Approximation of Large Dynamic Games

Aaron L. Bodoh-Creed - Cornell University *

February 8, 2012

Abstract

We provide a framework for simplifying the analysis and estimation of dynamic games with many agents by using nonatomic limit game approximations. We show that that the equilibria of a nonatomic approximation are epsilon-Bayesian-Nash Equilibria of the dynamic game with many players if the limit game is continuous. We also provide conditions under which the Bayesian-Nash equilibrium strategy correspondence of a large dynamic game is upper hemicontinuous in the number of agents. We use our results to show that repeated static Nash equilibria are the only equilibria of continuous repeated games of imperfect public or private monitoring in the limit as N approaches infinity. Extensions include: games with large players in the limit as N approaches infinity; games with (discontinuous) entry and exit decisions; Markov perfect equilibria of complete information stochastic games with aggregate shocks; and games with private and/or imperfect monitoring. Finally we provide an application of our framework to the analysis of large dynamic auction platforms such as E-Bay using the nonatomic limit model of Satterthwaite and Shneyerov [33].

1 Introduction

Our study provides a novel framework for employing dynamic nonatomic games to generate insights about game-theoretic models with a large, but finite, set of players. A number of

*I would like to thank Ilya Segal, Paul Milgrom, Doug Bernheim, Jon Levin, C. Lanier Benkard, Eric Budish, Matthew Jackson, Ramesh Johari, Ichiro Obara, Philip Reny, Gabriel Weintraub, 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

classic works have studied the links between large finite games and nonatomic games, but the goal is often to justify the use of tractable, preexisting competitive models by providing strategic underpinnings founded in large finite models (seminal examples include Roberts and Postlewaite [46] and Green [25]). In contrast, our work seeks to generate results that allow for us to study preexisting large finite games of interest using tractable nonatomic limit games. Given the broad array of strategy sets (e.g. continuous strategy spaces), information structures (e.g. imperfect public and private monitoring), and payoff structures (e.g. models with a discrete entry or exit payoff or convex investment costs) used in modern game-theoretic studies, one significant contribution of our work is providing foundations for these techniques that cover a broad class of game-theoretic structures currently in use. Our techniques enable us to simplify the analysis of otherwise intractable large finite games as well as provide general theoretical insights such as anti-folk theorems for large games with imperfect monitoring.

Games with a large but finite number of players are important in a variety of fields of economics, although particularly in the areas of industrial organization (e.g. Ericson and Pakes [22]), mechanism and market design (e.g. Cripps and Swinkels [20]), and macroeconomics. Computational methods are often employed to estimate and solve these models, but 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) for comparative statics exercises or as part of a structural estimation algorithm 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 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 a negligible effect on current or future prices.[1] 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.[2] In equilibrium the evolution

[1]This behavioral assumption is precisely correct for an agent in a nonatomic model.

[2]Throughout this work we use the term *strategic* to mean that the agents choose actions assuming that

of market aggregates, as well as the agents' beliefs about the evolution, are consistent with the distribution of actions taken by all agents.

The computational tractability of dynamic competitive models stems from the fact that the distribution of agent types and actions in each period of play is determined uniquely by the initial distribution of types and the aggregate states of the model (e.g. demand shocks to prices). By describing the path of equilibrium play in terms of these aggregate states, the dimension of the policy space is rendered independent of the number of agents and the curse of dimensionality is avoided. These nonatomic modeling techniques provide crucial tractability in applications ranging from the study of market frictions (e.g. Wolinsky [52]) 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 in complex market settings.

The primary goal of our study is to provide sufficient conditions under which a large finite game of interest can be approximated by a tractable nonatomic limit game. Our principal theoretical contribution is captured in two approximation theorems. The first theorem provides continuity conditions under which any equilibrium of the nonatomic limit game is an $\varepsilon-$Bayesian-Nash equilibrium 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 the nonatomic limit game that are sufficient for the set of Bayesian-Nash equilibrium strategies to be upper hemicontinuous in the limit as the number of agents approaches infinity. Therefore, any convergent sequence of equilibrium strategies of the N-agent game approaches an 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.[3] We interpret our first approximation result as a proof of approximation in terms of marginal incentives, while our second theorem provides a stronger result showing approximation in terms of observable actions and economic outcomes.

each agent's individual action can affect market aggregates. Price taking behavior in general equilibrium models is an example of *competitive* behavior.

[3] 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 for arbitrarily large N.

Existing approximation results in the literature usually focus on complete information games.[4] Our techniques extend the range of models for which approximation techniques are valid to include:

- Dynamic games of incomplete information
- Games with persistent private information
- Imperfect public or private monitoring structures
- Stochastic games with aggregate shocks
- Dynamic games with large players and a fringe of small agents
- Dynamic games with well-behaved discontinuities[5]

Our approximation theorems also intersect the literature on anti-folk theorems (e.g. Green [25], Sabourian [47]). We interpret folk theorem results as a statement regarding the limits of game-theoretic modeling - if few outcomes can be ruled out as equilibria, then few sharp predictions can be made. Anti-folk theorems for large games, such as the ones we provide for large games of imperfect public and/or private monitoring, prove that the set of equilibrium payoffs in the large game converges to the set of repeated static Nash payoffs as the number of players approaches infinity. Our anti-folk theorem suggests that imperfect monitoring may imply that sharp equilibrium predictions can be made in models with many players. 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.

We provide a number of extensions of our results including: games with large players in the limit as N approaches infinity; games with (discontinuous) entry and exit decisions; Markov perfect equilibria of complete information stochastic games with aggregate shocks; games with private and/or imperfect monitoring; coalition proof equilibria; asymmetric games with multiple player roles (e.g. buyers and sellers); and games with asynchronous actions.

[4]See section 7 for a discussion of the related literature.

[5]In our context "well-behaved" means the agent policy functions are almost everywhere continuous. This condition is typically satisfied in games with entry, exit, or lumpy investment. See assumption 6 of section 5.5.

Finally, we sketch an approximate model of a dynamic market comprised of simultaneous auctions inspired by the E-Bay online market in order to illustrate our analysis framework. We argue that a limit model of the form used by Satterthwaite and Shneyerov [48] can be used as a limit model to approximate the complex dynamic game faced by agents (and modelers) in the context of large online marketplaces (Hendricks and Porter [26]). Researchers could then apply the relatively tractable limit model to study issues such as the effect of aggregate shocks on the terms of trade, the impact of market frictions on the efficiency of trade, and whether participants account for the dynamic aspects of the market when placing their bids.

1.1 Outline of Paper

Section 2 sketches an application of our ideas to large, dynamic auction platforms such as E-Bay. Section 3 outlines the model framework we analyze, describes the stochastic evolution of the games, and defines the equilibrium concepts we employ. Section 4 states the approximation theorems, and section 5 provides applications of our approximation theorems. Sections 5.1 and 5.2 use our results to provide anti-folk theorems for large repeated games of imperfect public or private monitoring. Section 5.3 applies our approximation theorems to games with a few large and many small players, section 5.4 extends our result to perfect equilibria, and section 5.5 discusses games with entry, exit, and convex investment costs that admit discontinuous equilibrium strategies. Section 6 elaborates on the application of our framework to large, dynamic auction platforms such as E-Bay and suggests a direction for future empirical work. Section 7 relates our study to the existing literature, and section 8 concludes. All proofs are relegated to Appendix A. Appendix B includes section B.1 that provides approximations to large games with finite coalitions, as well as section B.2 that extends our methods to games with asymmetric roles and asynchronous actions.

2 An Initial Example

As noted by Hendricks and Porter [26], the advent of online auction systems such as E-Bay presents both tremendous opportunities and challenges to economic analysis. On the positive side, online auctions provide a unique opportunity to study an enormous data set covering a broad array of consumer goods that are bought and sold across the globe.

Unfortunately, due to the large, dynamic nature of the game played by buyers and sellers on the market platform, both theoretical and structural analyses of this market have been slow to develop. Agent decisions include when to enter, how much to bid, and how to adjust bids to account for continuation values from the option to return to bid next period.

As a first step, buyers and sellers must make a judgement as to whether to enter the auction or not based on their valuation for the goods in the market. Assuming a continuous valuation distribution, the entry decision is necessarily discontinuous. Intuitively one would expect the entry decision to be an increasing (decreasing) function of valuation for buyers (sellers), but most prior works assume that agent policies are continuous and it is possible that this discontinuity would invalid the use of a nonatomic limit game. We extend our approximation results in section 5.5 to cover models with equilibria that exhibit such well-behaved discontinuities.

Once the decision to enter has been made, buyers in the E-Bay market face a daunting task collecting, tracking, and analyzing the available market information to choose a bidding (or selling) strategy. Possible variables of interest include the number of buyers and sellers in the market, the duration of the outcome, information about seller reserve prices that might be revealed by past auction outcomes, and the valuations of rival biddersthat have been revealed by past bids. Furthermore, the agents must account for how the dynamics of past market conditions endogenously generated the information under consideration to make inferences regarding the exogenous economic primitives.

If the E-Bay market is studied using techniques from structural industrial organization, the natural approach would be to treat the market as a complete information game, analyze the Markov perfect equilibria of the game to deduce the underlying primitives, and then use the primitives to conduct counterfactual exercises. However, even if the primitives can be ascertained, solving such a complex game is computationally infeasible. Moreover, we would argue that is not entirely reasonable to believe that the typical bidder for even high value items such as automobiles is willing or capable of conducting the elaborate game-theoretic reasoning required to behave in accord with the equilibrium predictions.

An alternative, tractable model of the auction market is to assume that the economic environment is stationary with the same measure of buyers and sellers considering entry in each period. Furthermore, it is also plausible to assume that agents do not track the identify of the auctioneers over time and are unaware of the other bidders who participate in the market.[6] Finally, we assume that buyers are randomly assorted to auctions, which

[6]Although bidders on E-Bay can observe the identity of prior bidders in an auction, the prevalence of

is suggested by the buyers inability to assess the true number of competing bidders in each auction. The tractable model we have sketched here is (essentially) the model developed by Satterthwaite and Shneyerov [48] for the study of market frictions.

Satterthwaite and Shneyerov [48] were able to leverage the tractability of their model to provide characterizations of the buyer and seller entry and bidding strategies and show that the market price converges to the Walrasian price as $\delta \to 1$. Given the discontinuous entry decision and the potential discontinuities in the buyers' bidding strategy, it is not possible to use prior work to argue that the Satterthwaite and Shneyerov [48] model is a useful limit approximation of any large finite game, much less the intractable large, finite stochastic game outlined above. However, as we show in Section 5.5, we can use our framework to show that the equilibrium correspondence of the Satterthwaite and Shneyerov [48] can be used to approximate the equilibrium correspondence of a rich, computationally-intractable large finite model of E-Bay.

Given our approximation framework we can use the Satterthwaite and Shneyerov [48] to potentially study market design issues such as the effect of changing entry costs on participation, the effect of demand and supply shocks on seller revenue, and whether agents fully account for the possibility of bidding in successive periods. If one enriches the model of Satterthwaite and Shneyerov [48] to include a model of E-Bay's reputation model, it may be possible to formally model the process of acquiring a reputation as a reliable seller. One could also examine the market implications of the presence of buyers or sellers whom myopically defect, which causes an adverse selection problem for buyers and sellers. The problem of adverse selection then opens the possiblity of assessing market interventions for this selection such E-Bay's buyer insurance program or escrow services.

The E-Bay marketplace is but one example of interesting markets whose analysis has been halted by an inability to analyze the natural (i.e. large finite) economic models. By employing limit approximations we can not only extend game-theoretic techniques to new settings, a variety of pressing counterfactual questions (e.g. solving the adverse selection problem in online markets) can be addressed.

---

bidding in the last few minutes or hours of the auction limits bidders from seeing either the number or identity of those de facto competing for a good. The E-Bay interface does provide tools for examining past auctions of sellers, so this assumption is a behavioral assumption rather a statement regarding the economic primitives.

3 Model Framework

We employ two forms of models in our analysis. Our focus is on analyzing games with $N < \infty$ players, and most of our results focus on limits where $N \to \infty$. We use as a limit approximation a nonatomic game with a measure 1 continuum of players. Since many of the primitives (e.g. type and action spaces) are shared by all of the models, we define and discuss these aspects of the models simultaneously in section 2.1. Assumption 4 relates the N-player games and the limit game through a convergence assumption on the utility functions of the large finite games.

The model framework described in section 3.1 is required to be sufficiently abstract to capture the breadth of games that we approximate. Readers who desire examples are urged to read section 3.2, 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 3.3 describes the evolution and equilibrium notions we use in the large finite models we approximate. Section 3.4 defines the evolution and equilibria of the nonatomic games we use as limit approximations of large finite games.

3.1 Model

We assume the game occurs in discrete time with periods indexed by $t \in \mathbb{N}$. Agent actions are drawn from $\mathcal{A} \subset \mathbb{R}^d$ where $d < \infty$, and we endow the space of probability measures over $\mathcal{A}$, denoted $\Delta(\mathcal{A})$, with the weak-* topology. The set of possible agent types realized in period t is $\Theta_t \subset \mathbb{R}^{d_t}$, $d_t < \infty$, and an agent's type is the private information of that agent. The type space is then $\Theta = \cup_{t=0}^{\infty} \Theta_t$ and is endowed with the following metric

$$d_\Theta(\theta, \theta') = \{ \begin{array}{l} +\infty \text{ if there is not } t \text{ such that } \theta, \theta' \in \Theta_t \\ d_t(\theta, \theta') \text{ otherwise} \end{array}$$

where $d_t(\theta, \theta')$ is the Euclidean metric over Θ_t. This construction insures that at any time t the set of possible types is a finite dimensional set, but the possibility of expanding the dimension of the type space as $t \to \infty$ allows us to encode private histories into the agents' types. Probability measures over the space Θ are denoted $\Delta(\Theta)$ and endowed with the weak-* topology.

When making statements regarding uniform continuity of functions of $\Delta(X)$, we employ the Lévy-Prokhorov metric $d_{LP}^X : \Delta(X) \times \Delta(X) \to \mathbb{R}_+$.[7] 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_t^Θ and $\pi_t^\mathcal{A}$ 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 $y \in Y$.[8]

Let the metric space (Ω, d_Ω) be the aggregate states of the model and denote a generic element realized at time t as $\omega_t \in \Omega$. We endow Ω with sigma algebra generated by the Borel sets, $\mathcal{B}(\Omega)$, when dealing with measurability issues. The aggregate state could include aggregate payoff shocks, public information 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, $\{(\omega_\tau, \pi_\tau^\Theta, \pi_\tau^\mathcal{A})\}_{\tau<t}$.[9] The initial aggregate state ω_0 is distributed according to a measure ν over Ω.

We assume that the agent types and the state of the economy evolve in the measure space $(\Psi, \mathcal{B}(\Psi))$ 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.[10] Types for the agents are drawn at $t = 0$ according to the tight probability measure $\mu^\Theta(\omega_0)$. By letting μ^Θ be conditioned on ω_0, 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 semi-anonymous. Semi-anonymity is a symmetry assumption that requires agent utility, strategies, and the evolution of an agent's type depend on the actions and types of other agents only through the empirical distributions π_τ^Θ and $\pi_\tau^\mathcal{A}$. In addition, semi-anonymity allows us to elegantly embed the N agent games in the nonatomic limit game as $N \to \infty$.

The aggregate state evolves according to a Markov process with the associated transi-

[7]The Lévy-Prokhorov norm metricizes the weak-* topology.

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

[9]$(\omega_\tau, \pi_\tau^\Theta, \pi_\tau^\mathcal{A})$ denotes an aggregate state (ω_τ) and the distribution of types and actions (π_τ^Θ and $\pi_\tau^\mathcal{A}$ respectively) realized in period τ.

[10]Throughout this work we use the notation $\prod_{t=M}^{N} X_i$ to refer to the product space $X_M \times X_{M+1} \times ... \times X_N$.

tion probability function from ω_t to ω_{t+1} denoted

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

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 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. 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. 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

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

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

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

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.[11]

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

We require a condition that is intermediate between tightness and uniform tightness of the family[12]

$$\{T(\circ|\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})\}_{(\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}) \in \Theta \times \mathcal{A} \times \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})}$$

Note that compactness of Θ is sufficient for Assumption 2, but stronger than required.

Assumption 2. *For any $\gamma > 0$ we can choose for each $(\theta_t, a_t, \omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})$ 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 $R(\gamma)$.*[13]

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*[14] *and bounded.*[15]

Uniform continuity assumptions are required at several points in our study. The most economically substantive effect of these assumptions is to rule out economies that allow

[11] 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 $\omega_{t+1} = h(\omega_t)$. 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.

[12] This provides the necessary leverage to make our asymptotics claims uniform over initial aggregate states $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$.

[13] The radius of a set U is

$$\sup_{\theta, \theta' \in U} d_\Theta(\theta, \theta')$$

[14] We require no continuity conditions on w_N.

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

for increasing marginal returns to scale in the agents' types and actions. It would be surprising if our techniques applied in models that admitted increasing returns to scale - these are the sorts of models that one would expect would admit equilibria with outcomes are determined by the actions of a few large players. In settings with only a few large players, the curse of dimensionality is not an issue and our approximation techniques are not required.

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 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 $(\theta, a, \omega) \in \Theta \times \mathcal{A} \times \Omega$*

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

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 $\sigma, \sigma' \in \Sigma$

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

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.[16] 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 conveyed by 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

[16]See further examples in section 2.2.

define the stochastic evolution operators and equilibrium concepts in the large finite and nonatomic games.[17]

3.2 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* $T(\theta, U) = 1\{\theta \in U\}$.[18] *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 (anonymous) history of play is a unique information set. To allow an agent to condition on his own history of actions, we incorporate this history into his type so* $\Theta_t = \mathbb{R} \times \prod_{\tau=0}^{t-1} \mathcal{A}$ *where the first element is a payoff parameter and the remaining elements correspond to a history of actions in periods* 0 *through* $t-1$. *The agent's type evolution would then be* $T(\theta_t, a_t, U) = 1\{\theta_{t+1} \in U\}$ *where* $\theta_{t+1} = (\theta, a_1, ..., a_t)$ *and* θ *is the payoff parameter.*

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*

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

[17] 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 [45] or van der Vaart [51] for a full development of these technical issues.

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

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 2, 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*

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

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 in period* t *is* $\theta = (v, \phi_0, ..., \phi_{t-1})$ *is decomposed into a payoff type,* v*, and a length* t *history of private signals,* $(\phi_0, ..., \phi_{t-1})$*, 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, the agent's 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.*

3.3 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 $\Psi_N = \prod_{t=1}^{\infty} \Omega \times \Theta^N$. We assume that all random variables in the N-agent games are measurable with respect to a filtration $\{\mathcal{F}_t^N\}_{t=0}^{\infty}$.

The initial distribution of types in the N agent game is generated by N independent and identically distributed draws from $\mu^{\Theta}(\omega_0)$. The following lemma is immediate from the uniform law of large numbers stated in Corollaries 3 and 4 in the Appendix.

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

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 $\{\theta_{t+1}^1, ..., \theta_{t+1}^N\}$, and the value of $\pi_{t+1}^{\mathcal{A}} \in \Delta_N(\mathcal{A})$ generated by the strategy σ. The transition probability function for the large finite game is defined for any $U \subset \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$ and $W \subset \mathcal{B}(\Omega)$

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

$P_N^A(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ can be thought of either as a condition measure over the space $\Theta^N \times \mathcal{A}^N \times \Omega$ or as a conditional measure over the space $\Delta_N(\Theta) \times \Delta_N(\mathcal{A}) \times \Omega$. 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 $f : \prod_{\tau=t}^{\infty} \Omega \times \Theta^N \to \mathcal{X}$, we use the notation $E_t^{\Psi}[f]$ to denote an expectation with respect to $\prod_{\tau=t}^{\infty} \Omega \times \Theta^N$ 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 $\sigma \in \Sigma$,[19] we define intertem-

[19]By *symmetric* we mean that all asymmetries are reflected in Θ. This accomodates, for example, buyers

poral utility at time t as

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

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

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

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 $\theta_0^i\in supp[\pi_0^{\Theta}]$*[20]

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

3.4 Evolution and Equilibrium - Nonatomic Dynamic Games

The nonatomic dynamic game structure admits a measure 1 continuum of agents with an initial type distribution equal to $\mu^{\Theta}(\omega_0)$. 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 $T(\circ|\theta_t,a_t,\pi_t^{\Theta},\pi_t^{\mathcal{A}})$. Intuitions based on the law of large numbers for a countable set of random variables suggest that idiosyncratic shocks in the evolution of a single agent's type are eliminated in the aggregate.[21] In the nonatomic limit game the distribution of

and sellers in a two sided market (see Section B.2) or entrants and nonentrants in a market competition model.

[20] The notation supp[π] refers to the support of the probability measure π.

[21] 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 part of the definition of the limit game. We prove in Theorem 1 that the dynamics of the large finite game approach the nonatomic limit

$(\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A})$ is in effect determined by the realizations of $(\omega_0, \omega_1, .., \omega_{t-1})$ in conjunction with the initial distribution of types, $\mu^\Theta(\omega_0)$, and the strategies of the agents.

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 $P^C(\omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma) = (\pi_{t+1}^\Theta, \pi_{t+1}^\mathcal{A})$ where for any $V \in \mathcal{B}(\Theta), U \in \mathcal{B}(\mathcal{A})$[22]

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

Where confusion does, we suppress the notation for the strategy on which P^C is conditioned. The crucial conceptual difference between the large finite game and the nonatomic limit game is that the market aggregates may 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.

The state space for the optimization problem facing a single agent in the nonatomic limit game is $\prod_{t=1}^\infty \Omega \times \Theta$ where the second component is the agent's own type, which remains a random variable at the individual level. 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 paths $\prod_{\tau=t}^\infty \Omega \times \Theta$. Given some strategy $\sigma \in \Sigma$, the transition probability function of the nonatomic game has the form for $U \subset \Delta(\Theta) \times \Delta(\mathcal{A}), W \in \mathcal{B}(\Omega)$,

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

We view this as a description of the (Markov) evolution of $(\omega_0, \omega_1, ..)$ with the distributions $(\pi_t^\Theta, \pi_t^\mathcal{A})$ determined in each period by $(\omega_0, \omega_1, .., \omega_t)$, the endogenous agent strategies, and the initial distribution of types $\mu^\Theta(\omega_0)$. The $\tau < \infty$ fold iteration of this operator is denoted $(P_\infty^\mathcal{A})^\tau(U \times W|\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A})$.

game dynamics as $N \to \infty$.

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

The intertemporal utility in the nonatomic dynamic game is

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

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 agent i's utility in the nonatomic limit game when all agents play the symmetric strategy σ is

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

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

$$V_{\infty}(\theta_t^i, \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}|(\sigma_i', \sigma_{-i})) = (1-\delta) * \{E_t^{\Omega}[w(\theta_t^i, \sigma_i'(\theta_t^i, \omega_t, \pi_t^{\Theta}), \omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})] + \delta E_t^{\Omega}\left[V(\theta_{t+1}^i, \omega_{t+1}, \pi_{t+1}^{\Theta}, \pi_{t+1}^{\mathcal{A}}|(\sigma_i', \sigma_{-i}))\right]\}$$

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 $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ 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 $(\sigma^{DCE}, \omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}}) \in \Sigma \times \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$ such that for all $\theta_0^i \in \Theta$*

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

A stationary equilibria is a nonatomic equilibrium wherein the economy has 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 $\pi_{\infty}^{\Theta} \in \Delta(\Theta)$ such that*

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

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

2. For all $\theta \in \Theta$ we have

$$V_\infty(\theta_0^i, \omega, \pi_\infty^\Theta, \pi_\infty^{\mathcal{A}}|\sigma^{SE}) + \varepsilon \geq \sup_{\sigma_i' \in \Sigma} V_\infty(\theta_0^i, \omega, \pi_\infty^\Theta, \pi_\infty^{\mathcal{A}}|\sigma_i', \sigma_{-i}^{DCE})$$

Condition (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 $(\omega, \pi_\infty^\Theta, \pi_\infty^{\mathcal{A}})$.[23] The usefulness of this equilibrium concept lies in the fact the equilibrium strategy is finite dimensional and hence computationally tractable.[24]

If an analyst attempts to estimate or compute the equilibrium of a complete information stochastic game, the state space 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 any 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.[25] Any DCE strategy $\sigma^{DCE} : \Theta \times \Omega \times \Delta(\Theta) \to \Delta(\mathcal{A})$ can be written using an *indirect description* of the form $\sigma_{ID}^{DCE} : \Theta \times \Omega^\infty \to \Delta(\mathcal{A})$.

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

[24]See Bergin and Bernhardt [10] for existence results on exact DCE or SE in nonatomic games.

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

4 Results

The proofs of the approximation results of Section 3.2, Theorems 2 and 3, require that V_∞ be uniformly continuous in $\Theta \times \mathcal{A} \times \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$. 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 uniform continuity of w and a uniformly continuous strategy σ, are sufficient to prove that V_N uniformly converges to V_∞ and that V_∞ is uniformly continuous as required.

4.1 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*

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

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

Convergence rates can be bounded using asymptotic results from empirical process theory. When we claim the convergence rate is $O(N^{-0.5})$ we are asserting 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^*. Two equivalent ways of stating the convergence rate are:

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

Both γ and ρ determine the accuracy of our approximation results, and we cannot in general determine which of γ or ρ (or both) are limiting the approximation precision. We choose the statement emphasizing the convergence of γ 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.

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 $(\theta, a, \omega) \in \Theta \times \mathcal{A} \times \Omega$*

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

4.2 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. Therefore 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.[26]

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

Theorem 2. *Fix $\varepsilon > 0$. Assume that σ_{ID}^{DCE} is uniformly continuous.*[27] *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 game for $N > N^*$ players. Furthermore, N^* can be chosen uniformly across $(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ if if $\pi_0^{\Theta} \in \Delta_N(\Theta)$ and $\pi_0^{\mathcal{A}} \in \Delta_N(\mathcal{A})$.*

Theorem 2 proves that when an analyst 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[28] of $\frac{t*(t-1)}{2} * dim(\Theta) * dim(\Omega)$, independent of the number of agents, to describe the equilibrium behavior for the future t periods.[29] If an analyst attempts to solve for the policy function of a complete information stochastic game with N agents, the state space of the model would be $\Theta \times \Omega \times \Delta_N(\Theta)$ and the dimension of the policy space diverges exponentially as $N \to \infty$. Therefore, computing the exact equilibrium policy function in large complete information games is, in practice, computationally impossible.[30] 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 rare events.

The convergence rate of ε is important for the practical application of our theorems. If the utility function (w) is Lipschitz continuous, then it is straightforward to show that ε shrinks at a rate of $O(N^{-1/2})$. 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 static private values double auctions have been shown to converge at a rate equal

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

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

[29] The dimension of the space Ω^t diverges to infinity as $t \to \infty$. Some model assumptions must be made to regulate behavior for large t since any estimation is based on histories of finite length.

[30] One 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.

to $O(N^{-2+\alpha})$ for any $\alpha > 0$ (Cripps and Swinkels [20]).

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.[31]

Corollary 1. *Consider a Stationary Equilibrium of the nonatomic limit game, $(\sigma^{SE}, \pi_\infty^\Theta)$. Assume that:*

- *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 any convergent sequence of exact BNE has a limit in the set of DCE. 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.

Let the set of exact BNE of the N-agent game with initial state $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ be denoted by the correspondence $\mathcal{E}(\circ|\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A}) : \mathbb{N} \rightrightarrows \Sigma$, and denote the set of exact DCE strategies of the nonatomic limit game with initial state $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ as $\mathcal{E}^{NA}(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$.

Theorem 3. *Assume that*

- *Θ and Ω are compact*
- *There exists an N^* such that $\cup_{N=N^*}^\infty \mathcal{E}(N|\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ is equicontinuous in $\Theta \times \Omega \times \Delta(\Theta)$*

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

$$\underset{N\to\infty}{Lim}\, \mathcal{E}(N|\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A}) = \mathcal{E}^\infty \subset \mathcal{E}^{NA}(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$$

where $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A}) \in \Omega \times \Delta_{N^*}(\Theta) \times \Delta_{N^*}(\mathcal{A})$.

[31] This result is similar to the claim that stationary equilibria possess the asymptotic Markov equilibrium property of Benkard et al. [7]. See Section 5.5 for a comparison of our conditions to those in Benkard et al.

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 $N \to \infty$.[32] 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 proving theorem 3 is the necessity of verifying the continuity of the limit of arbitrary convergent sequences of the form $\{\sigma^N\}_{N=1}^{\infty}$, $\sigma^N \in \mathcal{E}(N|\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$).[33] We focus attention in our examples on cases where these restrictions are weak. For example, a game wherein $\Theta \times \Omega$ is discrete and $\Delta(\Theta)$ is unobservable has a uniformly equicontinuous strategy set. Applications with natural discontinuities such as games with entry and exit or convex investment costs require careful analysis, but these models can often be studied with our approximation framework (see Section 5.5).

If we do not assume the equicontinuity of $\cup_{N=N^*}^{\infty} \mathcal{E}(N|\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$), it is not clear that the limit $\sigma^N \to \sigma^{\infty}$ will be continuous in $\Theta \times \Omega \times \Delta(\Theta)$.[34] 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 2. *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} *in periods* 1 *through* τ. *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^{∞} *such that...*[35]

[32] Examples of cases in which ε-BNE 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 [46].

[33] 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.

[34] 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.

[35] Since our proof studies deviations of agent strategy that occur within a finite horizon, for any σ^{∞} we

- *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. These discontinuities may allow the agents to remain pivotal over outcomes even in economies where the number of agents is large.

5 Extensions and Applications

5.1 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 often complex and imperfect (but see Compte [19] and McLean et al. [40]). 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 result turns on the fact that continuous large games sharply limit the capacity for agents to detect defection from tacit collusion.[36] Folk theorems may be interpreted as providing the pessimistic result that sharp predictions are impossible in those games that admit such a result. Our result suggests that, at least in the context of games with many players, imperfect monitoring places a sharp constraint on the classes of admissible equilibria. Furthermore, our theorem 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

can define some $\tau^* < \infty$ such that we only require continuity over an open set U in $\prod_{t=0}^{\tau^*} \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$ containing Λ^{τ^*}.

[36] Most folk theorem results (for example Fudenberg et al. [23]) are inapplicable if this condition fails.

well as the assumptions required to generate predictions. 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 $t-1$.[37] An agent's type $\theta_t = (\phi_0, \phi_1, ..., \phi_{t-1})$ is 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 signal in period t is

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

which we assume is uniformly 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 of signals between 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.[38]

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$[\sigma^{DCE}(\theta)]$

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

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

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

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

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 $\delta < 1$, consider any convergent sequence of strategies $\{\sigma^N\}_{N=1}^{\infty}, \sigma^N \in \mathcal{E}(N)$. As $N \to \infty$, $\sigma^N \to \sigma^\infty$ where σ^∞ is a repeated static Nash equilibrium strategy.*

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 are the only equilibria of the nonatomic limit game. Our theorem immediately follows. □

Our result is driven by the infinitesimal effect of one agent's action on the signals observed by the other agents. Folk theorems typically generate results for the case wherein N is held fixed and $\delta \to 1$. 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 incentives for tacit collusion.

5.2 Large Games of Imperfect Public Monitoring

We focus on public equilibria[39] 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 $\rho(y_t|\pi_{t-1}^{\mathcal{A}})$. We assume that Y is finite and the aggregate state space is $\Omega = \bigcup_{t=0}^{\infty} Y^t$. 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 discreteness 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

[39] We do not impose perfection at this point, but our results carry over to perfect public equilibria (see Section 5.4).

Theorem 5. *Assume that ρ is continuous and $\delta < 1$ and consider any convergent sequence of strategies $\{\sigma^N\}_{N=1}^{\infty}, \sigma^N \in \mathcal{E}(N)$. As $N \to \infty$, $\sigma^N \to \sigma^\infty$ where σ^∞ is a repeated static Nash equilibrium strategy.*

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 2. Consider an N-agent Prisoner's Dilemma game. Let the action space be $\mathcal{A} = \{c, d\}$ and the public signal space be $Y = \{0, 1\}$. $\omega \in \Omega$ denotes an infinite history of public signals, while Θ denotes a privately observed randomization device drawn at $t = 0$ independently across agents. Let $f_t = \pi_t^{\mathcal{A}}[\{c\}]$ denote the fraction of agents 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 2 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 asymmetries mediated by the private information, Θ. The agents could also oscillate between the repeated static Nash equilibrium outcome ($f_t = 0$) and collusive behavior ($f_t = \frac{1}{2}$).

Corollary 2 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 noise 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); provide informative, noise-free private signals (T); and tightly control the aggregate noise present in the environment (G).

5.3 Games With a Few Large Players

There are many markets 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, and 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. [24], Levine and Pesendorfer [38]). 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 $N \to \infty$ while M remains fixed. In contrast, each of the large agents is treated as a measure 1 atom.[40]

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 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 $\omega_0 = \varnothing$.[41] Let the utility functions of the large players be denoted $\{u_m : \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})\}_{m=1}^{M}$.[42]

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 equilibrium set of the nonatomic agents in the 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

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

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

[42] We could extend our framework to employ utility function for the large players that change as $N \to \infty$, but we would require much additional notation as well as an analog of assumption 4 to insure the utility functions for the large players in the limit game are well defined.

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. □

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

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

- *The strategy of the small players, 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. If the utility functions of the large agents are continuous in the limit game, then a game-theoretic equilibrium strategy vector for the large players in the limit game will produce an $\varepsilon-$equilibrium strategy vector for the large players in the N-agent game for sufficiently large N.

5.4 Perfection in Games of Complete Information

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 4. *A symmetric $\varepsilon-$ Perfect Equilibrium ($\varepsilon-$PE) is a strategy $\sigma^{PE} \in \Sigma$ 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 5. *A symmetric $\varepsilon-$Perfect Competitive Equilibrium (ε-PCE) consists of a strategy $\sigma^{PCE} \in \Sigma$ such that for all $(\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}) \in \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$ and $\theta_t^i \in supp[\pi_t^\Theta]$*

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

We now generalize both Theorem 2 and Theorem 3 to the case where perfection is required.

Theorem 8. *Fix $\varepsilon > 0$. Assume that σ^{PCE} is uniformly continuous. Then we can choose N^*such that the PCE strategy σ^{PCE} is an ε- Perfect Equilibrium of the large finite game for $N > N^*$ players.*

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

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

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 with initial state $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ as $\mathcal{E}^P(N|\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ and the set of perfect competitive equilibrium strategies with initial state $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ as $\mathcal{E}^{PCE}(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$.

Theorem 9. *Assume that*

- *Θ and Ω are compact*
- *There exists an N^* such that $\cup_{N=N^*}^\infty \mathcal{E}^P(N|\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ is equicontinuous in $\Theta \times \Omega \times \Delta(\Theta)$*

Then the correspondence $\mathcal{E}^P(\circ|\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$ is upper hemicontinuous with

$$\underset{N\to\infty}{Lim} \mathcal{E}^P(N|\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A}) = \mathcal{E}^\infty \subset \mathcal{E}^{PCE}(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$$

where $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A}) \in \Omega \times \Delta_{N^*}(\Theta) \times \Delta_{N^*}(\mathcal{A})$.

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

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

By applying Theorem 3 to each element of this collection, we have our result. Note that since Theorem 1 and its corollaries imply uniform convergence, the convergence across this collection of DCE is uniform. □

For games of incomplete information, one would hope 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). Define *symmetric consistency* as consistency with respect to limits of symmetric, completely mixed strategies.[43] 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.

5.5 Entry and Exit Games and Discontinuities

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].[44] In this setting the entry and exit actions imply discontinuities in agent strategies. We focus our discussion on applying Theorem 1 in this setting as this is the natural extension of Theorem 5.4 of Benkard et al.

In the model of Benkard et al. [7], 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 of Benkard et al. [7]. *Uniform* continuity of w is assured by the bounded derivatives of w in assumption 5.1 of Benkard et al. [7]. Continuity of T and our assumption 2 are obtained from assumption 3.2 of Benkard et al. [7]. 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

[43]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.

[44]Convex investment costs as in Caballero and Engel [14] provide another important context with natural discontinuities.

mixed strategy, the firms mix between "Enter" and "Stay Out" with a probability that is conditional on the state of the economy, π_t^{Θ}. 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.[45]

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 is 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 Theorem 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 for each x_{it} the discontinuity is nongeneric. Denote $\pi^{\Theta} \in \Delta(\Theta)$ as *admissible* for a strategy $\sigma \in \Sigma$

[45] 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. Therefore the assumptions of our Theorem 1 hold and the oblivious equilibria Benkard et al. study are $\varepsilon-$Bayesian Nash equilibria.

if there is an admissible distribution $\widetilde{\pi}^{\Theta} \in \Delta(\Theta)$ such that for any $U \in \mathcal{B}(\Theta)$[46]

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

Assumption 6. *Consider any $\varepsilon > 0$. 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}.*

The definition of admissibility can be circular and endogenous, which would limit the usefulness of our results. However, we can typically verify that assumption 6 holds for all measures $\pi^{\Theta} \in \Delta(\Theta)$ except for some measures that are clearly not admissible. We demonstrate this application of assumption 6 in the discussion of Benkard et al. [7] at the end of this section.

For what follows, we employ an indirect representation of P^C as follows

$$\begin{aligned} \pi^{\mathcal{A}}[U] &= \int_{\Theta} \sigma(\theta, \omega_0, \pi^{\Theta})[U] * \pi^{\Theta}[d\theta] \text{ for any } U \in \mathcal{B}(\mathcal{A}) \\ P_{ID}^{C}(\circ|\pi^{\Theta}, \omega^t) &= (P^C)^{t+1}(\omega_0, \pi^{\Theta}, \pi^{\mathcal{A}}) \end{aligned}$$

We suppress the dependence of P_{ID}^C on σ throughout. The indirect representation is useful in that we no longer to specify admissible $\pi^{\mathcal{A}}$, but is restricted in that it does not describe evolution of the nonatomic limit game from any $(\omega, \pi^{\Theta}, \pi^{\mathcal{A}})$. It is straightforward to extend lemma 3, which proves that P_{ID}^C is continuous when σ_{ID}^{DCE} is continuous, to the case where the nonatomic limit game obeys assumption 6. Given this extension of lemma 3, proving analogs of Theorems 2 and 3 are immediate.[47]

Lemma 2. *Assume T is uniformly continuous, hold σ_{ID}^{DCE} fixed, and assumption 6 holds. For any $\varepsilon > 0$, we can find a $\gamma > 0$ such that if π^{Θ} and $\widetilde{\pi}^{\Theta}$ are admissible and $d_{LP}^{\Theta}(\pi^{\Theta}, \widetilde{\pi}^{\Theta}) < \gamma$, then $d_{LP}^{\Theta\times\mathcal{A}}(P_{ID}^C(\circ|\pi^{\Theta}, \omega^t), P_{ID}^C(\circ|\widetilde{\pi}^{\Theta}, \omega^t)) < \varepsilon$ uniformly over $\pi^{\Theta} \in \Delta(\Theta)$ for any sequence of aggregate states $\omega^t = (\omega_0, ..., \omega_t)$.*

Lemma 2 proves that well behaved discontinuities of the equilibrium strategy, strategies that satisfy assumption 6, do not cause the evolution of the nonatomic equilibrium (and hence the value function for the equilibrium) to be discontinuous. Using this result we can

46 Admissiblity is a requirement placed on π^{Θ} realized in the nonatomic limit game.
47 One need only replace references to lemma 3 with references to lemma 2.

replace the use of Lemma 3 with Lemma 2 in our proofs to generate an analog to Theorem 1, which allows us to state the following approximation theorem.

Theorem 10. *Fix $\varepsilon > 0$. Assume that:*

- σ_{ID}^{DCE} *obeys assumption 6*
- 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 distribution of exit payoffs is nonatomics, which implies that in any admissible π_y^{Θ} only a measure 0 of the agents are indifferent between choosing "Exit" and remaining in the industry. Therefore the exit strategy clearly obeys assumption 6. Theorem 10 implies that oblivious equilibria are ε-Nash equilibria. By representing a Perfect Competitive Equilibrium (see Section 5.4) 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 5.3) and assert that the DCE are proof against finite coalitions (see Section B.1) without loss of generality.

Finally, we can under certain conditions make the stringer statement that convergent sequences of equilibrium strategies from the large finite game converge to equilibrium strategies of the nontatomic limit game. The following is an analog to Theorem 3 stated without the compactness assumptions that ensure σ^{∞} is uniformly continuous.[48]

Theorem 11. *Assume that there exists an N^* such that $\cup_{N=N^*}^{\infty}\mathcal{E}(N|\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ is uniformly equicontinuous[49] in $\Theta \times \Omega \times \Delta(\Theta)$. Consider any convergent sequence of strategies*

[48]If one were to assume that Θ is compact, then we would not need to assume that σ^{∞} is uniformly continuous. This compactness assumption is not typically made in the structure of Ericson and Pakes [22].

[49]We assume only that the family of functions $\cup_{N=N^*}^{\infty}\mathcal{E}(N|\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ satisfy the same modulus of continuity over their respective domains. If the functions are uniformly equicontinuous with modulus of continuity $\delta(\varepsilon)$, then for any $\sigma \in \mathcal{E}(N|\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$, $N \geq N^*$, 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

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

$\{\sigma^N\}_{N=1}^{\infty}, \sigma^N \in \mathcal{E}(N)$. *As* $N \to \infty$, $\sigma^N \to \sigma^\infty$ *where* $\sigma^\infty \in \mathcal{E}^{NA}$.

Theorem 10 implies that the DCE strategies satisfy the exact requirements of a 0-BNE only asymptotically, whereas theorem 11 proves that any strategy that satisfies the 0-BNE conditions for a large N must be close to a 0-DCE strategy. One might conjecture the stronger conclusion of theorem 11 followed from the asymptotic satisfaction of equilibrium conditions, but in fact we require continuity conditions on the equilibria of the large finite game to insure that this result holds.

6 Large, Dynamic Auction Markets

In this section we elaborate on the discussion of section 2 and more formally describe a large finite model of the E-Bay platform that yields the nonatomic model of Satterthwaite and Shneyerov [48] as a limit approximation.[50] The game takes place in discrete periods of length δ indexed $t \in \mathbb{Z}$. Assume in the N agent game that S_N sellers and B_N buyers have the option of entering the online auction. Each seller is endowed with a single unit of a homogenous good with an opportunity cost of sale $c \in [0, 1]$ independently drawn from F_S, while the buyers have valuation $v \in [0, 1]$ for a single unit drawn from F_B. In this setting, there are no aggregate shocks and the agent types are fixed over time, so T and G are continuous. Furthermore, the agent type space $[0, 1]$ is compact.

Should the agents enter they suffer a participation cost equal to $\delta\kappa$. In each period the sellers are ordered randomly and sequentially matched with $k \in \mathbb{N}$ buyers with probability

$$\pi_k = \frac{\zeta^k}{k! e^\zeta}$$

where $\zeta = \frac{B_N}{S_N}$. Should the pool of buyers be depleted, all subsequent sellers are unmatched. Symmetrically, it may be the case that buyers remain unmatched. We assume that these matches are anonymous and so buyers and sellers cannot condition on their prior history with the particular auction participants.

In each auction, each buyer simultanenously submits a bid $B(v)$ while at the same time the sellers each name a reservation price $S(c)$. All bids are submitted prior to knowing the number of other bidders in the auction. Within each auction if any bids exceed the seller's reservation price, then the highest bidder obtains the good at the price bid. The

[50] The reader is referred to Satterthwaite and Shneyerov [48] for details.

seller and the winning bidder then exit the auction and receive the following payoffs

$$\begin{aligned} \text{Buyer:} \quad & v - p \\ \text{Seller:} \quad & p - c \end{aligned}$$

Remaining buyers (or in the case of no winner, the seller and all buyers) continue to the next period. We assume that traders discount all utility by $\beta\delta$ each period, where $\beta \geq 0$.

The agent strategies are composed of two decisions: whether to enter and then how much to bid (or name as the reservation price). We focus on symmetric cutoff strategies for the entry decision and stationary, symmetric bidding strategies for the agents. Denote the entry decision for the buyers (sellers) as $\varkappa_N^B : [0,1] \to \{\text{Enter, Stay Out}\}$ ($\varkappa_N^S[0,1] \to \{\text{Enter, Stay Out}\}$). Agents enter if and only if the expected utility of entry is weakly positive. If we assume that the expected utility of entering is strictly increasing in type for buyers and sellers, then assumption 6 of Section 5.5 holds. Alternately, we could assume that idiosyncratic sunk costs of entry are drawn, in which case assumption 6 is immediate. Denote the equilibrium bidding strategies of the buyer (seller) as $p_N^B : [0,1] \to [0,1]$ ($r_N^S : [0,1] \to [0,1]$). The buyer (seller) equilibrium strategies in the nonatomic limit game are denoted $\varkappa_\infty^B$ and p_∞^B ($\varkappa_\infty^S$ and r_∞^S).

We must demonstrate that agent utility in each period converges uniformly. We employ an outcome equivalent model in which all agents declare bids that are implemented only in the event that they are matched. Thus w_N represents the expected utility from stochastically being matched each period. To show uniform convergence, first we must assume that $\zeta = \frac{B_N}{S_N} \to a \in (0, \infty)$ as $N \to \infty$. Given the convergence of ζ, note that $\pi_k = \frac{\zeta^k}{k! e^\zeta}$ converges uniformly over all $k \in \mathbb{N}$ from the Glivenko-Cantelli lemma. Therefore we have that $w_N \to w$ uniformly.[51]

To apply theorem 2, we are required to assume that the bidding strategies, r_∞^S and p_∞^B, obey assumption 6. Lemma 5 of Satterthwaite and Shneyerov [48] proves that seller bidding functions are continuous. Lemma 4 of Satterthwaite and Shneyerov [48] proves the buyer's bidding function in the nonatomic limit game is strictly increasing (and hence there are at most a countable set of discontinuities), and since the buyer valuation distribution is nonatomic assumption 6 holds for p_∞^B.

To apply Theorem 11, we are required to consider a sequence $\{(p_N^B, r_N^S)\}_{N=N^*}^\infty$ that converges to a continuous strategy (p_∞^B, r_∞^S). Given this additional assumption, we have

[51]Although stated informally, the claim of uniform convergence is straightforward to prove.

that p_N^B and r_N^S converge uniformly to equilibrium bidding functions of the nonatomic limit game. We can then use Satterthwaite and Shneyerov's Theorems 1 and 2 to show that outcomes of any Bayesian-Nash equilibrium of the large finite game must converge to the Walrasian price with positive trade for β sufficiently small as $\delta \to 0$.

Given our convergence results, empirical auction studies could use the competitive Satterthwaite and Shneyerov [48] model as a benchmark for estimating the structural parameters of the E-Bay system. This could provide insights into the significance of market frictions (e.g. κ) that participants suffer. One could assess whether the buyers and sellers respond to continuation values as predicted, a behavioral anomaly that would be difficult (if not impossible) to isolate and study without a structural model.[52]

Given the freedom to solve for counterfactual equilibria computationally, one might extend the Satterthwaite and Shneyerov [48] model to include:

- Affiliated private values or interdependent values
- Selection into auctions by buyers
- "Buy It Now" sale prices
- Policies that eliminate or cap reservation values
- Large buyers or sellers who are continuously active in the market
- Imperfectly persistent buyer or seller valuations
- Aggregate shocks to supply or demand
- Seasonality or time-of-day effects

Given the challenge of solving a dynamic two-sided auction model with hundreds (or thousands) of simultaneous participants, these extensions and policy questions can only be addressed by solving and interpreting counterfactuals using techniques such as the approximation framework we propose.

[52]Given the evidence that bidders behave potentially irrationally in E-Bay auctions (e.g. Malmendier and Szeidl [42]), this outcome would not be (entirely) surprising. Computing the effect on bidder welfare and seller revenues remains an interesting subject for future work.

7 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 [25], and later Sabourian [47], provide sufficient conditions for the set of nonatomic subgame perfect equilibria (SPE) and the SPE of an analogous large finite repeated game to converge. Our contribution to this vein of literature is two-fold. First, Green [25] and Sabourian [47] study repeated games without private information. We analyze a broad class of games with 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 analyses of the model of industry competition formulated by Ericson and Pakes [22] using the notion of an oblivious equilibrium that is similar to our DCE. 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 Benkard 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. Adlakha et al. apply their results to a number of dynamic market games to analyze the structure of the stationary equilibria of models of large industries. Finally, the unpublished contemporaneous work Yang [53] develops a similar result to theorems 1 and 2 in a special case of our setting with complete information, no aggregate states, and a countable set of idiosyncratic shocks.

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 except under suitable continuity restrictions approximate equilibria need not resemble any exact equilibria, even when the equilibria become exact asymptotically. 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 in the strategy space to exact equilibrium policy functions. Third, even in the case where an approximate equilibrium faithfully approximates 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.[53]

A related literature on large games studies whether individual agents remain pivotal as the total number of agents grows (Fudenberg, Levine and Pesendorfer [24]; Al-Najjar and Smordinsky [4]; Al-Najjar [3], [2]). One 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.[54]

The prior literature on the relationship between large finite and nonatomic games generally focuses on games of complete information and does not emphasize the use of nonatomic games as a framework for analyzing the equilibria of game-theoretic models (examples include Housman [30], Khan and Sun [34], Carmona [15], and Carmona and Podczeck [16] and [17] amongst others). In addition some of these works focus on weaker approximate equilibrium concepts (e.g. the $(\varepsilon, \eta)-$equilibrium used by Carmona and Podczeck [16]), and it is not clear that these concepts yield our uniform equilibrium convergence results.

[53]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. 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 equilibria or a failure of lower hemicontinuity as $N \to \infty$.

[54]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.

Our work also relates to the work on convergence (or lack thereof) of markets to perfect competition as sources of friction disappear (Wolinsky [52], Jovanovic and Rosenthal [33], Hopenhayn [29], Serrano [49], Satterthwaite and Shneyerov [48]). 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.

A growing literature employs the techniques of game-theoretic analysis to analyze the performance of large markets (Pesendorfer and Swinkels [44], Swinkels [50], Cripps and Swinkels [20], McLean and Postlewaite [41], Budish [13], Kojima and Pathak [36], and Fudenberg et al. [24] amongst many others). Most of the works in this literature proceed by studying the properties of game-theoretic equilibria as the number of agents increases to infinity.

The large literature on Walrasian games addresses the question of whether finite games taken to represent strategic foundations for Walrasian equilibrium in fact achieve Walrasian equilibria in the limit as the number of agents increases (some examples of this literature include Hildenbrand [27] and [28], Roberts and Postlewaite [46], Otani and Sicilian [43], Jackson and Manelli [32]). In these games the agents in the economy declare a demand schedule, and the market price is defined by the price that clears the market given the aggregate declared demand. Jackson and Manelli provide the analysis closest in spirit to this paper. Jackson and Manelli prove that when agents believe they have little influence on the market-clearing price, then the agents' declared demands must converge to the Walrasian equilibrium demands as the economy grows in size. Our continuity condition, based on the exogenous economic primitives, is similar to the endogenous small influence condition of Jackson and Manelli.

Kubler and Schmedders [37] 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.

Macroeconomics papers use models and equilibrium concepts that are close kin to techniques used in our work. Chari and Kehoe [18] 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. Many of these studies can be interpreted as using a nonatomic model as an 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.

8 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 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 (e.g. 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 computation and estimation of structural models 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 equilibrium policy function can be described as a map from aggregate states and an individual agent's type to the action space. As a result the policy functions do not suffer from the curse of dimensionality. We 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 bound on the equilibrium outcomes of the large finite game.

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 [25], Sabourian [47], and Fudenberg, Levine and Pesendorfer [24] 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.

Once we establish our benchmark approximation theorems for the Bayesian-Nash equilibrium correspondence of the large finite games, we are able to provide a number of important extensions. Our first extension applies our theorems to games with a fixed set of large players that interact strategically with a competitive fringe of small players. This extension provides tools to simplify the analysis of models of industries with a large number of firms dominated by a small set of oligopolists as well as macroeconomic models with policy makers as large players. Our second primary extension is to approximate the set of perfect equilibria of complete information stochastic games. This allows us to approximate the set of Markov perfect equilibria, which are commonly employed in industrial organization models. In the appendix we provide extensions to the analysis of coalition-proof equilibria (Section B.1) and to models with multiple player roles and asynchronous actions (Section B.2).

Our final extension applies our techniques to the study of games with natural discontinuities. Our framework can be applied to models of entry and exit (e.g. Ericson and Pakes [22]) and models with convex investment costs that generate lumpy investment (e.g. Caballero and Engel [14]). In addition to demonstrating techniques for applying our ideas to important classes of computationally difficult models, these extensions show the robustness and generality of our approach.

As both a natural application and an example of how to apply our framework, we employ the nonatomic model of a decentralized auction market developed by Satterthwaite and Shneyerov [48] as a limit approximation of an online auction market such as E-Bay. We demonstrate that the model of Satterthwaite and Shneyerov [48] obeys the convergence and continuity conditions required to apply theorems 10 and 11. These theorems allow us to translate the results of Satterthwaite and Shneyerov [48] regarding the effects of market frictions on economics outcomes and convergence to Walrasian prices into statements about the efficiency properties and price stucture of the market platform. We discuss possible extensions and empirical results that could be discovered by calibrating the model of Satterthwaite and Shneyerov [48] and analyzing counterfactual market structures and sets of economic primitives.

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 5.4). Extending our results to refined equilibria of games with incomplete information requires formally defining these refinements in a nonatomic setting. Chari and Kehoe [18] provide one example of such an extension by placing 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 particular 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 5.3) to analyze structural models with a small number of dominant firms and a competitive fringe of small firms. Examining such hybrid models further using our framework, 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] Budish, E. (2008) "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes," *mimeo.*

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

[15] Carmona, G. (2004) "Nash Equilibria of Games with a Continuum of Players," *mimeo.*

[16] Carmona, G. and K. Podczeck (2009) "On the Existence of Pure-Strategy Equilibria in Large Games," *Journal of Economic Theory*, p.1300 - 1319.

[17] Carmona, G. and K. Podczeck (2011) "Approximation and Characterization of Nash Equilibria of Large Games," *mimeo.*

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

[19] Compte, O. (1998) "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66, pp. 597 - 626.

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

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

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

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

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

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

[26] Hendricks, K., and R. Porter (2007) "A Survey of Empirical Work in Auctions," in *Handbook of Industrial Organization*, vol. III: North-Holland.

[27] Hildenbrand, W. (1974) "On Economies with Many Agents," *Journal of Economic Theory*, 2, pp. 161-188.

[28] Hildenbrand, W. (1974) *Core and Equilibria of a Large Economy*, Princeton University Press: Princeton, New Jersey.

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

[30] Housman, D. (1988) "Infinite Player Noncooperative Games and the Continuity of the Nash Equilibrium Correspondence," *Mathematics of Operations Research*, 13, pp. 488 - 496.

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

[32] Jackson, M. and A. Manelli (1997) "Approximately Competitive Equilibria in Large Finite Economies," *Journal of Economic Theory*, 77, pp. 354 - 376.

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

[34] Khan, M. and Y. Sun (1997) "Non-Cooperative Games with Many Players," *Handbook of Game Theory with Economic Applications, Vol. 3*, North Holland: New York.

[35] Khan, M. and Y. Sun (1999) "Non-Cooperative Games with Many Players," *Journal of Mathematical Economics,* 31, pp. 455-492.

[36] Kojima, F. and P. Pathak (2007) "Incentives and Stability in Large Two Sided Matching Markets," *American Economic Review*, 99, pp. 608 - 627.

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

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

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

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

[41] McLean, R. and A. Postlewaite (2002) "Informational Size and Incentive Compatibility," *Econometrica*, 70, pp. 2421-2453.

[42] Malmendier, U. and A. Szeidl (2008) "Fishing for Fools," *mimeo.*

[43] Otani, Y. and J. Sicilian (1990) "Limit Properties of Equilibrium Allocations of Walrasian Strategic Games," *Journal of Economic Theory*, 51, pp. 295-312.

[44] Pesendorfer, W. and J. Swinkels (2000) "Efficiency and Information Aggregation in Auctions," *American Economic Review*, 90, pp. 499-525.

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

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

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

[48] 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.

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

[50] Swinkels, J. (2001) "Efficiency of Large Private Value Auctions," *Econometrica*, 69, pp. 37-68.

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

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

[53] Yang, J. (2011) "Connections between Finite- and Infinite-player Games: Normal- and Extended-form Analyses," *mimeo.*

A Proofs

We begin this section with some useful results from the theory of the weak convergence of empirical processes.

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*

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

Proof. (Proof of Theorem 12) Follows from van der Vaart et al. [51], p. 135 and noting that sets of the form $\{x : x \leq y\}$ for $y \in \mathbb{R}^d$ are lower contours and form a VC Class. □

Corollary 3. *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 3) 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 4. *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 4) This result follows directly from Theorems 2.6.7 and 2.14.9 of van der Vaart and Wellner [51]. □

A.1 Proofs from Section 3

We begin this section by proving that the nonatomic model transition operator

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

and the $\tau < \infty$ step ahead iterations of this operator are continuous. Standard proofs of continuity are contained in Appendix C.

Lemma 3. *If T, σ are (uniformly) continuous, then $(P^C)^\tau$ is (uniformly) continuous in $\Delta(\Theta) \times \Delta(\mathcal{A})$ for any $\tau < \infty$ conditional on any sequence of aggregate states $(\omega_{t+1}, ..., \omega_{t+\tau})$*

Proof. Contained in Appendix C. □

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

Proof. Contained in Appendix C. □

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 $P^C(\omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})$ as $N \to \infty$.

Lemma 5. *Fix $\gamma > 0$ 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 $(\omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})$ where π_t^Θ is a tight probability measure we have*

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

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

Proof. Our first step will be to approximate the measure π_t^Θ using a counting measure with a finite number of atoms. We can then apply the uniform law of large numbers to each one of these finite atoms to obtain uniform convergence of the transitions of the approximating measure as $N \to \infty$. From the uniform continuity of P^C (lemma 3), 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 $U \subset \Theta_t$ such that $\pi_t^\Theta[U] > 1 - \gamma$. Let $N(\theta, \gamma) = \{\theta' \in \Theta_t : d_\Theta(\theta, \theta') < \frac{\gamma}{2}\}$ be a $\frac{\gamma}{2}$-radius open neighborhood centered on $\theta \in \Theta_t$. 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 $C_k = N(\theta_k, \gamma)$ and $\overline{C}_k = C_k - \cup_{j>k} C_j$. Therefore, $\cup_{k=1}^K \overline{C}_k = U$ and $i \neq j$ implies $\overline{C}_i \cap \overline{C}_j = \varnothing$.

We create a γ approximation of π_t^Θ by defining $\widetilde{\pi}_t^\Theta \in \Delta_K(\Theta)$ where $\widetilde{\pi}_t^\Theta[\Theta_t] = 1$

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

where δ_{θ_k} denotes a measure 1 atom on θ_k.[55] Note that $d_{LP}^\Theta(\widetilde{\pi}_t^\Theta, \pi_t^\Theta) < \gamma$ since for any

[55] $\widetilde{\pi}_t^\Theta$ is not tight.

$A \in \mathcal{B}(\Theta)$ we have

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

where A^{γ} is a γ neighborhood of A. Note that

$$\alpha_k * 1\{A \cap C_k \neq \varnothing\} > \pi_t^{\Theta}[A \cap \overline{C}_k]$$

so $\widetilde{\pi}_t^{\Theta}[A^{\gamma}] + \gamma \geq \pi_t^{\Theta}[A]$ as required. To obtain the converse inequality, $\pi_t^{\Theta}[A^{\gamma}] + \gamma \geq \widetilde{\pi}_t^{\Theta}[A]$, we must show

$$\pi_t^{\Theta}[A^{\gamma}] = \sum_{k=1}^{K} \pi_t^{\Theta}[A^{\gamma} \cap \overline{C}_k] + \gamma \geq \sum_{k=1}^{K} \alpha_k * 1\{A \cap C_k \neq \varnothing\}$$

Note that $\{A \cap C_k \neq \varnothing\}$ implies that $\overline{C}_k \subseteq A^{\gamma}$ since $\overline{C}_k$ has diameter at most γ. Therefore $\pi_t^{\Theta}[A^{\gamma} \cap \overline{C}_k] = \alpha_k$, and our inequality holds (without the added γ term in fact).

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 $\pi_{t+1}^{\mathcal{A}}$.

Let $\mathcal{S}(\widetilde{\pi}_t^{\Theta})$ denote the support of $\widetilde{\pi}_t^{\Theta}$. Considered as an empirical distribution of types in the N-agent game,[56] $\widetilde{\pi}_t^{\Theta}$ has at least $N * \widetilde{\pi}_t^{\Theta}(\theta^*)$ realizations for each $\theta^* \in \mathcal{S}(\widetilde{\pi}_t^{\Theta})$. If $\widetilde{\pi}_t^{\Theta}(\theta^*) = 0$, then $\theta \in N(\theta^*, \gamma)$ are irrelevant for determining either π_{t+1}^{Θ} or $\widetilde{\pi}_{t+1}^{\Theta}$. Therefore assume $\widetilde{\pi}_t^{\Theta}(\theta^*) > 0$. Each θ^* transitions to a type in the next period according to the probability measure

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

[56] To avoid integer problems in this discussion, we could assume α_k is a rational number of the form $\alpha_k = \frac{a_k}{b_k}$ and choose N in multiples of $b_1 * b_2 * ... * b_{\overline{K}(\gamma)}$. To explicitly avoid these integer problems we could consider an ancillary approximation $\widehat{\pi}_t^{\Theta} \in \Delta_N(\Theta)$ that can be chosen to be arbitrarily close to $\widetilde{\pi}_t^{\Theta}$ in the limit as $N \to \infty$ since $\Delta_N(\Theta)$ is dense in $\Delta(\Theta)$ in this limit. This aspect of the proof is straightforward, so the argument is omitted.

Denote the CDF associated with the measure $P_1(\theta^*, \circ)$ as $J(\circ|\theta^*)$. Let $J_N(\circ|\theta^*)$ represent an empirical CDF of $N * \widetilde{\pi}_t^\Theta(\theta^*)$ realizations of $J(\circ|\theta^*)$. From Corollary 4, $J_N(\circ|\theta^*) \to J(\circ|\theta^*)$ almost surely in the weak-* topology 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 weakly to the true distribution defined by $P_1(\theta^*, \circ)$. 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 and uniformly to

$$\widetilde{\pi}_C^\Theta[\circ] = \sum_{\theta^* \in \mathcal{S}(\widetilde{\pi}^\Theta)} \widetilde{\pi}_t^\Theta[\theta^*] * P_1(\theta^*, \circ)$$

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

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

Note that $P^C(\omega_{t+1}, \widetilde{\pi}_t^\Theta, \pi_t^\mathcal{A}) = (\widetilde{\pi}_C^\Theta, \widetilde{\pi}_C^\mathcal{A})$.

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.[57] Therefore, we can choose $\gamma > 0$ (and hence large enough K) with probability[58] at least $1 - \rho$

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

From the continuity of P^C (lemma 3), we have for small enough γ that

$$\begin{aligned} d_{LP}^\Theta(\pi_C^\Theta, \widetilde{\pi}_C^\Theta) &< \frac{\gamma}{2} \\ d_{LP}^\mathcal{A}(\pi_C^\mathcal{A}, \widetilde{\pi}_C^\mathcal{A}) &< \frac{\gamma}{2} \end{aligned}$$

[57] Note that π_{t+1}^Θ and $\pi_{t+1}^\mathcal{A}$ are N-agent empirical distributions, so we can treat π_{t+1}^Θ and $\pi_{t+1}^\mathcal{A}$ as the realization of N conditionally independent draws from Θ and $\mathcal{A}$. The continuity of π_{t+1}^Θ and $\pi_{t+1}^\mathcal{A}$ in π_t^Θ and $\pi_t^\mathcal{A}$ follows immediately.

[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 $1 - \rho$.

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

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

□

We will now show that if the game starts at a tight measure π^{Θ}_0, then the evolution of the nonatomic game necessitates that the measure of types remain tight over any finite horizon.

Lemma 6. *Suppose that π^{Θ}_t is a tight measure and let $P^C(\omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t) = (\pi^{\Theta}_{t+1}, \pi^{\mathcal{A}}_{t+1})$. Then π^{Θ}_{t+1} is tight.*

Proof. Since π^{Θ}_t 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 5, which we denote $\widetilde{\pi}^{\Theta}_t$. Since T is tight, each of these finite atoms yields a tight measure in period $t+1$. Denote this step-ahead measure $\widetilde{\pi}^{\Theta}_{t+1}$. Since P^C is continuous in $\Delta(\Theta)$ by lemma 3 and $\widetilde{\pi}^{\Theta}_t$ is an arbitrarily close approximation of π^{Θ}_t, we have that the (tight) measure $\widetilde{\pi}^{\Theta}_{t+1}$ can be chosen to be arbitrarily close to π^{Θ}_{t+1}. Therefore π^{Θ}_{t+1} is tight. □

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

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

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

Proof. We first prove that for any we can choose N^* such that for all $N > N^*$

$$d^{\Omega\times\Theta\times\mathcal{A}}_{LP}(P^{\mathcal{A}}_N(\circ|\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t), P^{\mathcal{A}}_{\infty}(\circ|\widetilde{\omega}_t, \widetilde{\pi}^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t)) < \frac{\gamma}{6} \qquad \text{(i)}$$

with probability at least $(1-\rho)^{0.25}$. From Lemma 5 for any $\xi > 0$ we can choose N_1^* such that for all $N > N_1^*$ and uniformly over $(\omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$

$$d_{LP}^{\Theta\times\mathcal{A}}(\pi_{t+1}^{\Theta}, \widetilde{\pi}_{t+1}^{\Theta}) + d_{LP}^{\Theta\times\mathcal{A}}(\pi_{t+1}^{\mathcal{A}}, \widetilde{\pi}_{t+1}^{\mathcal{A}}) < \xi$$

with probability at least $(1-\rho)^{0.125}$ where $P^C(\omega_{t+1}, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}|\sigma^{DCE}) = (\widetilde{\pi}_{t+1}^{\Theta}, \widetilde{\pi}_{t+1}^{\mathcal{A}})$ and $(\pi_{t+1}^{\Theta}, \pi_{t+1}^{\mathcal{A}})$ are the (random) empirical distributions realized in the N-agent game. Since ω_{t+1} is distributed independently on N we have for sufficiently small ξ (or, alternately, N_1^* large enough) that

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}(P_N^{\mathcal{A}}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), P_\infty^{\mathcal{A}}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})) < \frac{\gamma}{18} \tag{A.1}$$

with probability at least $(1-\rho)^{0.125}$ uniformly over $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$. Lemma 6 implies that π_{t+1}^{Θ} is a tight measure.

Equation A.1 implies that for any $\gamma, \rho > 0$ we can choose N^* such that for $N > N^*$ we have with probability $(1-\rho)^{0.25}$

$$\begin{aligned} d_K^{\Omega\times\Theta\times\mathcal{A}}(P_N^{\mathcal{A}}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), P_\infty^{\mathcal{A}}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})) &< \frac{\gamma}{18} \\ d_K^{\Omega\times\Theta\times\mathcal{A}}(P_N^{\mathcal{A}}(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}}), P_\infty^{\mathcal{A}}(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}})) &< \frac{\gamma}{18} \end{aligned}$$

From the uniform continuity of $P_\infty^{\mathcal{A}}$ (lemma 4), 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

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}(P_\infty^{\mathcal{A}}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), P_\infty^{\mathcal{A}}(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}})) < \frac{\gamma}{18}$$

Combining these relations we have

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}(P_N^{\mathcal{A}}(\circ|\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}}), P_\infty^{\mathcal{A}}(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^{\Theta}, \widetilde{\pi}_t^{\mathcal{A}})) < \frac{\gamma}{6}$$

with probability $(1-\rho)^{0.25}$. This establishes our claim for $\tau = 1$, and we proceed with a proof by induction.

Now assume our claim holds for all $\tau' < \tau$ and consider the transition from $(\omega_t, \pi_t^{\Theta}, \pi_t^{\mathcal{A}})$ to $(\omega_{t+1}, \pi_{t+1}^{\Theta}, \pi_{t+1}^{\mathcal{A}})$. From our induction hypothesis we know that for any $\xi, \rho > 0$ we can choose N_1^* such that for all $N > N_1^*$ we have

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

with probability $(1-\rho)^{0.25}$. Consider a particular realization of $\omega_{t+1} \in \Omega$ and assume that $d_{LP}^\Theta(\pi_{t+1}^\Theta, \widetilde{\pi}_{t+1}^\Theta) + d_{LP}^{\mathcal{A}}(\pi_{t+1}^{\mathcal{A}}, \widetilde{\pi}_{t+1}^{\mathcal{A}}) < \xi$.[59] For $\xi > 0$ sufficiently small (alternately, N_1^* sufficiently large) we have from our induction hypothesis that for any $\gamma > 0$ we can choose N_2^* such that for all $N > N^* = \max\{N_1^*, N_2^*\}$ we have

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^{\tau-1}(\circ|\omega_{t+1}, \pi_{t+1}^\Theta, \pi_{t+1}^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^{\tau-1}(\circ|\omega_{t+1}, \widetilde{\pi}_{t+1}^\Theta, \widetilde{\pi}_{t+1}^{\mathcal{A}})) < \frac{\gamma}{3} \tag{A.2}$$

with probability at least $(1-\rho)^{0.25}$ and uniformly over ω_{t+1}. Since equation A.2 holds uniformly over ω_{t+1} and ω_{t+1} is distributed independently of N, we have that[60]

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^{\tau}(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^{\tau}(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}})) < \frac{\gamma}{3} \tag{A.3}$$

with probability at least $(1-\rho)^{0.5}$.

Therefore we have for any $\varepsilon, \rho > 0$ and $N > N^*$ we have with probability $1-\rho$

$$\begin{aligned} d_K^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^{\tau}(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^{\tau}(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}})) &< \frac{\gamma}{3} \\ d_K^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^{\tau}(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^{\tau}(\circ|\widetilde{\omega}_t, \widetilde{\pi}_t^\Theta, \widetilde{\pi}_t^{\mathcal{A}})) &< \frac{\gamma}{3} \end{aligned}$$

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

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

Combining these relations we have

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

with probability $1-\rho$. □

[59] Recollect that $(\pi_{t+1}^\Theta, \pi_{t+1}^{\mathcal{A}})$ are the (random) empirical distributions realized in the N-agent game and $(\widetilde{\pi}_{t+1}^\Theta, \widetilde{\pi}_{t+1}^{\mathcal{A}})$ are the empirical distributions realized in the nonatomic limit game.

[60] The added factor of $\frac{\gamma}{6}$ accounts for those low probability events where $d_{LP}^\Theta(\pi_{t+1}^\Theta, \widetilde{\pi}_{t+1}^\Theta) + d_{LP}^{\mathcal{A}}(\pi_{t+1}^{\mathcal{A}}, \widetilde{\pi}_{t+1}^{\mathcal{A}}) \geq \xi$

Before we prove our approximation theorems, we first prove a useful lemma regarding convergence of the continuation values of the game. Let $V_N \to V_\infty$ as $N \to \infty$ denote that for any $\varepsilon > 0$ we can choose N^* such that $\left\|V_N(\theta_t, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma) - V_\infty(\theta_t, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma)\right\| < \varepsilon$ for all $N > N^*$ and $(\theta_t, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}) \in \Theta \times \Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$.

Lemma 7. *Fix a uniformly continuous strategy σ. Then $V_N \to V_\infty$ as $N \to \infty$ and V_∞ is continuous in $(\theta_t, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A})$.*

Proof. Note that we can write

$$\begin{aligned} V_N(\theta_t, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma) &= (1-\delta) * E_t^\Psi \left[\sum_{\tau=0}^{\infty} \delta^\tau w_N(\theta_{t+\tau}^i, a_{t+\tau}, \omega_{t+\tau}, \pi_{t+\tau}^\Theta, \pi_{t+\tau}^\mathcal{A})\right] \\ V_\infty(\theta_t, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma) &= (1-\delta) * E_t^\Omega \left[\sum_{\tau=0}^{\infty} \delta^\tau w(\theta_{t+\tau}^i, a_{t+\tau}, \omega_{t+\tau}, \pi_{t+\tau}^\Theta, \pi_{t+\tau}^\mathcal{A})\right] \end{aligned}$$

From the the boundedness of w and assumption 4 we can define M such that

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

Choose τ^* such that $M * \delta^{\tau^*} < \frac{\varepsilon}{4}$. Note that through period τ^* we can choose N sufficiently large that for any realized path $((\theta_t, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}), .., (\theta_{t+\tau^*}, \omega_{t+\tau^*}, \pi_{t+\tau^*}^\Theta, \pi_{t+\tau^*}^\mathcal{A}))$ we have for $\tau \in \{0, .., \tau^*\}$

$$w_N(\theta_{t+\tau}, \omega_{t+\tau}, \pi_{t+\tau}^\Theta, \pi_{t+\tau}^\mathcal{A}) - w(\theta_{t+\tau}, \omega_{t+\tau}, \pi_{t+\tau}^\Theta, \pi_{t+\tau}^\mathcal{A})| < \frac{\varepsilon}{4}$$

Together these facts imply

$$\begin{aligned} ||V_N(\theta_t, &\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma) - V_\infty(\theta_t, \omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}|\sigma)|| < \\ &\frac{\varepsilon}{2} + (1-\delta) * \sum_{\tau=0}^{\tau^*} \delta^\tau ||E_t^\Psi \left[w(\theta_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi_{t+\tau}^\Theta, \pi_{t+\tau}^\mathcal{A})\right] - \\ &E_t^\Omega \left[w(\theta_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi_{t+\tau}^\Theta, \pi_{t+\tau}^\mathcal{A})\right] || \end{aligned}$$

Let $\mathcal{I}$ denote the event $((\omega_t, \pi_t^\Theta, \pi_t^\mathcal{A}), .., (\omega_{t+\tau^*}, \pi_{t+\tau^*}^\Theta, \pi_{t+\tau^*}^\mathcal{A}))$. We can decompose these expectations as follows

$$
\begin{aligned}
&E_t^{\Psi}[w(\theta^i_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})|\mathcal{I}] = \\
&\quad \int_{\Theta_{t+\tau}} \int_{\mathcal{A}} w(\theta_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau}) * T(d\theta_{t+\tau}|\theta_{t+\tau-1}, a_{t-1}, \omega_{t+\tau}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t) * ... * \\
&\quad T(d\theta_{t+1}|\theta_t, a_t, \omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t) * \sigma(\theta_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau})[da_{t+\tau}] * ... * \sigma(\theta_t, \omega_t, \pi^{\Theta}_t)[da_t]
\end{aligned}
$$

$$
\begin{aligned}
&E_t^{\Omega}[w(\theta^i_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})|\mathcal{I}] = \\
&\quad \int_{\Theta_{t+\tau}} \int_{\mathcal{A}} w(\theta_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau}) * T(d\theta_{t+\tau}|\theta_{t+\tau-1}, a_{t-1}, \omega_{t+\tau}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t) * ... * \\
&\quad T(d\theta_{t+1}|\theta_t, a_t, \omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t) * \sigma(\theta_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau})[da_{t+\tau}] * ... * \sigma(\theta_t, \omega_t, \pi^{\Theta}_t)[da_t]
\end{aligned}
$$

From the uniform continuity of w, T, and σ we have that $E_t^{\Psi}\left[w(\theta^i_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})|\mathcal{I}\right]$ is uniformly continuous in $\mathcal{I}$. Since theorem 1 implies $(P_N^{\mathcal{A}})^{\tau}$ converges to $(P_{\infty}^{\mathcal{A}})^{\tau}$ for all $\tau \in \{0, ..., \tau^*\}$ (and recollecting that $P_N^{\mathcal{A}}$ and $P_{\infty}^{\mathcal{A}}$ govern the distribution of $\mathcal{I}$ in the large finite and limit game respectively), we have that we can choose N large enough that

$$
\left\| E_t^{\Psi} w(\theta^i_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})] - E_t^{\Omega}\left[w(\theta^i_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})\right] \right\| < \frac{\varepsilon}{2}
$$

This then implies

$$
\begin{aligned}
\left\| V_N(\theta_t, \omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma) - V_{\infty}(\theta_t, \omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma) \right\| &< \frac{\varepsilon}{2} + (1-\delta) * \sum_{\tau=0}^{\tau^*} \delta^{\tau} \frac{\varepsilon}{2} \\
&< \varepsilon
\end{aligned}
$$

To see the proof of the continuity of V_{∞} it suffices to note that we can write

$$
V_{\infty}(\theta_t, \omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t|\sigma) = (1-\delta) * E_t^{\Omega}\left[\sum_{\tau=0}^{\infty} \delta^{\tau} E_t^{\Omega}\left[w(\theta^i_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})|\mathcal{I}\right]\right]
$$

Noting that $E_t^{\Omega}\left[w(\theta^i_{t+\tau}, a_{t+\tau}, \omega_{t+\tau}, \pi^{\Theta}_{t+\tau}, \pi^{\mathcal{A}}_{t+\tau})|\mathcal{I}\right]$ is continuous in $\mathcal{I}$ and that $P_{\infty}^{\mathcal{A}}$ (which governs the distribution of $\mathcal{I}$) is continuous in $(\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t)$. □

To prove Theorem 2, we argue that the equilibrium outcomes of the large finite game stay near the family of equilibrium outcomes generated in the nonatomic limit game for large N. Since σ_{ID}^{DCE}, the indirect representation of DCE strategy σ^{DCE}, is a best response to this family of paths, it must be an approximate best response to any nearby family of outcomes. Our mean field approximation result, Theorem 1, insures that the path of play realized in large finite game when all agents save one play σ_{ID}^{DCE} 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_{ID}^{DCE} yielding an ε-improvement would yield an improvement 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 high probability as $N \to \infty$ and the utility loss outside of this event is bounded, we can conclude then that σ_{ID}^{DCE} is an $\varepsilon-$BNE of the N-agent game for sufficiently large N. The proof is not trivial since optimal deviations from σ_{ID}^{DCE} need not be continuous, which makes the utility from such a deviation (potentially) discontinuous as P_N^A 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 game for $N > N^*$ if $\pi_0^\Theta \in \Delta_N(\Theta)$ and $\pi_0^\mathcal{A} \in \Delta_N(\mathcal{A})$. Furthermore, N^* can be chosen uniformly across $(\omega_0, \pi_0^\Theta, \pi_0^\mathcal{A})$.*

Proof. Suppose our claim is false. 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)

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

where $\widetilde{P}_N^A$ describes the evolution of the N player game where one agent uses strategy σ' and all others follow σ_{ID}^{DCE}. From the the boundedness of w and assumption 4 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

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

Therefore we can write[61]

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

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

$$\sigma'_{\tau^*}(\theta_t,\omega_t,\pi_t^{\Theta}) = \begin{cases} \sigma'(\theta_\tau,\omega_t,\pi_t^{\Theta}) \text{ for } t \leq \tau^* \\ \sigma_{ID}^{DCE}(\theta_t,\omega^t) \text{ for } t > \tau^* \end{cases}$$

Consider the problem facing the agent of type θ_0 in the N-agent game in period τ^* given strategy σ_{ID}^{DCE} is followed in all future periods and action a' is dictated by $\sigma_{ID}^{DCE}(\theta_\tau,\omega^\tau)$ in this period

$$\max_a w_N(\theta_{\tau^*},a,\omega_{\tau^*},\pi_{\tau^*}^{\Theta},\pi_{\tau^*}^{\mathcal{A}} + \frac{1}{N}[\varsigma_a - \varsigma_{a'}]) + \delta E_{\tau^*}^{\Psi}[V_N(\theta_{\tau^*+1}^i,\omega_{\tau^*+1},\pi_{\tau^*+1}^{\Theta},\pi_{\tau^*+1}^{\mathcal{A}}|\sigma_{ID}^{DCE})|a] \tag{OPT}$$

where we explicitly condition the expectation on the action of the player considered.

Let $Q \subset \Theta^{\tau^*}\times\Omega^{\tau^*}$ denote the support of the path of play realized in the nonatomic limit game in periods $0,1,...,\tau^*$ where all agents play σ_{ID}^{DCE}. Let Q^{ϕ} denote a ϕ neighborhood

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

of Q. Note that for $(\theta^{\tau^*}, \omega^{\tau^*}) \in Q$, where $\theta^{\tau^*} = (\theta_0, \theta_1, .., \theta_{\tau^*})$ is a path of evolution of agent type through period τ^*, the agent in the nonatomic limit game solves the problem in period τ^*

$$\sigma_{\tau^*}^{DCE}(\theta_{\tau^*}, \omega^{\tau^*}) \in \arg\max_a w(\theta_{\tau^*}, a, \omega_{\tau^*}, \pi_{\tau^*}^{\Theta}, \pi_{\tau^*}^{\mathcal{A}}) + \\ \delta E_{\tau^*}^{\Omega}[V_\infty(\theta_{\tau^*+1}^i, \omega_{\tau^*+1}, \pi_{\tau^*+1}^{\Theta}, \pi_{\tau^*+1}^{\mathcal{A}} | \sigma^{DCE}) | a]$$

Given the continuity of V_∞ and that $V_N \to V_\infty$ (lemma 7), the continuity of T and w, and the uniform convergence of w_N to w, there is no $\frac{\varepsilon}{2}$-improvement over σ_{ID}^{DCE} available in period τ^* in problem (OPT) for large enough N. The continuity and convergence properties also imply that in the event $(\theta^{\tau^*}, \omega^{\tau^*}) \in Q^\phi$ we have that σ_{ID}^{DCE} is an ε best response to (OPT) for $\phi > 0$ sufficiently small. Therefore it must be that there is an ε-improvement over σ_{ID}^{DCE} available in periods $\{1, ..., \tau^* - 1\}$. We can repeat our induction step, however, to show that there is no ε-improvement over σ_{ID}^{DCE} available in any period between $\{1, ..\tau^*\}$. Therefore, σ_{ID}^{DCE} is an $\varepsilon - BNE$ strategy in the N agent game at $t = 0$ conditional on the event Q^ϕ.

Note that since P_N^A converges weakly to P_∞^A and T and σ_{ID}^{DCE} are uniformly continuous, for any $\phi > 0$ we can choose N sufficiently large that Q^ϕ is realized with probability at least $1 - \phi$. From the boundedness of w_N we then have that σ_{ID}^{DCE} can be improved by at most

$$\varepsilon + \phi * M$$

For ϕ sufficiently small (and hence N sufficiently large) we have that $\varepsilon + \phi * M < 2\varepsilon$ and hence σ_{ID}^{DCE} is a 2ε−BNE. Our N^* can be chosen uniformly since since P_N^A converges to P_∞^A and w_N converges to w uniformly and T and σ_{ID}^{DCE} are uniformly continuous. □

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 and continuity properties rather than assume them.

Theorem 3. *Assume that*

- *Θ and Ω are compact*
- *There exists an N^* such that $\cup_{N=N^*}^{\infty}\mathcal{E}(N|\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$ is equicontinuous in $\Theta \times \Omega \times \Delta(\Theta)$*

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

$$\underset{N\to\infty}{Lim}\,\mathcal{E}(N|\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}}) = \mathcal{E}^{\infty}(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}}) \subset \mathcal{E}^{NA}(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}})$$

where $(\omega_0, \pi_0^{\Theta}, \pi_0^{\mathcal{A}}) \in \Omega \times \Delta_{N^*}(\Theta) \times \Delta_{N^*}(\mathcal{A})$.

Proof. (of Theorem 3) 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. By the Heine-Borel theorem, the set $\cup_{N=N^*}^{\infty}\mathcal{E}(N)$ is uniformly equicontinuous in $\Theta\times\Omega\times\Delta(\Theta)$. 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 [39] 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 $\Delta(\Theta)$.[62] This extension preserves the modulus of uniform 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 $\sigma^N \to \sigma^{\infty}$.[63] Since $\{\sigma^N\}_{N=1}^{\infty}$ is equicontinuous, the Arzelà-Ascoli theorem implies that σ^{∞} is continuous (Dunford and Schwartz [21]).

Let P_{∞}^{A} denote the the transition probability function for the nonatomic limit game given σ^{∞}, T, and G. Let $Q_N \subset \Theta^{\tau^*} \times \Omega^{\tau^*}$ denote the support of the path of play in the N agent game realized when all agents play σ^N, and let Q_N^{ϕ} denote a ϕ neighborhood of Q_N.

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

[62] An alternative proof technique would be to use Isbell [31] 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 $\Theta \times \Omega \times \Delta(\Theta)$.

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

condition

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

for some profitable deviation σ'. 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

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

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

$$\sigma'_{\tau^*}(\theta_\tau, \omega^\tau) = \begin{cases} \sigma'(\theta_\tau, \omega^\tau) \text{ if } \tau \leq \tau^* \\ \sigma^\infty(\theta_\tau, \omega^\tau) \text{ if } \tau > \tau^* \end{cases}$$

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

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

$$\max_a w(\theta_{\tau^*}, a, \omega_{\tau^*}, \pi_{\tau^*}^\Theta, \pi_{\tau^*}^\mathcal{A}) + E_{\tau^*}^\Omega[V_\infty(\theta_{\tau^*+1}^i, \omega_{\tau^*+1}, \pi_{\tau^*+1}^\Theta, \pi_{\tau^*+1}^\mathcal{A} | \sigma^{DCE}) | a] \tag{OPT}$$

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

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

The conditioning reflects both the uncertainty of the evolution of market aggregates, captured by P^A_N, as well as uncertainly regarding the future evolution of θ^i_t. We condition on the agent's action to explicitly denote that this choice has an effect on both $(\omega_{\tau^*+1}, \pi^{\Theta}_{\tau^*+1}, \pi^{\mathcal{A}}_{\tau^*+1})$ and $\theta^i_{\tau^*+1}$.

Given the continuity of V_∞ and that $V_N \to V_\infty$ as $N \to \infty$ (lemma 7)[64], the continuity of T and w, and the uniform convergence of w_N to w, we have that for sufficiently large N there is no $\frac{1}{4}\varepsilon$-improvement over σ^N available in period τ^* in problem (OPT). Since $\sigma^N \to \sigma^\infty$ uniformly, there is no $\frac{1}{2}\varepsilon$ improvement over σ^∞ available in period τ^* in problem (OPT). Furthermore, since this result holds for all $(\theta_{\tau^*}, \omega^{\tau^*}) \in Q_N$ and w and T are continuous, it must be that for $\phi > 0$ sufficiently small there is no $\frac{3}{4}\varepsilon$ improvement over σ^∞available in period τ^* for any $(\theta_{\tau^*}, \omega^{\tau^*}) \in Q^\phi_N$. Finally, note that since $P^A_N \to P^A_\infty$ in the weak-* topology, for any $\phi > 0$ we can choose N sufficiently large that $(\theta_{\tau^*}, \omega^{\tau^*}) \in Q^\phi_N$ with probability at least $1 - \phi$. Therefore for $\phi > 0$ sufficiently small, there can be at most a $\frac{3}{4}\varepsilon + \phi * M < \varepsilon$ improvement over σ^∞ in period τ^*. Thus it must be that there is an ε-improvement over σ^∞ available in periods $\{1, ..., \tau^* - 1\}$. However, we can repeat our induction step 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 in the limit as $N \to \infty$. The upper hemicontinuity of $\mathcal{E}(N|\omega_0, \pi^{\Theta}_0, \pi^{\mathcal{A}}_0)$ in the limit $N \to \infty$ is uniform over $\Omega \times \Delta_N(\Theta) \times \Delta_N(\mathcal{A})$ since P^A_N converges to P^A_∞ and w_N converges to w uniformly and T, σ^N, and σ^∞ are uniformly continuous. □

We now provide a proof of Corollary 5 using the machinery of Theorem 3.

Corollary 5. *Consider a convergent sequence of uniformly continuous Bayesian-Nash equilibrium strategies,* $\{\sigma^N\}^\infty_{N=1}$ *such that* $\sigma^N \to \sigma^\infty$*. Assume* σ^∞ *is uniformly continuous. Let* $\Lambda_\tau = \prod^\tau_{t=0} supp[(\omega_t, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t)]$ *be the support of the stochastic process induced*

[64] Note that σ'_{τ^*} yields the same outcomes as σ^∞ for periods $t > \tau^*$. Since all agents (effectively) follow σ^∞ following period τ^*, the evolution reverts back to being described by P^A_N.

by $P^{\mathcal{A}}_{\infty}$ in periods 1 through τ. 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_∞ 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. Let U^τ denote the projection of U onto $\prod_{t=0}^{\tau} \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$. 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

$$d_{LP}^{\Omega\times\Theta\times\mathcal{A}}((P_N^{\mathcal{A}})^\tau(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}}), (P_\infty^{\mathcal{A}})^\tau(\circ|\omega_t, \pi_t^\Theta, \pi_t^{\mathcal{A}})) < \gamma \tag{A.4}$$

for $\tau \in \{1, .., \tau^*\}$ with probability at least $1 - \rho$. Choose γ sufficiently small that equation A.4 implies that the path of $\{(\omega_{t+\tau}, \pi_{t+\tau}^\Theta, \pi_{t+\tau}^{\mathcal{A}})\}_{\tau=0}^{\tau^*}$ in the N-agent game lies in the open set U^{τ^*} with probability at least $1 - \rho$. From this point we the logic of Theorem 3 can be applied directly where the set U^{τ^*} replaces Q_N^ϕ. □

A.2 Proofs from Section 4

Lemma 2. *Assume T is uniformly continuous, σ_{ID}^{DCE} is fixed, and assumption 6 holds. For any $\varepsilon > 0$, we can find a $\gamma > 0$ such that if π^Θ and $\widetilde{\pi}^\Theta$ are admissible and $d_{LP}^\Theta(\pi^\Theta, \widetilde{\pi}^\Theta) < \gamma$, then $d_{LP}^{\Theta\times\mathcal{A}}(P_{ID}^C(\circ|\pi^\Theta, \omega^t), P_{ID}^C(\circ|\widetilde{\pi}^\Theta, \omega^t)) < \varepsilon$ uniformly over $\pi^\Theta \in \Delta(\Theta)$ for any sequence of aggregate states $\omega^t = (\omega_0, ..., \omega_t)$.*

Proof. It suffices to show that $P_{ID}^C(\circ|\pi^\Theta, \omega^t)$ is continuous to imply that for any $\tau < \infty$ we have $(P_{ID}^C)^\tau(\circ|\pi^\Theta, \omega^t)$ is continuous for any sequence of aggregate states $(\omega_{t+1}, ..., \omega_{t+\tau})$. We prove the one-step case since the τ step case follows by induction.

Fix $(\omega_{t+1}, \pi_t^\Theta, \pi_t^{\mathcal{A}}) \in \Omega \times \Delta(\Theta) \times \Delta(\mathcal{A})$. We want to show that for any $\varepsilon > 0$ we can choose $\gamma > 0$ so that if $d_{LP}(\pi_t^\Theta, \widetilde{\pi}_t^\Theta) < \gamma$ we have $d_{LP}(\pi_{t+1}^\Theta, \widetilde{\pi}_{t+1}^\Theta) + d_{LP}(\pi_{t+1}^{\mathcal{A}}, \widetilde{\pi}_{t+1}^{\mathcal{A}}) < \varepsilon$ where $P_{ID}^C(\circ|\pi^\Theta, \omega^t) = (\pi_{t+1}^\Theta, \pi_{t+1}^{\mathcal{A}})$ and $P_{ID}^C(\circ|\widetilde{\pi}^\Theta, \omega^t) = (\widetilde{\pi}_{t+1}^\Theta, \widetilde{\pi}_{t+1}^{\mathcal{A}})$.

We first show that $d_{LP}(\pi_{t+1}^\Theta, \widetilde{\pi}_{t+1}^\Theta) < \frac{\varepsilon}{2}$, which is equivalent to

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \pi^\Theta, \pi^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^\Theta(d\theta), \tag{A.5}$$
$$\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^\Theta, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \widetilde{\pi}^\Theta(d\theta)) < \frac{\varepsilon}{2}$$

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

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \pi^{\Theta}, \pi^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \pi^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)) < \frac{\varepsilon}{6} \tag{i}$$

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \pi^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)) \quad < \frac{\varepsilon}{6} \tag{ii}$$

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta))) \quad < \frac{\varepsilon}{6} \tag{iii}$$

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

Term (iii) requires assumption 6. Let $D = \{\theta : \sigma_{ID}^{DCE}(\theta, \omega^{t+1})$ is discontinuous at $\theta\}$ and note that assumption 6 implies that for any pair $\pi^{\Theta}, \widetilde{\pi}^{\Theta}$ we can choose a set $D^* \supseteq D$ such that $\pi^{\Theta}[D^*], \widetilde{\pi}^{\Theta}[D^*] < \frac{\varepsilon}{9}$ and T and σ_{ID}^{DCE} are uniformly continuous over $\Theta - D^*$. Therefore we can choose γ sufficiently small that

$$d_{LP}(\int_{\mathcal{A}\times(\Theta-D^*)} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)), \int_{\mathcal{A}\times(\Theta-D^*)} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta))) < \frac{\varepsilon}{18}$$

Note also that since $\int_{\mathcal{A}} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] \leq 1$ we have

$$d_{LP}(\int_{\mathcal{A}\times(\Theta-D^*)} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \pi^{\Theta}(d\theta)),) < \frac{\varepsilon}{18}$$

and

$$d_{LP}(\int_{\mathcal{A}\times(\Theta-D^*)} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta)),$$
$$\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}, \widetilde{\pi}^{\mathcal{A}}) * \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[da] * \widetilde{\pi}^{\Theta}(d\theta))) < \frac{\varepsilon}{18}$$

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

To close our argument we must show $d_{LP}(\pi_{t+1}^{\mathcal{A}}, \widetilde{\pi}_{t+1}^{\mathcal{A}}) < \frac{\varepsilon}{2}$. This is equivalent to

$$d_{LP}(\int_{\Theta} \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[\circ] * \pi_{t+1}^{\Theta}[d\theta], \int_{\Theta} \sigma_{ID}^{DCE}(\theta, \omega^{t+1})[\circ] * \widetilde{\pi}_{t+1}^{\Theta}[d\theta]) < \frac{\varepsilon}{2} \tag{A.6}$$

As this follows from an argument essentially identical to the proof that term (iii) above holds, it is omitted. Equations A.5 and A.6 together imply our desired result

$$d_{LP}(\pi_{t+1}^{\Theta}, \widetilde{\pi}_{t+1}^{\Theta}) + d_{LP}(\pi_{t+1}^{\mathcal{A}}, \widetilde{\pi}_{t+1}^{\mathcal{A}}) < \varepsilon$$

□

B Additional Extensions

B.1 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 $(\sigma'_{\mathcal{I}}, \sigma_{-\mathcal{I}})$.[65] Consider the following definition of K-coalition-proof equilibria in large finite games.

Definition 6. *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 $i \in \mathcal{I}, j \in \{1, ..., N\} \backslash \mathcal{I}$ and $\sigma'_{\mathcal{I}} \in \Sigma^M$, $\sigma'_j \in \Sigma$*

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

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

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. 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.[66]

Corollary 5. *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. Follows from a straighforward modification of Theorem 2. □

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.

B.2 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,

[66] The extension of Theorem 3 to coalition-proof equilibria 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.

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.[67] This notational artifice allows us to preserve the formal symmetry of the model while allowing agents to act asynchronously.

C Supplementary Proofs

Lemma 3. *If T, σ are (uniformly) continuous, then $(P^C)^\tau$ is (uniformly) continuous in $\Delta(\Theta) \times \Delta(\mathcal{A})$ for any $\tau < \infty$ conditional on any sequence of aggregate states $(\omega_{t+1}, ..., \omega_{t+\tau})$*

Proof. It suffices to show that $P^C(\omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})$ is continuous to imply that for any $\tau < \infty$ we have $(P^C)^\tau(\omega_{t+1}, \pi_t^\Theta, \pi_t^\mathcal{A})$ is continuous for any sequence of aggregate states $(\omega_{t+1}, ..., \omega_{t+\tau})$. We prove the one-step case since the τ step case follows by a straightforward induction step.

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

[67] 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 $\omega_t = "S"$ ($\omega_t = "B"$).

First we prove $d_{LP}(\pi^{\Theta}_{t+1}, \widetilde{\pi}^{\Theta}_{t+1}) < \frac{\varepsilon}{2}$. Note that this is equivalent to

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \pi^{\Theta}_t)[da] * \pi^{\Theta}(d\theta), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t)[da] * \widetilde{\pi}^{\Theta}_t(d\theta)) < \frac{\varepsilon}{2} \tag{C.1}$$

We proceed by analyzing the telescopic expansion of equation C.1 into the following terms

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \pi^{\Theta}_t, \pi^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \pi^{\Theta}_t)[da] * \pi^{\Theta}(d\theta), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \pi^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \pi^{\Theta}_t)[da] * \pi^{\Theta}_t(d\theta)) < \frac{\varepsilon}{10} \tag{i}$$

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \pi^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \pi^{\Theta}_t)[da] * \pi^{\Theta}_t(d\theta)), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \pi^{\Theta}_t)[da] * \pi^{\Theta}_t(d\theta)) \quad < \frac{\varepsilon}{10} \tag{ii}$$

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \pi^{\Theta}_t)[da] * \pi^{\Theta}_t(d\theta)), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t)[da] * \pi^{\Theta}_t(d\theta)) < \frac{\varepsilon}{5} \tag{iii}$$

$$d_{LP}(\int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t)[da] * \pi^{\Theta}_t(d\theta)), \int_{\mathcal{A}\times\Theta} T(\circ|\theta, a, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t, \widetilde{\pi}^{\mathcal{A}}_t) * \sigma(\theta, \omega_{t+1}, \widetilde{\pi}^{\Theta}_t)[da] * \widetilde{\pi}^{\Theta}_t(d\theta))) \quad < \frac{\varepsilon}{10} \tag{iv}$$

Continuity of terms (i) to (iv) clearly follows from the continuity of T and σ. Together these terms imply that equation C.1 holds.

Now we turn to the distribution of actions to show that $d_{LP}(\pi^{\mathcal{A}}_{t+1}, \widetilde{\pi}^{\mathcal{A}}_{t+1}) < \frac{\varepsilon}{2}$. This requires that

$$d_{LP}(\int_{\Theta}\sigma(\theta,\omega_{t+1},\pi^{\Theta}_{t+1})[\circ]*\pi^{\Theta}_{t+1}[d\theta],\\ \int_{\Theta}\sigma(\theta,\omega_{t+1},\widetilde{\pi}^{\Theta}_{t+1})[\circ]*\widetilde{\pi}^{\Theta}_{t+1}[d\theta])<\frac{\varepsilon}{2} \tag{C.2}$$

Performing the telescopic expansion yields

$$d_{LP}(\int_{\Theta}\sigma(\theta,\omega_{t+1},\pi^{\Theta}_{t+1})[\circ]*\pi^{\Theta}_{t+1}[d\theta],\\ \int_{\mathcal{A}\times\Theta}\int_{\Theta}\sigma(\theta,\omega_{t+1},\widetilde{\pi}^{\Theta}_{t+1})[\circ]*\pi^{\Theta}_{t+1}[d\theta])<\frac{\varepsilon}{6} \tag{i}$$

$$d_{LP}(\int_{\Theta}\sigma(\theta,\omega_{t+1},\widetilde{\pi}^{\Theta}_{t+1})[\circ]*\pi^{\Theta}_{t+1}[d\theta],\\ \int_{\mathcal{A}\times\Theta}\int_{\Theta}\sigma(\theta,\omega_{t+1},\widetilde{\pi}^{\Theta}_{t+1})[\circ]*\widetilde{\pi}^{\Theta}_{t+1}[d\theta])<\frac{\varepsilon}{6} \tag{ii}$$

Terms (i) and (ii) follow from the continuity of σ and together these terms imply that equation C.2 holds. Equations C.1 and C.2 together imply $d_{LP}(\pi^{\Theta}_t,\widetilde{\pi}^{\Theta}_t)+d_{LP}(\pi^{\mathcal{A}}_t,\widetilde{\pi}^{\mathcal{A}}_t)<\gamma$ (and hence the desired continuity of $P^C(\omega_{t+1},\pi^{\Theta}_t,\pi^{\mathcal{A}}_t)$ in the weak-* topology). ☐

Lemma 4. $(P^{\mathcal{A}}_{\infty})^{\tau}(\circ|\omega_t,\pi^{\Theta}_t,\pi^{\mathcal{A}}_t)$ *is uniformly continuous for any* $\tau<\infty$

Proof. We need to prove that expectations of uniformly continuous functions of $(\omega_{t+\tau},\pi^{\Theta}_{t+\tau},\pi^{\mathcal{A}}_{t+\tau})$ are uniformly continuous in $(\omega_t,\pi^{\Theta}_t,\pi^{\mathcal{A}}_t)$. 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 $E_t[h(\omega_{t+\tau},\pi^{\Theta}_{t+\tau},\pi^{\mathcal{A}}_{t+\tau})|\omega_t,\pi^{\Theta}_t,\pi^{\mathcal{A}}_t]$ is uniformly continuous in $(\omega_t,\pi^{\Theta}_t,\pi^{\mathcal{A}}_t)$. Note that the proposition is true for $\tau=1$ since P^C and G are uniformly continuous in the weak-* topology (the former from lemma 3 and the latter by assumption 1).

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

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

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 $\tau' = 1$. But then we can write

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

Since by assumption our proposition holds for $\tau - 1$, 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

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

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