Strategic Behavior in Non-Atomic Games*

Mehmet Barlo
Sabanci University

Guilherme Carmona
Universidade Nova de Lisboa

October, 2006

Preliminary Version: Please Do Not Quote

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 strategy is a *strategic* equilibrium if it is a limit point of a sequence of Nash equilibria 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 established that under the mass action interpretation of Nash [9], the concept of strategic equilibrium is the extension of Nash equilibrium in finite normal form games, to non-atomic games by also proving the following: Given any finite normal form game, we consider its symmetric replicas and non-atomic version and show that every strategic equilibrium can be approached by a sequence Nash equilibria of replica games, and every limit point of such a sequence constitutes a strategic equilibrium. Finally, this notion is applied to non-atomic games of voting and it is displayed that if the perturbed game is such that the societal choice does not involve discontinuities of high magnitudes, then each player choosing his favorite candidate is the unique strategic equilibrium.

*We are grateful to Narayana Kocherlakota, Andy McLennan, Aldo Rusticini and Jan Werner for their valuable advice and support. We thank Pedro Amaral, Kemal Badur, Han Ozsoylev, and especially David Schmeidler for helpful comments and suggestions. We benefited from discussions in the Mathematical Economics Workshop at the University of Minnesota. All remaining errors are ours.

1 Introduction

Typically economic situations featuring a large number of agents are modeled with non-atomic games. While their advantageous feature of players' inability to affect societal variables provides significant technical ease, it also 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 a compact set of integers (with more than 1 element), and their payoffs depend only on the average choice. Because that a player's action does not affect the societal 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.

In this study we propose the concept of *strategic equilibrium* for non-atomic games in which the payoff of each agent depends on what he chooses and a societal choice. For any non-atomic game and $\varepsilon > 0$, we define an ε–perturbed game by requiring each player to think that only he alone has an ε impact on the societal choice. ε–*strategic* equilibrium is a Nash equilibrium of the ε–perturbed game. Thus, the test that a strategy has to pass to be labeled as ε–strategic involves the following question: "Would I play my part of the proposed strategy if only I, and no others, were to have an ε impact on the societal choice?" Unless there are no positive fractions of players answering "no", then such a distribution is ε–strategic.[1] Finally, the set of strategic equilibria will be the set of limit points of a sequence of ε–strategic equilibria as ε goes to 0. After proving the existence of strategic equilibria under standard assumptions (e.g., Schmeidler [13], and Rath [12]), we show that strategic equilibrium is a refinement of Nash equilibrium.

In order to relate strategic equilibria with the limits of Nash equilibria of large games, we employ the *mass action interpretation* presented in Nash [9]. We consider any finite normal form game, and formulate its associated non-atomic game by symmetrically replicating players. Due to the mass action interpretation supplied by Nash in his Ph.D. thesis, it is known that for any $K \in \mathbb{N}$, the set of Nash equilibria of the K-replica game coincide with the Nash equilibria of the original one. We then prove that a strategy profile in the non-atomic version of the given finite normal form game is a strategic equilibrium if and only if the associated strategy profile in the finite normal form game is a Nash equilibrium. Thus, every strategic equilibrium can be approached by a sequence Nash equilibria of replicated games, and every limit point of such a sequence constitutes a strategic equilibrium. This, in turn, confirms that under the mass action interpretation the strategic equilibrium is the genuine extension of Nash equilibrium in finite normal form games to non-atomic games.

We apply the notion of strategic equilibrium to non-atomic games of voting with finitely many political parties (or candidates). In these games any voting profile is a Nash equilibrium, but we

[1] It needs to be pointed out that in the ε–perturbed game, agents are not rational, just like in Selten [14]. This is because an agent thinks that he alone has an ε impact on the societal choice, and does not contemplate that others do the same consideration.

show that the concept of strategic equilibrium performs a sharp refinement: When the societal choice in the ε–perturbed game is non-constant and features only small discontinuities in players' choices, then is a strategy profile under which every agent votes for his most favored candidate is strategic, even when obtaining majority counts.

The second application we provide is done in Cournot oligopoly setup. We formulate the non-atomic Cournot oligopoly, and show that the set of strategic equilibria contains only symmetric Nash equilibria. 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 order to facilitate an easier reading, we have chosen to present the applications in Section 2. The following section gives the formal definitions of and the assumptions on non-atomic games, and in Section 4 we define the concept of strategic equilibrium and prove that it exists and is a refinement of Nash equilibrium. Section 5 studies the strategic interaction under strategic equilibrium, and section 6 concludes.

2 Applications

In this section, we present 2 sets of examples in which the concept of strategic equilibrium eliminates implausible Nash equilibrium strategy profiles. The first types of examples concerns voting games, and the second Cournot competition.

2.1 Voting Games

We present 2 non-atomic voting games, a proportional and a presidential game of voting. In these games we will show that any strategy profile is Nash. In particular, not voting at all is shown to be an artifact of the specification of a non-atomic game (and is not due to strategic interaction among players).

In the first example the unique strategic equilibrium is one where each player votes for his favorite political party, even when a strategy profile in which nobody votes or in which an agent favoring an extreme right (left) political party votes for an extreme left (right, respectively) political party is Nash. In the second example such a strong conclusion still holds when the discontinuities in players' payoff functions are not too big. Nevertheless, in both of the examples not voting is not observed in any strategic equilibrium.

2.1.1 Proportional Voting

The set of agents is given by $[0,1]$, and the set of political parties by $M = \{1, ..., \bar{m}\}$. The action set of player $t \in [0,1]$ is given by $A = M \cup \{0\}$ where choosing 0 denotes not voting. Given a strategy profile $\mathbf{x} = (\mathbf{x}(t))_{t\in[0,1]}$ with $\mathbf{x}(t) \in A$ for all $t \in [0,1]$, political party $m \in M$ gets $\pi^m_{\mathbf{x}}$ portion of the

seats in the parliament defined by

$$\pi_{\mathbf{x}}^m = \begin{cases} \frac{\lambda(\{\tau:\mathbf{x}(\tau)=m\})}{1-\lambda^0} & \text{if } \lambda^0 < 1, \\ \frac{1}{\bar{m}} & \text{otherwise,} \end{cases}$$

where λ is the Lebesque measure defined on $[0,1]$, and $\lambda^0 = \lambda(\{\tau : \mathbf{x}(\tau) = 0\})$. That is, as long as the measure of players voting is strictly below 1, party m obtains seats relative to its performance. But, when almost all players choose not to vote, then each party gets equal shares in the parliament.

For simplicity of the argument we assume that player t has a strict preference ordering on M, characterized by a cardinal utility function $v_t : M \to \mathbb{R}_{++}$, 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 $v_t^m > v_t^{m'}$, and due to our assumption for all $m, m' \in M$, $v_t^m \neq v_t^{m'}$. Moreover, we let $m_t^\star$ be the favorite political party of agent t, i.e. $v_t^{m_t^\star} > v_t^m$ for all $m \in M$. The return of player t under strategy profile $\mathbf{x}$ then is

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

Given $\mathbf{x}$ an agent cannot affect the societal choice $(\pi_{\mathbf{x}}^m)_{m \in M}$, thus, any strategy profile is Nash. In particular, $\mathbf{x}(t) = 0$ for all t is a Nash equilibrium in which not voting is a society wide phenomenon.

At this stage it is useful to provide a detailed, yet non-technical, definition of the concept of strategic equilibrium. Given any non-atomic game and $\varepsilon > 0$, we define an ε–perturbed game, in which each player believes that he alone has an ε impact on the societal choice (alternatively, each player believes that he alone is an atom with ε mass). An ε–strategic equilibrium is a Nash equilibrium of the ε–perturbed game, and strategic equilibria belong to the set of limit points of a sequence of ε–strategic equilibria where ε goes to 0.

Moreover, as discussed in the introduction, the preferences of players need to be cardinal with the use of strategic equilibrium, even if attention were to be restricted to pure strategies. This is because, when considering the ε–perturbed game, $\varepsilon > 0$, we need to extend the given utility function of any player to its ε–perturbed one by only assigning ε mass to the associated player.

Given $\varepsilon > 0$ and $\mathbf{x}$, in the ε–perturbed game player t thinks that voting for political party m would make it obtain $\pi_{\mathbf{x}}^{m,\varepsilon}$ portion of the parliament where $\pi_{\mathbf{x}}^{m,\varepsilon}$ is

$$\pi_{\mathbf{x}}^{m,\varepsilon} = \begin{cases} \varepsilon + \pi_{\mathbf{x}}^m & \text{if } \lambda^0 < 1, \\ 1 & \text{otherwise.} \end{cases} \tag{1}$$

Consequently, when player t chooses political party m his return is

$$u_\varepsilon^m(t, \mathbf{x}) = \pi_{\mathbf{x}}^{m,\varepsilon} v_t^m + \sum_{m' \in M \setminus \{m\}} \pi_{\mathbf{x}}^{m'} v_t^{m'}.$$

For any $\varepsilon > 0$ by voting to his most favorite party $m_t^\star$ instead of choosing m, $m \neq m_t^\star$ or not voting at all, agent t would increase his expected utility by $\varepsilon(v_t^{m_t^\star} - v_t^m) > 0$ provided that $\lambda^0 < 1$. On the other hand, when almost all agents have not voted, the net gains of player t choosing $m_t^\star$

is $v_t^{m_t^\star} - \frac{1}{\bar{m}} \sum_{m \in M} v_t^m$ which clearly is strictly positive. Thus, for any $\varepsilon > 0$ agent t's best response is to vote for $m_t^\star$, establishing that the unique strategic equilibrium profile is where almost every player t votes only for his favorite political party.[2]

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, and would read as follows: 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.

2.1.2 Presidential Voting

The use of strategic equilibrium delivered a strong conclusion in the examples above. Yet, this conclusion cannot be easily generalized to other voting situations. The reason is that for the concept of strategic equilibrium to perform with its full power, in the ε–perturbed game the societal choice has to be affected in a non-constant and continuous fashion with a player's action. When there are discontinuities of that sort, strategic equilibrium still may produce a sharp refinement on Nash equilibria, but in general this is not guaranteed. Indeed, below we show that when attention is restricted to voting situations in which obtaining majority provides an excess return for a candidate or a political party (thus, a discontinuity as discussed above is present), the performance of the concept of strategic equilibrium will crucially depend on the magnitude of such discontinuities.

For the simplicity of the argument, assume that there are 2 candidates (or political parties), i.e. $M = \{1, 2\}$, and the action set of each player is $A = \{0, 1, 2\}$, where choosing 0 stands for not voting. Given a strategy $\mathbf{x}$, candidate (or political party) m gets $\pi_{\mathbf{x}}^m$ shares of the total vote where $\pi_{\mathbf{x}}^m$ is defined as above. We maintain the simplifying assumption that for all t, $v_t^1 \neq v_t^2$. The important modification to the game presented in the previous section is about the utility of players. Given $\mathbf{x}$, the utility of player t is

$$u(t, \mathbf{x}) = \sum_{m \in M} (\pi_{\mathbf{x}}^m v_t^m)\, \chi_{t,m}, \text{ where } \chi_{t,m} = \begin{cases} 1 & \text{if } \pi_{\mathbf{x}}^m \geq \pi_{\mathbf{x}}^{m'}, \\ \beta_t^m & \text{otherwise,} \end{cases}$$

with $\beta_t^m \in [0, 1)$ for all $t \in [0, 1]$.

The effect of a candidate (or political party) not being able to obtain the majority on a player's returns is modeled with the term β_t^m; a measure of the effect of the losing candidate on player t's utility. To see more consider the following example: Suppose that candidate 1 obtained %40, and candidate 2 %60 of the total votes, and player t's favorite is candidate 1. In this case the utility

[2] 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 $u(t, \mathbf{x}) = \sum_{m \in M} (B \pi_{\mathbf{x}}^m) v_t^m$. It is an easy exercise to show that the unique strategic equilibrium is one where almost all players choose their favorite project, even when all strategy profiles are Nash.

of player 1 would be $0.60v_t^2 + \beta_t^1 0.40v_t^1$. Letting $\beta_t^1 = 0$ depicts a situation where player t does not get any payoffs from his favorite candidate getting %40 of total votes. Note that depending on the voting situation under analysis, levels of β_t^m should have been determined, and thus, we treat them as primite variables of the model. For example, in a presidential voting, in which the losing candidate must end his political career, it may be natural to assume that $\beta_t^m = 0$ for all t, m. But in a voting situation such as the congressional elections $\beta_t^m > 0$ for some t and m will capture the situation when player t gets a return from party m being in the parliament even though it is in minority. Another important aspect that needs to be emphasized is that with $\beta_t^m = 1$ for all t, m, this model will be the identical to a 2 party version of the game in the previous section, and $\beta_t^m < 1$ for some t, m constitutes the discontinuity discussed in the first paragraph of this section. Finally, please notice that the only requirement we have imposed is that β_t^m is independent of $\pi_{\mathbf{x}}^m$.

It is easy to see that all strategy profiles are Nash, because none of the agents have any ability to affect the societal choice. In particular, all players choosing not to vote, i.e. $\mathbf{x}$ with $\mathbf{x}(\tau) = 0$ for all $\tau \in [0, 1]$, is a Nash equilibrium.

The below renders an analysis of the strategic equilibria: For any given $\varepsilon > 0$ and $\mathbf{x}$, in the ε–perturbed game player t voting for candidate m makes him think that m gets $\pi_{\mathbf{x}}^{m,\varepsilon}$ shares of total votes, where $\pi_{\mathbf{x}}^{m,\varepsilon}$ is as defined in equation 1. Consequently, when player t votes for party $m \in \{1, 2\}$ he obtains the following return:

$$u_\varepsilon^m(\mathbf{x}, t) = \begin{cases} v_t^m & \text{if } \lambda^0 = 1; \\ (\varepsilon + (1-\varepsilon)\pi_m^{\mathbf{x}})\, v_t^m + (1-\varepsilon)\pi_{\mathbf{x}}^{m'} v_t^{m'} \beta_t^{m'} & \text{if } \lambda^0 < 1, \text{ and } \pi_{\mathbf{x}}^m + \varepsilon > \pi_{\mathbf{x}}^{m'}, m \neq m'; \\ (\varepsilon + (1-\varepsilon)\pi_m^{\mathbf{x}})\, v_t^m + (1-\varepsilon)\pi_{\mathbf{x}}^{m'} v_t^{m'} & \text{if } \lambda^0 < 1, \text{ and } \pi_{\mathbf{x}}^m + \varepsilon = \pi_{\mathbf{x}}^{m'}, m \neq m'; \\ (\varepsilon + (1-\varepsilon)\pi_m^{\mathbf{x}})\, v_t^m \beta_t^m + (1-\varepsilon)\pi_{\mathbf{x}}^{m'} v_t^{m'} & \text{if } \lambda^0 < 1, \text{ and } \pi_{\mathbf{x}}^m + \varepsilon < \pi_{\mathbf{x}}^{m'}, m \neq m'. \end{cases}$$

Not voting delivers player t, $u_\varepsilon^0(\mathbf{x}, t) = u(\mathbf{x}, t)$.

The following Proposition provides a partial characterization for the set of strategic equilibria (with attention restricted to pure strategies):

Proposition 1 *Assume that $\beta_t^{m'} \leq \beta_t^{m_t^\star}$ for all t and $m' \neq m_t^\star$. Then, a pure strategy profile $\mathbf{x}^*$ is a strategic equilibrium of the presidential voting game if and only if:*

1. *$\mathbf{x}^*(t) = m_t^\star$, for almost every $t \in \{\tau : \beta_\tau^{m_\tau^\star} v_\tau^{m_\tau^\star} > v_\tau^{m'}, m_\tau^\star \neq m'\}$; and,*

2. *$\mathbf{x}^*(t) = m'$, $m_\tau^\star \neq m'$, for almost every $t \in \{\tau : \beta_\tau^{m_\tau^\star} v_\tau^{m_\tau^\star} < v_\tau^{m'}, m_\tau^\star \neq m'\}$, whenever $\lambda^0 < 1$ and $\pi_{m_t^\star}^{\mathbf{x}^*} < \pi_{m'}^{\mathbf{x}^*}$; and,*

3. *$\mathbf{x}^*(t) \in \{m_t^\star, m', 0\}$, $m_\tau^\star \neq m'$, for almost every $t \in \{\tau : \beta_\tau^{m_\tau^\star} v_\tau^{m_\tau^\star} = v_\tau^{m'}, m_\tau^\star \neq m'\}$, whenever $\lambda^0 < 1$ and $\pi_{m_t^\star}^{\mathbf{x}^*} < \pi_{m'}^{\mathbf{x}^*}$.*

Under the assumption of $\beta_t^{m'} \leq \beta_t^{m_t^\star}$ for all t and $m' \neq m_t^\star$, Proposition 1 establishes that a pure strategy profile $\mathbf{x}^*$ is a strategic equilibrium if and only if it satisfies all of the following 3 conditions: $\mathbf{x}^*(t)$ is is equal to $m_t^\star$ whenever $\beta_t^{m_t^\star} v_t^{m_t^\star} > v_t^{m'}$; and to m', where $m_t^\star \neq m'$, only if $\beta_t^{m_t^\star} v_t^{m_t^\star} < v_t^{m'}$, and the measure of players who voted is strictly positive; and finally, is in $\{m_t^\star, m', 0\}$, $m_t^\star \neq m'$ only

when $\beta_t^{m_t^\star} v_t^{m_t^\star} = v_t^{m'}$, and the measure of players choosing to vote is strictly positive. Therefore, in general we may let $\beta_t^m = 0$ for all t and $m \neq m_t^\star$, thus, only $\beta_t^{m_t^\star}$ matters.

The requirement of $\beta_t^{m'} \leq \beta_t^{m_t^\star}$ for all t and $m' \neq m_t^\star$, restricts attention to cases where for every player t his favorite candidate (political party) when in the minority, provides at least as high returns as the those that the other candidate would have. Moreover, the details of the below proof (in which we work without this assumption) discloses that this condition is relevant only when the two parties get exactly the same portion of votes. Therefore, we believe that this assumption does not eliminate any interesting cases.

The following result is an immediate consequence of Proposition 1, and displays that generically speaking in any strategic equilibrium each player votes:

Corollary 1 *Assume that $\beta_t^{m'} \leq \beta_t^{m_t^\star}$ for all t and $m' \neq m_t^\star$. Let $\mathbf{x}^*$ be a strategic equilibrium in pure strategies. If $\lambda(\{\tau : \beta_\tau^{m_\tau^\star} v_\tau^{m_\tau^\star} = v_\tau^{m'}, m_\tau^\star \neq m'\}) = 0$, then for almost all t, $\mathbf{x}^*(t) \neq 0$. Thus generically speaking, not voting does not happen in any strategic equilibrium (for a strictly positive measure of players).*

The proof of Proposition 1 follows from the next Lemma which provides the pure strategy best response of player t in the ε–perturbed game, ε arbitrarily small, without the assumption of $\beta_t^{m'} \leq \beta_t^{m_t^\star}$ for all t and $m' \neq m_t^\star$:

Lemma 1 *Let $\beta_t^m \in [0,1)$ (without any further requirements). For arbitrarily small $\varepsilon > 0$, the pure strategy best response correspondence of player t in the ε–perturbed game is: Let $m' \neq m_t^\star$, then*

$$BR_\varepsilon(\mathbf{x}, t) = \begin{cases} m_t^\star & \text{if } \lambda^0 = 1; \\ & \quad \text{or if } \lambda^0 < 1 \text{ and } \pi_{\mathbf{x}}^{m_t^\star} > \pi_{\mathbf{x}}^{m'}; \\ & \quad \text{or if } \lambda^0 < 1 \text{ and } \pi_{\mathbf{x}}^{m_t^\star} = \pi_{\mathbf{x}}^{m'} \text{ and } \beta_t^{m_t^\star} \geq \beta_t^{m'}; \\ & \quad \text{or if } \lambda^0 < 1 \text{ and } \pi_{\mathbf{x}}^{m_t^\star} = \pi_{\mathbf{x}}^{m'} \text{ and } \beta_t^{m_t^\star} < \beta_t^{m'} \\ & \qquad \text{and } (v_t^{m_t^\star} - v_t^{m'}) \geq (\beta_t^{m_t^\star} v_t^{m_t^\star} - \beta_t^{m'} v_t^{m'}); \\ & \quad \text{or if } \lambda^0 < 1 \text{ and } \pi_{\mathbf{x}}^{m_t^\star} < \pi_{\mathbf{x}}^{m'} \text{ and } \beta_t^{m_t^\star} v_t^{m_t^\star} > v_t^{m'}; \\ \{m_t^\star, m', 0\} & \text{if } \lambda(\{\tau : \mathbf{x}(\tau) = 0\}) < 1 \text{ and } \pi_{\mathbf{x}}^{m_t^\star} < \pi_{\mathbf{x}}^{m'} \text{ and } \beta_t^{m_t^\star} v_t^{m_t^\star} = v_t^{m'}; \\ m' & \text{otherwise.} \end{cases}$$

Proof of Lemma 1. It is clear that player t chooses $m_t^\star$ if $\lambda^0 = 1$, because $v_t^{m_t^\star} - 1/2(v_t^{m_t^\star} + v_t^{m'}) > 0$, $m' \neq m_t^\star$.

When $\lambda^0 < 1$ and $\pi_{\mathbf{x}}^{m_t^\star} > \pi_{\mathbf{x}}^{m'}$, there exists $\varepsilon > 0$ small enough so that no player can change the winning candidate in the ε–perturbed game. In this case the optimal choice of player t is $m_t^\star$, which is due to $(u_\varepsilon^{m_t^\star}(t, \mathbf{x}) - u_\varepsilon^{m'}(t, \mathbf{x}))$, $m_t^\star \neq m'$, being equal to

$$\left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\right) v_t^{m_t^\star} + (1-\varepsilon)\pi_{\mathbf{x}}^{m'} v_t^{m'} \beta_t^{m'}\right) - \left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m'}\right) \beta_t^{m'} v_t^{m'} + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star} v_t^{m_t^\star}\right)$$

which is equal to $\varepsilon(v_t^{m_t^\star} - \beta_t^{m'} v_t^{m'}) > 0$; and $(u_\varepsilon^{m_t^\star}(t, \mathbf{x}) - u_\varepsilon^0(t, \mathbf{x}))$, being equal to

$$\left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\right) v_t^{m_t^\star} + (1-\varepsilon)\pi_{\mathbf{x}}^{m'} v_t^{m'} \beta_t^{m'}\right) - \left(\pi_{\mathbf{x}}^{m_t^\star} v_t^{m_t^\star} + \pi_{\mathbf{x}}^{m'} v_t^{m'} \beta_t^{m'}\right)$$

equaling $\varepsilon\pi_{\mathbf{x}}^{m'}(v_t^{m_t^\star} - \beta_t^{m'}v_t^{m'}) > 0$, because $(1-\pi_{\mathbf{x}}^{m_t^\star}) = \pi_{\mathbf{x}}^{m'}$.

In the case of $\lambda^0 < 1$ and $\pi_{m_t^\star}^{\mathbf{x}} = \pi_{\mathbf{x}}^{m'}$, $m_t^\star \neq m'$, $(u_\varepsilon^{m_t^\star}(t,\mathbf{x}) - u_\varepsilon^{m'}(t,\mathbf{x}))$ is equal to

$$\begin{aligned}&\left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\right)v_t^{m_t^\star} + (1-\varepsilon)\pi_{\mathbf{x}}^{m'}\beta_t^{m'}v_t^{m'}\right) - \left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m'}\right)v_t^{m'} + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\beta_t^{m_t^\star}v_t^{m_t^\star}\right)\\ &\quad = \varepsilon\left(v_t^{m_t^\star} - v_t^{m'}\right) + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\left(\left(1-\beta_t^{m_t^\star}\right)v_t^{m_t^\star} - \left(1-\beta_t^{m'}\right)v_t^{m'}\right).\end{aligned}$$

Thus, if $\beta_t^{m_t^\star} \geq \beta_t^{m'}$, then $(u_\varepsilon^{m_t^\star}(t,\mathbf{x}) - u_\varepsilon^{m'}(t,\mathbf{x})) \geq (\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}(1-\beta_t^{m_t^\star}))(v_t^{m_t^\star} - v_t^{m'}) > 0$; and, in the case of $\beta_t^{m_t^\star} \geq \beta_t^{m'}$ the conclusion depends whether or not $(v_t^{m_t^\star} - v_t^{m'})$ is greater equal to $(\beta_t^{m_t^\star}v_t^{m_t^\star} - \beta_t^{m'}v_t^{m'})$. If $(v_t^{m_t^\star} - v_t^{m'}) \geq (\beta_t^{m_t^\star}v_t^{m_t^\star} - \beta_t^{m'}v_t^{m'})$, then $(u_\varepsilon^{m_t^\star}(t,\mathbf{x}) - u_\varepsilon^{m'}(t,\mathbf{x})) > 0$; and otherwise, m' must be the pure strategy ε–best response, when ε is arbitrarily small. Moreover, because that $\left((1-\varepsilon)\beta_t^{m'} - 1\right) < 0$ for all $\varepsilon > 0$, $(u_\varepsilon^{m_t^\star}(t,\mathbf{x}) - u_\varepsilon^0(t,\mathbf{x}))$ equals

$$\begin{aligned}&\left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\right)v_t^{m_t^\star} + (1-\varepsilon)\pi_{\mathbf{x}}^{m'}\beta_t^{m'}v_t^{m'}\right) - \left(\pi_{\mathbf{x}}^{m_t^\star}v_t^{m_t^\star} + \pi_{\mathbf{x}}^{m'}v_t^{m'}\right)\\ &\quad = \varepsilon\left(1-\pi_{\mathbf{x}}^{m_t^\star}\right)v_t^{m_t^\star} - \pi_{\mathbf{x}}^{m'}\left((1-\varepsilon)\beta_t^{m'} - 1\right)v_t^{m'} = \pi_{\mathbf{x}}^{m'}\left(\varepsilon v_t^{m_t^\star} - \left((1-\varepsilon)\beta_t^{m'} - 1\right)v_t^{m'}\right) > 0.\end{aligned}$$

For the final case, note that when $\pi_{\mathbf{x}}^{m_t^\star} < \pi_{\mathbf{x}}^{m'}$, $m_t^\star \neq m'$, there is an $\varepsilon > 0$, such that for all $\varepsilon' \leq \varepsilon$ player t will not be able to alter the winning candidate in any of the ε'–perturbed games. Therefore, it suffices to consider the case when $\lambda^0 < 1$ and $\pi_{m_t^\star}^{\mathbf{x}} < \pi_{\mathbf{x}}^{m'}$, $m_t^\star \neq m'$. In this case, for sufficiently small $\varepsilon > 0$, player t collects the net gains of

$$\begin{aligned}&\left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\right)\beta_t^{m_t^\star}v_t^{m_t^\star} + (1-\varepsilon)\pi_{\mathbf{x}}^{m'}v_t^{m'}\right) - \left(\pi_{\mathbf{x}}^{m_t^\star}\beta_t^{m_t^\star}v_t^{m_t^\star} + \pi_{\mathbf{x}}^{m'}v_t^{m'}\right)\\ &\quad = \varepsilon\beta_t^{m_t^\star}(1-\pi_{\mathbf{x}}^{m_t^\star})v_t^{m_t^\star} - \varepsilon\pi_{\mathbf{x}}^{m'}v_t^{m'} = \varepsilon\pi_{\mathbf{x}}^{m'}\left(\beta_t^{m_t^\star}v_t^{m_t^\star} - v_t^{m'}\right),\end{aligned}$$

provided that he chooses $m_t^\star$ instead of not voting; and

$$\left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\right)\beta_t^{m_t^\star}v_t^{m_t^\star} + (1-\varepsilon)\pi_{\mathbf{x}}^{m'}v_t^{m'}\right) - \left(\left(\varepsilon + (1-\varepsilon)\pi_{\mathbf{x}}^{m'}\right)v_t^{m'} + (1-\varepsilon)\pi_{\mathbf{x}}^{m_t^\star}\beta_t^{m_t^\star}v_t^{m_t^\star}\right),$$

which equals to $\varepsilon\left(\beta_t^{m_t^\star}v_t^{m_t^\star} - v_t^{m'}\right)$ when he votes for $m_t^\star$ instead of m', $m_t^\star \neq m'$. ■

Before concluding this section, the extent with which the performance of the concept of strategic equilibrium depends on the magnitude of discontinuities should be clear: In the presidential voting game, the measure for such a discontinuity is β_t^m. Indeed, if $\beta_t^m = 0$ for all t, m, then Proposition 1 exhibits that strategies in which a non-zero measure of players vote for their second best party can easily be sustained in strategic equilibrium. Indeed, the following is a strategic equilibrium when $\beta_t^m = 0$ for all t, m: For all t, $\mathbf{x}(t) = 1$ even when $m_t^\star = 2$. Nevertheless, even in this case not voting does not appear in any strategic equilibrium.

2.2 Cournot Oligopoly

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

The set of agents is given by $T = [0,1]$ and each of them can choose a quantity $\mathbf{x}(t) \in [0, \bar{q}]$, where $\bar{q} \geq 1$, and the symmetric unit cost of production for each $t \in [0,1]$ is 0. Given the quantity choices $\mathbf{x}$ the inverse demand is given by $p = 1 - \int \mathbf{x} \mathrm{d}\lambda$. Here, for simplicity we allow for negative prices which could be corrected by working with a symmetric and positive unit cost.

The profit function of firm t is

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

The set of Nash equilibria in this game is any strategy profile $\mathbf{x}$ satisfying $\int \mathbf{x} \mathrm{d}\lambda = 1$. The reason is that as long as $\int \mathbf{x} \mathrm{d}\lambda = 1$, $p = 0$, thus, any player would be indifferent between any of their choices, since each player is atomless. Moreover, if $\bar{q} = 1$ the Nash equilibrium is unique and is given by $\mathbf{x}(t) = 1$ for almost every player t in $[0,1]$.

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

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

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)

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

Obviously, this term converges to $\mathbf{x}^*(t) = 1$ for almost every player t. Hence, the set of strategic equilibria is given by $\mathbf{x}^* : \mathbf{x}^*(t) = 1$ for almost every $t \in [0,1]$. Thus, unlike for Nash equilibrium, there a unique strategic equilibrium (up to a measure zero set of agents).

3 Non-Atomic Games

Let A be a non-empty, compact metric space of *actions* and $\mathcal{M}(A)$ be the set of Borel probability measures on A endowed with the topology of the weak convergence of probability measures . By Parthasarathy [11, Theorem II.6.4], it follows that $\mathcal{M}(A)$ is a compact metric space. We use the following notation: we write $\mu_n \Rightarrow \mu$ whenever $\{\mu_n\}_{n=1}^\infty \subseteq \mathcal{M}(A)$ converges to μ and ρ denote the Prohorov metric on $\mathcal{M}(A)$, which is known to metricize the weak convergence topology. We let d_A denote the metric on A.

Let $L \in \mathbb{N}$ and $\mathcal{U}$ denote the space of continuous utility functions $u : A \times \mathcal{M}^L \to \mathbb{R}$ endowed with the supremum norm. The set $\mathcal{U}$ represents the space of *players' characteristics*; it is a complete, separable metric space.

The set of *players* is the unit interval $[0,1]$, endowed with the Lebesgue measure λ on the Lebesgue measurable sets. A *game with a continuum of players* is characterized by a measurable function $U : [0,1] \to \mathcal{U}$ and finite partition $\{T_i\}_{i=1}^L$ of $[0,1]$ such that T_i is measurable and $\lambda(T_i) > 0$ for all $i = 1, \ldots, L$. Each set T_i is interpreted as a group or an institution and is endowed with the following measure λ_i: for any measurable set $B \subseteq T_i$, $\lambda_i(B) = \lambda(B)/\lambda(T_i)$. We represent such game by $G = (\{T_i\}_{i=1}^L, U, A)$.

For convenience, let $U_i : T_i \to \mathcal{U}$ denote the restriction of U to T_i. A strategy $f = (f_1, \ldots, f_L)$ is a vector of measurable functions $f_i : T_i \to A$, $i = 1, \ldots, L$. A pair (U, f), where U is written as $(U_1, \ldots, U_L)$ and f is a strategy, induces a vector of probability measures on $\mathcal{U} \times A$ denoted by $(\lambda_1 \circ (U_1, f_1)^{-1}, \ldots, \lambda_1 \circ (U_L, f_L)^{-1}) \in \mathcal{M}(\mathcal{U} \times A)^L$. The payoff of player $t \in T_i$ is

$$U_t(f_i(t), \lambda_1 \circ f_1^{-1}, \ldots, \lambda_L \circ f_L^{-1}).$$

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

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

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

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 are implicitly assuming that the choice of any player does not affect the distribution of actions. It is in this sense that the notions of this section describe a game with a continuum of players.

A strategy $f = (f_1, \ldots, f_L)$ is a *Nash equilibrium of* G if $U_t(f(t), \lambda_1 \circ f_1^{-1}, \ldots, \lambda_L \circ f_L^{-1}) \geq U_t(a, \lambda_1 \circ f_1^{-1}, \ldots, \lambda_L \circ f_L^{-1})$ for all $t \in [0,1]$ and $a \in A$. Nash equilibria exist if either A or $U([0,1])$ is countable, but may fail to exist otherwise.

4 Strategic equilibria

4.1 Strategic Equilibrium Distributions

As we stressed in the introduction, we wish to consider those Nash equilibria that can be seen as a limit of equilibria in games in which players have a small, yet positive, impact in the societal choice of the group they belong to. Clearly, the reason why each player t does not have any impact on the societal choice of his group is because $\lambda_i(\{t\}) = 0$. The way we assign players weights to their group's societal choice is by considering the following measures: For each $1 \leq i \leq L$, $\varepsilon > 0$, and

$t \in [0,1]$, we define a measure $\lambda_{i,t,\varepsilon}$ in the following way: for all Borel-measurable set $B \subseteq T_i$,

$$\lambda_{i,t,\varepsilon}(B) = \begin{cases} \varepsilon + (1-\varepsilon)\lambda_i(B) & \text{if } t \in B \\ (1-\varepsilon)\lambda_i(B) & \text{otherwise.} \end{cases}$$

Thus, in $\lambda_{i,t,\varepsilon}$ player t is an atom in group i: In the game described by $\lambda_{i,t,\varepsilon}$, agent t believes that his choices can have an ε impact on the societal choice of group i. The following Lemma makes this precise.

Lemma 2 *For any measurable $f_i : T_i \to A$,*

$$\int f_i \mathrm{d}\lambda_{i,t,\varepsilon} = \varepsilon f_i(t) + (1-\varepsilon)\int_{T_i} f_i \mathrm{d}\lambda_i.$$

Proof. If f_i is simple, $f_i = \sum_{j=1}^{J} a_j 1_{A_j}$, with $t \in A_1$, then

$$\begin{aligned} \int f_i \mathrm{d}\lambda_{i,t,\varepsilon} &= \varepsilon a_1 + (1-\varepsilon)a_1\lambda_i(A_1) + (1-\varepsilon)\sum_{j=2}^{J} a_j\lambda_i(A_j) \\ &= \varepsilon f_i(t) + (1-\varepsilon)\int_{T_i} f_i \mathrm{d}\lambda_i. \end{aligned}$$

The general case follows from this by a limit argument. ■

Given $i \in \{1,\ldots,L\}$, $a \in A$, a strategy f_i and $t \in T_i$, let $f_i \setminus_t a$ denote the strategy g_i defined by $g_i(t) = a$, and $g_i(\tilde{t}) = f_i(\tilde{t})$, for all $\tilde{t} \neq t$, $\tilde{t} \in T_i$.

Lemma 3 *Let $i \in \{1,\ldots,L\}$, $\varepsilon > 0$, $a \in A$, $t \in T_i$, f_i and g_i be given. If $\lambda_i \circ f_i^{-1} = \lambda_i \circ g_i^{-1}$ then $\lambda_{i,t,\varepsilon} \circ (f_i \setminus_t a)^{-1} = \lambda_{i,t,\varepsilon} \circ (g_i \setminus_t a)^{-1}$.*

Proof. Let $B \subseteq T_i$ be measurable. Then,

$$\begin{aligned} \lambda_{i,t,\varepsilon} \circ (f_i \setminus_t a)^{-1} &= \begin{cases} \varepsilon + (1-\varepsilon)\lambda_i \circ f_i^{-1}(B) & \text{if } a \in B \\ (1-\varepsilon)\lambda_i \circ f_i^{-1}(B) & \text{otherwise.} \end{cases} \\ &= \begin{cases} \varepsilon + (1-\varepsilon)\lambda_i \circ g_i^{-1}(B) & \text{if } a \in B \\ (1-\varepsilon)\lambda_i \circ g_i^{-1}(B) & \text{otherwise.} \end{cases} \\ &= \lambda_{i,t,\varepsilon} \circ (g_i \setminus_t a)^{-1}(B). \end{aligned}$$

Thus, $\lambda_{i,t,\varepsilon} \circ (f_i \setminus_t a)^{-1} = \lambda_{i,t,\varepsilon} \circ (g_i \setminus_t a)^{-1}$. ■

For $i \in \{1,\ldots,L\}$, let $(\tau_j)_{j=1}^{L} \in \mathcal{M}(A)^L$. Then there exist a measurable function $f_i : T_i \to A$ such that $\tau_i = \lambda_i \circ f_i^{-1}$. For all $\varepsilon > 0$, $t \in T_i$ and $a \in A$, define

$$U_{i,\varepsilon}(t)\left(a, (\tau_j)_{j=1}^{L}\right) = U_i(t)\left(a, (\lambda_{i,t,\varepsilon} \circ (f_i \setminus_t a)^{-1}, \tau_{-i})\right).$$

By Lemma 3, $U_{i,\varepsilon}$ is well defined. We then define the *$\varepsilon-$perturbed game* G_ε of G as $G_\varepsilon = \left(\{T_i, U_{i,\varepsilon}\}_{i=1}^{L}, A\right)$. The $\varepsilon-$perturbed game has the same players, and actions spaces as the original game G, but differs from this because in G_ε every players believes that he has an ε impact on the distribution of actions of his group.

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 $\{\varepsilon_k\}_{k\in\mathbb{N}} \subseteq \mathbb{R}_{++}$ decreasing to zero and a sequence $\{\tau_k^*\}_{k\in\mathbb{N}}$ converging to τ^* such that τ_k^* is an equilibrium distribution of G_{ε_k}, for every $k \in \mathbb{N}$.

Conceptually, our approach is in the same spirit as Selten [14]. Given $\varepsilon \in (0, 1]$ and a non-atomic game, we define its ε-perturbed game, a modified version of the original non-atomic game, in which every player thinks he alone has ε impact on the distribution of actions of his group. Then, an *ε-strategic* equilibrium distribution is an equilibrium of the ε-game. Finally, a distribution is a strategic equilibrium distribution, if it is a limit point of a sequence of ε-strategic equilibrium distributions, where $\varepsilon > 0$ and $\varepsilon \searrow 0$.

Theorem 1 is on the existence of a strategic equilibrium distribution.

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

Proof. Note first that if $G = (\{T_i, U_i\}_i, A)$, $\tilde{G} = (\{T_i, V_i\}_i, A)$ and $U_i = V_i$ a.e. for all i, then τ is an equilibrium distribution of G if and only if τ is an equilibrium equilibrium of $\tilde{G}$. Hence, under the same hypothesis, τ is a strategic equilibrium distribution of G if and only if τ is a strategic equilibrium distribution of $\tilde{G}$. Then, without loss of generality, assume that U_i is Borel measurable for all i.

Let $\varepsilon > 0$ and $1 \leq i \leq L$. We claim that that $U_{i,\varepsilon}\ T_i \to \mathcal{U}$ is Borel measurable and $U_{i,\varepsilon}(t)$ is continuous for all $t \in T_i$.

Claim 1 *The function $U_{i,\varepsilon}(t) : A \times \mathcal{M}(A)^L \to$ is continuous for all $t \in T_i$.*

Proof. Let $a \in A$ and $\tau \in \mathcal{M}(A)^L$, and let $\{a_k\} \subseteq A$ and $\{\tau_k\} \subseteq \mathcal{M}(A)^L$ be such that $a_k \to a$ and $\tau_k \Rightarrow \tau$. Since A is a complete metric space, by Skorokhod's Theorem (see Hildenbrand [5, Theorem 37, p. 50]), there exist measurable functions f_i and $f_{i,k}$, $k \in \mathbb{N}$, of T_i into A such that $\tau_i = \lambda_i \circ f_i^{-1}$, $\tau_k = \lambda_i \circ f_{i,k}^{-1}$ and $\lim_k f_{i,k} = f_i$ a.e. in T_i. Thus,

$$U_{i,\varepsilon}(t)(a_k, \tau_k) = U_i(t)(a_k, (\lambda_{i,\varepsilon,t} \circ (f_{i,k} \setminus_t a_k)^{-1}, \tau_{-i,k})),$$

and so it is enough to show that $\lambda_{i,\varepsilon,t} \circ (f_{i,k} \setminus_t a_k)^{-1} \Rightarrow \lambda_{i,\varepsilon,t} \circ (f_i \setminus_t a)^{-1}$.

Let $h : A \to \mathbb{R}$ be continuous. Then, by the Change-of-variable formula (see Hildenbrand [5, Theorem 36, p. 50]) and Lemma 2, one obtains

$$\begin{aligned}
&\int_A h \mathrm{d}\lambda_{i,\varepsilon,t} \circ (f_{i,k} \setminus_t a_k)^{-1} = \int_{T_i} h \circ (f_{i,k} \setminus_t a_k) \mathrm{d}\lambda_{i,\varepsilon,t} = \\
&\varepsilon h(a_k) + (1-\varepsilon) \int_{T_i} h \circ f_{i,k} \mathrm{d}\lambda_i = \varepsilon h(a_k) + (1-\varepsilon) \int_A h \mathrm{d}\lambda_i \circ f_{i,k}^{-1} \to \\
&\varepsilon h(a) + (1-\varepsilon) \int_A h \mathrm{d}\lambda_i \circ f_i^{-1} = \int_A h \mathrm{d}\lambda_{i,\varepsilon,t} \circ (f_i \setminus_t a)^{-1}.
\end{aligned}$$

Thus, indeed we have $\lambda_{i,\varepsilon,t} \circ (f_{i,k} \setminus_t a_k)^{-1} \Rightarrow \lambda_{i,\varepsilon,t} \circ (f_i \setminus_t a)^{-1}$. ■

Claim 2 *The function $U_{i,\varepsilon} : T_i \to \mathcal{U}$ is Borel measurable.*

Proof. Note first that $U_i^{(a,\tau)}$ defined by $t \mapsto U_i(t)(a,\tau)$ is Borel measurable for all a and τ because $U_{i,(a,\tau)} = \pi^{(a,\tau)} \circ U_i$, where $\pi^{(a,\tau)}(u) = u(a,\tau)$ is continuous.

This fact implies that $U_{i,\varepsilon}^{(a,\tau)}$ is Borel measurable for all a and τ, as follows: note that $\lambda_{i,\varepsilon,t} \circ (f_i \setminus_t a)^{-1} = \lambda_{i,\varepsilon,\tilde{t}} \circ (f_i \setminus_{\tilde{t}} a)^{-1}$, for any measurable $f_i : T_i \to A$, $a \in A$ and $t, \tilde{t} \in T_i$. Thus, we can fix $t_0 \in T_i$ and write $U_{i,\varepsilon}^{(a,\tau)} = U_i^{(a,(\lambda_{i,\varepsilon,t_0} \circ (f_i \setminus_{t_0} a)^{-1}, \tau_{-i}))}$ if $\tau_i = \lambda_i \circ f_i^{-1}$, and so $U_{i,\varepsilon}^{(a,\tau)}$ is Borel measurable. It then follows from Lemma 4 that $U_{i,\varepsilon}$ is Borel measurable. ■

Due to claims 1 and 2, it follows by (a straightforward generalization of) Theorem 1 in Mas-Colell [8], it follows that G_ε, has an equilibrium distribution.

To finish the proof, we let τ_n^* be an equilibrium distribution of $G_{1/n}$. For all $i \in \{1, \ldots, L\}$, let

$$K_i = \{\lambda_i \circ U_i^{-1}, \lambda_i \circ U_{i,1}^{-1}, \lambda_i \circ U_{i,1/2}^{-1}, \ldots\}$$

and

$$C_i = \{\mu \in \mathcal{M}(\mathcal{U} \times A) : \mu_{i,\mathcal{U}} \in K_i\}.$$

It follows by [5, 32 and 33] that K_i is tight, and so [5, 34 and 35] implies that C_i is tight. Since $\{\tau_k^*\} \subseteq C_1 \times \cdots \times C_L$, it follows by [5, 31] that it has a converging subsequence. Hence, its limit point is a strategic equilibrium distribution of G. ■

We show that any strategic equilibrium is a Nash equilibrium.

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

Proof. Let τ^* be a strategic equilibrium distribution, and let $\{\varepsilon_k\}$ and $\{\tau_k^*\}$ be such that $\varepsilon_k \in \mathbb{R}_{++}$, $\lim_k \varepsilon_k = 0$, τ_k^* converges to τ^*, and τ_k^* is a Nash equilibrium distribution of G_{ε_k}, for all $k \in \mathbb{N}$. Since $\tau_{i,A,n} \Rightarrow \tau_{i,A}$, taking a subsequence if necessary, we may assume that $\rho(\tau_{i,A}, \tau_{i,A,n}) < 1/n$ for all $1 \leq i \leq L$.

Define, for each $u \in \mathcal{U}$,

$$\beta_n(u) = \sup_{a \in A, \nu \in \mathcal{M}(A)^L} \{|u(a,(\nu_i)_i) - u(a,(\tau_{i,A})_i)| : \rho(\nu_i, \tau_{i,A}) < 1/n \text{ for all } i\}.$$

Since u is continuous on $A \times \mathcal{M}(A)^L$, which is compact, it follows that u is uniformly continuous. Thus, $\beta_n(u) \searrow 0$ as $n \to \infty$. We claim that β_n is continuous in $\mathcal{U}$.

Let $\eta > 0$. Define $\delta < \eta/2$. Then if $||u - v|| < \delta$, we have for any $a \in A$, and $\nu \in \mathcal{M}(A)^L$ such that $\rho(\nu_i, \tau_{i,A}) < 1/n$ for all i,

$$\begin{aligned} |v(a,(\nu_i)_i) - v(a,(\tau_{i,A})_i)| &\leq |v(a,(\nu_i)_i) - u(a,(\nu_i)_i)| \\ &\quad + |u(a,(\nu_i)_i) + u(a,(\tau_{i,A})_i)| + |v(a,(\tau_{i,A})_i) - u(a,(\tau_{i,A})_i)| \\ &< \delta + \beta_n(u) + \delta, \end{aligned}$$

and so $\beta_n(v) \leq 2\delta + \beta_n(u) < \eta + \beta_n(u)$. By symmetry, $\beta_n(u) < \eta + \beta_n(v)$, and so $|\beta_n(u) - \beta_n(v)| < \eta$. Hence, β_n is continuous, as claimed.

Given the definition of β_n, $B_{\tau_n} \subseteq D_n$ where

$$D_n := \{(u,a) : u(a,(\tau_{i,A})_i) \geq u(A,(\tau_{i,A})_i) - 2\beta_n(u)\}.$$

Since β_n is continuous, we see that D_n is closed, and so Borel measurable. Thus, $\tau_{i,n}(D_n) = 1$, for all $i = 1, \ldots, L$. Also, $D_n \searrow B_\tau$. Hence, for fixed $n' \in \mathbb{N}, n' \geq n$, it follows that $\tau_{i,n'}(D_n) \geq \tau_{i,n'}(D_{n'}) = 1$, and so $\tau_i(D_n) \geq \limsup_{n'} \tau_{i,n'}(D_n) = 1$. Hence, $\tau_i(B_\tau) = \lim_n \tau_i(D_n) = 1$. Therefore, τ is an equilibrium distribution of G. ■

4.2 Strategic Equilibrium Strategies

We say that a strategy $f^* = (f_i^*)_{i=1}^L$ is a *strategic equilibrium strategy* of G if there exists a sequence $\{\varepsilon_k\}_{k\in\mathbb{N}} \subseteq \mathbb{R}_{++}$ decreasing to zero and a sequence $\{f_k^*\}_{k\in\mathbb{N}}$ converging to f^* in distribution such that f_k^* is an Nash equilibrium of G_{ε_k}, for every $k \in \mathbb{N}$.

Proposition 2 *Suppose that G is such that either A or $U([0,1])$ is countable. A strategy f is a strategic equilibrium if and only if $\lambda \circ f^{-1}$ is the marginal of some strategic equilibrium distribution.*

Proof. Suppose that f is a strategic equilibrium strategy of G. Let $\{e_k\}_k$ and $\{f_k\}_k$ be such that $\varepsilon_k \searrow 0$, f_k is a Nash equilibrium of G_{ε_k} and that $\lambda \circ f_k^{-1} \Rightarrow \lambda \circ f^{-1}$. Define $\tau_k = \lambda \circ (U, f_k)^{-1}$ for all k, and, taking a subsequence if necessary, assume that $\tau_k \Rightarrow \tau$. Then, τ is a strategic equilibrium distribution of G and $\tau_A = \lambda \circ f^{-1}$.

Conversely, let τ be a strategic equilibrium distribution such that $\tau_A = \lambda \circ f^{-1}$. Let $\{e_k\}_k$ and $\{\tau_k\}_k$ be such that $\varepsilon_k \searrow 0$, τ_k is a Nash equilibrium of G_{ε_k} and that $\tau_k \Rightarrow \tau$. Then, by Theorem 2 in Carmona [4], there exist f_k such that f_k is a Nash equilibrium of G_{ε_k} and $\lambda \circ f_k^{-1} = \tau_{A,k}$ for all k. Since $\lambda \circ f_k^{-1} = \tau_{A,k} \Rightarrow \tau_A = \lambda \circ f^{-1}$, it follows that f is a strategic equilibrium of G. ■

5 Strategic interaction

As we have noted in the introduction, the concept of strategic equilibrium is the correct non-atomic version of the Nash equilibrium in finite games. To show that formally, we will be using Nash's mass action interpretation taken from his Ph.D. thesis. For any given finite normal form game, using the mass action interpretation, we will formulate its associated non-atomic version. In theorem 3 we will prove that a strategy profile, in the non-atomic version of the given finite normal form game, is a strategic equilibrium if and only if the associated strategy profile in the finite normal form game is a Nash equilibrium.

In his Ph.D. dissertation, John Nash proposed two interpretations of his equilibrium concept, with the objective of showing how equilibrium points "(...) can be connected with observable phenomenon." (Nash [9, p. 21]) One interpretation is rationalistic: if we assume that players are rational, they 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, which Nash names mass-action interpretation, is less demanding on the players. In this interpretation, "[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

[3]For a formal discussion of these ideas, see Aumann and Brandenburger [1] and the Nobel Seminar 1994 [10].

complex reasoning processes." (Nash [9, p. 21]) 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 Nash equilibrium.

Below we present not only a new interpretation of Nash equilibrium but also prove that the notion of strategic equilibrium is the non-atomic extension of Nash equilibrium in finite normal form games.

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, $\Delta(A_i)$ 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 in n groups, from each of which a participant is draw at random. For concreteness, let $T_i = [0, 1]$, and $X_i = \Delta(A_i)$, for any $i \in N$; a player $t \in T_i$ chooses an element of $\Delta(A_i)$. From each T_i a player is selected according to the Lebesgue measure, and so the probability that the player selected from the ith group will play action $a_i \in A_i$ is $\int_{T_i} x_i^{a_i}$. We thus define $s_i(x_i) = \int_{T_i} x_i$, and

$$u_i(t, x) = v_i(s_1(x_1), \ldots, s_n(x_n)) = \sum_{a\in A} \prod_{i\in N} s_i^{a_i}(x_i) v_i(a).$$

We denote by G the game $(T_i, P_i, u_i)_{i\in N}$

Theorem 3 *A strategy profile $(x_1^*, \ldots, x_n^*)$ is a strategic equilibrium of G if and only if $(s_1(x_1^*), \ldots, s_n(x_n^*))$ is a Nash equilibrium of Γ.*

Proof. (Sufficiency) Let $(x_1^*, \ldots, x_n^*)$ be a strategy in G, and assume that $s^* := (s_1(x_1^*), \ldots, s_n(x_n^*))$ is a Nash equilibrium of Γ. Let $i \in N$. We have that $v_i(s^*) \geq v_i(s_i, s_{-i}^*)$ for all $s_i \in \Delta(A_i)$. This implies, in particular, that $v_i(s^*) \geq v_i(\varepsilon x(t) + (1 - \varepsilon)s_i^*, s_{-i}^*)$ for all $t \in T_i$, and $\varepsilon > 0$. Hence, $(x_1^*, \ldots, x_n^*)$ is a Nash equilibrium of G_ε for all $\varepsilon > 0$, and so a strategic equilibrium of G.

(Necessity) Let $(x_1^*, \ldots, x_n^*)$ be a strategic equilibrium of G, and let $s^* := (s_1(x_1^*), \ldots, s_n(x_n^*))$. We will show that for all $i \in N$, and $a_i \in A_i$ if $s_i^{a_i}(x_i^*) > 0$, then a_i maximizes $v_i(a_i, s_{-i}^*)$ in A_i.

Let $i \in N$, $a_i \in A_i$. If a_i does not maximize $v_i(a_i, s_{-i}^*)$ in A_i, then a_i does not maximize $v_i(a_i, s_{-i}^\varepsilon)$ in A_i for all $\varepsilon > 0$ sufficiently small, where $s^\varepsilon := (s_1(x_1^\varepsilon), \ldots, s_n(x_n^\varepsilon))$, and $(x_1^\varepsilon, \ldots, x_n^\varepsilon)$ is a Nash equilibrium of G_ε, $x^\varepsilon \to x^*$, and $\varepsilon \to 0$. Hence, $x_i^{\varepsilon,a_i}(t) = 0$ a.e. $t \in T_i$, and so $x_i^{*a_i}(t) = 0$ a.e. $t \in T_i$. Thus, $s_i^{a_i}(x_i^*) = 0$. ■

A Appendix

Lemma 4 *Let $T \subseteq [0, 1]$ and $g : T \to \mathcal{U}$ be such that $t \mapsto g(t)(a, \tau)$ is Borel measurable for all $a \in A$ and $\tau \in \mathcal{M}$. Then, g is Borel measurable.*

Proof. It is enough to show that $g^{-1}(B_\eta(u))$ is Borel measurable for all $\eta > 0$ and all $u \in \mathcal{U}$.

Let $\eta > 0$ and $u \in \mathcal{U}$. Consider first the closed balls: let $B = \{v \in \mathcal{U} : ||u - v|| \leq \eta\}$. Then,

$$\begin{aligned} g^{-1}(B) &= \{t : ||g(t) - u||_\infty \leq \eta\} \\ &= \{t : ||g(t)(a,\tau) - u(a,\tau)|| \leq \eta \text{ for all } a, \tau\} \\ &= \cap_{k=1}^{\infty} \{t : ||g(t)(a_k,\tau_k) - u(a_k,\tau_k)|| \leq \eta\}, \end{aligned}$$

where $\{(a_k, \tau_k)\}$ is a dense subset of $A \times \mathcal{M}$. Since $\{t : ||g(t)(a_k,\tau_k) - u(a_k,\tau_k)|| \leq \eta\}$ is Borel measurable for all k, it follows that $g^{-1}(B)$ is Borel measurable.

Let now $B = \{t : ||g(t) - u|| = \eta\}$. Then, since both $g(t)$ and u are continuous in $A \times \mathcal{M}$, it follows that

$$\begin{aligned} g^{-1}(B) &= \{t : ||g(t) - u||_\infty \leq \eta\} \cap \{t : ||g(t)(a,\tau) - u(a,\tau)|| = \eta \text{ for some } a, \tau\} \\ &= \{t : ||g(t) - u||_\infty \leq \eta\} \cap \cap_{j=1}^{\infty} \cup_{k=1}^{\infty} \{t : ||g(t)(a_k,\tau_k) - u(a_k,\tau_k)|| > \eta - \frac{1}{j}\}. \end{aligned}$$

Since $\{t : ||g(t)(a_k,\tau_k) - u(a_k,\tau_k)|| > \eta - \frac{1}{j}\}$ is Borel measurable for all k and j, it follows that $g^{-1}(B)$ is Borel measurable.

Thus, $g^{-1}(B_\eta(u)) = \{t : ||g(t) - u||_\infty \leq \eta\} \setminus \{t : ||g(t) - u||_\infty = \eta\}$ is Borel measurable. ■

References

[1] Aumann, R. and A. Brandenburger (1995), "Epistemic Conditions for Nash Equilibrium", *Econometrica,* **5,** 1161-1180.

[2] Carmona, G. (2003), "Equilibria in Economies with a Continuum of Agents: A Characterization in terms of Approximate Equilibria of Large Finite Economies," Universidade Nova de Lisboa.

[3] Carmona, G. (2003), "Nash and Limit Equilibria of Games with a Continuum of Players," Universidade Nova de Lisboa.

[4] Carmona, G. (2003), "On a Theorem by Mas-Colell," Universidade Nova de Lisboa.

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

[6] Hart, S., W. Hildenbrand, and E. Kohlberg (1974), "On Equilibrim Allocations as Distributions on the Commodity Space", *Journal of Mathematical Economics*, **1**, 159-166.

[7] Khan, M. A., K. P. 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.

[8] Mas-Colell, A. (1984) "On a Theorem by Schmeidler," *Journal of Mathematical Economics*, 13, 201-206.

[9] Nash, J. (1950), *Non-Cooperative Games*, Ph.D Thesis, Princeton University.

[10] Nobel Seminar, December 8, 1994, (1996) "The Work of John Nash in Game Theory", *Journal of Economic Theory,* **69,** 153-185.

[11] Parthasarathy, K. (1967) *Probability Measures on Metric Spaces*, New York: Academic Press.

[12] Rath, K.P. (1992), "A Direct Proof of the Existence of Pure Strategy Equilibria in Games with a Continuum of Players", *Economic Theory,* **2,** 427-433.

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

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