

Università degli Studi di Milano - Bicocca
Università Cattolica del Sacro Cuore

PhD THESIS

Essays on Strategic Interactions

CANDIDATE

Claudia Meroni

SUPERVISOR

Francesco De Sinopoli

Graduate School in Public Economics - DEFAP

Academic Year 2013/14

Abstract

In Poisson games, an extension of perfect equilibrium based on perturbations of the strategy space does not guarantee that players use admissible actions. This observation suggests that such a class of perturbations is not the correct one. In the first chapter, the right space of perturbations is characterized to give a definition of perfect equilibrium in Poisson games. Furthermore, such a space is used to define the corresponding strategically stable sets of equilibria. They are shown to satisfy existence, admissibility, and robustness against iterated deletion of dominated strategies and inferior replies.

In the second chapter, it is shown that in every Poisson game the number of connected components of Nash equilibria is finite. This result is obtained by exploiting the geometric structure of the Nash equilibrium set, which is shown to be a semi-analytic set; i.e., a set defined by a finite system of analytic inequalities. Furthermore, it is also shown that every Poisson game has a stable set contained in a connected component of equilibria.

In the third chapter, a different model of strategic interaction is analyzed, the double round-robin tournament. A tournament is a simultaneous n -player game that is built on a two-player game g . Each player meets all the other players in turn and in every match the game g is played. The winner of the tournament is the player who attains the highest total score, which is given by the sum of the payoffs that he gets in all the matches he plays. Arad and Rubinstein (2013) analyze tournaments of the *round-robin* type, where each player is matched with every other player once to play a symmetric game g . With the aim of extending the analysis of tournaments to asymmetric games, *double round-robins* are studied, where each player meets all his opponents twice to play g in alternating roles. In particular, the relationship between equilibria of the tournament and equilibria of the base game g is explored.

Contents

Abstract	ii
Chapter 1. Strategic Stability in Poisson Games	1
1.1. Introduction	1
1.2. Preliminaries	4
1.3. Admissibility	6
1.4. Perfection	9
1.5. Stability	17
1.6. Examples	21
Appendix A. Stable Sets in Generic Poisson Games	25
Appendix B. On the Inadequacy of Inner-Perfection	26
Chapter 2. The Structure of Nash Equilibria in Poisson Games	29
2.1. Introduction	29
2.2. Preliminaries	30
2.3. The Set of Nash Equilibria	32
2.4. Connectedness	35
Chapter 3. Double Round-Robin Tournaments	39
3.1. Introduction	39
3.2. The Model	43
3.3. The Two-Player Interaction	45
3.4. The Tournament	49
3.5. Examples	53
Bibliography	57

CHAPTER 1

Strategic Stability in Poisson Games*

1.1. Introduction

Poisson games belong to the general class of models characterized by population uncertainty (Myerson, 1998; Milchtaich, 2004). These games have been introduced by Myerson (1998), and have been widely used to model voting behavior (see, e.g., Myerson (2002); Bouton and Castanheira (2012); Núñez (2014)) as well as more general economic environments (see, e.g., Satterthwaite and Shneyerov (2007); Makris (2008, 2009); Ritzberger (2009); McLennan (2011); Jehiel and Lamy (2014)). In these models, every player does not know the exact number of other players in the population. Each player in the game, however, has probabilistic information about it and, given some beliefs about how the members of such a population behave, can compute the expected payoff that results from each of her available choices. Thus, a Nash equilibrium in this context is a description of behavior for the entire population that is consistent with the players' utility maximizing actions given that they use such a description to form their beliefs about the population's expected behavior.

Similarly to standard normal form and extensive form games, one can easily construct examples of Poisson games where not every Nash equilibrium is a plausible description of rational behavior. In particular, Nash equilibria in Poisson games can be in dominated strategies. Indeed, many applications of Poisson games (see, e.g., Myerson (2002); Maniquet and Morelli (2013); Bouton and Castanheira (2012); Núñez (2014); among others) focus on undominated strategies in their analysis. In addition, there are also examples in the applied

* This chapter is based on De Sinopoli et al. (2014).

literature of Poisson games that use some other kind of refinements (Huges (2012); Bouton (2013); Bouton and Gratton (forthcoming)). Hence, it seems worthwhile exploring, also in games with population uncertainty, what can be said from a theoretical standpoint about which Nash equilibria are the most reasonable and to propose a definition that selects such equilibria for us.

Following the main literature on equilibrium refinements, we start focusing our attention on *admissibility*. That is, the principle prescribing players not to play dominated strategies (Luce and Raiffa (1957, p.287, Axiom 5)). Furthermore, as in Kohlberg and Mertens (1986), we also require that the solution be robust against iterated deletion of dominated actions. Unfortunately, as it is already well known, such an iterative process can lead to different answers depending on which order is chosen to eliminate the dominated strategies. The response to this caveat is defining a set-valued solution concept and requiring that every solution to a Poisson game contain a solution to any game that can be obtained by eliminating dominated strategies. Of course, a definition of such a concept for Poisson games should be guided by the literature on Strategic Stability for finite games (Kohlberg and Mertens (1986); Mertens (1989, 1991); Hillas (1990); Govindan and Wilson (2008)). In broad terms, a strategically stable set is a subset of Nash equilibria that is robust against every element in some given space of perturbations. The choice of such a space determines the properties that the final concept satisfies and the perturbations are just a means of obtaining the game theoretical properties that we desire (Kohlberg and Mertens (1986, p. 1005, footnote 3)). As argued above, a strategically stable set of equilibria should only contain undominated strategies. Furthermore, it should always contain a strategically stable set of any game obtained by eliminating a dominated strategy. However, De Sinopoli and Pimienta (2009) show that the main instrument used to define strategic stability in normal form games—i.e. Nash equilibria of strategy perturbed games—fails to guarantee that players only use undominated strategies when applied to Poisson games.

Thus, before defining strategically stable sets of equilibria in Poisson games we need to find the appropriate space of perturbations that guarantees that every member of the stable set is undominated. It turns out that the “right” space of perturbations is of the same nature as the one used in infinite normal-form games (Simon and Stinchcombe (1995); Al-Najjar (1995); Carbonell-Nicolau (2011)) and different from the one used in finite games (Selten (1975)) even if players have finite action sets. Once this class of perturbations has been identified, it can be reinterpreted as a collection of perturbations of the best response correspondence. Then, a stable set is defined as a minimal subset of fixed points of the best response correspondence with the property that every correspondence that can be obtained using such perturbations has a fixed point close to it.

As an illustration of stable sets in Poisson games we construct a referendum game with a threshold for implementing a new policy (see Example 1.5). In this example, every voter prefers the new policy over the status quo but some voters incur a cost in supporting it. Given the parametrization that we use, the game has three equilibria which can be ranked according to the probability of implementation of the new policy: *zero*, *low*, and *high*. We show that the first equilibrium is dominated because, in particular, voters who do not incur any cost do not support the new policy. In the second equilibrium, only voters who incur the cost do not support the new policy, even if they are indifferent between supporting it or not. Furthermore, every such a voter would strictly prefer supporting the new policy and paying the cost if the share of voters supporting the new policy was slightly higher than the equilibrium one. We show that this equilibrium is undominated and perfect but becomes unstable once dominated strategies are eliminated. Hence, the unique stable set of the game is the equilibrium in which the new policy is implemented with *high* probability.

We review the general description of Poisson games in the next Section. We then discuss the admissibility postulate in Section 1.3 and the definition of

perfection in Section 1.4. The space of perturbations used to define perfect equilibria is used to describe, in Section 1.5, the stable sets of equilibria in Poisson games. We show that they satisfy existence, admissibility and iterated deletion of dominated strategies. Section 3.5 contains some applications of stability. In Appendix A we show that, in generic Poisson games, every Nash equilibrium is a singleton stable set.

1.2. Preliminaries

We begin fixing a *Poisson game* $\Gamma \equiv (n, \mathcal{T}, r, C, (C_t)_{t \in \mathcal{T}}, u)$. The number of *players* is distributed according to a Poisson random variable with parameter n . Hence, the probability that there are k players in the game is equal to

$$P(k | n) = \frac{e^{-n} n^k}{k!}.$$

The set $\mathcal{T} = \{1, \dots, T\}$ is the set of player types. The probability that a randomly selected player is of each type is given by the vector $r = (r_1, \dots, r_T) \in \Delta(\mathcal{T})$.¹ That is, a player is of type $t \in \mathcal{T}$ with probability r_t .

The finite set of *actions* is C . However, we allow that not every action be available to type t players. The set of actions that are in fact available to players of type t is $C_t \subset C$.² An action profile $x \in Z(C)$ specifies for each action $c \in C$ the number of players $x(c)$ that have chosen that action. The set of *action profiles* is $Z(C) \equiv \mathbb{Z}_+^C$. Players' preferences in the game are summarized by $u = (u_1, \dots, u_T)$. It is assumed that each function $u_t : C_t \times Z(C) \rightarrow \mathbb{R}$ be bounded. We interpret $u_t(c, x)$ as the payoff accrued by a type t player when she chooses action c and the realization resulting from the rest of the population's behavior is the action profile $x \in Z(C)$.

The set of *mixed actions* for players of type t is $\Delta(C_t)$. If $\alpha \in \Delta(C_t)$ the *carrier* of α is the subset $\mathcal{C}(\alpha) \subset C_t$ of pure actions that are given strictly positive probability by α . We identify the mixed action that attaches probability one

¹ For any finite set K we write $\Delta(K)$ for the set of probability distributions on K .

² Given two sets E and F , we use the expression $E \subset F$ allowing for set equality.

to action $c \in C$ with the pure action c . As in Myerson (1998), a *strategy function* σ is an element of $\Sigma \equiv \{\sigma \in \Delta(C)^{\mathcal{T}} : \sigma_t \in \Delta(C_t) \text{ for all } t\}$. That is, a strategy function maps types to the set of mixed actions available to the corresponding type. We always write strategy functions as bracketed arrays $(\sigma_1, \dots, \sigma_T)$ where $\sigma_t \in \Delta(C_t)$ for $t = 1, \dots, T$. Furthermore, we may also refer to strategy functions simply as *strategies*. The “average” behavior induced by the strategy function σ is represented by $\tau(\sigma) \in \Delta(C)$ and it is defined by $\tau(\sigma)(c) \equiv \sum_{t \in \mathcal{T}} r(t) \sigma_t(c)$. Construct the set $\tau(\Sigma) \equiv \{\tilde{\tau} \in \Delta(C) : \tilde{\tau} = \tau(\sigma) \text{ for some } \sigma \in \Sigma\}$.³ When the population’s aggregate behavior is summarized by $\tau \in \tau(\Sigma)$, the probability that the action profile $x \in Z(C)$ is realized is equal to

$$\mathbf{P}(x | \tau) \equiv \prod_{c \in C} \left(e^{-n\tau(c)} \frac{(n\tau(c))^{x(c)}}{x(c)!} \right).$$

The expected payoff to a type t player who plays $c \in C_t$ is computed as usual,

$$U_t(c, \tau) \equiv \sum_{x \in Z(C)} \mathbf{P}(x | \tau) u_t(c, x). \quad (1.2.1)$$

Note that, for each type $t \in \mathcal{T}$, each action $c \in C_t$ defines a bounded and continuous function $U_t(c, \cdot) : \Delta(C) \rightarrow \mathbb{R}$.

Action $c \in C_t$ is a *pure best response* against $\tau \in \Delta(C)$ for players of type t if $c \in \arg\max_{c' \in C_t} U_t(c', \tau)$. The finite set of such actions is written $\text{PBR}_t(\tau)$. The set of *best responses* against τ is $\text{BR}_t(\tau) \equiv \Delta(\text{PBR}_t(\tau))$. We write $\text{BR}(\tau) \subset \Sigma$ for the collection of strategy functions σ that satisfy $\sigma_t \in \text{BR}_t(\tau)$ for every t .

³ In the usual description of Poisson games (Myerson (1998)), the set $\tau(\Sigma)$ coincides with $\Delta(C)$ because every type has the same action set. We need to relax this assumption because after eliminating dominated actions different types can end up with different action sets. In these cases, $\tau(\Sigma)$ is a subset of $\Delta(C)$. Take as an example a plurality voting game with three candidates (a , b and c) and two types of voters, each type of voter having the same ex-ante probability. Let type 1 voters have preferences $a >_1 b >_1 c$ and let type 2 voters have preferences $c >_2 b >_2 a$. There is no cost of voting and abstention is not possible. Consider the game obtained after eliminating dominated actions so that type 1 voters *cannot* vote for candidate c and type 2 voters *cannot* vote for candidate a . In this game, $\tau(\Sigma) \equiv \{(\tilde{\tau}(a), \tilde{\tau}(b), \tilde{\tau}(c)) \in \Delta(C) : \tilde{\tau}(a), \tilde{\tau}(c) \leq 1/2\}$.

DEFINITION 1.1 (Nash equilibrium). The strategy function σ is a Nash equilibrium of the Poisson game Γ if $\sigma \in \text{BR}(\tau(\sigma))$.

Since Σ is compact and convex and $\text{BR} \circ \tau$ is upper semicontinuous and convex valued, every Poisson game has a Nash equilibrium (Myerson (1998)). Furthermore, once we fix n , \mathcal{T} , r , C and $(C_t)_{t \in \mathcal{T}}$, standard arguments show that the Nash equilibrium correspondence (mapping utilities to equilibria) is upper semicontinuous.

1.3. Admissibility

Consider a referendum where voters have only two options, voting *yes* or *no* to some policy question. For the policy to be implemented the law requires that at least $K > 1$ voters vote *yes*, otherwise the policy is not implemented. Every voter in the game wants the policy to be implemented. The strategy that prescribes every player to vote *no* is a Nash equilibrium, however, it is clear that such a strategy is dominated. Similar examples can be easily constructed.

We now introduce the standard concept of dominated actions and dominated strategies.

DEFINITION 1.2 (Dominated actions). Action α is dominated by β for players of type t if $U_t(\alpha, \tau) \leq U_t(\beta, \tau)$ for every $\tau \in \tau(\Sigma)$ and $U_t(\alpha, \tau') < U_t(\beta, \tau')$ for some $\tau' \in \tau(\Sigma)$.

That is, an action α is dominated if there is another action such that, *regardless of what other players do*, always gives higher utility than α and, sometimes, strictly higher. We say that an action α is *strictly dominated* by β if the inequality is strict for every $\tau \in \tau(\Sigma)$. Following from this concept, there is a definition of dominated strategies.

DEFINITION 1.3 (Dominated strategies). The strategy function σ is dominated if there is a $t \in \mathcal{T}$ such that σ_t is a dominated action for players of type t .

Likewise, a strategy function is *strictly dominated* if it prescribes a strictly dominated action for some type. De Sinopoli and Pimienta (2009) prove that every Poisson game has a Nash equilibrium in undominated strategies.

In an attempt to capture undominated behavior, we can also give a straightforward extension of the definition of perfection to Poisson games. If E is a finite set, let us denote by $\Delta^\circ(E)$ the set of completely mixed probability distributions on E . This is the set of distributions that give strictly positive probability to every element in E . We now define a perturbation as a pair $(\varepsilon, \sigma^\circ)$ where $\varepsilon > 0$ and σ° is a completely mixed, i.e., a strategy function such that $\sigma_t^\circ \in \Delta^\circ(C_t)$ for every $t \in \mathcal{T}$. In a perturbed game and under the perturbation $(\varepsilon, \sigma^\circ)$, if the strategy function σ is played then, for each type t , the action σ_t is substituted by $(1 - \varepsilon)\sigma_t + \varepsilon\sigma_t^\circ$. Given a *strategy-perturbation* $(\varepsilon, \sigma^\circ)$ we denote the corresponding *strategy-perturbed Poisson game* by $\Gamma_{\varepsilon, \sigma^\circ}$. We can now give the usual definition of perfect equilibrium. For the time being, we call it *inner-perfection*.

DEFINITION 1.4 (Inner-perfection). The strategy function σ is an *inner-perfect equilibrium* if there is a sequence of perturbations $\{(\varepsilon^k, \sigma^k)\}_k$ and a sequence of strategy functions $\{\zeta^k\}_k$ such that $\{\varepsilon^k\}_k$ converges to zero, $\{\zeta^k\}_k$ converges to σ , and ζ^k is a Nash equilibrium of $\Gamma_{\varepsilon, \sigma^k}$ for every k .

Using standard arguments, De Sinopoli and Pimienta (2009) show that every Poisson game has an inner-perfect equilibrium and that the usual alternative definitions (based on, e.g., ε -perfect equilibria) are also equivalent in the context of Poisson games. It is also showed there that, contrary to well known results for normal form games, inner-perfect equilibria can be in dominated strategies. The following example illustrates why.

EXAMPLE 1.1. Let Γ be a Poisson game with expected number of players equal to $n = 2$, set of types $\mathcal{T} = \{1\}$, set of actions $C = \{a, b\}$, and utility function

$$u(a, x) = e^{-2} \quad \text{for every } x \in Z(C),$$

$$u(b, x) = \begin{cases} 1 & \text{if } x(a) = x(b) = 1, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Notice that e^{-2} is the probability that $x(a) = x(b) = 1$ under the strategy $\sigma = (\frac{1}{2}a + \frac{1}{2}b)$. Also notice that action b is dominated by action a , the former only does as good as the latter against the strategy $\sigma = (\frac{1}{2}a + \frac{1}{2}b)$, and does strictly worse for any other strategy $\sigma' \neq \sigma$. The action $\gamma = \frac{1}{2}a + \frac{1}{2}b$ is also dominated by a . Nevertheless, it is a best response against σ . Finally, since σ is completely mixed, we can conclude that the dominated strategy σ is an inner-perfect equilibrium.

In order to see where the difference with respect to normal form games is coming from, it is useful to plot how the players' utility varies as the opponents change their behavior. We do that in Figure 1.1, where we represent utilities with respect to the probability attached to action a by an average member of the population. (There is only one type of player so, in this example, the sets Σ and $\tau(\Sigma)$ coincide.) The first thing to notice is that $U(b, \cdot)$ is not linear in Σ and that it attains its maximum at the completely mixed strategy $\sigma = (\frac{1}{2}a + \frac{1}{2}b)$. At that point, $U(a, \cdot)$ coincides with $U(b, \cdot)$. If we were to integrate $U(a, \cdot)$ and $U(b, \cdot)$ over the domain of strategies, the integral of $U(a, \cdot)$ would always be larger than the integral of $U(b, \cdot)$. Of course, not only is this true when we integrate with respect to the Lebesgue measure, but also when we integrate with respect to any Borel probability measure that does not give probability one to $\{\sigma\}$. Hence, if we approach σ by an arbitrary sequence of "sufficiently mixed" Borel probability measures over Σ , action b would always be an inferior response to every element of such a sequence. In the next section we formalize and generalize this intuition.




FIGURE 1.1. Utility functions in Example 1.1.

1.4. Perfection

The set $\tau(\Sigma)$ is equipped with the Euclidean distance d , so $(\tau(\Sigma), d)$ is a compact metric space. The distance between τ and an arbitrary subset $A \subset \tau(\Sigma)$ is $d(\tau, A) \equiv \inf\{d(\tau, a) : a \in A\}$.

We let \mathcal{B} denote the σ -algebra of Borel sets in $\tau(\Sigma) \subset \Delta(C)$. The set of all Borel probability measures over the measurable space $(\tau(\Sigma), \mathcal{B})$ is denoted \mathcal{M} . We topologize \mathcal{M} with the weak* topology. This topology is characterized by the following: a sequence of measures $\{\mu^k\} \subset \mathcal{M}$ converges (weakly) to μ if for every continuous function $f : \tau(\Sigma) \rightarrow \mathbb{R}$ the sequence of real numbers $\int_{\tau(\Sigma)} f d\mu^k$ converges to $\int_{\tau(\Sigma)} f d\mu$. It can be showed (Billingsley (1968, pg. 239)) that \mathcal{M} is a compact metrizable space and that a sequence $\{\mu^k\}$ converges to μ if and only if it converges with respect to the Prokhorov metric.

Let $\delta : \tau(\Sigma) \rightarrow \mathcal{M}$ be the function that maps each $\tau \in \tau(\Sigma)$ to the Dirac measure $\delta(\tau) \in \mathcal{M}$ that assigns probability one to $\{\tau\}$. With abuse of notation, if $\sigma \in \Sigma$ we write $\delta(\sigma)$ instead of $\delta(\tau(\sigma))$. Denote by \mathcal{M}° the subset of measures $\mu \in \mathcal{M}$ that satisfy $\mu(O) > 0$ for every nonempty open set $O \subset \tau(\Sigma)$.

We extend the domain of the utility functions to \mathcal{M} :

$$\bar{U}_t(c, \mu) \equiv \int_{\tau(\Sigma)} U_t(c, \tau) d\mu.$$

PROPOSITION 1.1. *The utility functions $\bar{U}_t(c, \cdot) : \mathcal{M} \rightarrow \mathbb{R}$ are continuous and linear in \mathcal{M} .*

PROOF. Consider a sequence $\{\mu^k\} \rightarrow \mu$. Since the functions $U_t(c, \cdot)$ are continuous on $\tau(\Sigma)$, from weak convergence we obtain $\bar{U}_t(c, \mu^k) \rightarrow \bar{U}_t(c, \mu)$.

To prove linearity, take some $\alpha \in [0, 1]$, some $\mu, \mu' \in \mathcal{M}$, and note that

$$\begin{aligned} \overline{U}_t(c, \alpha\mu + (1-\alpha)\mu') &= \int_{\tau(\Sigma)} U_t(c, \tau) d(\alpha\mu + (1-\alpha)\mu') = \\ &= \alpha \int_{\tau(\Sigma)} U_t(c, \tau) d\mu + (1-\alpha) \int_{\tau(\Sigma)} U_t(c, \tau) d\mu' = \alpha \overline{U}_t(c, \mu) + (1-\alpha) \overline{U}_t(c, \mu'). \end{aligned}$$

□

REMARK 1.1. Recall that the utility functions $U_t(c, \cdot)$ are continuous but, typically, *not* linear in $\tau(\Sigma)$.

Given any $\mu \in \mathcal{M}$ we write $\overline{\text{PBR}}_t(\mu)$ for the set of actions $c \in C_t$ that maximize $\overline{U}_t(c, \mu)$. As usual, we also define the set of mixed actions $\overline{\text{BR}}_t(\mu) \equiv \Delta(\overline{\text{PBR}}_t(\mu))$. The sets $\overline{\text{PBR}}(\mu)$ and $\overline{\text{BR}}(\mu)$ are defined accordingly. The correspondence $\overline{\text{BR}}$ is upper semicontinuous and convex valued.

The following result follows directly from the definitions.

PROPOSITION 1.2. *The strategy function σ is a Nash equilibrium of the Poisson game Γ if and only if $\sigma \in \overline{\text{BR}}(\delta(\sigma))$.*

It is convenient to recast the definition of dominated actions using the extension of the utility functions to \mathcal{M} . We do so in the next proposition and state it without proof.

PROPOSITION 1.3. *Action α is dominated by β for players of type t if and only if $\overline{U}_t(\alpha, \mu) \leq \overline{U}_t(\beta, \mu)$ for every $\mu \in \mathcal{M}$ and $\overline{U}_t(\alpha, \mu') < \overline{U}_t(\beta, \mu')$ for some $\mu' \in \mathcal{M}$.*

Moreover, an action α is strictly dominated by β if the strict inequality holds for every $\mu \in \mathcal{M}$.

We are now in a position to characterize the set of dominated actions for a given type. The next theorem is reminiscent of classical results that hold in finite normal form games (see Gale and Sherman (1950); Bohnenblust et al. (1950); Pearce (1984)).

THEOREM 1.1. *An action $\alpha \in \Delta(C_t)$ is undominated for a player of type t if and only if there is a $\mu^\circ \in \mathcal{M}^\circ$ such that $\alpha \in \overline{\text{BR}}_t(\mu^\circ)$.*

PROOF. If there is a measure μ° that assigns positive probability to every open set in $\tau(\Sigma)$ and $\alpha \in \overline{\text{BR}}_t(\mu^\circ)$ then action α cannot be dominated.

Suppose now that $\alpha \notin \overline{\text{BR}}_t(\mu^\circ)$ for every $\mu^\circ \in \mathcal{M}^\circ$. Fix some $\rho^\circ \in \mathcal{M}^\circ$, some $0 < \varepsilon < 1$, and construct the infinite two-player zero-sum game $\Gamma(t, \alpha, \rho^\circ, \varepsilon) \equiv (\Delta(C_t), \mathcal{M}, V_{\alpha, \rho^\circ}^\varepsilon)$ where, for any $\beta \in \Delta(C_t)$ and $\mu \in \mathcal{M}$, player one's payoff function $V_{\alpha, \rho^\circ}^\varepsilon$ is given by:

$$V_{\alpha, \rho^\circ}^\varepsilon(\beta, \mu) \equiv \overline{U}_t(\beta, \varepsilon \rho^\circ + (1 - \varepsilon)\mu) - \overline{U}_t(\alpha, \varepsilon \rho^\circ + (1 - \varepsilon)\mu).$$

Let $(\beta^\varepsilon, \mu^\varepsilon)$ be a Nash equilibrium of $\Gamma(t, \alpha, \rho^\circ, \varepsilon)$.⁴ We have

$$0 = V_{\alpha, \rho^\circ}^\varepsilon(\alpha, \mu^\varepsilon) < V_{\alpha, \rho^\circ}^\varepsilon(\beta^\varepsilon, \mu^\varepsilon) \leq V_{\alpha, \rho^\circ}^\varepsilon(\beta^\varepsilon, \mu) \text{ for every } \mu \in \mathcal{M}. \quad (1.4.1)$$

The weak inequality follows from player two's Nash equilibrium conditions and the strict inequality follows because α is never a best response against any element in \mathcal{M}° . Hence, β^ε dominates α in the zero-sum game $\Gamma(t, \alpha, \rho^\circ, \varepsilon)$. Passing to a subsequence if necessary, consider the limit β^* of $\{\beta^\varepsilon\}$ as ε goes to zero. Define the function V_α as follows:

$$V_\alpha(\beta, \mu) \equiv \overline{U}_t(\beta, \mu) - \overline{U}_t(\alpha, \mu).$$

From (1.4.1) we know that $V_{\alpha, \rho^\circ}^\varepsilon(\beta^\varepsilon, \mu) > 0$ for every $\mu \in \mathcal{M}$. Hence, by continuity, $V_\alpha(\beta^*, \mu) \geq 0$ for every $\mu \in \mathcal{M}$. Now we only need to find $\mu' \in \mathcal{M}$ such that $V_\alpha(\beta^*, \mu') > 0$.

For ε small enough the carrier $\mathcal{C}(\beta^*)$ is a subset of the carrier $\mathcal{C}(\beta^\varepsilon)$, therefore, for such small values of ε we also have $V_{\alpha, \rho^\circ}^\varepsilon(\beta^*, \mu) > 0$ for every $\mu \in \mathcal{M}$.

⁴ There is always a Nash equilibrium. In particular, if d_p is the Prokhorov metric on \mathcal{M} , Prokhorov's Theorem implies that (\mathcal{M}, d_p) is a compact metric space because $(\tau(\Sigma), d)$ is also a compact metric space (Billingsley (1968, pg. 37)). The set \mathcal{M} is also a nonempty and convex subset of a normed vector space. Furthermore, the payoff function $V_{\alpha, \rho^\circ}^\varepsilon$ is linear in both arguments, making the associated best response correspondence convex valued (and upper semicontinuous). Existence of Nash equilibrium follows from the Fan-Glicksberg fixed point theorem.

Since, by definition, $V_\alpha(\beta^*, \varepsilon\rho^\circ + (1-\varepsilon)\mu) = V_{\alpha, \rho^\circ}^\varepsilon(\beta^*, \mu)$, we can conclude that α is dominated by β^* in the original Poisson game. \square

This result implies that a definition of perfection that guarantees that players do not play dominated actions needs to be based on elements of the set \mathcal{M}° . Hence, we define a *perturbation* as a pair $(\varepsilon, \mu^\circ) \in (0, 1) \times \mathcal{M}^\circ$. The interpretation is that with vanishing probability ε , the average behavior of the population is perturbed towards the completely mixed measure μ° . Thus, a Nash equilibrium of such a perturbed game is a strategy function σ that satisfies $\sigma \in \overline{\text{BR}}((1-\varepsilon)\delta(\sigma) + \varepsilon\mu^\circ)$. Moreover, a strategy function satisfies this property if and only if it is a Nash equilibrium of a suitably defined utility-perturbed Poisson game.

Given a Poisson game $\Gamma = (n, \mathcal{T}, r, C, (C_t)_{t \in \mathcal{T}}, u)$ and a perturbation (ε, μ°) we define the *perturbed Poisson game* $\Gamma_{\varepsilon, \mu^\circ} \equiv (n, \mathcal{T}, r, C, (C_t)_{t \in \mathcal{T}}, u(\cdot | \varepsilon, \mu^\circ))$ where the utility functions are given, for every type $t \in \mathcal{T}$ and every action $c \in C_t$, by

$$u_t(c, x | \varepsilon, \mu^\circ) \equiv (1-\varepsilon)u_t(c, x) + \varepsilon \int_{\tau(\Sigma)} U_t(c, \tau) d\mu^\circ. \quad (1.4.2)$$

PROPOSITION 1.4. *Given a perturbation (ε, μ°) , the strategy function σ is a Nash equilibrium of $\Gamma_{\varepsilon, \mu^\circ}$ if and only if $\sigma \in \overline{\text{BR}}((1-\varepsilon)\delta(\sigma) + \varepsilon\mu^\circ)$.*

PROOF. Just notice that for every $t \in \mathcal{T}$ and every $c \in C_t$,

$$\begin{aligned} \overline{U}_t(c, (1-\varepsilon)\delta(\sigma) + \varepsilon\mu^\circ) &= (1-\varepsilon)\overline{U}_t(c, \delta(\sigma)) + \varepsilon\overline{U}_t(c, \mu^\circ) \\ &= (1-\varepsilon)U_t(c, \tau(\sigma)) + \varepsilon \int_{\tau(\Sigma)} U_t(c, \tau) d\mu^\circ \\ &= (1-\varepsilon) \sum_{x \in Z(C)} \mathbf{P}(x | \tau(\sigma)) u_t(c, x) + \varepsilon \int_{\tau(\Sigma)} U_t(c, \tau) d\mu^\circ \\ &= \sum_{x \in Z(C)} \mathbf{P}(x | \tau(\sigma)) \left[(1-\varepsilon)u_t(c, x) + \varepsilon \int_{\tau(\Sigma)} U_t(c, \tau) d\mu^\circ \right] \\ &= \sum_{x \in Z(C)} \mathbf{P}(x | \tau(\sigma)) u_t(c, x | \varepsilon, \mu^\circ) = U_t(c, \tau(\sigma) | \varepsilon, \mu^\circ). \end{aligned}$$

\square

Note that, given a perturbation (ε, μ°) , we can first normalize utility functions in $\Gamma_{\varepsilon, \mu^\circ}$ by dividing them by $(1-\varepsilon)$ and think of the perturbation as adding,

for each type $t \in \mathcal{T}$ and each action $c \in C_t$, the constant value $\frac{\varepsilon}{1-\varepsilon} \int_{\tau(\Sigma)} U_t(c, \tau) d\mu^\circ$ to the function $U_t(c, \cdot)$. In Example 1.1, for instance, for any perturbation (ε, μ°) the value that is added to $U_t(a, \cdot)$ by the perturbation in the corresponding perturbed Poisson game is always strictly larger than the value added to $U_t(b, \cdot)$ (see Figure 1.1). This “lifts” the expected utility function $U_t(a, \cdot)$ more than $U_t(b, \cdot)$ and makes action b strictly dominated in the perturbed Poisson game. Note as well that, given a Poisson game Γ , the set of all perturbed Poisson games (as defined above) is a strict subset of the set of Poisson games that can be generated by perturbing the utility functions in Γ .

Taking the perturbations to zero, we introduce a new definition of perfection for Poisson games.

DEFINITION 1.5 (Outer-perfection). The strategy function σ is an *outer-perfect equilibrium* if there is a sequence of perturbations $\{(\varepsilon^k, \mu^k)\}_k$ and a sequence of strategy functions $\{\sigma^k\}_k$ such that $\{\varepsilon^k\}_k$ converges to zero, $\{\sigma^k\}_k$ converges to σ , and σ^k is a Nash equilibrium of $\Gamma_{\varepsilon^k, \mu^k}$ for every k .

Every perturbed Poisson game has a Nash equilibrium. For any sequence of Poisson games we can construct an associated sequence of Nash equilibria. Such a sequence is contained in the compact set Σ so it has a subsequence that converges. Hence, every Poisson game has an outer-perfect equilibrium. (Furthermore, it can also be proved that if the sequence of strategies $\{\sigma^k\}_k$ supports an outer-perfect equilibrium σ given the sequence of perturbations $\{(\varepsilon^k, \mu^k)\}_k$ then the sequence of perturbed equilibria $\{(1 - \varepsilon^k)\delta(\sigma^k) + \varepsilon^k \mu^k\}_k$ converges weakly to $\delta(\sigma)$.) The major difference between this concept and the usual implementation of perfection in finite or infinite normal form games (Selten (1975); Simon and Stinchcombe (1995); also, inner-perfect equilibria in the present paper) is that the set of perturbations *is not* a subset of the set of mixed strategies in the game. Thus, a perturbed Poisson game *cannot* be interpreted as a game where players make mistakes when implementing their intended actions. We adhere to the view that, if the possibility of mistakes is real then

it should be properly modeled in the game (Kohlberg and Mertens (1986, p. 1005, footnote 3)). Perturbations are just technical devices used to obtain desirable game theoretical properties. We complete this argument presenting in Appendix B an example that illustrates the inadequacy of a concept of strategic stability based on inner-perfect equilibria (which is the concept that, in the current context, does admit a motivation based on players' mistakes when playing the game). We also show at the end of this section that outer-perfection neither implies nor is implied by inner-perfection.

The following two corollaries follow from Theorem 1.1.

COROLLARY 1.1. *Every outer-perfect equilibrium is a Nash equilibrium in undominated strategies.*

PROOF. Theorem 1.1 implies that, for any perturbation (ε, μ°) , every dominated action in the Poisson game Γ becomes strictly dominated in $\Gamma_{\varepsilon, \mu^\circ}$. Hence, it is used with probability zero in every Nash equilibrium of $\Gamma_{\varepsilon, \mu^\circ}$. Since an outer-perfect equilibrium is the limit point of a sequence of Nash equilibria of perturbed Poisson games, such a dominated action is also used with probability zero in any outer-perfect equilibrium of Γ . \square

COROLLARY 1.2. *If $\#\mathcal{T} = 1$ then every undominated equilibrium is outer-perfect.*

PROOF. Let σ be an undominated equilibrium. Because σ is an equilibrium, $\sigma \in \overline{\text{BR}}(\delta(\sigma))$. Because σ is undominated, there is a measure $\mu^\circ \in \mathcal{M}^\circ$ such that $\sigma \in \overline{\text{BR}}(\mu^\circ)$. Therefore, Proposition 1.1 implies that $\sigma \in \overline{\text{BR}}((1-\varepsilon)\delta(\sigma) + \varepsilon\mu^\circ)$. Taking ε to zero proves the result. \square

In the next example we show that Corollary 1.2 does not generalize to Poisson games with more than two types.

EXAMPLE 1.2. Take a Poisson game with expected number of players $n = 2$, set of types $\mathcal{T} = \{1, 2\}$, and set of actions $C = C_1 = C_2 = \{a, b\}$. The probability of




FIGURE 1.2. Utility functions in Example 1.2.

each type is $r_1 = 2/3$ and $r_2 = 1/3$. Utility functions are as follows:⁵

$$u_1(a, x) = \begin{cases} 1 & \text{if } x(a) = x(b) = 1, \\ 0 & \text{otherwise,} \end{cases} \quad u_2(a, x) = e^{-2} \quad \text{for every } x,$$

$$u_1(b, x) = (2 - x(a))e^{-2}, \quad u_2(b, x) = (2 - x(a))e^{-2}.$$

The corresponding expected utility functions U_1 and U_2 are plotted in Figure 1.2. As we can see, no type has a dominated action. It is easy to see that the strategy function $\sigma = (\frac{3}{4}a + \frac{1}{4}b, b)$ is an undominated Nash equilibrium such that $\tau(\sigma) = \frac{1}{2}a + \frac{1}{2}b$. However, it is not outer-perfect. Given that $U_1(a, \cdot)$ is always below $U_2(a, \cdot)$ and that $U_1(b, \cdot) = U_2(a, \cdot)$, for any $\mu \in \mathcal{M}^\circ$ such that type 1 players are indifferent between a and b , necessarily, players of type 2 strictly prefer a to b . Hence, we cannot construct a sequence of perturbed Poisson games whose associated sequence of Nash equilibria converges to σ . (In turn, the Nash equilibrium $(\frac{1}{4}a + \frac{3}{4}b, a)$ is indeed outer-perfect.)

We now explore further the relationship between inner-perfect and outer-perfect equilibria. We have already seen above that an inner-perfect equilibrium can be a dominated strategy. Therefore, not every inner-perfect equilibrium is outer-perfect. We can also easily illustrate this last fact here with the strategy function $\sigma = (\frac{3}{4}a + \frac{1}{4}b, b)$ in Example 1.2. Indeed, the sequence

⁵ Note, however, that the utility functions u_i are *not* bounded, contrary to our assumption when we defined Poisson games. We chose unbounded utility functions only for the sake of simplicity in the exposition of the result.




FIGURE 1.3. Utility functions in Example 1.3.

of completely mixed strategies $\sigma^\varepsilon = ((\frac{3-2\varepsilon}{4})a + (\frac{1+2\varepsilon}{4})b, \varepsilon a + (1-\varepsilon)b)$ converges to σ . Given that $\tau(\sigma^\varepsilon) = \tau(\sigma)$ for every small enough ε , the strategy function σ is a best response against every element in such a sequence. Thus, σ is an inner-perfect equilibrium (in undominated strategies).

On the other hand, as we show in the next example, not every outer-perfect equilibrium is inner-perfect.

EXAMPLE 1.3. Let the Poisson game Γ have expected number of players equal to $n = 4$, only one type, set of actions $C = \{a, b\}$ and utility function:

$$u(a, x) = \begin{cases} 2 & \text{if } x(a) = 0, \\ 8 & \text{if } x(a) = 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$u(b, x) = 2 \quad \text{for every } x.$$

We represent the corresponding utility functions $U(a, \cdot)$ and $U(b, \cdot)$ in Figure 1.3.⁶ Strategies a and b are both undominated. Furthermore, (b) is a Nash equilibrium of the game and, since $\#\mathcal{T} = 1$, Corollary 1.2 implies that it is also an outer-perfect equilibrium. However, it is not inner-perfect as for any completely mixed strategy close to strategy (b) action a is strictly preferred to action b .

⁶ An analogous picture can be obtained from a Poisson model of a congestion problem such as the Farol Bar game proposed by Arthur (1994). Each agent has two alternatives: drinking a beer at home (action b) or at a bar (action a). The utility of drinking in the bar alone is the same as the one from drinking at home. Furthermore, the utility of drinking in the bar is increasing in the company up to a point where the bar is too crowded and it starts to decline.

We summarize these observations in the next proposition.

PROPOSITION 1.5. *A Nash equilibrium in undominated actions is not necessarily outer-perfect even if it is also an inner-perfect equilibrium. Moreover, an outer-perfect equilibrium is not necessarily inner-perfect.*

1.5. Stability

In the following example, we show that the process of iterated deletion of dominated actions can lead to different solutions depending on the order of elimination.

EXAMPLE 1.4 (Iterated dominance). This example shows why iterated dominance and existence force us to use a set valued solution concept. Consider a Poisson game with set of types $\mathcal{T} \equiv \{1, 2\}$ with probabilities $r_1 = 1/4$ and $r_2 = 3/4$, and set of actions $C_1 = C_2 = C \equiv \{a, b, c, d\}$. Preferences are given by the following utility functions.

$$\begin{aligned}
 u_1(a, x) &= \begin{cases} 2 & \text{if } x(a) \geq x(b), \\ 0 & \text{otherwise,} \end{cases} & u_2(a, x) &= \begin{cases} 1 & \text{if } x(c) > x(d), \\ 0 & \text{otherwise,} \end{cases} \\
 u_1(b, x) &= \begin{cases} 1 & \text{if } x(b) \geq x(a), \\ 0 & \text{otherwise,} \end{cases} & u_2(b, x) &= \begin{cases} 1 & \text{if } x(d) > x(c), \\ 0 & \text{otherwise,} \end{cases} \\
 u_1(c, x) &= 0 \text{ for all } x \in Z(C), & u_2(c, x) &= 0 \text{ for all } x \in Z(C), \\
 u_1(d, x) &= -1 \text{ for all } x \in Z(C). & u_2(d, x) &= -1 \text{ for all } x \in Z(C).
 \end{aligned}$$

Actions c and d are dominated for both types. If we first eliminate d from C_1 and C_2 then b is dominated for type 2 players. Eliminating b from C_2 and c from C_1 and C_2 , we see that at least $3/4$ of the population choose a . Correspondingly, the best action for players of type 1 for every remaining strategy of the population is to also play a . We obtain that (a, a) survives the process of iterated deletion of dominated actions.

On the other hand, if we first eliminate c from C_1 and C_2 then a is dominated for type 2 players. We can eliminate a for type 2 players and d for every player in the game to conclude that at least $3/4$ of the population choose b . Provided the expected number of players n is large enough, choosing b dominates choosing a for players of type 1. With this order of elimination of dominated actions, only the strategy function (b, b) survives. Note that the two equilibria that we obtain through the process of iterated deletion of dominated actions also induce different expected utility to the players in the game.

Thus, if we want to provide a definition of equilibrium that is robust against iterated deletion of dominated actions we are led to define a set-valued concept. In the previous example, e.g., such an equilibrium concept would have to include both (a, a) and (b, b) .

Following Kohlberg and Mertens (1986) we say that a set of equilibria is *stable* if it is minimal with respect to the following property:

PROPERTY (S). $S \subset \Sigma$ is a closed set of Nash equilibria of Γ satisfying: for any $\varepsilon > 0$ there is a $\bar{\eta} > 0$ such that for any perturbation (η, μ°) with $0 < \eta < \bar{\eta}$ we can find a σ that is ε -close to S and satisfies $\sigma \in \overline{\text{BR}}((1 - \eta)\delta(\sigma) + \eta\mu^\circ)$.

REMARK 1.2. Property (S) in Kohlberg and Mertens (1986) requires that every close by strategy perturbed game have a Nash equilibrium close to S . Taking into account the space of perturbations used here (see page 12), we instead directly perturb the best response correspondence by perturbing the aggregate behavior of the population. As a result, Property (S) here requires that every close by perturbed correspondence have a fixed point close to S . Of course, Proposition 1.4 implies that such a fixed point is a Nash equilibrium of a close by (payoff) perturbed Poisson game.

PROPOSITION 1.6. *Every Poisson game has a stable set.*

Existence of stable sets in Poisson games is a particular case of a more general existence result. Stable sets in Poisson games are an example of *Q-robust*




FIGURE 1.4. Utility functions in a game with a disconnected stable set.

sets of fixed points (McLennan (2012, Definition 8.3.5)). Loosely speaking, a set of fixed points X of a correspondence F is *essential* if every correspondence “close” to F has a fixed point close to X . This means that X is stable against every small perturbation of F . For instance, we can show that the set of *all* Nash equilibria of a (Poisson) game is essential. We can weaken this concept and restrict the set of allowed perturbations to those belonging to some class Q . (In our case, the characterization given in Theorem 1.1 indicates that we only consider perturbations caused by altering the average behavior of the population towards some $\mu^\circ \in \mathcal{M}^\circ$.) Then we say that a set of fixed points X of a correspondence F is *Q-robust* if every correspondence “close” to F that can be obtained through a perturbation in Q has a fixed point close to X . McLennan (2012) shows that, if F is an upper semicontinuous and closed valued correspondence, every Q -robust set contains a minimal Q -robust set and that every connected Q -robust set contains a minimal connected Q -robust set (Theorem 8.3.8). However, not every stable set is necessarily connected.

Indeed, let us modify the utility functions in Example 1.3 so that the expected utility functions are those depicted in Figure 2.1. (Utility functions $u(a, \cdot)$ and $u(b, \cdot)$ can be found that generate such $U(a, \cdot)$ and $U(b, \cdot)$.) In this new game, $U(a, \cdot)$ and $U(b, \cdot)$ coincide in three isolated points. Furthermore, there are 2 stable sets, $\{\sigma^*\}$ and $\{(a), (b)\}$. Not only is the latter stable set disconnected but also its members belong to different connected components of Nash equilibria. Hence, $\{\sigma^*\}$ is the unique connected stable set.

We now prove that stable sets satisfy admissibility.

PROPOSITION 1.7. *Every point of a stable set is an outer-perfect, hence, undominated, equilibrium.*

PROOF. Let S be a stable set and let $\sigma \in S$ be a strategy function that is not an outer-perfect equilibrium. Therefore, there is some $\bar{\varepsilon} > 0$ and some $\bar{\eta} > 0$ such that for every $\eta < \bar{\eta}$ and every μ° we have $\zeta \notin \overline{\text{BR}}((1-\eta)\delta(\zeta) + \eta\mu^\circ)$ whenever $d(\sigma, \zeta) < \bar{\varepsilon}$. This implies that no strategy in the open ball $B(\sigma, \bar{\varepsilon}/2) \equiv \{\zeta : d(\sigma, \zeta) < \bar{\varepsilon}/2\}$ is an outer-perfect equilibrium either. Thus, it follows that $S \setminus \{B(\sigma, \bar{\varepsilon}/2)\}$ satisfies Property **(S)**, so either it is a stable set or it contains one. By minimality, S is not a stable set. \square

Furthermore, stable sets are robust to elimination of dominated actions in the following sense:

PROPOSITION 1.8. *A stable set contains a stable set of any game obtained by deletion of a dominated pure action or a pure action that is an inferior response to any strategy function in the stable set.*

PROOF. Take a stable set S of the Poisson game Γ . Let $c \in C_t$ be an action that is either dominated for players of type t or satisfies $c \notin \text{BR}_t(\tau(\sigma))$ for every $\sigma \in S$. Let $\tilde{\Gamma}$ be the reduced game obtained from Γ by deleting c from C_t . We know that $\sigma \in S$ implies $\sigma_t(c) = 0$. Therefore, every strategy function in S can be considered as a strategy function in the smaller game $\tilde{\Gamma}$. Let $\tilde{\Sigma} \subset \Sigma$ be the resulting space of mixed strategies, so that $\tau(\tilde{\Sigma}) \subset \tau(\Sigma)$. Furthermore, let $\tilde{\mathcal{M}}$ be the set of Borel measures on $\tau(\tilde{\Sigma})$.

Fix ε and choose an η as in Property **(S)**. Consider the measures $\mu \in \mathcal{M}^\circ$ and $\tilde{\mu} \in \tilde{\mathcal{M}}^\circ$. For any $0 < \kappa < 1$, we have $\mu^\kappa \equiv \kappa\mu + (1-\kappa)\tilde{\mu} \in \mathcal{M}^\circ$. Hence, there is a σ^κ that is ε -close to S such that $\sigma^\kappa \in \overline{\text{BR}}((1-\eta)\delta(\sigma^\kappa) + \eta\mu^\kappa)$. Taking the limit as κ approaches zero gives us, by continuity, a strategy function $\tilde{\sigma}$ that is ε -close to S and satisfies $\tilde{\sigma} \in \overline{\text{BR}}((1-\eta)\delta(\tilde{\sigma}) + \eta\tilde{\mu})$. We conclude that S satisfies Property **(S)** in $\tilde{\Gamma}$. Thus, either S is a stable set of $\tilde{\Gamma}$ or it contains one. \square

REMARK 1.3. Robustness against elimination of inferior responses has been used to formalize *forward induction* (Kohlberg and Mertens (1986); Mertens (1989)). Once we fix a solution of the game, players should consider as “certain not to be employed” those behaviors of the opponents in which some players use an inferior response against every member of the solution. If we accept that, we can ask that the solution be robust against the deletion of such inferior responses (Kohlberg (1990, p. 13); Hillas and Kohlberg (2002, p. 1645)). A similar reasoning may be used in Poisson games. Therefore, in this sense, we can say that stable sets in a Poisson game satisfy forward induction.

1.6. Examples

We now compute the stable sets of the past examples. In Example 1.1 the only undominated action is a so, by admissibility, $\{(a)\}$ is the unique stable set.

In Example 1.2 the strict equilibrium (b, b) is, of course, a singleton stable set. The strategy function (a, a) is a Nash equilibrium such that, for every small perturbation, the corresponding perturbed Poisson game has a Nash equilibrium close to (a, a) . To see this note that action a is a strict best response for type 2 players. For players of type 1, if a perturbation “lifts” $U(a, \cdot)$ more than $U(b, \cdot)$ then a is a strict best response in the perturbed game. (See the discussion following Equation (1.4.2).) On the other hand, if a perturbation “lifts” $U(b, \cdot)$ more than $U(a, \cdot)$ then both functions cross at some point close to $\tau(a) = 1$. Hence, $\{(a, a)\}$ is also a stable set. Finally, we can also see that the strategy function $(\frac{1}{4}a + \frac{3}{4}b, a)$ is *strictly outer-perfect* and, consequently, also a singleton stable set.⁷

The Poisson game in Example 1.3 (see also the game described in footnote 6) has two Nash equilibria that are also outer-perfect, the pure strategy (b) and a

⁷ A strictly outer-perfect equilibrium of a Poisson game Γ is a Nash equilibrium σ^* with the property that *every* perturbed Poisson game sufficiently close to Γ has a Nash equilibrium close to σ^* .

mixed strategy σ^* that satisfies $U(a, \sigma^*) = U(b, \sigma^*)$. The set $\{(b)\}$ is not stable because for those perturbations that, in the resulting perturbed Poisson game, “lift” $U(a, \cdot)$ more than $U(b, \cdot)$ there is no Nash equilibrium close to (b) . In turn, $\{\sigma^*\}$ is clearly stable.

Example 1.4 has a one dimensional and connected set of Nash equilibria that goes from (a, a) to (b, b) . (In every point of this set type 2 players are indifferent between a and b and, in a subset of it, type 1 players are also indifferent between a and b). We already argued that a stable set must contain (a, a) and (b, b) . In both equilibria type 1 players play a strict best response. Given that neither a nor b are dominated, there are close by perturbed games that “lift” $U_2(a, \cdot)$ more than $U_2(b, \cdot)$ as well as perturbed games that “lift” $U_2(b, \cdot)$ more than $U_2(a, \cdot)$. The strategy function (a, a) is a Nash equilibrium of every game in the first class of perturbed games while (b, b) is a Nash equilibrium of every game in the second class of perturbed games. Therefore, by minimality, the unique stable set is $\{(a, a), (b, b)\}$.

We conclude this section analyzing a variation of the referendum example proposed at the beginning of Section 3.

EXAMPLE 1.5 (A voting example). There is a referendum where voters have to vote either *yes* or *no* to some new policy and at least $K > 1$ voters should vote *yes* for the policy to be implemented. For concreteness, let us assume $K = 2$ and that the expected number of players is 4. Every voter prefers the new policy to the status quo. Let us fix players’ payoff from the outcome of the election equal to 1 if the policy is implemented and equal to 0 if it is not. Suppose further that there are two types of voters. Type 1 voters incur a cost $c = \frac{1}{2}e^{-\frac{1}{2}} (\approx 0.3)$ if they vote *yes* whereas type 2 voters do not have any cost of voting. Let the probability that a player is of type 2 be equal to $\frac{1}{8}$.

Voting *yes* is a weakly dominant action for type 2 players. They are only indifferent between *yes* and *no* under the strategy function (no, no) . The expected utility functions of type 1 players are depicted in Figure 1.5. Given the utility




FIGURE 1.5. Utility functions for type 1 players in Example 1.5.

values chosen, $U_1(no, \cdot)$ represents the probability that *two or more* voters vote *yes* for each value of $\tau(yes)$. On the other hand, $U_1(yes, \cdot)$ is equal to the probability that *one or more* voters vote *yes* minus the cost c . The two functions cross in two isolated points, $\frac{1}{8}$ and $\tau^* (\approx 0.44)$. When $\tau(yes) < \frac{1}{8}$, few other voters are expected to vote *yes* and type 1 voters prefer to vote *no* because their probability of being pivotal is not enough to overcome their cost to voting *yes*. As the number of other voters who are expected to vote *yes* increases, the probability of being pivotal in the referendum increases and type 1 voters start preferring voting *yes* to voting *no*. When the number of other voters who are expected to vote *yes* grows even larger so that $\tau(yes) > \tau^*$, the probability that the other voters meet the threshold necessary for the policy to be implemented increases, making *yes* again an inferior response for type 1 voters.

Hence, this game has three isolated Nash equilibria. There is a dominated Nash equilibrium where every player votes *no*, a *low support* equilibrium where only type 2 players play *yes*, and a *high support* equilibrium where type 2 voters play *yes* and type 1 voters play *yes* with probability $\frac{8}{7}(\tau^* - \frac{1}{8})$. The last two Nash equilibria are both outer-perfect, however, the high support equilibrium is the unique stable set of the game. Indeed, consider the low support equilibrium and eliminate the dominated action *no* for type 2 players. In the reduced game at least $\frac{1}{8}$ of the population vote *yes*. Thus, if we only consider values of τ such that $\tau(yes) \geq \frac{1}{8}$, we obtain a picture similar to the one in Figure 1.3. In

the same fashion as in that example, it can be seen that there are close by perturbed games that do not have a Nash equilibrium close to the low support equilibrium. From this we conclude that the unique stable set of the game consists only of the high support equilibrium.

Appendix A. Stable Sets in Generic Poisson Games

We show that for generic Poisson games every Nash equilibrium is a singleton stable set. We do this in a similar fashion to Carbonell-Nicolau (2010) who uses Fort's Theorem (Fort (1951)) to show that, for some large families of infinite normal-form games, generic members are such that every Nash equilibrium is *essential*.⁸ We point out, however, that the same caveat that is usually raised upon this type of genericity results applies here. The examples of Poisson games that we find in applications are nongeneric: there typically is a non-injective function mapping action profiles to events (in the case of voting games, e.g., pivotal events) where utilities are defined instead.

Once we fix n, \mathcal{T}, r, C and $(C_t)_{t \in \mathcal{T}}$, a Poisson game is given by a function $u : \mathcal{T} \times C \times Z(C) \rightarrow \mathbb{R}$. Since \mathcal{T} and C are finite and $Z(C)$ is countable, we can see such a function u as a point in the space of all bounded sequences ℓ^∞ . Thus, the Nash equilibrium correspondence NE can be thought of as $\text{NE} : \ell^\infty \rightarrow \Sigma$. Such a correspondence is upper semicontinuous and compact valued.

Recall that a G_δ set is a countable intersection of open sets. A topological space is called a *Baire space* if the union of any countable collection of closed sets with empty interior has empty interior. Since $(\ell^\infty, \|\cdot\|_\infty)$ is a Banach space,⁹ the Baire Category Theorem implies that it is also a Baire space.

THEOREM 1.2 (Fort (1951)). *If $F : X \rightarrow Y$ is an upper semicontinuous and compact valued correspondence from a Baire space to a metric space then F is both upper and lower semicontinuous at every point of a dense G_δ subset of X .*

At a lower semicontinuity point u of the Nash equilibrium correspondence, for every Nash equilibrium σ of u and every sequence of Poisson games $\{u^k\}_k$ converging to u , there is an associated sequence $\{\sigma^k\}_k$ converging to σ such that

⁸ A Nash equilibrium σ of a game Γ is essential (Wu Wen-Tsun (1963)) if every game close to Γ has a Nash equilibrium close to σ . Of course, an essential Nash equilibrium is a singleton stable set.

⁹ In our context, we can write $\|u\|_\infty = \max_{t \in \mathcal{T}} \max_{c \in C_t} \sup_{x \in Z(C)} |u_t(c, x)|$.

σ^k is a Nash equilibrium of u^k for every k . This is, in particular, true for every sequence of perturbed Poisson games converging to u . We thus conclude:

COROLLARY 1.3. *For every game in a dense G_δ set of Poisson games every Nash equilibrium is a singleton stable set.*

Appendix B. On the Inadequacy of Inner-Perfection

In the following example we illustrate why a definition of stability based on inner-perfect equilibrium perturbations is not adequate even if it is accompanied by a restriction that only allows to select undominated actions.

EXAMPLE 1.6. Consider the Poisson game Γ with $n = 2$, only one type, set of actions $C = \{a, b, c\}$, and utility function

$$\begin{aligned} u(a, x) &= \begin{cases} x(c) + 1 & \text{if } x(a) + x(c) = 1, \\ x(c) & \text{otherwise,} \end{cases} \\ u(b, x) &= e^{-1} && \text{for all } x \in Z(C), \\ u(c, x) &= -1 && \text{for all } x \in Z(C). \end{aligned}$$

Action $\gamma = \frac{1}{2}a + \frac{1}{2}b$ is not dominated: it does better than action a against the strategy (b) and it does better than action b against the strategy $(\frac{1}{2}b + \frac{1}{2}c)$. For any strategy-perturbation $(\varepsilon, \sigma^\circ)$ such that ε is close enough to zero, consider the strategy-perturbed game $\Gamma_{\varepsilon, \sigma^\circ}$ and the strategy

$$\zeta_{\varepsilon, \sigma^\circ} \equiv \frac{1}{1 - \varepsilon} \left(\left(\frac{1 - z(\varepsilon\sigma^\circ(c))}{2} - \varepsilon\sigma^\circ(c) - \varepsilon\sigma^\circ(a) \right) a + \left(\frac{1 + z(\varepsilon\sigma^\circ(c))}{2} - \varepsilon\sigma^\circ(b) \right) b \right),$$

where the correcting factor $z(\varepsilon\sigma^\circ(c)) > 0$ is chosen so that $P(1 | 1 - z(\varepsilon\sigma^\circ(c))) = e^{-1} - 2\varepsilon\sigma^\circ(c)$. The strategy $\zeta_{\varepsilon, \sigma^\circ}$ is an undominated Nash equilibrium of the strategy-perturbed game $\Gamma_{\varepsilon, \sigma^\circ}$ which is close to $\sigma = (\frac{1}{2}a + \frac{1}{2}b)$. (Note that under the corresponding perturbed strategy $\sigma_{\varepsilon, \sigma^\circ} \equiv (1 - \varepsilon)\zeta_{\varepsilon, \sigma^\circ} + \varepsilon\sigma^\circ$ the expected value of $x(c)$ is $2\varepsilon\sigma^\circ(c)$ and, therefore, $U(a, \sigma_{\varepsilon, \sigma^\circ}) = e^{-1}$.)

Hence, $\{(\frac{1}{2}a + \frac{1}{2}b)\}$ would be a stable set of the Poisson game according to a definition of stability based on strategy-perturbations (i.e. inner-perfection).

However, after eliminating the strictly dominated action c , action a and, consequently, action $\gamma = \frac{1}{2}a + \frac{1}{2}b$ become weakly dominated.

On the other hand, one can show that in this game the only stable set (according to our definition) is the set made of the strict equilibrium (b).

CHAPTER 2

The Structure of Nash Equilibria in Poisson Games*

2.1. Introduction

The geometric structure of Nash equilibria has been exploited to obtain several game-theoretical results. In particular, the graph of the Nash equilibrium correspondence of a finite game is a *semi-algebraic set*, i.e., it is defined by a finite system of polynomial inequalities. This fact allows the application of the semi-algebraic apparatus to the analysis of finite non-cooperative games. Kohlberg and Mertens (1986), for instance, show that the set of Nash equilibria of a game consists of a finite number of connected components. This follows because every semi-algebraic set has a finite triangulation, i.e. it is homeomorphic to a finite union of compact polyhedra (van der Waerden (1939), Satz 1, p. 123). Blume and Zame (1994) exploit another fundamental result of semi-algebraic geometry, the Tarski-Seidenberg theorem, to show that the perfect and sequential equilibrium correspondences have a semi-algebraic structure and to obtain the generic equivalence between these two equilibrium concepts. Moreover, the Generic Local Triviality theorem (Hardt (1980), Bochnak et al. (1987)) has been used to provide generic-finiteness results for equilibria of finite games (Govindan and McLennan (2001), Govindan and Wilson (2001)).

The objective of this paper is to examine the geometric structure of Nash equilibria in Poisson games and to analyze the corresponding game-theoretical consequences. Unfortunately, the same tools of semi-algebraic geometry that are useful to examine finite games cannot be used to obtain analogous results

* This chapter is based on a joint work with Carlos Pimienta (The University of New South Wales).

for Poisson games. In a Poisson game, the set of Nash equilibria is not defined by a system of polynomial inequalities precisely because utility functions are not polynomials. Nonetheless, we show that in a Poisson game utility functions are real-analytic and, correspondingly, the set of Nash equilibria is a *semi-analytic set* (that is, a set that can be written as the solution set to a finite system of analytic inequalities). The special structure of semi-analytic sets has a direct implication for Poisson games, viz., the finiteness of the number of connected components of equilibria. This result allows us to further develop the analysis of stable sets of Poisson games and, as long as each connected component maps into a unique outcome, implies that the set of equilibrium outcomes is finite.¹

In the next section, we review the general description of Poisson games. We discuss the geometric structure of the Nash equilibrium set in Section 2.3. We exploit such a structure in Section 2.4 to prove that stable sets in Poisson games satisfy the same version of connectedness as Kohlberg-Mertens stable sets.

2.2. Preliminaries

We adopt the same notation used in De Sinopoli et al. (2014), where the description of Poisson games closely follows the one introduced by Myerson (1998). We repeat it here for the sake of completeness and clarity.

We define a *Poisson game* as a tuple $\Gamma \equiv (n, \mathcal{T}, r, C, (C_t)_{t \in \mathcal{T}}, u)$. The number of *players* is a Poisson random variable with parameter n . Given n , the probability that there are k players in the game is therefore equal to

$$P(k | n) = \frac{e^{-n} n^k}{k!}.$$

¹In finite games, Kreps and Wilson (1982) prove that for each generic tree, the set of equilibrium outcomes is finite. In particular, this implies that each component of equilibria has a unique outcome distribution (Kohlberg and Mertens, 1986). In Poisson games, the generic finiteness of equilibrium outcomes is still an open problem.

The set $\mathcal{T} = \{1, \dots, T\}$ is the non-empty finite set of possible types of players. A player is of type $t \in \mathcal{T}$ with probability r_t . The probabilities that a player is of each type are listed in the vector $r = (r_1, \dots, r_T) \in \Delta(\mathcal{T})$.²

We let C be the finite set of *actions* and $C_t \subset C$ be the set of actions that are available to players of type t .³ An *action profile* $x \in Z(C)$ is a vector that specifies for each action $c \in C$ the number of players $x(c)$ who choose that action. The set of action profiles is $Z(C) \equiv \mathbb{Z}_+^C$. Players' preferences in the game are summarized by $u = (u_1, \dots, u_T)$, where each function $u_t : C_t \times Z(C) \rightarrow \mathbb{R}$ is assumed to be bounded. Then $u_t(c, x)$ denotes the payoff to a player of type t who chooses action c , when x is the profile of actions chosen by the rest of the population.

The set $\Delta(C_t)$ is the set of *mixed actions* for type t players. We identify the mixed action that assigns probability one to action $c \in C$ with the pure action c . A *strategy function* σ is an element of $\Sigma \equiv \{\sigma \in \Delta(C)^{\mathcal{T}} : \sigma_t \in \Delta(C_t) \text{ for all } t\}$, i.e. a mapping from the set of types to the set of mixed actions available to the corresponding type. We may sometimes refer to strategy functions simply as *strategies*. Let $\tau(\sigma) \in \Delta(C)$ be the ‘‘average’’ behavior induced by the strategy σ , which is given by $\tau(\sigma)(c) \equiv \sum_{t \in \mathcal{T}} r(t) \sigma_t(c)$. Moreover, we define the set $\tau(\Sigma) \equiv \{\bar{\tau} \in \Delta(C) : \bar{\tau} = \tau(\sigma) \text{ for some } \sigma \in \Sigma\}$. When the population's aggregate behavior is given by $\tau \in \tau(\Sigma)$, the probability that the action profile $x \in Z(C)$ is realized is equal to

$$\mathbf{P}(x \mid \tau) \equiv \prod_{c \in C} \left(e^{-n\tau(c)} \frac{(n\tau(c))^{x(c)}}{x(c)!} \right).$$

The expected payoff to a player of type t who plays action $c \in C_t$ is then

$$U_t(c, \tau) \equiv \sum_{x \in Z(C)} \mathbf{P}(x \mid \tau) u_t(c, x).$$

Note that, for each type $t \in \mathcal{T}$, each action $c \in C_t$ defines a bounded and continuous function $U_t(c, \cdot) : \Delta(C) \rightarrow \mathbb{R}$.

We recall now some further notation that will be needed in Section 2.4, in particular in the definition of stable sets.

² For any finite set S , we write $\Delta(S)$ for the set of probability distributions on S .

³ Given two sets A and B , we use the expression $A \subset B$ allowing for set equality.

The set $\tau(\Sigma)$ is equipped with the Euclidean distance d , so $(\tau(\Sigma), d)$ is a compact metric space. Given a subset $A \subset \tau(\Sigma)$, the distance between τ and A is $d(\tau, A) \equiv \inf\{d(\tau, a) : a \in A\}$. Let \mathcal{B} be the σ -algebra of Borel sets in $\tau(\Sigma) \subset \Delta(C)$. We denote by \mathcal{M} the set of all Borel probability measures over the measurable space $(\tau(\Sigma), \mathcal{B})$. Moreover, we denote by \mathcal{M}° the subset of measures $\mu \in \mathcal{M}$ that satisfy $\mu(O) > 0$ for every nonempty open set $O \subset \tau(\Sigma)$. Let $\delta : \tau(\Sigma) \rightarrow \mathcal{M}$ be the function that maps each $\tau \in \tau(\Sigma)$ to the Dirac measure $\delta(\tau) \in \mathcal{M}$ that assigns probability one to $\{\tau\}$. With abuse of notation, if $\sigma \in \Sigma$ we write $\delta(\sigma)$ instead of $\delta(\tau(\sigma))$.

The domain of the utility functions can be extended to \mathcal{M} :

$$\bar{U}_t(c, \mu) \equiv \int_{\tau(\Sigma)} U_t(c, \tau) d\mu.$$

Note that the functions $\bar{U}_t(c, \cdot) : \mathcal{M} \rightarrow \mathbb{R}$ are continuous and linear in \mathcal{M} .

2.3. The Set of Nash Equilibria

We start proving that the expected utility functions $U_t(c, \cdot) : \Delta(C) \rightarrow \mathbb{R}$ of a Poisson game are real analytic functions, i.e., functions that are locally given by a convergent power series.

LEMMA 2.1. *Given a Poisson game Γ , for every type $t \in \mathcal{T}$ and every choice $c \in C_t$ the expected utility function $U_t(c, \cdot)$ is real analytic.*

PROOF. If $\#C = K$ we first note that $U_t(c, \cdot)$ can be considered as a function $U_t(c, \cdot) : \mathbb{C}^K \rightarrow \mathbb{C}$. We show that $U_t(c, \cdot)$ is a complex analytic function, i.e. a complex function that is locally given by a convergent power series. The sum, product and composition of complex analytic functions are complex analytic, and the limit of a sequence of complex analytic functions that converges uniformly on a compact subset of the domain is complex analytic in that subset.⁴ Examples of complex analytic functions are polynomials and the exponential function,

⁴The latter result follows from the Weierstrass approximation theorem. This is in contrast with the situation in real analysis, where the limit of a sequence of real analytic functions that

hence, once we fix $n \in \mathbb{N}_+$ then, for any $y \in \mathbb{N}_+$, the function $z \mapsto e^{-nz}(nz)^y(y!)^{-1}$ is complex analytic because it is the product of compositions of complex analytic functions.

Let $C = \{c_1, \dots, c_K\}$, take some $x = (x(c_1), \dots, x(c_K)) \in \mathbb{N}_+^C$, and consider some number $u_t(c, x) \in \mathbb{R}$. The function

$$(\tau(c_1), \dots, \tau(c_K)) \mapsto \prod_{i=1}^K \left(e^{-n\tau(c_i)} \frac{(n\tau(c_i))^{x(c_i)}}{x(c_i)!} \right) u_t(c, x)$$

is also complex analytic because it is a composition of complex analytic functions.

We now index the countable set $Z(C)$ by taking some bijective function $\phi : \mathbb{N}_{++} \rightarrow Z(C)$ and letting $x^j \equiv \phi(j)$.

Consider some bounded function $u_t(c, \cdot) : Z(C) \rightarrow \mathbb{R}$. The function $U_t^m(c, \cdot)$ given by

$$(\tau(c_1), \dots, \tau(c_K)) \mapsto \sum_{j=1}^m \left(\prod_{i=1}^K e^{-n\tau(c_i)} \frac{(n\tau(c_i))^{x^j(c_i)}}{x^j(c_i)!} \right) u_t(c, x^j)$$

is complex analytic because it is the finite sum of complex analytic functions. The sequence of complex analytic functions $\{U_t^m(c, \cdot)\}_{m=1}^\infty$ is said to converge uniformly to $U_t(c, \cdot)$ on a compact subset T of \mathbb{C}^K if, for each $\varepsilon > 0$, there exists a $N \in \mathbb{N}_{++}$ such that

$$|U_t^m(c, \cdot) - U_t(c, \cdot)| < \varepsilon \text{ for all } m \geq N \text{ and } \tau \in T.$$

Let $T = \tau(\Sigma)$. Since we have

$$\lim_{j \rightarrow \infty} \prod_{i=1}^K e^{-n\tau(c_i)} \frac{(n\tau(c_i))^{x^j(c_i)}}{x^j(c_i)!} = 0$$

for every $(\tau(c_1), \dots, \tau(c_K)) \in \tau(\Sigma)$, the sequence $\{U_t^m(c, \cdot)\}_{m=1}^\infty$ converges uniformly to $U_t(c, \cdot)$ in $\tau(\Sigma)$ and, therefore, the expected utility function $U_t(c, \cdot)$ is complex analytic in $\tau(\Sigma)$. Thus, we can conclude that $U_t(c, \cdot)$ is real analytic when restricted to $\tau(\Sigma) \subseteq \mathbb{R}^K$. \square

converges uniformly on a compact subset of the domain is not necessarily real analytic in that subset. This is the reason why we extend the utility functions to the complex domain.

We can now show that the set of Nash equilibria of a Poisson game is a semi-analytic set, since the Nash equilibrium conditions can be written as a finite system of real analytic inequalities.

DEFINITION 2.1. A set $X \subset \mathbb{R}^n$ is *semi-analytic* if it is the finite union of sets of the form

$$\{x \in \mathbb{R}^n : f_1(x) = 0, \dots, f_k(x) = 0 \text{ and } g_1(x) > 0, \dots, g_m(x) > 0\}$$

where f_i and g_j are real analytic functions.

Of course, the class of semi-analytic sets includes also sets defined by weak inequalities. This class is closed under finite union, finite intersection, finite product and complementation (see Lojasiewicz (1964, 1965)). Note that, since any polynomial is an analytic function, f_i and g_j can be real polynomials.

Now, recall the definition of Nash equilibria in Poisson games:

DEFINITION 2.2. The strategy function $\sigma \in \Sigma$ is a *Nash equilibrium* of the Poisson game Γ if

$$U_t(\sigma_t, \tau(\sigma)) \geq U_t(\sigma'_t, \tau(\sigma)) \quad \text{for all } t \in \mathcal{T}, \sigma'_t \in \Delta(C_t).$$

Given a Poisson game Γ , $\sigma \in \mathbb{R}^{\mathcal{T}}$ is a Nash equilibrium of Γ if and only if it is a solution to the following system of real analytic equalities and inequalities:

$$\sum_{c_t \in C_t} \sigma_t(c_t) - 1 = 0 \quad \text{for all } t \in \mathcal{T}, \quad (2.3.1)$$

$$\sigma_t(c_t) \geq 0 \quad \text{for all } t \in \mathcal{T}, c_t \in C_t. \quad (2.3.2)$$

$$\sigma_t(c_t)[U_t(c_t, \tau(\sigma)) - U_t(k_t, \tau(\sigma))] \geq 0 \quad \text{for all } t \in \mathcal{T}, c_t, k_t \in C_t, \quad (2.3.3)$$

Note that conditions (2.3.1) and (2.3.2) define the set Σ of strategy functions which is, therefore, a semi-analytic set.⁵ Then, condition (2.3.3) ensures that, when every player is playing according to the strategy function σ , the action of each type gives positive probability only to the pure actions that are best

⁵In particular, the set Σ is semi-algebraic, since it is defined by a finite system of polynomial inequalities.

replies against the actions of the others. Thus, the third condition selects all the strategies σ that are best replies against themselves. It follows that the three conditions together define the set of all Nash equilibria of the Poisson game Γ , which is therefore semi-analytic.

LEMMA 2.2. *The set of Nash equilibria of a Poisson game Γ is a semi-analytic set.*

A property of semi-analytic sets that has an immediate implication for Poisson games is that the number of connected components of a relatively compact semi-analytic set is finite (Gabriélov (1968), Bierstone and Milman (1988)).⁶ Hence, we obtain:

THEOREM 2.1. *The set of Nash equilibria of any Poisson game has finitely many connected components.*

2.4. Connectedness

In standard games, a *connected component* of the equilibrium set (i.e., a maximal connected set of equilibria) does not substantially differ from a single-valued equilibrium inasmuch as, for any generic game, each component of equilibria generates a unique payoff distribution (Kohlberg and Mertens, 1986). In particular, this result follows from the generic finiteness of equilibrium payoffs (Kreps and Wilson, 1982, Theorem 2). Thus, for any generic extensive form game, all equilibria in the same connected component differ only out of the equilibrium path. This means that they are indistinguishable to an external observer that repeatedly watches agents playing the game, who can see only the probability distribution on terminal nodes.

The notion of stability defined by Kohlberg and Mertens (1986) does not satisfy connectedness, and a stable set may even contain points from different

⁶A *relatively compact* subset X of a topological space Y is a subset that has a compact closure.

connected components of the equilibrium set. This is a reason why alternative definitions of stability have been proposed (Hillas (1990); Mertens (1989, 1991)). However, Kohlberg-Mertens stable sets satisfy a weaker version of connectedness. Namely, every game has a stable set which is contained in a single connected component of the set of Nash equilibria.

Now, let us recall the definition of stable sets in Poisson games. We define a *perturbation* as a pair $(\varepsilon, \mu^\circ) \in (0, 1) \times \mathcal{M}^\circ$. The interpretation is that, with vanishing probability ε , the average behavior of the population is perturbed towards the completely mixed measure μ° . Hence, a Nash equilibrium of such a perturbed game is a strategy function σ that satisfies $\sigma \in \overline{\text{BR}}((1-\varepsilon)\delta(\sigma) + \varepsilon\mu^\circ)$, where the set $\overline{\text{BR}}(\mu)$ is defined in the usual way.

DEFINITION 2.3. (De Sinopoli et al. (2014)). A set of equilibria of a Poisson game Γ is *stable* if it is minimal with respect to the following property:

PROPERTY (S). $S \subset \Sigma$ is a closed set of Nash equilibria of Γ satisfying: for any $\varepsilon > 0$ there is a $\bar{\eta} > 0$ such that for any perturbation (η, μ°) with $0 < \eta < \bar{\eta}$ we can find a σ that is ε -close to S and satisfies $\sigma \in \overline{\text{BR}}((1-\eta)\delta(\sigma) + \eta\mu^\circ)$.

The fact that stable sets in Poisson games do not satisfy connectedness is shown in De Sinopoli et al. (2014, p. 58) by means of the following example. Consider a Poisson game with one type of players that can choose between two actions, a and b , with expected utility functions depicted in Figure 2.1. The game has two stable sets, σ^* and $\{(a), (b)\}$. The latter set is disconnected, and its elements belong to different connected components of the set of Nash equilibria.

We can exploit the analytic structure of Nash equilibria to prove that stable sets of Poisson games satisfy the same weaker version of connectedness as stable sets of standard games. This result is relevant per se as an enhancement of the analysis of stable sets. Furthermore, it implies that every game has a stable payoff as long as every connected component of equilibria maps into a




FIGURE 2.1. Utility functions in a Poisson game with a disconnected stable set.

unique outcome. As mentioned above, this is the case for generic finite games, while it is still an open problem to prove an analogous result for Poisson games.

Fix n , \mathcal{T} , r , C and $(C_t)_{t \in \mathcal{T}}$, and identify a Poisson game with its utility function $u : \mathcal{T} \times C \times Z(C) \rightarrow \mathbb{R}$.

THEOREM 2.2. *Given a Poisson game u , at least one of the connected components of its Nash equilibrium set is such that, for every Poisson game \hat{u} sufficiently close to u , there is a Nash equilibrium of \hat{u} close to this component.*

PROOF. Note that the whole set of fixed points N of an upper semi-continuous correspondence F is an *essential* collection of fixed points, i.e. every correspondence close to F has a fixed point close to N . Kinoshita (1952) shows that if a set of fixed points N is essential and N_1, \dots, N_k is a partition of N into disjoint compact sets, then some N_j is essential. Since we can choose such compact sets to be connected, we know that every Poisson game has a connected set of Nash equilibria that is essential. \square

Note that, given a Poisson game Γ , the set of the perturbed Poisson games in the definition of stable set is a strict subset of the set of Poisson games that can be generated by perturbing the utility functions in Γ . Thus, Theorem 2.2 implies the first part of the following corollary:

COROLLARY 2.1. *Every Poisson game has a stable set contained in a connected component of equilibria. Moreover, every Poisson game has a minimal connected set of Nash equilibria that satisfies Property (S).*

PROOF. Stable sets in Poisson games are an example of Q -robust sets of fixed points (McLennan, 2012, Definition 8.3.5). In broad terms, a set of fixed points X of a correspondence F is Q -robust if every correspondence close to F that can be obtained through a perturbation in some class Q has a fixed point close to X . (For stable sets in Poisson games, Q is the class of perturbations induced by altering the average behavior of the population towards some $\mu^\circ \in \mathcal{M}^\circ$.) McLennan (2012, Theorem 8.3.8) shows that, if F is an upper semi-continuous and closed valued correspondence, every Q -robust set contains a minimal Q -robust set and that every connected Q -robust set contains a minimal connected Q -robust set. \square

CHAPTER 3

Double Round-Robin Tournaments*

3.1. Introduction

Consider the following two-player game:

	<i>C</i>	<i>D</i>
<i>A</i>	6, 5	4, 3
<i>B</i>	4, 0	3, 1

If two individuals were matched to play such a game with the objective to maximize their expected payoff, they would play the unique Nash equilibrium $\{A, C\}$. Now, think of a situation in which the two individuals are matched twice to play simultaneously the same game in alternating roles, with the objective to maximize the sum of the payoffs in the two matches. If both of them play A when in the role of the first player and C when in the role of the second player, they both get a payoff of 11. None of them has an incentive to deviate, since any other choice leads to a lower payoff, and $\{(A, C), (A, C)\}$ is the unique equilibrium of the game. Consider the same situation with the difference that now a prize is awarded to the individual with the highest total payoff, and the two individuals want to maximize their probability of winning the prize. In the event of a tie, each of them would be the winner with probability $1/2$. Notice that in this case $\{(A, C), (A, C)\}$ is no longer an equilibrium. Indeed, if one of the two individuals is playing (A, C) , the other one has an incentive to deviate, for instance, to (B, C) , losing two points but lowering the payoff of his opponent by five points, becoming in this way the sole winner. It is easy to see that the

* This chapter is based on a joint work with Francesco De Sinopoli (University of Verona) and Carlos Pimienta (The University of New South Wales).

unique (symmetric) equilibrium of the game is $\{(B, D), (B, D)\}$ and both players win with probability $1/2$. Then, we can think of the same competitive situation extended to any number of individuals, where everyone meets all his opponents twice to play the same two-player game. Such a competition is called *double round-robin tournament*.

A double round-robin tournament is a simultaneous n -player game that is built on a two-player game g . Each player is matched with every other player twice and in every match the game g is played. In the two matches with the same opponent, a player plays once in the role of the first player and once in the role of the second player of g . Using sports terminology, we say that a player plays once “at home” and once “away”. The winner of the tournament is the player with the highest score, where a player’s score is the sum of the payoffs that he gets in all the matches he plays. This implies that players do not care about their absolute total score and maximize their probability of winning the tournament.

Arad and Rubinstein (2013) analyze tournaments of the *round-robin* type, where each player meets every other player once. In their analysis, the two-player game on which the tournament is built is a symmetric game. Moreover, the solution concept is that of a symmetric mixed strategy equilibrium, where a mixed strategy is executed only once and the player employs the resulting action in all his matches. Such a concept is appropriate to a situation in which individuals are drawn at random from a large population and are matched in pairs to play the same game anonymously. In this case, indeed, a mixed strategy can be interpreted as a distribution of actions in the population.

With our model, we aim to extend the analysis of tournaments to asymmetric games. Since any two individuals play the asymmetric game twice, the symmetry among players is restored. Hence, we also employ the solution concept of a symmetric Nash equilibrium, where players randomize independently at home and away. Following Arad and Rubinstein (2013), we assume that each player employs always the same action when he is playing at home, as well as

when he is playing away. Consider again the example at the beginning of this section. Under the assumptions of our model, the tournament with three players has a unique symmetric equilibrium in which every player chooses B at home and D away (the same as the two-individual case). For $n \geq 4$, instead, the only symmetric Nash equilibrium of the tournament prescribes to choose A when playing at home and C when playing away, which are the Nash equilibrium strategies of the two-player game. This is not a coincidence, as we will see when analyzing the relationship between equilibria of the tournament and equilibria of the two-player game on which it is built. Intuitively, as n grows, deviating from the equilibrium strategies of the base game becomes less and less profitable. Indeed, the loss inflicted to each of the other players becomes negligible with respect to the higher payoffs that the other players obtain in the increasing number of matches among them.

An alternative tournament model is the one in Laffond et al. (2000). As in our model, each player chooses one action and employs it in all his interactions. However, a player's payoff is given by the sum of the payoffs that he gets in all the (symmetric) games he plays, so players do care about their absolute total score.¹

Note the difference between the tournament model that we adopt and the classic model of contests. In the classic contest model, the players compete for a given prize by exerting an effort that increases their probability of winning (see, e.g., Green and Stokey (1983), Dixit (1987), Konrad (2009); among others). Each player's utility depends on his probability of winning, which is a function also of the other players' efforts, and on the cost of his own effort. In the tournament model, instead, the ranking of a player depends on the combination between his choice and the choices of all the other players, and actions are costless. A

¹The same assumption is examined in the example at the beginning, when individuals have the objective to maximize the sum of the payoffs in the two matches. In that case, the unique equilibrium prescribes the Nash equilibrium strategies of the base game. With the appropriate modifications, this is consistent with the results in Laffond et al. (2000).

particular contest model is that of the elimination tournament, which consists of several rounds in which individuals play pair-wise matches (see, e.g., Rosen (1986), Konrad (2004), Groh et al. (2012)). Differently from our model, the winner of a match advances to the next round of the tournament, while the loser is eliminated from the competition.

Double round-robins are common in several sports competitions, especially in those with a large number of matches per season. Most professional association football leagues in the world are based on a double round-robin,² as are most basketball leagues outside the United States.³ In such competitions, the assumptions of our model fit for instance the case of the teams, which have to make several choices to comply with throughout the entire tournament (e.g. the players, the coach, the home field).

Besides this straightforward interpretation, we can think also of another interpretation of tournaments. Namely, if we focus on the Nash equilibria of a two-player game g , we can use tournaments as an equilibrium refinement. As a matter of fact, a mixed-strategy Nash equilibrium of g can be interpreted as a stable distribution of pure strategies in a large population, where individuals play g over time and maximize their expected payoff. One might ask whether a Nash equilibrium of g can also be interpreted as a stable distribution of actions in situations where many individuals are matched to play g in alternating roles and maximize their probability of winning. It turns out that only some of the equilibria of g are “stable” in this sense. Thus, tournaments provide a refinement criterion, which selects all the equilibria of g that are limit points of equilibria of the tournament built on g as the number of players goes to infinity.

²Examples are the top european national leagues, like Spain’s La Liga, England’s Premier League, Germany’s Bundesliga, Italy’s Serie A. Moreover, double round-robins are used during the qualification phases of the FIFA World Cup and the respective continental leagues, and during the group phases of the UEFA Champions League and the Copa Libertadores de América.

³The qualification stages of the Euroleague are an example.

Finally, as Arad and Rubinstein (2013) point out, the analysis of the relationship between the equilibria of the tournament and the equilibria of the base game can be useful for experimental design. Indeed, the tournament structure has been used in some experiments to study the agents' behavior in the game g . Such a design has been criticized because, in the case in which a prize is awarded to the participant with the highest score, individuals may have different incentives from those in the base game. In this work, we aim to examine theoretically whether the equilibria of a game g are a good approximation of the equilibria of the tournament based on g , extending the analysis of Arad and Rubinstein (2013) to asymmetric games.

We describe the model in the next Section. In Section 3.3, we analyze the interaction between any two players of the tournament. We examine the relationship between equilibria of the tournament and equilibria of the game on which it is built in Section 3.4, and we discuss it in some examples in Section 3.5.

3.2. The Model

A *double round-robin tournament* $D(g, n)$, simply referred to as *tournament* in the following, is a simultaneous n -player game built on a two-player game $g = (S_h, S_a, u_h, u_a)$. Each player plays g with every other player twice, *at home* and *away*. The sets S_h and S_a are the finite sets of *actions*, the former available to the player who is playing at home and the latter available to the player who is playing away. Every player is assumed to employ the same action $s_h \in S_h$ in all the matches he plays at home and the same action $s_a \in S_a$ in all the matches he plays away. The real-valued functions u_h and u_a are defined on $S_h \times S_a$. When the player who is at home plays s_h and the player who is away plays s_a , $u_h(s_h, s_a)$ and $u_a(s_h, s_a)$ are the payoffs they respectively obtain in the match. A player's *score* is the sum of the payoffs he obtains in the $2(n-1)$ matches he participates in. The player with the highest score wins the tournament. In

the case of a tie, the winner is chosen randomly among the set of top-scoring players. We assume that each player's objective is to maximize his probability of winning the tournament.

A *pure strategy* s^i of player i is a mapping which assigns an action s_h to the matches he plays at home and an action s_a to the matches he plays away. The set of all pure strategies of each player is $S \equiv S_h \times S_a$. A *mixed strategy* σ^i of player i is an element of $\Sigma \equiv \Delta(S)$, the set of all probability distributions on S . A *strategy profile* $\sigma = (\sigma^1, \dots, \sigma^n)$ is an element of Σ^n . We define a *b-strategy* $b^i = (b_h^i, b_a^i)$ of player i as a pair of probability distributions, the first on S_h and the second on S_a . The set of all *b-strategies* of each player is $B \equiv \Delta(S_h) \times \Delta(S_a)$. Given a *b-strategy* b^i of player i , the corresponding *product mixed strategy* is the mixed strategy σ^i defined by $\sigma^i(s) = b^i(s_h) \cdot b^i(s_a)$ for every $s \in S$. A *b-strategy profile* $b = (b^1, \dots, b^n)$ is an element of B^n .

Following Arad and Rubinstein (2013), we assume that mixed strategies and *b-strategies* are executed only once and the player employs the resulting actions s_h and s_a in all the matches he plays at home and away respectively. Under this assumption, a mixed strategy corresponds to a single randomization, while a *b-strategy* corresponds to two independent randomizations, one at home and one away.

Given the structure of the problem, the tournament is a symmetric game (but the match g is usually not) and we focus on symmetric Nash equilibria.⁴ Let $P(\sigma^i, \sigma)$ be the probability that player i wins the tournament when he plays the mixed strategy σ^i and his $(n - 1)$ opponents play according to σ .

DEFINITION 3.1. A strategy profile $\sigma = (\sigma^*, \dots, \sigma^*)$ is a *symmetric Nash equilibrium* of the tournament $D(g, n)$ if

$$P(\sigma^*, \sigma) \geq P(\sigma', \sigma) \quad \text{for all } \sigma' \in \Sigma.$$

⁴For the sake of notation, we will often denote a symmetric Nash equilibrium with the (mixed or *b*-) strategy that every player chooses.

The set Σ is nonempty, compact, and convex, and the function $P(\sigma^i, \sigma)$ is linear in σ^i and continuous in σ , making the associated best response correspondence convex-valued and upper semicontinuous. Thus, by standard fixed point theorems applied to finite symmetric games, a symmetric Nash equilibrium in mixed strategies always exists.

Notice that, if $\sigma = (\sigma^*, \dots, \sigma^*)$ is a symmetric Nash equilibrium of the tournament, every action in the support of σ^* wins with probability $1/n$ when all the other $(n-1)$ players play σ^* . On the contrary, every action that is not in the support of σ^* wins the tournament with probability not greater than $1/n$.

We define symmetric Nash equilibria of tournaments also in terms of b -strategies. Let $P(b^i, b)$ be the probability that player i wins the tournament when he plays the b -strategy b^i and his opponents play according to the b -strategy profile b .

DEFINITION 3.2. A b -strategy profile $b = (b^*, \dots, b^*)$ is a *symmetric Nash equilibrium* of the tournament $D(g, n)$ if

$$P(b^*, b) \geq P(b', b) \quad \text{for all } b' \in B.$$

3.3. The Two-Player Interaction

We now examine the relationship between mixed strategies and b -strategies. In particular, we prove that mixed and b -strategies are related through an analogue of Kuhn's theorem (Kuhn, 1953). As we will see in the next section, this result allows us to extend some general properties of round-robin tournaments to our tournaments.

Given the match $g = (S_h, S_a, u_h, u_a)$, construct the two-player symmetric game $G = (S, u)$, where $u(s, s') = u_h(s_h, s'_a) + u_a(s'_h, s_a)$ is the utility of a player who plays s when his opponent plays s' . G summarizes the two matches that each player plays with every other player, one at home and one away, therefore

it describes any two-player interaction of the tournament. Clearly, the double round-robin tournament built on top of g coincides with the round-robin tournament built on top of G . Moreover, note that the set of b -strategies B is the strategy space of the match g , while the set of mixed strategies Σ is each player's strategy set in the game G . Thus, with slight abuse of notation, we will denote a strategy profile of g with b^i and a strategy of G with σ^i .

Let two (mixed or b -) strategies of player i be *outcome-equivalent in the two-player interaction* of the tournament if, for every (mixed or b -) strategy of the opponent, they induce the same probability distributions on the payoffs that each player can get at home and on the payoffs that each player can get away.

PROPOSITION 3.1. *In the two-player interaction of the tournament, for any mixed strategy $\sigma^i \in \Sigma$ of player i there is an outcome-equivalent b -strategy $b^i \in B$, and vice versa.*

PROOF. Consider the two matches that player i plays against player j , at home and away. Let $C(s_h)$ and $C(s_a)$ denote respectively the set of all pure strategies that choose s_h in the match at home and the set of all pure strategies that choose s_a in the match away.

For every mixed strategy σ^k of player k , $k = i, j$, consider the b -strategy b^k defined by $b^k(s_h) = \sum_{s \in C(s_h)} \sigma^k(s)$ and $b^k(s_a) = \sum_{s \in C(s_a)} \sigma^k(s)$. Fix a mixed strategy σ^j of player j . For any $s, s' \in S$, the weight that σ^i induces on player i 's utility $u^i(s, s') = u_h^i(s_h, s'_a) + u_a^i(s'_h, s_a)$ is equal to $\sigma^i(s)\sigma^j(s')$.⁵ Thus, for any $s_h, s'_h \in S_h$ and $s_a, s'_a \in S_a$, σ^i induces a weight of $\sum_{s \in C(s_h)} \sum_{s' \in C(s'_a)} \sigma^i(s)\sigma^j(s')$ on $u_h^i(s_h, s'_a)$ and a weight of $\sum_{s \in C(s_a)} \sum_{s' \in C(s'_h)} \sigma^i(s)\sigma^j(s')$ on $u_a^i(s'_h, s_a)$. By construction, these weights are respectively equal to $b_h^i(s_h)b_a^j(s'_a)$ and to $b_a^i(s_a)b_h^j(s'_h)$, which are exactly the weights of $u_h^i(s_h, s'_a)$ and $u_a^i(s'_h, s_a)$ according to b^i , given b^j . In the same way, the equivalence holds also for the weights induced by σ^i and b^i on player j 's utilities.

⁵Note that, given σ^j , $\sigma^i(s)\sigma^j(s')$ is also the weight induced by σ^i on player j 's utility $u^j(s', s) = u_h^j(s'_h, s_a) + u_a^j(s_h, s'_a)$.

For the other way round, it is enough to consider for any b -strategy b^k of player k , $k = i, j$, the corresponding product mixed strategy σ^k . Given a (mixed or b -) strategy of player j , for any $s_h, s'_h \in S_h$ and $s_a, s'_a \in S_a$, the weights of $u_h^i(s_h, s'_a)$, $u_a^i(s'_h, s_a)$, $u_h^j(s'_h, s_a)$, and $u_a^j(s_h, s'_a)$ according to b^i and σ^i are clearly the same. \square

We consider now the Nash equilibria of the game G . Given the equivalence result stated in Proposition 3.1, we can define these equilibria also in terms of b -strategies.

First, note that a pair of mixed strategies (σ^i, σ^j) is a Nash equilibrium of G if and only if the following conditions are both satisfied:

$$\text{if } \sigma^i(s) > 0, \text{ then } u(s, \sigma^j) = \max_{s' \in S} u(s', \sigma^j), \text{ and} \quad (3.3.1)$$

$$\text{if } \sigma^j(s) > 0, \text{ then } u(s, \sigma^i) = \max_{s' \in S} u(s', \sigma^i) \quad (3.3.2)$$

where, given $s \in S$ and $\sigma^j \in \Sigma$, $u(s, \sigma^j) = \sum_{s' \in S} \sigma^j(s') u(s, s')$.

Let $b^i, b^j \in B$. We define $u_h(s_h, b_a^j) = \sum_{s_a \in S_a} b^j(s_a) u_h(s_h, s_a)$ and $u_a(b_h^j, s_a) = \sum_{s_h \in S_h} b^j(s_h) u_a(s_h, s_a)$. Then, $u_h(b_h^i, b_a^j) = \sum_{s_h \in S_h} \sum_{s_a \in S_a} b^i(s_h) b^j(s_a) u_h(s_h, s_a)$ and $u_a(b_h^j, b_a^i) = \sum_{s_h \in S_h} \sum_{s_a \in S_a} b^i(s_a) b^j(s_h) u_a(s_h, s_a)$. Finally, we have $u(b^i, b^j) = u_h(b_h^i, b_a^j) + u_a(b_h^j, b_a^i)$.

DEFINITION 3.3. A pair of b -strategies (b^i, b^j) is a *Nash equilibrium* of G if

$$u(b^i, b^j) \geq u(b', b^j) \quad \text{for all } b' \in B, \text{ and}$$

$$u(b^j, b^i) \geq u(b^j, b') \quad \text{for all } b' \in B.$$

Recall that each player randomizes independently in the match he plays at home and in the match he plays away. Moreover, it is easy to see that in the game G a player has a profitable deviation from a given strategy if and only if it is profitable for him to deviate either in the match at home, or in the match away, or in both. It follows that (b^i, b^j) is a Nash equilibrium of G if and only if the following conditions are satisfied:

$$\text{if } b^i(s_h) > 0, \text{ then } u_h(s_h, b_a^j) = \max_{s'_h \in S_h} u_h(s'_h, b_a^j), \quad (3.3.3)$$

$$\text{if } b^i(s_a) > 0, \text{ then } u_a(b_h^j, s_a) = \max_{s'_a \in S_a} u_a(b_h^j, s'_a), \quad (3.3.4)$$

and the two analogous conditions for b^j .

Note that conditions (3.3.3) and (3.3.4) imply that

$$\text{if } b^i(s_h) \cdot b^i(s_a) > 0, \text{ then } u_h(s_h, b_a^j) + u_a(b_h^j, s_a) = \max_{s'_h \in S_h, s'_a \in S_a} u_h(s'_h, b_a^j) + u_a(b_h^j, s'_a),$$

that is,

$$\text{if } b^i(s_h) \cdot b^i(s_a) > 0, \text{ then } u(s, b^j) = \max_{s' \in S} u(s', b^j), \quad (3.3.5)$$

and, analogously,

$$\text{if } b^j(s_h) \cdot b^j(s_a) > 0, \text{ then } u(s, b^i) = \max_{s' \in S} u(s', b^i). \quad (3.3.6)$$

COROLLARY 3.1. *If (b^i, b^j) is a Nash equilibrium of G , then every equivalent pair of mixed strategies (σ^i, σ^j) is a Nash equilibrium of G . Vice versa, if (σ^i, σ^j) is a Nash equilibrium of G , then the equivalent pair of b -strategies (b^i, b^j) is a Nash equilibrium of G .*

PROOF. Let (σ^i, b^i) and (σ^j, b^j) be two couples of equivalent strategies. By definition, for every $s \in S$ and $k \in \{i, j\}$, $u(s, \sigma^k) = u(s, b^k)$. Moreover, if $\sigma^k(s) > 0$ then $b^k(s_h) \cdot b^k(s_a) > 0$, while if $b^k(s_h)$ then $\sigma^k(s_h s'_a) > 0$ for at least one $s'_a \in S_a$ and if $b^k(s_a) > 0$ then $\sigma^k(s'_h s_a) > 0$ for at least one $s'_h \in S_h$. Therefore, if (b^i, b^j) is a Nash equilibrium of G , then by conditions (3.3.5) and (3.3.6) also (σ^i, σ^j) is a Nash equilibrium of G . Conversely, if (σ^i, σ^j) is a Nash equilibrium of G then so is (b^i, b^j) , since by condition (3.3.1) $b^i(s_h) > 0$ implies $u_h(s_h, b_a^j) = \max_{s'_h \in S_h} u_h(s'_h, b_a^j)$ and $b^i(s_a) > 0$ implies $u_a(b_h^j, s_a) = \max_{s'_a \in S_a} u_a(b_h^j, s'_a)$, and by condition (3.3.2) the same holds for b^j . \square

Similar arguments imply the following equivalence result about Nash equilibria of g and symmetric Nash equilibria of G :

PROPOSITION 3.2. *If $b^i = (b_h^i, b_a^i)$ is a Nash equilibrium of the game g , then (b^i, b^i) is a symmetric Nash equilibrium of the corresponding game G , and vice versa.*

Let the strategies b^i and σ^i be outcome-equivalent in the two-player interaction. Clearly, it follows from Corollary 3.1 that if $b^i = (b_h^i, b_a^i)$ is a Nash equilibrium of g , then (σ^i, σ^i) is a symmetric Nash equilibrium of G , and vice versa.

3.4. The Tournament

In this section, we analyze the relationship between equilibria of the game g and equilibria of the tournament built on g . First, we explore further the relation between mixed and b -strategies in the tournament. Note that, given a b -strategy b^i , there may be more than one mixed strategy that is outcome-equivalent to b^i in the two-player interaction. Hence, we start by examining whether all the strategies that are equivalent in the two-player interaction are also equivalent in the tournament.

Two (mixed or b -) strategies of player i are *outcome-equivalent in the tournament* if, for every n and for every (mixed or b -) strategy of the other $(n - 1)$ players, they induce the same probability of winning for each action. The following example shows that two strategies that are outcome-equivalent in the two-player interaction are not necessarily outcome-equivalent in the tournament.

EXAMPLE 3.1 (Battle of the sexes). Let the match g be the “battle of the sexes” game

	L	R
T	2, 1	0, 0
B	0, 0	1, 2

which generates the following two-player symmetric game G :

	TL	TR	BL	BR
TL	3,3	1,2	2,1	0,0
TR	2,1	0,0	4,2	2,1
BL	1,2	2,4	0,0	1,2
BR	0,0	1,2	2,1	3,3

Take the b -strategy $b^i = (\frac{2}{3}T + \frac{1}{3}B, \frac{1}{3}L + \frac{2}{3}R)$, which is a Nash equilibrium of g . The corresponding product mixed strategy $\sigma^i = (\frac{2}{9}TL + \frac{4}{9}TR + \frac{1}{9}BL + \frac{2}{9}BR)$ and the strategy $\tilde{\sigma}^i = (\frac{2}{3}TR + \frac{1}{3}BL)$ are both outcome-equivalent to b^i in the two-player interaction, and they are both symmetric Nash equilibria of G . However, they are not outcome-equivalent in the tournament. To see this, consider for instance the tournament $D = (g, 4)$ where player i faces three opponents playing strategy TR . If player i plays σ^i , he wins the tournament with probability $\frac{2}{3}$, while if he plays $\tilde{\sigma}^i$ he wins with probability $\frac{1}{2}$.

Henceforth, we refer to strategies that are outcome-equivalent in the two-player interaction simply as *equivalent* strategies. The previous example implies that if a (mixed or b -) strategy is a Nash equilibrium of the tournament $D(g, n)$, an equivalent strategy is not necessarily an equilibrium of $D(g, n)$. However, the relationship between equilibria of the tournament and equilibria of the game on which it is built suggests that the “equilibrium behavior” of strategies should be studied at the limit, as $n \rightarrow \infty$. Thus, let us now analyze such a relationship.

For the time being, we consider Nash equilibria of the tournament in terms of mixed strategies. First, recall that the double round-robin tournament built on g coincides with the round-robin tournament built on the corresponding game G . Arad and Rubinstein (2013) prove two main results about the relationship between the equilibria of a round-robin tournament and the equilibria of the symmetric game on which it is built. Thus, we can apply directly their results to the relationship between the equilibria of a tournament $D(g, n)$ and the equilibria of the corresponding game G . As a consequence of Proposition

3.2, we can then extend such results to the equilibria of the game g , which is not required to be symmetric.

The following proposition is a direct extension of Proposition 1 in Arad and Rubinstein (2013), so we state it without proof.

PROPOSITION 3.3. *Let σ^i be the limit point of a subsequence of symmetric Nash equilibria of $D(g, n)$ as $n \rightarrow \infty$. Then σ^i is a (symmetric) Nash equilibrium of G , and the equivalent b -strategy $b^i = (b_h^i, b_a^i)$ is a Nash equilibrium of g .*

PROPOSITION 3.4. *Let $b^i = (b_h^i, b_a^i)$ be a Nash equilibrium of g and let σ^i be an equivalent symmetric Nash equilibrium of G . The strategy σ^i is not necessarily a limit point of a sequence of symmetric Nash equilibria of $D(g, n)$ as $n \rightarrow \infty$.*

PROOF. Consider the following game g :⁶

	A	B
A	1, 1	0, 1
B	1, 0	0, 0

and the double round-robin tournament $D(g, n)$ built on it. The b -strategy (A, A) is a Nash equilibrium of g . However, the unique equivalent mixed strategy, AA , is dominated in the tournament by any other strategy. Indeed, for any choice of the opponents, a player has always the incentive to deviate from AA both at home and away, and BB is the only symmetric equilibrium of $D(g, n)$ for any n . Therefore, there is no Nash equilibrium of the tournament that assigns positive probability to AA . □

Note that Propositions 3.3 and 3.4 suggest a refinement criterion, according to which the equilibria of g that are equivalent to those equilibria of G with a close-by equilibrium of the tournament with a large number of players, are “more stable” than others.

⁶Note that this is the same “degenerate” game used to prove Proposition 2 in Arad and Rubinstein (2013).

Let $b^i = (b_h^i, b_a^i)$ be a Nash equilibrium of g . One may conjecture that if a mixed strategy that is equivalent to b^i is the limit point of a sequence of equilibria of the tournament as $n \rightarrow \infty$, then all the equivalent mixed strategies are. The following example shows that, if one of the equivalent mixed strategies is such a limit point, the other equivalent mixed strategies are not necessarily equilibria of the tournament for any n .

EXAMPLE 3.2. Let the match g be:

	C	D
A	2, 0	0, 2
B	0, 1	1, 0

and the corresponding two-player symmetric game G be:

	AC	AD	BC	BD
AC	2, 2	0, 4	3, 0	1, 2
AD	4, 0	2, 2	2, 1	0, 3
BC	0, 3	1, 2	1, 1	2, 0
BD	2, 1	3, 0	0, 2	1, 1

The b -strategy $b^i = (\frac{1}{3}A + \frac{2}{3}B, \frac{1}{3}C + \frac{2}{3}D)$ is the unique Nash equilibrium of g . Proposition 3.3 implies that at least one of the mixed strategies that are equivalent to b^i (and, hence, symmetric equilibria of G) must be a limit point of a subsequence of symmetric Nash equilibria of $D(g, n)$ as $n \rightarrow \infty$. Now, take the equivalent mixed strategy $\tilde{\sigma}^i = (\frac{1}{3}AC + \frac{2}{3}BD)$. For any n , if the other $(n - 1)$ players play according to $\tilde{\sigma}^i$, action BD wins the tournament with probability always higher than that of action AC , since BD yields always one additional point. It follows that $\tilde{\sigma}^i$ is never an equilibrium of $D(g, n)$ for any n .

Nevertheless, the result of the example does not exclude the possibility that $\tilde{\sigma}^i$ is the limit point of a subsequence of symmetric Nash equilibria of the tournament as $n \rightarrow \infty$. Up to now, whether the conjecture mentioned before is true or not is still an open problem.


Notice that, if the conjecture were refuted, the analysis should be restricted to b -strategies and to corresponding product mixed strategies. Recall, indeed, that the object of the analysis are the equilibria of the game g , and the equilibria of G are just a means to extend directly the results of Arad and Rubinstein (2013) to our model. To this extent, an appropriate proof of the existence of symmetric Nash equilibria of tournaments in b -strategies should be done.⁷

3.5. Examples

We present now some examples, where the relationship between equilibria of the base game and equilibria of the tournament is discussed. In particular, we compare the equilibrium refinement of tournaments to the standard refinements in the literature.

The first example shows that a dominated Nash equilibrium of g can be a symmetric equilibrium of $D(g, n)$ for every n . Moreover, it shows that a stable set of g in the sense of Kohlberg and Mertens (1986) does not necessarily contain an equilibrium of the tournament $D(g, n)$ for any n .

EXAMPLE 3.3 (Ultimatum game). Consider the following ultimatum game, where Player 1 can offer a fair (F) or unfair (U) proposal about how to split 10 dollars, and Player 2 can either accept (A) or reject (R) it:




Let the match g be the corresponding normal form game:

⁷In particular, we have that the set B is a nonempty, compact, and convex set. However, the function $P(b^i, b)$ is continuous in b but it is not quasi-concave in b^i , so we cannot use standard arguments to state that the best response correspondence is convex-valued.

	<i>AA</i>	<i>AR</i>	<i>RA</i>	<i>RR</i>
<i>U</i>	10,0	10,0	0,0	0,0
<i>F</i>	5,5	0,0	5,5	0,0

Note that the mixed strategy FRA is a symmetric equilibrium of the tournament for every n , and the equivalent b -strategy (F, RA) is an undominated Nash equilibrium of g . Also URR is a symmetric equilibrium of the tournament for every n , but the equivalent b -strategy (U, RR) is a dominated Nash equilibrium of g (however, the strategy URR is not dominated in the tournament). Moreover, note that the b -strategy (U, AA) is a strictly perfect equilibrium of g , therefore a Kohlberg-Mertens stable set. The equivalent product mixed strategy UAA , however, is not a symmetric equilibrium of the tournament for any n . Indeed, in the matches in which he moves second, each player has the incentive to deviate and reject the unfair offer, in order to inflict a loss of 10 to his opponent.

EXAMPLE 3.4 (Modified ultimatum game). Consider now a modified version of the ultimatum game, in which Player 1 can make also an intermediate (M) offer that gives an amount z to Player 2, with $0 < z < 5$:




Let the match $g(z)$ be the corresponding normal form game:

	<i>AAA</i>	<i>AAR</i>	<i>ARA</i>	<i>ARR</i>	<i>RAA</i>	<i>RAR</i>	<i>RAA</i>	<i>RRR</i>
<i>U</i>	10,0	10,0	10,0	10,0	0,0	0,0	0,0	0,0
<i>M</i>	10-z,z	10-z,z	0,0	0,0	10-z,z	10-z,z	0,0	0,0
<i>F</i>	5,5	0,0	5,5	0,0	5,5	0,0	5,5	0,0

As before, to offer the fair proposal when moving first and to accept only it when moving second is a symmetric Nash equilibrium of the tournament for every n and every z . Now, however, to accept also an unfair proposal can be part of an equilibrium strategy of the tournament. Indeed, for each z , $MRAA$ is a symmetric Nash equilibrium of the tournament $D(g(z), n)$ if and only if $n \geq n_z = \lceil \frac{10}{z} \rceil$. Note that $\lim_{z \rightarrow 0} n_z = \infty$.

Lastly, we present a further example in which dominated equilibria of g are limit points of symmetric equilibria of the tournament as $n \rightarrow \infty$.

EXAMPLE 3.5 (Entry game). Consider the following “entry game”, in which Firm 1 has to decide whether or not to enter the market, and Firm 2 has to decide how to compete, either aggressively (Fight) or not (Accomodate):



Let the match g be the corresponding normal form game:

	F	A
N	0, 2	0, 2
E	-1, -1	1, 1

The game g has an undominated Nash equilibrium, (E, A) . The corresponding product mixed strategy, EA , is a symmetric Nash equilibrium of the tournament $D(g, n)$ for every $n \geq 3$. Moreover, g has a continuum of dominated Nash equilibria, $\{(N, \alpha F + (1 - \alpha)A) : \frac{1}{2} \leq \alpha \leq 1\}$. Note that NF is a symmetric equilibrium of the tournament $D(g, n)$ for $n \geq 4$. To see whether the other mixed strategies in the continuum $\{(\alpha NF + (1 - \alpha)NA) : \frac{1}{2} \leq \alpha \leq 1\}$ are symmetric equilibria of the tournament for large values of n , note first that the only profitable deviation to consider is playing E instead of N when in the role of the first player.

Then, for a given $\alpha \in [\frac{1}{2}, 1)$, consider a player who plays $(\alpha EF + (1-\alpha)EA)$ while all the other players are playing $(\alpha NF + (1-\alpha)NA)$. Let x be the number of players that play NF in equilibrium. The player who plays E attains a score of $3n - 3 - 2x$, the players that play NA get $2n - 3$, while the players that play NF get $2n - 5$ (so they never win the tournament). Of course, when all the players play the same strategy, they all win with probability $1/n$. Thus, playing E when in the role of the first player is a profitable deviation if and only if

$$\mathbb{P}\left(x = \frac{n}{2}\right) \frac{1}{n-x} + \mathbb{P}\left(x < \frac{n}{2}\right) > \frac{1}{n}, \quad (3.5.1)$$

where $\mathbb{P}\left(x = \frac{n}{2}\right) = \binom{n-1}{n/2} \alpha^{n/2} (1-\alpha)^{n-1-n/2}$ is positive only if n is an even number, and $\mathbb{P}\left(x < \frac{n}{2}\right) = \sum_{k=0, \dots, m-1} \binom{n-1}{k} \alpha^k (1-\alpha)^{n-1-k}$, with $m = \lceil \frac{n}{2} \rceil$.

When $\alpha = \frac{1}{2}$, E is always a profitable deviation, so $(\frac{1}{2}NF + \frac{1}{2}NA)$ is never a symmetric Nash equilibrium of the tournament for any n . For a fixed $n > 4$, the lhs of (3.5.1) is decreasing in α and equals the rhs at a value $\alpha_n^* \in (\frac{1}{2}, 1)$. It follows that, for every $n > 4$, $\{(\alpha NF + (1-\alpha)NA) : \alpha_n^* \leq \alpha \leq 1\}$ is a continuum of Nash equilibria of the tournament $D(g, n)$. In particular, α_n^* is decreasing in n and approaches $\frac{1}{2}$ as n goes to infinity.⁸ We can thus conclude that all the strategies in the continuum, included $(\frac{1}{2}NF + \frac{1}{2}NA)$, are limit points of equilibria of the tournament as $n \rightarrow \infty$.

⁸For $n = 10, 100, 1000, 10000$, $\alpha_n^* \approx 0.72, 0.61, 0.55, 0.51$.

Bibliography

- Al-Najjar, N. Strategically stable equilibria in games with infinitely many pure strategies. *Mathematical Social Sciences*, 29(2):151–164, 1995.
- Arad, A. and Rubinstein, A. Strategic tournaments. *American Economic Journal: Microeconomics*, 5(4):31–54, 2013.
- Arthur, W. B. Inductive reasoning and bounded rationality. *American Economic Review*, 84(2):406–11, 1994.
- Bierstone, E. and Milman, P. D. Semianalytic and subanalytic sets. *Publications Mathématiques de l’I.H.É.S.*, 67:5–42, 1988.
- Billingsley, P. *Convergence of Probability Measures*. Wiley, 1968.
- Blume, L. E. and Zame, W. R. The algebraic geometry of perfect and sequential equilibrium. *Econometrica*, 62(4):783–94, 1994.
- Bochnak, J., Coste, M., and Roy, M.-F. *Géométrie Algébrique Réelle*. Berlin: Springer-Verlag, 1987.
- Bohnenblust, H., Karlin, S., and Shapley, L. Solutions of discrete two-person games. In Kuhn, H. and Tucker, A., editors, *Contributions to the Theory of Games*, volume 1, pages 51–72. 1950.
- Bouton, L. A theory of strategic voting in runoff elections. *American Economic Review*, 103(4):1248–88, 2013.
- Bouton, L. and Castanheira, M. One person, many votes: divided majority and information aggregation. *Econometrica*, 80(1):43–87, 2012.
- Bouton, L. and Gratton, G. Majority runoff elections: strategic voting and Duverger’s hypothesis. *Theoretical Economics*, forthcoming.
- Carbonell-Nicolau, O. Essential equilibria in normal-form games. *Journal of Economic Theory*, 145(1):421–431, 2010.
- Carbonell-Nicolau, O. On the existence of pure-strategy perfect equilibrium in discontinuous games. *Games and Economic Behavior*, 71(1):23–48, 2011.

- De Sinopoli, F. and Pimienta, C. Undominated (and) perfect equilibria in Poisson games. *Games and Economic Behavior*, 66(2):775–784, 2009.
- De Sinopoli, F., Meroni, C., and Pimienta, C. Strategic stability in Poisson games. *Journal of Economic Theory*, 153(C):46–63, 2014.
- Dixit, A. K. Strategic behavior in contests. *American Economic Review*, 77(5):891–98, 1987.
- Fort, M. K. Points of continuity of semicontinuous functions. *Publicationes Mathematicae Debrecen*, 2:100–102, 1951.
- Gabriélov, A. M. Projections of semi-analytic sets. *Functional Analysis and Its Applications*, 2(4):282–291, 1968.
- Gale, D. and Sherman, S. Solutions of finite two-person games. In Kuhn, H. and Tucker, A., editors, *Contributions to the Theory of Games*, volume 1, pages 37–49. 1950.
- Govindan, S. and McLennan, A. On the generic finiteness of equilibrium outcome distributions in game forms. *Econometrica*, 69(2):455–71, 2001.
- Govindan, S. and Wilson, R. Direct proofs of generic finiteness of Nash equilibrium outcomes. *Econometrica*, 69(3):765–69, 2001.
- Govindan, S. and Wilson, R. Metastable equilibria. *Mathematics of Operations Research*, 33(4):787–820, 2008.
- Green, J. R. and Stokey, N. L. A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3):349–64, 1983.
- Groh, C., Moldovanu, B., Sela, A., and Sunde, U. Optimal seedings in elimination tournaments. *Economic Theory*, 49(1):59–80, 2012.
- Hardt, R. M. Semi-algebraic local triviality in semi-algebraic mappings. *American Journal of Mathematics*, 102(2):291–302, 1980.
- Hillas, J. On the definition of the strategic stability of equilibria. *Econometrica*, 58(6):1365–90, 1990.
- Hillas, J. and Kohlberg, E. Foundations of strategic equilibria. In Aumann, R. J. and Hart, S., editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 42, pages 1597–1663. Elsevier Science, 2002.
- Huges, N. Voting in legislative elections under plurality rule. mimeo, 2012.
- Jehiel, P. and Lamy, L. On discrimination in procurement auctions. mimeo, 2014.

- Kinoshita, S. On essential components of the set of fixed points. *Osaka Mathematical Journal*, 4(1):19–22, 1952.
- Kohlberg, E. Refinement of Nash equilibrium: the main ideas. In Ichiishi, T., Neyman, A., and Tauman, Y., editors, *Game Theory and Applications*, pages 3–45. San Diego: Academic Press, 1990.
- Kohlberg, E. and Mertens, J.-F. On the strategic stability of equilibria. *Econometrica*, 54(5):1003–37, 1986.
- Konrad, K. A. Bidding in hierarchies. *European Economic Review*, 48(6):1301–1308, 2004.
- Konrad, K. A. *Strategy and Dynamics in Contests*. New York: Oxford University Press, 2009.
- Kreps, D. M. and Wilson, R. Sequential equilibria. *Econometrica*, 50(4):863–94, 1982.
- Kuhn, H. W. Extensive games and the problem of information. In Kuhn, H. and Tucker, A., editors, *Contributions to the Theory of Games*, volume 2, pages 193–216. Princeton: Princeton University Press, 1953.
- Laffond, G., Laslier, J.-F., and Breton, M. L. K-player additive extension of two-player games with an application to the Borda electoral competition game. *Theory and Decision*, 48(2):129–137, 2000.
- Lojasiewicz, S. Triangulation of semi-analytic sets. *Annali della Scuola Normale Superiore di Pisa*, 18(4):449–474, 1964.
- Lojasiewicz, S. *Ensembles Semi-Analytique*. I.H.É.S. Bures-sur-Yvette, 1965.
- Luce, R. D. and Raiffa, H. *Games and Decisions: Introduction and Critical Survey*. Courier Dover Publications, 1957.
- Makris, M. Complementarities and macroeconomics: Poisson games. *Games and Economic Behavior*, 62(1):180–189, 2008.
- Makris, M. Private provision of discrete public goods. *Games and Economic Behavior*, 67(1):292–299, 2009.
- Maniquet, F. and Morelli, M. Approval quorums dominate participation quorums. mimeo, 2013.
- McLennan, A. Manipulation in elections with uncertain preferences. *Journal of Mathematical Economics*, 47(3):370–375, 2011.
- McLennan, A. *Advanced Fixed Point Theory for Economics*. mimeo, 2012.

- Mertens, J.-F. Stable equilibria - a reformulation. Part I. Definition and basic properties. *Mathematics of Operations Research*, 14(4):575–625, 1989.
- Mertens, J.-F. Stable equilibria - a reformulation. Part II. Discussion of the definition, and further results. *Mathematics of Operations Research*, 16(4):694–753, 1991.
- Milchtaich, I. Random-player games. *Games and Economic Behavior*, 47(2):353–388, 2004.
- Myerson, R. B. Population uncertainty and Poisson games. *International Journal of Game Theory*, 27(3):375–392, 1998.
- Myerson, R. B. Comparison of scoring rules in Poisson voting games. *Journal of Economic Theory*, 103(1):219–251, 2002.
- Núñez, M. The strategic sincerity of approval voting. *Economic Theory*, 56(1):157–189, 2014.
- Pearce, D. G. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–50, 1984.
- Ritzberger, K. Price competition with population uncertainty. *Mathematical Social Sciences*, 58(2):145–157, 2009.
- Rosen, S. Prizes and incentives in elimination tournaments. *American Economic Review*, 76(4):701–15, 1986.
- Satterthwaite, M. and Shneyerov, A. Dynamic matching, two-sided incomplete information, and participation costs: existence and convergence to perfect competition. *Econometrica*, 75(1):155–200, 2007.
- Selten, R. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):24–55, 1975.
- Simon, L. K. and Stinchcombe, M. B. Equilibrium refinement for infinite normal-form games. *Econometrica*, 63(6):1421–43, 1995.
- van der Waerden, B. L. *Einführung in die Algebraische Geometrie*. Berlin-Heidelberg-New York: Springer Verlag, 1939.
- Wu Wen-Tsun, J. J.-H. Essential equilibrium points of n-person non-cooperative games. *Scientia Sinica*, 11(5):1307–1322, 1963.