^\Theta, \pi_{t+1}^{\mathcal{A}})$$

and the $\tau < \infty$ step ahead iterations of this operator are continuous.

Lemma 4. *If T,σ are (uniformly) continuous, then $(P^C)^\tau$ is (uniformly) continuous in $\Delta(\Theta)\times\Delta(\mathcal{A})\times\Sigma$ for any $\tau<\infty$ conditional on any sequence of aggregate states $(\omega_{t+1},...,\omega_{t+\tau})$*

Proof. It suffices to show that $P^C(\omega_{t+1},\pi_t^\Theta,\pi_t^{\mathcal{A}}|\sigma)$ is continuous to imply that for any $\tau<\infty$ we have $(P^C)^\tau(\omega_{t+1},\pi_t^\Theta,\pi_t^{\mathcal{A}}|\sigma)$ is continuous for any sequence of aggregate states $(\omega_{t+1},...,\omega_{t+\tau})$.

Fix $(\omega_{t+1},\pi_t^\Theta,\pi_t^{\mathcal{A}},\sigma)\in\Omega\times\Delta(\Theta)\times\Delta(\mathcal{A})\times\Sigma$. We want to show that for any $\varepsilon>0$ we can choose $\gamma>0$ so that if $d_{LP}(\pi_t^\Theta,\widetilde{\pi}_t^\Theta)+d_{LP}(\pi_t^{\mathcal{A}},\widetilde{\pi}_t^{\mathcal{A}})<\gamma$ and

$$\sup_{(\theta,\omega,\pi^\Theta)\in\Theta\times\Omega\times\Delta(\Theta)} d_{LP}^{\mathcal{A}}(\sigma(\theta,\omega,\pi^\Theta),\sigma'(\theta,\omega,\pi^\Theta))<\gamma$$

we have

$$(*)\quad d_{LP}(\int_{\mathcal{A}\times\Theta}T(\circ|\theta,a,\omega_{t+1},\pi_t^\Theta,\pi_t^{\mathcal{A}})*\sigma(\theta,\omega_{t+1},\pi_t^\Theta)[da]*\pi^\Theta(d\theta),$$
$$\int_{\mathcal{A}\times\Theta}T(\circ|\theta,a,\omega_{t+1},\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}})*\sigma'(\theta,\omega_{t+1},\widetilde{\pi}_t^\Theta)[da]*\widetilde{\pi}_t^\Theta(d\theta))<\varepsilon$$

We proceed by analyzing the telescopic expansion of (*) into the following terms

$$\text{(i)}\quad d_{LP}(\int_{\mathcal{A}\times\Theta}T(\circ|\theta,a,\omega_{t+1},\pi_t^\Theta,\pi_t^{\mathcal{A}})*\sigma(\theta,\omega_{t+1},\pi_t^\Theta)[da]*\pi^\Theta(d\theta),$$
$$\int_{\mathcal{A}\times\Theta}T(\circ|\theta,a,\omega_{t+1},\pi_t^\Theta,\pi_t^{\mathcal{A}})*\sigma'(\theta,\omega_{t+1},\pi_t^\Theta)[da]*\pi_t^\Theta(d\theta))<\frac{\varepsilon}{5}$$

$$\text{(ii)}\quad d_{LP}(\int_{\mathcal{A}\times\Theta}T(\circ|\theta,a,\omega_{t+1},\pi_t^\Theta,\pi_t^{\mathcal{A}})*\sigma'(\theta,\omega_{t+1},\pi_t^\Theta)[da]*\pi^\Theta(d\theta),$$
$$\int_{\mathcal{A}\times\Theta}T(\circ|\theta,a,\omega_{t+1},\pi_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}})*\sigma'(\theta,\omega_{t+1},\pi_t^\Theta)[da]*\pi_t^\Theta(d\theta))<\frac{\varepsilon}{5}$$

(iii)

$$d_{LP}(\int_{\mathcal{A}\times\Theta}T(\circ|\theta,a,\omega_{t+1},\pi_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}})*\sigma'(\theta,\omega_{t+1},\pi_t^\Theta)[da]*\pi_t^\Theta(d\theta)),$$
$$\int_{\mathcal{A}\times\Theta}T(\circ|\theta,a,\omega_{t+1},\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}})*\sigma'(\theta,\omega_{t+1},\pi_t^\Theta)[da]*\pi_t^\Theta(d\theta))\quad<\frac{\varepsilon}{5}$$

(iv) $d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}}) * \sigma'(\theta, \omega_{t+1}, \pi_t^{\Theta})[da] * \pi_t^{\Theta}(d\theta)),$

$$\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}}) * \sigma'(\theta, \omega_{t+1}, \widetilde{\pi}_t^{\Theta})[da] * \pi_t^{\Theta}(d\theta)) < \frac{\varepsilon}{5}$$

(v)

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}}) * \sigma'(\theta, \omega_{t+1}, \widetilde{\pi}_t^{\Theta})[da] * \pi_t^{\Theta}(d\theta)),$$

$$\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}}) * \sigma'(\theta, \omega_{t+1}, \widetilde{\pi}_t^{\Theta})[da] * \widetilde{\pi}_t^{\Theta}(d\theta))) \quad < \frac{\varepsilon}{5}$$

Continuity of each term clearly follows from the continuity of T, σ and σ'. Together these terms imply that (*) holds. From the definition of P^C, we then have the desired continuity of $P^C(\omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}|\sigma)$ in the weak-* topology. □

Lemma 4 establishes the uniform continuity of $(P_\infty^{\mathcal{A}})^\tau(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ conditional on $(\omega_{t+1}, ..., \omega_{t+\tau})$. Now we prove that the distribution of $(\omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})$ is uniformly continuous in $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$.

Lemma 5. *$(P_\infty^{\mathcal{A}})^\tau(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ is uniformly continuous for any $\tau < \infty$*

Proof. [New proof in v27] To show our lemma we need to prove that expectations of continuous functions of $(\omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})$ are continuous in $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ and that expectations of uniformly continuous functions of $(\omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})$ are uniformly continuous in $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$. Since the proofs are essentially identical, we describe both inductive arguments at the same time.

We will proceed by considering an arbitrary (uniformly) continuous function $h : \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A}) \to \mathbb{R}$ and showing that $E[h(\omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}]$ is (uniformly) continuous in $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$. Not initially, that the proposition is clearly true for $\tau = 1$ since P^C and G are)uniformly) continuous in the weak-* topology.

Assume that the proposition is trust for all $\tau' < \tau < \infty$. Using the law of iterated expectations we can write

$$E[h(\omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}] =$$
$$E_t[E_{t+\tau-1}[h(\omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})|(\omega_{t+\tau-1}, \pi_{t+\tau-1}^{\Theta}, \pi_{t+\tau-1}^{\mathcal{A}})]|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}]$$

If we define

$$f(\omega_{t+\tau-1}, \pi^{\Theta}_{t+\tau-1}, \pi^{\mathcal{A}}_{t+\tau-1}) = E_{t+\tau-1}[h(\omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})|(\omega_{t+\tau-1}, \pi^{\Theta}_{t+\tau-1}, \pi^{\mathcal{A}}_{t+\tau-1})]$$

then f is (uniformly) continuous in $(\omega_{t+\tau-1}, \pi^{\Theta}_{t+\tau-1}, \pi^{\mathcal{A}}_{t+\tau-1})$ since we have assumed our proposition holds for $t = 1$. But then we can write

$$E[h(\omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})|\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t] = E_t[f(\omega_{t+\tau-1}, \pi^{\Theta}_{t+\tau-1}, \pi^{\mathcal{A}}_{t+\tau-1})|\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t]$$

Since by assumption our proposition holds for the $\tau - 1$ ahead operator, we have that

$$E_t[f(\omega_{t+\tau-1}, \pi^{\Theta}_{t+\tau-1}, \pi^{\mathcal{A}}_{t+\tau-1})|\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t]$$

is (uniformly) continuous in $(\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t)$. Therefore

$$E[h(\omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})|\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t]$$

is (uniformly) continuous in $(\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t)$ and our proof is complete. □

We now prove Theorem 1 through a series of lemmas. First we prove that conditional on the aggregate state in period $t + 1$, the empirical distributions of actions and types in period $t + 1$ converge to the deterministic distributions described by $P^C(\omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma)$.

Lemma 6. *Fix $\delta > 0$ and $\rho \in [0, 1)$. Assume $\sigma \in \Sigma$ is uniformly continuous. Then there exists N^* such that for any $N > N^*$ and for any $(\omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t)$ where π^{Θ}_t is a tight probability measure we have*

$$\begin{aligned} d^{\Theta}_{LP}(\pi^{\Theta}_{t+1}, \pi^{\Theta}_C) &< \delta \\ d^{\mathcal{A}}_{LP}(\pi^{\mathcal{A}}_{t+1}, \pi^{\mathcal{A}}_C) &< \delta \end{aligned}$$

with probability at least $1 - \rho$ where $P^C(\omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma) = (\pi^{\Theta}_C, \pi^{\mathcal{A}}_C)$. Furthermore, convergence is uniform over $(\omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t)$ and at the rate $O(N^{-0.5})$.

Proof. Our first step will be to approximate the measure π^{Θ}_t using a counting measure with a finite number of atoms that is independent of N. We can then apply the uniform law of large numbers to each one of these finite atoms to get uniform convergence of the transitions of the approximating measure as $N \to \infty$. From the continuity of our transition operator, this then implies the uniform convergence of π_t^Θ.

Since π_t^Θ is a tight measure, for any $\gamma > 0$ we can choose a compact set $U \subset \Theta$ such that $\pi_t^\Theta[U] > 1 - \gamma$. Let $N(\theta, \gamma) = \{\theta' \in \Theta : d_\Theta(\theta, \theta') < \gamma\}$ be a γ open neighborhood centered on θ. From the compactness of U, for any collection of γ-neighborhoods of every point in U we can choose a finite open cover. Let $C(\gamma) = \{N(\theta_k, \gamma) : \theta_k \in U\}_{k=1}^K$ denote such an open cover with the minimal number of covering sets, and let $K = |C(\gamma)|$. Note that from Assumption 2 we have that $K < \left(\sqrt{d}\frac{R(\gamma)}{\gamma}\right)^d < \overline{K}(\gamma) < \infty$. Let $C_i = N(\theta_k, \gamma)$ and $\overline{C}_i = C_i - \cup_{j>i} C_j$. Therefore, $\cup_{i=1}^K \overline{C}_i = U$ and $i \neq j$ implies $\overline{C}_i \cap \overline{C}_j = \varnothing$.

We define our a γ approximation of π_t^Θ by defining $\widetilde{\pi}_t^\Theta \in \Delta_K(\Theta)$ where

$$\begin{aligned} \widetilde{\pi}_t^\Theta &= \sum_{k=1}^K \alpha_k * \delta_{\theta_k} \\ \alpha_k &= \pi_t^\Theta[\overline{C}_k] \end{aligned}$$

where δ_{θ_k} denotes a measure 1 atom on θ_k. Note that $d_{LP}^\Theta(\widetilde{\pi}_t^\Theta, \pi_t^\Theta) < \gamma$ since for any $A \in \mathcal{B}(\Theta)$ we have

$$\begin{aligned} \widetilde{\pi}_t^\Theta(A^\gamma) &= \sum_{k=1}^K \alpha_k * 1\{A \cap \overline{C}_k \neq \varnothing\} \\ \pi_t^\Theta(A) &= \sum_{k=1}^K \pi_t^\Theta(A \cap \overline{C}_k) \end{aligned}$$

where U^γ is a γ neighborhood of the set U. Note that for each k, $\alpha_k * 1\{A \cap \overline{C}_k \neq \varnothing\} > \pi_t^\Theta(U \cap \overline{C}_k)$ so $\widetilde{\pi}_t^\Theta(A^\gamma) + \gamma \geq \pi_t^\Theta(A)$ as required. The converse inequality holds by a similar analysis.

Let the empirical distribution of types and actions in period $t+1$ of the N agent game conditional on $(\omega_{t+1}, \widetilde{\pi}_t^\Theta, \pi_t^\mathcal{A})$ be denoted as $\widetilde{\pi}_{t+1}^\Theta$ and $\widetilde{\pi}_{t+1}^\mathcal{A}$ respectively. Similarly, denote the empirical distribution of types and actions in period $t+1$ of the N agent game conditional on $(\omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})$ be denoted as π_{t+1}^Θ and $\pi_{t+1}^\mathcal{A}$ respectively

Let $\mathcal{S}(\pi_t^\Theta)$ denote the support of π_t^Θ. Considered as an empirical distribution of types in the N-agent game, $\widetilde{\pi}_t^\Theta$ has at least $\frac{N}{K}$ realizations for each $\theta^* \in \mathcal{S}(\widetilde{\pi}^\Theta)$. Each

θ^* transitions to a type in the next period according to the probability measure

$$P_1(\theta^*, \circ) = \int_{\mathcal{A}} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}_t^{\Theta}, \pi_t^{\mathcal{A}}) * \sigma(\theta, \omega_{t+1}, \widetilde{\pi}_t^{\Theta})[da]$$

Denote the CDF associated with the measure $P_1(\theta^*, \circ)$ as $J(\circ|\theta^*)$. Let $J_{N/K}(\circ|\theta^*)$ represent an empirical CDF of $\frac{N}{K}$ realizations of $J(\circ|\theta^*)$. From Corollary 8, $J_{N/K}(\circ|\theta^*) \to J(\circ|\theta^*)$ almost surely in the Kolmogorov metric as $N \to \infty$ at a rate of $O(N^{-0.5})$. Hence the empirical distribution of transitions of $\theta_t = \theta^*$ to θ_{t+1} converges in the Kolmogorov topology (and hence weakly) to the true distribution defined by

$$\int_{\mathcal{A}} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}_t^{\Theta}, \pi_t^{\mathcal{A}}) * \sigma(\theta, \omega_{t+1}, \widetilde{\pi}_t^{\Theta})[da]$$

Since this convergence is uniform over $\mathcal{S}(\widetilde{\pi}^{\Theta})$, we have that the empirical distribution of types in period $t+1$, $\widetilde{\pi}_{t+1}^{\Theta}$, converges almost surely to

$$\widetilde{\pi}_C^{\Theta}[\circ] = \sum_{\theta^* \in \mathcal{S}(\widetilde{\pi}^{\Theta})} \widetilde{\pi}_t^{\Theta}[\theta^*] * \int_{\mathcal{A}} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}_t^{\Theta}, \pi_t^{\mathcal{A}}) * \sigma(\theta, \omega_{t+1}, \widetilde{\pi}_t^{\Theta})[da]$$

An essentially identical argument proves that $\widetilde{\pi}_{t+1}^{\mathcal{A}}$ converges almost surely to

$$\widetilde{\pi}_C^{\mathcal{A}}[\circ] = \sum_{\theta^* \in \mathcal{S}(\widetilde{\pi}^{\Theta})} \widetilde{\pi}_{t+1}^{\Theta}[\theta^*] * \sigma(\theta^*, \omega_{t+1}, \widetilde{\pi}_{t+1}^{\Theta})[\circ]$$

Further note that $P^C(\omega_{t+1}, \widetilde{\pi}_t^{\Theta}, \pi_t^{\mathcal{A}}|\sigma) = (\widetilde{\pi}_C^{\Theta}, \widetilde{\pi}_C^{\mathcal{A}})$. Finally, note that the aggregate convergence rate is $O(N^{-0.5})$ by taking the maximum of the convergence rate over the at most $\overline{K}(\gamma)$ points in the support of $\widetilde{\pi}_t^{\Theta}$.

From the continuity of T and σ, the distributions of π_{t+1}^{Θ} and $\pi_{t+1}^{\mathcal{A}}$ must be approximately equal to the distribution of $\widetilde{\pi}_{t+1}^{\Theta}$ and $\widetilde{\pi}_{t+1}^{\mathcal{A}}$ for small enough γ. Therefore, we have for small enough γ (and hence large enough K) with probability[58] at least $1-\rho$

$$\begin{aligned} d_{LP}^{\Theta}(\pi_{t+1}^{\Theta}, \widetilde{\pi}_{C}^{\Theta}) &< \frac{\delta}{2} \\ d_{LP}^{\mathcal{A}}(\pi_{t+1}^{\mathcal{A}}, \widetilde{\pi}_{C}^{\mathcal{A}}) &< \frac{\delta}{2} \end{aligned}$$

From the continuity of P^C, we have for small enough γ that

$$\begin{aligned} d_{LP}^{\Theta}(\pi_{C}^{\Theta}, \widetilde{\pi}_{C}^{\Theta}) &< \frac{\delta}{2} \\ d_{LP}^{\mathcal{A}}(\pi_{C}^{\mathcal{A}}, \widetilde{\pi}_{C}^{\mathcal{A}}) &< \frac{\delta}{2} \end{aligned}$$

Combining these relations yields with probability at least $1-\rho$

$$\begin{aligned} d_{LP}^{\Theta}(\pi_{t+1}^{\Theta}, \pi_{C}^{\Theta}) &< \delta \\ d_{LP}^{\mathcal{A}}(\pi_{t+1}^{\mathcal{A}}, \pi_{C}^{\mathcal{A}}) &< \delta \end{aligned}$$

□

We will now show that if the game starts at a tight measure π_0^{Θ}, then the evolution of the nonatomic game necessitates that the measure of types remain tight over finite horizons.

Lemma 7. *Suppose that π_t^{Θ} is a tight measure and let $P^C(\omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}|\sigma) = (\pi_{t+1}^{\Theta}, \pi_{t+1}^{\mathcal{A}})$. Then π_{t+1}^{Θ} is tight.*

Proof. Since π_t^{Θ} is a tight measure we can generate an approximating measure composed of a convex combination of a finite set of atoms as per our proof of Lemma 6, which we denoted $\widetilde{\pi}_t^{\Theta}$. Since T is tight, each of these finite atoms yields a tight measure in period $t+1$ given the nonatomic limit game dynamics defined by P^C. Denote this step-ahead measure $\widetilde{\pi}_{t+1}^{\Theta}$. Since P^C is continuous in $\Delta(\Theta)$ and $\widetilde{\pi}_t^{\Theta}$ is an

[58]Since the convergence of $(\widetilde{\pi}_{t+1}^{\Theta}, \widetilde{\pi}_{t+1}^{\mathcal{A}})$ to $(\widetilde{\pi}_{C}^{\Theta}, \widetilde{\pi}_{C}^{\mathcal{A}})$ is almost sure in N, for any $\rho, \delta > 0$ we can find a sufficiently large $\overline{N}$ such that for all $N > \overline{N}$ we have

$$\begin{aligned} d_{LP}^{\Theta}(\widetilde{\pi}_{t+1}^{\Theta}, \widetilde{\pi}_{C}^{\Theta}) &< \frac{\delta}{2} \\ d_{LP}^{\mathcal{A}}(\widetilde{\pi}_{t+1}^{\mathcal{A}}, \widetilde{\pi}_{C}^{\mathcal{A}}) &< \frac{\delta}{2} \end{aligned}$$

with probability at least $1-\rho$. This "largeness" requirement on N is independent of the requirement that N be large enough that we can choose a sufficiently small γ.

arbitrarily close approximation of π_t^Θ, we have that the (tight) measure $\widetilde{\pi}_{t+1}^\Theta$ can be chosen to be arbitrarily close to π_{t+1}^Θ. Therefore π_{t+1}^Θ is tight. □

From the continuity of P^C, we have that if $d(\omega^r, \omega^s) \to 0$ that $d_{LP}^\Theta(\pi_t^\Theta(\omega^r), \pi_t^\Theta(\omega^t)) \to 0$ and $d_{LP}^\mathcal{A}(\pi_t^\mathcal{A}(\omega^r), \pi_t^\mathcal{A}(\omega^t)) \to 0$. The following lemma proves that conditional on ω^t, we get weak-* convergence of $(\pi_\tau^\Theta, \pi_\tau^\mathcal{A}) \to (\pi_\tau^\Theta(\omega^\tau), \pi_\tau^\mathcal{A}(\omega^\tau))$ uniformly for all $\tau \leq t$.

Lemma 8. [59] *Assume all of the agents but agent i follows continuous strategy σ, play commences at $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ where π_0^Θ is tight. For any $\gamma, \rho > 0$ we have uniformly over $\omega^t \in \Omega^t$ that for large enough N with probability at least $1-\rho$ that conditional on ω^t*

$$\begin{aligned} d_{LP}^\Theta(\pi_t^\Theta, \pi_t^\Theta(\omega^t)) &< \delta \\ d_{LP}^\mathcal{A}(\pi_t^\mathcal{A}, \pi_t^\mathcal{A}(\omega^t)) &< \delta \end{aligned}$$

Furthermore the convergence is uniform over $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ and at rate of $O(N^{-0.5})$.

Proof. Let $\omega^t = (\omega_0, ..., \omega_t)$. Choose any sequence $(\delta_1, ..., \delta_t)$ such that $\delta_\tau < \delta_{\tau+1}$, $\delta_\tau > 0$, and $\delta_t < \delta$.

Suppose all agents play according to continuous strategy σ. Conditional on starting at $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ and ω_1 of the sequence ω^t, Lemma 6 proves that we can choose N_1 sufficiently large such that with probability at least $(1-\rho)^{1/t}$ we have

$$\begin{aligned} d_{LP}^\Theta(\pi_1^\Theta, \pi_1^\Theta(\omega^1)) &< \delta_1 \\ d_{LP}^\mathcal{A}(\pi_1^\mathcal{A}, \pi_1^\mathcal{A}(\omega^1)) &< \delta_1 \end{aligned} \tag{1}$$

where $\omega^1 = (\omega_0, \omega_1)$. Furthermore, Lemma 7 implies that $\pi_1^\Theta(\omega^1)$ is tight. Conditional on $\omega^2 = (\omega_0, \omega_1, \omega_2)$ and equation (1) holding, Lemmas 6 and 7 imply that we can choose N_2 sufficiently large that

$$\begin{aligned} d_{LP}^\Theta(\pi_2^\Theta, \pi_2^\Theta(\omega^2)) &< \delta_2 \\ d_{LP}^\mathcal{A}(\pi_2^\mathcal{A}, \pi_2^\mathcal{A}(\omega^2)) &< \delta_2 \end{aligned}$$

[59] In the case that Θ^t asymptotically becomes infinite dimensional as $t \to \infty$, we have to employ the modified Convergence Corollary that holds uniformly over w^t where $t < \tau^*$ and $\tau^* < \infty$ is chosen arbitrarily. This restriction is immaterial in our proofs of Theorems 2 and 3 since we deal with improvements over σ_{ID}^{DCE} that occur over a finite horizon.

with probability $(1-\rho)^{2/t}$ where $\pi_2^{\Theta}(\omega^2)$ is tight. By iterating this argument we have our conclusion where $N = \max\{N_1, ..., N_t\}$. Finally note that the convergence rate is $O(N^{-0.5})$ in each of these t steps, implying a uniform convergence rate of $O(N^{-0.5})$ over all periods between 1 and t.

Suppose that agent i plays according to $\sigma' \neq \sigma$ (σ' may be discontinuous). Then we have that at each stage of the argument that the realized π_τ^{Θ}, $\tau \leq t$, is perturbed by a measure equal to $\frac{1}{N}[\zeta_\theta - \zeta_{\theta'}]$ where ζ_a is a measure one atom on $\theta \in \Theta$. Since

$$\pi_\tau^{\Theta} + \frac{1}{N}[\zeta_\theta - \zeta_{\theta'}] \to \pi_\tau^{\Theta}$$

in the Kolmogorov topology (and hence weakly) as $N \to \infty$, the continuity of T and σ implies that we can choose N sufficiently large that this perturbation does not alter our earlier conclusion. □

We now proceed to prove the theorems from the body of Section 3. Each theorem is restated with its original numbering for convenience

Theorem 1. *Fix $\tau^* < \infty$, $\gamma > 0$ and $\rho \in [0,1)$. Assume $\sigma \in \Sigma$ is uniformly continuous. Then there exists $N^*, \overline{\gamma} > 0$ such that for any $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), (\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}})$ where $d_\Omega(\omega_t, \widetilde{\omega}_t) + d_{LP}^{\Theta}(\pi_t^{\Theta}, \widetilde{\pi}_t^{\Theta}) + d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}}, \widetilde{\pi}_t^{\mathcal{A}}) < \overline{\gamma}$, any $N > N^*$, and all $\tau \in \{1, .., \tau^*\}$ we have that*

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^\tau(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}})) < \gamma$$

with probability at least $1-\rho$. Furthermore the convergence rate is $O(N^{-0.5})$ and uniform over $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$.

Proof. Consider the probability of a measurable event of the form $Q \subset \prod_{\tau=0}^{\tau=\tau^*} \Omega$. Lemma 8 proves that for any $\delta > 0$, we can choose N^* sufficiently large such that conditional on $\omega^{\tau^*} \in Q$ and for all $\tau \leq \tau^*$

$$\begin{aligned} d_{LP}^{\Theta}(\pi_\tau^{\Theta}, \pi_\tau^{\Theta}(\omega^\tau)) &< \delta \\ d_{LP}^{\mathcal{A}}(\pi_\tau^{\mathcal{A}}, \pi_\tau^{\mathcal{A}}(\omega^\tau)) &< \delta \end{aligned} \qquad (*)$$

with probability at least $1-\rho$. Let us consider only events in the N-agent game where $(*)$ holds. Once we note that $G(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ is uniformly continuous in $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$, for δ sufficiently small (and hence N sufficiently large) we have that

$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^\tau(\circ|\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}),(P_\infty^{\mathcal{A}})^\tau(\circ|\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}))<\gamma$ uniformly.[60] The convergence rate is $O(N^{-0.5})$ is inherited from Lemma 8.

Therefore we have for any $\varepsilon,\rho>0$ we can choose N^* such that for $N>N^*$ we have with probability $\sqrt{1-\rho}$

(1) $d_K^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^\tau(\circ|\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}),(P_\infty^{\mathcal{A}})^\tau(\circ|\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}))<\frac{\varepsilon}{3}$
(2) $d_K^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^\tau(\circ|\widetilde{\omega}_t,\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}}),(P_\infty^{\mathcal{A}})^\tau(\circ|\widetilde{\omega}_t,\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}}))<\frac{\varepsilon}{3}$

From the uniform continuity of $P_\infty^{\mathcal{A}}$, we can choose $\overline{\gamma}>0$ sufficiently small that if $d_\Omega(\omega_t,\widetilde{\omega}_t)+d_{LP}^\Theta(\pi_t^\Theta,\widetilde{\pi}_t^\Theta)+d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}},\widetilde{\pi}_t^{\mathcal{A}})<\overline{\gamma}$, then

3. $d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_\infty^{\mathcal{A}})^\tau(\circ|\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}),(P_\infty^{\mathcal{A}})^\tau(\circ|\widetilde{\omega}_t,\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}}))<\frac{\varepsilon}{3}$

Combining these relations we have $d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^\tau(\circ|\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}),(P_\infty^{\mathcal{A}})^\tau(\circ|\widetilde{\omega}_t,\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}}))<\varepsilon$ with probability $1-\rho$. From the finiteness of τ^*, it follows that we can choose $N^*,\overline{\gamma}>0$ so that this conclusion holds uniformly for all $\tau\in\{1,...,\tau^*\}$. □

We can strengthen Theorem 1 and it's resultant corollaries to uniform convergence (a choice of $N^*,\overline{\gamma}>0$ that holds uniformly over $(\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}),(\widetilde{\omega}_t,\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}}))$ if the condition of Lemma **??** is satisfied.

We now provide a proof of our extension of Theorem 1 to the case where Θ^t is a finite dimensional Euclidean space.

Corollary 1. *Fix $\tau^*<\infty$, $\gamma>0$ and $\rho\in[0,1)$. Assume:*

- $\{T_{t+\tau}(\circ|\theta_t,a_t,\omega_{t+1},\pi_t^\Theta,\pi_t^{\mathcal{A}})\}_{\tau=1}^{\tau^*}$ *are uniformly continuous*
- $\sigma\in\Sigma$ *is uniformly continuous over* $\Theta^{\tau^*}\times\Omega\times\Delta(\Theta^{\tau^*})$

There exists $N^*,\overline{\gamma}>0$ such that for any $(\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}),(\widetilde{\omega}_t,\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}})$ where $d_\Omega(\omega_t,\widetilde{\omega}_t)+d_{LP}^\Theta(\pi_t^\Theta,\widetilde{\pi}_t^\Theta)+d_{LP}^{\mathcal{A}}(\pi_t^{\mathcal{A}},\widetilde{\pi}_t^{\mathcal{A}})<\overline{\gamma}$, any $N>N^*$, and all $\tau\in\{1,..,\tau^*\}$ we have that

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^\tau(\circ|\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}}),(P_\infty^{\mathcal{A}})^\tau(\circ|\widetilde{\omega}_t,\widetilde{\pi}_t^\Theta,\widetilde{\pi}_t^{\mathcal{A}}))<\gamma$$

with probability at least $1-\rho$.

[60] If the evolution of ω_t were independent of $(\pi_t^\Theta,\pi_t^{\mathcal{A}})$, we would simply note that since $(\pi_t^\Theta,\pi_t^{\mathcal{A}})$ converges to $(\pi_{t+\tau}^\Theta(\omega^\tau),\pi_{t+\tau}^{\mathcal{A}}(\omega^\tau))$ conditional on and uniformly over ω^τ, the weak-* convergence of $(\omega^\tau,\pi_{t+\tau}^\Theta,\pi_{t+\tau}^{\mathcal{A}})$ to $(\omega^\tau,\pi_{t+\tau}^\Theta(\omega^\tau),\pi_{t+\tau}^{\mathcal{A}}(\omega^\tau))$ must hold.
Instead, we note that if we focus on the high probability event that $(\pi_{t+\tau}^\Theta,\pi_{t+\tau}^{\mathcal{A}})$ in the N agent game remains close to (within δ of) $(\pi_{t+\tau}^\Theta(\omega^\tau),\pi_{t+\tau}^{\mathcal{A}}(\omega^\tau))$ conditional on ω^τ for all $\tau\le\tau^*$, then the continuity of $G(\circ|\omega_t,\pi_t^\Theta,\pi_t^{\mathcal{A}})$ insures the distribution of ω^{τ^*} will converge weakly as $N\to\infty$. This then implies the weak-* convergence of $(\omega_{t+\tau},\pi_{t+\tau}^\Theta,\pi_{t+\tau}^{\mathcal{A}})$ for all $\tau\le\tau^*$.

Proof. Note that in Lemma 6 our argument proving that the empirical distribution of the transitions of $\theta_t = \theta^*$ to θ_{t+1} converges weakly to the true distribution defined by

$$\int_{\mathcal{A}} T_\tau(U|\theta, a, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}) * \sigma(\theta, \omega_{t+1}, \pi_t^\Theta)[da]$$

applies separately for each $\tau \in \{t+1, ..., t+\tau^*\}$ since each Θ^τ is a finite dimensional Euclidean space. Further, this convergence can be taken to be uniform since $\{t+1, ..., t+\tau^*\}$ is finite. The remainder of our arguments for Lemma 8 and Theorem 1 can be applied directly. □

The following two corollaries provide mean field theorems for the case where a single agent deviates a significant amount (Corollary 2) and to the case where all of the agents deviate a small amount (Corollary 3).

Corollary 2. *Fix $\tau^* < \infty$, $\gamma > 0$ and $\rho \in [0,1)$. Assume $\sigma \in \Sigma$ is uniformly continuous. Suppose agent i deviates from σ to $\sigma' \in \Sigma$ in period t. Then there exists $N^*, \overline{\gamma} > 0$ such that for any $(\omega_t, \pi_t^\Theta), (\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta)$ where $d_\Omega(\omega_t, \widetilde{\omega}_t) + d_{LP}^\Theta(\pi_t^\Theta, \widetilde{\pi}_t^\Theta) < \overline{\gamma}$, any $N > N^*$, and all $\tau \in \{1, .., \tau^*\}$ we have that*

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((\widetilde{P}_N^A)^\tau(\circ|\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}), (P_\infty^A)^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^\mathcal{A})) < \gamma$$

with probability at least $1-\rho$.

Proof. Note that Lemma 8 allows for the case where a single agent i deviates from the continuous strategy σ. We can reapply the argument used in Theorem 1 to prove our result. □

Corollary 3. *Fix $\tau^* < \infty$, $\gamma > 0$ and $\rho \in [0,1)$. Assume $\sigma, \sigma' \in \Sigma$ are uniformly continuous. Let P_N^A denote the market evolution generated by σ in conjunction with T and G, and let $\widehat{P}_N^A$ reflect the market evolution generated by σ', T, and G. Then there exists $N^*, \overline{\gamma} > 0$ such that for any $(\omega_t, \pi_t^\Theta), (\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta)$ where $d_\Omega(\omega_t, \widetilde{\omega}_t) + d_{LP}^\Theta(\pi_t^\Theta, \widetilde{\pi}_t^\Theta) + d_{LP}^\mathcal{A}(\pi_t^\mathcal{A}, \widetilde{\pi}_t^\mathcal{A}) + d_\Sigma(\sigma, \sigma') < \overline{\gamma}$, any $N > N^*$, and all $\tau \in \{1, .., \tau^*\}$ we have that*

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((\widehat{P}_N^A)^\tau(\circ|\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}), (P_\infty^A)^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^\mathcal{A})) < \gamma$$

with probability at least $1-\rho$.

Proof. From Theorem 1, for $N < \infty$ sufficiently large and $\overline{\gamma} > 0$ sufficiently small we have for $d_\Omega(\omega_t, \widetilde{\omega}_t) + d^\Theta_{LP}(\pi^\Theta_t, \widetilde{\pi}^\Theta_t) + d^\mathcal{A}_{LP}(\pi^\mathcal{A}_t, \widetilde{\pi}^\mathcal{A}_t) < \overline{\gamma}$ that with probability at least $1 - \rho$

$$\begin{aligned} d^{\Omega\times\Theta\times\mathcal{A}}_{LP}((P^A_N)^\tau(\circ|\omega_t, \pi^\Theta_t, \pi^\mathcal{A}_t), (P^A_\infty)^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}^\Theta_t, \widetilde{\pi}^\mathcal{A}_t)) &< \frac{\gamma}{3} \\ d^{\Omega\times\Theta\times\mathcal{A}}_{LP}((\widehat{P}^A_N)^\tau(\circ|\omega_t, \pi^\Theta_t, \pi^\mathcal{A}_t), (\widehat{P}^A_\infty)^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}^\Theta_t, \widetilde{\pi}^\mathcal{A}_t)) &< \frac{\gamma}{3} \end{aligned}$$

where $\widehat{P}^A_\infty$ denotes the transition probability function associated with strategy σ' and evolution operators T and G. From the continuity of P^C with respect to σ, for $\overline{\gamma}$ sufficiently small and $d_\Sigma(\sigma, \sigma') < \overline{\gamma}$ we have

$$^{\Omega\times\Theta\times\mathcal{A}}_{LP}((P^A_\infty)^\tau(\circ|\omega_t, \pi^\Theta_t, \pi^\mathcal{A}_t), (\widehat{P}^A_\infty)^\tau(\circ|\widetilde{\omega}_t, \widetilde{\pi}^\Theta_t, \widetilde{\pi}^\mathcal{A}_t)) < \frac{\gamma}{3}$$

Combining these results yields our conclusion. □

To prove Theorem 2, we argue that the large finite game stays near the family of paths generated in the nonatomic limit game for large N. Since σ^{DCE}_{ID} is a best response to this family of paths, it must be an approximate best response to any "nearby" family of paths. Our mean field approximation results, Theorem 1 and Corollary 2, insure that the path of play realized in large finite game when all agents save one play σ^{DCE}_{ID} is such a "nearby" family of paths with high probability. From the boundedness of the utility functions, the benefit to optimizing in the event the path of play deviates significantly from the path of the equilibrium of the nonatomic limit game is bounded and vanishes in expectation as $N \to \infty$.

To make our argument precise, we derive a form of the One-Shot Deviation Principal to show a deviation from σ^{DCE}_{ID} yielding an ε-improvement would occur within time $\tau^* < \infty$. We work backward from the terminal information sets at time τ^* to show that no such profitable deviation exists in the event that play stays near the support of outcomes realized in the nonatomic game. Since this event occurs with very high probability as $N \to \infty$ and the utility loss outside of this event is bounded, we can conclude then that σ^{DCE}_{ID} is an $\varepsilon-$BNE of the N-agent game for sufficiently large N. The proof is not trivial since optimal deviations from σ^{DCE}_{ID} need not be continuous, which makes the utility from such a deviation (potentially) discontinuous as P^A_N converges to P^A_∞.

Theorem 2. *Fix $\varepsilon > 0$. Assume that σ_{ID}^{DCE} is uniformly continuous. Then we can choose N^*such that the DCE $(\sigma_{ID}^{DCE}, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ is an ε-Bayesian-Nash Equilibrium of the large finite stochastic game for $N > N^*$. Furthermore, N^* can be chosen uniformly across starting points $(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}}) \in \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$, $N > N^*$.*

Proof. Since the problem facing a single agent in the equilibrium of nonatomic game formed by σ_{ID}^{DCE} is a single person stochastic decision problem, we can without loss of generality find an outcome equivalent (potentially discontinuous) perfect competitive equilibrium strategy $\sigma_{\tau^*}^{PCE}$ where for any $\tau^* < \infty$ and all $\omega^\tau \in \Omega^\tau$ and $\theta_\tau \in \Theta$, $\tau \in \{1, ..., \tau^*\}$, we define[61]

$$V_\infty(\theta_\tau^i, \omega_\tau, \pi_\tau^{\Theta}, \pi_\tau^{\mathcal{A}} | \sigma_i^{PCE}, \sigma_{-i}^{DCE}) \geq \sup_{\sigma_i' \in \Sigma} V_\infty(\theta_\tau^i, \omega_\tau, \pi_\tau^{\Theta}, \pi_\tau^{\mathcal{A}} | \sigma_i', \sigma_{-i}^{DCE})$$

Let $Q \subset \Theta^{\tau^*} \times \Omega^{\tau^*}$ denote the support of the path of play realized in the equilibrium where all agents play σ_{ID}^{DCE}. Let Q^δ denote a δ neighborhood of Q. By definition, in the event $(\theta^{\tau^*}, \omega^{\tau^*}) \in Q$ where $\theta^{\tau^*} = (\theta_1, ..., \theta_{\tau^*})$ we have $\sigma_{\tau^*}^{PCE}(\theta_{\tau^*}, \omega^{\tau^*}) = \sigma_{ID}^{DCE}(\theta_{\tau^*}, \omega^{\tau^*})$. For $\tau > \tau^*$ we let $\sigma_{\tau^*}^{PCE}(\theta_{\tau^*}, \omega^{\tau^*}) = \sigma_{ID}^{DCE}(\theta_{\tau^*}, \omega^{\tau^*})$.

Suppose that σ_{ID}^{DCE} is not an ε-Bayesian-Nash for some $\varepsilon > 0$ for all $N > N^*$ and any sufficiently large N^*. Then it must be the case that for any $N^* > 0$ there exists $N > N^*$ and $\sigma' \in \Sigma$ such that for a measure $\rho > 0$ of $\theta_0 \in \Theta$ (with respect to π_0^{Θ})

$$\sum_{\tau=0}^{\infty} \delta^\tau E_0^{\Psi}[w_N(\theta_\tau, \sigma_{ID}^{DCE}(\theta_\tau, \omega^\tau), \omega_\tau, \pi_\tau^{\Theta}, \pi_\tau^{\mathcal{A}}) | P_N^A] + 2\varepsilon \leq$$
$$\sum_{\tau=0}^{\infty} \delta^\tau E_0^{\Psi}[w_N(\theta_\tau, \sigma'(\theta_\tau, \omega, \pi_\tau^{\Theta}), \omega_\tau, \pi_\tau^{\Theta}, \pi_\tau^{\mathcal{A}}) | \widetilde{P}_N^A]$$

From the uniform convergence of w_N to w we can conclude

$$\sum_{\tau=0}^{\infty} \delta^\tau E_0^{\Psi}[w(\theta_\tau, \sigma_{ID}^{DCE}(\theta_\tau, \omega^\tau), \omega_\tau, \pi_\tau^{\Theta}, \pi_\tau^{\mathcal{A}}) | P_N^A] + \frac{3}{2}\varepsilon \leq$$
$$\sum_{\tau=0}^{\infty} \delta^\tau E_0^{\Psi}[w(\theta_\tau, \sigma'(\theta_\tau, \omega, \pi_\tau^{\Theta}), \omega_\tau, \pi_\tau^{\Theta}, \pi_\tau^{\mathcal{A}}) | \widetilde{P}_N^A]$$

[61]Recollect that in the nonatomic game the initial state $(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ and ω^τ uniquely define $(\omega_\tau, \pi_\tau^{\Theta}, \pi_\tau^{\mathcal{A}})$.

From the the boundedness of w we can define M such that

$$\sup_{(\theta,a,\omega,\pi^{\Theta},\pi^{\mathcal{A}})\in\Theta\times\mathcal{A}\times\Omega\times\Delta_N(\Theta)\times\Delta_N(\mathcal{A})} w_N(\theta,a,\omega,\pi^{\Theta},\pi^{\mathcal{A}})- \inf_{(\theta,a,\omega,\pi^{\Theta},\pi^{\mathcal{A}})\in\Theta\times \mathrm{S}\times\mathcal{A}} w_N(\theta,a,\omega,\pi^{\Theta},\pi^{\mathcal{A}}) < M$$

This implies that there exists $\tau^* < \infty$ such that

$$\delta^{\tau^*} * M < \frac{\varepsilon}{8}$$

Therefore we can write[62]

$$\sum_{\tau=0}^{\tau^*} \delta^{\tau} E_0^{\Psi}[w(\theta_\tau,\sigma_{ID}^{DCE}(\theta_\tau,\omega^\tau),\omega_\tau,\pi_\tau^{\Theta},\pi_\tau^{\mathcal{A}})|P_N^A] + \frac{5}{4}\varepsilon \leq \sum_{\tau=0}^{\tau^*} \delta^{\tau} E_0^{\Psi}[w(\theta_\tau,\sigma'(\theta_\tau,\omega,\pi_\tau^{\Theta}),\omega_\tau,\pi_\tau^{\Theta},\pi_\tau^{\mathcal{A}})|\widetilde{P}_N^A]$$

Therefore it must be the case that a deviation to strategy σ' yields at least an ε benefit within $\tau^* < \infty$ periods.

Let $\sigma'_{ID} : \Theta \times \Omega^\infty \to \Delta(\mathcal{A})$ be defined so that $\sigma'_{ID}(\theta_t,\omega^t)$ generates the same distribution of actions as σ' conditional on the event (θ_t,ω^t). Since σ' is an ε improvement over σ_{ID}^{DCE}, we have from the continuity of w, T, and G and Lemma 8 that by choosing ρ,δ sufficiently small that σ'_{ID} is an ε improvement over σ_{ID}^{DCE} within τ^* periods for sufficiently large N.

Consider the problem facing the agent in the N-agent game in period τ^* given strategy σ_{ID}^{DCE} is followed in all future periods

$$\text{(OPT)} \quad \max_a w_N(\theta_{\tau^*},a,\omega_{\tau^*},\pi_{\tau^*}^{\Theta},\pi_{\tau^*}^{\mathcal{A}}) + E_{\tau^*}^{\Psi}[V_N(\theta_{\tau^*+1}^i,\omega_{\tau^*+1},\pi_{\tau^*+1}^{\Theta},\pi_{\tau^*+1}^{\mathcal{A}}|\sigma^{DCE})|a]$$

From the continuity of w, T, and G; the uniform convergence of w_N to w; the continuity of $\sigma_{\tau^*}^{PCE}$ for periods $\tau > \tau^*$; and the convergence of P_N^A and $\widetilde{P}_N^A$ to P_∞^A as $N \to \infty$ proven by Theorem 1 and Corollary 2, we have for $t > \tau^*$

$$\begin{aligned} V_N(\theta_t^i,\omega_t,\pi_t^{\Theta},\pi_t^{\mathcal{A}}|\widetilde{\sigma}_i,\sigma_{-i}^{DCE}) &\to V_\infty(\theta_t^i,\omega_t,\pi_t^{\Theta},\pi_t^{\mathcal{A}}|\sigma^{DCE}) \qquad \text{(VFE)} \\ V_N(\theta_t^i,\omega_t,\pi_t^{\Theta},\pi_t^{\mathcal{A}}|\sigma^{DCE}) &\to V_\infty(\theta_t^i,\omega_t,\pi_t^{\Theta},\pi_t^{\mathcal{A}}|\sigma^{DCE}) \end{aligned}$$

[62]This is a *continuity at infinity* assumption required for the One Step Deviation principle to apply.

Note that the agent in the nonatomic limit game solves the problem

$$\sigma^{PCE}_{\tau^*}(\theta_t, \omega^t) \in \arg\max_a \, w(\theta_{\tau^*}, a, \omega_{\tau^*}, \pi^{\Theta}_{\tau^*}, \pi^{\mathcal{A}}_{\tau^*}) + E^{\Omega}_{\tau^*}[V_\infty(\theta^i_{\tau^*+1}, \omega_{\tau^*+1}, \pi^{\Theta}_{\tau^*+1}, \pi^{\mathcal{A}}_{\tau^*+1} | \sigma^{DCE}) | a]$$

Given the continuity of $V_N \to V_\infty$ as $N \to \infty$, the continuity of w, and the uniform convergence of w_N to w, the Theorem of the Maximum (Berge [9]) implies that there is no ε-improvement over σ^{PCE} available in period τ^* in problem (OPT). Similarly, in the event $(\theta^{\tau^*}, \omega^{\tau^*}) \in Q^\delta$ where $\theta^{\tau^*} = (\theta_1, ..., \theta_{\tau^*})$ we have that σ^{DCE} is an ε best response to (OPT). Therefore it must be that there is an ε-improvement over σ^{PCE} available in periods $\{1, ..., \tau^* - 1\}$. We can repeat our induction step, however, to show that there is no ε-improvement over σ^{PCE} available in any period between $\{1, ..\tau^*\}$. Therefore, σ^{PCE} is an $\varepsilon - BNE$ and σ^{DCE} is an $\varepsilon - BNE$ in the N agent game conditional on the event Q^δ at $t = 0$.

Note that since P^A_N converges weakly to P^A_∞, for any $\delta > 0$ we can choose N sufficiently large that Q^δ is realized with probability at least $1 - \delta$.[63] Since σ^{DCE}_{ID} is an ε optima conditional on event Q^δ. From the boundedness of w_N we then have that σ^{DCE}_{ID} can be improved by at most

$$\varepsilon * (1 - \delta) + \delta * M$$

For δ sufficiently large (and hence N sufficiently large) we have that $\varepsilon*(1-\delta)+\delta*M < 2\varepsilon$ and hence σ^{DCE}_{ID} is an $2\varepsilon-$BNE. Our N^* can be chosen uniformly since since P^A_N converges to P^A_∞ and w_N converges to w uniformly. □

Theorem 3 requires additional equicontinuity and compactness assumptions to insure that the limit of a sequence of exact equilibria of the N-agent game are continuous. Alternately we could assume in Theorem 3 that $\sigma^N \to \sigma^\infty$ and that σ^∞ is uniformly continuous. Under these alternative assumptions, the remainder of the proof of Theorem 3 would go through by replacing references to continuous extensions $\widetilde{\sigma}^N$ with references to σ^∞. Since our goal is to limit assumptions on endogenous objects such as equilibrium strategies, we prefer to prove these extension

[63]The occurence of Q^δ with probability at leats $1 - \delta$ is a consequence of the convergence of P^A_N and $\widetilde{P}^A_N$ in the Lévy-Prokhorov metric that metricizes the weak-* topology.

and continuity properties rather than assume them. We note, however, that we have not found assumptions on the model primitives that ensure uniform continuity of the equilibrium strategy.

Theorem 3. *Assume that*

- Θ *and* Ω *are compact*
- *There exists an* N^* *such that* $\cup_{N=N^*}^{\infty}\mathcal{E}(N)$ *is uniformly equicontinuous in* $\Theta \times \Omega \times \Delta(\Theta)$

Then the correspondence $\mathcal{E}$ *is upper hemicontinuous with*

$$\underset{N\to\infty}{Lim}\mathcal{E}(N) = \mathcal{E}^{\infty} \subset \mathcal{E}^{NA}$$

where the equilibria have initial state $(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}}) \in \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$.

Proof. (of Theorem 3 - new as of version 31) By using an equivalent metric on $\Delta_N(\Theta)$, we can convert any uniformly continuous function $\sigma^N \in \mathcal{E}(N)$ into a Lipschitz continuous function (p. 79 Aliprantis and Border [5]). Lemma 2 of Lindenstrauss [26] shows that if σ^N is Lipschitz continuous in $\Delta_N(\Theta)$, then there exists a modulus preserving Lipschitz continuous extension of σ^N from $\Delta_N(\Theta)$ to $\Delta(\Theta)$.[64] This extension preserves the modulus of continuity under the original metric over $\Delta(\Theta)$. The extension becomes unique in the limit as $N \to \infty$ since $\Delta_N(\Theta)$ is dense in $\Delta(\Theta)$ in this limit.

Consider a sequence of exact BNE strategies, $\{\sigma^N\}_{N=1}^{\infty}$, and in an abuse of notation we also denote an arbitrary sequence of modulus preserving extensions of each element of this sequence to $\Theta \times \Omega \times \Delta(\Theta)$ as $\{\sigma^N\}_{N=1}^{\infty}$. Assume that the sequence is convergent, so $\sigma^N \to \sigma^{\infty}$.[65] Note that the compactness of Θ implies $\Delta(\Theta)$ is compact, and hence the product space $\Theta \times \Omega \times \Delta(\Theta)$ is compact under the product topology. Since $\{\sigma^N\}_{N=1}^{\infty}$ is equicontinuous, the Arzelà-Ascoli theorem prove that σ^{∞} is continuous (Dunford and Schwartz [16]). Let $\widehat{P}_N^A$ denote the transition

[64] An alternative proof technique would be to use Isbell [22] Corollary 1.3 (the uniformity of $\Delta(\mathcal{A})$ implies $\Delta(\mathcal{A})$ is an extension space) and Theorem 1.2 to prove that the family $\{\sigma_N^{BNE} : \Theta \times \Omega \times \Delta_N(\Theta) \to \Delta(\mathcal{A})\}_{N=1}^{\infty}$ of exact equilibria of the N-agent game can be equiuniformly extended to a family of functions $\{\widetilde{\sigma}_N^{BNE}\}_{N=1}^{\infty}$ with domain $\Theta \times \Omega \times \Delta(\Theta)$.

[65] Since the extensions are unqiue in the limit as $N \to \infty$, it does not matter which extensions of the original sequence of BNE strategies we study.

probability function for the N-agent game given σ^N, T, and G. Let $\widetilde{P}_N^A$ denote the transition probability function for the N-agent game given T, G, that all but one agent plays σ^N, and the remaining agent . Finally, let P_∞^A denote the the transition probability function for the nonatomic limit game given σ^∞, T, and G.

Let $Q \subset \Theta^{\tau^*} \times \Omega^{\tau^*}$ denote the support of the path of play realized in the equilibrium where all agents play σ^∞. Let Q^δ denote a δ neighborhood of Q.

Suppose $\sigma^N \to \sigma^\infty \notin \mathcal{E}^{NA}$. This can be the case only if there exists $\varepsilon > 0$ such that for a measure $\rho > 0$ of $\theta_0^i \in \Theta$ (with respect to π_0^Θ) satisfies the following suboptimality condition

$$\text{(Suboptimality)} \qquad V_\infty(\theta_0^i, \omega_0, \pi_0^\Theta, \pi_0^{\mathcal{A}} | \sigma^\infty) + 2\varepsilon < \sup_{\sigma_i' \in \Sigma} V_\infty(\theta_0^i, \omega_0, \pi_0^\Theta, \pi_0^{\mathcal{A}} | \sigma_i', \sigma_{-i}^\infty)$$

Denote a strategy that achieves such a profitable deviation σ_D. Throughout this argument, we will employ the indirect representations of all of the strategies we consider. From the the boundedness of w we can define M such that

$$\sup_{(\theta,a,\omega,\pi^\Theta,\pi^{\mathcal{A}}) \in \Theta \times \mathcal{A} \times \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})} w(\theta, a, \omega, \pi^\Theta, \pi^{\mathcal{A}}) - \inf_{(\theta,a,\omega,\pi^\Theta,\pi^{\mathcal{A}}) \in \Theta \times \mathrm{S} \times \mathcal{A}} w(\theta, a, \omega, \pi^\Theta, \pi^{\mathcal{A}}) < M$$

This implies that there exists $\tau^* < \infty$ such that

$$\delta^{\tau^*} * M < \frac{\varepsilon}{4}$$

Therefore we can write

$$\sum_{\tau=0}^{\tau^*} \delta^\tau E_0^\Omega[w(\theta_\tau, \sigma^\infty(\theta_\tau, \omega^\tau), \omega, \pi_\tau^\Theta, \pi_\tau^{\mathcal{A}}) | P_\infty^A] + \frac{3}{2}\varepsilon \leq \sum_{\tau=0}^{\tau^*} \delta^\tau E_0^\Omega[w(\theta_\tau, \sigma_D(\theta_\tau, \omega^\tau), \omega, \pi_\tau^\Theta, \pi_\tau^{\mathcal{A}}) | P_\infty^A]$$

Therefore, the following strategy yields a $\frac{3}{2}\varepsilon$ benefit (relative to σ^∞) for a postive measure of types in the nonatomic limit game in one of the first τ^* periods.

$$\sigma^{\tau^*}(\theta_\tau, \omega^\tau) = \{ \begin{array}{l} \sigma_D(\theta_\tau, \omega^\tau) \text{ if } \tau \leq \tau^* \\ \sigma^\infty(\theta_\tau, \omega^\tau) \text{ if } \tau > \tau^* \end{array}$$

Note that the deviation σ' is continuous for all $\tau > \tau^*$.

Consider the problem facing the agent in the nonatomic limit game in period τ^* for some $(\theta_{\tau^*}, \omega^{\tau^*}) \in Q$.

$$\text{(OPT)} \quad \max_a w(\theta_{\tau^*}, a, \omega_{\tau^*}, \pi^{\Theta}_{\tau^*}, \pi^{\mathcal{A}}_{\tau^*}) + E^{\Omega}_{\tau^*}[V_\infty(\theta^i_{\tau^*+1}, \omega_{\tau^*+1}, \pi^{\Theta}_{\tau^*+1}, \pi^{\mathcal{A}}_{\tau^*+1}|\sigma^{DCE})|a]$$

From the continuity of w, T, and G; the uniform convergence of w_N to w; the continuity of σ^{τ^*} for periods $\tau > \tau^*$; and the convergence of $\widetilde{P}^A_N$ and $\widehat{P}^A_N$ to P^A_∞ as $N \to \infty$ proven by Theorem 1 and Corollary 2, we have for $t > \tau^*$

$$\text{(A.1)} \qquad \begin{aligned} V_N(\theta^i_t, \omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma^{\tau^*}_i, \sigma^N_{-i}) &\to V_\infty(\theta^i_t, \omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma^{DCE}) \\ V_N(\theta^i_t, \omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma^N) &\to V_\infty(\theta^i_t, \omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma^{DCE}) \end{aligned}$$

Note that the agent in the N-agent game solves the problem

$$\sigma^N(\theta_{\tau^*}, \omega_{\tau^*}, \pi^{\Theta}_{\tau^*}) \in \arg\max_a w_N(\theta_{\tau^*}, a, \omega_{\tau^*}, \pi^{\Theta}_{\tau^*}, \pi^{\mathcal{A}}_{\tau^*}) + E^{\Psi}_{\tau^*}[V_N(\theta^i_{\tau^*+1}, \omega_{\tau^*+1}, \pi^{\Theta}_{\tau^*+1}, \pi^{\mathcal{A}}_{\tau^*+1}|\sigma^N)|a]$$

Given the continuity of $V_N \to V_\infty$ as $N \to \infty$, the continuity of w, and the uniform convergence of w_N to w, the Theorem of the Maximum (Berge [9]) implies that for sufficiently large N there is no $\frac{\varepsilon}{2}$-improvement over σ^N available in period τ^* in problem (OPT). Since $\sigma^N \to \sigma^\infty$ uniformly, there is no $\frac{3}{4}\varepsilon$ improvement over σ^∞available in period τ^*. Furthermore, since this result holds for all $(\theta_{\tau^*}, \omega^{\tau^*}) \in Q$ and w and T are continuous, it must be that for δ sufficiently small there is no ε improvement over σ^∞available in period τ^* for any $(\theta_{\tau^*}, \omega^{\tau^*}) \in Q^\delta$. Finally, note that since $\widehat{P}^A_N \to P^A_\infty$ in the weak-* topology, for any $\delta, \rho > 0$ we can choose N sufficiently large that $(\theta_{\tau^*}, \omega^{\tau^*}) \in Q^\delta$ with probability at least $1 - \rho$. Therefore, there cannot be an ε improvement over σ^∞ for a measure ρ of agents in period τ^*. It must be that there is an ε-improvement over σ^∞ available in periods $\{1, ..., \tau^* - 1\}$. However, we can repeat our induction steo to show that there is no ε-improvement over σ^∞ available in any period between $\{1, ..\tau^*\}$. Therefore, σ^∞ is an $\varepsilon-$BNE of the nonatomic limit game.

Since our argument holds for all convergent sequences of strategies, we have from Theorem 17.16 of Border et al. [5] that the equilibrium correspondence is upper hemicontinuous. The upper hemicontinuity of the convergence of $\mathcal{E}(N)$ is uniform over $\Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$ since $\widehat{P}^A_N$ converges to P^A_∞ and w_N converges to w uniformly and the continuous approximation, $\sigma_c^{'}$, can be made uniformly continuous by the compactness of $\Theta \times \Omega^{\infty}$and the Heine-Borel theorem. □

We now provide a proof of Corollary 5, which assumes directly that our sequence of exact BNE strategies, $\{\sigma^N\}_{N=1}^{\infty}$, converges to a uniformly continuous limit strategy σ^{∞}. In addition, we weaken our continuity conditions to require continuity only within an open set containing the equilibrium path. This weakening makes it clear that only discontinuities that are on (or arbitrarily close to) the equilibrium path of play can support equilibrium strategies that are not approximated by any DCE.

Corollary 5. *Consider a convergent sequence of uniformly continuous Bayesian-Nash equilibrium strategies, $\{\sigma^N\}_{N=1}^{\infty}$ such that $\sigma^N \to \sigma^{\infty}$. Assume σ^{∞} is uniformly continuous. Let $\Pi^{\tau} = \prod_{t=0}^{\tau} supp[(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})]$ be the support of the stochastic process induced by P_{∞}^{A}. For N sufficiently large, if there exists an open set U in $\prod_{t=0}^{\infty} \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$ containing Π^{τ} such that...*

- *T is uniformly continuous over $\Theta \times \mathcal{A} \times U$*
- *G is uniformly continuous over U*

Then σ^{∞} is a Dynamic Competitive Equilibrium of the limit game.

Proof. From Theorem 1 for any $\tau^* < \infty$, $\gamma > 0$, and $\rho \in (0,1]$, there exists N^* such that for economies with $N > N^*$ agents

$$(*) \qquad d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{A})^{\tau}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), (P_{\infty}^{A})^{\tau}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})) < \gamma$$

for $\tau \in \{1, .., \tau^*\}$ with probability at least $1 - \rho$. Choose γ sufficiently small that (*) implies that the path of $\{(\omega_{t+\tau}, \pi_{t+\tau}^{\Theta}, \pi_{t+\tau}^{\mathcal{A}})\}_{\tau=0}^{\tau^*}$ lies in the open set U with probability at least $1 - \rho$. From this point we can proceed with the logic of Theorem 3 can be applied directly. □

A.2. **Proofs from Section 4.**

Lemma 1. *Assume T is uniformly continuous, hold σ_{ID}^{DCE} fixed, and assume Condition A holds. For any $\varepsilon > 0$, we can find a $\gamma > 0$ such that $d_{LP}^{\Theta}(\pi^{\Theta}, \widetilde{\pi}^{\Theta}) < \gamma$ implies $d_{LP}^{\Theta}(P^{ID}(\circ|\pi^{\Theta}, \omega^t), P^{ID}(\circ|\widetilde{\pi}^{\Theta}, \omega^t)) < \varepsilon$ uniformly over $\pi^{\Theta} \in \Delta(\Theta)$.*

Proof. From the definition of P^{ID}, we can write the conclusion of our theorem as follows

$$(*)\quad d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta,a,\omega_{t+1},\pi^{\Theta},\pi^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \pi^{\Theta}(d\theta), \int_{\mathcal{A}\times\Theta} T(\circ|\theta,a,\omega_{t+1},\widetilde{\pi}^{\Theta},\widetilde{\pi}^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta)) < \varepsilon$$

We proceed by analyzing the telescopic expansion of this into the following terms

$$\text{(i)}\quad d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta,a,\omega_{t+1},\pi^{\Theta},\pi^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \pi^{\Theta}(d\theta), \int_{\mathcal{A}\times\Theta} T(\circ|\theta,a,\omega_{t+1},\pi^{\Theta},\widetilde{\pi}^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \pi^{\Theta}(d\theta)) < \frac{\varepsilon}{3}$$

$$\text{(ii)}\quad d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta,a,\omega_{t+1},\pi^{\Theta},\widetilde{\pi}^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \pi^{\Theta}(d\theta)), \int_{\mathcal{A}\times\Theta} T(\circ|\theta,a,\omega_{t+1},\widetilde{\pi}^{\Theta},\widetilde{\pi}^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \pi^{\Theta}(d\theta)) \quad < \frac{\varepsilon}{3}$$

(iii)

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta,a,\omega_{t+1},\widetilde{\pi}^{\Theta},\widetilde{\pi}^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \pi^{\Theta}(d\theta)), \int_{\mathcal{A}\times\Theta} T(\circ|\theta,a,\omega_{t+1},\widetilde{\pi}^{\Theta},\widetilde{\pi}^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta))) \quad < \frac{\varepsilon}{3}$$

Terms (i) and (ii) follow from the continuity of T.

Term (iii) requires Condition A. Let $D = \{\theta : \sigma^{DCE}_{ID}(\theta,\omega^{t+1})$ is discontinuous at $\theta\}$ and note that Condition A implies we can choose a set $D^* \supseteq D$ such that $\pi^{\Theta}[D^*], \widetilde{\pi}^{\Theta}[D^*] < \frac{\varepsilon}{9}$ and T and σ^{DCE}_{ID} are uniformly continuous over $\Theta - D^*$. From the properties of the weak-* topology we can choose γ sufficiently small that

$$d_{LP}(\int_{\mathcal{A}\times(\Theta-D^*)} T(\circ|\theta,a,\omega_{t+1},\widetilde{\pi}^{\Theta},\widetilde{\pi}^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \pi^{\Theta}(d\theta)), \int_{\mathcal{A}\times(\Theta-D^*)} T(\circ|\theta,a,\omega_{t+1},\widetilde{\pi}^{\Theta},\widetilde{\pi}^{\mathcal{A}}) * \sigma^{DCE}_{ID}(\theta,\omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta))) < \frac{\varepsilon}{9}$$

uniformly over π^{Θ}. Note then that

$$d_{LP}\Big(\int_{\mathcal{A}\times(\Theta-D^*)} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)),$$

$$\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)), \Big) < \frac{\varepsilon}{9}$$

and

$$d_{LP}\Big(\int_{\mathcal{A}\times(\Theta-D^*)} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta)),$$

$$\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta))\Big) < \frac{\varepsilon}{9}$$

Together these imply (iii) holds. Taken together (i), (ii), and (iii) imply (*). □