

Strategic Behavior in Non-Atomic Games*

Mehmet Barlo

Guilherme Carmona[†]

Sabanci University

University of Cambridge

December, 2011

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. 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 and Jan Werner 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.

[†]Faculty of Economics and King's College, University of Cambridge and Nova School of Business and Economics. Address: University of Cambridge, Faculty of Economics, Sidgwick Avenue, Cambridge CB3 9DD, UK; phone: (44) 122333 5225; email: gmc42@cam.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 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.

Another interesting example is provided in the framework of the mass action interpretation of Nash equilibrium given in Nash (1950): 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. Therefore, 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 a Nash equilibrium of the k -replica game if and only if the vector consisting of the average choice across players of a given island is a Nash equilibrium of the original game. However, this equivalence

fails to hold in the limit case of a continuum of players in 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.

Such failures of (lower hemi) continuity of the equilibrium correspondence in non-atomic games are well documented in the literature. Indeed, as noted in Levine and Pesendorfer (1995) “equilibria can be radically different in a model with a finite number of agents than in a model with a continuum of agents”. We refer the reader to Carmona and Podczeck (2011) and the references therein for more on this subject.

In this paper we propose a refinement of Nash equilibrium in non-atomic games designed to alleviate that problem 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, because the definition of the refined equilibrium concept involves only non-atomic games, its analysis is relatively easier compared with that of limit points of equilibria of large finite games.

This study presents and analyzes the concept of *strategic equilibrium* (henceforth to be abbreviated by 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 (henceforth referred 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 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 finite; 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 two example above: 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. We prove that a strategy profile in the non-atomic version is a SE if and only if the vector of the average across players on the same island is a Nash equilibrium of the original normal form game.

Hence, in these 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.

Similar conclusions are reached in other applications we consider. We display that the notion of SE provides a sharp refinement on non-atomic games of voting with finitely many political parties (or candidates). Even though in these games any voting profile is a Nash equilibrium, we prove that that the concept of SE eliminates almost all implausible Nash equilibria: When players' payoff functions depend continuously on the distribution of seats parties get, the set of SE consists of strategy profiles under which almost every agent votes for his most favored political party. Moreover, abstention (by a strictly positive measure of players) is not observed in any of the SE.

The second application we provide is a symmetric Cournot oligopoly (i.e., all cost functions are alike). There we show that the set of strategic equilibria contains only

symmetric Nash equilibria, as it would happen when the set of players is finite. On the other hand, Nash equilibria are characterized by any strategy that yields zero profits. Technically, this example is of interest as it involves a uncountable actions space and displays the non-linearities in an agent's individual maximization problem in the perturbed game.

In the third application, we demonstrate that in the optimal taxation game of Levine and Pesendorfer (1995) the use of SE, instead of Nash equilibrium, makes sure that the first-best can be obtained even with non-atomic players. Indeed, using the concept of Nash equilibrium 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 Nash equilibrium in finite player versions of the same game, it is well known that the second-best, the Ramsey Equilibrium, is the best possible with the use of Nash 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 in a sequential strategic interaction.

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.

We have chosen to present the formal definitions in the context of applications in

Section 2 to ease the exposition. Section 3 describes the general framework of non-atomic games, and in Section 4 we define the concept of SE and prove that it exists and is a refinement of Nash equilibrium. Finally, Section 5 presents our non-atomic version of Nash's mass action game.

2 Applications

In this section, we present three sets of examples in which the concept of SE eliminates implausible Nash equilibrium outcomes that arise in non-atomic games. The first example concerns voting games, the second Cournot competition, and the third optimal taxation.

2.1 Proportional Voting

We present a non-atomic game of proportional voting in which each player has a preference ordering on the set of political parties represented by a cardinal utility function. Moreover, the payoff a player obtains is the weighted average of his utilities on parties with the weights equal the fractions of seats parties obtain.

First we show that any strategy profile is Nash. In particular, not voting at all and every fan of an extreme right (left) political party choosing an extreme left (right, respectively) political party are among Nash equilibria.

The concept of SE provides a sharp refinement: Under the assumption that there are at least two or more parties who are most favored for a strictly positive measure of agents, we prove that the unique SE is one where each player votes for his favorite political party.

The set of agents is given by $[0, 1]$, which is endowed with the Lebesgue measure λ , and the set of political parties by $M = \{1, \dots, \bar{m}\}$. The action set of player $t \in [0, 1]$ is given by $A = M \cup \{0\}$ where choosing 0 denotes not voting.

The seats in the parliament that a party receives after an election depends on the fraction of the population voting for it. Let $\mu = (\mu_0, \mu_1, \dots, \mu_{\bar{m}})$ be a probability

distribution on A with μ_0 representing the fraction of people that has abstained and μ_m denoting the fraction of people who has voted for party $m \in M$. Given such a distribution, each party receives a portion of the parliament equal to the fraction of votes it receives. Since only a fraction of $1 - \mu_0$ of the people has voted, then party $m \in M$ receives a fraction $\mu_m/(1 - \mu_0)$ of the parliament. Clearly, this formula only makes sense if $\mu_0 < 1$. If $\mu_0 = 1$, an alternative definition must be given and we assume that each party is assigned an equal share of the parliament. Since μ is a probability distribution, then $\mu_a \geq 0$ for all $a \in A$ and $\sum_{a \in A} \mu_a = 1$. Let $\mathcal{M}(A)$ denote the space of all probability distributions on A ; also, let $\mathcal{M}(M)$ denote the space of all probability distributions on M . The above rule governing how to split the parliament across parties defines the following function $\pi : \mathcal{M}(A) \rightarrow \mathcal{M}(M)$:

$$\pi(\mu) = \begin{cases} \left(\frac{\mu_1}{1-\mu_0}, \dots, \frac{\mu_m}{1-\mu_0} \right) & \text{if } \mu_0 < 1, \\ \left(\frac{1}{m}, \dots, \frac{1}{m} \right) & \text{otherwise,} \end{cases} \quad (1)$$

for all $\mu \in \mathcal{M}(A)$.

Naturally, we are interested in voting distributions arising from players' choices. A strategy is a measurable function $\mathbf{x} : [0, 1] \rightarrow A$, assigning an action to each player. Then, a strategy \mathbf{x} induces a distribution on A which will be used to measure the set of players playing each action. Indeed, percentage of abstention is equal to $\lambda(\{t \in [0, 1] : \mathbf{x}(t) = 0\}) = \lambda(\mathbf{x}^{-1}(0))$ (denoted by $\lambda \circ \mathbf{x}^{-1}(0)$); and the percentage of the population who voted for party $m \in M$ is equal to $\lambda(\{t \in [0, 1] : \mathbf{x}(t) = m\}) = \lambda(\mathbf{x}^{-1}(m))$ (similarly, denoted by $\lambda \circ \mathbf{x}^{-1}(m)$). It should be pointed out that since \mathbf{x} is measurable, then for all a in A we have that $\mathbf{x}^{-1}(a) = \{t \in [0, 1] : \mathbf{x}(t) = a\}$ is measurable. Hence, $\lambda \circ \mathbf{x}^{-1} = (\lambda \circ \mathbf{x}_0^{-1}, \dots, \lambda \circ \mathbf{x}_m^{-1}) \in \mathcal{M}(A)$. That is, the distribution on A induced by \mathbf{x} is simply the Lebesgue measure of the inverse image of \mathbf{x} . Therefore, (1) delivers the distribution of the seats in the parliament induced by \mathbf{x} . That is, we will restrict attention to $\pi(\lambda \circ \mathbf{x}^{-1}) \in \mathcal{M}(M)$. Finally, for notational purposes, let $\pi_{\mathbf{x}} = \pi(\lambda \circ \mathbf{x}^{-1})$; and, for all $a \in A$, $\pi_{\mathbf{x}}^a$ denotes the a th coordinate of $\pi_{\mathbf{x}}$.

Each player $t \in [0, 1]$ is assumed to have a preference ordering on M , characterized

by a (cardinal) utility function $v_t : M \rightarrow \mathbb{R}_{++}$, where v_t^m denotes the utility of player t when the political party m obtains all the seats in the parliament. We say player t strictly prefers m to m' if and only if $v_t^m > v_t^{m'}$. For simplicity we restrict attention to strict preferences, and thus, for all $m, m' \in M$, $v_t^m \neq v_t^{m'}$. Moreover, we let m_t^* be the favorite political party of agent t , i.e. $v_t^{m_t^*} > v_t^m$ for all $m \in M$. The return of player t under strategy profile \mathbf{x} then is

$$u(t, \mathbf{x}) = \sum_{m \in M} \pi_{\mathbf{x}}^m v_t^m.$$

We can define a function, U , assigning each player to a utility function, i.e. $U(t) = v_t$. Then, $U^{-1}(v_t)$ identifies the set of players who have the utility function v_t , and we say all players $\tau \in U^{-1}(v_t)$ are of the same *type* as player t , because $v_\tau = v_t$ for all $\tau \in U^{-1}(v_t)$. In order to abstract from non-fruitful technicalities, we assume the set of players' types is finite. Therefore, we can partition $[0, 1]$ into finite number of subsets $U_{j=1}^J$, thus, attention is restricted only to a finite number of payoff functions.

Because that the Lebesgue measure assigns measure zero to any one of the players, for any given \mathbf{x} no single player can affect the distribution $\pi_{\mathbf{x}}$. Thus, it follows that any strategy is a Nash equilibrium. In particular, the strategy defined by $\mathbf{x}(t) = 0$ for all $t \in [0, 1]$ is a Nash equilibrium in which abstention is a society wide phenomenon.

In order to overcome this unpleasant feature, we propose the concept of SE. Given any non-atomic game and $\varepsilon > 0$, we define the ε -perturbed game, in which each player believes that he alone has an ε impact on the distribution of actions resulting from a strategy (alternatively, each player believes that he alone is an atom with ε mass). Then, a strategy \mathbf{x} is a SE if there exists a sequence $\{\varepsilon_k, \mathbf{x}_k\}_{k=1}^\infty$ with $\varepsilon_k \searrow 0$ and \mathbf{x}_k a Nash equilibrium of ε_k -perturbed game, and $\lambda \circ (U, \mathbf{x}_k)^{-1}$ converges to $\lambda \circ (U, \mathbf{x})^{-1}$. Convergence of $\lambda \circ (U, \mathbf{x}_k)^{-1}$ to $\lambda \circ (U, \mathbf{x})^{-1}$ means that $\lambda(\{t : U(t) = v_j \text{ and } \mathbf{x}_k(t) = m\})$ converges to $\lambda(\{t : U(t) = v_j \text{ and } \mathbf{x}(t) = m\})$ for all $j = 1, \dots, J$ and $m = 0, \dots, M$.¹

¹This characterization of the convergence of $\lambda \circ (U, \mathbf{x}_k)^{-1}$ to $\lambda \circ (U, \mathbf{x})^{-1}$ holds only because both the set of types and the set of action are finite. See Section 3 for the general definition.

Given $\varepsilon > 0$, the ε -perturbed game is defined by modifying each player's payoff function. This is done by altering the distribution of actions that a strategy induces in the following way: Player t thinks that his choice has an ε impact on the distribution, which implies that he thinks that all the other players have an impact of only $(1 - \varepsilon)$ in total. This, in turn, implies that if he votes (i.e., chooses $m \in M$), then he believes that the fraction of voters is at least ε . In fact, he believes that the abstention is only $(1 - \varepsilon)\pi_{\mathbf{x}}^0$. If he does not vote, then the abstention is $\varepsilon + (1 - \varepsilon)\pi_{\mathbf{x}}^0$. Furthermore, if he votes for party m , then he believes that the fraction of the population voting for it is $\varepsilon + (1 - \varepsilon)\pi_{\mathbf{x}}^m$, while the fraction of party $m' \neq m$ is $(1 - \varepsilon)\pi_{\mathbf{x}}^{m'}$. Formally, the above distribution can be obtained as follows: Let δ_t be probability measure on $[0, 1]$ that assigns probability one to t . Then, it is clear that player t is computing the distribution induced by \mathbf{x} using the measure $\lambda_{\varepsilon,t} = \varepsilon\delta_t + (1 - \varepsilon)\lambda$. That is, player t believes that the distribution on A is $\lambda_{\varepsilon,t} \circ \mathbf{x}^{-1}$ instead of $\lambda \circ \mathbf{x}^{-1}$. Note that $\lambda_{\varepsilon,t} \circ \mathbf{x}^{-1} = \varepsilon\delta_a + (1 - \varepsilon)\lambda \circ \mathbf{x}^{-1}$, where $a = \mathbf{x}(t)$ and δ_a is the probability measure on A concentrated at a .

As a result of employing (1) with $\lambda_{\varepsilon,t} \circ \mathbf{x}^{-1}$ we obtain a distribution $\pi(\lambda_{\varepsilon,t} \circ \mathbf{x}^{-1})$ in $\mathcal{M}(M)$. We denote $\pi(\lambda_{\varepsilon,t} \circ \mathbf{x}^{-1})$ by $\pi_{\mathbf{x},\varepsilon,a}$, where $a = \mathbf{x}(t)$. Thus, for all $m \in M$, $\pi_{\mathbf{x},\varepsilon,a}^m$ denotes the fraction of total votes that party m gets when player t chooses $a \in A$ in the ε -perturbed game.

When player t chooses 0, the distribution on parties remains the same, i.e.

$$\pi_{\mathbf{x},\varepsilon,0} = \pi_{\mathbf{x}}. \quad (2)$$

This is because, when $\pi_{\mathbf{x}}^0 < 1$ and $\mathbf{x}(t) = 0$, player t contemplates the fraction of players (including himself) who do not vote to be given by $1 - (\varepsilon + (1 - \varepsilon)\pi_{\mathbf{x}}^0) = (1 - \varepsilon)(1 - \pi_{\mathbf{x}}^0)$. Moreover, because $\mathbf{x}(t) = 0$, player t thinks that the measure of players voting to party m is $(1 - \varepsilon)\pi_{\mathbf{x}}^m$. On the other hand, when $\pi_{\mathbf{x}}^0 = 1$, $\varepsilon + (1 - \varepsilon)\pi_{\mathbf{x}}^0 = 1$ showing that in the ε -perturbed game the measure of players not voting does not change.

In the ε -perturbed game player t thinks that voting to party m would result in

political parties obtain the following portions of the parliament:

$$\pi_{\mathbf{x},\varepsilon,m} = \left(\frac{(1-\varepsilon)\pi_{\mathbf{x}}^1}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0}, \dots, \frac{(1-\varepsilon)\pi_{\mathbf{x}}^{m-1}}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0}, \frac{\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^m}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0}, \frac{(1-\varepsilon)\pi_{\mathbf{x}}^{m+1}}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0}, \dots, \frac{(1-\varepsilon)\pi_{\mathbf{x}}^{\bar{m}}}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0} \right) \quad (3)$$

When player t 's choice is m , for $m' \neq m$, we can write $\pi_{\mathbf{x},\varepsilon,m}^{m'} = [(1-\varepsilon)\pi_{\mathbf{x}}^{m'}]/[1-\pi_{\mathbf{x}}^0 + \varepsilon\pi_{\mathbf{x}}^0]$, thus, $\pi_{\mathbf{x},\varepsilon,m}^{m'} < \pi_{\mathbf{x}}^{m'}$ if and only if $\pi_{\mathbf{x}}^{m'} > 0$; and, furthermore, $\pi_{\mathbf{x},\varepsilon,m}^{m'} = \pi_{\mathbf{x}}^{m'} = 0$ if $\pi_{\mathbf{x}}^{m'} = 0$. Similarly, $\pi_{\mathbf{x},\varepsilon,m}^m > \pi_{\mathbf{x}}^m$ if and only if $\pi_{\mathbf{x}}^m + \pi_{\mathbf{x}}^0 < 1$ or $\pi_{\mathbf{x}}^0 = 1$. Moreover, $\pi_{\mathbf{x},\varepsilon,m}^m = \pi_{\mathbf{x}}^m = 1$ if $\pi_{\mathbf{x}}^m + \pi_{\mathbf{x}}^0 = 1$ and $\pi_{\mathbf{x}}^0 < 1$.

Player t 's payoff is defined using the same expression as before, but with $\pi(\lambda_{\varepsilon,t} \circ \mathbf{x}^{-1})$ instead of $\pi_{\mathbf{x}}$: The payoff of player t choosing $a \in A$ in the ε -perturbed game for a given strategy \mathbf{x} is

$$u_{\varepsilon}^a(t, \mathbf{x}) = \sum_{m \in M} \pi_{\mathbf{x},\varepsilon,a}^m v_t^m.$$

For all $\varepsilon > 0$ and all strategies \mathbf{x} , by voting to his most favorite party m_t^* instead of choosing m , $m \neq m_t^*$, player t would strictly increase his expected utility, unless $\pi_{\mathbf{x}}^0 < 1$ and $\pi_{\mathbf{x}}^{m_t^*} + \pi_{\mathbf{x}}^0 = 1$. This is because, for all ε , \mathbf{x} and $\hat{m} \neq m_t^*$,

$$\begin{aligned} u_{\varepsilon}^{m_t^*}(t, \mathbf{x}) - u_{\varepsilon}^{\hat{m}}(t, \mathbf{x}) &= \frac{\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^*}}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0} v_t^{m_t^*} + \sum_{m' \neq m_t^*} \frac{(1-\varepsilon)\pi_{\mathbf{x}}^{m'}}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0} v_t^{m'} \\ &\quad - \frac{\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{\hat{m}}}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0} v_t^{\hat{m}} - \sum_{m' \neq \hat{m}} \frac{(1-\varepsilon)\pi_{\mathbf{x}}^{m'}}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0} v_t^{m'} \\ &= \frac{\varepsilon}{1-(1-\varepsilon)\pi_{\mathbf{x}}^0} (v_t^{m_t^*} - v_t^{\hat{m}}) > 0. \end{aligned}$$

Furthermore, player t would also increase his utility by voting to m_t^* instead of not voting provided that $\pi_{\mathbf{x}}^0 + \pi_{\mathbf{x}}^{m_t^*} < 1$ or $\pi_{\mathbf{x}}^0 = 1$. This is because in the first case for

every ε , \mathbf{x} satisfying either of these conditions, $u_\varepsilon^{m_t^*}(t, \mathbf{x}) - u_\varepsilon^0(t, \mathbf{x})$ is given by

$$\begin{aligned}
\sum_{m \in M} \left(\pi_{\mathbf{x}, \varepsilon, m_t^*}^m - \pi_{\mathbf{x}, \varepsilon, 0}^m \right) v_t^m &= \frac{\varepsilon}{1 - (1 - \varepsilon)\pi_{\mathbf{x}}^0} v_t^{m_t^*} \\
&\quad - \sum_{m \in M} \left(\frac{1}{1 - \pi_{\mathbf{x}}^0} - \frac{1 - \varepsilon}{1 - (1 - \varepsilon)\pi_{\mathbf{x}}^0} \right) \pi_{\mathbf{x}}^m v_t^m \\
&= \frac{\varepsilon}{1 - (1 - \varepsilon)\pi_{\mathbf{x}}^0} v_t^{m_t^*} \\
&\quad - \sum_{m \in M} \left(\frac{\varepsilon}{(1 - \pi_{\mathbf{x}}^0)(1 - (1 - \varepsilon)\pi_{\mathbf{x}}^0)} \right) \pi_{\mathbf{x}}^m v_t^m \\
&= \left(\frac{\varepsilon}{1 - (1 - \varepsilon)\pi_{\mathbf{x}}^0} \right) \left(v_t^{m_t^*} - \sum_{m \in M} \left(\frac{\pi_{\mathbf{x}}^m}{1 - \pi_{\mathbf{x}}^0} \right) v_t^m \right) > 0,
\end{aligned}$$

because $v_t^{m_t^*} > v_t^m$ for all $m \neq m_t^*$; The second case, i.e. when $\pi_{\mathbf{x}}^0 = 1$, $(u_\varepsilon^{m_t^*}(t, \mathbf{x}) - u_\varepsilon^0(t, \mathbf{x}))$ equals $v_t^{m_t^*} - \frac{1}{m} \sum_{m \in M} v_t^m$, clearly strictly positive. Finally, when $\pi_{\mathbf{x}}^0 < 1$ and $\pi_{\mathbf{x}}^{m_t^*} + \pi_{\mathbf{x}}^0 = 1$, then party m_t^* already has all the seats at the parliament, and so player t is indifferent between voting for m_t^* and not voting.

Thus, we have two cases: The first and interesting case happens when there is no $m \in M$ such that $\lambda(\{t \in [0, 1] : v_t^m > v_t^{m'} \text{ for all } m' \in M\}) = 1$. In this case, \mathbf{x} is a SE if and only if $\mathbf{x}(t) = m_t^*$ for almost all t , establishing that the unique SE profile is where almost every player t votes only for his favorite political party.

The second case happens when there exists $m \in M$ such that $\lambda(\{t \in [0, 1] : v_t^m > v_t^{m'} \text{ for all } m' \in M\}) = 1$. In this case, for all $\varepsilon > 0$, if \mathbf{x} is an ε -strategic equilibrium, then $\pi_{\mathbf{x}}^m = 1 - \pi_{\mathbf{x}}^0$ and $\pi_{\mathbf{x}}^0 < 1$. Hence, \mathbf{x} is a SE if and only if

$$\pi_{\mathbf{x}}^m = 1 - \pi_{\mathbf{x}}^0.$$

The assumption that no agent can be indifferent between two political parties is just to simplify the argument. If we were to allow indifference relations on the set of political parties by some (possibly all of the) players, the result would essentially be the same (provided that there are no two or more parties each of which is strictly preferred to all the others by almost every agent): The set of strategic equilibria will be strategy profiles in which every agent t would not vote for any of the political

parties $m \in M \setminus M(t)$, where $M(t) = \{\tilde{m} \in M : v_t(\tilde{m}) \geq v_t(m), \forall m \in M\}$. That is, in any strategic equilibrium agents choose one of their favorite political parties.

A slightly modified version of this game can be used in the analysis of allocating public resources on projects: A fixed amount of perfectly divisible public resources $B \in \mathbb{R}_{++}$, is to be allocated on M projects, all of which do not require any capital investments. Each player $t \in [0, 1]$ chooses an action in $A = M \cup \{0\}$, where 0 denotes not voting. Given \mathbf{x} , m gets $\pi_{\mathbf{x}}^m$ of B as defined above, and player t 's utility function is $u(t, \mathbf{x}) = \sum_{m \in M} (B\pi_{\mathbf{x}}^m)v_t^m$. Due to the above, unless all players strictly prefer the same project, the unique SE is one where almost all players choose their favorite project despite the fact that all strategy profiles are Nash equilibria.

2.2 Cournot Oligopoly

In this section we formulate and analyze a symmetric Cournot oligopoly with a continuum of players, and demonstrate that SE rules out all the non-symmetric Nash equilibria.

The set of agents is given by $[0, 1]$, endowed with the Lebesgue measure; and each of them can choose a quantity $\mathbf{x}(t) \in [0, \bar{q}]$, where $\bar{q} > 1$. The symmetric unit cost of production for each $t \in [0, 1]$ is 1. Given the quantity choices \mathbf{x} the inverse demand is given by $p = 2 - \int \mathbf{x} d\lambda$.

The profit function of firm t is

$$\Pi(t, \mathbf{x}) = \left(1 - \int \mathbf{x}\right) \mathbf{x}(t).$$

Let $U(t) = \Pi(t, \cdot)$ for all $t \in [0, 1]$ be the function assigning payoff functions to all players in the game.

The set of Nash equilibria in this game is any strategy profile \mathbf{x} satisfying $\int \mathbf{x} d\lambda = 1$. The reason is that as long as $\int \mathbf{x} d\lambda = 1$, then $p = 1$, thus, any player would be indifferent between any of their choices, since each player is atomless.

Given a profile \mathbf{x} and $\varepsilon > 0$, the profit of t in the ε -perturbed game is

$$\Pi_\varepsilon(t, \mathbf{x}) = \left(1 - (1 - \varepsilon) \int \mathbf{x} - \varepsilon \mathbf{x}(t)\right) \mathbf{x}(t).$$

Thus, the best response of t is

$$\mathbf{x}_\varepsilon(t) = \frac{1 - (1 - \varepsilon) \int \mathbf{x}}{2\varepsilon}.$$

In equilibrium,

$$\int \mathbf{x}_\varepsilon = \int \left(\frac{1 - (1 - \varepsilon) \int \mathbf{x}_\varepsilon}{2\varepsilon} \right) = \frac{1 - (1 - \varepsilon) \int \mathbf{x}_\varepsilon}{2\varepsilon}.$$

Thus, $\int \mathbf{x}_\varepsilon = \frac{1}{1+\varepsilon}$ which gives us (by substituting back to the best response function)

$$\mathbf{x}_\varepsilon(t) = \frac{1}{1 + \varepsilon}.$$

Letting $\mathbf{x}^*(t) = 1$ for every player $t \in [0, 1]$, it follows easily that $\lambda \circ (U, \mathbf{x}_\varepsilon)^{-1}$ converges to $\lambda \circ (U, \mathbf{x}^*)^{-1}$. Conversely, if $\tau = \lim_\varepsilon \lambda \circ (U, \mathbf{x}_\varepsilon)^{-1}$, then $\tau(\{(\Pi, 1)\}) = 1$ and so for all SE \mathbf{x} , we have $\mathbf{x}(t) = \mathbf{x}^*(t)$ for almost every t . Hence, the set of strategic equilibria consists of strategies \mathbf{x} such that $\mathbf{x}(t) = 1$ for almost every $t \in [0, 1]$. Thus, unlike for Nash equilibrium, there is a unique SE (up to a measure zero set of players).

2.3 Optimal Taxation

The strategic interaction analyzed in this section concerns the optimal taxation game, example 3, of Levine and Pesendorfer (1995). We will show that the use of SE, instead of Nash equilibrium, will make sure that the first-best can be obtained even with a continuum of non-atomic players. Moreover, this game is of additional interest as it involves the use of SE in conjunction with sequential rationality.

The strategic interaction between the government, the large player L , and large number of identical small players, where a representative individual is denoted by S , takes place over three periods, 0, 1, and 2. The government who can precommit in the initial period to a reaction (to the choices of the households) in the final period,

must choose whether to place a tax on capital or use a distortionary tax in order to raise adequate revenue. Households must choose an action after the precommitment, but before the actual move of the large player.

The set of households $S = [0, 1]$ and is endowed with the Lebesgue measure λ . Each household is endowed with 1 unit of capital which can be invested, $0 \leq x_S \leq 1$ to deliver $(1+r)x_S$, $r > 0$. We denote the set of actions of any one of the households by $A_S = [0, 1]$. The households care about the total amount of capital at the end of the game. The government has to collect some amount of resources which is strictly higher than $1+r$. In order to raise that amount of resources, the government may use one of the following two tax schemes. The first tax scheme consists of a *non-distortionary tax* on the investments (collecting all the investments plus its interests). In this situation the tax collected does not suffice to cover the needed resources of the government. Hence, the government incurs a loss due to the revenue shortfall in the amount of $p(1-x_S)$. The utility of household S would then be given by $(1-x_S)$, the amount of capital net of the investment. The government's utility is the utility of the household S minus the penalty resulting from the revenue shortfall: $(1-x_S)(1-p)$. The second tax scheme consists of a *distortionary tax* on some other resource in the economy (say, labor), which will cover the amount of needed resources for the government. Because that it is distortionary, each household incurs a cost of $c > 1+r$. The utility of the household at the end of the game consists of their endowment net of investment, i.e. $(1-x_S)$, plus the proceeds from their investment, i.e. $(1+r)x_S$, and finally minus the cost from the use of a distortionary tax, c . Hence, is equal to $(1+rx_S-c)$. As there are no revenue shortfalls, the government obtains a utility, equal to that of the households, i.e. $(1+rx_S-c)$.

The government implements the non-distortionary tax with a probability of x_L . The set of actions of the government is denoted by $A_L = [0, 1]$.

The government's payoff function, $u_L : A_S \times A_L \rightarrow \mathbb{R}$, is defined by

$$u_L(x_L, x_S) = (1-x_L)(1+rx_S-c) + x_L(1-x_S)(1-p)$$

where $c > 1 + r$ and $p > 1$. Each household's payoff function is $u_S : A_S \times A_L \rightarrow \mathbb{R}$

$$u_S(x_L, x_S) = (1 - x_L)(1 + rx_S - c) + x_L(1 - x_S).$$

Define $\underline{u}_S = \min_{x_L} \max_{x_S} u_S(x_S, x_L)$, denoting the payoff that S can guarantee for himself. In his best response, S chooses $x_S = 1$ if $x_L < r/(1 + r)$, anything in $[0, 1]$ if $x_L = r/(1 + r)$, and $x_S = 0$ otherwise. Thus, the minmax is obtained when $x_L = 0$ and $x_S = 1$, and delivers $\underline{u}_S = (1 + r - c) < 0$. Let, $u_L^* = \max_{x_S, x_L} u_L(x_S, x_L)$, subject to $u_S(x_S, x_L) \geq \underline{u}_S$. In this situation, $x_L = 1$ and $x_S = 1$ solves this problem, and renders $u_L^* = 0$; and, because that the utility of S would be 0, this arrangement is also individually rational.

Going into the non-atomic case, it is worthwhile to point out that in this game there is only one type of households. I.e. $\lambda \circ U^{-1}(u_S) = 1$. Let $\mathbf{x}_L : A_S \rightarrow A_L$ be a strategy of the government and $\mathbf{x}_S : [0, 1] \rightarrow A_S$ a strategy for the households. Let, for all $a, \alpha \in A_S$,

$$U_{t,S}(a, \mathbf{x}_L(\alpha)) = (1 - \mathbf{x}_L(\alpha))(1 + ra - c) + \mathbf{x}_L(\alpha)(1 - a).$$

Then, each household t 's payoff function is defined by $U_{t,S}(\mathbf{x}_S(t), \mathbf{x}_L(\int \mathbf{x}_S))$ and the government's payoff function is defined by

$$\begin{aligned} U_L\left(\mathbf{x}_L, \int \mathbf{x}_S\right) &= \left(1 - \mathbf{x}_L\left(\int \mathbf{x}_S\right)\right)\left(1 + r \int \mathbf{x}_S - c\right) \\ &\quad + \mathbf{x}_L\left(\int \mathbf{x}_S\right)\left(1 - \int \mathbf{x}_S\right)(1 - p), \end{aligned}$$

where $c > 1 + r$ and $p > 1$.

A pair $(\mathbf{x}_S, \mathbf{x}_L)$ is a *Stackelberg response* if

$$U_{t,S}\left(\mathbf{x}_S(t), \mathbf{x}_L\left(\int \mathbf{x}_S\right)\right) \geq U_{t,S}\left(a, \mathbf{x}_L\left(\int \mathbf{x}_S\right)\right)$$

for all $a \in A_S$ and almost all $t \in [0, 1]$. A pair $(\mathbf{x}_S, \mathbf{x}_L)$ is a *precommitment equilibrium* if it is a Stackelberg response and if

$$U_L\left(\mathbf{x}_L, \int \mathbf{x}_S\right) \geq U_L\left(\tilde{\mathbf{x}}_L, \int \tilde{\mathbf{x}}_S\right) \tag{4}$$

for all Stackelberg responses $(\tilde{\mathbf{x}}_S, \tilde{\mathbf{x}}_L)$.

Note that Assumptions 1 – 3 of Levine and Pesendorfer (1995) are satisfied. Thus, by Theorem 1 of Levine and Pesendorfer (1995), it is known that a precommitment equilibrium exists, and the unique amount received by the government is strictly lower than the first-best level 0.² In what follows, we show that there exists a strategic precommitment equilibrium and the unique amount received by the large player is $U_L^* = 0$.

Let $\varepsilon > 0$. In words, in the ε -perturbed game a.e. 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: Chose of 1 whenever the societal choice in the ε -perturbed game is 1; otherwise, the government “punishes” the small players by choosing 0. This, in turn, will make sure that the first-best can be obtained in equilibrium in the ε -perturbed game, $\varepsilon > 0$. And, it is the unique SE payoff because the government chooses first and the best responses of the households in the any ε -perturbed game is uniquely determined. This, then, clearly implies that the limit as ε tends to 0 (i.e. the strategic precommitment equilibrium) is one in which the unique SE amount received by the government is the first-best. These are formally presented below.

A pair $(\mathbf{x}_S, \mathbf{x}_L)$ is an ε -Stackelberg response if

$$U_{t,S} \left(\mathbf{x}_S(t), \mathbf{x}_L \left(\int \mathbf{x}_S \right) \right) \geq U_{t,S} \left(a, \mathbf{x}_L \left(\varepsilon a + (1 - \varepsilon) \int \mathbf{x}_S \right) \right)$$

for all $a \in A_S$ and a.e. $t \in [0, 1]$. A pair $(\mathbf{x}_S, \mathbf{x}_L)$ is an ε -precommitment equilibrium if it is an ε -Stackelberg response and if condition (4) holds for all ε -Stackelberg responses $(\tilde{\mathbf{x}}_S, \tilde{\mathbf{x}}_L)$. Finally, a pair (\mathbf{x}_S, x_L) is a *strategic precommitment equilibrium* distribution if there exists $\{\varepsilon_k\}_{k=1}^\infty$ such that $\varepsilon_k \rightarrow 0$, $(\mathbf{x}_S^k, \mathbf{x}_L^k)$ is an ε_k -precommitment equilibrium for all $k \in \mathbb{N}$, $\lambda \circ (U, \mathbf{x}_S^k)^{-1}$ converges to $\lambda \circ (U, \mathbf{x}_S)^{-1}$ and \mathbf{x}_L^k converges

²Clearly, the first best strategies cannot be Stackelberg responses. This is because, when $\mathbf{x}_L(1) = 1$, due to $\int \mathbf{x}_S \setminus_t x'_{t,S} = \int \mathbf{x}_S$, $x'_{t,S} \in A_S$, a.e. t 's best response is $x'_{t,S} = 0$, and not $\mathbf{x}_S(t) = 1$.

uniformly to \mathbf{x}_L .

Claim 1 *There exists a strategic precommitment equilibrium and the unique payoff received by the government is 0 (the first-best).*

Proof. Consider \mathbf{x}_L^* defined by

$$\mathbf{x}_L^*(z) = \begin{cases} 1 & \text{if } z = 1, \\ 0 & \text{if } z < 1 \end{cases}$$

and \mathbf{x}_S^* defined by $\mathbf{x}_S^*(t) = 1$ for all $t \in [0, 1]$. For all $\varepsilon > 0$, $\varepsilon a + (1 - \varepsilon) \int \mathbf{x}_S^* = 1$ if and only if $a = 1$. Thus,

$$U_{t,S} \left(a, \mathbf{x}_L^* \left(\varepsilon a + (1 - \varepsilon) \int \mathbf{x}_S^* \right) \right) = \begin{cases} 0 & \text{if } a = 1, \\ 1 + ra - c & \text{if } a < 1. \end{cases}$$

Because that $1 + ra - c < 0$ for all $a \leq 1$, we have that $(\mathbf{x}_S^*, \mathbf{x}_L^*)$ is a ε -Stackelberg response for all $\varepsilon > 0$. Since $U_L(\mathbf{x}_L^*, \int \mathbf{x}_S^*) = 0 \geq U_L(\mathbf{x}_L, \int \mathbf{x}_S)$ for all $(\mathbf{x}_L, \mathbf{x}_S)$, it follows that $(\mathbf{x}_S^*, \mathbf{x}_L^*)$ is a ε -precommitment equilibrium for all $\varepsilon > 0$. It is then clear that $(\mathbf{x}_S^*, \mathbf{x}_L^*)$ is a strategic precommitment equilibrium. This establishes existence of a strategic precommitment equilibrium.

We next show the uniqueness of the strategic precommitment equilibrium payoff for the government. Let $(\mathbf{x}_S, \mathbf{x}_L)$ be a strategic precommitment equilibrium and $(\mathbf{x}_L^k, \mathbf{x}_S^k)$ be a sequence of ε_k -precommitment equilibria satisfying the above conditions. Fix $k \in \mathbb{N}$ and note that if $U_L(\mathbf{x}_L^k, \int \mathbf{x}_S^k) < 0$, then $U_L(\mathbf{x}_L, \int \mathbf{x}_S) < U_L(\mathbf{x}_L^k, \int \mathbf{x}_S^k)$ and, therefore, $(\mathbf{x}_S^k, \mathbf{x}_L^k)$ is not an ε_k -precommitment equilibrium. Hence, $U_L(\mathbf{x}_L^k, \int \mathbf{x}_S^k) = 0$ for all $k \in \mathbb{N}$.

We have that $\int \mathbf{x}_S^k \rightarrow \int \mathbf{x}_S$ since $\lambda \circ (U, \mathbf{x}_S^k)^{-1}$ converges to $\lambda \circ (U, \mathbf{x}_S)^{-1}$. We also have that $\mathbf{x}_L^k(\int \mathbf{x}_S^k) \rightarrow \mathbf{x}_L(\int \mathbf{x}_S)$ since \mathbf{x}_L^k converges uniformly to \mathbf{x}_L . Thus,

$$\begin{aligned} U_L \left(\mathbf{x}_L, \int \mathbf{x}_S \right) &= \left(1 - \mathbf{x}_L \left(\int \mathbf{x}_S \right) \right) \left(1 + r \int \mathbf{x}_S - c \right) \\ &\quad + \mathbf{x}_L \left(\int \mathbf{x}_S \right) \left(1 - \int \mathbf{x}_S \right) (1 - p) \\ &= \lim_k U_L \left(\mathbf{x}_L^k, \int \mathbf{x}_S^k \right) = 0. \end{aligned}$$

■

3 Non-Atomic Games

In this section, we formally describe non-atomic games in which each player has a compact set of actions and a continuous payoff function that depends only on his choice and on the distribution of actions. The set of *players* is a probability space (T, Σ, λ) such that $\{t\} \in \Sigma$ for all $t \in T$. The set of *actions* is denoted by A , and we assume that it is a non-empty, compact metric space. By a distribution of actions we mean a Borel probability measure on A . Let $\mathcal{M}(A)$ be the space of Borel probability measures on A endowed with the topology of the weak convergence of probability measures. By Parthasarathy (1967, Theorem II.6.4), it follows that $\mathcal{M}(A)$ is a compact metric space. We write $\mu_n \Rightarrow \mu$ whenever $\{\mu_n\}_{n=1}^\infty \subseteq \mathcal{M}(A)$ converges to μ , which happens if $\int_A h d\mu_n$ converges to $\int_A h d\mu$ for all bounded, continuous functions $h : A \rightarrow \mathbb{R}$. Let ρ denote the Prohorov metric on $\mathcal{M}(A)$, which is known to metricize the weak convergence topology.

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. Formally, each player's payoff depends on his choice $a \in A$ and on a L - dimensional vector (τ_1, \dots, τ_L) , $L \in \mathbb{N}$, of distributions on A . Let $L \in \mathbb{N}$ and \mathcal{U} denote the space of real-valued continuous payoff functions defined on $A \times (\mathcal{M}(A))^L$. The set \mathcal{U} represents the space of *players' characteristics* or *types*. We endow it with the supremum norm, thus, making it a complete and separable metric space.

A *game with a continuum of players* is defined by assigning a payoff function to each player and defining a partition of the set of players into the relevant subgroups. Thus, it is characterized by a measurable function $U : T \rightarrow \mathcal{U}$ and a finite partition $\{T_i\}_{i=1}^L$ of T such that T_i is measurable and $\lambda(T_i) > 0$ for all $i = 1, \dots, L$. Each set T_i is interpreted as a group or an institution and is endowed with the following measure $\lambda_i = \lambda/\lambda(T_i)$. We represent such game by $G = (\{T_i\}_{i=1}^L, U, A)$.

For convenience, let $U_i : T_i \rightarrow \mathcal{U}$, $i = 1, \dots, L$, denote the restriction of U to

T_i . A strategy $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_L)$ is a vector of measurable functions $\mathbf{x}_i : T_i \rightarrow A$, $i = 1, \dots, L$. A pair $((U_i)_{i=1}^L, (\mathbf{x}_i)_{i=1}^L)$ induces a vector of probability measures on $\mathcal{U} \times A$ denoted by $(\lambda_1 \circ (U_1, \mathbf{x}_1)^{-1}, \dots, \lambda_L \circ (U_L, \mathbf{x}_L)^{-1}) \in (\mathcal{M}(\mathcal{U} \times A))^L$. The payoff of player $t \in T_i$ is

$$U_t(\mathbf{x}_i(t), \lambda_1 \circ \mathbf{x}_1^{-1}, \dots, \lambda_L \circ \mathbf{x}_L^{-1}).$$

Given a vector of Borel probability measures $(\tau_1, \dots, \tau_L) \in (\mathcal{M}(\mathcal{U} \times A))^L$, we denote by $\tau_{i\mathcal{U}}$ and $\tau_{i,A}$ the marginals of τ_i on \mathcal{U} and A respectively. The expression $u(a, \tau_{1,A}, \dots, \tau_{L,A}) \geq u(A, \tau_{1,A}, \dots, \tau_{L,A})$ means $u(a, \tau_{1,A}, \dots, \tau_{L,A}) \geq u(a', \tau_{1,A}, \dots, \tau_{L,A})$ for all $a' \in A$.

Given a game $G = (\{T_i, U_i\}_{i=1}^L, A)$, a vector of Borel probability measures $(\tau_1, \dots, \tau_L) \in \mathcal{M}(\mathcal{U} \times A)^L$ is an *equilibrium distribution for G* if for all $i = 1, \dots, L$,

1. $\tau_{i\mathcal{U}} = \lambda_i \circ U_i^{-1}$, and
2. $\tau_i(\{(u, a) \in \mathcal{U} \times A : u(a, \tau_{1,A}, \dots, \tau_{L,A}) \geq u(A, \tau_{1,A}, \dots, \tau_{L,A})\}) = 1$.

We will use the following notation: $B_\tau = \{(u, a) \in \mathcal{U} \times A : u(a, (\tau_{i,A})_i) \geq u(A, (\tau_{i,A})_i)\}$. Note that B_τ is closed, and so a Borel set; hence $\tau_i(B_\tau)$ is well defined. Also, if (u, a) belong to B_τ , then a maximizes the function $\tilde{a} \mapsto u(\tilde{a}, (\tau_{i,A})_i)$. Thus, we implicitly assume that no player can affect the distribution of actions, and in this sense the above describe a game with a continuum of players.

A strategy $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_L)$ is a *Nash equilibrium of G* if $U_t(\mathbf{x}(t), \lambda_1 \circ \mathbf{x}_1^{-1}, \dots, \lambda_L \circ \mathbf{x}_L^{-1}) \geq U_t(a, \lambda_1 \circ \mathbf{x}_1^{-1}, \dots, \lambda_L \circ \mathbf{x}_L^{-1})$ for almost all $t \in T$ and all $a \in A$. Nash equilibria exist if either A or $U(T)$ (or both) are countable, or if (T, Σ, λ) 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).³

³Formally, (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).

4 Strategic equilibria

4.1 Strategic Equilibrium Distributions

As was stressed in the introduction, we wish to consider those Nash equilibria that can be seen as a limit of equilibria in games in which each 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 , $\lambda_i(\{t\}) = 0$.

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

$$\lambda_{i,t,\varepsilon} \circ \mathbf{x}_i^{-1} = \varepsilon\delta_{\mathbf{x}_i(t)} + (1 - \varepsilon)\lambda_i \circ \mathbf{x}_i^{-1}. \quad (5)$$

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:

For all $\varepsilon > 0$, $t \in T_i$, $a \in A$ and $\tau = (\tau_j)_{j=1}^L \in (\mathcal{M}(A))^L$, define

$$U_{i,\varepsilon}(t)(a, \tau) = U_i(t)(a, (\varepsilon\delta_a + (1 - \varepsilon)\tau_i, \tau_{-i})).$$

We then define the ε -perturbed game G_ε of G as $G_\varepsilon = (\{T_i, U_{i,\varepsilon}\}_{i=1}^L, A)$. Note that the ε -perturbed game G_ε has the same players, and actions spaces as the original game G and that, for every strategy \mathbf{x} ,

$$U_{i,\varepsilon}(t)(\mathbf{x}) = U_i(\mathbf{x}(t), (\varepsilon\delta_{\mathbf{x}(t)} + (1 - \varepsilon)\lambda_i \circ \mathbf{x}_i^{-1}, (\lambda_j \circ \mathbf{x}_j^{-1})_{j \neq i})).$$

We say that a distribution $\tau^* = (\tau_i^*)_{i=1}^L \in \mathcal{M}(\mathcal{U} \times A)^L$ is a *strategic equilibrium distribution* of G if there exists a sequence $\{\varepsilon_k\}_{k=1}^\infty \subseteq (0, 1)$ decreasing to zero and a sequence $\{\tau_k^*\}_{k=1}^\infty$ converging to τ^* such that τ_k^* is an equilibrium distribution of G_{ε_k} , for every $k \in \mathbb{N}$.

Theorem 1 establishes the existence of SE distributions.

Theorem 1 *Every game with a continuum of players has a SE distribution.*

Proof. Let $\varepsilon > 0$ and $1 \leq i \leq L$. Note first that $U_{i,\varepsilon} : T_i \rightarrow \mathcal{U}$ is Borel measurable. To see this define, for all $a \in A$ and $\tau \in \mathcal{M}(A)^L$, $U_{i,\varepsilon}^{(a,\tau)}$ by $t \mapsto U_{i,\varepsilon}(t)(a, \tau)$ and $U_i^{(a,\tau)}$ by $t \mapsto U_i(t)(a, \tau)$. Since $U_{i,\varepsilon}^{(a,\tau)} = U_i^{(a,(\varepsilon\delta_a + (1-\varepsilon)\tau_i, \tau_{-i}))}$ and $U_i^{(a,(\varepsilon\delta_a + (1-\varepsilon)\tau_i, \tau_{-i}))}$ is measurable by Carmona (2009, Proposition 1), it follows that $U_{i,\varepsilon}^{(a,\tau)}$. Then, it follows again by Carmona (2009, Proposition 1) that $U_{i,\varepsilon}$ is measurable.

Next, we show that $U_{i,\varepsilon}(t)$ is continuous for all $t \in T_i$. In fact, if $a \in A$, $\tau \in \mathcal{M}(A)^L$, $\{a_k\}_{k=1}^\infty \subseteq A$ and $\{\tau_k\}_{k=1}^\infty \subseteq \mathcal{M}(A)^L$ are such that $\lim_k a_k = a$ and $\lim_k \tau_k = \tau$ then $\varepsilon\delta_{a_k} + (1-\varepsilon)\tau_i^k \Rightarrow \varepsilon\delta_a + (1-\varepsilon)\tau_i$ and the continuity of U_i implies that $\lim_k U_{i,\varepsilon}(t)(a_k, \tau_k) = \lim_k U_i(a_k, (\varepsilon\delta_{a_k} + (1-\varepsilon)\tau_i^k, \tau_{-i}^k)) = U_i(a, (\varepsilon\delta_a + (1-\varepsilon)\tau_i, \tau_{-i})) = U_{i,\varepsilon}(t)(a, \tau)$.

Since $U_{i,\varepsilon}$ is measurable and $U_{i,\varepsilon}(t)$ is continuous for all $i \in \{1, \dots, L\}$ and $t \in T_i$, it follows by (a straightforward generalization of) Theorem 1 in Mas-Colell (1984) that G_ε has an equilibrium distribution.

To finish the proof, we let τ_n^* be an equilibrium distribution of $G_{1/n}$. Since $\{U_{i,1/n}\}_n$ converges uniformly to U_i , then it follows that $\lambda_i \circ U_{i,1/n}^{-1}$ converges to $\lambda_i \circ U_i^{-1}$ for all $i = 1, \dots, L$.

For all $i \in \{1, \dots, L\}$, let $K_i = \{\lambda_i \circ U_i^{-1}, \lambda_i \circ U_{i,1}^{-1}, \lambda_i \circ U_{i,1/2}^{-1}, \dots\}$ and $C_i = \{\mu \in \mathcal{M}(\mathcal{U} \times A) : \mu_{i,\mathcal{U}} \in K_i\}$. It follows by Hildenbrand (1974, Theorems 32 and 33) that K_i is tight, and so again by Hildenbrand (1974, Theorems 34, and 35) implies that C_i is tight. Since $\{\tau_k^*\} \subseteq C_1 \times \dots \times C_L$, it follows by Hildenbrand (1974, Theorem 31) that it has a converging subsequence. Hence, its limit point is a SE distribution of G . ■

Next, we show that any SE distribution is an equilibrium distribution.

Theorem 2 *Every SE distribution is an equilibrium distribution.*

Proof. Let τ be a SE distribution, and let $\{\varepsilon_k\}$ and $\{\tau_k\}$ be such that $\varepsilon_k \in (0, 1)$, $\lim_k \varepsilon_k = 0$, τ_k converges to τ , and τ_k is an equilibrium distribution of G_{ε_k} , for all $k \in \mathbb{N}$.

Note that, for all $k \in \mathbb{N}$, $\{(u, a) \in \mathcal{U} \times A : u(a, \tau_A) \geq u(A, \tau_A)\}$ is closed and $\tau_k(\{(u, a) \in \mathcal{U} \times A : u(a, \tau_A^k) \geq u(A, \tau_A^k)\}) = 1$. Hence, $\text{supp}(\tau_k) \subseteq \{(u, a) \in \mathcal{U} \times A : u(a, \tau_A^k) \geq u(A, \tau_A^k)\}$.

We next show that $\text{supp}(\tau) \subseteq \{(u, a) \in \mathcal{U} \times A : u(a, \tau_A) \geq u(A, \tau_A)\}$, which implies that τ is an equilibrium distribution as desired.

Let $(u^*, a^*) \in \text{supp}(\tau)$. By Carmona and Podczeck (2009, Lemma 12), there exists a subsequence $\{\tau_{k_m}\}_m$ of $\{\tau_k\}_k$ and, for each $m \in \mathbb{N}$, $(u_m, a_m) \in \text{supp}(\tau_{k_m})$ such that $\lim_m (u_m, a_m) = (u^*, a^*)$. Hence, for all $m \in \mathbb{N}$ and $a' \in A$, $u_m(a_m, \tau_A^m) \geq u_m(a', \tau_A^m)$ and so $u^*(a^*, \tau_A) \geq u^*(a', \tau_A)$. Thus, $(u^*, a^*) \in \{(u, a) \in \mathcal{U} \times A : u(a, \tau_A) \geq u(A, \tau_A)\}$.

■

4.2 Strategic Equilibrium Strategies

We say that a strategy $\mathbf{x} = (\mathbf{x}_i)_{i=1}^L$ is a *SE strategy* of G if there exists a sequence $\{\varepsilon_k\}_{k=1}^\infty \subseteq (0, 1)$ decreasing to zero and a sequence $\{\mathbf{x}_k\}_{k=1}^\infty$ such that \mathbf{x}_k is a Nash equilibrium of G_{ε_k} for every $k \in \mathbb{N}$ and $\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$.

Proposition 1 *Let G be a non-atomic game. Then the following conditions are equivalent:*

- (a) \mathbf{x} is a SE.
- (b) there exists a sequence $\{\varepsilon_k\}_{k=1}^\infty \subseteq (0, 1)$ decreasing to zero and a sequence $\{\mathbf{x}_k\}_{k=1}^\infty$ such that \mathbf{x}_k is a Nash equilibrium of G_{ε_k} for every $k \in \mathbb{N}$ and $\lambda \circ (U, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$.

Furthermore, if either A is countable or $U(T)$ is countable or (T, Σ, λ) is super-atomless, then both (a) and (b) are equivalent to

(c) $\lambda \circ (U, \mathbf{x})^{-1}$ is a SE distribution.

Proof. The equivalence between (a) and (b) follows from the fact that $\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$ if and only if $\lambda \circ (U, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$. To see the latter equivalence, suppose first that $\lambda \circ (U, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$.

Let $\varepsilon > 0$ and $h : \mathcal{U} \times A \rightarrow \mathbb{R}$ be a bounded uniformly continuous real-valued function. We will show that there exists $K \in \mathbb{N}$ such that $|\int_{\mathcal{U} \times A} h d\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} - \int_{\mathcal{U} \times A} h d\lambda \circ (U, \mathbf{x})^{-1}| < \varepsilon$ 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 $u, u' \in \mathcal{U}$ and $a, a' \in A$ such that $\|u - u'\| < \delta$ and $d(a, a') < \delta$. Since $U_{\varepsilon_k}(t)$ converges uniformly to $U(t)$, then the function $f_k : T \rightarrow \mathbb{R}$ defined by $f_k(t) = \|U_{\varepsilon_k}(t) - U(t)\|$ for all $k \in \mathbb{N}$ and $t \in T$ converges pointwise to zero. Hence, by Ergorov's Theorem, there exists a measurable $F \subseteq T$ and $K' \in \mathbb{N}$ such that $\lambda(T \setminus F) < \eta$ and $|f_k(t)| \leq \delta/2$ for all $t \in F$ and $k \geq K'$.

Since $\lambda \circ (U, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$, there exists $K \in \mathbb{N}$ such that $K \geq K'$ such that $|\int_{\mathcal{U} \times A} h d\lambda \circ (U, \mathbf{x}_k)^{-1} - \int_{\mathcal{U} \times A} h d\lambda \circ (U, \mathbf{x})^{-1}| < \varepsilon/2$ for all $k \geq K$.

Hence, for all $k \geq K$,

$$\begin{aligned} & \left| \int_{\mathcal{U} \times A} h d\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} - \int_{\mathcal{U} \times A} h d\lambda \circ (U, \mathbf{x})^{-1} \right| \\ & \leq \left| \int_{\mathcal{U} \times A} h d\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} - \int_{\mathcal{U} \times A} h d\lambda \circ (U, \mathbf{x}_k)^{-1} \right| \\ & \quad + \left| \int_{\mathcal{U} \times A} h d\lambda \circ (U, \mathbf{x}_k)^{-1} - \int_{\mathcal{U} \times A} h d\lambda \circ (U, \mathbf{x})^{-1} \right| \\ & < \int_T |h(U_{\varepsilon_k}(t), \mathbf{x}_k(t)) - h(U(t), \mathbf{x}_k(t))| d\lambda(t) + \frac{\varepsilon}{2} \\ & = \int_{T \setminus F} |h(U_{\varepsilon_k}(t), \mathbf{x}_k(t)) - h(U(t), \mathbf{x}_k(t))| d\lambda(t) \\ & \quad + \int_F |h(U_{\varepsilon_k}(t), \mathbf{x}_k(t)) - h(U(t), \mathbf{x}_k(t))| d\lambda(t) + \frac{\varepsilon}{2} < 2B\eta + \eta + \frac{\varepsilon}{2} < \varepsilon. \end{aligned}$$

Note that a similar argument to the above show that $\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$ implies $\lambda \circ (U, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$.

We finally turn to the proof of the equivalence between (a) and (c). Suppose that \mathbf{x} is a SE strategy of G . Let $\{\varepsilon_k\}_k$ and $\{\mathbf{x}_k\}_k$ be such that $\varepsilon_k \searrow 0$, \mathbf{x}_k is a Nash equilibrium of G_{ε_k} and that $\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} \Rightarrow \lambda \circ (U, \mathbf{x})^{-1}$. Since $\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1}$ is an equilibrium distribution of G_{ε_k} , then $\lambda \circ (U, \mathbf{x})^{-1}$ is a SE distribution of G .

Conversely, let $\tau = \lambda \circ (U, \mathbf{x})^{-1}$ be a strategic equilibrium distribution. Let $\{\varepsilon_k\}_k$ and $\{\tau_k\}_k$ be such that $\varepsilon_k \searrow 0$, τ_k is an equilibrium distribution of G_{ε_k} and that $\tau_k \Rightarrow \tau$. Then, since A is countable or $U(T)$ is countable or (T, Σ, λ) is superatomless, 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}_k such that \mathbf{x}_k is a Nash equilibrium of G_{ε_k} and $\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} = \tau_k$ for all k . Since $\lambda \circ (U_{\varepsilon_k}, \mathbf{x}_k)^{-1} = \tau_k \Rightarrow \tau = \lambda \circ (U, \mathbf{x})^{-1}$, it follows that \mathbf{x} is a strategic equilibrium of G . ■

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.⁴

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

⁴For a formal discussion of these ideas, see Aumann and Brandenburger (1995) and Kuhn (1996).

constitutes a 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) Nash equilibria of the normal-form game are exactly the profiles of distributions over actions induced by the strategic equilibria 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 $N = \{1, \dots, n\}$ is the set of positions, $\Delta(A_i)$ 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 , where the set of players is $N_k = \{(i, j) : 1 \leq i \leq n, 1 \leq j \leq k\}$ and player (i, j) has $\Delta(A_i)$ as his action space. Under the assumption that all matchings are equally likely, the probability that an action $a \in A = A_1 \times \dots \times A_n$ is played when players are using a strategy $\sigma = (\sigma_{i,j})_{i \in N, j=1, \dots, k}$ is

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

Let $\bar{\sigma} = (\bar{\sigma}_1, \dots, \bar{\sigma}_n) \in \Delta(A_1) \times \dots \times \Delta(A_n)$ 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

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

It is then easy to see (after going over the proof of Theorem 3) that for any $k \in \mathbb{N}$, σ is a Nash equilibrium of G_k if and only if $\bar{\sigma}$ is a Nash equilibrium of Γ . In words, Nash equilibria of G_k are precisely those strategies under which the average behavior in all positions is part of the same 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: let $T_i = [0, 1]$ for all $i \in N$; a player $t \in T_i$ chooses an element of A_i . Let $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ be a strategy. A player is selected from each T_i according to the Lebesgue measure, and thus, the probability that the player selected from the i th group will play action $a_i \in A_i$ is $\lambda_i \circ \mathbf{x}_i^{-1}(a_i)$. We thus define $\bar{\mathbf{x}}_i = \lambda_i \circ \mathbf{x}_i^{-1}$ and

$$U_i(t)(a_i, \lambda_1 \circ \mathbf{x}_1^{-1}, \dots, \lambda_n \circ \mathbf{x}_n^{-1}) = v_i(\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n) = \sum_{a \in A} \left(\prod_{i' \in N} \bar{\mathbf{x}}_{i'}(a_{i'}) \right) v_i(a).^5$$

We denote by G the game $(T_i, A_i, U_i)_{i \in N}$.

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

Theorem 3 *A strategy profile $\mathbf{x}^* = (\mathbf{x}_1^*, \dots, \mathbf{x}_n^*)$ is a SE of G if and only if $(\bar{\mathbf{x}}_1^*, \dots, \bar{\mathbf{x}}_n^*)$ is a Nash equilibrium of Γ .*

Proof. (Sufficiency) Let $(\mathbf{x}_1^*, \dots, \mathbf{x}_n^*)$ be a strategy in G , and assume that $\bar{\mathbf{x}}^* = (\bar{\mathbf{x}}_1^*, \dots, \bar{\mathbf{x}}_n^*)$ is a Nash equilibrium of Γ . Let $i \in N$. We have that $v_i(\bar{\mathbf{x}}^*) \geq v_i(\sigma_i, \bar{\mathbf{x}}_{-i}^*)$ for all $\sigma_i \in \Delta(A_i)$. This implies, in particular, that $v_i(\bar{\mathbf{x}}^*) \geq v_i(\varepsilon \mathbf{x}(t) + (1 - \varepsilon)\bar{\mathbf{x}}_i^*, \bar{\mathbf{x}}_{-i}^*)$ for all $t \in T_i$, and $\varepsilon > 0$. Hence, $(\mathbf{x}_1^*, \dots, \mathbf{x}_n^*)$ is a Nash equilibrium of G_ε for all $\varepsilon > 0$, thus, a SE of G .

(Necessity) Let $(\mathbf{x}_1^*, \dots, \mathbf{x}_n^*)$ be a SE of G , and let $\bar{\mathbf{x}}^* = (\bar{\mathbf{x}}_1^*, \dots, \bar{\mathbf{x}}_n^*)$. 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 Nash equilibrium of Γ . Let $i \in N$ and $a_i \in A_i$. If a_i

⁵The above notation is appropriate in the following sense: represent A_i by the unit vectors $\{e_1^i, \dots, e_{|A_i|}^i\}$ in $\mathbb{R}^{|A_i|}$ and define $\hat{x}_i(t) = e_j^i$ if and only if $x_i(t) = a_j \in A_i$. Then, $\lambda_i \circ x^{-1} = \int_{T_i} \hat{x}_i d\lambda_i$. Hence, $\lambda_i \circ x^{-1}$ can, in fact, be understood as an average.

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 $\bar{\mathbf{x}}^\varepsilon := (\lambda_1 \circ (\mathbf{x}_1^\varepsilon)^{-1}, \dots, \lambda_n \circ (\mathbf{x}_n^\varepsilon)^{-1})$, and $(\mathbf{x}_1^\varepsilon, \dots, \mathbf{x}_n^\varepsilon)$ is a Nash equilibrium of G_ε , \mathbf{x}_i^ε converges in distribution to \mathbf{x}_i^* for all $i \in N$ and $\varepsilon \rightarrow 0$. Since

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

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

Theorem 3 provides a new interpretation of 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 Nash equilibrium can be understood as a “stable” average behavior in a large society. However, since every SE is a Nash equilibrium, 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).

References

- 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., 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.
- 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.
- 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. E. A. (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. (1984): “On a Theorem by Schmeidler,” *Journal of Mathematical Economics*, 13, 201–206.
- NASH, J. (1950): “Non-Cooperative Games,” Ph.D. thesis, Princeton University.
- PARTHASARATHY, K. (1967): *Probability Measures on Metric Spaces*. Academic Press, New York.

- PODCZECK, K. (2008): “On the Convexity and Compactness of the Integral of a Banach Space Valued Correspondence,” *Journal of Mathematical Economics*, 44, 836–852.
- SABOURIAN, H. (1990): “Anonymous Repeated Games with a Large Number of Players and Random Outcomes,” *Journal of Economic Theory*, 51, 92–110.
- SELTEN, R. (1975): “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory*, 4, 25–55.