

Equilibria of nonatomic anonymous games*

Simone Cerreia-Vioglio^a, Fabio Maccheroni^a, David Schmeidler^b

^aUniversità Bocconi and Igier, ^bTel Aviv University

November 2019

Abstract

We add here another layer to the literature on nonatomic anonymous games started with the 1973 paper by Schmeidler. More specifically, we define a new notion of equilibrium which we call ε -estimated equilibrium and prove its existence for any positive ε . This notion encompasses and brings to nonatomic games recent concepts of equilibrium such as self-confirming, peer-confirming, and Berk–Nash. This augmented scope is our main motivation. At the same time, our approach also resolves some conceptual problems present in Schmeidler (1973), pointed out by Shapley. In that paper the existence of pure-strategy Nash equilibria has been proved for any nonatomic game with a continuum of players, endowed with an atomless countably additive probability. But, requiring Borel measurability of strategy profiles may impose some limitation on players' choices and introduce an exogenous dependence among players' actions, which clashes with the nature of noncooperative game theory. Our suggested solution is to consider every subset of players as measurable. This leads to a nontrivial purely finitely additive component which might prevent the existence of equilibria and requires a novel mathematical approach to prove the existence of ε -equilibria.

1 Introduction

The original framework of Schmeidler. Games with a continuum of anonymous players were introduced by Schmeidler in [30] where he also proved the existence of pure-strategy Nash equilibria for these games.¹ At the time, there were models of markets and cooperative games with infinitely many players, but not of noncooperative games. In [30], the players' space is modelled to be the unit interval endowed with the Borel σ -algebra and the Lebesgue measure, where there is a finite set of actions and each player chooses an action from this set. The utility of each player depends on the distribution of actions across all players and the action he chooses. The interpretation is that the same game is repeated in

*Acknowledgments to be added.

¹At the end of the paper, we recall how this problem naturally arose.

each period. The payoff, in utils, is received at the end of the period. At the same time, because of the anonymity assumption, the strategic complications of repeated games are meaningless here. A paradigmatic example is that of daily commuters driving downtown (or back home) and having to choose a bridge (or tunnel) to enter the city. Thus, in each period they play a one-shot game, analyzed in [30]. Here, the metaphysical assumption of correctly guessing what other players will do, required for playing a Nash equilibrium strategy in one-shot games, is mitigated by two factors. The first is minor: each player has to guess correctly the distribution of the strategy (the same guess for all). The second is major: there is regularity in the daily traffic of commuters. Schmeidler [30] formalizes these intuitions. The limitations of this model are discussed below.

Our motivations. The goal of our paper is to generalize the above finding in several directions. We are motivated by three main observations:

- (i) In recent years, alternative, and perhaps more realistic, notions of equilibrium have been developed for noncooperative games with finitely many players. At the same time, these notions have not been considered for nonatomic anonymous games. In particular, we have in mind equilibrium concepts which allow for beliefs to be not necessarily correct, but nonetheless consistent with the information possessed by each player whether it is endogenously or exogenously generated.
- (ii) In a personal conversation with Schmeidler (in the early 1970s), Shapley pointed out a problem with the modelling of a nonatomic population of players as the unit interval with the Lebesgue measure on Borel sets. As in some mathematical sense there are more nonmeasurable sets than measurable sets in the unit interval, the game, that is the payoff function, may not be defined out of equilibrium. In a similar vein, as later formalized in a general equilibrium framework, Dubey and Shapley [11] raise another issue with the measurability assumption. The measurability of a strategy profile (and similarly of the profile of utilities, which is a common assumption) yields its “near” continuity.² This in turn clashes with the noncooperative idea of strictly independent decision-making, since “close players” tend to play “close strategies”.
- (iii) In modelling a large population of players in which each agent “has the same negligible weight”, Schmeidler opted for the infinite set of points in the unit interval endowed with the Lebesgue measure. At the same time, as noted by Aumann [3], in analyzing economies with a continuum of traders, “the choice of the unit interval as a model for the set of [players] is of no particular significance. A planar or spatial region would

²This intuition is based on Lusin’s Theorem which states that for each $\varepsilon > 0$ each measurable function is continuous when restricted to a suitable compact set which has a measure of at least $1 - \varepsilon$ (see, e.g., Aliprantis and Border [1, Theorem 12.8]).

have done just as well. In technical terms, [the players’ space] can be any measure space without atoms.” Thus, for example, one could alternatively model the players’ space as the set of natural numbers endowed with a natural density. More generally, we suggest using Savage’s structure of nonatomic probabilities defined on the power set of the space of players (Section 2).

Our contributions. Our second and third motivation bring us to model the players’ space as a set T endowed with a nonatomic probability λ defined on all subsets T . Using a measure over the power set takes care of both Shapley’s and Aumann’s comments. In particular, by considering the power set, we allow for the most permissive measurable structure possible, since any profile of strategies or utilities becomes automatically measurable. Measuring the subsets of players/coalitions according to a nonatomic probability on the power set is consistent with Savage’s [29] approach and equivalent to having a qualitative probability on the players’ space, satisfying axiom P6’. However, modelling the players’ space in this generality implies that Nash equilibria might fail to exist (see Example 1 based on Khan, Qiao, Rath, and Sun [18]).

This naturally brings us to look at ε -equilibria and to our first motivation. We introduce a concept of approximate equilibrium for nonatomic anonymous games, which we call ε -estimated equilibrium. This notion of ε -equilibrium encompasses several approximate equilibrium concepts: ε -self-confirming (ε -SCE), ε -peer-confirming (ε -PCE), and ε -Berk–Nash (ε -BNE). These equilibria and their ε -versions are formally defined and discussed in the relevant sections, Sections 3.1, 3.2, and 3.3 (see also the related literature below). They were mostly developed for finite games and, inter alia, in this paper we extend them to nonatomic games. Nevertheless, the principles behind their definitions in a finite-players framework naturally translate to a nonatomic setup. The common thread behind ε -SCE, ε -PCE, and ε -BNE in an anonymous nonatomic game is the following scheme, which is also the basis for our ε -estimated equilibria:

1. Every player best-responds to his beliefs (optimality);
2. The belief of every player is consistent with what he can observe (ε -discrepancy).³

Where these types of equilibrium differ is how point 2 is formalized, since point 1 is translated in the same way for all of them. In particular, in SCE, each player receives a message which is a function of the action he takes and the distribution of actions of the other players. In equilibrium, almost every player best-responds to a distribution that generates a message which is ε -close to the message generated by the true distribution of the actions. In PCE, the message each player receives is the distribution of the actions conditional on

³More precisely, we require points 1 and 2 to hold for every player except a null set of them.

a subset of players: his peers. Thus, almost all the players best-respond to a distribution which is ε -close to the true distribution of actions of their peers, not of all the players. In both ε -SCE and ε -PCE the distributions to which players best-respond are ε -close in terms of observables to the true one; thus they are *endogenously* generated. By contrast, in BNE, each player t entertains an *exogenous* set of possible distributions of actions, denoted by Q_t , that he believes are accurate in describing other players' behavior. Moreover, he is not willing to depart from Q_t . So in equilibrium, almost every player best-responds to a distribution which is ε -close to the best estimate in Q_t of the true distribution of actions, according to a statistical measure.

Our notion of ε -estimated equilibrium provides a framework where we can account for all the three different features described above: that is, the distribution of actions used by each player in equilibrium is ε -close, whether in statistical terms or proper distance, to the set of all distributions which are compatible with the true one. This latter set can be exogenously determined as in BNE or endogenously generated as in SCE or PCE.

In Theorem 1, under mild assumptions, we prove that ε -estimated equilibria always exist. As particular cases, we obtain the existence of self-confirming ε -equilibria (Corollary 1), peer-confirming ε -equilibria (Corollary 3), and Berk–Nash ε -equilibria (Corollary 4). Finally, despite the fact that standard Nash equilibria might fail to exist, we prove that ε -Nash equilibria do exist (Corollary 2).

Related literature. The seminal contribution of Aumann [3] (in a general equilibrium framework), followed by Schmeidler [30] (in a game-theoretic framework), initiated a large literature where the negligibility of agents is modelled via a nonatomic probability players' space (see, e.g., Khan and Sun [22] for a survey).⁴ We will next discuss the relevant literature by connecting it to our three main motivations/contributions.

- (i) Our definition of ε -estimated equilibrium seems to be new. At the same time, it encompasses three types of equilibrium: self-confirming (SCE), peer-confirming (PCE), and Berk–Nash (BNE) which were developed almost exclusively for games with finitely many players, respectively, by Battigalli [6] as well as Fudenberg and Levine [14] (SCE), Lipnowski and Sadler [23] (PCE), and Esponda and Pouzo [12] (BNE). The only exceptions seem to be SCE and BNE, which were also studied for population games, where the latter can be seen as a very special form of nonatomic games. Moreover, we also consider ε -versions of the above three concepts of equilibrium. In discussing ε -SCE of course, two approaches are available. The first assumes that: a) players best-

⁴Many subsequent papers extended Schmeidler's results to more general players' spaces, but where λ is always assumed to be countably additive and A is allowed to be infinite: see, e.g., Balder [5], Khan and Sun [20], Khan, Rath, and Sun [19], Rath [26], and the references therein. The scope of this type of result is analyzed in Carmona and Podczeck [9].

respond to their beliefs, but b) beliefs are only ε -consistent with evidence. The second requires that: a') players ε -best-respond to their beliefs, but b') beliefs are perfectly consistent. For games with finitely many players, the first approach was introduced by Battigalli [6] and Kalai and Lehrer [16] and [17], while the second was proposed for pure equilibria by Azrieli [4]. For nonatomic games, other than population games, the first approach seems to be unexplored, while the second was studied by Azrieli [4]. Using the same setting as Schmeidler, that is, assuming that the players' space is the unit interval with the Lebesgue measure, Azrieli shows that self-confirming equilibria exist (that is, when $\varepsilon = 0$), but when utility depends on the entire profile of strategies and the message feedback is the distribution of actions.⁵ Moreover, in trying to obtain the nonatomic games of Schmeidler as a limit of finite-players games which become arbitrarily large, he shows that self-confirming ε -equilibria eventually exist.⁶ Finally, Azrieli limits his analysis to the case where there is nonmanipulable information also known as own-action independence of feedback. Loosely speaking, this is the case when the feedback each player receives does not depend on the action taken by the player. This rules out several interesting cases.

In our work, we opt for a definition of ε -SCE which requires rational optimization on the players' side, but allows them to entertain ε -consistent beliefs. We do not assume own-action independence. The assumption of ε being strictly positive is due to two reasons: one mathematical and one conceptual. Mathematically, by considering players' spaces which involve finitely additive measures λ , one can show that self-confirming equilibria might fail to exist (Example 1). Conceptually, we take the point of view of Kalai and Lehrer [16] and [17]: we impose rational behavior on players, but allow for slightly inconsistent beliefs. The latter assumption can be justified by interpreting the belief of each player as the belief entertained after many rounds of play, so that learning yields approximately correct predictions about observables. At the limit, beliefs would be perfectly consistent with observations, but before that they might be just ε -consistent.

- (ii) The issue of measurability in nonatomic economies and games has been raised and dealt with by several authors in the past. Khan and Sun [21] proposed to replace the unit interval with the Lebesgue measure with a generic Loeb space. In this way, players (resp., coalitions) are represented as hyperreals (resp., sets of hyperreals). Their approach is mathematically very elegant, but very different from ours. Ours is conceptually simpler: we simply remove any measurability constraint by replacing the Borel σ -algebra with the power set. This comes at a cost: the loss of countable additivity of λ . This not only complicates the technical analysis, but generates a conceptual loss. In fact, in an independent paper, Khan, Qiao, Rath, and Sun [18] show that the existence

⁵In our specification, this would collapse to a Nash equilibrium.

⁶For a related concept and result see also Section 5 of Fudenberg and Kamada [13].

of Nash equilibria for any game with players' space $(T, \mathcal{T}, \lambda)$ is equivalent to the countable additivity of λ . Since the existence of Nash equilibria cannot be guaranteed with mere finite additivity, they study the existence of ε -Nash equilibria, thus overlapping our Corollary 2.

- (iii) The issue of modelling the players' space as a continuum or as a discrete space has also been discussed by Al-Najjar [2], who considers as competing models the continuum space $[0, 1]$ versus a dense countable grid of $[0, 1]$. This paper also shares some of the motivation coming from the Dubey-Shapley's remark on measurability (see Dubey and Shapley [11] as well as Khan and Sun [21]). Thus, in trying to build a link between these two conceptually equivalent approaches countable additivity is necessarily lost, as in our case. The main results of Al-Najjar [2] show that, under suitable conditions, the two approaches to modelling the players' space, that is, a continuum versus a discrete dense grid, are equivalent. In order to achieve this result, Al-Najjar shows that all his Nash-type equilibria, for his class of discrete games, can be purified. Compared to our work, Al-Najjar is not concerned with any other form of equilibrium other than Nash equilibria. Moreover, he establishes the existence of a form of ε -equilibrium for those discrete nonatomic games that arise as limits of proper sequences of finite-players games. Example 2 shows that for our more general class of games these ε -equilibria are not always guaranteed to exist.

We conclude by mentioning one more work. One of the important papers on nonatomic games which introduces a novel approach is Mas-Colell [25]. His approach is based on distributions of strategies,⁷ which allows for not considering strategy profiles. In this way, issues of measurability can be partially overridden in the proofs. It is an alternative framework which permits the discussion of players' negligibility. In this framework though, Shapley's observation would still apply and the assumption of countable additivity still seems to be playing a major role. Finally, we are not aware of refinements of and variations on this distributional concept of equilibrium.

Roadmap. In Sections 2 and 3, we formally introduce nonatomic players' spaces, nonatomic games with estimation feedback, and the definition of ε -estimated equilibrium whose existence is proven in Theorem 1. In Sections 3.1, 3.2, and 3.3, as a by-product, we obtain the existence of self-confirming, Nash, peer-confirming, and Berk–Nash ε -equilibria. Proofs are relegated to the appendices. In particular, in Appendix A.1, Lemma 1 generalizes Theorem 7 of Khan and Sun [20] which deals with the set of distributions induced by all the selections of a correspondence. In Appendix A.2, we provide a brief summary of how the

⁷This reformulation is connected to the distributional approach for Bayesian games with a continuum of types (see [25, Remarks 3 and 4]).

main proofs are carried out and prove all the results contained in the main text. In Appendix A.3, we conclude the paper with one of the authors explaining the origin of nonatomic games.

2 Nonatomic players' spaces

A *players' space* is a pair (T, λ) where T is a set of players and λ is a (finitely additive) *probability* on the power set of T .⁸ When $T = \mathbb{N}$, a fundamental class of probabilities that are not countably additive are *natural densities*, that is, probabilities λ such that

$$\lambda(E) = \lim_{k \rightarrow \infty} \frac{|E \cap \{1, \dots, k\}|}{k}$$

whenever the limit exists. As is well known, there are many natural densities and all of them satisfy the following property:

Strong continuity (Savage's nonatomicity) *For each $\varepsilon > 0$ there exists a finite partition $\{F_1, F_2, \dots, F_k\}$ of T such that $\lambda(F_i) < \varepsilon$ for all $i = 1, \dots, k$.*

Under strong continuity, any singleton (i.e., any single player) has measure 0 and for each $F \subseteq T$ and $\beta \in (0, 1)$ there exists $E \subseteq F$ such that $\lambda(E) = \beta\lambda(F)$.⁹ This is the class of probabilities introduced by Savage [29] when he solved De Finetti's open problem on the representation of qualitative probabilities (see also Samet and Schmeidler [28]).

3 Nonatomic games and their equilibria

Nonatomic games are games where each single player has no influence on the strategic interaction, but only the aggregate behavior of "large" sets of players can change the players' payoffs. Formally, a *nonatomic (anonymous) game* is a triplet $G = ((T, \lambda), A, u)$ where (T, λ) is the players' space, A is the space of players' actions/strategies and u is their profile of utilities.¹⁰ Below, we discuss in detail these mathematical objects and their interpretations.

- $A = \{1, \dots, n\}$ is the set of pure *strategies/actions*.
- $\Delta = \{x \in \mathbb{R}_+^n \mid \sum_{i=1}^n x_i = 1\}$ is the $n - 1$ dimensional simplex. We denote by d_Δ the distance on Δ induced by the Euclidean norm. This set represents all possible

⁸Recall that λ is a finitely additive probability if and only if λ is a positive finitely additive set function such that $\lambda(T) = 1$.

⁹See Maharam [24, Example 2.1 and Theorem 2] and Bhaskara Rao and Bhaskara Rao [7, Theorem 5.1.6 and Remark 5.1.7]. In this literature, natural densities are called density measures or density charges.

¹⁰In the paper, given a generic set B , we use the term profile to refer to a function from the set of players T to B . We will denote a profile by either $b : T \rightarrow B$ or by $b = (b_t)_{t \in T}$. The latter notation will allow us, with a small abuse, to treat $(b_t)_{t \in T}$ also as a set.

distributions of players' strategies. Note that an element in Δ can actually take two possible interpretations. In fact, given a player t and an element of Δ , this element can either be interpreted as a subjective belief of player t (in this case, we often denote it by β_t) or be interpreted as an objective distribution of players' strategies (in this case, we typically denote it by x).

- $u = (u_t)_{t \in T}$ is a profile of functions $u_t : A \times \Delta \rightarrow \mathbb{R}$. For each t in T , $u_t(a, \beta_t)$ represents the ex-ante *utility* of player t , when he chooses strategy a , if his belief about the distribution of opponents' strategies is β_t .

As mentioned in the Introduction, nonatomic games were first studied by Schmeidler [30]. In this paper, we consider a class of games which we term *nonatomic games with estimation feedback*. It has a richer structure and nonatomic games can be seen as a specific parametrization.

Formally, a nonatomic game with estimation feedback is a quintet $G = ((T, \lambda), A, u, (\Pi, \pi), f)$ where $((T, \lambda), A, u)$ is a nonatomic game defined as above, (Π, π) is a neighborhood structure, and f is a profile of estimation feedback functions which discipline the beliefs' formation of agents in equilibrium. Formally, we have that:

- (Π, π) is a *neighborhood structure* if and only if $\Pi = \{T_j\}_{j=1}^m$ is a finite cover of T whose elements have strictly positive measure and π is a function from T to $\{1, \dots, m\}$. In particular, each T_j is a nonempty subset of T such that $\lambda(T_j) > 0$ and $\cup_{j=1}^m T_j = T$. An important example of finite covers are finite partitions of the players' space. We interpret an element of Π , T_j , as the j -th subpopulation of T and for each $t \in T$ the value $\pi(t)$ will denote which subpopulation player t observes.¹¹
- $f = (f_t)_{t \in T}$ is a profile of (*estimation*) *feedback functions* $f_t : A \times \Delta \times \Delta \rightarrow [0, \infty)$. Each f_t is assumed to be such that for each $y \in \Delta$ there exists $x_y \in \Delta$ for which it holds that

$$f_t(a, x_y, y) = 0 \quad \forall a \in A \tag{1}$$

For each t in T , $f_t(a, \beta_t, x)$ represents a measure of consistency between the belief β_t (entertained by player t) and the actual distribution of players' strategies x within the subpopulation observed by player t , with the idea that the larger $f_t(a, \beta_t, x)$ is the greater is the discrepancy between the player's belief and the subpopulation actions' distribution x . In line with this interpretation, property (1) says that for each possible true model x there exists a belief β_t such that this discrepancy is minimal, no matter what action a is chosen by player t . To better understand (1), we next state a stronger

¹¹Despite being a natural requirement, we can dispense with the assumption that $t \in T_{\pi(t)}$. In other words, we do not need to assume that any player t belongs to the subpopulation he observes.

property which implies (1) and has a more immediate interpretation. In all our specifications, with the exception of (11), it will be satisfied: for each $t \in T$ and for each $a \in A$

$$x = y \implies f_t(a, x, y) = 0 \quad (2)$$

In words, this latter property says that discrepancy is minimal provided the belief β_t is indeed correct, that is $\beta_t = x$.¹² Under (1) or (2), we deliberately allow for the possibility that $f_t(a, \beta_t, x) = 0$, but $\beta_t \neq x$: a belief might be consistent with evidence, but still incorrect.

Finally, we need three extra mathematical objects:

- $\Sigma = A^T$ is the set of all functions σ from T to A . Each $\sigma \in \Sigma$ represents a *strategy profile* in which the generic player t chooses strategy $\sigma(t)$.
- Given $j \in \{1, \dots, m\}$, λ^j denotes the probability on the power set of T defined by

$$\lambda^j(E) = \frac{\lambda(E \cap T_j)}{\lambda(T_j)} \quad \forall E \subseteq T$$

In other words, λ^j is the players' conditional measure in the subpopulation j . Note that if λ is strongly continuous, so is each λ^j .

- Given $\sigma \in \Sigma$ and $j \in \{1, \dots, m\}$, $\lambda_\sigma^j \in \Delta$ is the distribution of σ on A in the j -th subpopulation, that is,

$$\lambda_\sigma^j = \left(\lambda^j(\{t \in T_j \mid \sigma(t) = a\}) \right)_{a \in A}$$

The vector λ_σ^j represents the *true* distribution of players' pure strategies in the j -th subpopulation,¹³ when they all play according to σ . When Π is trivial, that is, $\Pi = \{T\}$, then Π contains only one element and $\lambda = \lambda^1$. In this case, we write λ_σ in place of λ_σ^1 . Similarly, the vector λ_σ represents the true distribution of players' pure strategies in the entire population.

We can now introduce our most general concept of equilibrium. Its interpretation is postponed to the next three sections where we discuss three particular and important specifications (see also the Introduction).

¹²Note that (2) implies (1). Fix $t \in T$. For each $y \in \Delta$, set $x_y = y$. By (2), it follows that $f_t(a, x_y, y) = 0$ for all $a \in A$.

¹³By definition of λ_σ^j , note that

$$\lambda_\sigma^j = \left(\lambda^j(\{t \in T \mid \sigma(t) = a\}) \right)_{a \in A}$$

for all $j \in \{1, \dots, m\}$.

Definition 1 Let $\varepsilon \geq 0$. An ε -estimated equilibrium (in pure strategies) for the nonatomic game with estimation feedback $G = ((T, \lambda), A, u, (\Pi, \pi), f)$ is a strategy profile $\sigma \in \Sigma$ such that there exists a profile of beliefs $\beta \in \Delta^T$ satisfying

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ f_t(\sigma(t), \beta(t), \lambda_\sigma^{\pi(t)}) \leq \varepsilon \end{array} \right\} \right) = 1 \quad (3)$$

We are ready to state our main result.

Theorem 1 Let $G = ((T, \lambda), A, u, (\Pi, \pi), f)$ be a nonatomic game with estimation feedback and $\varepsilon > 0$. If λ is strongly continuous and $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument,¹⁴ then G has an ε -estimated equilibrium.

3.1 Self-confirming and Nash equilibria

An interesting class of nonatomic games with estimation feedback arises when the feedback function of player t is generated by a message function $m_t : A \times \Delta \rightarrow M$, where M is a metric space with distance d .¹⁵ For each t in T , $m_t(a, x)$ represents the *message* player t receives when he chooses strategy a and the distribution of players' strategies is x . In games with finitely many players, typically the message function m_t depends on the action chosen by the player and the profile of actions chosen by the opponents. Nevertheless, given our underlying assumption of anonymity, it seems natural to replace the latter with the actions' distribution in the population.

With this in mind, the next type of equilibrium models a situation in which the belief β_t adopted by each agent t in equilibrium is consistent/confirmed with/by the message received. More formally, β_t is such that the expected message $m_t(\sigma(t), \beta_t)$ is ε -close to the received message $m_t(\sigma(t), \lambda_\sigma)$.

We define a *nonatomic game with message feedback* to be a quartet $G = ((T, \lambda), A, u, m)$ where $((T, \lambda), A, u)$ is a nonatomic game and $m = (m_t)_{t \in T}$ is a profile of message functions. Note that a nonatomic game with *message* feedback can be mapped into a nonatomic game with *estimation* feedback. In fact, it is enough to consider (Π, π) to be trivial, that is

¹⁴We say that $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument if and only if for each $\varepsilon > 0$ there exists $\delta_\varepsilon > 0$ such that

$$d_\Delta(x, y) < \delta_\varepsilon \implies |f_t(a, \gamma, x) - f_t(a, \gamma, y)| < \varepsilon \quad \forall t \in T, \forall a \in A, \forall \gamma \in \Delta$$

In other words, the family of functions $\{f_t(a, \gamma, \cdot)\}_{t \in T, a \in A, \gamma \in \Delta}$ from Δ to $[0, \infty)$ is equicontinuous.

¹⁵To simplify notation, we assume that the message space is the same for all players. This is without loss of generality. We could have equivalently assumed that each player has his own message space M_t , and in the proofs embed this set into a larger common message space M . Our assumptions of equicontinuity on the message functions m_t (cf. Corollary 1) would seamlessly pass through the embedding as well.

$\Pi = \{T\}$, and set the profile of feedback functions to be such that:¹⁶

$$f_t(a, x, y) = d(m_t(a, x), m_t(a, y)) \quad \forall t \in T, \forall a \in A, \forall x, y \in \Delta \quad (4)$$

It can be seen immediately that each f_t satisfies (2), and thus (1). We can define our concept of self-confirming ε -equilibrium which we discuss below.

Definition 2 *Let $\varepsilon \geq 0$. A self-confirming ε -equilibrium (in pure strategies) for the nonatomic game with message feedback $G = ((T, \lambda), A, u, m)$ is a strategy profile $\sigma \in \Sigma$ such that there exists a profile of beliefs $\beta \in \Delta^T$ satisfying*

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d(m_t(\sigma(t), \beta(t)), m_t(\sigma(t), \lambda_\sigma)) \leq \varepsilon \end{array} \right\} \right) = 1 \quad (5)$$

In words, a strategy profile $\sigma \in \Sigma$ is a self-confirming ε -equilibrium (ε -SCE) if and only if

1. Almost all players best-respond to their beliefs (optimality);
2. Beliefs are not significantly refuted by what they can observe (ε -confirmation).

As noted in the Introduction, self-confirming equilibria were introduced for games with finitely many players by Battigalli [6] and Fudenberg and Levine [14], and also ε -confirmation was introduced by Battigalli [6] and Kalai and Lehrer [16] and [17]. To the best of our knowledge, the above definition of ε -equilibrium seems to be novel for nonatomic games and also natural (cf. the related literature section). Furthermore, it encompasses the notions of self-confirming equilibrium and ε -Nash equilibrium (a fortiori, Nash equilibrium). To see this, we begin by observing that if $\varepsilon = 0$ and $m_t : A \times \Delta \rightarrow \Delta$ is such that

$$m_t(a, x) = x \quad \forall t \in T, \forall a \in A, \forall x \in \Delta \quad (6)$$

that is, $(M, d) = (\Delta, d_\Delta)$ and feedback is (statistically) *perfect*, then (5) becomes

$$\lambda(\{t \in T \mid u_t(\sigma(t), \lambda_\sigma) \geq u_t(a, \lambda_\sigma) \quad \forall a \in A\}) = 1$$

which means that σ is a Nash equilibrium. In this case, beliefs are not only *perfectly* consistent with observations but also correct. Maintaining the perfect feedback assumption (6), but allowing for $\varepsilon > 0$, (5) becomes

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d_\Delta(\beta(t), \lambda_\sigma) \leq \varepsilon \end{array} \right\} \right) = 1$$

¹⁶In this case, note that π can only take one value.

Under a suitable assumption of continuity of u (see Corollary 2 and its proof), we can show that σ is an ε -Nash equilibrium for some suitable $\hat{\varepsilon} > 0$, that is,

$$\lambda(\{t \in T \mid u_t(\sigma(t), \lambda_\sigma) \geq u_t(a, \lambda_\sigma) - \hat{\varepsilon} \quad \forall a \in A\}) = 1$$

The intuition is simple: if beliefs are “close” to the true distribution, players are not far from objective maximization.

Finally, if we remove perfect feedback but maintain $\varepsilon = 0$, (5) becomes

$$\lambda\left(\left\{t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ m_t(\sigma(t), \beta(t)) = m_t(\sigma(t), \lambda_\sigma) \end{array} \right\}\right) = 1$$

which is arguably the nonatomic anonymous games counterpart of the definition of self-confirming equilibrium (SCE).

Starting with ε -estimated equilibria, most of our analysis deals with the case in which $\varepsilon > 0$. There are two reasons why we do so. First, conceptually, $\varepsilon > 0$ allows beliefs to be only imperfectly confirmed, mirroring the fact that players’ observations might be noisy and learning slow. Second, self-confirming equilibria and Nash equilibria might not exist, as the following examples show. In a nutshell, Example 1 provides an instance where Nash and SCE equilibria do not exist, but their ε -versions do. Example 2 provides an instance where ε -uniform Nash equilibria *à la* Al-Najjar [2] do not exist, but standard ε -Nash equilibria do.

Example 1 The next example builds on Khan, Qiao, Rath, and Sun [18].¹⁷ Consider $T = \mathbb{N}$ and let λ be a natural density. Consider two strategies, that is, $A = \{1, 2\}$. Assume that for each $t \in T$

$$u_t(a, x) = \begin{cases} \frac{1}{t} - x_1 & a = 1 \\ x_1 - \frac{1}{t} & a = 2 \end{cases} \quad \forall x \in \Delta$$

Let $m_t = u_t$ for all $t \in T$. This amounts to the standard assumption of mere payoff observability. Assume that $\sigma \in \Sigma$ is an SCE, that is, there exists $\beta \in \Delta^T$ such that

$$\lambda(\{t \in T \mid u_t(\sigma(t), \lambda_\sigma) = u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A\}) = 1$$

For ease of notation, set $\lambda_\sigma = x$ and define the set of “optimizing” players by

$$O = \{t \in T \mid u_t(\sigma(t), \lambda_\sigma) = u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A\}$$

We have two cases:

¹⁷The example of Khan, Qiao, Rath, and Sun [18] seems to be the first one in the literature to exhibit a well-behaved nonatomic game which does not have any Nash equilibrium, be it pure or mixed.

1. $x_1 > 0$. Since λ is a natural density and O has mass 1, then O is infinite. Thus, there exists $\bar{t} \in \mathbb{N}$ such that $\frac{2}{t} - x_1 < 0$ for all $t \in O \cap \{1, \dots, \bar{t}\}^c$. Consider $t \in O \cap \{1, \dots, \bar{t}\}^c \neq \emptyset$. By contradiction, assume that $\sigma(t) = 1$. The SCE conditions imply that

$$\frac{1}{t} - x_1 = \frac{1}{t} - \beta(t)_1 \geq \beta(t)_1 - \frac{1}{t}$$

yielding that $0 \leq \beta(t)_1 \leq \frac{2}{t} - x_1 < 0$, a contradiction. Since t was arbitrarily chosen in $O \cap \{1, \dots, \bar{t}\}^c$, it follows that $\sigma(t) = 2$ for all $t \in O \cap \{1, \dots, \bar{t}\}^c$. Since λ is a natural density and O and $O \cap \{1, \dots, \bar{t}\}^c$ differ by a finite set $\lambda(O \cap \{1, \dots, \bar{t}\}^c) = 1$, we have that $\lambda_\sigma = x$ is such that $x_2 = 1$, a contradiction with $0 = 1 - x_2 = x_1 > 0$.

2. $x_1 = 0$. Consider $t \in O$. By contradiction, assume that $\sigma(t) = 2$. The SCE conditions imply that

$$x_1 - \frac{1}{t} = \beta(t)_1 - \frac{1}{t} \geq \frac{1}{t} - \beta(t)_1$$

yielding that $0 = x_1 = \beta(t)_1$ and $0 \geq \frac{2}{t} > 0$, a contradiction. Since t was arbitrarily chosen in O , $\sigma(t) = 1$ for all $t \in O$, yielding that $\lambda_\sigma = x$ is such that $x_1 = 1$, a contradiction with $x_1 = 0$.

To sum up, we have just shown that the nonatomic game with message feedback above does not have any self-confirming equilibrium and, in particular, any Nash equilibrium.¹⁸ This happens despite the fact that the profile of message functions is extremely well-behaved being $m = (m_t)_{t \in T}$ equicontinuous with respect to the second argument (cf. Corollary 1).¹⁹ At the same time, consider $\varepsilon > 0$. Let $\bar{t} \in \mathbb{N}$ be such that $\min\{1, \varepsilon\} > \frac{1}{t}$ for all $t \in \mathbb{N}$ such that $t > \bar{t}$. Set $\bar{\varepsilon} = \min\{1, \varepsilon\}$. Consider a strategy profile $\sigma \in \Sigma$ and a belief profile $\beta \in \Delta^T$ such that $\sigma(t) = 2$ and $\beta(t)_1 = \frac{\bar{\varepsilon} + \frac{1}{t}}{2} \in (\frac{1}{t}, \bar{\varepsilon}) \subseteq (0, 1)$ for all $t \in \mathbb{N}$ such that $t > \bar{t}$.

¹⁸Two extra observations are in order:

a. In the nonatomic game above, SCE equilibria and Nash equilibria coincide. This is by chance, as the next point shows.

b. Khan, Qiao, Rath, and Sun [18] consider $T = \mathbb{N}$ and let λ be a natural density. They assume $A = \{1, 2\}$ and \hat{u} to be such that for each $t \in T$

$$\hat{u}_t(a, x) = \begin{cases} \frac{1}{t} - x_1 & a = 1 \\ 0 & a = 2 \end{cases} \quad \forall x \in \Delta$$

With similar arguments, they prove that the nonatomic game $((T, \lambda), A, \hat{u})$ does not have any Nash equilibrium. At the same time, if we consider the augmented nonatomic game with message feedback $((T, \lambda), A, \hat{u}, m)$ where $m_t = \hat{u}_t$ for all $t \in T$, then we can show that there exists an SCE equilibrium. In fact, if $\sigma \in \Sigma$ is such that $\sigma(t) = 2$ for all $t \in T$, by setting $\beta \in \Delta^T$ such that $\beta(t)_1 = 1$ for all $t \in T$, we obtain the result.

¹⁹Indeed, note that for each $\varepsilon > 0$ we can set $\delta_\varepsilon = \varepsilon$ and get

$$d_\Delta(x, y) < \varepsilon \implies |m_t(a, x) - m_t(a, y)| = |x_1 - y_1| \leq d_\Delta(x, y) < \varepsilon \quad \forall t \in T, \forall a \in A$$

Since $\{1, \dots, \bar{t}\}$ is finite and λ is a natural density, we have that $\lambda_\sigma = x$ is such that $x_2 = 1$, that is, $x_1 = 0$. It follows that for each $t \in \{1, \dots, \bar{t}\}^c$

$$|m_t(\sigma(t), \beta(t)) - m_t(\sigma(t), \lambda_\sigma)| = \left| \frac{\bar{\varepsilon} + \frac{1}{t}}{2} - \frac{1}{t} - x_1 + \frac{1}{t} \right| = \frac{\bar{\varepsilon} + \frac{1}{t}}{2} < \bar{\varepsilon} \leq \varepsilon$$

and

$$u_t(\sigma(t), \beta(t)) = \beta(t)_1 - \frac{1}{t} = \frac{\bar{\varepsilon} - \frac{1}{t}}{2} > 0 > \frac{1}{t} - \beta(t)_1 = u_t(1, \beta(t))$$

Since $\{1, \dots, \bar{t}\}^c$ has mass 1, we can conclude that $\sigma \in \Sigma$ is an ε -SCE. \blacktriangle

Example 2 Al-Najjar [2] (cf. the Introduction) also deals with the lack of countable additivity and studies the following equilibrium: a strategy $\sigma \in \Sigma$ is an Al-Najjar equilibrium (in pure strategies) if and only if for each $\varepsilon > 0$

$$\lambda(\{t \in T \mid u_t(\sigma(t), \lambda_\sigma) \geq u_t(a, \lambda_\sigma) - \varepsilon \quad \forall a \in A\}) > 1 - \varepsilon \quad (7)$$

We next show that also these equilibria might fail to exist. In what follows, it will often be useful to set

$$O_\varepsilon = \{t \in T \mid u_t(\sigma(t), \lambda_\sigma) \geq u_t(a, \lambda_\sigma) - \varepsilon \quad \forall a \in A\}$$

Two observations are in order. First, compared to the ε -Nash equilibria we study (Corollary 2), the key difference is that, in our case, σ might depend on the given ε , while in Al-Najjar's case, σ must work with any ε . In particular, one can easily show that:²⁰ $\sigma \in \Sigma$ is an Al-Najjar equilibrium if and only if for each $\varepsilon > 0$

$$\lambda(\{t \in T \mid u_t(\sigma(t), \lambda_\sigma) \geq u_t(a, \lambda_\sigma) - \varepsilon \quad \forall a \in A\}) = 1$$

Second, by taking the intersection of the sets $O_{1/n}$, this allows us to conclude easily that an Al-Najjar equilibrium is a Nash equilibrium, provided λ is countably additive. We consider the nonatomic game $((T, \lambda), A, \tilde{u})$ where (T, λ) and A are as in Example 1 and for each $t \in T$

$$\tilde{u}_t(a, x) = \begin{cases} \frac{1}{t} - x_1 & a = 1 \text{ and } x_1 > 0 \\ x_1 - \frac{1}{t} & a = 2 \text{ and } x_1 > 0 \\ 1 & a = 1 \text{ and } x_1 = 0 \\ \frac{1}{t} & a = 2 \text{ and } x_1 = 0 \end{cases} \quad \forall x \in \Delta$$

Assume that $\sigma \in \Sigma$ satisfies (7). For ease of notation, set $\lambda_\sigma = x$. As before, we have two cases:

²⁰It is easy to see that if $0 < \varepsilon < \varepsilon'$, then $O_\varepsilon \subseteq O_{\varepsilon'}$, thus

$$\lambda(O_{\varepsilon'}) \geq \lambda(O_\varepsilon) > 1 - \varepsilon \quad \forall \varepsilon' > 0, \forall \varepsilon \in (0, \varepsilon')$$

yielding that $\lambda(O_{\varepsilon'}) = 1$ for all $\varepsilon' > 0$.

1. $x_1 > 0$. Fix $\varepsilon > 0$. Since λ is a natural density, the set O_ε has mass 1, and $\lambda(\{t \in T \mid \sigma(t) = 1\}) > 0$, we have that $O_\varepsilon \cap \{t \in T \mid \sigma(t) = 1\}$ is infinite. Since ε was arbitrarily chosen, this implies that we can construct a strictly increasing sequence $\{t_k\}_{k \in \mathbb{N}} \subseteq \mathbb{N}$ such that $t_k \in O_{1/k} \cap \{t \in T \mid \sigma(t) = 1\}$ for all $k \in \mathbb{N}$. Since $t_k \in O_{1/k}$, $\sigma(t_k) = 1$, and $x_1 > 0$, we have that for each $k \in \mathbb{N}$

$$\frac{1}{t_k} - x_1 \geq x_1 - \frac{1}{t_k} - \frac{1}{k} \implies 0 < x_1 \leq \frac{1}{t_k} + \frac{1}{2k}$$

By passing to the limit, we obtain that $0 < x_1 \leq 0$, a contradiction.

2. $x_1 = 0$. Fix $\varepsilon > 0$. Since λ is a natural density, the set O_ε has mass 1, and $\lambda(\{t \in T \mid \sigma(t) = 2\}) > 0$, we have that $O_\varepsilon \cap \{t \in T \mid \sigma(t) = 2\}$ is infinite. Since ε was arbitrarily chosen, this implies that we can construct a strictly increasing sequence $\{t_k\}_{k \in \mathbb{N}} \subseteq \mathbb{N}$ such that $t_k \in O_{1/k} \cap \{t \in T \mid \sigma(t) = 2\}$ for all $k \in \mathbb{N}$. Since $t_k \in O_{1/k}$, $\sigma(t_k) = 2$, and $x_1 = 0$, we have that for each $k \in \mathbb{N}$

$$\frac{1}{t_k} \geq 1 - \frac{1}{k}$$

By passing to the limit, we obtain that $0 \geq 1$, a contradiction.

To sum up, we have just shown that the nonatomic game $((T, \lambda), A, \tilde{u})$ does not have any equilibrium as defined in (7). At the same time, it is not hard to see that this game admits an ε -Nash equilibrium for every $\varepsilon > 0$. One way to observe this is to consider the augmented game $((T, \lambda), A, \tilde{u}, m)$ in which each player has perfect statistical feedback: that is

$$m_t(a, x) = x \quad \forall t \in T, \forall a \in A, \forall x \in \Delta$$

Since $m = (m_t)_{t \in T}$ is equicontinuous with respect to the second argument (cf. Corollary 1), we have that for each $\varepsilon > 0$ there exists an ε -SCE. Given $\varepsilon \in (0, 1)$, it can immediately be proved that a strategy profile σ is an ε -SCE if and only if $\lambda(\{t \in T : \sigma(t) = 1\}) \in [0, \varepsilon/\sqrt{2}]$. Given our choice of m , following the intuition that “if beliefs are close to the true distribution, players are not far from objective maximization”, we can prove that, given $\varepsilon \in (0, 1)$, if σ is an $\frac{\varepsilon}{2\sqrt{2}}$ -SCE and $\lambda(\{t \in T : \sigma(t) = 1\}) > 0$, then σ is an ε -Nash equilibrium. In other words, $((T, \lambda), A, \tilde{u})$ does not have any equilibrium as defined in (7), but for each $\varepsilon > 0$ it has an ε -Nash equilibrium.²¹ ▲

We are ready to state the main results of this section.

²¹Since λ is strongly continuous, note that, given $\varepsilon \in (0, 1)$, we can always find $\sigma \in \Sigma$ such that $\lambda(\{t \in T : \sigma(t) = 1\}) \in (0, \varepsilon/4]$. In other words, in light of the above characterization, we can always find a strategy profile σ which is an $\frac{\varepsilon}{2\sqrt{2}}$ -SCE such that $\lambda(\{t \in T : \sigma(t) = 1\}) > 0$.

Corollary 1 *Let $G = ((T, \lambda), A, u, m)$ be a nonatomic game with message feedback and $\varepsilon > 0$. If λ is strongly continuous and $m = (m_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument,²² then G has an ε -SCE.*

In particular, under the assumption of payoff observability, that is, $m_t(a, x) = u_t(a, x)$ for all $t \in T, a \in A, x \in \Delta$, Corollary 1 yields that if λ is strongly continuous and $u = (u_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument, then there exists an ε -SCE strategy profile $\sigma \in \Sigma$ such that

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ |u_t(\sigma(t), \beta(t)) - u_t(\sigma(t), \lambda_\sigma)| \leq \varepsilon \end{array} \right\} \right) = 1$$

where $\beta \in \Delta^T$. In this case, the objective observed payoff substantially matches the expected one. Building on Corollary 1 and following a similar intuition, we also obtain the existence of ε -Nash equilibria.

Corollary 2 *Let $G = ((T, \lambda), A, u)$ be a nonatomic game and $\varepsilon > 0$. If λ is strongly continuous and $u = (u_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument,²³ then G has an ε -Nash equilibrium, that is, there exists a strategy profile $\sigma \in \Sigma$ such that*

$$\lambda(\{t \in T \mid u_t(\sigma(t), \lambda_\sigma) \geq u_t(a, \lambda_\sigma) - \varepsilon \quad \forall a \in A\}) = 1$$

At this point, the reader might wonder how restrictive are our assumptions of equicontinuity.²⁴ At first sight, it might appear that the degree of similarity among players, imposed by a measurable structure as in the original framework of Schmeidler, is here replaced by equicontinuity. The following example should clarify that this is far from being the case.

Example 3 Assume that players have expected-utility like preferences, namely, for each $t \in T$

$$u_t(a, x) = \sum_{b \in A} v_t(a, b) x_b \quad \forall a \in A, \forall x \in \Delta$$

²²We say that $m = (m_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument if and only if for each $\varepsilon > 0$ there exists $\delta_\varepsilon > 0$ such that

$$d_\Delta(x, y) < \delta_\varepsilon \implies d(m_t(a, x), m_t(a, y)) < \varepsilon \quad \forall t \in T, \forall a \in A$$

²³We say that $u = (u_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument if and only if for each $\varepsilon > 0$ there exists $\delta_\varepsilon > 0$ such that

$$d_\Delta(x, y) < \delta_\varepsilon \implies |u_t(a, x) - u_t(a, y)| < \varepsilon \quad \forall t \in T, \forall a \in A$$

²⁴See also the discussion following Corollary 3.

where $v_t : A \times A \rightarrow \mathbb{R}$. As is well known, each v_t can be normalized to be taking values in $[0, 1]$, without altering the player's preferences. In light of this, an immediate application of the Cauchy-Schwarz inequality yields that

$$|u_t(a, x) - u_t(a, y)| \leq \sqrt{n}d_\Delta(x, y) \quad \forall t \in T, \forall a \in A$$

proving that $u = (u_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument. Thus, preferences can be extremely different within the above class and yet satisfy equicontinuity. \blacktriangle

As mentioned in the Introduction, Khan, Qiao, Rath, and Sun [18] showed that in the absence of countable additivity the existence of Nash equilibria is not guaranteed. They also reported an independent result of existence of an ε -Nash equilibrium. Their definition is weaker than ours. In their case, a strategy profile $\sigma \in \Sigma$ is an ε -equilibrium if and only if

$$\lambda(\{t \in T \mid u_t(\sigma(t), \lambda_\sigma) \geq u_t(a, \lambda_\sigma) - \varepsilon \quad \forall a \in A\}) \geq 1 - \varepsilon$$

3.2 Peer-confirming equilibria

Lipnowski and Sadler [23] propose a notion of equilibrium in which players best-respond to beliefs which are required to be correct only when it comes to the behavior of opponents who belong to the same neighborhood. Moreover, they further require that this is correctly and commonly believed by players. Formally, the collection of neighborhoods is a partition of the players in terms of the connected components of an underlying undirected network. They study games with finitely many players. For simultaneous-move games, peer-confirming equilibrium is an example of rationalizable self-confirming equilibrium (see also Rubinstein and Wolinsky [27] as well as Fudenberg and Kamada [13]). In what follows, we dispense with the assumption of correct and common belief. Moreover, given anonymity, we require that players' observations are only about the actions' distributions in the subpopulation they face.

We define a *nonatomic game with a neighborhood structure* to be a quartet $G = ((T, \lambda), A, u, (\Pi, \pi))$ where $((T, \lambda), A, u)$ is a nonatomic game and (Π, π) is a neighborhood structure. Note that a nonatomic game with a *neighborhood structure* can be mapped into a nonatomic game with *estimation feedback*. In fact, it is enough to set the profile of feedback functions to be such that

$$f_t(a, x, y) = d_\Delta(x, y) \quad \forall t \in T, \forall a \in A, \forall x, y \in \Delta \quad (8)$$

It can be seen immediately that each f_t satisfies (2), and thus (1). We can define the version of peer-confirming ε -equilibrium that we analyze below.

Definition 3 *Let $\varepsilon \geq 0$. A peer-confirming ε -equilibrium (in pure strategies) for the nonatomic game with a neighborhood structure $G = ((T, \lambda), A, u, (\Pi, \pi))$ is a strategy profile*

$\sigma \in \Sigma$ such that there exists a profile of beliefs $\beta \in \Delta^T$ satisfying

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d_\Delta(\beta(t), \lambda_\sigma^{\pi(t)}) \leq \varepsilon \end{array} \right\} \right) = 1 \quad (9)$$

In words, a strategy profile $\sigma \in \Sigma$ is a peer-confirming ε -equilibrium (ε -PCE) if and only if

1. Almost all players best-respond to their beliefs (optimality);
2. Beliefs are almost correct in terms of the subpopulation observed (ε -neighborhood confirmation).

Corollary 3 *Let $G = ((T, \lambda), A, u, (\Pi, \pi))$ be a nonatomic game with a neighborhood structure and $\varepsilon > 0$. If λ is strongly continuous, then G has an ε -PCE.*

It is important to note how the corollary above does not require any extra property of continuity. For, in such a case feedback is perfect, when restricted to each subpopulation, and action independent, automatically satisfying the requirement of equicontinuity in Theorem 1. Conceptually, this confirms that, in contrast with measurability assumptions, our properties of equicontinuity do not impose automatically that “close players” have similar preferences/behavior (cf. Example 3).

3.3 Berk–Nash equilibria

Esponda and Pouzo [12] propose a notion of equilibrium that allows for players’ beliefs to be possibly misspecified. It is a different way, compared to self-confirming equilibria, to allow for potentially incorrect beliefs in equilibrium. They term their notion of equilibrium Berk–Nash. Berk–Nash equilibria are based on the assumption that each player has a set of probabilistic models over the payoff-relevant features, in our case $\{Q_t\}_{t \in T} \subseteq \Delta^\circ$,²⁵ and:

1. All players best-respond to their beliefs (optimality);
2. Each player’s belief is restricted to be the best fit (in terms of Kullback–Leibler distance) among the set of beliefs he considers possible.

In our setup, this would mean that each player t has a (possibly misspecified) set of models $Q_t \subseteq \Delta^\circ$. A strategy profile $\sigma \in \Sigma$ is a Berk–Nash equilibrium if and only if there

²⁵As usual, Δ° denotes the set

$$\{x \in \Delta \mid x_i > 0 \quad \forall i \in \{1, \dots, n\}\}$$

In other words, Δ° is the relative interior of Δ .

exists a profile of beliefs $\beta \in \Delta^T$ such that the set of all players that satisfy the following two conditions has full measure:²⁶

1. $u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t))$ for all $a \in A$;
2. $\beta(t) \in \operatorname{argmin}_{z \in Q_t} K(\lambda_\sigma || z)$ (where K is the Kullback–Leibler distance).

In what follows, we offer a more general version for nonatomic games of the above equilibrium. In order to do so, we define a *nonatomic game with model misspecification* to be a quintet $G = ((T, \lambda), A, u, \mathcal{Q}, D)$ where

- a. $((T, \lambda), A, u)$ is a nonatomic game;
- b. $\mathcal{Q} = (Q_t)_{t \in T}$ is a profile of sets of actions' distributions, that is, Q_t is a nonempty, compact, and convex subset of Δ^o for all $t \in T$;
- c. $D : \Delta \times \Delta^o \rightarrow [0, \infty)$ is a statistical divergence, that is, a jointly convex and continuous function such that for each $x, y \in \Delta^o$

$$x = y \iff D(x || y) = 0. \quad (10)$$

The next example describes a class of widely used statistical divergences.

Example 4 The most classic statistical divergences are ϕ -divergences which have the form

$$D_\phi(x || z) = \sum_{i=1}^n z_i \phi\left(\frac{x_i}{z_i}\right)$$

where ϕ is a positive, continuous, strictly convex function on \mathbb{R}_+ such that $\phi(1) = 0$. For example, for $\phi(s) = s \log s - s + 1$,²⁷ D_ϕ is the Kullback–Leibler distance, for $\phi(s) = (s - 1)^2 / 2$, D_ϕ is the χ^2 -distance, and for $\phi(s) = (\sqrt{s} - 1)^2$, D_ϕ is the Hellinger distance. In all these specifications, D_ϕ satisfies (10) and it is jointly convex and continuous. \blacktriangle

Note that a nonatomic game with *model misspecification* can be mapped into a nonatomic game with *estimation feedback*. In fact, it is enough to consider (Π, π) to be trivial, that is $\Pi = \{T\}$, and set the profile of feedback functions to be such that:²⁸

$$f_t(a, x, y) = d_\Delta(x, \operatorname{argmin}_{z \in Q_t} D(y || z)) \quad \forall t \in T, \forall a \in A, \forall x, y \in \Delta \quad (11)$$

²⁶Compared to Esponda and Pouzo [12], we do not assume that players' are expected utility and have a prior μ over $\operatorname{argmin}_{z \in Q_t} K(\lambda_\sigma || z)$. In other words, players are only allowed to consider degenerate priors. A priori, this makes it more difficult to obtain an existence result. Moreover, we are also not considering any extra form of feedback.

²⁷Here, it is assumed implicitly that $\phi(0) = 1$ which is obtained by taking the limit for $s \rightarrow 0$.

²⁸In this case, note that π can only take one value. Moreover, when $x \in \Delta$ and Y is a nonempty subset of Δ , $d_\Delta(x, Y)$ denotes the distance of x from the set Y , that is,

$$d_\Delta(x, Y) = \inf_{y \in Y} d_\Delta(x, y)$$

In our case, $Y = \operatorname{argmin}_{z \in Q_t} D(y || z)$.

It is not hard to show that each f_t satisfies (1), but might fail to satisfy (2). We can define our version of Berk–Nash ε -equilibrium which we discuss below.

Definition 4 *Let $\varepsilon \geq 0$. A Berk–Nash ε -equilibrium (in pure strategies) for the nonatomic game with model misspecification $G = ((T, \lambda), A, u, \mathcal{Q}, D)$ is a strategy profile $\sigma \in \Sigma$ such that there exists a profile of beliefs $\beta \in \Delta^T$ satisfying*

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d_\Delta(\beta(t), \operatorname{argmin}_{z \in Q_t} D(\lambda_\sigma || z)) \leq \varepsilon \end{array} \right\} \right) = 1 \quad (12)$$

Note that a strategy profile $\sigma \in \Sigma$ is a Berk–Nash ε -equilibrium (ε -BNE) if and only if

1. Almost all players best-respond to their beliefs (optimality);
2. Beliefs are ε -close to the set of probabilistic models which are the best fit in the primitive set Q_t of the realized distribution (ε -fit).

Although prima facie they might appear similar, the notion of ε -BNE is conceptually and formally very different from that of ε -SCE.²⁹ The next result proves that, under suitable conditions, the former type of equilibria always exists. To do so, we need a last piece of notation. Given $\delta > 0$, we denote

$$\Delta_\delta = \{x \in \Delta \mid x_i \geq \delta \quad \forall i \in \{1, \dots, n\}\}$$

In words, Δ_δ is the set of all actions' distributions which are uniformly bounded away from zero by δ .

Corollary 4 *Let $G = ((T, \lambda), A, u, \mathcal{Q}, D)$ be a nonatomic game with model misspecification and $\varepsilon > 0$. If λ is strongly continuous, D is strictly convex in the second argument, and there exists $\delta > 0$ such that $Q_t \subseteq \Delta_\delta$ for all $t \in T$, then G has an ε -BNE.*

Remark 1 Three observations are in order:

1. In Definition 4, we allow each player's equilibrium belief $\beta(t)$ to be possibly outside the set Q_t . This could be interpreted as allowing for the possibility that each player fears model misspecification and willingly considers probability models that are outside his posited set Q_t (see Cerreia-Vioglio, Hansen, Maccheroni, and Marinacci [10]). At the same time, we could have considered the following more stringent definition of ε -BNE where this is not allowed. In this case, $\sigma \in \Sigma$ would be an ε -BNE if and only if there exists a profile of beliefs $\beta \in \Delta^T$ satisfying

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d_\Delta(\beta(t), \operatorname{argmin}_{z \in Q_t} D(\lambda_\sigma || z)) \leq \varepsilon \text{ and } \beta(t) \in Q_t \end{array} \right\} \right) = 1$$

²⁹The two equilibrium notions are distinct, but share some overlap (see Esponda and Pouzo [12]).

Under the same exact assumptions of Corollary 4, we can show that also these ε -equilibria exist.

2. Our results do not directly apply to the case in which D is the Kullback–Leibler distance K . In fact, in this case, $K(x|\cdot)$ might fail to be strictly convex.³⁰ At the same time, any perturbation $\kappa > 0$ of a statistical divergence D , that is $D + \kappa d_{\Delta}^2$, is a statistical divergence and satisfies the condition of strict convexity in Corollary 4.
3. The assumption “there exists $\delta > 0$ such that $Q_t \subseteq \Delta_{\delta}$ for all $t \in T$ ” is equivalent to the condition “each Q that belongs to the Hausdorff distance closure of \mathcal{Q} is a subset of Δ° ”. In other words, it is an assumption of relative compactness. \blacktriangle

A Appendix

In what follows, we first provide the proofs of the results in the main text and then conclude with one of the authors explaining the origin of nonatomic games. We begin with Appendix A.1 where we discuss a result which is key in proving Theorem 1. Appendix A.2 contains the remaining proofs. In a nutshell, this latter section is divided into two parts. First, we deal with the proof of existence of ε -estimated equilibria. Second, we prove the existence of ε -SCE, ε -NE, ε -PCE, and ε -BNE by showing that they are all particular cases or consequence of the existence of ε -estimated equilibria.

In the appendix, the vector space we use is the cartesian product of m copies of \mathbb{R}^n , that is $(\mathbb{R}^n)^m$, where n is given by the cardinality of the space of actions A and m is given by the cardinality of the neighborhood structure (Π, π) . We denote the elements of $(\mathbb{R}^n)^m$ by bold letters, that is \mathbf{x} and \mathbf{y} , while x_j will be the vector in \mathbb{R}^n which is the j -th component of \mathbf{x} . If $m = 1$, then we denote \mathbf{x} and \mathbf{y} simply by x and y . We endow $(\mathbb{R}^n)^m$ with the topology induced by the norm $\|\mathbf{x}\| = \max_{j \in \{1, \dots, m\}} \|x_j\|_2$ where $\|\cdot\|_2$ is the Euclidean norm. Finally, we denote the cartesian product of m copies of Δ by Δ^m . Note that Δ^m is a nonempty, convex, and compact subset of $(\mathbb{R}^n)^m$ and we endow it with the distance induced by $\|\cdot\|$.

A.1 A key general result

The next lemma uses the terminology of Bhaskara Rao and Bhaskara Rao [7]. Before discussing it, we need a piece of notation which will turn out to be useful in our later analysis. If T and A are two generic nonempty sets and $\Gamma : T \rightrightarrows A$ is a (nonempty valued) correspondence, we denote by $\text{Sel}(\Gamma)$ the set of all selections of Γ , that is, the set of all functions $\gamma : T \rightarrow A$ such that $\gamma(t) \in \Gamma(t)$ for all $t \in T$. Just for this section, \mathcal{T} is an arbitrary σ -algebra of subsets of T .³¹ Finally, given a \mathcal{T} -measurable map $\gamma : T \rightarrow A$ and a probability

³⁰In fact, $K(x|\cdot)$ is strictly convex if $x \in \Delta^{\circ}$, but it might fail to be so if $x \in \Delta \setminus \Delta^{\circ}$.

³¹In the rest of the paper, \mathcal{T} is the power set.

$\mu : \mathcal{T} \rightarrow [0, 1]$, recall that

$$\mu_\gamma = (\mu(\{t \in T \mid \gamma(t) = a\}))_{a \in A}$$

Lemma 1 *Let (T, \mathcal{T}) be a measurable space, A a finite set with n elements, and $\boldsymbol{\lambda} = (\lambda^1, \dots, \lambda^m)$ a vector of strongly continuous probabilities on \mathcal{T} . If $\Gamma : T \rightrightarrows A$ is a correspondence, then*

$$\{\boldsymbol{\lambda}_\gamma = (\lambda_\gamma^1, \dots, \lambda_\gamma^m) \mid \gamma \in \text{Sel}(\Gamma) \text{ and } \gamma \text{ is } \mathcal{T}\text{-measurable}\}$$

is a convex subset of Δ^m .

Proof. If $\phi, \gamma \in \text{Sel}(\Gamma)$ and are \mathcal{T} -measurable, for each $\alpha \in (0, 1)$, we want to construct $\psi \in \text{Sel}(\Gamma)$ which is \mathcal{T} -measurable and such that $\boldsymbol{\lambda}_\psi = \alpha \boldsymbol{\lambda}_\phi + (1 - \alpha) \boldsymbol{\lambda}_\gamma$.

Set $S_{ij} = \phi^{-1}(i) \cap \gamma^{-1}(j)$ for all $i, j \in A$. Then $\{S_{ij}\}_{i, j \in A}$ forms a partition of T (with possibly some empty elements) and all its elements belong to \mathcal{T} , because $\phi^{-1}(i), \gamma^{-1}(j) \in \mathcal{T}$ for all $i, j \in A$.

Since $\lambda^1, \dots, \lambda^m$ are strongly continuous and \mathcal{T} is a σ -algebra, for any $i, j \in A$, there are $T_{ij}, U_{ij} \in \mathcal{T}$ such that $S_{ij} = T_{ij} \sqcup U_{ij}$,³² $\boldsymbol{\lambda}(T_{ij}) = \alpha \boldsymbol{\lambda}(S_{ij})$, and $\boldsymbol{\lambda}(U_{ij}) = (1 - \alpha) \boldsymbol{\lambda}(S_{ij})$. This is trivial if S_{ij} is empty, else set

$$\begin{aligned} \mathcal{T}_{ij} &= \mathcal{T} \cap S_{ij} \\ \lambda_{ij}^k(S) &= \lambda^k(S) \quad \forall S \in \mathcal{T}_{ij}, \forall k = 1, \dots, m \end{aligned}$$

and notice that $\lambda_{ij}^1, \dots, \lambda_{ij}^m$ are strongly continuous, positive, and bounded charges on the σ -algebra \mathcal{T}_{ij} . By Bhaskara Rao and Bhaskara Rao [7, Theorem 11.4.9], the set

$$R(\boldsymbol{\lambda}_{ij}) = \{(\lambda_{ij}^1(S), \dots, \lambda_{ij}^m(S)) \mid S \in \mathcal{T}_{ij}\}$$

is convex in \mathbb{R}^m . Moreover, both $\mathbf{0} = (\lambda_{ij}^1(\emptyset), \dots, \lambda_{ij}^m(\emptyset))$ and $\boldsymbol{\lambda}_{ij}(S_{ij}) = (\lambda_{ij}^1(S_{ij}), \dots, \lambda_{ij}^m(S_{ij}))$ belong to $R(\boldsymbol{\lambda}_{ij})$. By convexity of the latter, there exists $T_{ij} \in \mathcal{T}_{ij}$ such that $\boldsymbol{\lambda}_{ij}(T_{ij}) = \alpha \boldsymbol{\lambda}_{ij}(S_{ij})$. But then $T_{ij}, U_{ij} = S_{ij} \setminus T_{ij} \in \mathcal{T}$, $S_{ij} = T_{ij} \sqcup U_{ij}$, $\boldsymbol{\lambda}(T_{ij}) = \boldsymbol{\lambda}_{ij}(T_{ij}) = \alpha \boldsymbol{\lambda}_{ij}(S_{ij}) = \alpha \boldsymbol{\lambda}(S_{ij})$, and $\boldsymbol{\lambda}(U_{ij}) = (1 - \alpha) \boldsymbol{\lambda}(S_{ij})$ by additivity of $\boldsymbol{\lambda}$.

The function $\psi : T \rightarrow A$ defined by

$$\psi(t) = \begin{cases} \phi(t) = i & \text{if } t \in T_{ij} \\ \gamma(t) = j & \text{if } t \in U_{ij} \end{cases}$$

³² \sqcup denotes the disjoint union.

is well defined and $\psi(t) \in \{\phi(t), \gamma(t)\} \subseteq \Gamma(t)$ for all $t \in T$, so that $\psi \in \text{Sel}(\Gamma)$. For each $k \in A$,

$$\begin{aligned} \psi^{-1}(k) &= \{t \in T \mid \psi(t) = k\} = \left\{ t \in \left(\bigsqcup_{i,j \in A} T_{ij} \right) \sqcup \left(\bigsqcup_{i,j \in A} U_{ij} \right) \mid \psi(t) = k \right\} \\ &= \left(\bigsqcup_{i,j \in A} \{t \in T_{ij} \mid \psi(t) = k\} \right) \sqcup \left(\bigsqcup_{i,j \in A} \{t \in U_{ij} \mid \psi(t) = k\} \right) \\ &= \left(\bigsqcup_{i,j \in A} \{t \in T_{ij} \mid \phi(t) = k\} \right) \sqcup \left(\bigsqcup_{i,j \in A} \{t \in U_{ij} \mid \gamma(t) = k\} \right) \end{aligned}$$

but, for all $t \in T_{ij}$, $\phi(t) = i$, then

- if $i = k$, $\{t \in T_{ij} \mid \phi(t) = k\} = T_{ij}$,
- else $i \neq k$ and $\{t \in T_{ij} \mid \phi(t) = k\} = \emptyset$,

thus $\bigsqcup_{i,j \in A} \{t \in T_{ij} \mid \phi(t) = k\} = \bigsqcup_{i,j \in A: i=k} T_{ij} = \bigsqcup_{j \in A} T_{kj}$; analogously, for all $t \in U_{ij}$, $\gamma(t) = j$; then

- if $j = k$, $\{t \in U_{ij} \mid \gamma(t) = k\} = U_{ij}$,
- else $j \neq k$ and $\{t \in U_{ij} \mid \gamma(t) = k\} = \emptyset$,

thus $\bigsqcup_{i,j \in A} \{t \in U_{ij} \mid \gamma(t) = k\} = \bigsqcup_{i,j \in A: j=k} U_{ij} = \bigsqcup_{i \in A} U_{ik}$; therefore,

$$\psi^{-1}(k) = \left(\bigsqcup_{j \in A} T_{kj} \right) \sqcup \left(\bigsqcup_{i \in A} U_{ik} \right) \in \mathcal{T}$$

As a consequence, ψ is \mathcal{T} -measurable and, for each $k \in A$, and for each $l = 1, \dots, m$,

$$\begin{aligned} \lambda^l(\psi^{-1}(k)) &= \sum_{j \in A} \lambda^l(T_{kj}) + \sum_{i \in A} \lambda^l(U_{ik}) = \sum_{j \in A} \alpha \lambda^l(S_{kj}) + \sum_{i \in A} (1 - \alpha) \lambda^l(S_{ik}) \\ &= \alpha \lambda^l \left(\bigsqcup_{j \in A} S_{kj} \right) + (1 - \alpha) \lambda^l \left(\bigsqcup_{i \in A} S_{ik} \right) \\ &= \alpha \lambda^l \left(\bigsqcup_{j \in A} (\phi^{-1}(k) \cap \gamma^{-1}(j)) \right) + (1 - \alpha) \lambda^l \left(\bigsqcup_{i \in A} (\phi^{-1}(i) \cap \gamma^{-1}(k)) \right) \\ &= \alpha \lambda^l \left(\phi^{-1}(k) \cap \bigsqcup_{j \in A} \gamma^{-1}(j) \right) + (1 - \alpha) \lambda^l \left(\gamma^{-1}(k) \cap \bigsqcup_{i \in A} \phi^{-1}(i) \right) \\ &= \alpha \lambda^l(\phi^{-1}(k) \cap T) + (1 - \alpha) \lambda^l(\gamma^{-1}(k) \cap T) \\ &= \alpha \lambda^l(\phi^{-1}(k)) + (1 - \alpha) \lambda^l(\gamma^{-1}(k)) \end{aligned}$$

thus $\lambda_\psi^l = \alpha\lambda_\phi^l + (1 - \alpha)\lambda_\gamma^l$. Since this is true for each $l = 1, \dots, m$, then $\lambda_\psi = \alpha\lambda_\phi + (1 - \alpha)\lambda_\gamma$, as wanted. \blacksquare

Building on this lemma, Gilboa, Maccheroni, Marinacci, and Schmeidler [15] prove that, when $m = 1$, $\{\lambda_\gamma \mid \gamma \in \text{Sel}(\Gamma) \text{ and } \gamma \text{ is } \mathcal{T}\text{-measurable}\}$ is indeed the core of the belief function

$$\text{Bel}(I) = \lambda(\{t \in T \mid \Gamma(t) \subseteq I\}) \quad \forall I \subseteq A$$

and they characterize its extreme points *à la* Shapley [31].

A.2 Proofs and related material

In what follows and up to the proof of Corollary 1, we consider a nonatomic game with estimation feedback $G = ((T, \lambda), A, u, (\Pi, \pi), f)$. Recall that Π is a collection of nonempty subsets of T , $\{T_j\}_{j=1}^m$, such that $\lambda(T_j) > 0$ for all $j \in \{1, \dots, m\}$ and $T = \cup_{j=1}^m T_j$. The proof of existence of ε -estimated equilibria rests on two key ideas which we formally develop below:

1. We first consider different correspondences and study their properties. This study culminates with the correspondence $\widetilde{\mathbf{BR}}_{f,\varepsilon} : \Delta^m \rightrightarrows \Delta^m$ defined in (14) below. All of these correspondences are basically ε -consistent/confirmed best-reply correspondences. To fix ideas, for the case $\Pi = \{T\}$ and $m = 1$, in words, given $x \in \Delta$ and $y \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(x)$, y is a possible distribution of strategies in the population, which arises if the players' distribution of actions was x and players best-responded to it using a belief which was ε -consistent with respect to x .
2. We then show that $\widetilde{\mathbf{BR}}_{f,\varepsilon}$ has a fixed point by using Browder's fixed point theorem. This will give us the equilibrium in pure strategies that we are after.

Consider $\varepsilon > 0$. First, let $\mathbf{BR}_{f,\varepsilon} : T \times \Delta^m \rightrightarrows A$ be defined by

$$\begin{aligned} \mathbf{BR}_{f,\varepsilon}(t, \mathbf{x}) \\ = \{b \in A \mid \exists \beta_t \in \Delta \text{ s.t. } f_t(b, \beta_t, x_{\pi(t)}) < \varepsilon \text{ and } u_t(b, \beta_t) \geq u_t(a, \beta_t) \quad \forall a \in A\} \end{aligned}$$

for all $(t, \mathbf{x}) \in T \times \Delta^m$. Clearly, $\mathbf{BR}_{f,\varepsilon}(t, \mathbf{x})$ is the set of all pure strategies which are a best-reply of player t to some belief β_t where β_t is ε -consistent when assuming the true distribution restricted to the subpopulation $T_{\pi(t)}$ is $x_{\pi(t)}$. One can derive several related “ ε -consistent best-reply” correspondences from this basic one. For each $\mathbf{x} \in \Delta^m$, denote the \mathbf{x} -section $\mathbf{BR}_{f,\varepsilon}(\cdot, \mathbf{x}) : T \rightrightarrows A$ of $\mathbf{BR}_{f,\varepsilon}$ by $\mathbf{BR}_{f,\varepsilon}^{\mathbf{x}}$. Next, let $\Phi_{f,\varepsilon} : \Delta^m \rightrightarrows \Sigma$ be defined as $\Phi_{f,\varepsilon}(\mathbf{x}) = \text{Sel}(\mathbf{BR}_{f,\varepsilon}^{\mathbf{x}})$ for all $\mathbf{x} \in \Delta^m$ where $\text{Sel}(\mathbf{BR}_{f,\varepsilon}^{\mathbf{x}})$ is the set of all selections of $\mathbf{BR}_{f,\varepsilon}^{\mathbf{x}}$.

Thus, for a strategy profile $\sigma \in \Sigma$, we have that

$$\begin{aligned}
& [\sigma \in \Phi_{f,\varepsilon}(\mathbf{x})] \\
& \iff [\forall t \in T, \sigma(t) \in \mathbf{BR}_{f,\varepsilon}^{\mathbf{x}}(t)] \\
& \iff [\forall t \in T, \sigma(t) \in \mathbf{BR}_{f,\varepsilon}(t, \mathbf{x})] \\
& \iff [\forall t \in T \exists \beta_t \in \Delta \text{ s.t. } f_t(\sigma(t), \beta_t, x_{\pi(t)}) < \varepsilon \text{ and } u_t(\sigma(t), \beta_t) \geq u_t(a, \beta_t) \quad \forall a \in A]
\end{aligned}$$

Remark 2 If there exists $\mathbf{x} \in \Delta^m$ such that $\sigma \in \Phi_{f,\varepsilon}(\mathbf{x})$ and $x_{\pi(t)} = \lambda_{\sigma}^{\pi(t)}$ for all $t \in T$, then σ is an ε -estimated equilibrium. In fact, we have

1. For each $t \in T$ there exists $\beta_t \in \Delta$ such that $f_t(\sigma(t), \beta_t, x_{\pi(t)}) < \varepsilon$ and $u_t(\sigma(t), \beta_t) \geq u_t(a, \beta_t)$ for all $a \in A$;
2. We can define $\beta : T \rightarrow \Delta$ by $\beta(t) = \beta_t$ for all $t \in T$.

This implies that for each $t \in T$

- a) $u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t))$ for all $a \in A$ (optimality);
- b) $f_t(\sigma(t), \beta(t), \lambda_{\sigma}^{\pi(t)}) < \varepsilon$ (strict ε -consistency),

that is, $\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ f_t(\sigma(t), \beta(t), \lambda_{\sigma}^{\pi(t)}) < \varepsilon \end{array} \right\} = T$. In particular, it holds that

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ f_t(\sigma(t), \beta(t), \lambda_{\sigma}^{\pi(t)}) \leq \varepsilon \end{array} \right\} \right) = 1$$

▲

Next, consider the correspondence $\mathbf{B}_{f,\varepsilon} : \Delta^m \rightrightarrows \Sigma$ defined by

$$\mathbf{B}_{f,\varepsilon}(\mathbf{x}) = \left\{ \sigma \in \Sigma \mid \exists \beta \in \Delta^T \text{ s.t. } \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A, \forall t \in T \\ \sup_{t \in T} f_t(\sigma(t), \beta(t), x_{\pi(t)}) < \varepsilon \end{array} \right\} \quad \forall \mathbf{x} \in \Delta^m$$

Lemma 2 $\mathbf{B}_{f,\varepsilon}(\mathbf{x}) = \bigcup_{\eta \in (0, \varepsilon)} \Phi_{f,\eta}(\mathbf{x}) \subseteq \Phi_{f,\varepsilon}(\mathbf{x})$ for all $\mathbf{x} \in \Delta^m$.

Proof. Fix $\mathbf{x} \in \Delta^m$. Consider $\sigma \in \bigcup_{\eta \in (0, \varepsilon)} \Phi_{f,\eta}(\mathbf{x})$. It follows that $\sigma \in \Phi_{f,\eta}(\mathbf{x})$ for some $\eta \in (0, \varepsilon)$. This implies that $\sigma \in \Sigma$ and $\sigma(t) \in \mathbf{BR}_{f,\eta}(t, \mathbf{x})$ for all $t \in T$, that is, for each $t \in T$ there exists $\beta_t \in \Delta$ such that $f_t(\sigma(t), \beta_t, x_{\pi(t)}) < \eta$ and $u_t(\sigma(t), \beta_t) \geq u_t(a, \beta_t)$ for all $a \in A$. In particular, if we define $\beta \in \Delta^T$ as $\beta(t) = \beta_t$ for all $t \in T$, we have that $u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t))$ for all $a \in A$ and for all $t \in T$, and

$$\sup_{t \in T} f_t(\sigma(t), \beta(t), x_{\pi(t)}) \leq \eta < \varepsilon$$

yielding that $\sigma \in \mathbf{B}_{f,\varepsilon}(\mathbf{x})$. *Vice versa*, if $\sigma \in \mathbf{B}_{f,\varepsilon}(\mathbf{x})$, then there exists $\beta \in \Delta^T$ such that

$$\sup_{t \in T} f_t(\sigma(t), \beta(t), x_{\pi(t)}) < \varepsilon \quad (13)$$

and $u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t))$ for all $a \in A$ and for all $t \in T$. It follows that there exists $\bar{\eta} \in (0, \varepsilon)$ such that (13) holds with $\bar{\eta}$ in place of ε . This implies that $\sigma \in \Phi_{f,\bar{\eta}}(\mathbf{x}) \subseteq \bigcup_{\eta \in (0, \varepsilon)} \Phi_{f,\eta}(\mathbf{x})$.

Obviously, if $0 < \eta < \eta'$, then $\mathbf{BR}_{f,\eta}(t, \mathbf{x}) \subseteq \mathbf{BR}_{f,\eta'}(t, \mathbf{x})$ for all $t \in T$ and for all $\mathbf{x} \in \Delta^m$ and, in particular, $\Phi_{f,\eta}(\mathbf{x}) \subseteq \Phi_{f,\eta'}(\mathbf{x})$. This implies that $\bigcup_{\eta \in (0, \varepsilon)} \Phi_{f,\eta}(\mathbf{x}) \subseteq \Phi_{f,\varepsilon}(\mathbf{x})$. ■

Remark 2 above will be useful to justify the following last correspondence: $\widetilde{\mathbf{BR}}_{f,\varepsilon} : \Delta^m \rightrightarrows \Delta^m$ defined by

$$\widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x}) = \{\mathbf{y} \in \Delta^m \mid \exists \sigma \in \mathbf{B}_{f,\varepsilon}(\mathbf{x}) \text{ s.t. } \lambda_{\sigma}^j = y_j \quad \forall j \in \{1, \dots, m\}\} \quad \forall \mathbf{x} \in \Delta^m \quad (14)$$

In other words, $\widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x})$ is the collection of actions' distributions $\mathbf{y} = (y_j)_{j=1}^m$ on the subpopulations of players, which can be induced by the β -optimal choice of strategies σ where beliefs $\beta = (\beta_t)_{t \in T}$ are close enough in terms of feedback to $\mathbf{x} = (x_j)_{j=1}^m$. Note that

$$\widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x}) = \left\{ (\lambda_{\sigma}^j)_{j=1}^m \mid \sigma \in \mathbf{B}_{f,\varepsilon}(\mathbf{x}) \right\} = \left\{ (\lambda_{\sigma}^j)_{j=1}^m \mid \sigma \in \bigcup_{\eta \in (0, \varepsilon)} \Phi_{f,\eta}(\mathbf{x}) \right\} \quad (15)$$

$$= \bigcup_{\eta \in (0, \varepsilon)} \left\{ (\lambda_{\sigma}^j)_{j=1}^m \mid \sigma \in \Phi_{f,\eta}(\mathbf{x}) \right\} = \bigcup_{\eta \in (0, \varepsilon)} \left\{ (\lambda_{\sigma}^j)_{j=1}^m \mid \sigma \in \text{Sel}(\mathbf{BR}_{f,\eta}^{\mathbf{x}}) \right\} \quad (16)$$

An immediate implication of definition (14) is the next result.

Lemma 3 *If $\mathbf{x} \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x})$, then there exists an ε -estimated equilibrium σ such that $\lambda_{\sigma}^j = x_j$ for all $j \in \{1, \dots, m\}$.*

Proof. By Lemma 2 and the definition of $\widetilde{\mathbf{BR}}_{f,\varepsilon}$, if $\mathbf{x} \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x})$, then there exists $\sigma \in \mathbf{B}_{f,\varepsilon}(\mathbf{x}) \subseteq \Phi_{f,\varepsilon}(\mathbf{x})$ such that $\lambda_{\sigma}^j = x_j$ for all $j \in \{1, \dots, m\}$. Remark 2 yields that σ is an ε -estimated equilibrium. ■

Lemma 4 *If λ is strongly continuous, then $\widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x})$ is nonempty and convex for all $\mathbf{x} \in \Delta^m$.*

Proof. Fix $\mathbf{x} \in \Delta^m$ and $\eta \in (0, \varepsilon)$. Since f satisfies (1), recall that

$$\forall t \in T, \forall z \in \Delta, \exists \gamma_{t,z} \in \Delta \text{ s.t. } \forall a \in A \quad f_t(a, \gamma_{t,z}, z) = 0 \quad (17)$$

Since \mathbf{x} is given, define $\beta \in \Delta^T$ to be such that $\beta(t) = \gamma_{t, x_{\pi(t)}}$ for all $t \in T$. Note that $\beta(t) \in \Delta$ satisfies $f_t(a, \beta(t), x_{\pi(t)}) = 0 < \eta$ for all $a \in A$ and for all $t \in T$. Since A is finite,

for each $t \in T$ choose $\sigma(t) \in A$ such that $u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t))$ for all $a \in A$. This defines a function $\sigma : T \rightarrow A$, that is $\sigma \in \Sigma$, such that $\sigma \in \Phi_{f,\eta}(\mathbf{x})$. By Lemma 2, we have that $\sigma \in \mathbf{B}_{f,\varepsilon}(\mathbf{x})$ and $(\lambda_\sigma^j)_{j=1}^m \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x})$. Convexity is a consequence of the following two observations:

1. By Lemma 1 and since each λ^j is strongly continuous, recall that $\{(\lambda_\sigma^j)_{j=1}^m \mid \sigma \in \text{Sel}(\mathbf{BR}_{f,\eta}^{\mathbf{x}})\}$ is a convex subset of Δ^m for all $\eta \in (0, \varepsilon)$.
2. By (16), we have that

$$\widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x}) = \bigcup_{\eta \in (0, \varepsilon)} \left\{ (\lambda_\sigma^j)_{j=1}^m \mid \sigma \in \text{Sel}(\mathbf{BR}_{f,\eta}^{\mathbf{x}}) \right\}$$

It follows that $\widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x})$ is the union of a chain of convex sets,³³ proving convexity. \blacksquare

For the next result recall that a) d_Δ is the distance on Δ induced by the Euclidean norm; b) we say that $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument if and only if for each $\varepsilon > 0$ there exists $\delta_\varepsilon > 0$ such that

$$d_\Delta(x, y) < \delta_\varepsilon \implies |f_t(a, \gamma, x) - f_t(a, \gamma, y)| < \varepsilon \quad \forall t \in T, \forall a \in A, \forall \gamma \in \Delta$$

The intuition behind the proof of the next lemma is that if a strategy σ was β -optimal and β was ε -consistent, given \mathbf{x} , small perturbations of \mathbf{x} do not disrupt optimality and ε -consistency.

Lemma 5 *If $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument, then $\widetilde{\mathbf{BR}}_{f,\varepsilon}^{-1}(\mathbf{y})$ is open for all $\mathbf{y} \in \Delta^m$.*

Proof. Fix $\mathbf{y} \in \Delta^m$. Recall that $\widetilde{\mathbf{BR}}_{f,\varepsilon}^{-1}(\mathbf{y}) = \left\{ \mathbf{x} \in \Delta^m \mid \mathbf{y} \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x}) \right\}$. Note that

$$\mathbf{x} \in \widetilde{\mathbf{BR}}_{f,\varepsilon}^{-1}(\mathbf{y}) \iff \mathbf{y} \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x})$$

and $\widetilde{\mathbf{BR}}_{f,\varepsilon}^{-1}(\mathbf{y})$ is open if and only if “for each $\bar{\mathbf{x}}$ such that $\mathbf{y} \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(\bar{\mathbf{x}})$, there exists a ball in Δ^m of radius δ and center $\bar{\mathbf{x}}$ such that $\mathbf{y} \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(\mathbf{x})$ for all $\mathbf{x} \in B_\delta(\bar{\mathbf{x}})$ ”.

Now arbitrarily choose $\bar{\mathbf{x}}$ such that $\mathbf{y} \in \widetilde{\mathbf{BR}}_{f,\varepsilon}(\bar{\mathbf{x}})$. By definition of $\widetilde{\mathbf{BR}}_{f,\varepsilon}(\bar{\mathbf{x}})$, there exist $\sigma \in \mathbf{B}_{f,\varepsilon}(\bar{\mathbf{x}}) \subseteq \Sigma$ and $\beta \in \Delta^T$ such that

1. $\lambda_\sigma^j = y_j$ for all $j \in \{1, \dots, m\}$;

³³Recall that if $0 < \eta < \eta'$, then

$$\mathbf{BR}_{f,\eta}(t, \mathbf{x}) \subseteq \mathbf{BR}_{f,\eta'}(t, \mathbf{x}) \quad \forall t \in T, \forall \mathbf{x} \in \Delta^m$$

This implies that $\text{Sel}(\mathbf{BR}_{f,\eta}^{\mathbf{x}}) = \Phi_{f,\eta}(\mathbf{x}) \subseteq \Phi_{f,\eta'}(\mathbf{x}) = \text{Sel}(\mathbf{BR}_{f,\eta'}^{\mathbf{x}})$ for all $\mathbf{x} \in \Delta^m$.

2. $\sup_{t \in T} f_t(\sigma(t), \beta(t), \bar{x}_{\pi(t)}) < \varepsilon$;
3. $u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t))$ for all $a \in A$ and for all $t \in T$.

By point 2, there exists $\bar{\varepsilon} \in (0, \varepsilon)$ such that

$$\sup_{t \in T} f_t(\sigma(t), \beta(t), \bar{x}_{\pi(t)}) < \bar{\varepsilon} < \varepsilon$$

Let $\hat{\varepsilon} \in (0, \frac{\varepsilon - \bar{\varepsilon}}{2})$. Since $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument, there exists $\delta_{\hat{\varepsilon}} > 0$ such that

$$d_{\Delta}(x, y) < \delta_{\hat{\varepsilon}} \implies |f_t(a, \gamma, x) - f_t(a, \gamma, y)| < \hat{\varepsilon} \quad \forall t \in T, \forall a \in A, \forall \gamma \in \Delta$$

For each $\mathbf{x} \in B_{\delta_{\hat{\varepsilon}}}(\bar{\mathbf{x}})$ note that $d_{\Delta}(x_j, \bar{x}_j) < \delta_{\hat{\varepsilon}}$ for all $j \in \{1, \dots, m\}$. This implies that for each $t \in T$ and for each $\mathbf{x} \in B_{\delta_{\hat{\varepsilon}}}(\bar{\mathbf{x}})$

$$|f_t(\sigma(t), \beta(t), x_{\pi(t)}) - f_t(\sigma(t), \beta(t), \bar{x}_{\pi(t)})| < \hat{\varepsilon}$$

Since $f_t \geq 0$ for all $t \in T$, it follows that for each $t \in T$ and for each $\mathbf{x} \in B_{\delta_{\hat{\varepsilon}}}(\bar{\mathbf{x}})$

$$\begin{aligned} f_t(\sigma(t), \beta(t), x_{\pi(t)}) &= |f_t(\sigma(t), \beta(t), x_{\pi(t)})| \\ &\leq |f_t(\sigma(t), \beta(t), \bar{x}_{\pi(t)})| + |f_t(\sigma(t), \beta(t), x_{\pi(t)}) - f_t(\sigma(t), \beta(t), \bar{x}_{\pi(t)})| \\ &= f_t(\sigma(t), \beta(t), \bar{x}_{\pi(t)}) + |f_t(\sigma(t), \beta(t), x_{\pi(t)}) - f_t(\sigma(t), \beta(t), \bar{x}_{\pi(t)})| \\ &< \bar{\varepsilon} + \hat{\varepsilon} \end{aligned}$$

This implies that

$$\sup_{t \in T} f_t(\sigma(t), \beta(t), x_{\pi(t)}) \leq \bar{\varepsilon} + \hat{\varepsilon} < \frac{\bar{\varepsilon} + \varepsilon}{2} < \varepsilon \quad \forall \mathbf{x} \in B_{\delta_{\hat{\varepsilon}}}(\bar{\mathbf{x}})$$

In other words, for each $\mathbf{x} \in B_{\delta_{\hat{\varepsilon}}}(\bar{\mathbf{x}})$ we have that $\sigma \in \Sigma$ is such that the same $\beta \in \Delta^T$ of above satisfies points 2 and 3, but with \mathbf{x} in place of $\bar{\mathbf{x}}$. This yields that $\sigma \in \mathbf{B}_{f, \varepsilon}(\mathbf{x})$ for all $\mathbf{x} \in B_{\delta_{\hat{\varepsilon}}}(\bar{\mathbf{x}})$. Since $\mathbf{y} = (y_j)_{j=1}^m = (\lambda_{\sigma}^j)_{j=1}^m$, we obtain that $\mathbf{y} \in \widetilde{\mathbf{BR}}_{f, \varepsilon}(\mathbf{x})$ for all $\mathbf{x} \in B_{\delta_{\hat{\varepsilon}}}(\bar{\mathbf{x}})$, proving the statement. \blacksquare

Proof of Theorem 1. By Lemma 3, it is enough to show that $\widetilde{\mathbf{BR}}_{f, \varepsilon} : \Delta^m \rightrightarrows \Delta^m$ has a fixed point. Clearly, $\Delta^m \subseteq (\mathbb{R}^n)^m$ is nonempty, compact, and convex. By Lemmas 4 and 5, $\widetilde{\mathbf{BR}}_{f, \varepsilon}$ has nonempty and convex values and $\widetilde{\mathbf{BR}}_{f, \varepsilon}^{-1}(\mathbf{y})$ is open for all $\mathbf{y} \in \Delta^m$. By Browder's Fixed Point Theorem for correspondences (see Theorem 1 of Browder [8]), $\widetilde{\mathbf{BR}}_{f, \varepsilon}$ has a fixed point. \blacksquare

We next prove the remaining results of the main text.

Proof of Corollary 1. It is enough to observe that a nonatomic game with *message* feedback can be mapped into a nonatomic game with *estimation* feedback where f is defined as in (4) and $\Pi = \{T_1\}$.³⁴ With this identification, an ε -estimated equilibrium is a self-confirming ε -equilibrium. By Theorem 1, it is then enough to show that $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument. Since $m = (m_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument, we have that for each $\varepsilon > 0$ there exists $\delta_\varepsilon > 0$ such that

$$d_\Delta(x, y) < \delta_\varepsilon \implies d(m_t(a, x), m_t(a, y)) < \varepsilon \quad \forall t \in T, \forall a \in A \quad (18)$$

Since for each $t \in T$ we have that $f_t(a, x, y) = d(m_t(a, x), m_t(a, y))$ for all $a \in A$ and for all $x, y \in \Delta$, observe that

$$\begin{aligned} |f_t(a, \gamma, x) - f_t(a, \gamma, y)| &= |d(m_t(a, \gamma), m_t(a, x)) - d(m_t(a, \gamma), m_t(a, y))| \\ &\leq d(m_t(a, x), m_t(a, y)) \quad \forall t \in T, \forall a \in A, \forall x, y, \gamma \in \Delta \end{aligned}$$

By (18), we can conclude that for each $\varepsilon > 0$ there exists $\delta_\varepsilon > 0$ such that

$$\begin{aligned} d_\Delta(x, y) < \delta_\varepsilon &\implies \\ |f_t(a, \gamma, x) - f_t(a, \gamma, y)| &\leq d(m_t(a, x), m_t(a, y)) < \varepsilon \quad \forall t \in T, \forall a \in A, \forall \gamma \in \Delta \end{aligned}$$

proving equicontinuity with respect to the third argument of f . ■

Proof of Corollary 2. Consider the nonatomic game $G = ((T, \lambda), A, u)$ and $\varepsilon > 0$. Since $u = (u_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument, we have that for each $\hat{\varepsilon} > 0$ there exists $\delta_{\hat{\varepsilon}} > 0$ such that

$$d_\Delta(x, y) < \delta_{\hat{\varepsilon}} \implies |u_t(a, x) - u_t(a, y)| < \hat{\varepsilon} \quad \forall t \in T, \forall a \in A \quad (19)$$

Consider the profile $m = (m_t)_{t \in T}$ of message functions such that each $m_t : A \times \Delta \rightarrow \Delta$ is defined to be such that

$$m_t(a, x) = x \quad \forall a \in A, \forall x \in \Delta$$

Note that in this case $(M, d) = (\Delta, d_\Delta)$. Clearly, $m = (m_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the second argument. Given $\varepsilon > 0$, consider $\delta_{\varepsilon/2} > 0$ as in (19). By Corollary 1, we have that there exists a self-confirming $\delta_{\varepsilon/2}/2$ -equilibrium $\sigma \in \Sigma$, that is, there exists $\beta \in \Delta^T$ such that

$$\begin{aligned} 1 &= \lambda \left(\left\{ t \in T \left| \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d(m_t(\sigma(t), \beta(t)), m_t(\sigma(t), \lambda_\sigma)) \leq \delta_{\varepsilon/2}/2 \end{array} \right. \right\} \right) \\ &= \lambda \left(\left\{ t \in T \left| \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d_\Delta(\beta(t), \lambda_\sigma) \leq \delta_{\varepsilon/2}/2 \end{array} \right. \right\} \right) \end{aligned}$$

³⁴Thus, $m = 1$, $T_1 = T$, and $\pi(t) = 1$ for all $t \in T$.

Define by O the set of “optimizing” players

$$O = \left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d_\Delta(\beta(t), \lambda_\sigma) \leq \delta_{\varepsilon/2}/2 \end{array} \right\}$$

Since u satisfies (19), note that if $t \in O$, then we have that $d_\Delta(\beta(t), \lambda_\sigma) \leq \delta_{\varepsilon/2}/2 < \delta_{\varepsilon/2}$ which implies that for each $a \in A$

$$|u_t(\sigma(t), \beta(t)) - u_t(\sigma(t), \lambda_\sigma)| < \frac{\varepsilon}{2} \text{ and } |u_t(a, \beta(t)) - u_t(a, \lambda_\sigma)| < \frac{\varepsilon}{2}$$

Since $t \in O$, we can conclude that

$$\begin{aligned} u_t(\sigma(t), \lambda_\sigma) &> u_t(\sigma(t), \beta(t)) - \frac{\varepsilon}{2} \geq u_t(a, \beta(t)) - \frac{\varepsilon}{2} \\ &> u_t(a, \lambda_\sigma) - \frac{\varepsilon}{2} - \frac{\varepsilon}{2} = u_t(a, \lambda_\sigma) - \varepsilon \quad \forall a \in A \end{aligned}$$

Since t was arbitrarily chosen in O , we have that

$$O \subseteq \{t \in T \mid u_t(\sigma(t), \lambda_\sigma) \geq u_t(a, \lambda_\sigma) - \varepsilon \quad \forall a \in A\}$$

Since O has mass 1, it follows that $\sigma \in \Sigma$ is an ε -Nash equilibrium. ■

Proof of Corollary 3. It is enough to observe that a nonatomic game with a *neighborhood structure* can be mapped into a nonatomic game with *estimation* feedback where f is defined as in (8). With this identification, an ε -estimated equilibrium is a peer-confirming ε -equilibrium. By Theorem 1, it is then enough to show that $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument. But, note that

$$|f_t(a, \gamma, x) - f_t(a, \gamma, y)| = |d_\Delta(\gamma, x) - d_\Delta(\gamma, y)| \leq d_\Delta(x, y) \quad \forall t \in T, \forall a \in A, \forall \gamma \in \Delta$$

trivially proving equicontinuity with respect to the third argument of f . ■

We conclude by proving Corollary 4. But, before doing so, we need to make an intermediate observation. Consider a statistical divergence D . Recall that $D : \Delta \times \Delta^\circ \rightarrow [0, \infty)$ is a jointly convex and continuous function. Denote by \mathcal{K} the collection of all nonempty compact sets of Δ . We endow \mathcal{K} with the Hausdorff distance (see, e.g., Aliprantis and Border [1, Chapter 3, Sections 16 and 17]). We denote by $\bar{\mathcal{Q}}$ a compact set of \mathcal{K} such that each $Q \in \bar{\mathcal{Q}}$ is a nonempty, convex, and compact subset of Δ° . Given $x \in \Delta$ and $Q \in \bar{\mathcal{Q}}$, consider the minimization problem

$$\min D(x||y) \text{ sub to } y \in Q$$

Define $\mu : \Delta \times \bar{\mathcal{Q}} \rightrightarrows \Delta$ to be the solution correspondence of this minimization problem, that is, for each $x \in \Delta$ and for each $Q \in \bar{\mathcal{Q}}$,

$$\mu(x, Q) = \left\{ z \in \Delta : z \in Q \text{ and } D(x||z) = \min_{y \in Q} D(x||y) \right\}$$

By Berge's maximum theorem, note that μ is upper hemicontinuous when $\Delta \times \bar{\mathcal{Q}}$ is endowed with the product topology. In particular, if D is strictly convex with respect to the second argument, μ is single-valued, that is, μ is a continuous function. Finally, define the map $g : \Delta \times \Delta \times \bar{\mathcal{Q}} \rightarrow [0, \infty)$ by

$$g(\beta, x, Q) = d_{\Delta}(\beta, \mu(x, Q)) \quad \forall \beta, x \in \Delta, \forall Q \in \bar{\mathcal{Q}}$$

Since μ is a continuous function, it follows that g is continuous when $\Delta \times \Delta \times \bar{\mathcal{Q}}$ is endowed with the product topology. By Aliprantis and Border [1, Corollary 3.31] and since $\Delta \times \Delta \times \bar{\mathcal{Q}}$ is a compact metric space, g is uniformly continuous.

Proof of Corollary 4. Set $\bar{\mathcal{Q}} = \text{cl } \mathcal{Q}$. By point 3 of Remark 1, note that $\bar{\mathcal{Q}}$ is a compact subset of \mathcal{K} such that each $Q \in \bar{\mathcal{Q}}$ is a nonempty, convex, and compact subset of Δ^o . For each $t \in T$ define $f_t : A \times \Delta \times \Delta \rightarrow [0, \infty)$ by

$$f_t(a, \gamma, x) = g(\gamma, x, Q_t) \quad \forall a \in A, \forall \gamma, x \in \Delta \quad (20)$$

It is then enough to observe that a nonatomic game with *model misspecification* can be mapped into a nonatomic game with *estimation* feedback where f is defined as in (20) and $\Pi = \{T_1\}$.³⁵ With this identification, an ε -estimated equilibrium is an ε -BNE. By Theorem 1, it is then enough to show that $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument. Since g is uniformly continuous, the statement is trivially true. ■

The proof of the last two points of Remark 1 is routine. Thus, we conclude by only proving point 1.

Proof of point 1 of Remark 1. Set $\bar{\mathcal{Q}} = \text{cl } \mathcal{Q}$. Note that $\bar{\mathcal{Q}}$ is a compact subset of \mathcal{K} such that each $Q \in \bar{\mathcal{Q}}$ is a nonempty, convex, and compact subset of Δ^o . For each $t \in T$ define $f_t : A \times \Delta \times \Delta \rightarrow [0, \infty)$ as in the proof of Corollary 4, that is,

$$f_t(a, \gamma, x) = g(\gamma, x, Q_t) \quad \forall t \in T, \forall a \in A, \forall \gamma, x \in \Delta$$

Since g is continuous and $\Delta \times \Delta \times \bar{\mathcal{Q}}$ is compact, observe that $g \geq 0$ takes a maximum value $M \geq 0$. Define the profile of feedback functions \tilde{f} to be such that for each $t \in T$

$$\tilde{f}_t(a, \gamma, x) = \begin{cases} f_t(a, \gamma, x) & \gamma \in Q_t \\ M + 1 & \gamma \notin Q_t \end{cases} \quad \forall a \in A, \forall \gamma, x \in \Delta$$

Note that each \tilde{f}_t satisfies (1). By the proof of Corollary 4, $f = (f_t)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument. It follows that for each $\varepsilon > 0$ there exists $\delta_{\varepsilon} > 0$ such that

$$d_{\Delta}(x, y) < \delta_{\varepsilon} \implies |f_t(a, \gamma, x) - f_t(a, \gamma, y)| < \varepsilon \quad \forall t \in T, \forall a \in A, \forall \gamma \in \Delta$$

³⁵Thus, $m = 1$, $T_1 = T$, and $\pi(t) = 1$ for all $t \in T$.

Consider $x, y \in \Delta$ such that $d_\Delta(x, y) < \delta_\varepsilon$ and consider $t \in T$, $a \in A$, and $\gamma \in \Delta$. We have two cases, either $\gamma \in Q_t$ or $\gamma \notin Q_t$. In the first case, $\left| \tilde{f}_t(a, \gamma, x) - \tilde{f}_t(a, \gamma, y) \right| = |f_t(a, \gamma, x) - f_t(a, \gamma, y)| < \varepsilon$, while in the second case $\left| \tilde{f}_t(a, \gamma, x) - \tilde{f}_t(a, \gamma, y) \right| = |M + 1 - (M + 1)| = 0 < \varepsilon$. Since t , a , and γ were chosen arbitrarily, it follows that $\tilde{f} = \left(\tilde{f}_t \right)_{t \in T}$ is a family of functions which is equicontinuous with respect to the third argument. Next, we can consider the nonatomic game with estimation feedback $\left((T, \lambda), A, u, (\Pi, \pi), \tilde{f} \right)$, where $\Pi = \{T_1\}$.³⁶ By Theorem 1, we have that for each $\tilde{\varepsilon} > 0$ there exists an $\tilde{\varepsilon}$ -estimated equilibrium σ for this game, that is, there exists $\beta \in \Delta^T$ such that

$$\lambda \left(\left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ \tilde{f}_t(\sigma(t), \beta(t), \lambda_\sigma) \leq \tilde{\varepsilon} \end{array} \right\} \right) = 1$$

If given $\varepsilon > 0$ we define $\tilde{\varepsilon} = \frac{\min\{M+1, \varepsilon\}}{2} > 0$, since $\tilde{\varepsilon} < M + 1, \varepsilon$, then we have that

$$\begin{aligned} & \left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ \tilde{f}_t(\sigma(t), \beta(t), \lambda_\sigma) \leq \tilde{\varepsilon} \end{array} \right\} \\ & \subseteq \left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ f_t(\sigma(t), \beta(t), \lambda_\sigma) \leq \tilde{\varepsilon} \text{ and } \beta(t) \in Q_t \end{array} \right\} \\ & \subseteq \left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d_\Delta(\beta(t), \operatorname{argmin}_{z \in Q_t} D(\lambda_\sigma || z)) \leq \tilde{\varepsilon} \text{ and } \beta(t) \in Q_t \end{array} \right\} \\ & \subseteq \left\{ t \in T \mid \begin{array}{l} u_t(\sigma(t), \beta(t)) \geq u_t(a, \beta(t)) \quad \forall a \in A \\ d_\Delta(\beta(t), \operatorname{argmin}_{z \in Q_t} D(\lambda_\sigma || z)) \leq \varepsilon \text{ and } \beta(t) \in Q_t \end{array} \right\} \end{aligned}$$

yielding the statement. ■

A.3 A comment by David Schmeidler: How I arrived at the congestion model

During the late sixties while working on my PhD thesis I found that the research/teaching assistant grant did not satisfy my material needs and soon found a moonlighting job as a member of a team advising on transportation. There I was exposed (mainly in the Tel Aviv area) to the daily traffic commuters congestion (and bottlenecks), and to the fact that transportation experts techniques were restricted to optimization and almost all were unaware of noncooperative equilibria and disregarded the possibilities of inefficient equilibria. It occurred to me that a model with a nonatomic continuum of players and finitely many pure strategies may describe the daily commuter's game. At the time there were cooperative games and general equilibrium models with such a set of players, but no noncooperative games with infinitely many players. I needed a fourth paper to complete my thesis so I

³⁶Thus, $m = 1$, $T_1 = T$, and $\pi(t) = 1$ for all $t \in T$.

constructed the model. It had to be submitted in Hebrew so this was the first and last research work I wrote in Hebrew. I translated the previously completed three papers and submitted my thesis. It was approved before the summer of 1969. I translated this work into English and submitted it for publication during 1970/71. The only difference between the English and the Hebrew version was that the latter was slightly more general. In the latter there was a finite set of linear functionals instead of the unique expectation functional in the published version.³⁷ Finally, when teaching noncooperative games to math students or graduates, the congestion example was most helpful and made the Nash or Harsanyi-Nash equilibrium a less metaphysical concept.

References

- [1] C. D. Aliprantis and K. C. Border, *Infinite Dimensional Analysis*, 3rd ed., Springer-Verlag, Berlin, 2006.
- [2] N. I. Al-Najjar, Large games and the law of large numbers, *Games and Economic Behavior*, 64, 1–34, 2008.
- [3] R. J. Aumann, Markets with a continuum of traders, *Econometrica*, 32, 39–50, 1964.
- [4] Y. Azrieli, On pure conjectural equilibrium with non-manipulable information, *International Journal of Game Theory*, 38, 209–219, 2009.
- [5] E. J. Balder, A unifying pair of Cournot-Nash equilibrium existence results, *Journal of Economic Theory*, 102, 437–470, 2002.
- [6] P. Battigalli, Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali, mimeo, 1987.
- [7] K. P. S. Bhaskara Rao and M. Bhaskara Rao, *Theory of Charges: A Study of Finitely Additive Measures*, Academic Press, London, 1983.
- [8] F. E. Browder, The fixed point theory of multi-valued mappings in topological vector spaces, *Mathematische Annalen*, 177, 283–301, 1968.
- [9] G. Carmona and K. Podczeck, On the existence of pure-strategy equilibria in large games, *Journal of Economic Theory*, 144, 1300–1319, 2009.
- [10] S. Cerreia-Vioglio, L. P. Hansen, F. Maccheroni, and M. Marinacci, Decision under model misspecification, mimeo, 2019.

³⁷See Remark 2 of Schmeidler [30].

- [11] P. Dubey and L. S. Shapley, Noncooperative general exchange with a continuum of traders: Two models, *Journal of Mathematical Economics*, 23, 253–293, 1994.
- [12] I. Esponda and D. Pouzo, Berk–Nash equilibrium: A framework for modeling agents with misspecified models, *Econometrica*, 84, 1093–1130, 2016.
- [13] D. Fudenberg and Y. Kamada, Rationalizable partition-confirmed equilibrium with heterogeneous beliefs, *Games and Economic Behavior*, 109, 364–381, 2018.
- [14] D. Fudenberg and D. K. Levine, Self-confirming equilibrium, *Econometrica*, 61, 523–545, 1993.
- [15] I. Gilboa, F. Maccheroni, M. Marinacci, and D. Schmeidler, A theory of evidence and measurement, mimeo, 2018.
- [16] E. Kalai and E. Lehrer, Subjective equilibrium in repeated games, *Econometrica*, 61, 1231–1240, 1993.
- [17] E. Kalai and E. Lehrer, Subjective games and equilibria, *Games and Economic Behavior*, 8, 123–163, 1995.
- [18] M. A. Khan, L. Qiao, K. P. Rath, and Y. Sun, Modeling infinitely many agents: Why countable additivity is necessary, mimeo, 2019.
- [19] M. A. Khan, K. P. Rath, and Y. Sun, On the existence of pure strategy equilibria in games with a continuum of players, *Journal of Economic Theory*, 76, 13–46, 1997.
- [20] M. A. Khan and Y. Sun, Pure strategies in games with private information, *Journal of Mathematical Economics*, 24, 633–653, 1995.
- [21] M. A. Khan and Y. Sun, Nonatomic games on Loeb spaces, *Proceedings of the National Academy of Sciences*, 93, 15518–15521, 1996.
- [22] M. A. Khan and Y. Sun, Non-cooperative games with many players, *Handbook of Game Theory with Economic Applications*, 3, 1761–1808, 2002.
- [23] E. Lipnowski and E. Sadler, Peer-confirming equilibrium, *Econometrica*, 87, 567–591, 2019.
- [24] D. Maharam, Finitely additive measures on the integers, *Sankhyā, Series A*, 38, 44–59, 1976.
- [25] A. Mas-Colell, On a theorem of Schmeidler, *Journal of Mathematical Economics*, 13, 201–206, 1984.

- [26] K. P. Rath, A direct proof of the existence of pure strategy equilibria in games with a continuum of players, *Economic Theory*, 2, 427–433, 1992.
- [27] A. Rubinstein and A. Wolinsky, Rationalizable conjectural equilibrium: Between Nash and rationalizability, *Games and Economic Behavior*, 6, 299–311, 1994.
- [28] D. Samet and D. Schmeidler, Desirability relations in Savage’s model of decision making, <https://www.tau.ac.il/~samet/papers/desirability.pdf>, 2019.
- [29] L. J. Savage, *The Foundations of Statistics*, Wiley, New York, 1954.
- [30] D. Schmeidler, Equilibrium points of nonatomic games, *Journal of Statistical Physics*, 7, 295–300, 1973.
- [31] L. S. Shapley, Cores of convex games, *International Journal of Game Theory*, 1, 11–26, 1971.