Les Cahiers du GERAD ISSN: 0711–2440

The Return Function: A New Computable Perspective on Bayesian-Nash Equilibria

L.N. Hoang, F. Soumis,
G. Zaccour

G–2014–07

February 2014

The Return Function: A New Computable Perspective on Bayesian-Nash Equilibria

Lê Nguyên Hoang
François Soumis
GERAD & Polytechnique Montréal
Montréal (Québec) Canada, H3C 3A7
le.nguyen.hoang@gerad.ca
francois.soumis@gerad.ca

Georges Zaccour
GERAD & Chair in Game Theory and Management
HEC Montréal
Montréal (Québec) Canada, H3T 2A7
georges.zaccour@gerad.ca

February 2014

Les Cahiers du GERAD
G–2014–07

Abstract: In this paper, we suggest a new approach called the *return function* to deal with the determination of Bayesian-Nash equilibria in games of incomplete information. Whereas in the traditional approach players reply to each others' strategies, here each player replies to his own return function. In short, given a player's choice of action and the other players' strategies, the return function of that given player is the probability distribution of the outcome. Interestingly, we show that the dynamics of best-reply strategies, which are hard to compute in practice, are mapped to an observable and easy-to-compute dynamics of return functions. We propose a new algorithm for computing Bayesian-Nash equilibria, and illustrate its implementation on a cake-cutting problem. Finally, we prove the convergence of the dynamics of return functions to the Bayesian-Nash equilibrium under fairly general topological assumptions.

Key Words: Games/group decisions; Noncooperative.

Résumé : On propose dans un cet article une nouvelle méthode, appelée la fonction retour, pour la détermination des équilibres bayésiens de Nash dans des jeux à information incomplète. Tandis que traditionnellement on calcule ces équilibres en cherchant pour chaque joueur la stratégie qui constitue la meilleure réponse aux stratégies des autres, notre approche consiste à déterminer pour chaque joueur sa meilleure réponse aux fonctions retour des autres joueurs. En bref, à une action donnée d'un joueur et un profil de stratégies des autres joueurs, la fonction retour pour ce joueur correspond à la distribution de probabilités sur les gains. On montre que cette approche est plus simple à utiliser que l'approche traditionnelle. On illustre son utilisation sur un problème de partage de gâteau, un problème largement étudié dans la littérature. Pour terminer, on fournit une analyse détaillée de la convergence de cette méthode vers l'équilibre bayésien.

Mots clés : Jeux non coopératifs; équilibres bayésiens; fonction retour; partage de gâteau.

1 Introduction

Mechanism (or market) design has proven to be a successful approach for efficiently determining the value of a product or a service when there is no natural price that can be posted or negotiated for that product, as is the case for, e.g., a painting by Picasso, energy prices in a deregulated electricity market, or the exploitation rights for a hydrocarbon-rich basin. These examples, and many others, share the following features: (i) There is a finite number of strategic agents (players, bidders or claimers) interested in acquiring the object (product, service or resource). (ii) Each agent has a *private* value for the object under consideration, and does not know how much the other agents value the same object. For instance, the cost of producing a kilowatt is not the same for all electricity companies in a given market, and each company knows its own cost but only has incomplete knowledge of its competitors' costs. Similarly, a Picasso painting does not have the same value for all art collectors. (iii) The rules of the game are not given in advance, but are designed by an agent, called a mechanism designer, principal or regulator, who has an interest in the outcome. For instance, a public commission may auction off television airwaves to wireless carriers to create faster and more reliable networks, or to maximize its own revenues. A parent may ask children at a party their flavour preference to fairly allocate a heterogeneous birthday cake.

During the last two decades or so, important developments in market design have taken place for three main reasons: "(i) The creation by government agencies, private firms or industrial associations of a number of markets to privatize public assets, restructure deregulated industries, or enhance inter-firm relations; (ii) a renewed focus on strategic analysis and game theory that together with the emergence of experimental economics contributed to the establishment of market design as a serious research field in economics; (iii) and, most importantly, the explosive development of electronic business, e-business tools that can embed the most complex market rules and facilitate their deployment." (Bourbeau et al. (2005)). The recent operations-research literature includes work on assignment problems: see, e.g., Su and Zenios (2006) for the kidney-transplant trade-off; Abdulkadirolu and Sönmez (2003) or Pathak (2011) for school choice; in supply chains, see, e.g., Jain and Raghavan (2009), Chen and Cheng (2012), Mes et al. (2011); and in revenue management, see, e.g., Vulcano et al. (2002), Manelli and Vincent (2007), Devenur and Hayes (2009).

One approach to market design is Bayesian mechanism design, where the information about the types of agents is incomplete, but all agents and the principal have some beliefs about others' types. A belief is a probability distribution on the agents' types. This setting corresponds to a game with incomplete information, also known as a Bayesian game (see Harsanyi (1968b,c,a)), whose solution is called a Bayesian-Nash (BN) equilibrium. Finding BN equilibria is a very difficult task as it involves solving for agents' best-response strategies and for the best inference from what is possibly a strategic lie. However, due to the revelation principle[1] (Gibbard (1973), Holmstrom (1977), Myerson (1979), Dasgupta et al. (1979)), one can only confine one's attention to equilibria in which agents truthfully report their types.

In this article, we propose a new approach, which we call the return function, to compute BN equilibria in mechanism design. In a nutshell, given a player's choice of action, the other players' strategies and the mechanism chosen by the market designer, the return function of that given player, is the probability distribution of the outcome. Given this return function, the expected utility of a player's outcome is then defined as a function of this return function and of the player's type. Our formulation is fairly general and accounts for outcomes that cannot be defined deterministically. In short, our approach for computing Bayesian-Nash equilibria is based on determining the return function that is a best-reply to itself, instead of looking for a strategy profile that is a best reply to itself. Interestingly, this simplification in the formulation comes at no cost in terms of the strategic aspect of the game, and lends itself to an efficient algorithm for computing BN equilibria. To illustrate, we provide such an algorithm, and we describe the results of instances from the cake-cutting problem.

Our approach shares some similarities with the one followed in mean-field games (Huang et al. (2006), Lasry and Lions (2007), Achdou and Capuzzo-Dolcetta (2010)). These games involve a large number of players, where each one reacts to the mass of other players over time, with each player's trajectory being

[1] The revelation principle states: "For any Bayesian-Nash equilibrium there corresponds a Bayesian game with the same equilibrium outcome but in which players truthfully report type."

written as a function of only the distribution of the trajectory of the mass. In the particular case of linear-quadratic models with Brownian motions, each player's best trajectory is further simplified to only depend on the average and variance of the distribution of the trajectory for the mass. Importantly, mean-field games then rely on assuming a certain trajectory of the mass, computing individual best-replies to this trajectory, and deriving the trajectory of the mass from these best-replies. In other words, in mean-field games, we proceed by updating the trajectory of the mass, rather than each player's strategy. There is also a link with Rabinovich et al. (2013), where the fictitious-play algorithm is extended to compute pure-strategy Bayesian equilibria for games with continuous sets of types. Later, we show that the updating approaches used in mean-field games and in Rabinovich et al. (2013) can be interpreted as special cases of updating of our return function.

The rest of the paper is organized as follows: In Section 2, we introduce the model and the return function. Section 3 is devoted to the computation of Bayesian-Nash equilibria with the return function. In Section 4, we provide an illustration in the context of a cake-cutting problem. Section 5 discusses theoretical convergence of the return function, and Section 6 concludes.

2 Model and Equilibrium

Let $N = \{1, \ldots, n\}$ be the set of players and A the set of actions of player $j \in N$. Denote by a_j an action of j, by $a = (a_1, \ldots, a_n) \in A^n$ the vector of all players' actions, and by $a_{-j} = (a_1, \ldots, a_{j-1}, a_{j+1,\ldots}, a_n) \in A^{n-1}$ the vector of the actions of all players other than j. Similar notations will be used throughout the paper for vectors of objects that refer to all players or to $n-1$ players.

The action profile $a \in A^n$ induces an outcome x from the set of outcomes X. This mapping is called a mechanism $\mathcal{M}$, that is, $\mathcal{M}(a) = x \in X$. To illustrate, x could be, e.g., the workers' schedule for a given week depending on their requests, the shares of a cake allocated to the different claimers according to their stated preferences, or the quantity of energy to be supplied the next day by the bidding electricity companies.

In practice, this outcome may not be deterministically specified, because of some inherent random events. For instance, the next day's electricity demand depends on temperature, which cannot be predicted with certainty. To reflect this, we let mechanism $\mathcal{M}$ be a function in $\Delta(X)$, where $\Delta(X)$ is the set of probability distribution on the set of outcomes X. If the mechanism is deterministic, then $\mathcal{M}(a)$ would have a Dirac distribution.

Each player $j \in N$ is defined by his type θ_j. For each player j, a utility function matches his type and the outcome with a real number, as follows:

$$u_j(\theta_j, x) \geq u_j(\theta_j, x') \quad \text{for} \quad x \succeq x',$$

where the symbol $\succeq$ means preferred to. Let $\theta = (\theta_1, \ldots, \theta_n) \in \Theta^n$, where Θ is the set of types, assumed to be the same for all players. As we are dealing with mechanisms $\mathcal{M}(a) \in \Delta(X)$, we extend the domain of the definition of $u_j(\theta_j, .)$ to the set $\Delta(X)$ of probabilities on outcomes by considering that, for all $\theta_j \in \Theta$ and all $\tilde{x} \in \Delta(X)$,

$$u_j(\theta_j, \tilde{x}) = \mathbb{E}_{x \sim \tilde{x}}[u_j(\theta_j, x)],$$

where $x \sim \tilde{x}$ means that the random variable x follows the probability $\tilde{x}$.

We assume that each player knows his type and has incomplete knowledge of the other players' types. Denote by $\tilde{\theta} \in \Delta(\Theta^n)$ the probability distribution on types of all players, and by $\tilde{\theta}_{-j} \in \Delta(\Theta^{-j})$ the probability distribution on all players' types but j's. These probabilities are called beliefs.

Remark 1 *For clarity of exposition, we are assuming here that the set of actions, the set of types and the beliefs are the same for all players. Admittedly, this is not the most general formulation, but in principle, there is no conceptual difficulty in extending the analysis to the case where the players have different action sets. In particular, the assumption that the set of types is the same for all players can easily be relaxed by defining a set Θ_j for each $j \in N$. As we will be considering beliefs, this would then be equivalent to saying that each player has a set of types $\Theta = \bigcup_{j \in N} \Theta_j$ with a nil distribution over $\Theta - \Theta_j$.*

Remark 2 *Classical modeling of Bayesian games fix $X = A^n$ and $\mathcal{M} = id_{|A^n}$, i.e., the mechanism is simply the identity of A^n. Utilities $u_j(\theta_j, a_1, \ldots, a_n)$ are then functions of the type and the profile of actions. However, for many problems such as the cake-cutting problem, it is more natural to consider a set of outcomes, which is really what players are interested in, i.e., the actual parts of the cake that are allocated, rather than the announced preferences of the other players.*

Denote by σ_j a strategy of player j, that is, a mapping that associates an action a_j to a type, i.e., $\sigma_j(\theta_j) = a_j$. Denote by $\Sigma = \big(\Delta(A)\big)^{\Theta}$ the set of strategies of this player, where $\Delta(A)$ denotes the space of probability distributions on A. As for the set of actions, we assume, without any loss of generality, that the set of strategies is the same for all players. Let $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Sigma^n$.

The relationships among all the variables defined so far are shown in Figure 1.

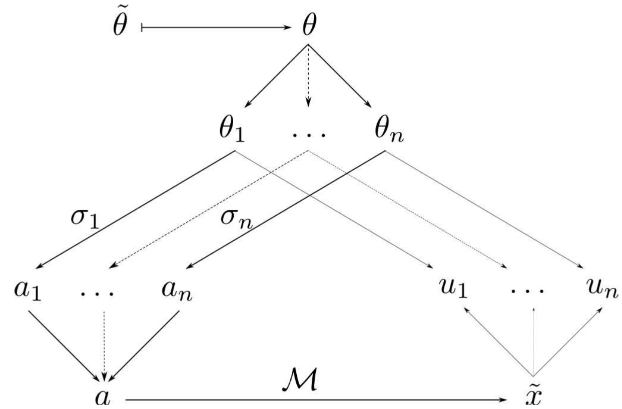

Figure 1: Variables

We end this section by recalling the definitions of best-reply (BR) strategies and Bayesian-Nash (BN) equilibria.

Definition 1 *The set of best-reply strategies σ_j^{BR} for player j to strategy profile σ_{-j} is given by*

$$BR_{\Sigma,j}(\sigma_{-j}) = \arg\max_{\sigma_j^{BR} \in \Sigma} \ \mathbb{E}_{\theta_j \sim \tilde{\theta}_j}\left[u_j\left(\theta_j, \mathcal{M}(\sigma_j^{BR}(\theta_j), \sigma_{-j}(\theta_{-j}))\right)\right]. \tag{1}$$

The set $BR_\Sigma(\sigma)$ of vectors σ^{BR} of best-reply strategy profiles for all players is given by

$$BR_\Sigma(\sigma) = \Big\{\sigma^{BR} \in \Sigma^n \mid \forall j \in N, \sigma_j^{BR} \in BR_{\Sigma,j}(\sigma_{-j})\Big\}. \tag{2}$$

The notation BR_Σ is used to denote that this is an operator that takes elements from Σ and outputs a subset of Σ. This remark will be useful when we get to the return-function best-replies.

Definition 2 *The set BN_Σ of Bayesian-Nash equilibria is the set of strategy profiles that are best replies against themselves, i.e.,*

$$BN_\Sigma = \Big\{\sigma^{BN} \in \Sigma^n \mid \sigma^{BN} \in BR_\Sigma(\sigma^{BN})\Big\}. \tag{3}$$

2.1 Return Function

For a strategy profile $\sigma_{-j} \in \Sigma^{-j}$, we associate to player j the return function $\varphi_j^\sigma(\cdot)$. This function maps an action a_j of player j to the induced probability of outcome. More precisely, $\varphi_j^\sigma(a_j)$ is the average of the probabilities of outcomes when player j plays action a_j and other players use strategies σ_{-j}, and when other player types follow the belief $\tilde{\theta}_{-j}$, i.e.,

$$\varphi_j^\sigma(a_j) = \mathbb{E}_{\theta_{-j} \sim \tilde{\theta}_{-j}}[\mathcal{M}(a_j, \sigma_{-j}(\theta_{-j}))]. \tag{4}$$

Intuitively, the return function $\varphi_j^\sigma(\cdot)$ answers the following question asked by player j: *If I do that, what could happen and with what probability?* Answering this question then enables the player to choose the best-reply action to his return function. This means, as we shall see, that the function contains all the information required to define the concept of a best-reply, and thus of Bayesian-Nash equilibria.

Note that the return function is itself defined by the other players' strategies, as well as by the belief $\tilde{\theta}$ and the mechanism $\mathcal{M}$. However, since the belief $\tilde{\theta}$ and the mechanism $\mathcal{M}$ are fixed once and for all (i.e., no updating is required), we simplify the notation by not expliciting the dependencies of the return function on these objects.

As each return function $\varphi_j^\sigma = \varphi_j^\sigma(\cdot)$ is a function that to an action a_j associates a probability $\tilde{x} \in \Delta(X)$ on outcomes, it is an object of the space $\Phi = \Delta(X)^A$. In fact, we can extend our concept of return functions to any function $\varphi \in \Phi$ that maps actions to probabilities on outcomes, even when φ is not obtained from a strategy profile σ. This will be of interest to us later, as we will be focusing on the natural topology of the space Φ of return functions to prove the convergence of our algorithms for the computation of Bayesian-Nash equilibria.

The fact that some return functions are deduced from a strategy profile can then be reinterpreted by the mapping $\phi : \Sigma^n \to \Phi^n$, which maps strategy profiles to return-function profiles for all players. This mapping is defined by

$$\phi : \sigma = (\sigma_1, \ldots, \sigma_n) \mapsto \varphi^\sigma = (\varphi_1^\sigma(\cdot), \ldots, \varphi_n^\sigma(\cdot)). \tag{5}$$

Before proceeding further, we would like to highlight three important features of the return functions: First, they avoid some of the inherent complexities related to the actions profile and to the mechanism. More precisely, given the knowledge of σ, it is in practice very hard to compute best-reply strategies, especially if the number of players is large, the beliefs are not simple distributions or if the mechanism is non-analytical. We stress that even in such complex (and relevant in practice) cases, the return functions can still perform the task of computing best-reply strategies. Second, if the actions and outcomes are public information, then the return functions become observable objects. Consequently, it is easy for a player to iteratively approximate his own return function. Put differently, our approach does not require to estimate other players' strategies, which is a hard task as types often remain private information, even after the end of the game. Third, from Figure 2, we clearly see that there is a correspondence (mapping ϕ) between the best-reply strategies and the best-reply return functions. The same can be noted for the Bayesian-Nash equilibrium (see Theorem 1 below for a more precise statement). This means that all the classical descriptions of Bayesian games using strategy profiles can be translated in terms of return functions. Let us now construct this translation.

Given a return function profile φ, the set of best-reply strategies for player j is given by

$$BR_{\Phi\to\Sigma,j}(\varphi) = \arg\max_{\sigma_j \in \Sigma} \mathbb{E}_{\theta_j \sim \tilde{\theta}_j}\left[u_j\left(\theta_j, \varphi_j(\sigma_j(\theta_j))\right)\right]. \tag{6}$$

From this, similarly to what has been done earlier, we define the set of best-reply strategy profiles $BR_{\Phi\to\Sigma}(\varphi)$. The notation $BR_{\Phi\to\Sigma}$ highlights the fact that it is a correspondence from Φ to Σ as depicted in Figure 2. Observe that this definition is consistent with previous ones for best-reply strategy profiles, because we have $BR_{\Phi\to\Sigma}(\phi(\sigma)) = BR_\Sigma(\sigma)$.

We can now reformulate the Bayesian game by focusing on return functions only, that is, by defining the set of best-reply return-function profiles to a return-function profile as follows:

$$BR_\Phi(\varphi) = \phi(BR_{\Phi\to\Sigma}(\varphi)) = \{\phi(\sigma) \in \Phi^n \mid \sigma \in BR_{\Phi\to\Sigma}(\varphi)\}. \tag{7}$$

A return-function profile $\varphi^{BN} \in BN_\Phi$ is then a Bayesian-Nash equilibrium if $\varphi^{BN} \in BR_\Phi(\varphi^{BN})$. Since $BR_{\Phi\to\Sigma}(\phi(\sigma)) = BR_\Sigma(\sigma)$, our construction leads to ϕ being a sort of morphism that preserves best-replies, in the sense that

$$\sigma^{BR} \in BR_\Sigma(\sigma) \Rightarrow \phi(\sigma^{BR}) \in BR_\Phi(\phi(\sigma)). \tag{8}$$

This property is represented graphically by the commutativity of the diagram of Figure 2, where double arrows represent correspondences, i.e., matching onto the power set of the output set.

In particular, this property implies the preservation of Bayesian-Nash equilibrium.

$$\begin{tikzcd}
\Sigma \arrow[r, Rightarrow, "BR_\Sigma"] \arrow[d, "\phi"'] & \Sigma \arrow[d, "\phi"] \\
\Phi \arrow[r, Rightarrow, "BR_\Phi"] \arrow[ur, Rightarrow, "BR_{\Phi\to\Sigma}"] & \Phi
\end{tikzcd}$$

Figure 2: Diagram depicting the best-reply correspondences

Theorem 1 *A return-function profile φ^{BN} is a Bayesian-Nash equilibrium if and only if there exists a Bayesian-Nash strategy profile $\sigma^{BN} \in BN_\Sigma$ such that $\varphi^{BN} = \phi(\sigma^{BN})$. In other words, $BN_\Phi = \phi(BN_\Sigma)$.*

Proof. Let $\varphi \in BN_\Phi$. Then, $\varphi \in BR_\Phi(\varphi)$, which means that $\varphi \in \phi(BR_{\Phi\to\Sigma}(\varphi))$. Thus, there exists $\sigma \in BR_{\Phi\to\Sigma}(\varphi)$ such that $\varphi = \phi(\sigma)$. But $BR_{\Phi\to\Sigma}(\varphi) = BR_{\Phi\to\Sigma}(\phi(\sigma))$, thus $\sigma \in BR_{\Phi\to\Sigma}(\phi(\sigma)) = BR_\Sigma(\sigma)$. Therefore, $\sigma \in BN_\Sigma$, which means that $\varphi \in \phi(BN_\Sigma)$, and proves the first inclusion $BN_\Phi \subseteq \phi(BN_\Sigma)$.

Reciprocally, assume $\varphi \in \phi(BN_\Sigma)$. Then, $\varphi = \phi(\sigma)$, where $\sigma \in BN_\Sigma$. Thus, $\sigma \in BR_\Sigma(\sigma)$, which means that $\sigma \in BR_{\Phi\to\Sigma}(\phi(\sigma)) = BR_{\Phi\to\Sigma}(\varphi)$. Thus, $\phi(\sigma) = \varphi \in \phi(BR_{\Phi\to\Sigma}(\varphi)) = BR_\Phi(\varphi)$. This proves the second inclusion $BN_\Phi \supseteq \phi(BN_\Sigma)$, and concludes the proof. □

What is more, a sequence $\{\sigma^i\}_{i\in\mathbb{N}}$, where $\sigma^{i+1} \in BR_\Sigma(\sigma^i)$, is associated to the sequence $\varphi^i = \phi(\sigma^i) = \varphi_{\sigma^i}$, which then satisfies $\varphi^{i+1} \in BR_\Phi(\varphi^i)$. This sequence defines a best-reply dynamics for return functions, which will be studied in the last section of this paper. It is this dynamics that is a generalization of mean-field games and of the approach used by Rabinovich et al. (2013).

Remark 3 *An example of the use of return functions that has been widely studied recently are mean-field games. In particular, in linear-quadratic mean-field games, the utility function of a player only depends on the average actions of the other players. The method used to solve these games can be seen as using the return function $\varphi_j(a_j) = (a_j, m)$, where m is the average trajectory of the other players and a_j is the trajectory of player j. What is then done is to compute the average trajectory of best-reply trajectories to the trajectory m. In other words, what is computed is precisely the best-reply return function. In such a case, the large-number-of-players hypothesis and the simplicity of the mechanism make the return function very simple. The difficulty of mean-field games lies rather in the complexity of the action space.*

Remark 4 *In Rabinovich et al. (2013), the authors compute Bayesian-Nash equilibria when the set of actions is finite and the game is symmetric. Their approach consists in computing the probability distribution of each action played. This corresponds to defining $X = A^n$, $\mathcal{M} = id_{|A^n}$ and $\varphi(a) = (a, h_\sigma(\cdot))$, where $h_\sigma(\cdot)$ maps each action a' to the probability $h_\sigma(a')$ that a player will choose this action. Later on, the authors retain the case where the utility functions are affine in the type $\theta \in [0,1]$. In such a setting, the outcome, that is, the relevant information for the players, is reduced to the slope s and the intercept ι of the utility function. Thus, their approach corresponds to using the return function $\varphi : a \mapsto (s, \iota) \in \mathbb{R}^2$. Since the set of actions is finite, the return function is equivalently represented by the vector $L = (s_a, \iota_a)_{a\in A} \in (\mathbb{R}^2)^A$. Next, the authors provide an algorithm to compute ϵ-equilibria of a large class of auction problems.*

Further, if ϕ is continuous, then the convergence of the return-function sequence is implied by the convergence of the strategy-profile sequence.

Theorem 2 *If the mapping ϕ is continuous and if the sequence σ^i converges to a Bayesian-Nash equilibrium σ^{BN}, then the sequence $\varphi^i = \phi(\sigma^i)$ also converges to a Bayesian-Nash equilibrium.*

Proof. Since ϕ is continuous, $\varphi^i = \phi(\sigma^i)$ converges towards $\phi(\sigma^{BN})$. Since $\sigma^{BN} \in BN_\Sigma$, we have $\phi(\sigma^{BN}) \in \phi(BN_\Sigma) = BN_\Phi$, which proves the theorem. □

In the last section of this paper, we will show that for natural topologies on the two spaces (strategies and return functions), and given assumptions on the mechanism $\mathcal{M}$, the function ϕ is continuous.

Remark 5 *In some Bayesian games, the types of the players are interdependent, and therefore the type of player could be an argument of the return function. For simplicity, we will not explicit such a dependence here, but there is a priori no conceptual difficulty to include such a relation in this model. Clearly, this would complicate the estimation of the return function, but, depending on the structure of the problem, its interpolation might still not be very difficult.*

Remark 6 *When all players j have the same beliefs θ_{-j} about the other players' types, it is natural to assume that $\varphi_j(\cdot) = \varphi(\cdot), \forall j \in N$. This would occur in, e.g., games with a large number of players, where excluding a player will not fundamentally affect the probabilistic representation of the game for the other players. Clearly, in such a case, the mechanism must be symmetric. We will provide an illustrative example with symmetric players.*

3 Computation of Equilibria with the Return Function

The objective of this section is to show how the Bayesian-Nash equilibrium can be computed using the return function. To do this, we will iteratively compute the return function using a learning process. Our approach is similar to the idea of fictitious play, however with an important difference. Indeed, whereas in fictitious play, we determine a player's best reply to the other players' mixed strategies, here, we compute the best action for each player, given his evaluation of how his actions affected the outcome in the previous iteration.

3.1 Computing the Return Function

To compute Bayesian-Nash equilibria, we propose to compute a sequence $(\varphi^i)_{i\in\mathbb{N}}$ of return function profiles, following the idea of having $\varphi^{i+1} \in \Phi^{BR}(\varphi^i)$. Instead of directly considering the players' strategies, our approach consists, at each iteration, in choosing a type profile θ according to the belief $\tilde{\theta}$, and then maximizing $u_j\left(\theta_j, \varphi_j^i(a_j)\right)$ with respect to a_j. Algorithm 1 develops this idea, and in particular, the learning aspect involved in the computations:

Algorithm 1 Optimization of the parameterized mechanism

while the convergence criterion is not verified **do**
- (1) Generate a type profile θ^i according to belief $\tilde{\theta}$.
- (2) Compute the actions a_j^i maximizing the utility of the approximated return function $u_j^i\left(\theta_j, \varphi_j^i(a_j^i)\right)$.
- (3) Given actions a^i, compute an outcome x^i, according to probability $\mathcal{M}\left(a^i\right)$.
- (4) Update the return functions φ_j^{i+1} given the actions a_j^i and the outcome x^i.
- (5) Increment i.

end while

Remark 7 *In the second step, the maximizer may not be unique. As we shall see later in Section 5.3, such cases are in fact rare, and when they occur, we will retain the first maximizer found by the algorithm.*

If the convergence is met, then the resulting Bayesian-Nash equilibrium is the best-reply strategy to the computed return function.[2] The first and third steps of this algorithm are only as difficult as the problem's inputs, that is, the belief $\tilde{\theta}$ and the mechanism $\mathcal{M}$. In some problems, the mechanism may require solving a large optimization problem, which may be quite difficult in itself. That being said, it is the second step that is most demanding in most applications, as it involves n non-linear optimization problems with a potentially

[2] It is worth mentioning that this algorithm can easily be parallelized, to speed up the learning process.

large action space. To get around some of these difficulties, we will use some local optimization tools (e.g., a gradient-based method), interpolation and perturbation schemes. Finally, the fourth step represents the learning process used by a player to update his return function.

3.2 Implementation

Let us first note that the estimation implementation procedure for the return function is simple and could be highly improved by using sophisticated optimization methods, but this is not a main concern in this paper. Technically speaking, the return function φ_j for player $j \in N$ has been computed by a mapping container with keys $a^i_j \in A$. Each key a^i_j is associated to a set of observations $(a^i_j, x^i) \in A \times X$ for player j at iteration i. In addition, each observation is associated to a reliability R^i. The more recent the observation is, the greater the reliability.

By never excluding past observations, 1 does not really produce the sequence whose convergence will be studied in Section 4, as φ^{i+1} is to be considered an average of φ^i and $BR_\Phi(\varphi^i)$. However, by choosing $R^i = i$, the relative weight of new observations increases, without being so overwhelming as to impose a huge variance due to the randomness of each measure. For these reasons, we should expect a similar convergence to the one we proved for the best-reply dynamics. This intuition is backed up by the following theorem.

Theorem 3 *Let U^i be a sequence of independent random variables on a bounded interval $[0, U]$, V^i the sequence of weighted averages of the first i terms given by*

$$V^i = \Big(\sum_{k=1}^{i} R^i U^i\Big) \Big/ \Big(\sum_{k=1}^{i} R_i\Big). \tag{9}$$

If $R^i = i$ and the sequence $\mathbb{E}[U^i]$ converges, then, almost surely, $\lim V^i = \lim_{i\to\infty} \mathbb{E}[U^i]$.

Proof. The proof is given in Appendix. It is a variant of Kolmogorov's strong law. □

Now, the return function φ_j is essentially used to compute the utility function $u_j(\theta_j, \varphi_j(a_j))$. Since no observations for action a_j may be available, and, even if they were, we might not have enough of them to have a reliable measure of the return function, we therefore propose to interpolate the return function to compute $u_j(\theta_j, \varphi_j(a_j))$. This interpolation relies on an assumption of continuity of the return function φ_j in a_j. This is a realistic hypothesis, as the probability on types obtained through $\theta_{-j} \sim \tilde{\theta}_{-j}$ will smooth the return function through averaging. Denote by $d_A(a_j, a^i_j)$ a distance between the actions a_j and a^i_j, and by $w(d_A(a_j, a^i_j), R^i)$ a weight assigned to this distance, which is decreasing in $d_A(a_j, a^i_j)$.[3]

The interpolation we are using here is a mix of Radius-Based Function and nearest point interpolation. Let

$$B^i_A(a_j, \alpha) = \{a^k_j \in A \mid k \leq i \text{ and } d_A(a_j, a^k_j) \leq \alpha\},$$

be the ball of center a_j and radius α, where α is an arbitrary positive number. Given $\theta_j \in \Theta$, the interpolated utility of player j when he chooses action a_j is estimated by

$$\hat{u}_j(\theta_j, \varphi^i_j(a_j)) = \frac{\sum\limits_{a^k_j \in B^i_A(a_j,\alpha)} w(d_A(a_j, a^k_j), R^k)\ u_j(\theta_j, \varphi^i_j(a^k_j))}{\sum\limits_{a^k_j \in B^i_A(a_j,\alpha)} w(d_A(a_j, a^k_j), R^k)}. \tag{10}$$

Remark 8 *To avoid being stuck at a non-optimal solution, we perturb the action profile obtained at the last iteration, and use it to update the return function. This will be done by simply adding a random number drawn from a given interval to each element of the last action vector. As observations pile up, this interval is narrowed.*

[3]To illustrate, in Section 4, we will take $w(d\left(a, a^i, R^i\right)) = R^i/(1 + d\left(a, a^i\right))^2$.

4 Illustrative Example: A Cake-Cutting Problem

To illustrate the theory developed above, we provide an example of a cake-cutting problem, which is of great use in operations research, as it allows the modelling of an assignment problem that includes the preferences of the agents involved. For instance, some shift-scheduling and matching problems exactly fit the cake-cutting problem (CCP) formalism. This problem can be stated as follows: given a cake and a set of players having additive utility functions over the subsets of the cake, the problem is how to allocate the cake to optimize a certain objective, while satisfying some constraints, such as fairness. This class of problem has been the subject of numerous papers; see, e.g., Brams and Taylor (1995), Brams and Taylor (1996), Robertson and Webb (1998), and recent papers, e.g., Mossel and Tamuz (2010) and Chen et al. (2010).

4.1 The Model

To simplify the computation, while still being able to illustrate our approach, we suppose that the cake initially contains a set K of homogenous portions. An outcome is a matrix

$$x = \{x_{jk}\} \in [0,1]^{N\times K},$$

where x_{jk} is the portion $k \in K$ allocated to player $j \in N$. Note that the set X of admissible outcomes (allocations) is given by

$$X = \{x \in [0,1]^{N\times K} | \forall k \in K, \sum_{j\in N} x_{jk} \leq 1\}.$$

Denote by $\theta_{jk} \geq 0$ the utility of player $j \in N$ for portion $k \in K$. Thus, the type of player j is the vector θ_j, and his utility function over any subset x_j of the cake is $u_j(\theta_j, x_j) = \sum_{k\in K} \theta_{jk} x_{jk}$. We normalize the utility function of each player for the whole cake to one. Therefore, if $1_K = \{1\}_{k\in K}$ is the allocation of the cake, then $u_j(\theta_j, 1_K) = \sum_{k\in K} \theta_{jk} = 1$. As a result, the set Θ of types is the polytope $\Theta = \{\theta_j \in \mathbb{R}_+^K \mid \sum_{k\in K} \theta_{jk} = 1\}$.

One approach to allocating the cake would be to solve the following linear optimization problem:

$$\begin{aligned} &\max_x \; u \\ \text{subject to} : &\sum_{k\in K} a_{jk} x_{jk} = u, \quad \forall j \in N, \\ &x \in X. \end{aligned} \tag{11}$$

We will refer to this approach as mechanism $\mathcal{M}$, that is, the mechanism that associates an outcome x to a vector a of actions. We define an admissible action for player j to be one that satisfies the constraints $a_{jk} \geq 0, k \in K$ and $\sum_{k\in K} a_{jk} = 1$. The normalization of the admissibility constraints on the types and actions implies that the set of actions and types are the same, i.e, $A = \Theta$. Note that if the above optimization problem has multiple solutions, then we will simply choose one of them randomly.

An example of a belief $\tilde{\theta}$ is obtained through the following process: for each player $j \in N$, we randomly generate a vector in $[0,1]^K$ according to the uniform distribution. Then, we normalize this vector by dividing each component by the sum of all components, hence obtaining a vector $\theta_j \in \Theta$. Note that this mechanism $\mathcal{M}$ is symmetric, and the beliefs follow identical probability distributions due to our way of generating the vector $\theta_j, j \in N$. A consequence of these symmetries is that the return function φ_j is independent of j, and therefore, we have $\varphi_j = \varphi$ for all $j \in N$.

4.2 A Simple Two-Player Example

In principle, the above problem can be solved for any number of players and any set K. However, for the sake of graphical representation and analytical calculation, we first focus on the simplest possible setting of two players, and $|K| = 2$. Let $a_1 = (a_{11}, a_{12})$ and $a_2 = (a_{21}, a_{22})$. In this context, problem 11 becomes

$$\max \ u \tag{12}$$
$$\text{subject to :} u = a_{11}x_{11} + (1-a_{11})x_{12}, \tag{13}$$
$$u = a_{21}(1-x_{11}) + (1-a_{21})(1-x_{12}), \tag{14}$$
$$x_{11}, x_{12} \geq 0, \tag{15}$$

where x_{jk} represents the portion $k = 1, 2$ allocated to player $j = 1, 2$. The description of the type θ_j and the action a_j of each player j can now be reduced to one real variable each in $[0, 1]$ that is, θ_{j1} and a_{j1}. As a result, a strategy σ_j is a mapping from $[0, 1]$ into $[0, 1]$.

To simplify the theoretical analysis, we assume that the beliefs $\tilde{\theta}$ about types are obtained by drawing uniformly and randomly θ_{j1} in $[0, 1]$, hence obtaining $\theta_{j2} = 1 - \theta_{j1}$. Also, we are re-scaling values of θ and a by multiplying them all by 100. The equations (13, 14) then imply the equalities

$$(a_{11} + a_{21})x_{11} = a_{21} + (100 - a_{21})(1 - x_{12}) - (100 - a_{11})x_{12}, \text{ and} \tag{16}$$
$$(200 - a_{11} - a_{21})x_{12} = (100 - a_{21}) + a_{21}(1 - x_{11}) - a_{11}x_{11} \tag{17}$$

Using Algorithm 1, we compute a Bayesian-Nash equilibrium. We have proceeded to 500 iterations. In the 10 first iterations, players played truthfully, so that we could compute the truthful return function. Afterwards, at each iteration, players' types were drawn randomly, and their actions were computed by choosing the best-replies to the return function, given their types. Actions were then used to compute the outcome, which was then used to feed the return function.

Figure 3 represents the equilibrium strategies σ^{BN}. More precisely, it stands for the best-reply strategy to the return function obtained after 500 iterations. In Figure 3, 200 dots have been obtained by randomly drawing a player's type, and computing the best-reply action to the return function obtained after 500 iterations. The x-axis is the player's weight for a portion, while the y-axis is his announced weight for the portion. Different portions $k \in K$ of the cake are depicted by different colors of dots. The dotted diagonal line corresponds to the truthful strategy σ_1^{truth}. As we can see in this figure, the equilibrium strategy σ_j^{BN} consists of overvaluing the portion he desires less, and undervaluing the portion he prefers. Let $\varphi^{BN} = \phi(\sigma^{BN})$.

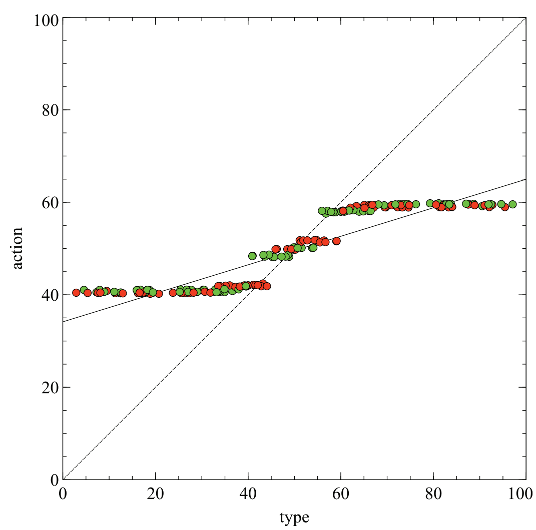

Figure 3: Bayesian-Nash Equilibrium with 2 Players and 2 Attributes

The plain line in the figure represents a linear regression of the equilibrium strategy. It is given by $\sigma_j^{BN}(\theta_j) = 35 + 0.3\theta_j$. Interestingly, this gives us a measure to analyze the convergence of the algorithm. Figure 4 displays the evolution of the linear regression slopes for every 10 iterations, as well as the corresponding average sum of squares.

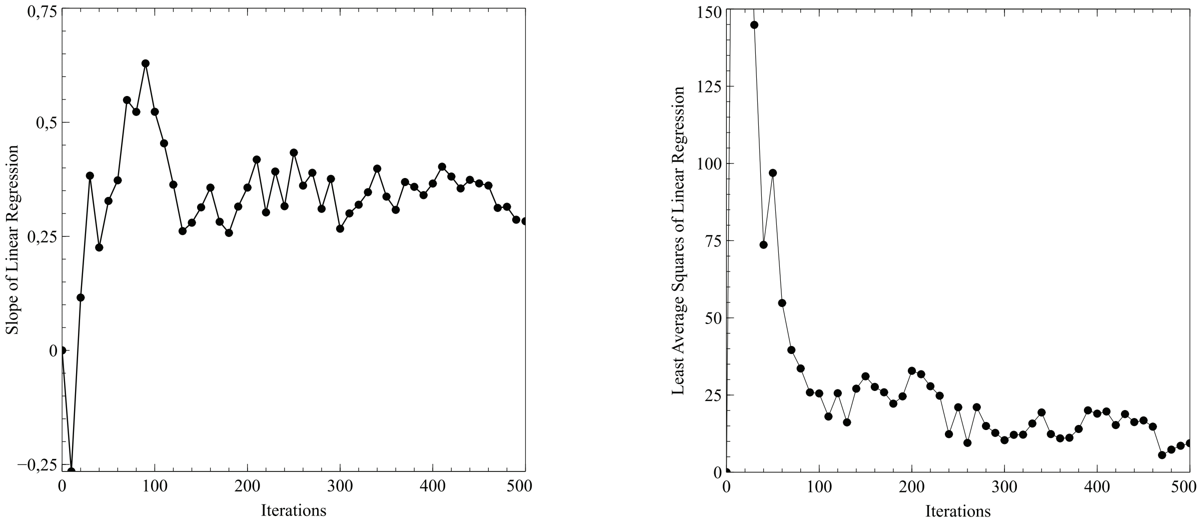

Figure 4: Slopes of Linear Regressions of Best-Replies to Return Function as Iterations Increase

It is important to notice that these figures show some convergence of the algorithm. While the slopes stabilize at about 0.3, the average sum of squares decreases to close to 0. However, it is to be expected that this average sum will actually never reach zero. After all, it is clear from Figure 3 that the Bayesian-Nash equilibrium does not involve affine strategies.

4.3 Analytical Analysis

In this simple example of two players and two attributes, using equations (16, 17), the mechanism can be drawn as in Figure 5. The main diagonal of the square corresponds to the two players having identical utility functions, in which case there are numerous solutions, all yielding utilities of 50 for both players. The other diagonal is another discontinuity of the mechanism, which appears when the first player's preference for the first portion equals the second player's preference for the second portion. The two diagonals define four areas. In each area, the mechanism is defined analytically as a function of the players' actions, as represented in Figure 5. This will enable us to verify the validity of our numerical determination of the Bayesian-

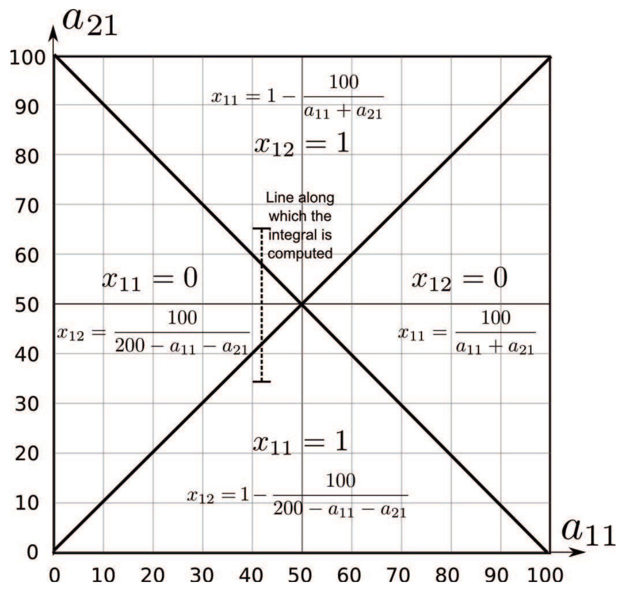

Figure 5: Mechanism of the Cake-Cutting Problem with 2 Attributes and 2 Players

Nash equilibrium; first, by analytically computing the best-reply strategy to the interpolated Bayesian-Nash strategy, and next, by showing that this analytical best-reply strategy is close to the interpolated Bayesian-Nash strategy.

A linear approximation of the strategy in Figure 3 consists of choosing the action a_{11} such that

$$a_{11} = \sigma^{BN}(\theta_{11}) = a^*_{min} + (100 - 2a^*_{min})\theta_{11}.$$

The numerical value of a^*_{min} is approximately 35.

Now, supposing that player 2 chooses to play the computed Bayesian-Nash equilibrium strategy, we can evaluate the return function for player 1, and write his payoff if he chooses action a_{11} and has a type θ_1 :

$$u_1(\theta_1, \varphi^{BN}(a_1)) = \int_{\theta_2} u_1(\theta_{11}, \mathcal{M}(a_1, \sigma^{BN}(\theta_2)))d\tilde{\theta}_2, \tag{18}$$

$$= \int_{a_2=a^*_{min}}^{100-a^*_{min}} \frac{u_1(\theta_{11}, \mathcal{M}(a_1, a_2))}{100 - 2a^*_{min}} da_2. \tag{19}$$

If $a_{11} \leq 50$, we can separate the integral into the cases where $a_{21} \leq a_{11}$, $a_{11} \leq a_{21} \leq 100 - a_{11}$ and $a_{21} \geq 100 - a_{11}$. The three resulting terms are

$$\frac{1}{100 - 2a^*_{min}} \left[\int_{a^*_{min}}^{\theta_{11}} \left[\theta_{11} + (100 - \theta_{11})(1 - \frac{100}{200 - a_{11} - a_{21}}) \right] da_{21} \right], \tag{20}$$

$$\frac{1}{100 - 2a^*_{min}} \left[\int_{a_{11}}^{100-a_{11}} (100 - \theta_{11}) \frac{100}{200 - a_{11} - a_{21}} da_{21} \right], \tag{21}$$

$$\frac{1}{100 - 2a^*_{min}} \left[\int_{100-a_{11}}^{100-a^*_{min}} \left[\theta_{11}(1 - \frac{100}{a_{11} + a_{21}}) + (100 - \theta_{11}) \right] da_{21} \right]. \tag{22}$$

Calculating these integrals yields

$$(100 - 2a^*_{min})u_1(\theta_{11}, \varphi^{BN}(a_1)) =$$
$$2(a_{11} - a^*_{min}) + \ln \frac{4(100 - a_{11})^2}{200 - a_{11} - a^*_{min}} + \theta_{11} \ln \frac{200 - a_{11} - a^*_{min}}{4(100 - a_{11})^2(100 + a_{11} - a^*_{min})}. \tag{23}$$

Therefore, the level curves, for which $u_1(\theta_{11}, \varphi^{BN}(a_1))$ is constant, are described by the following equation:

$$\theta_{11} = \frac{k - 2a_{11} - \ln \frac{4(100-a_{11})^2}{200-a_{11}-a^*_{min}}}{\ln \frac{200-a_{11}-a^*_{min}}{4(100-a_{11})^2(100+a_{11}-a^*_{min})}}, \qquad \text{for } a_{11} \leq 1/2, \tag{24}$$

with $k = (100 - 2a^*_{min})u_1(\theta_{11}, \varphi^{BN}(a_1)) + 2a^*_{min}$. Therefore, we get $u_1(\theta_{11}, \varphi^{BN}(a_1)) = \frac{k-2a^*_{min}}{100-2a^*_{min}}$. The level curves for expected utility $u_1(\theta_{11}, \varphi^{BN}(a_1))$, as functions of θ_{11} and a_{11}, at the Bayesian-Nash equilibrium are shown in Figure 6, where an approximated best-reply affine strategy is drawn in red.

The best-reply action of player 1 is the action that maximizes the expected value of $u_1(\theta_{11}, \varphi^{BN}(a_1))$, given a value of θ_{11}. Therefore, if $\theta_{11} = 20$ for instance, the largest value he can obtain for $u_1(\theta_{11}, \varphi^{BN}(a_1))$ is 75, which is reached by choosing action $a_{11} = 35$. As we can easily see, the best-reply strategy does not coincide with the computed Bayesian-Nash equilibrium, but is very close to it. This shows that the use of the return-function method has been efficient for computing the Bayesian-Nash equilibrium.

4.4 The Non-Linear Component

An interesting observation can be made regarding a pattern in the Bayesian-Nash equilibrium strategies. Indeed, a closer look at the best-reply curves in Figure 3 reveals that they are S-shaped (like an $arctan(\cdot)$ function). Computing the best-reply strategy to the interpolated strategy in Figure 3, we again get this S-shaped curve as shown in Figure 7.

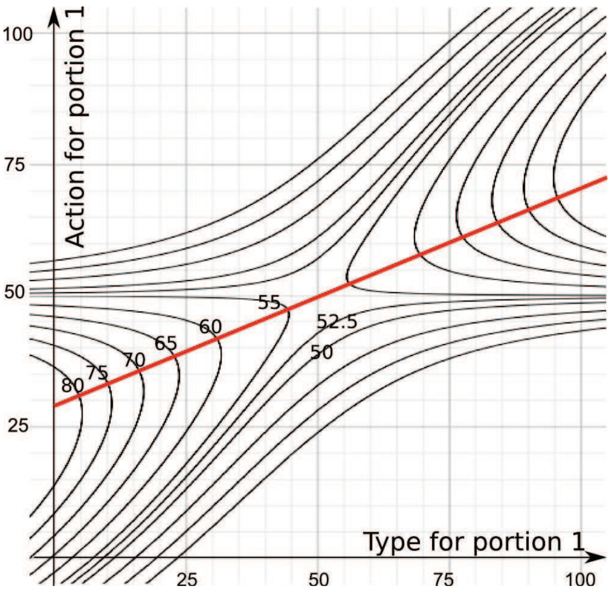

Figure 6: Level Curves of a Player's Expected Utilities at a Bayesian-Nash Equilibrium

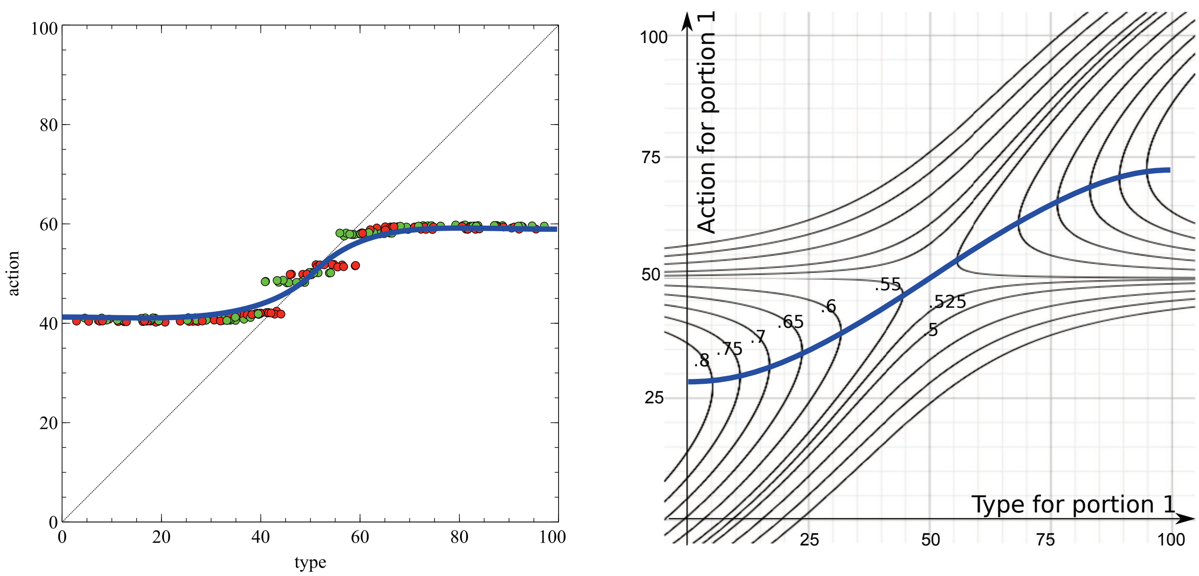

Figure 7: S-Shape of Computed Bayesian-Nash Equilibrium and the Best-Reply Strategy to a Linear Regression of the Computed Bayesian-Nash Equilibrium

4.5 A More General Example

Now, we consider a more general case with 20 players and 3 attributes. As before, we use Algorithm 1 to derive a Bayesian-Nash return function after 100 iterations. Figure 8 displays the computed best-reply to this return function, and thus represents the Bayesian-Nash equilibrium strategy. Once again, we have plotted the announced weights as a function of the type weights. However, this time, because the sets of types and of actions are 2-dimensional (they are the simplex of the 3-dimensional space), the figure only displays a projection of the full strategy at Bayesian-Nash equilibria. This explains the fuzziness of the cloud of dots.

Again, the Bayesian-Nash equilibrium shows an overbid for portions of the cake that the players do not want, while the portions they do want are underbid. This shift between underbidding and overbidding occurs at 100/3, which corresponds to the player liking the portion just as much as the average of other players.

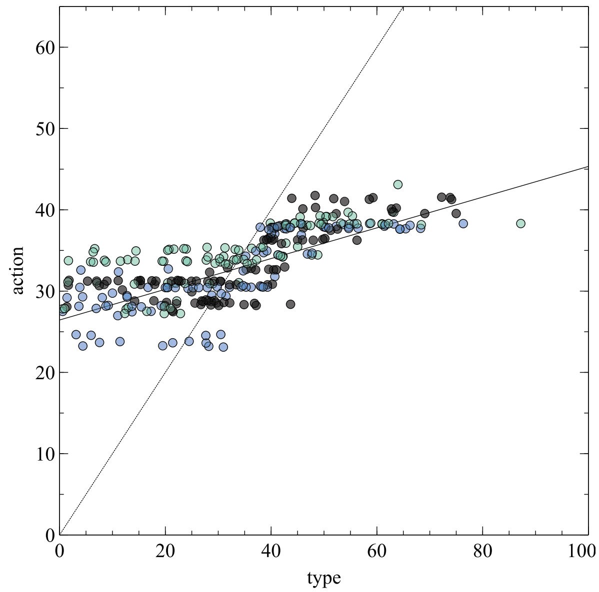

Figure 8: Bayesian-Nash Equilibrium with 20 Players and 3 Attributes

In this more complex setting, analytical calculations are too complicated. Indeed, both the belief and the mechanism are hard to write algebraically, and consequently, it is extremely difficult to algebraically compute this Bayesian-Nash equilibrium.

Similarly to what we did in the 2-player 2-attribute case, we computed the slope of the linear regression of Figure 8, as well as the average of the squares in Figure 9.

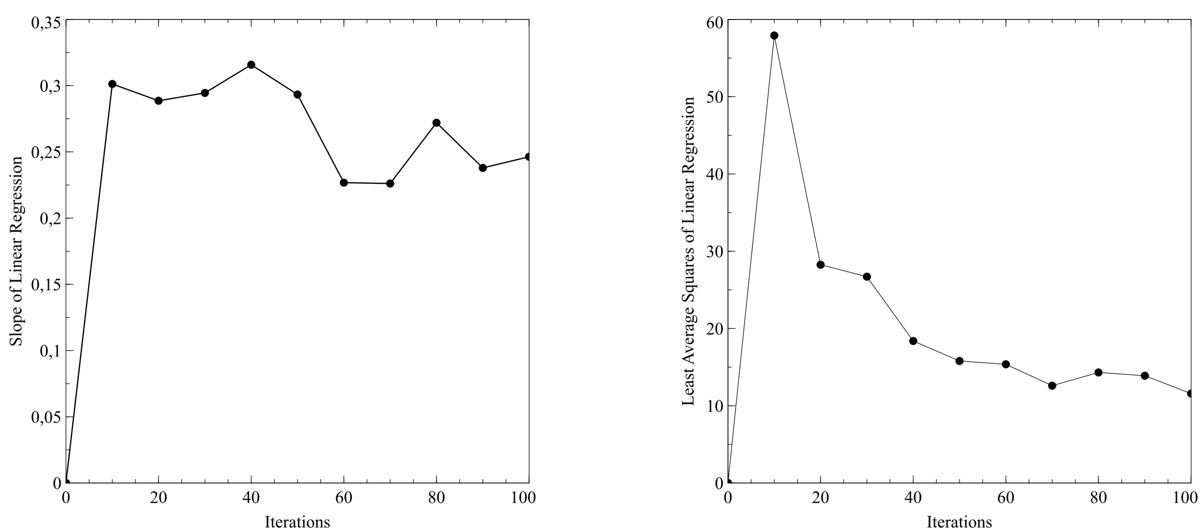

Figure 9: Slopes and Least Average Squares of Linear Regressions for 20 Players and 3 Attributes

We finish by commenting on the complexity of our algorithm. First, the algorithm's complexity time does not increase much with the number of players. Indeed, although the time it takes to perform one iteration of the loop of Algorithm 1 increases (only linearly) in the number of players, we generate many more observations at each iteration, which speeds up the approximation of the return function. Therefore, we expect to need fewer iterations in total to reach a Bayesian-Nash equilibrium. On the other hand, the complexity time increases more rapidly with the number of attributes. This is due to the the algorithm's second step, which involves an optimization over the set of actions. Obviously, the higher-dimensional is this set, the more time it takes to solve this optimization problem. Again, we mention that advanced methods for nonlinear optimization could be used to speed up the computation at this step.

5 Theoretical Convergence

In this last section, we provide a theoretical proof of the iterative algorithm's convergence to a Bayesian-Nash equilibrium. In particular, we show that, despite the cumulative error due to the approximations made at each successive iteration of the return function, we can guarantee its convergence to a Bayesian-Nash equilibrium.

We start with some necessary preliminaries of topology, define the metric topology in which the convergence is defined, and conclude with a convergence theorem. The next proposition introduces a metric topology on the set of outcome probabilities, which is the output set of the return function.

Proposition 1 *The following distance $d_{\Delta(X)}$ is a pseudometric*[4] *on $\Delta(X)$:*

$$d_{\Delta(X)}(\tilde{x}^1, \tilde{x}^2) = \sup_{j\in N, \theta_j \in \Theta_j} |u(\theta_j, \tilde{x}^1) - u(\theta_j, \tilde{x}^2)|. \tag{25}$$

Proof. The function $d_{\Delta(X)}$ is obviously non-negative and symmetric, and it clearly satisfies the triangle inequality. □

The natural topology on return-function profiles is given by the following metric:

Proposition 2 *The following distance d_Φ is a pseudometric on Φ:*

$$d_\Phi(\varphi^1, \varphi^2) = \sup_{j\in N, a_j \in A} d_{\Delta(X)}(\varphi^1(a_j), \varphi^2(a_j)). \tag{26}$$

Proof. The function d_Φ is obviously non-negative and symmetric, and it clearly satisfies the triangle inequality. □

From now on, Φ will refer to the space occurring after the metric identification corresponding to d_Φ is done. In particular, (Φ, d_Φ) is a metric space.

5.1 Continuity of ϕ

When we introduced the return function, we pointed out that the convergence of a sequence of best-replying strategy profiles could induce the convergence of the associated sequence of return-function profiles, provided that the mapping ϕ is continuous. In this section, we characterize the conditions under which this mapping is indeed continuous.

To talk about continuity of ϕ, we first need some topology on the set of strategy profiles. An important remark to be made first is that many strategy profiles yield the same outcomes. More precisely, if two strategy profiles lead to different actions for a set of types of probability zero, then they will be virtually identical. A usual way of formalizing this idea is by using quotient spaces, i.e., given a set of strategies, we group virtually identical strategies into a set. By doing this for all strategies, we divide the set of strategies into subsets, each containing virtually identical strategies. The quotient space is defined as the set of all these subsets. This construction is characterized by the following theorem.

Theorem 4 *Let us denote $\sigma^1 \sim \sigma^2$ if $\sigma^1(\theta)$ differs from $\sigma^2(\theta)$ on a set of nil probability according to $\tilde{\theta}$. Then, $\sim$ is a well-defined equivalence relation on Σ, and it defines the quotient space $\Sigma_\sim = \Sigma/\sim$. Moreover, the spaces $BR_\Sigma(\sigma)/\sim$ are well-defined in $\Sigma_\sim$ for all $\sigma \in \Sigma$ and do not depend on the representative σ of its equivalence class.*

[4]We use the term pseudometric to highlight the fact that the distance could be equal to zero, without necessarily involving the same object.

Proof. This is an immediate consequence of defining best-replies as the maximization of an expectation. Thus, nil probability spaces have no influence on a strategy's optimality. □

This theorem indicates that, to study best-reply dynamics and Bayesian-Nash equilibria, the right topology of strategy profiles should be defined on $\Sigma_\sim$.

Now, to go further, let us assume that there is some metric $d_{\Delta(A)}$ on $\Delta(A)$. There is no canonical way of choosing such a metric on the space of probability distribution $\Delta(A)$ (see, e.g., Gibbs and Su (2002) for a list of such metrics). For this reason, we will remain vague regarding the choice of the metric and simply ask that the following property be satisfied:

Hypothesis 1 *For any parameterized probabilities on actions $\tilde{a}^1(t)$ and $\tilde{a}^2(t)$, where the parameter t follows some probability distribution $\tilde{t}$, the metric $d_{\Delta(A)}$ satisfies*

$$d_{\Delta(A)}\Big(\mathbb{E}_{t\sim\tilde{t}}[\tilde{a}^1(t)], \mathbb{E}_{t\sim\tilde{t}}[\tilde{a}^2(t)]\Big) \leq \mathbb{E}_{t\sim\tilde{t}}\Big[d_{\Delta(A)}(\tilde{a}^1(t), \tilde{a}^2(t))\Big]. \tag{27}$$

This property is a sort of convexity property. It states that the distance between two means of probabilities on actions is smaller that the average distance between these two actions.

Remark 9 *If A is finite, then the space $\Delta(A)$ of probability distributions on A can be embedded into $\mathbb{R}^A$. Now, for any norm $||.||$ on $\mathbb{R}^A$, we have*

$$||\ \mathbb{E}[\tilde{a}^1(t)] - \mathbb{E}[\tilde{a}^2(t)]\ || = ||\ \mathbb{E}[\tilde{a}^1(t) - \tilde{a}^2(t)]\ || \leq \mathbb{E}\big[\ ||\tilde{a}^1(t) - \tilde{a}^2(t)||\ \big]. \tag{28}$$

This proves that, if A is finite, all norms on $\mathbb{R}^A$ define metrics on $\Delta(A)$ that satisfy Hypothesis 1.

We can now define a metric on $\Delta(A^n)$ using, for instance,

$$d_{\Delta(A^n)}(\tilde{a}^1, \tilde{a}^2) = \sup_{j\in N}\ d_{\Delta(A)}(\tilde{a}^1_j, \tilde{a}^2_j). \tag{29}$$

Based on this (or on a similar metric that would rather involve an expectation), we can define the following metric on $\Sigma_\sim$.

Proposition 3 *The distance d_Σ is a pseudometric on $\Sigma_\sim$ by*

$$d_\Sigma(\sigma^1, \sigma^2) = \sup_{j\in N}\ \mathbb{E}_{\theta_j\sim\tilde{\theta}_j}\Big[d_{\Delta(A)}(\sigma^1_j(\theta_j), \sigma^2_j(\theta_j))\Big].$$

Proof. The function d_Σ is obviously non-negative and symmetric, and it clearly satisfies the triangle inequality. □

This construction leads us to a framework to define the continuity of ϕ, and a sufficient condition to prove it. To do so, we need the following lemma.

Lemma 1 *For any strategies σ^1, σ^2, the following relation holds:*

$$\sup_{j\in N}\ d_{\Delta(A^n)}\Big((a_j, \sigma^1_{-j}(\tilde{\theta}_{-j})), (a_j, \sigma^2_{-j}(\tilde{\theta}_{-j}))\Big) \leq d_\Sigma(\sigma^1, \sigma^2). \tag{30}$$

Proof. First note that, for any player j, we have

$$d_{\Delta(A^n)}\Big((a_j, \sigma^1_{-j}(\tilde{\theta}_{-j})), (a_j, \sigma^2_{-j}(\tilde{\theta}_{-j}))\Big) = \sup_{i\neq j}\ d_{\Delta(A)}\big(\sigma^1_i(\tilde{\theta}_i), \sigma^2_i(\tilde{\theta}_i)\big). \tag{31}$$

Thus, by taking the supremum over all players j, we have

$$\sup_{j\in N}\ d_{\Delta(A^n)}\Big((a_j,\sigma^1_{-j}(\tilde\theta_{-j})),(a_j,\sigma^2_{-j}(\tilde\theta_{-j}))\Big)=\sup_{j\in N}\ d_{\Delta(A)}\big(\sigma^1_j(\tilde\theta_j),\sigma^2_j(\tilde\theta_j)\big). \tag{32}$$

Now, using Hypothesis 1 on the metric $d_{\Delta(A)}$, given that $\sigma^1_j(\tilde\theta_j)=\mathbb{E}_{\theta_j\sim\tilde\theta_j}[\sigma^1_j(\theta_j)]$, we have, for all j,

$$d_{\Delta(A)}\big(\sigma^1_j(\tilde\theta_j),\sigma^2_j(\tilde\theta_j)\big)\le\mathbb{E}_{\theta_j\sim\tilde\theta_j}\Big[d_{\Delta(A)}\big(\sigma^1_j(\theta_j),\sigma^2_j(\theta_j)\big)\Big]. \tag{33}$$

By taking the supremum over j in both sides, we obtain the exact formula in the lemma. □

Theorem 5 *If the mechanism $\mathcal{M}$ is uniformly continuous from $(\Delta(A^n),d_{\Delta(A^n)})$ to $(\Delta(X),d_{\Delta(X)})$, then so is the mapping ϕ from $(\Sigma_\sim,d_\Sigma)$ to (Φ,d_Φ).*

Proof. Let $\epsilon>0$. Since $\mathcal{M}$ is uniformly continuous, then there exists $\delta>0$ such that, for any probabilities $\tilde a^1,\tilde a^2$ on actions, we have

$$d_{\Delta(A)}(\tilde a^1,\tilde a^2)\le\delta\quad\Rightarrow\quad d_{\Delta(X)}(\mathcal{M}(\tilde a^1),\mathcal{M}(\tilde a^2))\le\epsilon. \tag{34}$$

Note that we have

$$d_{\Delta(X)}\Big(\big(\phi_j(\sigma^1)\big)(a_j),\big(\phi_j(\sigma^2)\big)(a_j)\Big)=d_{\Delta(X)}\Big(\mathcal{M}\big(a_j,\sigma^1_{-j}(\tilde\theta_{-j})\big),\mathcal{M}\big(a_j,\sigma^2_{-j}(\tilde\theta_{-j})\big)\Big) \tag{35}$$

Now, using Lemma 1 and equation (34), if $d_\Sigma(\sigma^1,\sigma^2)\le\delta$, then the right-hand side of equation (35) can be upper-bounded by ϵ. By taking the supremum over all players j and all actions $a_j\in A$, we obtain the inequality

$$d_\Phi(\phi(\sigma^1),\phi(\sigma^2))\le\epsilon. \tag{36}$$

This proves the uniform continuity of ϕ. □

The continuity of ϕ then implies that any convergence of a sequence $(\sigma^i)_{i\in\mathbb{N}}$ of strategy profiles implies the convergence of the sequence $(\phi(\sigma^i))_{i\in\mathbb{N}}$ of return function profiles. In other words, ϕ does not only preserve best-replies and Bayesian-Nash equilibria, it also preserves the topology of these spaces.

5.2 Best-Reply Dynamics

Let us now turn our attention to the best-reply dynamics. To start with, we assume that there exists some subspace U of Φ such that, for any $\varphi\in U$, the best-reply return function is uniquely defined and that $BR_\Phi(\varphi)\subset U$. With a slight abuse of language, this unique best-reply return function[5] is simply denoted by $BR_\Phi(\varphi)$. We will discuss the case of best-reply return functions that are not uniquely defined in section 5.3.

In our present setting, we thus have a mapping $BR_\Phi:U\to U$ that defines the best-reply dynamics by $(BR^i_\Phi(\varphi))_{i\in\mathbb{N}}$, where BR^i_Φ is the composition of i best-replies BR_Φ, i.e., $BR^i_\Phi(\varphi)=\underbrace{BR_\Phi(BR_\Phi(\ldots BR_\Phi}_{i\text{ times}}(\varphi)\ldots))$.

Remark 10 *The Bayesian-Nash return functions in U are the fixed points of BR.*

Theorem 6 *Let $(\sigma^i)_{i\in\mathbb{N}}$ be a sequence of uniquely defined best-reply strategy profiles that converges to a Bayesian-Nash equilibrium $\sigma^{BN}\in BN_\Sigma$, and $\varphi^i=\phi(\sigma^i)$. Assume ϕ continuous. Then, $BR(\varphi^i)$ is always uniquely well-defined, we have $\varphi^i=BR^i_\Phi(\varphi^0)$ and the sequence $(BR^i_\Phi(\varphi^0))_{i\in\mathbb{N}}$ converges towards a Bayesian-Nash return function.*

[5] To improve the chances of uniqueness, Φ is identified to the space separated by the pseudometric d_Φ, i.e. the space $\Phi/\sim$ where $\varphi^1\sim\varphi^2$ if $d_\Phi(\varphi^1,\varphi^2)=0$.

Proof. First, $BR_\Phi(\varphi^i) = \phi(BR_{\Phi\to\Sigma}(\phi(\sigma^i))) = \phi(BR_\Sigma(\sigma^i))$. Since $BR_\Sigma(\sigma^i)$ is uniquely well-defined by assumption, so is $\phi(BR_\Sigma(\sigma^i)) = BR_\Phi(\varphi^i)$.

Now, let us show by induction that $\varphi^i = BR_\Phi^i(\varphi^0)$. This equality holds trivially for $i = 0$. Assume it holds for i. We have $\varphi^{i+1} = \phi(\sigma^{i+1}) \in \phi(BR_\Sigma(\sigma^i)) = \phi(BR_{\Phi\to\Sigma}(\phi(\sigma^i))) = \phi(BR_{\Phi\to\Sigma}(\varphi^i)) = BR_\Phi(\varphi^i)$, which we know to be a singleton. Thus, $\varphi^{i+1} = BR_\Phi(\varphi^i) = BR_\Phi(BR_\Phi^i(\varphi^0)) = BR_\Phi^{i+1}(\varphi^0)$. This concludes the proof by induction.

Finally, the last part of the theorem is an immediate consequence of Theorem 2. □

This theorem shows that the return function is just as good as strategy profiles to describe best-reply strategies, Bayesian-Nash equilibria and iterations of best-reply strategies towards equilibria. From now on, we will only be working on return functions. Thus, we denote $BR = BR_\Phi$ and $BN = BN_\Phi$.

In Theorem 6, we assumed that the return function was computed exactly at each iteration. This is in fact not true as at each iteration an error is made. However, by sampling more and more at each iteration, we can guarantee that this error goes to 0. For the convergence to still hold under sampling errors, we need to introduce additional topological properties on the best-reply dynamics.

Hypothesis 2 *The best-reply function BR is a uniformly continuous function from $U \subseteq \Phi$ into itself.*

This hypothesis will be useful when we consider reciprocal images of convergence spaces.

Hypothesis 3 *All sequences $(BR^i(\varphi))_{i\in\mathbb{N}} \in U^{\mathbb{N}}$ converge uniformly, i.e.,*

$$\forall \epsilon > 0, K_\epsilon = \inf\Big\{K \in \mathbb{N}, \forall \varphi \in U, \exists \varphi^\infty \in BN \cap U, \forall i \geq K, BR^i(\varphi) \in B(\varphi^\infty, \epsilon)\Big\} < \infty. \tag{37}$$

Let us define BN_ϵ by

$$BN_\epsilon = \Big(\bigcup_{\varphi^{BN}\in BN\cap U} B(\varphi^{BN}, \epsilon)\Big) \cap U. \tag{38}$$

Then, Hypothesis 3 can be reinterpreted as $U = BR^{-K_\epsilon}(BN_\epsilon)$ for any $\rho > 0$, where $BR^{-i}(V)$ denotes the reciprocal image of V by the mapping BR^i.

Hypothesis 4 *Bayesian-Nash return functions of U are locally contracting points of BR, i.e.,*

$$\exists \nu, \rho > 0, \forall \varphi^{BN} \in BN, \forall \varphi \in B(\varphi^{BN}, \rho), d_\Phi(BR(\varphi), \varphi^{BN}) \leq (1-\nu) d_\Phi(\varphi, \varphi^{BN}). \tag{39}$$

This hypothesis is necessary to ensure that once the approximated sequence gets close to the Bayesian-Nash Equilibrium, it will remain close and actually converge.

Proposition 4 *Assume Hypotheses 3-4 hold true. Then, the sequence $(BR^i(\varphi))_{i\in\mathbb{N}}$ converges towards $\varphi^{BN} \in BN$ if and only if $\varphi \in BR^{-K_\rho}(B(\varphi^{BN}, \rho))$.*

Proof. If $\varphi \in BR^{-K_\rho}(B(\varphi^{BN}, \rho))$, then $BR^{K_\rho}(\varphi) \in B(\varphi^{BN}, \rho)$. The contracting hypothesis allows us to conclude the convergence of $BR^i(\varphi)$ towards φ^{BN}. Reciprocally, assume that $BR^i(\varphi)$ converges to φ^{BN}. Since $BR^{K_\rho}(\varphi)$ is necessarily in a contracting area, it will necessarily converge. For the limit to be φ^{BN}, it needs to be the contracting area of φ^{BN}. Thus, $BR^{K_\rho}(\varphi) \in B(\varphi^{BN}, \rho)$. This proves the inverse inclusion, and concludes the proof. □

Corollary 1 *Assume Hypotheses 3-4 hold true. If BR is continuous from U into itself, then the limit of sequences $BR^i(\varphi)$ depends only on the connected component of φ in U.*

Proof. The set of points that converge to an equilibrium φ^{BN} is the set $BR^{-K_\rho}(B(\varphi^{BN}, \rho))$. Now, all points of $BR^{-K_\rho}(\bar{B}(\varphi^{BN}, \rho))$ also lead to a convergence to φ^{BN}, which implies that

$$BR^{-K_\rho}(B(\varphi^{BN}, \rho)) = BR^{-K_\rho}(\bar{B}(\varphi^{BN}, \rho)). \tag{40}$$

Since the left-hand term is an open set, while the other is a closed set, this means that the set of all points converging to φ^{BN} is both open and closed in U. Now, from Hypothesis 3, we know that U is the union of all $BR^{-K_\rho}(B(\varphi^{BN},\rho))$. This proves that each connected component of U belongs to one of the $BR^{-K_\rho}(B(\varphi^{BN},\rho))$, for $\varphi^{BN} \in BN$, and proves the corollary. □

This means that every connected component of U is matched with a Bayesian-Nash Equilibrium.

5.3 Discussion of Hypotheses

In this section, we briefly discuss the hypothesis needed to guarantee the theoretical convergence of the return functions towards a Bayesian-Nash equilibrium. Let us first give an intuition on why a non-unique best reply is rare. Suppose that φ had two best-replies, $\varphi^1 \neq \varphi^2$; then these would be images by ϕ of two best-replies $\sigma^1 \neq \sigma^2$. Consequently, this means that there is a set of types of strictly positive probability for which $u_j(\theta_j, \varphi(\sigma_j^1(\theta_j))) = u_j(\theta_j, \varphi(\sigma_j^2(\theta_j)))$. In other words, a return function profile φ that has two best-reply strategies must be a zero of the mappings

$$\varphi \mapsto u_j(\theta_j, \varphi(\sigma_j^1(\theta_j))) - u_j(\theta_j, \varphi(\sigma_j^2(\theta_j))), \tag{41}$$

for a set of θ_j of positive probability. Now, this mapping is continuous, and from our construction of the topology of φ, we expect that a slight change of φ will yield a non-zero value for this function. If this were the case, then the set of φ that has both σ^1 and σ^2 as best-replies would be some strict closed submanifold of the space of return functions. We acknowledge the fact that it is hard to rigorously state these intuitions, partly because, in general, these spaces are of infinite dimensions. Still, we may expect the set of return functions exhibiting multiple best-replies to be a union of such strict closed submanifolds. Consequently, we can expect its complement to be a dense open set of Φ.

Now, observe that if U is a compact set, then a sufficient condition for Hypothesis 2 to be satisfied is for BR to be continuous. Hypothesis 3 holds if the union of $BR^{-k}(B(\varphi^\infty,\epsilon))$ covers U for some integer k. Of course the actual satisfaction of this hypothesis depends on the choice of the set U. Here, there is a trade-off: the larger the set U, the easier it is to achieve the convergence (see Theorem 7 below), but the harder it is to satisfy Hypothesis 3. In any event, choosing a very large k should help to ensure the coverage of U.

Considering U to be the set of all return functions whose best-replies are uniquely defined is not sufficient to guarantee a continuous function BR to be uniformly continuous. However, BR may still be uniformly continuous, as we expect BR to be naturally extended by continuity on each closure of each connected component of U. Each closure could then be compact, which would prove BR to be uniformly continuous. Now if this closure leads to discontinuities on BR, it would be still possible to cut the edges of the connected components to define a compact set U, which would be a strict subset of the set of return functions whose best-replies are uniquely defined. This is illustrated in Figure 10 where curved lines represent the submanifolds mentioned above, and where the shaded area corresponds to the chosen U.

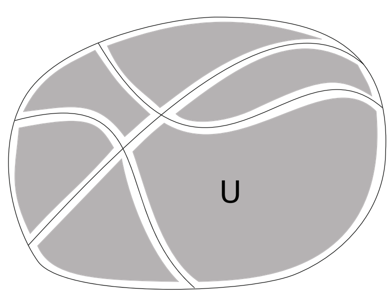

Figure 10: Choice of the set U, represented as the shaded area

Finally, Hypothesis 4 is only slightly stronger than the asymptotic stability, which is almost necessary for any convergence. In particular, in the context of the next section, where the approximation errors are interpreted as the addition of small perturbations to the trajectory of sequences $BR^i(\varphi)$, this strong stability will be essential to guarantee any convergence.

5.4 Approached Best-Reply Dynamics

In this section, we deal with the convergence of the algorithm, where the approximation $\varphi^{i+1} \approx BR(\varphi^i)$ is used. To have convergence, we assume that the distance between φ^{i+1} and $BR(\varphi^i)$ is a decreasing function of i that goes to 0. We suppose that this can be achieved by more sub-iterations at iteration i to better approximate φ^{i+1}.

Hypothesis 5 *The sequence (φ^i) satisfies $d_\Phi(\varphi^{i+1}, BR(\varphi^i)) \to 0$.*

Theorem 7 *Assume Hypotheses 4-5 are satisfied. Consider the values $\nu, \rho > 0$ defined by Hypothesis 4. If for a sufficiently large i, $\varphi^i \in BN_\rho$, then the sequence (φ^i) converges towards a Bayesian-Nash Equilibrium.*

Proof. Under the assumptions, there exists $K \in \mathbb{N}$ such that $\varphi^K \in BN_\rho$ and $d_\Phi(\varphi^{i+1}, BR(\varphi^i)) \leq \nu\rho/2$ for all $i \geq K+1$. Let φ^{BN} be the equilibrium corresponding to φ^K. We have the following inequalities:

$$\begin{aligned}\forall i \geq K, d_\Phi(\varphi^{i+1}, \varphi^{BN}) &\leq d_\Phi(\varphi^{i+1}, BR(\varphi^i)) + d_\Phi(BR(\varphi^i), \varphi^{BN}) &(42)\\ &\leq d_\Phi(\varphi^{i+1}, BR(\varphi^i)) + (1-\nu) d_\Phi(\varphi^i, \varphi^{BN}) &(43)\\ &\leq \nu\rho/2 + (1-\nu) d_\Phi(\varphi^i, \varphi^{BN}). &(44)\end{aligned}$$

We used an induction argument to justify that φ^i belonged to BN_ρ. Another induction argument implies that $d_\Phi(\varphi^i, \varphi^{BN}) \leq \rho$, which leads to

$$d_\Phi(\varphi^{i+1}, \varphi^{BN}) \leq (1-\nu/2) d_\Phi(\varphi^i, \varphi^{BN}), \tag{45}$$

and completes the convergence proof. □

Finally, we state this section's main theoretical result.

Theorem 8 *Assume Hypotheses 2, 3, 4 and 5 are satisfied. If for a sufficiently large i, $\varphi^i \in U$, then the sequence (φ^i) converges to a Bayesian-Nash Equilibrium.*

Proof. See Appendix. □

Intuitively, we have contracting neighborhoods of Nash equilibria, whose inverse images by the best-reply dynamics take over nearly all the space of the return functions. The trouble though comes from errors. While the image of a neighborhood contracts it, errors dilate it. So, we need to make sure we can guarantee falling strictly enough in the interior of the neighborhood through contractions. This is where the uniform continuity of BR is required.

6 Conclusion

In this paper, we introduced the return function as a new approach to study Bayesian games and to compute Bayesian-Nash equilibria. The main novelty consists in studying best-reply dynamics in the space of return functions instead of doing it in the space of strategies. We theoretically showed the near equivalence of analyses in the different spaces. Then, we provided an algorithm that exploits this equivalence by using approximations of best-reply return functions, with samplings and interpolations. This algorithm turned out to be extremely efficient at computing Bayesian-Nash equilibria in settings that are hardly analytically accessible. In the simple 2-player case of cake-cutting, where some analytical remarks could be made, we showed how accurate numerical computations are, as they revealed patterns in the Bayesian-Nash equilibria that could hardly have been anticipated otherwise. In the much more complex setting of 20 players, we achieved successful computations with our methods. Finally, we put these experimental results on solid mathematical foundations in our last section, as we proved convergence despite inherent cumulative computational errors due to samplings and interpolations. We believe this work to be a breakthrough in computational approaches to Bayesian games, opening up new research orientations in further understanding Bayesian-Nash equilibria, with improvements in computational performance and with applications to diverse fields like Bayesian mechanism design.

Appendix: Proofs of Statements

A.1 Proof of Theorem 3

First, let us prove that the Cesaro series of the sequence U^i converges.

Lemma A.1 *Given the hypothesis of Theorem 3, almost surely, we have*

$$\frac{1}{i}\sum_{k=1}^{i} U^i \to \lim_{k\to\infty} \mathbb{E}[U^k]. \tag{A.1}$$

Proof. Since $U^i \in [0, U]$, we have $Var(U^i) \leq U^2$. Thus, the variances are finite and

$$\sum_{i=1}^{\infty} \frac{Var(U^i)}{i^2} \leq U^2 \sum_{i=1}^{\infty} \frac{1}{i^2} < \infty. \tag{A.2}$$

Thus, the lemma is guaranteed by Kolmogorov's strong law. □

Next, we need the following lemma.

Lemma A.2 *Assume that $\frac{1}{n}\sum\limits_{k=1}^{n} a_k \to 0$ and b_n is an monotonic sequence of positive real numbers. Also assume that $nb_n = O(c_n)$, where $c_n = \sum\limits_{k=1}^{n} b_k$, and that $\lim c_n = \infty$. Then, $\frac{1}{c_n}\sum\limits_{k=1}^{n} a_k b_k \to 0$.*

Proof. Let $A_n = \sum\limits_{k=1}^{n} a_k$. Without loss of generality, we can assume that $b_0 = 0$. Also, since $nb_n = O(c_n)$, we can define $K > 0$ such that $nb_n \leq Kc_n$. Let $\epsilon > 0$, and N such that for all $n \geq N$, we have $|A_n| \leq n\epsilon$. Using Abel's transform, we have

$$\frac{1}{c_n}\sum_{k=1}^{n} a_k b_k = \frac{1}{s_n}\sum_{k=1}^{n}(A_k - A_{k-1})b_k = \frac{A_n b_n}{c_n} + \frac{1}{c_n}\sum_{k=1}^{n-1} A_k(b_{k+1} - b_k). \tag{A.3}$$

Thus, we can deduce that, for $n \geq N$,

$$\Big|\frac{1}{c_n}\sum_{k=1}^{n} a_k b_k\Big| \leq \frac{nb_n}{c_n}\epsilon + \frac{1}{c_n}\sum_{k=1}^{N}|A_k||b_{k+1} - b_k| + \frac{\epsilon}{c_n}\Big|\sum_{k=N}^{n-1} k(b_{k+1} - b_k)\Big|. \tag{A.4}$$

By noting $R_N = \sum\limits_{k=1}^{N} |A_k||b_{k+1} - b_k|$ which does not depend on n, and by using another Abel's transform, we have

$$\Big|\frac{1}{c_n}\sum_{k=1}^{n} a_k b_k\Big| \leq \frac{nb_n}{c_n}\epsilon + \frac{R_N}{c_n} + \frac{\epsilon}{c_n}\Big|(n-1)b_n - \sum_{k=N}^{n-1} b_k\Big| \leq \frac{2nb_n + R_N + c_n}{c_n}\epsilon \leq (2K+2)\epsilon, \tag{A.5}$$

for n high enough such that $R_N/c_n \leq 1$. This proves that $\frac{1}{c_n}\sum\limits_{k=1}^{n} a_k b_k \to 0$. □

Finally, we get to the proof of Theorem 3.

Proof. We apply the previous lemma for $a_i = U^i - \lim_{k\to\infty} \mathbb{E}[U^k]$ and $b_i = R^i = i$. The sequence $b_i = i$ is clearly monotonic and positive. Plus, $c_i = i(i-1)/2 \sim i^2/2$. Thus, $ib_i = i^2 = O(i^2) = O(c_i)$, and, clearly, $\lim c_i = \infty$. Given that Lemma A.1 guarantees that $(\sum a_i)/i \to 0$, then, almost surely, we have

$$\frac{1}{c_i}\sum_{k=1}^{i} b_i a_i = \frac{\sum_{k=1}^{i} R^i U^i}{\sum_{k=1}^{i} R^i} - \lim_{k\to\infty} \mathbb{E}[U^k] \to 0. \tag{A.6}$$

In other words, the weighted average sequence V^i converges almost surely towards $\lim_{k\to\infty} \mathbb{E}[U^k]$. □

A.2 Proof of Theorem 8

In order to control deviations of sequences, let us introduce the narrowing concept.

Definition A.1 *The α-narrowing $N_\alpha(U)$ of a subset U of a metric set X is defined as follows*

$$N_\alpha(U) = \{x \in X, B(x,\alpha) \subseteq U\}. \tag{A.7}$$

Since all open sets contain balls, for all open sets U, there always is $\mu > 0$ such that the μ-narrowing of U is a non-empty. Any non-empty narrowing therefore contains balls, which means that it is a neighbourhood of some of its points.

Definition A.2 *A set V is a strict narrowing of U if it is included in an α-narrowing of U, and we denote this by $V \sqsubset U$, i.e.,*

$$V \sqsubset U \Leftrightarrow \exists \alpha > 0, V \subseteq N_\alpha(U). \tag{A.8}$$

Proposition A.1 *Assume $V \subseteq U$. Then, we have the following equality:*

$$\inf_{v\in V, x\notin U} d(v,x) = \sup\{\alpha \geq 0, V \subseteq N_\alpha(U)\}. \tag{A.9}$$

Proof. Let $\alpha > 0$ such that $V \subseteq N_\alpha(U)$. Then, for any $x_V \in V$, $B(x_V,\alpha) \subseteq U$. If $x \notin U$, then $x \notin B(x_V,\alpha)$, which proves that $d(x_V,x) > \alpha$. This shows that the first term is greater than the second.

Now, let $\beta = \inf_{x\in V, x\notin U} d(x_V,x) > 0$. Let $0 < \alpha < \beta$. If $x_V \in V$, any point x such that $d(x_V,x) \leq \alpha$ cannot be outside of U. Thus, $x \in U$, which shows that $B(x_V,\alpha) \subseteq U$. As a result $V \subseteq N_\alpha(U)$. By having α tending towards β, this shows that the second term is at least β. This proves the other inequality, and concludes the proof. □

Lemma A.3 *If $\alpha > \beta > 0$, then $N_\alpha(U)$ is included in $N_\beta(U)$.*

Proof. Let $x \in N_\alpha(U)$, then $B(x,\alpha) \subseteq U$. But since $\alpha > \beta > 0$, $B(x,\beta) \subseteq B(x,\alpha) \subseteq U$. This proves that $x \in N_\beta(U)$, and concludes the proof. □

Corollary A.1 *$N_\alpha(U)$ is a strict narrowing of U.*

Proof. It suffices to take $\beta = \alpha/2$ and to apply the previous lemma. □

Lemma A.4 *Let $\alpha,\beta > 0$, and $x \in X$. Then $B(x,\beta) \subseteq N_\alpha(B(x,\alpha+\beta))$.*

Proof. Let $y \in B(x,\beta)$, and $z \in B(y,\alpha)$. We have

$$d(x,z) \leq d(x,y) + d(y,z) \leq \beta + \alpha, \tag{A.10}$$

thus $z \in B(x,\alpha+\beta)$, which proves that $B(y,\alpha) \subseteq B(x,\alpha+\beta)$. Incidentally, $B(x,\beta) \subseteq N_\alpha(B(x,\alpha+\beta))$. □

Lemma A.5 *For any $\alpha, \beta > 0$, we have $N_{\alpha+\beta}(U) \subseteq N_\alpha(N_\beta(U))$.*

Proof. Let $x \in N_{\alpha+\beta}(U)$. Then $B(x, \alpha+\beta) \subseteq U$. But since

$$B(x, \alpha+\beta) \supseteq \bigcup_{y \in B(x,\alpha)} B(y, \beta), \tag{A.11}$$

we know that for all $y \in B(x, \alpha)$, $y \in N_\beta(U)$. This shows that $x \in N_\alpha(N_\beta(U))$, and concludes the proof. □

Corollary A.2 *If V is a strict narrowing of U, then there exists an α-narrowing of U for which V is still a strict narrowing, i.e.,*

$$V \sqsubset U \Rightarrow \exists \alpha > 0, V \sqsubset N_\alpha(U). \tag{A.12}$$

Proof. Let $\beta > 0$ such that $V \subseteq N_\beta(U)$, and let $\alpha = \beta/2$. Then,

$$V \subseteq N_\beta(U) = N_{\alpha+\alpha}(U) \subseteq N_\alpha(N_\alpha(U)), \tag{A.13}$$

which proves that V is a strict narrowing of $N_\alpha(U)$. □

Theorem A.1 *Let f be a uniformly continuous function from X into itself. If V is a strict narrowing of U, then there exists an α-narrowing of U such that the reciprocal image of V is a strict narrowing of the reciprocal image of $N_\alpha(U)$, i.e.,*

$$f \textit{ uniformly continuous and } V \sqsubset U \Rightarrow \exists \alpha > 0, f^{-1}(V) \sqsubset f^{-1}(N_\alpha(U)). \tag{A.14}$$

Proof. Let α such that $V \sqsubset N_\alpha(U)$, which exists because of Corollary A.2. Let $\epsilon = d(V, N_\alpha(U)^c)$. We know that $\epsilon > 0$ because of Proposition A.1. Now, since f is uniformly continuous, there exists $\delta > 0$ such that whenever $d(f(x), f(y)) \geq \epsilon$, we have $d(x, y) \geq \delta$. If we take $x \in f^{-1}(V)$ and $y \in f^{-1}(N_\alpha(U))$, this implies that $d(f^{-1}(V), f^{-1}(N_\alpha(U))^c) \geq \delta > 0$. Proposition A.1 enables us to conclude that $f^{-1}(V) \sqsubset f^{-1}(N_\alpha(U))$. □

If U is a subset, then we denote $f_\mu^{-1}(U) = f^{-1}(N_\mu(U))$. If $\mu \in (\mathbb{R}_+^*)^n$ is a vector of positive numbers, then we define $f_\mu^{-n}(U) = f_{\mu_n}^{-1}(f_{\mu_{n-1}}^{-1}(f_{\mu_{n-2}}^{-1}(\ldots f_{\mu_1}^{-1}(U))))$.

Corollary A.3 *Assume that U has a non-empty narrowing and that f is uniformly continuous. Then,*

$$\forall \epsilon > 0, \forall n \in \mathbb{N}, \exists \mu \in (\mathbb{R}_+^*)^n, \forall i \in \{1, \ldots, n\}, f^{-i}(N_\epsilon(U)) \subseteq f_{\mu_1,\ldots,\mu_i}^{-i}(U). \tag{A.15}$$

Proof. It suffices to apply Theorem A.1 to construct the sequence of $(\mu_i)_{1 \leq i \leq n}$. □

Finally, we can prove Theorem 8.

Proof. Consider $\nu, \rho > 0$ which corresponds to Hypothesis 4. Let $\epsilon = \rho/2$. According to Lemma A.4, we have $BN_\epsilon \subseteq N_\epsilon(BN_\rho)$. Now, because of Hypothesis 2 of uniform continuity, we can apply Corollary A.3 for $n = K_\epsilon$, which provides the existence of $(\mu_i)_{1 \leq i \leq K_\epsilon} \in (\mathbb{R}_+^*)^{K_\epsilon}$ such that

$$\forall i \in \{1, \ldots, K_\epsilon\}, BR^{-i}(N_\epsilon(BN_\rho)) \subseteq BR_{\mu_1,\ldots,\mu_i}^{-i}(BN_\rho). \tag{A.16}$$

Thus,

$$BR_\mu^{-K_\epsilon}(BN_\rho) \supseteq BR^{-K_\epsilon}(N_\epsilon(BN_\rho)) \supseteq BR^{-K_\epsilon}(BN_\epsilon) = U, \tag{A.17}$$

according to Hypothesis 3. Now, let $\alpha = \inf\{\mu_i, 1 \leq i \leq K_\epsilon\} \cup \{\nu\rho/2\}$. We know that $\alpha > 0$, because all μ_i are positive, and because there is a finite number of them. Hypothesis 5 now implies that there is a rank K_0 such that, for all $i \geq K_0$, we have $d_\Phi(\varphi^{i+1}, BR(\varphi^i)) \leq \alpha$. Now, if (φ^i) comes an infinite number of times inside U, then there is $K_1 \geq K_0$ such that $\varphi^{K_0} \in U$.

Now, let us show by induction over i that $\varphi^{K_1+i} \in BR^{-K_\epsilon+i}_{\mu_1,\ldots,\mu_{K_\epsilon-i}}(BN_\rho)$ for $0 \leq i \leq K_\epsilon$. We know that $\varphi^{K_1+0} = \varphi^{K_1} \in U = BR^{-K_\epsilon+0}_{\mu}(BN_\rho)$. Let us now assume that the induction is true for i. Then, using the definition of BR^{-n}_{μ},

$$\varphi^{K_1+i} \in BR^{-K_\epsilon+i}_{\mu_1,\ldots,\mu_{K_\epsilon-i}} = BR^{-1}(N_{\mu_{K_\epsilon-i}}(BR^{-K_\epsilon+i+1}_{\mu_1,\ldots,\mu_{K_\epsilon-i-1}}(BN_\rho)))). \tag{A.18}$$

This means that $BR(\varphi^{K_1+i}) \in N_{\mu_{K_\epsilon-i}}(BR^{-K_\epsilon+i+1}_{\mu_1,\ldots,\mu_{K_\epsilon-i-1}}(BN_\rho)))$, and implies that all points within a distance less than or equal to $\mu_{K_\epsilon-i}$ are included in $BR^{-K_\epsilon+i+1}_{\mu_1,\ldots,\mu_{K_\epsilon-i-1}}(BN_\rho))$. Since $d_\Phi(BR(\varphi^{K_1+i}), \varphi^{K_1+i+1}) \leq \alpha \leq \mu_{K_\epsilon-i}$, we deduce that

$$\varphi^{K_1+i+1} \in BR^{-K_\epsilon+i+1}_{\mu_1,\ldots,\mu_{K_\epsilon-i-1}}(BN_\rho)), \tag{A.19}$$

which proves the induction proof. Eventually, $\varphi^{K_1+K_\epsilon} \in BN_\rho$.

Since, K_1 was chosen such that $d_\Phi(\varphi^{i+1}, BR(\varphi^i)) \leq \alpha \leq \nu\rho/2$ for all $i \geq K_1$, we can then apply Theorem 7, which proves that the sequence φ^i converges towards a Bayesian-Nash Equilibrium. □

References

Abdulkadirolu, A., T. Sönmez. 2003. School choice: A mechanism design approach. *American Economic Review* **93** 729–747.

Achdou, Y., I. Capuzzo-Dolcetta. 2010. Mean field games: Numerical methods. *SIAM Journal on Numerical Analysis* **48**(3) 1136–1162.

Bourbeau, B., T.G. Crainic, M. Gendreau, J. Robert. 2005. Design for optimized multi-lateral multi-commodity markets. *European Journal of Operational Research* **163** 503–529.

Brams, J.B., A.D. Taylor. 1995. An envy-free cake division protocol. *The American Mathematical Monthly* **102** 9–12.

Brams, J.B., A.D. Taylor. 1996. *Fair Division - From Cake-cutting to Dispute Resolution*. Cambridge University Press.

Chen, J.-M., H.-L. Cheng. 2012. Effect of the price-dependent revenue-sharing mechanism in a decentralized supply chain. *Central European Journal of Operations Research* **20** 299–317.

Chen, Y., J.K. Lai, D.C. Parkes, A.D. Procaccia. 2010. Truth, justice, and cake cutting. *Proceedings of the 24th AAAI Conference on Artificial Intelligence (to appear)* .

Dasgupta, P., P. Hammond, E. Maskin. 1979. The implementation of social shoice rules: Some results on incentive compatibility. *Review of Economic Studies* **46** 185216.

Devenur, N.R., T.P. Hayes. 2009. The adwords problem: Online keyword matching with budgeted bidders under random permutations. *Proceedings of the 19th ACM Conference on Electronic Commerce*. 71–78.

Gibbard, A. 1973. Manipulation of voting schemes: A general result. *Econometrica* **41** 587–601.

Gibbs, A.L., F.E. Su. 2002. On choosing and bounding probability metrics. *International Statistical Review* **70**(3) 419–435.

Harsanyi, J.C. 1968a. Games with incomplete information played by "Bayesia" players, part III. The basic probability distribution of the game. *Management Science* **14**(7) 486–502.

Harsanyi, J.C. 1968b. Games with incomplete information played by "Bayesian" players part I. The basic model. *Management Science* **14**(3) 159–183.

Harsanyi, J.C. 1968c. Games with incomplete information played by "Bayesian" players part II. Bayesian equilibrium points. *Management Science* **14**(5) 320–334.

Holmstrom, B. 1977. On incentives and control in organizations. Ph.D. thesis, Stanford University.

Huang, M., R.P. Malhamé, P.E. Caines. 2006. Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems* **6**(3) 221–252.

Jain, S., S. Raghavan. 2009. A queuing approach for inventory planning with batch ordering in multi-echelon supply chains. *Central European Journal of Operations Research* **17** 95–110.

Lasry, J.-M., P.-L. Lions. 2007. Mean field games. *Japanese Journal of Mathematics* **2**(1) 229–260.

Manelli, A.M., D.R. Vincent. 2007. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Central European Journal of Operations Research* **137** 153–185.

Mes, M., M. van der Heijdenet, P. Shurr. 2011. Interaction between intelligent agent strategies for real-time transportation planning. *Central European Journal of Operations Research* .

Mossel, E., O. Tamuz. 2010. Truthful fair division. *Proceedings of the 3rd International Symposium on Algorithmic Game Theory*. 288–299.

Myerson, R. 1979. Incentive-compatibility and the bargaining problem. *Econometrica* **47** 6173.

Pathak, P.A. 2011. The mechanism design approach to student assignment. *Annual Review of Economics* **3** 513–536.

Rabinovich, Z., V. Naroditskiy, E.H. Gerding, N.R. Jennings. 2013. Computing pure Bayesian-Nash equilibria in games with finite actions and continuous types. *Artificial Intelligence* **195**(0) 106 – 139. doi:http://dx.doi.org/10.1016/j.artint.2012.09.007. URL http://www.sciencedirect.com/science/article/pii/S0004370212001191.

Robertson, J., W. Webb. 1998. *Cake-Cutting Algorithms: Be Fair If You Can*. AK Peters Ltd.

Su, X., S.A. Zenios. 2006. Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Science* **52** 1647–1660.

Vulcano, G., G. van Ryzin, C. Maglaras. 2002. Optimal dynamic auctions for revenue management. *Management Science* **48** 1388–1407.