Strategic Behavior in Non-Atomic Games*

Mehmet Barlo[†] Guilherme Carmona[‡]

June 26, 2015

Abstract

In order to remedy the possible loss of strategic interaction in non-atomic games with a societal choice, this study proposes a refinement of Nash equilibrium, *strategic* equilibrium. Given a non-atomic game, its perturbed game is one in which every player believes that he alone has a small, but positive, impact on the societal choice; and a distribution is a *strategic* equilibrium if it is a limit point of a sequence of Nash equilibrium distributions of games in which each player's belief about his impact on the societal choice goes to zero. After proving the existence of strategic equilibria, we show that all of them must be Nash. We also show that all regular equilibria of smooth non-atomic games are strategic. Moreover, it is displayed that in many economic applications, the set of strategic equilibria coincides with that of Nash equilibria of large finite games.

Journal of Economic Literature Classification Numbers: C72

Keywords: Nash equilibrium; Strategic equilibrium; Games with a continuum of players; Equilibrium distributions

---

*This is a revised version of chapter 4 in Barlo (2003). We thank Pedro Amaral, Kemal Badur, Ehud Kalai, Narayana Kocherlakota, Andy McLennan, Han Ozsoylev, Aldo Rusticini, David Schmeidler, Jan Werner and, specially, the editor Atsushi Kajii for helpful comments and suggestions. We benefited from discussions in the Mathematical Economics Workshop at the University of Minnesota. Financial support from Fundação para a Ciência e a Tecnologia is gratefully acknowledged. All remaining errors are ours.

[†]**Corresponding Author**: Faculty of Arts and Social Sciences, Sabanci University, Orhanli, Tuzla, 34956, Istanbul, Turkey; Phone: +90 216 483 9284; Fax: +90 216 483 9250 (CC. M. Barlo); email: barlo@sabanciuniv.edu.

[‡]School of Economics, University of Surrey, Guildford, GU2 7XH UK; phone: +44 1483 683475; email: g.carmona@surrey.ac.uk.

1 Introduction

Modeling economic situations featuring a large number of agents with non-atomic games is especially convenient because the inability of players to affect societal variables provides significant technical ease. However, this advantageous feature may result in the dismissal of the strategic behavior desired to be depicted. Although admittedly extreme, the following example delivers a clear portrait of this point: Consider a game where players' choices have to be in $\{0, 1\}$, and their payoffs depend only on the average choice. Because that a player's action does not affect the average choice and, thus, his own payoff, any player is indifferent between any of his choices, and as a result any strategy profile is a Nash equilibrium. On the other hand, the unique plausible Nash equilibrium is one where each player chooses the highest integer, because this strategy is the unique Nash equilibrium of the finite, but arbitrarily large, player version of the same game.

Such failure of (lower hemi) continuity of the equilibrium correspondence in non-atomic games casts some doubts on the usefulness of the continuum model. Indeed, Aumann (1964) regarded it as a mathematically convenient approximation to the "true" model featuring a finite number of players. But, unlike the non-atomic model in Aumann (1964) which provides a clean solution to the core-equivalence problem that would work only in an approximately way in finite models, the above example shows that, in some games, the continuum model is not a good approximation to the finite one. Further examples are given in Novshek and Sonnenschein (1983).

Naturally, this issue has been widely investigated and several reassuring results have been obtained (see, among many others, Hildenbrand (1974), Postlewaite and Schmeidler (1978) and Mas-Colell (1983)). However, for the class of games we consider, in general, Nash equilibria of non-atomic games correspond to limit points of approximate equilibria of sequences of finite-player games converging to the original (see Carmona and Podczeck (2011)). In an approximate equilibrium the action played by each one of a large fraction of players must yield a payoff close to the maximum he or she can achieve. And, in general, it is not possible to obtain a similar result using exact equilibria of the approximating large finite games even for regular equilibria. We show this using a notion of regular equilibria analogous to those of Harsanyi (1973) and van Damme (1991).

Given the above difficulties, the current paper proposes a refinement of Nash equilibrium in non-atomic games, *strategic equilibrium* (henceforth to be abbreviated by SE), designed to alleviate these problems in a tractable way. In fact, our goal is to develop an equilibrium concept for non-atomic games that intuitively has the same properties of the limit points of equilibria of large finite games (the precise meaning of this will be illustrated below) and, at the same time, its existence is generally guaranteed. Furthermore, the identification of SE is relatively easier compared with that of limit points of equilibria of large finite games (henceforth, limit equilibria). In other words, as in Aumann (1964), we want to keep the analytical convenience of non-atomic games and, at the same time, to focus on equilibria of non-atomic games that provide a more accurate approximation to the equilibria of finite-player versions of these games. Perhaps more importantly, we show that in non-atomic games with finitely many actions and payoff functions, the latter being sufficiently smooth (such a game is henceforth referred to as a smooth game), every regular equilibrium is a SE. In this light, SE can be regarded as an extension of regular equilibrium for general non-atomic games.

This study presents and analyzes the concept of SE for non-atomic games in which the payoff of each agent depends on what he chooses and on the distribution of actions chosen by the other players (which we refer to as the societal choice). For any non-atomic game and $\varepsilon > 0$, we define an ε-perturbed game by requiring each player to imagine that he alone has an ε impact on the societal choice. Then, the set of SE consists of limits of Nash equilibrium distributions of ε-perturbed games when ε tends to 0. It needs to be pointed out that in the ε-perturbed game, players are not rational as in Selten (1975). This is because each player thinks that he alone has an ε impact on the societal choice, and does not contemplate that others do the same consideration.

After proving the existence of SE distributions under standard assumptions (e.g., Mas-Colell (1984)) we show that the SE is a refinement of Nash equilibrium. Moreover, using the representation results of Khan and Sun (1995), Carmona (2008) and Carmona and Podczeck (2009), it is established that this analysis can be extended to strategy profiles whenever either one of the following holds: (1) the action space of every player is countable; or (2) the set of possible types of players is countable; or (3) the space of players is super-atomless.

The impact of focusing on SE is well illustrated in the above example: In the game where players choose either 0 or 1, there is only one SE which consists of almost all players choosing 1. Hence, the distribution of actions induced by the SE coincides with the distribution induced by the unique Nash equilibrium of the same game when played by a finite number of players.

A similar strong conclusion holds in the Nash's mass action game as well: A (finite) normal-form game is interpreted to consist of a finite number of positions (or islands), each characterized by a finite action space and a payoff function on the joint action space. One, then, imagines that the actual players in this game reside on one of those islands, players on the same island have identical payoffs and are equally likely to be chosen to play the game. Therefore, starting from the case where there is only one player on each island, we formulate associated replicas by symmetrically multiplying players on each island and assuming that each player on an island is equally likely to be selected. Hence, for any $k \in \mathbb{N}$, the k-replica game is one in which there are k players on each island who are equally likely to be selected to play the original game, and the payoff function and the action set of every player on an island are identical. It is, then, not difficult to see that for any $k \in \mathbb{N}$, a strategy is an equilibrium of the k-replica game if and only if the vector consisting of the average choices across players of a given island is a mixed strategy Nash equilibrium of the original game. However, this equivalence fails to hold in the limit case of a continuum of players on each island, each of whom are selected according to the Lebesgue measure. Indeed, in this case, no player can affect the average choice of the island they reside on, and thus, every strategy is a Nash equilibrium. However, when SE is employed, this equivalence is restored: We prove that a strategy profile in the non-atomic version is a SE if and only if the vector of the average choices across players on the same island is a mixed strategy Nash equilibrium of the original normal-form game.

Similar conclusions are reached in dynamic situations as well. After presenting the notion of strategic subgame perfect equilibrium (henceforth SSPE), we demonstrate that its use in the optimal taxation game of Levine and Pesendorfer (1995), instead of subgame perfect equilibrium (abbreviated by SPE), makes sure that the first-best can be obtained even with non-atomic players. Indeed, using the concept of SPE in non-atomic optimal taxation games, e.g. Chari and Kehoe (1989), the government cannot detect (thus, punish) individual deviations because one single agent cannot affect the societal choice, a phenomenon labeled as the "disappearance of information" by Levine and Pesendorfer (1995). Even though, the first-best is uniquely obtained in SPE in finite player versions of the same (extensive-form) game, it is well known that the second-best, the Ramsey Equilibrium, is the best possible with the use of SPE in non-atomic formulations. This, in turn, gives rise to discussions about whether or not the government may commit in order to achieve this particular payoff. Besides delivering a sharper conclusion that is not in "paradoxical" terms with that from finite player cases, this game is also of interest as it involves the use of SE with sequential rationality.

However, the set of SE does not equal the set of limit equilibria in general. In fact, we provide an example of a regular equilibrium of a smooth non-atomic game, hence of a SE, which fails to be a limit equilibrium. On the other hand, in the above examples, the notion of SE meets our desiderata of always existing and reproducing the (limit) properties of equilibria of the same game played by a large finite number of players.

It should be emphasized that our analysis is related to, but differs from that of Green (1980), Sabourian (1990), Levine and Pesendorfer (1995), and Carmona and Podczeck (2011) who try to justify the set of Nash equilibria of non-atomic games as limits of equilibria of large finite games with either noisy observations about deviating players or employing the ε-equilibrium concept. That is, we are not asking "when agents are negligible in large finite games", but rather analyzing equilibria of non-atomic games that are limits of equilibria of games where each player thinks that he alone is not negligible.

Section 2 describes the general framework of non-atomic games. In Section 3 we define the concept of SE and prove that it exists and is a refinement of Nash equilibrium. Section 4 considers regular equilibria of smooth non-atomic games. Finally, Section 5 involves Nash's mass action interpretation while Section 6 formalizes the notion of SSPE and displays its use in the optimal taxation game of Levine and Pesendorfer (1995).

2 Games with a measure space of players

In this section, we formally describe a class of games with a measure space of players. This class of games is a particular case of the model in Carmona and Podczeck (2014) although we follow Mas-Colell's (1984) distributional approach.

The set of players consists of a finite set $\bar{T}$ and a probability space $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ such that $\{t\} \in \hat{\Sigma}$ for all $t \in \hat{T}$ and $\bar{T} \cap \hat{T} = \emptyset$. The set of atomic players is $\bar{T}$ and the set of atomless players is $\hat{T}$. Let $T = \bar{T} \cup \hat{T}$.

The action set of each player $t \in \bar{T}$ is denoted by X_t and we let $X = \prod_{t \in \bar{T}} X_t$ and also

$X_{-t} = \prod_{t' \in \bar{T} \setminus \{t\}} X_{t'}$. We assume that X_t is a nonempty, compact and convex subset of a locally convex topological vector space for each $t \in \bar{T}$.

In order to accommodate general examples, such as the Nash's mass action game, we allow players' payoff functions to depend on the distribution of choices made by a finite number of subgroups of players in $\hat{T}$. Specifically, we consider a finite partition $\{\hat{T}_l\}_{l=1}^L$ of $\hat{T}$ such that $\hat{T}_l$ is measurable for all $l = 1, \ldots, L$. Players in $\hat{T}_l$ have a common action space A_l, assumed to be a nonempty compact metric space.

Let $l \in \{1, \ldots, L\}$. The payoff of each player $t \in \hat{T}_l$ depends on his choice $a \in A_l$, on the profile of choices $x \in X$ and on the vector $(\pi_1, \ldots, \pi_L)$ of distributions on $A_1, \ldots, A_L$ and it is assumed to be continuous. Given a complete separable metric space Y, let $M(Y)$ denote the space of Borel probability measures on Y endowed with the topology of the weak convergence of probability measures. For convenience, let $\mathcal{M} = M(A_1) \times \cdots \times M(A_L)$. Furthermore, let $\mathcal{U}_l$ denote the space of real-valued continuous payoff functions defined on $A_l \times X \times \mathcal{M}$ endowed with the sup norm. Payoff functions of players in $\hat{T}_l$ are described by a measurable function $\hat{U}_l : \hat{T}_l \to \mathcal{U}_l$.

Similarly, the payoff of each player $t \in \bar{T}$ depends on his choice $x \in X_t$, on the profile of choices $x_{-t} \in X_{-t}$ made by the other players in $\bar{T}$ and on the vector $(\pi_1, \ldots, \pi_L) \in \mathcal{M}$. We let u_t denote player t's payoff function and assume that u_t is continuous and also that the mapping $x \mapsto u_t(x, x_{-t}, \pi_1, \ldots, \pi_L)$ is quasi-concave for each $(x_{-t}, \pi_1, \ldots, \pi_L) \in X_{-t} \times \mathcal{M}$.

We summarize a game by a list $G = (\bar{T}, (\hat{T}, \hat{\Sigma}, \hat{\nu}), (A_l, \hat{U}_l)_{l=1}^L, (X_t, u_t)_{t \in \bar{T}})$. We say that G is *non-atomic* if $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ is atomless.

Let $G = (\bar{T}, (\hat{T}, \hat{\Sigma}, \hat{\nu}), (A_l, \hat{U}_l)_{l=1}^L, (X_t, u_t)_{t \in \bar{T}})$ be a non-atomic game and $\mathcal{C} = \prod_{l=1}^L M(\mathcal{U}_l \times A_l)$. Given a vector of Borel probability measures $(\tau_1, \ldots, \tau_L) \in \mathcal{C}$, we denote by $\tau_{l,\mathcal{U}_l}$ and τ_{l,A_l} the marginals of τ_l on $\mathcal{U}_l$ and A_l respectively. Given a game G, we say that $(x^*, \tau_1, \ldots, \tau_L) \in X \times \mathcal{C}$ is an *equilibrium distribution of* G if, for each $t \in \bar{T}$ and $x_t \in X_t$,

$$u_t(x^*, \tau_{1,A_1}, \ldots, \tau_{L,A_L}) \geq u_t(x_t, x^*_{-t}, \tau_{1,A_1}, \ldots, \tau_{L,A_L}), \tag{1}$$

and, for each $l = 1, \ldots, L$,

$$\hat{\nu}(\hat{T}_l)\tau_{l,\mathcal{U}_l}(B) = \hat{\nu}(\{t \in \hat{T}_l : U_l(t) \in B\}), \tag{2}$$

for each Borel measurable $B \subseteq \mathcal{U}_l$ and

$$\tau_l(\{(u, a) \in \mathcal{U}_l \times A_l : u(a, x^*, \tau_{1,A_1}, \ldots, \tau_{L,A_L}) \geq u(a', x^*, \tau_{1,A_1}, \ldots, \tau_{L,A_L}) \text{ for each } a' \in A\}) = 1. \tag{3}$$

Note that the above definition of an equilibrium distribution make sense only in the case where $\hat{\nu}$ is atomless because in (3) it is assumed that no player $t \in \hat{T}$ can affect the distribution of actions.

Note also that the above definition allows for the case where $\hat{T} = \emptyset$, in which case condition (2) holds trivially for any $(\tau_1, \ldots, \tau_L) \in \mathcal{C}$ and, consequently, neither (2) nor (3) imposes any meaningful restriction on the vector $(\tau_1, \ldots, \tau_L)$ of payoff-action distributions; in fact, these distributions play no role. Thus, if, in addition, u_t is independent of the vector of action distributions $(\pi_1, \ldots, \pi_L) \in \mathcal{M}$ for each $t \in \bar{T}$, then the game G is nothing but a standard finite-player normal-form game and (1) simply states that x^* is a Nash equilibrium of this finite-player normal-form game. In this case, we abuse on our terminology by saying that x^* is an equilibrium of G. Overall, the case where $\hat{T} = \emptyset$ and u_t is independent of $(\pi_1, \ldots, \pi_L)$ for each $t \in \bar{T}$ is rather special and trivial, but nevertheless it useful to keep it as a special case of our framework: this allows us to discuss issues regarding the relationship between SE and limit equilibria in a unified way (see Sections 4 and 5).

An alternative case arises when $\hat{\nu}(\hat{T}_l) > 0$ for all $l = 1, \ldots, L$. In this case, it is convenient to let $\hat{\nu}_l$ be defined by $\hat{\nu}_l(B) = \hat{\nu}(B)/\hat{\nu}(T_l)$ for each Borel measurable $B \subseteq \hat{T}_l$ and write (2) as

$$\tau_{l,\mathcal{U}_l} = \hat{\nu}_l \circ \hat{U}_l^{-1}$$

for each $l = 1, \ldots, L$.

It is also convenient to consider equilibrium strategies, which are defined as follows. Let $\mathbf{x}_l : \hat{T}_l \to A_l$ be measurable for each $l = 1, \ldots, L$. We say that $(x^*, \mathbf{x}_1, \ldots, \mathbf{x}_L)$ is an *equilibrium strategy of* G if $(x^*, \hat{\nu}_1 \circ (\hat{U}_1, \mathbf{x}_1)^{-1}, \ldots, \hat{\nu}_L \circ (\hat{U}_L, \mathbf{x}_L)^{-1})$ is an equilibrium distribution of G.

It is well known that every non-atomic game G as defined above has an equilibrium distribution (this also follows from Theorems 1 and 2 below). Furthermore, equilibrium strategies exist if either A_l or $\hat{U}_l(\hat{T}_l)$ (or both) are countable for all $l = 1, \ldots, L$, or if $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ is super-atomless (see, respectively, Khan and Sun (1995), Carmona (2008) and Carmona and Podczeck (2009)) but may fail to exist otherwise as shown by Khan, Rath, and Sun (1997).[1]

[1] Formally, a measure space (T, Σ, φ) is *super-atomless* if for every $E \in \Sigma$ with $\varphi(E) > 0$, the subspace of $L^1(\varphi)$ consisting of the elements of $L^1(\varphi)$ vanishing off E is non-separable. This notion was first introduced by Podczeck (2008); see also Podczeck (2009).

3 Strategic equilibria

As was stressed in the introduction, we wish to consider those equilibria that can be seen as a limit of equilibria in games in which each non-atomic player imagines that he alone has a small, yet positive, impact on the distribution of actions (societal choice) of the group he belongs to. Clearly, the need for a modification arises because for each player $t \in \hat{T}$, $\hat{\nu}(\{t\}) = 0$.

Associating a player with such a weight on his group's societal choice is done with the help of the following measures: For each $1 \leq l \leq L$, $\varepsilon > 0$, and $t \in \hat{T}_l$, let δ_t be the probability measure on T concentrated at t (i.e. $\delta_t(\{t\}) = 1$), and define a measure $\hat{\nu}_{l,t,\varepsilon} = \varepsilon\delta_t + (1-\varepsilon)\hat{\nu}_l$. Thus, under $\hat{\nu}_{l,t,\varepsilon}$ player t alone is an atom in group l with mass ε. In other words, in the game described by $\hat{\nu}_{l,t,\varepsilon}$, t believes that he alone has an ε impact on the societal choice of group l. In fact, for all strategies $\mathbf{x}_l : \hat{T}_l \to A_l$,

$$\hat{\nu}_{l,t,\varepsilon} \circ \mathbf{x}_l^{-1} = \varepsilon\delta_{\mathbf{x}_l(t)} + (1-\varepsilon)\hat{\nu}_l \circ \mathbf{x}_l^{-1}. \tag{4}$$

In order to construct a game where each player imagines that he, but no other player, has an ε impact on the distribution of the choices of the type he belongs to, we define the ε-perturbed game by altering players' payoff functions using the above measures.

Given a game $G = (\bar{T}, (\hat{T}, \hat{\Sigma}, \hat{\nu}), (A_l, \hat{U}_l)_{l=1}^L, (X_t, u_t)_{t\in\bar{T}})$, define, for all $\varepsilon > 0$, $1 \leq l \leq L$, $t \in \hat{T}_l$, $a \in A_l$, $x \in X$ and $\pi = (\pi_j)_{j=1}^L \in \mathcal{M}$,

$$\hat{U}_{l,\varepsilon}(t)\,(a, x, \pi) = \hat{U}_l(t)\,(a, x, (\varepsilon\delta_a + (1-\varepsilon)\pi_l, \pi_{-l}))\,. \tag{5}$$

We have that, for each $l = 1, \ldots, L$, $\hat{U}_{l,\varepsilon} : \hat{T}_l \to \mathcal{U}_l$ is measurable and that $U_{l,\varepsilon}(t)$ is continuous for each $t \in \hat{T}_l$ (see Section A.1). We then define the *ε-perturbed game G_ε* of G as $G_\varepsilon = (\bar{T}, (\hat{T}, \hat{\Sigma}, \hat{\nu}), (A_l, \hat{U}_{l,\varepsilon})_{l=1}^L, (X_t, u_t)_{t\in\bar{T}})$. Note that the ε-perturbed game G_ε has the same players and action spaces as the original non-atomic game G but different payoff functions for the players in $\hat{T}$, namely those defined in (5). Therefore, G_ε is a non-atomic game as defined in Section 2.

Definition 1 *We say that $(x^*, \tau_1^*, \ldots, \tau_L^*) \in X \times \mathcal{C}$ is a* strategic equilibrium in distribution *of G if there exists a sequence $\{\varepsilon_k\}_{k=1}^\infty \subseteq (0,1)$ decreasing to zero and a sequence $\{(x^k, \tau_1^k, \ldots, \tau_L^k)\}_{k=1}^\infty \subseteq X \times \mathcal{C}$ converging to $(x^*, \tau_1^*, \ldots, \tau_L^*)$ such that $(x^k, \tau_1^k, \ldots, \tau_L^k)$ is an equilibrium distribution of G_{ε_k}, for every $k \in \mathbb{N}$.*

In words, a strategic equilibrium in distribution (henceforth, to be abbreviated by the SED) can be approximated as a limit point of equilibrium distributions of games where each non-atomic player's belief about his sui generis ability to affect the societal choice converges to zero.

Theorem 1 establishes the existence of SED (see Section A.2 for its proof).

Theorem 1 *Every non-atomic game has a SED.*

Next, we show that any SED is an equilibrium distribution (see Section A.3 for its proof).

Theorem 2 *Every SED of a non-atomic game is an equilibrium distribution.*

We have achieved so far some of our goals, namely to define an equilibrium concept for games having a non-atomic (sub)set of players (i) without replacing it with a discrete approximation, (ii) such that it generally exists and (iii) refines the standard notion of an equilibrium distribution. In Sections 5 and 6.1 we show our other goal of having it in line with limit equilibria also holds in some interesting examples.

In general, working with equilibrium strategies in non-atomic games instead of equilibrium distributions is more intuitive and easier to formalize as such an approach involves an immediate generalization of finite player cases. Indeed, it has been employed by Schmeidler (1973), the initial paper to formalize non-atomic games. However, in general settings, the existence of an equilibrium strategy is not guaranteed (see Khan, Rath, and Sun (1997)) due to measurability constraints. For this reason, as Mas-Colell (1984) have shown, it is convenient to focus on equilibrium distributions as this approach eliminates the measurability problems involved in establishing the existence of an equilibrium strategy. Nevertheless, under some additional assumptions, equilibrium strategies do exist. As we will show, strategic equilibrium in strategies also exist under the same additional conditions.

The notion of a strategic equilibrium in strategy is as follows.

Definition 2 *A strategy* $(x, \mathbf{x}_1, \ldots, \mathbf{x}_L)$ *is a* strategic equilibrium in (behavioral) strategy *of* G *if there exists a sequence* $\{\varepsilon_k\}_{k=1}^{\infty} \subseteq (0, 1)$ *decreasing to zero and a sequence of strategies* $\{(x^k, \mathbf{x}_1^k, \ldots, \mathbf{x}_L^k)\}_{k=1}^{\infty}$ *such that* $(x^k, \mathbf{x}_1^k, \ldots, \mathbf{x}_L^k)$ *is an equilibrium strategy of* G_{ε_k} *for every* $k \in \mathbb{N}$ *and* $\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1}$ *converges to* $\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$ *for each* $l = 1, \ldots, L$.

This definition requires that the distribution on actions implied by a strategic equilibrium in strategy (henceforth, to be abbreviated by the SES) must coincide with a limit distribution on actions associated with a sequence of equilibrium distributions of games where each non-atomic player's belief about his sui generis ability to affect the societal choice converges to zero.

It is often more convenient to consider, in the above definition and for each $l = 1, \dots, L$, the sequence $\{\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l^k)^{-1}\}_{k=1}^{\infty}$ instead of $\{\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1}\}_{k=1}^{\infty}$ as the former does not require changes in the payoff functions of the players in $\hat{T}$. Theorem 3 shows that such replacement causes no change to the definition of a SES (see Section A.4 for its proof). Furthermore, it also shows that SED and SES are equivalent in the case where either A_l is countable or $\hat{U}_l(\hat{T}_l)$ is countable for all $l = 1, \dots, L$, or $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ is super-atomless – these are the cases where existence results for equilibrium strategies are known to hold.

Theorem 3 *Let G be a non-atomic game. Then the following conditions are equivalent:*

(a) $(x, \mathbf{x}_1, \dots, \mathbf{x}_L)$ *is a SES.*

(b) There exists a sequence $\{\varepsilon_k\}_{k=1}^{\infty} \subseteq (0,1)$ *decreasing to zero and a sequence of strategies* $\{(x^k, \mathbf{x}_1^k, \dots, \mathbf{x}_L^k)\}_{k=1}^{\infty}$ *such that* $(x^k, \mathbf{x}_1^k, \dots, \mathbf{x}_L^k)$ *is an equilibrium strategy of* G_{ε_k} *for every* $k \in \mathbb{N}$ *and* $\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l^k)^{-1}$ *converges to* $\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$ *for each* $l = 1, \dots, L$.

Furthermore, if either A_l *is countable or* $\hat{U}_l(\hat{T}_l)$ *is countable for all* $l = 1, \dots, L$, *or* $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ *is super-atomless, then both (a) and (b) are equivalent to*

(c) $(x, \hat{\nu}_1 \circ (\hat{U}_1, \mathbf{x}_1)^{-1}, \dots, \hat{\nu}_L \circ (\hat{U}_L, \mathbf{x}_L)^{-1})$ *is a SED.*

4 Games with finite characteristics

In this section we consider the case where players in $\hat{T}$ have a finite action space and a finite set of payoff functions. We adapt Harsanyi's (1973) notion of regular equilibrium to non-atomic games and show that every regular equilibrium is a SE. Furthermore, we show that this conclusion may fail regarding limit points of Nash equilibria of large finite-player games approaching a given non-atomic game.

For simplicity, we consider only one group of players in $\hat{T}$, ignore the set $\bar{T}$ and assume that $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ is atomless. We say that a game $G = ((\hat{T}, \hat{\Sigma}, \hat{\nu}), A, \hat{U})$ has *finite characteristics* if both A and $\hat{U}(\hat{T})$ are finite. Let $\mu = \hat{\nu} \circ \hat{U}^{-1}$ and $\mathbb{S}$ be the support of μ. In such game, the set $M(A)$ is identified with $\Delta_A = \{\pi \in \mathbb{R}_+^{|A|} : \sum_{a \in A} \pi_a = 1\}$ and, likewise, $M(\mathcal{U} \times A)$ is identified with $\Delta_{\mathbb{S} \times A} = \{\tau \in \mathbb{R}_+^{|\mathbb{S}||A|} : \sum_{(u,a) \in \mathbb{S} \times A} \tau_{(u,a)} = 1\}$.

We say that G is *smooth* if G has finite characteristics and, for each $u \in \mathbb{S}$, there exists an open set O_u in $\mathbb{R}^{|A|}$ and $\hat{u} : A \times O_u \to \mathbb{R}$ such that $\Delta_A \subseteq O_u$, $u = \hat{u}|_{A\times\Delta_A}$ and $\pi \mapsto \hat{u}(a, \pi)$ is continuously differentiable for each $a \in A$.

We turn to the definition of a regular equilibrium of a smooth game G. Let $O = \cap_{u\in S} O_u$ and $W = \{\tau \in \mathbb{R}^{|\mathbb{S}||A|} : \tau_A \in O\}$. Note that O is an open subset of $\mathbb{R}^{|A|}$ such that $\Delta_A \subseteq O$ and that W is an open subset of $\mathbb{R}^{|\mathbb{S}||A|}$ such that $\Delta_{\mathbb{S}\times A} \subseteq W$. Let $a^* = (a^*_u)_{u\in\mathbb{S}} \in A^{|\mathbb{S}|}$ and define $F^{a^*} : W \to \mathbb{R}^{|\mathbb{S}||A|}$ by setting, for each $\tau \in W$ and $(u, a) \in \mathbb{S} \times A$,

$$F^{a^*}_{u,a}(\tau) = \tau(u, a)[u(a, \tau_A) - u(a^*_u, \tau_A)] \text{ if } a \neq a^*_u$$

and

$$F^{a^*}_{u,a^*_u}(\tau) = \sum_{a\in A} \tau(u, a) - \mu(u).$$

Let $J^{a^*}(\tau^*)$ be the Jacobian of F^{a^*} at τ^*, i.e. $J^{a^*}(\tau^*) = \frac{\partial F^{a^*}(\tau^*)}{\partial \tau}$. We say that τ^* is a *regular equilibrium of* G if τ^* is an equilibrium of G and $J^{a^*}(\tau^*)$ is non-singular for some $a^* \in A^{|\mathbb{S}|}$ such that $\tau^*(u, a^*_u) > 0$ for each $u \in \mathbb{S}$.

Because we only consider cases where players in $\hat{T}$ have a finite action space and a finite set of payoff functions, by Theorem 3, SED and SES are equivalent; thus, in what follows we simply refer to any one of them as SE.

The main result of this section states that every regular equilibrium of a smooth game is a SE (see Section A.5 for its proof).

Theorem 4 *If $G = ((\hat{T}, \hat{\Sigma}, \hat{\nu}), A, \hat{U})$ is smooth, then every regular equilibrium of G is a SE.*

The above theorem tells that the well behaved and robust behavior associated with regular equilibrium prevails under the SE. Thus, the SE can be viewed as a coarse generalization of regular equilibrium (which demands that attention be restricted to smooth games) to non-smooth games. On the other hand, as we will show in what follows, an analog of Theorem 4 does not hold regarding limit points of equilibria of large finite-player games. Therefore, the notion of limit equilibrium is more demanding as it does not necessarily allow the behavior compatible with the regular equilibrium even when the game under consideration is smooth.

Let $G = ((\hat{T}, \hat{\Sigma}, \hat{\nu}), A, \hat{U})$ be smooth. We say that $\tau \in M(\mathbb{S} \times A)$ is a *limit equilibrium of* G if there exists a sequence $\{(G_k, f_k)\}_{k=1}^{\infty}$ such that, for each $k \in \mathbb{N}$,

(a) $G_k = (\bar{T}_k, (\hat{T}, \hat{\Sigma}, \hat{\nu}), A, \hat{U}, (X_t^k, u_t^k)_{t\in\bar{T}_k})$ is a game such that $\hat{T} = \emptyset$, $|\bar{T}_k| = k$, $X_t^k = A$ for each $t \in \bar{T}_k$ and there exists $U_k : \bar{T}_k \to \mathcal{U}$ such that

$$u_t^k(x) = U_k(t)(x_t, \nu_k \circ x^{-1})$$

for each $t \in \bar{T}_k$ and $x \in X_k$, where ν_k denotes the uniform distribution on $\bar{T}_k$,

(b) f_k is an equilibrium of G_k, and

(c) $\nu_k \circ (U_k, f_k)^{-1} \to \tau$.

The following is an example of a regular equilibrium of a smooth game which is not a limit equilibrium.[2] The example is a coordination game with two actions such that the payoff each player obtains from each of the two actions equals the fraction of those players who choose that action. Thus, with non-atomic players there is an equilibrium where exactly half of the players choose each one of the two actions. There is, however, no equilibrium with such property in the case of a finite space of players. For simplicity, suppose that there is a finite and even number n of players and exactly half of the population chooses each one of the two actions. Then each player can profitably deviate in order to reside with a strict majority: As a result of a deviation, say from a to b, there is a fraction of $1/2 + 1/n$ players choosing b whereas with no deviations there is a fraction of $1/2$ players choosing a. Since each player's payoff of each action equals the fraction of those choosing the same action, this deviation yields a gain of $1/n > 0$.

Example 1 *Let* $G = ((\hat{T}, \hat{\Sigma}, \hat{\nu}), A, \hat{U})$ *be such that* $A = \{\alpha, \beta\}$, $\hat{\nu} \circ \hat{U}^{-1}(u) = 1$ *where* $u : A \times M(A) \to \mathbb{R}$ *is such that, for each* $a \in A$ *and* $\pi \in M(A)$,

$$u(a, \pi) = \begin{cases} \pi(\alpha) & \text{if } a = \alpha, \\ 1 - \pi(\alpha) & \text{if } a = \beta. \end{cases}$$

We have that τ^* *such that* $\tau^*(u, \alpha) = 1/2 = \tau^*(u, \beta)$ *is a regular equilibrium of* G. *However,* τ^* *is not a limit equilibrium of* G *(see Section A.6 for a proof).*

[2]This example first appeared in Carmona, Páscoa, and Podczeck (2008) and it is presented in the current paper with their consent.

5 Mass-action interpretation of Nash equilibria

In his Ph.D. dissertation (see Nash (1950)), John Nash proposed two interpretations of his equilibrium concept, with the objective of showing how equilibrium points "(...) can be connected with observable phenomenon." One interpretation is rationalistic: if we assume that players are rational, know the full structure of the game, the game is played just once, and there is just one Nash equilibrium, then players will play according to that equilibrium.[3]

A second interpretation, that Nash referred to by *the mass action interpretation*, is less demanding on players: "[i]t is unnecessary to assume that the participants have full knowledge of the total structure of the game, or the ability and inclination to go through any complex reasoning processes." What is assumed is that there is a population of participants for each position in the game, which will be played throughout time by participants drawn at random from the different populations. If there is a stable average frequency with which each pure strategy is employed by the "average member" of the appropriate population, then this stable average frequency constitutes a mixed strategy Nash equilibrium.

Below we consider a continuum-of-player mass-action version of a normal-form game and we present a new interpretation of Nash equilibrium: The mixed strategy Nash equilibria of a given finite normal-form game are exactly the profiles of distributions over actions induced by the SE of its continuum-of-player mass-action version.

Consider a finite normal-form game $\Gamma = (N, (\Delta_{A_i}, v_i)_{i \in N})$, where $N = \{1, \ldots, n\}$ is the set of positions, Δ_{A_i} is the set of mixed strategies over the finite action set A_i, and v_i is the usual extension to mixed strategies of the payoff function. As in Nash's mass action interpretation, imagine that this game is played in a large society divided into n groups, from each of which a participant is drawn at random.

For any $k \in \mathbb{N}$, we define the k-replica game as follows: There are k players in each position, and we assume that each player is matched with $n-1$ players selected from the other positions. This gives rise to the k-replica game $G_k = (\bar{T}_k, (\hat{T}, \hat{\Sigma}, \hat{\nu}), A, \hat{U}, (X_t^k, u_t^k)_{t \in \bar{T}_k})$. In the game G_k, $\hat{T} = \emptyset$ and $\bar{T}_k = \{(i,j) : 1 \leq i \leq n, 1 \leq j \leq k\}$. Each player $t \in \bar{T}_k$ such that $t = (i,j)$ for some $1 \leq i \leq n$ and $1 \leq j \leq k$ has $X_t^k = \Delta_{A_i}$ as his action space. Under the assumption that all matchings are equally likely, the probability that an action $a \in A := A_1 \times \cdots \times A_n$ is played when players are using a strategy $\sigma = (\sigma_{i,j})_{i \in N, j=1,\ldots,k} \in X_k$

[3] For a formal discussion of these ideas, see Aumann and Brandenburger (1995) and Kuhn and et al. (1996).

is

$$\prod_{i=1}^{n} \sum_{j=1}^{k} \frac{\sigma_{i,j}(a_i)}{k}.$$

Let $\bar{\sigma} = (\bar{\sigma}_1, \ldots, \bar{\sigma}_n) \in \Delta_{A_1} \times \cdots \times \Delta_{A_n}$ be defined by $\bar{\sigma}_i(a_i) = \sum_{j=1}^{k} \sigma_{i,j}(a_i)/k$ and let the payoff of a player in position i be defined by

$$v_i^k(\sigma) = \sum_{a \in A} \left(\prod_{i' \in N} \bar{\sigma}_{i'}(a_{i'}) \right) v_i(a).$$

Then set, for each $t \in \bar{T}_k$ such that $t = (i, j)$ for some $1 \leq i \leq n$ and $1 \leq j \leq k$,

$$u_t^k(\sigma) = v_i^k(\sigma)$$

for each $\sigma \in X_k$. It is then easy to see (after going over the proof of Theorem 5) that for any $k \in \mathbb{N}$, σ is an equilibrium of G_k if and only if $\bar{\sigma}$ is a mixed strategy Nash equilibrium of Γ. In words, equilibria of G_k are precisely those strategies under which the average behavior in all positions is part of the same mixed strategy Nash equilibrium of the original game Γ. I.e., on average, every position is best-replying to the others.

Even though this equivalence holds for every $k \in \mathbb{N}$, it fails to do so in the limit case of a continuum of players. To see this consider the non-atomic game given by $G = (\bar{T}, (\hat{T}, \hat{\Sigma}, \hat{\nu}), (A_i, \hat{U}_i)_{i \in N}, (X_t, u_t)_{t \in \bar{T}})$. The set of players in G are such that $\bar{T} = \emptyset$, $\hat{T} = \cup_{i=1}^{n} [2i - 1, 2i]$ with Lebesgue measure on the Lebesgue measurable sets. Furthermore, let $\hat{T}_i = [2i-1, 2i]$ and $\hat{\nu}_i$ be the Lebesgue measure on $\hat{T}_i$ for all $i \in N$. Each player $t \in \hat{T}_i$ chooses an element of A_i. Regarding payoffs, for each $i \in N$, $t \in \hat{T}_i$, $a \in A_i$ and $(\pi_1, \ldots, \pi_n) \in \mathcal{M}$, let

$$\hat{U}_i(t) = v_i(\pi_1, \ldots, \pi_n).$$

The intuition behind the above definition of players' payoff functions is easily seen by considering the case of a strategy in G. Let $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)$ be such a strategy. A player is selected from each $\hat{T}_i$ according to the Lebesgue measure, and thus, the probability that the player selected from the ith group will play action $a_i \in A_i$ is $\hat{\nu}_i \circ \mathbf{x}_i^{-1}(a_i)$. We thus define $\bar{\mathbf{x}}_i = \hat{\nu}_i \circ \mathbf{x}_i^{-1}$ and then

$$\hat{U}_i(t)(a_i, \hat{\nu}_1 \circ \mathbf{x}_1^{-1}, \ldots, \hat{\nu}_n \circ \mathbf{x}_n^{-1}) = v_i(\bar{\mathbf{x}}_1, \ldots, \bar{\mathbf{x}}_n) = \sum_{a \in A} \left(\prod_{i' \in N} \bar{\mathbf{x}}_{i'}(a_{i'}) \right) v_i(a).^4$$

[4]The above notation is appropriate in the following sense: represent A_i by the unit vectors $\{e_1^i, \ldots, e_{|A_i|}^i\}$ in $\mathbb{R}^{|A|}$ and define $\hat{x}_i(t) = e_j^i$ if and only if $x_i(t) = a_j \in A_i$. Then, $\hat{\nu}_i \circ x^{-1} = \int_{T_i} \hat{x}_i \mathrm{d}\hat{\nu}_i$. Hence, $\hat{\nu}_i \circ x^{-1}$ can, in fact, be understood as an average.

It is easy to see that every strategy is a Nash equilibrium of G, because no $t \in T_i$ can affect $\bar{\mathbf{x}}_i$, $i = 1, \ldots, n$. On the other hand, the following theorem shows that the SE of G are characterized by the property that, on average, every position is best-replying to the others.[5] Hence, the distribution of actions induced by the SE of G correspond to the limit points of the corresponding distributions of equilibria of G_k when k converges to infinity.

Theorem 5 *A strategy profile $\mathbf{x}^* = (\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ is a SE of G if and only if $(\bar{\mathbf{x}}_1^*, \ldots, \bar{\mathbf{x}}_n^*)$ is a mixed strategy Nash equilibrium of Γ.*

Theorem 5, the proof of which can be found in Section A.7, provides a new interpretation of mixed strategy Nash equilibria: they constitute precisely the vector of distributions of actions, one for each position, that are induced by a (pure strategy) SE. Similarly as in Nash's mass action interpretation, a mixed strategy Nash equilibrium can be understood as a "stable" average behavior in a large society. However, since every SE is a Nash equilibrium (of the associated non-atomic game), our interpretation is rationalistic and so different from Nash's. Nevertheless, it is interesting to see that for our interpretation one needs to regard full rationality as a limit case of incomplete rationality as in Selten (1975).

6 Strategic subgame perfect equilibrium

After formally describing a class of extensive-form games with a measure space of players and finite periods, we introduce the notion of strategic subgame perfect equilibrium (SSPE) and display its use in the optimal taxation game of Levine and Pesendorfer (1995).

In order to sustain some ease of exposition, our formulation involves a setting where all players are assumed to choose actions in every period but the set of actions available to them may vary depending on the time index. Indeed, their available set of actions are allowed to be singleton sets. Moreover, in the current section we do not consider subgroups of non-atomic players. Additionally, as in Sabourian (1990) we concentrate on non-atomic extensive form games in which the choice of each non-atomic player does not depend on his own past choices.

The set of players consists of a finite set $\bar{T}$ and a probability space $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ such that $\{t\} \in \hat{\Sigma}$ for all $t \in \hat{T}$ and $\bar{T} \cap \hat{T} = \emptyset$. The set of atomic players equals $\bar{T}$ and here we focus

[5] As A_i is finite for all $i \in N$, Theorem 3 implies that SED and SES are equivalent; so we simply refer to any one of them as SE.

only on the cases where $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ is atomless. Let $T = \bar{T} \cup \hat{T}$.

The time is discreet and its index is given by $k = 1, \ldots, K$. Given $1 \leq k \leq K$, the set of actions allowed in period k is denoted by C_k and defined as follows: For any $t \in T$

$$C_{k,t} = \begin{cases} X_{k,t} & \text{if } t \in \bar{T} \\ A_k & \text{if } t \in \hat{T}. \end{cases}$$

where $X_{k,t}$ is a nonempty topological space for each $t \in \bar{T}$ and players in $\hat{T}$ have a common action space A_k which is assumed to be a nonempty compact metric space. We denote $X_k = \prod_{t \in \bar{T}} X_{k,t}$ and $X_{k,-t} = \prod_{t' \in \bar{T} \setminus \{t\}} X_{k,t'}$. Moreover, for any given $1 \leq k \leq K$, we denote $X^k = \prod_{l=1}^{k} X_k$ and $A^k = \prod_{l=1}^{k} A_k$. Recall that $M(A_k)$ denotes the space of Borel probability measures on A_k endowed with the topology of the weak convergence of probability measures. Similarly, denoting $M(A_k)$ by M_k, we let $M^k = \prod_{l=1}^{k} M_k$.

Now we describe the set of (public) *histories* which are denoted by H. The initial history is denoted by $\emptyset$ and we let $H_0 = \{\emptyset\}$. Furthermore, for each $1 \leq k \leq K$, we let $H_k = X^k \times M^k$ be the set of histories with length k. The set of histories is given by $H = \cup_{k=0}^{K} H_k$ while H_K is referred to as the set of terminal histories and $H \setminus H_K$ as the set of non-terminal histories. Moreover, for any $h \in H$, we say that the length of h, $\ell(h)$, equals k whenever $h \in H_k$.

In what follows, we specify players' *payoffs*. The payoff of each player $t \in \hat{T}$ depends on (i) his own choices and (ii) the profile of choices of players in $\bar{T}$ and (iii) the vector (taken across time) of distributions on the action sets of players in $t \in \hat{T}$ and it is assumed to be bounded. The payoff of $t \in \hat{T}$ is formally defined as follows: Let $\mathcal{P}$ denote the space of real-valued bounded payoff functions defined on $A^K \times H_K$ and endowed with the sup norm. Payoff functions of players in $\hat{T}$ are described by a measurable function $\hat{U} : \hat{T} \to \mathcal{P}$. Similarly, the payoff of each player $t \in \bar{T}$ is determined by (i) the profile of choices made by the players in $\bar{T}$ (including himself) and (ii) the vector of distributions on the action sets of players in $t \in \hat{T}$. We let u_t denote player t's payoff function, a bounded real-valued function defined on H_K.

A *non-atomic extensive-form game* with $K \in \mathbb{N}$ periods is, therefore,

$$\Gamma = (\bar{T}, (\hat{T}, \hat{\Sigma}, \hat{\nu}), (C_k)_{k=1}^{K}, \hat{U}, (u_t)_{t \in \bar{T}}).$$

Thus, a strategy of player $t \in \bar{T}$ is denoted by $\bar{s}_t$ and maps every $h \in H \setminus H_K$ into $X_{\ell(h)+1,t}$. We let the set of strategies of player $t \in \bar{T}$ be denoted by $\bar{S}_t$ and use the following notation: $\bar{s} = (\bar{s}_t)_{t\in\bar{T}}$ and $\bar{s}_{-t} = (\bar{s}_{t'})_{t'\in\bar{T}\setminus\{t\}}$ and $\bar{S} = \prod_{t\in\bar{T}} \bar{S}_t$ and $\bar{S}_{-t} = \prod_{t'\in\bar{T}\setminus\{t'\}} \bar{S}_{t'}$. Similarly, a strategy of player in $t \in \hat{T}$ is a function $\hat{\mathbf{s}}(t) : H \setminus H_K \to \cup_{l=1}^{K} A_k$ such that, for any $h \in H \setminus H_K$, $\hat{\mathbf{s}}_t(h) \in A_{\ell(h)+1}$ for all $t \in \hat{T}$ and $\hat{\mathbf{s}}(h) : \hat{T} \to A_{\ell(h)+1}$ is measurable. We denote such a strategy profile by $s = (\bar{s}, \hat{\mathbf{s}})$ and let S stand for the set of such strategy profiles.

Given a strategy profile $s \in S$, the (public) outcome path induced is given as follows: $\omega^1(s) = (\bar{s}(\emptyset), \hat{\nu}\circ\hat{\mathbf{s}}(\emptyset)^{-1}) \in H_1$ and $\omega^k(s) = (\bar{s}(\omega^1(s), \ldots, \omega^{k-1}(s)), \hat{\nu}\circ\hat{\mathbf{s}}(\omega^1(s), \ldots, \omega^{k-1}(s))^{-1}) \in X_k \times M_k$ for each $1 < k \le K$. Moreover, the outcome path induced by a strategy profile $s \in S$ at a given non-terminal public history $h \in H \setminus H_K$ is denoted by $\omega(s,h) = (\omega^l(s,h))_{l=1}^{K-\ell(h)}$ and defined by $\omega^1(s,h) = (\bar{s}(h), \hat{\nu}\circ\hat{\mathbf{s}}(h)^{-1}) \in X_{\ell(h)+1}\times M_{\ell(h)+1}$ and $\omega^k(s,h) = (\bar{s}(h, \omega^1(s,h), \ldots, \omega^{k-1}(s,h)), \hat{\nu}\circ\hat{\mathbf{s}}(h, \omega^1(s,h), \ldots, \omega^{k-1}(s,h))^{-1}) \in X_{\ell(h)+k} \times M_{\ell(h)+k}$ for any $1 < k \le K - \ell(h)$. It may be useful to point out that $\omega(s, \emptyset) = \omega(s)$.

The action path of a non-atomic player $t \in \hat{T}$ induced by a strategy profile $s \in S$ at a given non-terminal history $h \in H \setminus H_K$ is denoted by $a_t(s,h) = (a_t^l(s,h))_{l=1}^{K-\ell(h)}$ and defined by $a_t^1(s,h) = \hat{\mathbf{s}}_t(h)$ and $a_t^k(s,h) = \hat{\mathbf{s}}_t(h, \omega^1(s,h), \ldots, \omega^{k-1}(s,h))$ for any $1 < k \le K - \ell(h)$.

A *subgame perfect equilibrium* (SPE) of Γ is a strategy profile $s^* \in S$ such that for each $h \in H \setminus H_K$,

1. for all $t \in \bar{T}$, $u_t(\omega(s^*,h)) \ge u_t(\omega((\bar{s}'_t, s^*_{-t}), h))$ for all $\bar{s}'_t \in \bar{S}_t$, and
2. for all $t \in \hat{T}$, $\hat{U}_t(a_t(s^*,h), \omega(s^*,h)) \ge \hat{U}_t(\tilde{a}_t, \omega(s^*,h))$ for every profile $\tilde{a}_t = (\tilde{a}_t^k)_{k=\ell(h)+1}^{K} \in \prod_{m=\ell(h)+1}^{K} A_m$.

This definition makes sense only in the case where $\hat{\nu}$ is atomless because in the second condition it is assumed that no player $t \in \hat{T}$ can affect choices of players in $\bar{T}$ and the distribution of actions of players in $\hat{T}$.

The notion of SSPE consists of those equilibria that can be seen as a limit of SPE in non-atomic extensive-form games in which each non-atomic player imagines that he alone has a small, yet positive, impact on the distribution of actions.

Towards this regard consider the following measures: For each $\varepsilon > 0$ and $t \in \hat{T}$, let δ_t be the probability measure on $\hat{T}$ concentrated at t (i.e. $\delta_t(\{t\}) = 1$), and define a measure $\hat{\nu}_{t,\varepsilon} = \varepsilon\delta_t + (1-\varepsilon)\hat{\nu}$. Thus, under $\hat{\nu}_{t,\varepsilon}$ player t alone is an atom with mass ε. So for any given $h \in H \setminus Z$ and all strategies $\hat{\mathbf{s}}(h) : \hat{T} \to A_{\ell(h)+1}$,

$$\hat{\nu}_{t,\varepsilon} \circ (\hat{\mathbf{s}}(h))^{-1} = \varepsilon\delta_{\hat{\mathbf{s}}(h)(t)} + (1-\varepsilon)\left(\hat{\nu}\circ\hat{\mathbf{s}}(h)^{-1}\right).$$

Given a non-atomic extensive-form game with $K \in \mathbb{N}$ periods Γ, define, for all $\varepsilon > 0$ and $t \in \hat{T}$ and $(a_k, x_k, \pi_k)_{k=1}^K \in A^K \times X^K \times M^K$,

$$\hat{U}_\varepsilon(t)\left((a_k)_{k=1}^K, (x_k)_{k=1}^K, (\pi_k)_{k=1}^K\right) = \hat{U}(t)\left((a_k)_{k=1}^K, (x_k)_{k=1}^K, (\varepsilon\delta_{a_k} + (1-\varepsilon)\pi_k)_{k=1}^K\right).$$

Note that $\hat{U}_\varepsilon : \hat{T} \to \mathcal{P}$ is measurable (analogous to what is shown in Section A.1).

The *ε-perturbed non-atomic extensive-form game* Γ_ε of Γ is defined by

$$\Gamma_\varepsilon = (\bar{T}, (\hat{T}, \hat{\Sigma}, \hat{\nu}), (C_k)_{k=1}^K, \hat{U}_\varepsilon, (u_t)_{t\in\bar{T}}).$$

Note that the difference between Γ_ε and Γ involves only the payoff functions of players in $\hat{T}$. Therefore, Γ_ε is a non-atomic extensive-forme game as defined above.

Definition 3 *A* strategic subgame perfect equilibrium *of a given non-atomic extensive-form game Γ is a strategy profile $s^* = (\bar{s}^*, \hat{\mathbf{s}}^*) \in S$, if there are sequences $\{\varepsilon_n\}_{n=1}^\infty \subseteq (0,1)$ decreasing to zero and $\{s^n\}_{n=1}^\infty \subseteq S$ such that $\omega(s^n) \to \omega(s^*)$ and $s^n = (\bar{s}^n, \hat{\mathbf{s}}^n) \in S$ is a subgame perfect equilibrium strategy of Γ_{ε_n} for every $n \in \mathbb{N}$.*

In words, a SSPE is a strategy profile such that the outcome path induced equals to one obtained from a limit point of a sequence of SPE strategies of games where each non-atomic player's belief about his sui generis ability to affect the societal choice converges to zero.

6.1 Optimal taxation

The strategic interaction analyzed in this section concerns the optimal taxation game, example 3 of Levine and Pesendorfer (1995) (henceforth, LP). We show that the use of SSPE, instead of SPE (closely related to notion of precommitment equilibrium proposed by LP), makes sure that the first-best can be obtained even with a continuum of non-atomic players.[6] Indeed, it is well known that with finitely many players, the first-best can be supported via the unique SPE. However, when a continuum of agents is considered, LP shows that the unique precommitment equilibrium is not the first-best.

In what follows, we provide the nature of the strategic interaction and then define the game formally and describe the relationship between our formalization and that of LP. Finally, we present our result.

[6] The qualifier "precommitment" refers to the specification of the game in LP in which the large player (i.e. the government) precommits in the initial period to a reaction to the choices of the households.

The strategic interaction happens between the government, the large player L, and a continuum of identical households (small players) where the set of households is $\hat{T} = [0,1]$ and is endowed with the Lebesgue measure $\hat{\nu}$.

The interaction takes place over two periods: 1, 2. In the first period L has to commit itself to a reaction function mapping households' aggregate investment (which will be made in period 2) into the probability of implementing one of the two taxation schemes available to the government, a non-distortionary tax and a distortionary tax (these two tax schemes will be explained in the next paragraph). This reaction function of L is denoted by $\bar{\mathbf{x}} : [0,1] \to [0,1]$ where, for each $\alpha \in [0,1]$, $\bar{\mathbf{x}}(\alpha)$ denotes the probability of implementing the non-distortionary tax when aggregate investment equals α. Meanwhile, the households do not choose in the first period. In the second period, L does not choose while the non-atomic players choose how much of their unit endowment of capital to invest. Their investment decision is described by a measurable function $\hat{\mathbf{x}} : [0,1] \to [0,1]$ where, for each function $\bar{\mathbf{x}} : [0,1] \to [0,1]$, $\hat{\mathbf{x}}_t(\bar{\mathbf{x}}) \in [0,1]$ denotes the investment of $t \in [0,1]$ when the government has chosen $\bar{\mathbf{x}}$. Such an investment, in turn, delivers a gross return of $(1+r)\hat{\mathbf{x}}_t(\bar{\mathbf{x}})$, where $r > 0$ denotes the interest rate.

Given $\bar{\mathbf{x}}$ and $\hat{\mathbf{x}}$, L has to implement his precommitment decision concerning how to tax the households. Specifically, given the average investment behavior of households, $\int \hat{\mathbf{x}}$, taxation must be used in order to raise adequate revenue to cover a given amount of spending equal to $1+r$ and the following two tax schemes are available. The first tax scheme consists of a *non-distortionary tax* on the investments (collecting all the investments plus its interests). The second tax scheme consists of a *distortionary tax* on some other resource in the economy (say, labor). The government implements the non-distortionary tax with a probability determined by his precommitment and defined by $\bar{\mathbf{x}}(\int \mathbf{x}) \in [0,1]$.

The payoffs are described as follows. When L, the government, employs the non-distortionary tax on households' investments it collects all of households' investments and interests. In this case, the payoff of household t is his after-tax income which equals $1-\hat{\mathbf{x}}_t \geq 0$, the amount of capital net of the investment. With non-distortionary taxes, the tax collected suffices to cover the needed resources of the government if and only if (almost) every household invests at the maximum amount 1. Thus, when $\int \hat{\mathbf{x}} = 1$, the payoff of L equals 0, which is the same as the average (representative) households' payoff. However, when $\int \hat{\mathbf{x}} < 1$, the revenues collected does not suffice and due to a revenue shortfall L incurs a loss of $p > 1$ for every unit of revenue shortfall and this results in a total penalty the amount of $p(1-\int \hat{\mathbf{x}})$. So the government's payoff is the payoff of the average household minus the penalty resulting from the revenue shortfall: $(1-\int \hat{\mathbf{x}})(1-p) \leq 0$. Notice that with non-distortionary taxes $(1-\int \hat{\mathbf{x}})(1-p)$ equals the maximum payoff of L, i.e. 0, if and only if $\int \hat{\mathbf{x}} = 1$; otherwise, it is strictly negative.

If L uses distortionary taxation, then the amount of needed resources for the government are fully covered. However, because that this tax is distortionary, each household incurs a cost of $c > 1+r$. Therefore, with such taxes payoffs of household t consists of their endowment net of investment, i.e. $1-\hat{\mathbf{x}}_t$, plus the proceeds from their investment, i.e. $(1+r)\hat{\mathbf{x}}_t$ minus the cost from the use of a distortionary tax, c; hence, is equal to $(1+r\hat{\mathbf{x}}_t - c)$. Notice that this amount is strictly negative for every $\hat{\mathbf{x}}_t \in [0,1]$. As there are no revenue shortfalls with the use of distortionary taxes, the government obtains a payoff, equal to that of the average (representative) household, i.e. $1+r\int \hat{\mathbf{x}} - c$, which is guaranteed to be strictly negative.

In this setting, it is clear that the first-best is given by an outcome where $\bar{\mathbf{x}}(1) = 1$ and $\int \hat{\mathbf{x}} = 1$ meaning that the government chooses the non-distortionary taxes and the average households investment equals 1; in turn, delivering a payoff of 0 to all the players.

Next we describe this game in the formal setting of Section 6. Let $\bar{T} = \{L\}$ and $\hat{T} = [0,1]$ where $\hat{\nu}$ is given by the Lebesgue measure on $[0,1]$.

The time index is $k = 1,2$. The actions allowed in period k, $(C_k)_{k=1,2}$ are defined as follows. Let $\mathbb{X}$ be the space of (bounded) functions mapping $[0,1]$ into itself, endowed with the sup norm. We then let $C_{1,L}$ be a nonempty subset of $\mathbb{X}$ and $A_1 = \{0\}$. In this formalization, the assumption that non-atomic players do not choose in period 1 is modeled by letting their allowed set of actions equal to this singleton set; the reason for not specifying $C_{1,L} = \mathbb{X}$ will be clear shortly. In period 2, $C_{2,L} = \{0\} \subseteq \mathbb{R}$ and $A_2 = [0,1]$.

In this setting, due to the specific nature of the game, the set of histories takes a simple form: $H_1 = \{h \in H : h = (\bar{\mathbf{x}}, 0) \text{ with } \bar{\mathbf{x}} \in C_{1,L}\}$ as player's in $[0,1]$ are restricted to choose 0. In what follows, we will suppress non-atomic players' choices for any history in H_1 and simply refer such histories by $h = \bar{\mathbf{x}}$. The terminal histories take the form $H_2 = \{h' \in H : h' = (h,(0,\pi)) \text{ with } h \in H_1 \text{ and } \pi \in M([0,1])\}$. Since L is restricted to choose 0 in period two, we suppress such degenerate choices and denote any $h' \in H_2$ by $h' = (\bar{\mathbf{x}}, \pi)$.

Given any terminal history $(\bar{\mathbf{x}}, \pi) \in H_2$ and letting $i : [0,1] \to [0,1]$ denote the identity, the payoff of L is

$$u_L(\bar{\mathbf{x}}, \pi) = \left(1 - \bar{\mathbf{x}}\left(\int i \mathrm{d}\pi\right)\right)\left(1 + r\int i\mathrm{d}\pi - c\right) + \bar{\mathbf{x}}\left(\int i\mathrm{d}\pi\right)\left(1 - \int i\mathrm{d}\pi\right)(1-p).$$

The payoff of household $t \in [0,1]$ is, for each $a \in A_2$ and $(\bar{\mathbf{x}}, \pi) \in H_2$,

$$\hat{U}(t)(a, \bar{\mathbf{x}}, \pi) = \left(1 - \bar{\mathbf{x}}\left(\int i\mathrm{d}\pi\right)\right)(1 + ra - c) + \bar{\mathbf{x}}\left(\int i\mathrm{d}\pi\right)(1-a).$$

The above defines a non-atomic extensive-form game Γ_{ot}.

A delicate issue regarding the game Γ_{ot} concerns the specification of the government's choice set $C_{1,L}$ in period 1. Indeed, if we were to set $C_{1,L}$ equal to $\mathbb{X}$, then a SPE fails to exist. Indeed, there are $\bar{\mathbf{x}} \in \mathbb{X}$ such that there is no Nash equilibrium in the subgame induced by $\bar{\mathbf{x}}$, e.g. if $\bar{\mathbf{x}} = 1_{\{\alpha \in [0,1]:\alpha > 3/4\}}$.

LP addressed the above issue in an ingenious way by introducing the notion of precommitment equilibrium defined as follows. A pair $(\bar{\mathbf{x}}, \hat{\mathbf{x}})$ is a *Stackelberg response* if it satisfies for all $t \in [0,1]$

$$\hat{U}(t)\left(\hat{\mathbf{x}}_t, \bar{\mathbf{x}}, \hat{\nu} \circ \hat{\mathbf{x}}^{-1}\right) \geq \hat{U}(t)\left(a, \bar{\mathbf{x}}, \hat{\nu} \circ \hat{\mathbf{x}}^{-1}\right),$$

for all $a \in [0,1]$. Furthermore, a pair $(\bar{\mathbf{x}}^*, \hat{\mathbf{x}}^*)$ is a *precommitment equilibrium* if it is a Stackelberg response and

$$u_L\left(\bar{\mathbf{x}}^*, \hat{\nu} \circ (\hat{\mathbf{x}}^*)^{-1}\right) \geq u_L\left(\bar{\mathbf{x}}, \hat{\nu} \circ \hat{\mathbf{x}}^{-1}\right)$$

for all Stackelberg responses $(\bar{\mathbf{x}}, \hat{\mathbf{x}})$. Thus, only those strategies $\bar{\mathbf{x}}$ for which there exists an equilibrium in the subgame it induces are considered in the notion of precommitment equilibrium. Consequently, by letting $C_{1,L}$ be the set of $\bar{\mathbf{x}} \in \mathbb{X}$ such that $(\bar{\mathbf{x}}, \hat{\mathbf{x}})$ is a Stackelberg response for some measurable $\hat{\mathbf{x}} : [0,1] \to [0,1]$, we obtain that the set of SPE outcomes of Γ_{ot} coincides with its precommitment equilibria.

However, our goal is to consider SSPE and, therefore, we need to restrict the set $C_{1,L}$ further. Indeed, we need to consider $C_{1,L}$ so that, for some sequence $\{\varepsilon_n\}_{n=1}^{\infty} \subseteq (0,1)$ converging to zero, every $\bar{\mathbf{x}} \in C_{1,L}$ is such that there exists a Nash equilibrium in the subgame it induces in the ε-perturbed non-atomic extensive form game $\Gamma_{ot,\varepsilon}$ of Γ_{ot}. To be concrete, we set $C_{1,L}$ to be such set corresponding to the sequence $\{1/n\}_{n=1}^{\infty}$. The important remark to make is that $C_{1,L}$ is nonempty and, indeed, contains the strategies used in LP's analysis, namely the set of all constant functions $\bar{\mathbf{x}} : [0,1] \to [0,1]$ and the function $\bar{\mathbf{x}} = 1_1$ (see the proof of their Theorem 1).

Clearly, the first-best outcome cannot be supported by a SPE: When $\int \hat{\mathbf{x}}(\bar{\mathbf{x}}) = 1$ and $\bar{\mathbf{x}}(1) = 1$, then $(\bar{\mathbf{x}}, \hat{\mathbf{x}})$ is not a SPE: the best response of every $t \in [0,1]$ equals 0, and not 1. In fact, as LP have shown, the government's SPE payoff is strictly lower than the first-best level 0. In what follows, we show that the use of SSPE enables us to sustain the first-best payoff for L uniquely (see Section A.8 for its proof).

Theorem 6 *The non-atomic extensive-form game Γ_{ot} has a SSPE. Furthermore, player L's payoff in any SSPE of Γ_{ot} equals 0 (the first-best).*

The intuition behind this result is as follows. When $\varepsilon > 0$, in the ε-perturbed game every non-atomic player imagines that his deviation would be affecting the societal choice, thus, deviations would be identifiable by the government. Indeed, the government, in contrast to the non-atomic case, may employ the following strategy: Choose non-distortionary taxation whenever the societal choice with ε-perturbations is 1; otherwise, the government "punishes" the small players by choosing the distortionary tax. This, in turn, will make sure that the first-best can be obtained by SSPE. And it is the unique such equilibrium payoff because the government has the option to deviate to this strategy and the best responses of the households with ε-perturbations is uniquely determined as described above. Thus, the unique SSPE payoff of the government equals 0, the first-best.

A Appendix

A.1 Proof of the claims in Section 3

Claim 1 *$\hat{U}_{l,\varepsilon} : \hat{T}_l \to \mathcal{U}_l$ is measurable for each $l = 1, \ldots, L$.*

Proof. Fix $l \in \{1, \ldots, L\}$ and define, for all $a \in A_l$, $x \in X$ and $\pi \in \mathcal{M}$, $\hat{U}_{l,\varepsilon}^{(a,x,\pi)}$ by $t \mapsto U_{l,\varepsilon}(t)(a,x,\pi)$ and $U_l^{(a,x,\pi)}$ by $t \mapsto U_l(t)(a,x,\pi)$. Since $U_{l,\varepsilon}^{(a,x,\pi)} = U_l^{(a,(\varepsilon\delta_a + (1-\varepsilon)\pi_l, \pi_{-l}))}$ and $U_l^{(a,(\varepsilon\delta_a + (1-\varepsilon)\pi_l, \pi_{-l}))}$ is measurable by Carmona (2009, Proposition 1), it follows that $U_{l,\varepsilon}^{(a,x,\pi)}$ is also measurable. Then, it follows again by Carmona (2009, Proposition 1) that $U_{l,\varepsilon}$ is measurable. ■

Claim 2 *$\hat{U}_{l,\varepsilon}(t)$ is continuous for each $l = 1, \ldots, L$ and $t \in \hat{T}_l$.*

Proof. Let $l \in \{1, \ldots, L\}$ and $t \in T_l$. If $a \in A_l$, $x \in X$, $\pi \in \mathcal{M}$, $\{a_k\}_{k=1}^{\infty} \subseteq A_l$, $\{x_k\}_{k=1}^{\infty} \subseteq X$ and $\{\pi_k\}_{k=1}^{\infty} \subseteq \mathcal{M}$ are such that $\lim_k a_k = a$, $\lim_k x_k = x$ and $\lim_k \tau_k = \tau$ then $\varepsilon\delta_{a_k} +$

$(1-\varepsilon)\pi_l^k \to \varepsilon\delta_a + (1-\varepsilon)\pi_l$ and the continuity of $U_l(t)$ implies that $\lim_k U_{l,\varepsilon}(t)(a_k, x_k, \pi_k) = \lim_k U_l(t)(a_k, x_k, (\varepsilon\delta_{a_k} + (1-\varepsilon)\pi_l^k, \pi_{-l}^k)) = U_l(t)(a, (\varepsilon\delta_a + (1-\varepsilon)\pi_l, \pi_{-l})) = U_{l,\varepsilon}(t)(a, x, \pi)$. This concludes the proof. ■

A.2 Proof of Theorem 1

Let $\varepsilon > 0$. We start by showing that G_ε has an equilibrium distribution using an analogous argument to Mas-Colell (1984). Let $M = \{(\tau_1, \ldots, \tau_L) \in \mathcal{C} : \tau_{l,\mathcal{U}_l} = \hat{\nu}_l \circ \hat{U}_l^{-1}$ for each $l = 1, \ldots, L\}$, a nonempty, compact and convex subset of $\mathcal{C}$. Define a correspondence $\Phi : X \times M \rightrightarrows X \times M$ by setting, for each $(x, \tau_1, \ldots, \tau_L) \in X \times M$, $\Phi(x, \tau_1, \ldots, \tau_L)$ be the set of $(x', \tau_1', \ldots, \tau_L') \in X \times M$ such that

$$u_t(x_t', x_{-t}, \tau_{1,A_1}, \ldots, \tau_{L,A_L}) \geq u_t(\tilde{x}_t, x_{-t}, \tau_{1,A_1}, \ldots, \tau_{L,A_L})$$

for all $t \in \bar{T}$ and $\tilde{x}_t \in X_t$, and

$$\tau_l'(\{(u,a) \in \mathcal{U}_l \times A_l : u(a, x, \tau_{1,A_1}, \ldots, \tau_{L,A_L}) \geq u(\tilde{a}, x, \tau_{1,A_1}, \ldots, \tau_{L,A_L}) \text{ for each } \tilde{a} \in A_l\}) = 1$$

for each $l = 1, \ldots, L$. It is straightforward to show that Φ has a fixed point, which is an equilibrium distribution of G_ε.

To finish the proof, we let $(x^k, \tau_1^k, \ldots, \tau_L^k)$ be an equilibrium distribution of $G_{1/k}$ for each $k \in \mathbb{N}$. Since $\{\hat{U}_{l,1/k}\}_{k=1}^\infty$ converges uniformly to $\hat{U}_l$, then it follows that $\hat{\nu}_l \circ U_{l,1/k}^{-1}$ converges to $\hat{\nu}_l \circ U_l^{-1}$ for all $l = 1, \ldots, L$. For all $l \in \{1, \ldots, L\}$, let $K_l = \{\hat{\nu}_l \circ \hat{U}_l^{-1}, \hat{\nu}_l \circ \hat{U}_{l,1}^{-1}, \hat{\nu}_l \circ \hat{U}_{l,1/2}^{-1}, \ldots\}$ and $C_l = \{\mu \in M(\mathcal{U}_l \times A_l) : \mu_{l,\mathcal{U}_l} \in K_l\}$. It follows by Hildenbrand (1974, Theorems 32 and 33) that K_l is tight, and by Hildenbrand (1974, Theorems 34 and 35) that C_l is tight. Recall that X is compact. Since $\{x^k\}_{k=1}^\infty \subseteq X$ and $\{(\tau_1^k, \ldots, \tau_L^k)\}_{k=1}^\infty \subseteq C_1 \times \cdots \times C_L$, it follows (using Hildenbrand (1974, Theorem 31)) that $\{(x^k, \tau_1^k, \ldots, \tau_L^k)\}_{k=1}^\infty$ has a converging subsequence. Hence, its limit point is a SED of G.

A.3 Proof of Theorem 2

To simplify the notation, for each $\tau = (\tau_1, \ldots, \tau_L) \in \mathcal{C}$, $l \in \{1, \ldots, L\}$, $u \in \mathcal{U}_l$, $x \in X$ and $a \in A_l$, we let $\tau_A = (\tau_{1,A_1}, \ldots, \tau_{L,A_L})$ and write $u(a, x, \tau_A) \geq u(A, x, \tau_A)$ whenever $u(a, x, \tau_A) \geq u(a', x, \tau_A)$ for all $a' \in A_l$.

Let $(x, \tau_1, \ldots, \tau_L)$ be a SED and let $\{\varepsilon_k\}_k$ and $\{(x^k, \tau_1^k, \ldots, \tau_L^k)\}_k$ be such that $\varepsilon_k \in (0,1)$, $\lim_k \varepsilon_k = 0$, $(x^k, \tau_1^k, \ldots, \tau_L^k)$ converges to $(x, \tau_1, \ldots, \tau_L)$, and $(x^k, \tau_1^k, \ldots, \tau_L^k)$ is an equilibrium distribution of G_{ε_k}, for all $k \in \mathbb{N}$.

Note that $\{(u,a) \in \mathcal{U}_l \times A_l : u(a, x, \tau_A) \geq u(A, x, \tau_A)\}$ is closed and, for all $k \in \mathbb{N}$, $\tau_l^k(\{(u,a) \in \mathcal{U}_l \times A_l : u(a, x^k, \tau_A^k) \geq u(A, x^k, \tau_A^k)\}) = 1$ for each $l = 1, \ldots, L$. Hence, $\mathrm{supp}(\tau_l^k) \subseteq \{(u,a) \in \mathcal{U}_l \times A_l : u(a, x^k, \tau_A^k) \geq u(A, x^k, \tau_A^k)\}$ for each $l = 1, \ldots, L$.

We next show that $\mathrm{supp}(\tau_l) \subseteq \{(u,a) \in \mathcal{U}_l \times A_l : u(a, x, \tau_A) \geq u(A, x, \tau_A)\}$ for each $l = 1, \ldots, L$. Let $l \in \{1, \ldots, L\}$ and $(u^*, a^*) \in \mathrm{supp}(\tau_l)$. By Carmona and Podczeck (2009, Lemma 12), there exists a subsequence $\{\tau_l^{k_m}\}_m$ of $\{\tau_l^k\}_k$ and, for each $m \in \mathbb{N}$, $(u^m, a^m) \in \mathrm{supp}(\tau_l^{k_m})$ such that $\lim_m (u^m, a^m) = (u^*, a^*)$. Hence, for all $m \in \mathbb{N}$ and $a' \in A$, $u^m(a^m, x^{k_m}, \tau_A^{k_m}) \geq u^m(a', x^{k_m}, \tau_A^{k_m})$ and so $u^*(a^*, x, \tau_A) \geq u^*(a', x, \tau_A)$. Thus, $(u^*, a^*) \in \{(u,a) \in \mathcal{U}_l \times A_l : u(a, x, \tau_A) \geq u(A, x, \tau_A)\}$. It then follows that $\tau_l(\{(u,a) \in \mathcal{U}_l \times A_l : u(a, x, \tau_A) \geq u(A, x, \tau_A)\}) = 1$.

Furthermore, we have that $u_t(x^k, \tau_A^k) \geq u_t(x_t', x_{-t}^k, \tau_A^k)$ for all $t \in \bar{T}$, $x_t' \in X_t$ and $k \in \mathbb{N}$, which implies that $u_t(x, \tau_A) \geq u_t(x_t', x_{-t}, \tau_A)$ for all $t \in \bar{T}$ and $x_t' \in X_t$. This, together with what has been shown in the preceding paragraph, establishes that $(x, \tau_1, \ldots, \tau_L)$ is an equilibrium distribution of G.

A.4 Proof of Theorem 3

The equivalence between (a) and (b) follows from the fact that $\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1} \to \hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$ if and only if $\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l^k)^{-1} \to \hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$. To see the latter equivalence, fix $l \in \{1, \ldots, L\}$ and suppose first that $\hat{\nu}_l \circ (U_l, \mathbf{x}_l^k)^{-1} \to \hat{\nu}_l \circ (U_l, \mathbf{x}_l)^{-1}$.

Let $\varepsilon > 0$ and $h : \mathcal{U}_l \times A_l \to \mathbb{R}$ be a bounded uniformly continuous real-valued function. We will show that there exists $K \in \mathbb{N}$ such that $|\int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1} - \int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}| < \varepsilon$ for all $k \geq K$.

Since h is bounded, there exists $B > 0$ such that $||h|| \leq B$. Let $\eta > 0$ be such that $\eta < \varepsilon/[2(1 + 2B)]$. Since h is uniformly continuous, there exists $\delta > 0$ such that $|h(u,a) - h(u', a')| < \eta$ for all $u, u' \in \mathcal{U}_l$ and $a, a' \in A_l$ such that $||u - u'|| < \delta$ and $d(a, a') < \delta$. Since $\hat{U}_{l,\varepsilon_k}(t)$ converges uniformly to $\hat{U}_l(t)$ for each $t \in \hat{T}_l$, the function $f_k : \hat{T}_l \to \mathbb{R}$ defined by $f_k(t) = ||\hat{U}_{l,\varepsilon_k}(t) - \hat{U}_l(t)||$ for all $k \in \mathbb{N}$ and $t \in \hat{T}_l$ converges pointwise to zero. Hence, by Ergorov's Theorem, there exists a measurable $F \subseteq \hat{T}_l$ and $K' \in \mathbb{N}$ such that $\hat{\nu}_l(\hat{T}_l \setminus F) < \eta$ and $|f_k(t)| \leq \delta/2$ for all $t \in F$ and $k \geq K'$.

Since $\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l^k)^{-1} \to \hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$, there exists $K \in \mathbb{N}$ such that $K \geq K'$ and $|\int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l^k)^{-1} - \int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}| < \varepsilon/2$ for all $k \geq K$.

Hence, for all $k \geq K$,

$$
\begin{aligned}
&\left| \int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1} - \int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1} \right| \\
&\leq \left| \int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1} - \int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l^k)^{-1} \right| \\
&\quad + \left| \int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l^k)^{-1} - \int_{\mathcal{U}_l \times A_l} h \mathrm{d}\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1} \right| \\
&< \int_{\hat{T}_l} |h(\hat{U}_{l,\varepsilon_k}(t), \mathbf{x}_l^k(t)) - h(\hat{U}_l(t), \mathbf{x}_l^k(t))| \mathrm{d}\hat{\nu}_l(t) + \frac{\varepsilon}{2} \\
&= \int_{\hat{T}_l \setminus F} |h(\hat{U}_{l,\varepsilon_k}(t), \mathbf{x}_l^k(t)) - h(\hat{U}_l(t), \mathbf{x}_l^k(t))| \mathrm{d}\hat{\nu}_l(t) \\
&\quad + \int_F |h(\hat{U}_{l,\varepsilon_k}(t), \mathbf{x}_l^k(t)) - h(\hat{U}_l(t), \mathbf{x}_l^k(t))| \mathrm{d}\hat{\nu}_l(t) + \frac{\varepsilon}{2} < 2B\eta + \eta + \frac{\varepsilon}{2} < \varepsilon.
\end{aligned}
$$

Note that a similar argument to the above show that $\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1} \to \hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$ implies $\hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l^k)^{-1} \to \hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$.

We finally turn to the proof of the equivalence between (a) and (c). Suppose that $(x, \mathbf{x}_1, \ldots, x_L)$ is a SES of G. Let $\{\varepsilon_k\}_k$ and $\{(x_k, \mathbf{x}_1^k, \ldots, \mathbf{x}_L^k)\}_k$ be such that $\varepsilon_k \searrow 0$, $(x_k, \mathbf{x}_1^k, \ldots, \mathbf{x}_L^k)$ is a equilibrium strategy of G_{ε_k}, $x_k \to x$ and $\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1} \to \hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$ for each $l = 1, \ldots, L$. Since $(x_k, \hat{\nu}_1 \circ (\hat{U}_{1,\varepsilon_k}, \mathbf{x}_1^k)^{-1}, \ldots, \hat{\nu}_L \circ (\hat{U}_{L,\varepsilon_k}, \mathbf{x}_L^k)^{-1})$ is an equilibrium distribution of G_{ε_k}, then $(x, \hat{\nu}_1 \circ (\hat{U}_1, \mathbf{x}_1)^{-1}, \ldots, \hat{\nu}_L \circ (\hat{U}_L, \mathbf{x}_L)^{-1})$ is a SED of G.

Conversely, let $(x, \hat{\nu}_1 \circ (\hat{U}_1, \mathbf{x}_1)^{-1}, \ldots, \hat{\nu}_L \circ (\hat{U}_L, \mathbf{x}_L)^{-1})$ be a SED of G and let $\tau_l = \hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$ for each $l = 1, \ldots, L$. Let $\{\varepsilon_k\}_k$ and $\{(x_k, \tau_1^k, \ldots, \tau_L^k)\}_k$ be such that $\varepsilon_k \searrow 0$, $(x_k, \tau_1^k, \ldots, \tau_L^k)$ is an equilibrium distribution of G_{ε_k} and $(x_k, \tau_1^k, \ldots, \tau_L^k) \to (x, \tau_1, \ldots, \tau_L)$. Then, since either A_l is countable or $\hat{U}_l(\hat{T}_l)$ is countable for all $l = 1, \ldots, L$, or $(\hat{T}, \hat{\Sigma}, \hat{\nu})$ is super-atomless, it follows by Khan and Sun (1995, Theorem 2), Carmona (2008, Theorem 1) or Carmona and Podczeck (2009, Lemma 7) respectively that there exist $(\mathbf{x}_1^k, \ldots, \mathbf{x}_L^k)$ such that $(x_k, \mathbf{x}_1^k, \ldots, \mathbf{x}_L^k)$ is an equilibrium strategy of G_{ε_k} and $\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1} = \tau_l^k$ for all $k \in \mathbb{N}$ and $l = 1, \ldots, L$. Since $\hat{\nu}_l \circ (\hat{U}_{l,\varepsilon_k}, \mathbf{x}_l^k)^{-1} = \tau_l^k \to \tau_l = \hat{\nu}_l \circ (\hat{U}_l, \mathbf{x}_l)^{-1}$ for each $l = 1, \ldots, L$, it follows that $(x, \mathbf{x}_1, \ldots, \mathbf{x}_L)$ is a SES of G.

A.5 Proof of Theorem 4

The proof follows those of Lemma 2.5.2 and Theorem 2.5.5 in van Damme (1991).

Let G be smooth, τ be a regular equilibrium of G and a_u^* be such that $\tau_{u,a_u^*} > 0$ for each $u \in \mathbb{S}$. We first show that

$$\{a \in A : \tau_{u,a} = 0\} = \{a \in A : u(a, \tau_A) < u(a_u^*, \tau_A)\} \text{ for each } u \in \mathbb{S}. \tag{6}$$

This can be established as follows. Since τ is an equilibrium of G, then $\{a \in A : u(a, \tau_A) < u(a_u^*, \tau_A)\} \subseteq \{a \in A : \tau_{u,a} = 0\}$ for each $u \in \mathbb{S}$. To show the additional inclusion, let $\tilde{J}$ be obtained from $J^{a^*}(\tau)$ by crossing out the rows and columns corresponding to $(u, a) \in \mathbb{S} \times A$ such that $\tau_{u,a} = 0$. Since $\tau_A = (\sum_{u \in A} \tau_{u,a})_{a \in A}$, we have that, for each $(u, a) \in \mathbb{S} \times A$ such that $\tau_{u,a} = 0$,

$$\frac{\partial F_{u,a}^{a^*}(\tau)}{\partial \hat{\tau}_{u,a}} = u(a, \tau_A) - u(a_u^*, \tau_A) + \tau_{u,a} \left[\frac{\partial u(a, \tau_A)}{\partial \pi_a} - \frac{\partial u(a_u^*, \tau_A)}{\partial \pi_a}\right] = u(a, \tau_A) - u(a_u^*, \tau_A), \text{ and}$$
$$\frac{\partial F_{u,a}^{a^*}(\tau)}{\partial \hat{\tau}_{u',a'}} = \tau_{u,a} \left[\frac{\partial u(a, \tau_A)}{\partial \pi_{a'}} - \frac{\partial u(a_u^*, \tau_A)}{\partial \pi_{a'}}\right] = 0$$

for any $(u', a') \in \mathbb{S} \times A$ such that $(u', a') \neq (u, a)$. This implies that

$$\det(J^{a^*}(\tau)) = \det(\tilde{J}) \prod_{(u,a) \in \mathbb{S} \times A : \tau_{u,a} = 0} [u(a, \tau_A) - u(a_u^*, \tau_A)]\,. \tag{7}$$

Now consider $(u, a) \in \mathbb{S} \times A$ such that $u(a, \tau_A) = u(a_u^*, \tau_A)$. If $\tau_{u,a} = 0$, then (7) implies $\det(J^{a^*}(\tau)) = 0$, a contradiction. Hence, it must be $\tau_{u,a} > 0$. Thus, $\{a \in A : u(a, \tau_A) = u(a_u^*, \tau_A)\} \subseteq \{a \in A : \tau_{u,a} > 0\}$ for each $u \in \mathbb{S}$, i.e. $\{a \in A : \tau_{u,a} = 0\} \subseteq \{a \in A : u(a, \tau_A) < u(a_u^*, \tau_A)\}$ for each $u \in \mathbb{S}$. This concludes the proof of (6).

Let V be a neighborhood of 0 in $\mathbb{R}$ such that, for each $\varepsilon \in V$ and $a \in A$, $\varepsilon\delta_a + (1-\varepsilon)\pi \in W$, where δ_a is now regarded as the vector in Δ_A with its coordinate corresponding to a being equal to 1. Define $F : W \times V \to \mathbb{R}^{|\mathbb{S}||A|}$ by setting, for each $(\tau, \varepsilon) \in W \times V$ and $(u, a) \in \mathbb{S} \times A$,

$$F_{u,a}(\tau, \varepsilon) = \tau(u, a)[u(a, \varepsilon\delta_a + (1-\varepsilon)\tau_A) - u(a_u^*, \varepsilon\delta_a + (1-\varepsilon)\tau_A)] \text{ if } a \neq a_u^*$$

and

$$F_{u,a_u^*}(\tau) = \sum_{a \in A} \tau(u, a) - \mu(u).$$

Since τ is an equilibrium of G and $\tau_{u,a_u^*} > 0$ for each $u \in \mathbb{S}$, then $F(\tau, 0) = 0$. Since τ is a regular equilibrium of G, then

$$\frac{\partial F(\tau)}{\partial \hat{\tau}} = \frac{\partial F^{a^*}(\tau)}{\partial \hat{\tau}}$$

is non-singular.

It then follows by the implicit function theorem (see, e.g., Rudin (1976, Theorem 9.28, p.224)) that there exists an open neighborhood W^* of τ, an open neighborhood V^* of 0 and a differentiable function $f : V^* \to W^*$ such that, for each $(\tau', \varepsilon) \in W^* \times V^*$,

$$F^*(\tau', \varepsilon) = 0 \text{ if and only if } \tau' = f(\varepsilon). \tag{8}$$

Shrinking W^* and V^* if necessary, we may assume that

$$\tau'_{u,a} > 0 \text{ for all } (u, a) \in \mathbb{S} \times A \text{ such that } \tau_{u,a} > 0 \tag{9}$$

holds for each $\tau' \in W^*$ and that

$$u(a, \varepsilon\delta_a + (1-\varepsilon)\tau'_A) < u(a^*_u, \varepsilon\delta_a + (1-\varepsilon)\tau'_A) \text{ for all } (u, a) \in \mathbb{S} \times A \text{ such that } \tau_{u,a} = 0 \tag{10}$$

holds for each $(\tau', \varepsilon) \in W^* \times V^*$. Note that (10) is possible due to (6).

Hence, for each $\varepsilon \in V^*$, letting $\tau' = f(\varepsilon) \in W^*$, we have that the following holds. First, by (9),

$$\tau'_{u,a} > 0 \text{ for all } (u, a) \in \mathbb{S} \times A \text{ such that } \tau_{u,a} > 0. \tag{11}$$

Second, by (8) and (11),

$$u(a, \varepsilon\delta_a + (1-\varepsilon)\tau'_A) = u(a^*_u, \varepsilon\delta_a + (1-\varepsilon)\tau'_A) \text{ for all } (u, a) \in \mathbb{S} \times A \text{ such that } \tau_{u,a} > 0. \tag{12}$$

Third, by (10),

$$u(a, \varepsilon\delta_a + (1-\varepsilon)\tau'_A) < u(a^*_u, \varepsilon\delta_a + (1-\varepsilon)\tau'_A) \text{ for all } (u, a) \in \mathbb{S} \times A \text{ such that } \tau_{u,a} = 0. \tag{13}$$

Forth, and finally, by (8) and (13),

$$\tau'_{u,a} = 0 \text{ for all } (u, a) \in \mathbb{S} \times A \text{ such that } \tau_{u,a} = 0. \tag{14}$$

Let $\mathbf{x} : \hat{T} \to A$ be an equilibrium strategy such that $\hat{\nu} \circ (\hat{U}, \mathbf{x})^{-1} = \tau$. Let $\{\varepsilon_k\}_{k=1}^{\infty} \subseteq (0, 1)$ be such that $\varepsilon_k \downarrow 0$ and $\varepsilon_k \in V^*$ for all $k \in \mathbb{N}$. Fix $k \in \mathbb{N}$ and let $\tau_k = f(\varepsilon_k)$ and $\mathbf{x}_k : \hat{T} \to A$ be such that $\hat{\nu} \circ (\hat{U}, \mathbf{x}_k)^{-1} = \tau_k$. Such $\mathbf{x}_k$ exists because A is finite and because, by (8), $\tau_{k,\mathbb{S}}(u) = \sum_{a \in A} \tau_k(u, a) = \mu(u)$ for each $u \in \mathbb{S}$. Then (11)–(14) imply that $\mathbf{x}_k$ is an equilibrium strategy of G_{ε_k}. Furthermore,

$$\hat{\nu} \circ (\hat{U}, \mathbf{x}_k)^{-1} = \tau_k = f(\varepsilon_k) \to f(0) = \tau = \hat{\nu} \circ (\hat{U}, \mathbf{x})^{-1}.$$

Hence, by Theorem 3, $\mathbf{x}$ is a SES and τ is a SED. These complete the proof.

A.6 Proof of the claims made in Example 1

It is clear that G is a smooth game by letting $\hat{u}(\alpha, \pi) = \pi(\alpha)$ and $\hat{u}(\beta, \pi) = 1 - \pi(\alpha)$ for each $\pi \in \mathbb{R}^{|A|}$. It is also clear that τ^* is an equilibrium of G.

Let $a^* = \beta$ and note that, for each $\tau \in \mathbb{R}^{|S||A|}$,

$$F^{a^*}_{u,\alpha}(\tau) = \tau_{u,\alpha}[2\tau_{u,\alpha} - 1], \text{ and}$$
$$F^{a^*}_{u,\beta}(\tau) = \tau_{u,\alpha} + \tau_{u,\beta} - 1.$$

Then

$$J^{a^*}(\tau^*) = \begin{pmatrix} 4\tau^*_{u,\alpha} - 1 & 0 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

and is, therefore, non-singular. This establishes that τ^* is a regular equilibrium of G.

We next show that τ^* is not a limit equilibrium of G. Suppose not. Then there exists $\{(G_k, f_k)\}_{k=1}^{\infty}$ satisfying the conditions in the definition of limit equilibrium. Let $k \in \mathbb{N}$ be sufficiently large such that $\tau_k(u, \alpha) > 0$ and $\tau_k(u, \beta) > 0$, where $\tau_k = \nu_k \circ (U_k, f_k)^{-1}$. For convenience, let $\pi_k = \tau_k(u, \alpha)$. Since f_k is a Nash equilibrium of G_k, then (1) applied to $t \in \bar{T}_k$ such that $f_k(t) = \alpha$ implies that

$$\pi_k \geq 1 - (\pi_k - 1/k) \Leftrightarrow 2\pi_k - 1 \geq 1/k. \tag{15}$$

Similarly, (1) applied to $t \in \bar{T}_k$ such that $f_k(t) = \beta$ implies that

$$1 - \pi_k \geq \pi_k + 1/k \Leftrightarrow 2\pi_k - 1 \leq -1/k. \tag{16}$$

Combining (15) and (16), we obtain $-1/k \geq 1/k$, a contradiction. This contradiction shows that τ^* is not a limit equilibrium of G.

A.7 Proof of Theorem 5

(Sufficiency) Let $(\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ be a strategy in G, and assume that $\bar{\mathbf{x}}^* = (\bar{\mathbf{x}}_1^*, \ldots, \bar{\mathbf{x}}_n^*)$ is a mixed strategy Nash equilibrium of Γ. Let $i \in N$, $t \in \hat{T}_i$, $\varepsilon > 0$ and $a \in A$. We have that

$$\hat{U}_{i,\varepsilon}(\mathbf{x}^*(t), \bar{\mathbf{x}}_1^*, \ldots, \bar{\mathbf{x}}_n^*) \geq \hat{U}_{i,\varepsilon}(a, \bar{\mathbf{x}}_1^*, \ldots, \bar{\mathbf{x}}_n^*). \tag{17}$$

Indeed, we have that $v_i(\bar{\mathbf{x}}^*) \geq v_i(\sigma_i, \bar{\mathbf{x}}_{-i}^*)$ for all $\sigma_i \in \Delta_{A_i}$, which implies, in particular, that $v_i(\bar{\mathbf{x}}^*) \geq v_i(\varepsilon a + (1-\varepsilon)\bar{\mathbf{x}}_i^*, \bar{\mathbf{x}}_{-i}^*)$ for all $a \in A_i$. It follows from (17) that $(\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ is an equilibrium strategy of G_ε for all $\varepsilon > 0$ and, thus, a SE of G.

(Necessity) Let $(\mathbf{x}_1^*, \ldots, \mathbf{x}_n^*)$ be a SE of G, and let $\bar{\mathbf{x}}^* = (\bar{\mathbf{x}}_1^*, \ldots, \bar{\mathbf{x}}_n^*)$. We show that for all $i \in N$, and $a_i \in A_i$ if $\bar{\mathbf{x}}_i^*(a_i) > 0$, then a_i maximizes $\hat{a}_i \mapsto v_i(\hat{a}_i, \bar{\mathbf{x}}_{-i}^*)$ in A_i, which then implies that $\bar{\mathbf{x}}^*$ is a mixed strategy Nash equilibrium of Γ. Let $i \in N$ and $a_i \in A_i$. If a_i does not maximize $\hat{a}_i \mapsto v_i(\hat{a}_i, \bar{\mathbf{x}}_{-i}^*)$ in A_i, then a_i does not maximize $\hat{a}_i \mapsto v_i(\hat{a}_i, \bar{\mathbf{x}}_{-i}^\varepsilon)$ in A_i for all $\varepsilon > 0$ sufficiently small, where $\bar{\mathbf{x}}^\varepsilon := (\hat{\nu}_1 \circ (\mathbf{x}_1^\varepsilon)^{-1}, \ldots, \hat{\nu}_n \circ (\mathbf{x}_n^\varepsilon)^{-1})$, and $(\mathbf{x}_1^\varepsilon, \ldots, \mathbf{x}_n^\varepsilon)$ is an equilibrium strategy of G_ε, $\mathbf{x}_i^\varepsilon$ converges in distribution to $\mathbf{x}_i^*$ for all $i \in N$ and $\varepsilon \to 0$. Since

$$\hat{U}_{i,\varepsilon}(t)(\hat{a}_i, \bar{\mathbf{x}}^\varepsilon) = \varepsilon v_i(\hat{a}_i, \bar{\mathbf{x}}_{-i}^\varepsilon) + (1-\varepsilon) v_i(\bar{\mathbf{x}}^\varepsilon)$$

and $\mathbf{x}^\varepsilon$ is an equilibrium strategy of G_ε, then $\mathbf{x}_i^\varepsilon(t) \neq a_i$ a.e. $t \in \hat{T}_i$ and so $\bar{\mathbf{x}}_i^\varepsilon(a_i) = 0$. Thus, $\bar{\mathbf{x}}_i^*(a_i) = 0$ since $\mathbf{x}_i^\varepsilon$ converges in distribution to $\mathbf{x}_i^*$.

A.8 Proof of Theorem 6

Let $\varepsilon > 0$ and consider $(\bar{\mathbf{x}}^*, \hat{\mathbf{x}}^{*,\varepsilon})$ defined by $\bar{\mathbf{x}}^* = 1_1$, $\hat{\mathbf{x}}^{*,\varepsilon}(\bar{\mathbf{x}}^*) \equiv 1$ and, for each $\bar{\mathbf{x}} \in C_{1,L} \setminus \{\bar{\mathbf{x}}^*\}$, let $\hat{\mathbf{x}}_t^{*,\varepsilon}(\bar{\mathbf{x}})$ be any equilibrium of the subgame induced by $\bar{\mathbf{x}}$ in $\Gamma_{ot,\varepsilon}$.

Since $\varepsilon a + (1-\varepsilon)\int \hat{\mathbf{x}}^* = 1$ if and only if $a = 1$, we have that

$$\hat{U}_\varepsilon(t)\left(a, \bar{\mathbf{x}}^*, \hat{\nu} \circ \hat{\mathbf{x}}^{*,\varepsilon}(\bar{\mathbf{x}}^*)^{-1}\right) = \begin{cases} 0 & \text{if} \quad a = 1, \\ 1 + ra - c & \text{if} \quad a < 1. \end{cases}$$

Since $1 + ra - c < 0$, $a \in [0,1]$, then $\hat{\mathbf{x}}_t^{*,\varepsilon}(\bar{\mathbf{x}}^*)$ solves $\max_{a \in A_2} \hat{U}_\varepsilon(t)\left(a, \bar{\mathbf{x}}^*, \hat{\nu} \circ \hat{\mathbf{x}}^{*,\varepsilon}(\bar{\mathbf{x}}^*)^{-1}\right)$. Furthermore, $u_L\left(\bar{\mathbf{x}}^*, \hat{\nu} \circ \hat{\mathbf{x}}^{*,\varepsilon}(\bar{\mathbf{x}}^*)^{-1}\right) = 0 \geq u_L\left(\bar{\mathbf{x}}, \hat{\nu} \circ \hat{\mathbf{x}}(\bar{\mathbf{x}})^{-1}\right)$ for all $\bar{\mathbf{x}} \in C_{1,L}$. Hence, $(\bar{\mathbf{x}}^*, \hat{\mathbf{x}}^{*,\varepsilon})$ is a SPE of $\Gamma_{ot,\varepsilon}$. It is then clear that $(\bar{\mathbf{x}}^*, \hat{\mathbf{x}}^*)$ with $\hat{\mathbf{x}}^*(\bar{\mathbf{x}}^*) \equiv 1$ is a SSPE of Γ_{ot}. This establishes existence of a SSPE of Γ_{ot}.

We next show the uniqueness of the SSPE payoff of L. Let $(\bar{\mathbf{x}}, \hat{\mathbf{x}})$ be any SSPE and let $\{\varepsilon_n\}_{n=1}^\infty$ with $\varepsilon_n \searrow 0$ and $\{\bar{\mathbf{x}}^n, \hat{\mathbf{x}}^n\}_{n=1}^\infty$ be such that $\omega(\bar{\mathbf{x}}^n, \hat{\mathbf{x}}^n) \to \omega(\bar{\mathbf{x}}^n, \hat{\mathbf{x}}^n)$ and $(\bar{\mathbf{x}}^n, \hat{\mathbf{x}}^n)$ is a SPE of $\Gamma_{ot,\varepsilon_n}$ for all $n \in \mathbb{N}$. Then we have that, for each n, $0 \geq u_L\left(\bar{\mathbf{x}}^n, \int i\mathrm{d}\hat{\nu} \circ (\hat{\mathbf{x}}^n)^{-1}\right) \geq u_L\left(\bar{\mathbf{x}}^*, \int i\mathrm{d}\hat{\nu} \circ (\hat{\mathbf{x}}^*)^{-1}\right) = 0$. Hence, $u_L\left(\bar{\mathbf{x}}^n, \int i\mathrm{d}\hat{\nu} \circ \hat{\mathbf{x}}^n(\bar{\mathbf{x}}^n)^{-1}\right) = 0$ for all $n \in \mathbb{N}$. This, in turn, implies that $\int \hat{\mathbf{x}}^n(\bar{\mathbf{x}}^n) = \int i\mathrm{d}\hat{\nu} \circ \hat{\mathbf{x}}^n(\bar{\mathbf{x}}^n)^{-1} = 1$, $\hat{\nu} \circ \hat{\mathbf{x}}^n(\bar{\mathbf{x}}^n)^{-1}(\{1\}) = 1$ and $\bar{\mathbf{x}}_n(1) = 1$. Since $\omega(\bar{\mathbf{x}}^n, \hat{\mathbf{x}}^n) \to \omega(\bar{\mathbf{x}}^n, \hat{\mathbf{x}}^n)$, then $\bar{\mathbf{x}}(1) = 1$, $\hat{\nu} \circ \hat{\mathbf{x}}(\bar{\mathbf{x}})^{-1}(\{1\}) = 1$ and $\hat{\mathbf{x}}_t(\bar{\mathbf{x}}) = 1$. Thus, $u_L(\bar{\mathbf{x}}, \hat{\nu} \circ \hat{\mathbf{x}}(\bar{\mathbf{x}})^{-1}) = 0$ as claimed. This concludes the proof.

References

Aumann, R. (1964): "Markets with a Continuum of Traders," *Econometrica*, 32, 39–50.

Aumann, R., and A. Brandenburger (1995): "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63, 1161–1180.

Barlo, M. (2003): "Essays in Game Theory," Ph.D. thesis, University of Minnesota.

Carmona, G. (2008): "Large Games with Countable Characteristics," *Journal of Mathematical Economics*, 44, 344–347.

——— (2009): "A Remark on the Measurability of Large Games," *Economic Theory*, 39, 41–44.

Carmona, G., M. Páscoa, and K. Podczeck (2008): "On the Existence of Limit Equilibria in Large Games," Universidade Nova de Lisboa and Universität Wien.

Carmona, G., and K. Podczeck (2009): "On the Existence of Pure Strategy Nash Equilibria in Large Games," *Journal of Economic Theory*, 144, 1300–1319.

——— (2011): "Approximation and Characterization of Nash Equilibria of Large Games," University of Cambridge and Universität Wien.

——— (2014): "Existence of Nash Equilibrium in Games with a Measure Space of Players and Discontinuous Payoff Functions," *Journal of Economic Theory*, 152, 130–178.

Chari, V. V., and P. J. Kehoe (1989): "Sustainable Plans," *Journal of Political Economy*, 98(4), 783–802.

Green, E. (1980): "Non-Cooperative Price Taking in Large Dynamic Markets," *Journal of Economic Theory*, 22, 155–182.

Harsanyi, J. (1973): "Oddness of the Number of Equilibrium Points: A New Proof," *International Journal of Game Theory*, 2, 235–250.

Hildenbrand, W. (1974): *Core and Equilibria of a Large Economy*. Princeton University Press, Princeton.

Khan, M., K. Rath, and Y. Sun (1997): "On the Existence of Pure Strategy Equilibria in Games with a Continuum of Players," *Journal of Economic Theory*, 76, 13–46.

Khan, M., and Y. Sun (1995): "The Marriage Lemma and Large Anonymous Games with Countable Actions," *Mathematical Proceedings of the Cambridge Philosophical Society*, 117, 385–387.

Kuhn, H., and et al. (1996): "The Work of John Nash in Game Theory," *Journal of Economic Theory*, 69, 153–185.

Levine, D., and W. Pesendorfer (1995): "When Are Agents Negligible?," *American Economic Review*, 85(5), 1160–1170.

Mas-Colell, A. (1983): "Walrasian Equilibria as Limits of Noncooperative Equilibria. Part I: Mixed Strategies," *Journal of Economic Theory*, 30, 153–170.

——— (1984): "On a Theorem by Schmeidler," *Journal of Mathematical Economics*, 13, 201–206.

Nash, J. (1950): "Non-Cooperative Games," Ph.D. thesis, Princeton University.

Novshek, W., and H. Sonnenschein (1983): "Walrasian Equilibria as Limits of Noncooperative Equilibria. Part II: Pure strategies," *Journal of Economic Theory*, 30, 171–187.

Podczeck, K. (2008): "On the Convexity and Compactness of the Integral of a Banach Space Valued Correspondence," *Journal of Mathematical Economics*, 44, 836–852.

——— (2009): "On Purification of Measure-valued Maps," *Econ. Theory*, 38, 399–418.

Postlewaite, A., and D. Schmeidler (1978): "Approximate Efficiency of Non-Walrasian Nash Equilibria," *Econometrica*, 46, 127–135.

Rudin, W. (1976): *Principles of Mathematical Analysis*. McGraw-Hill, Singapore.

Sabourian, H. (1990): "Anonymous Repeated Games with a Large Number of Players and Random Outcomes," *Journal of Economic Theory*, 51, 92–110.

Schmeidler, D. (1973): "Equilibrium Points of Non-atomic Games," *Journal of Statistical Physics*, 4, 295–300.

Selten, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25–55.

van Damme, E. (1991): *Perfection and Stability of Nash equilibrium*. Springer Verlag, Berlin.