

A new refinement of Nash equilibrium concept in discontinuous games.*

Philippe Bich[†]

May 16, 2017

Abstract

We introduce the new concept of prudent equilibrium to model strategic uncertainty, and prove it exists in large classes of discontinuous games. When the game is better-reply secure, we show that prudent equilibrium refines Nash equilibrium. In contrast with the current literature, we don't use probabilities to model players' strategies and beliefs about other players' strategies. We provide examples (first-price auctions, location game, Nash demand game, etc.) where prudent equilibrium concept removes most non-intuitive solutions of the game.

JEL classification: C02, C62, C72, L13.

Keywords: prudent equilibrium, Nash equilibrium, refinement, strategic uncertainty, better-reply secure, discontinuous games.

1 Introduction

Consider a first-price sealed-bid auction with complete information between two bidders. The players are characterized by their valuation v_1 and v_2 of the item for sale, $v_1 < v_2$, and they are supposed to choose bids x_1 and x_2 in $[0, M]$, $M > 0$. Assume that in case of ties, i.e. if $x_1 = x_2$, then the winner is the player with the highest value. An easy computation proves that for every $x \in [v_1, v_2]$, the strategy profile (x, x) is a Nash equilibrium of this strategic game. Yet, for $x > v_1$, these equilibria represent fragile situations, because of *strategic uncertainty*: if player 2 does not respect his equilibrium strategy and decreases slightly his bid, then player 1 gets the item for a price x higher than his valuation v_1 . As a matter of fact, any strategy $x_1 \leq v_1$ is also a best-reply of player 1 if player 2 plays $x > v_1$, but it is also immune to a small modification of player 2's strategy. Thus, if player 2 is supposed to play $x > v_1$, and if he predicts that his opponent should play $x_1 \leq v_1$, then he could be tempted to lower his bid x in order to increase his payoff. Finally, playing $x > v_1$ for player 1 seems definitely a bad choice, even if the other player is assumed to play the same strategy.

*I thank Lisa Morhaim for useful comments. I thank some participants of the following seminars or conferences for their valuable comments: French Symposium on Games, May 2015 (Paris), SWET, June 2016 (Paris), SAET, July 2016 (Rio).

[†]Paris School of Economics, Centre d'Economie de la Sorbonne UMR 8174, Université Paris I Panthéon/ Sorbonne, 106/112 Boulevard de l'Hôpital 75647 Paris Cedex 13. E-mail: bich@univ-paris1.fr

This example illustrates that Nash equilibrium concept has sometimes to be refined in order to keep some predictive power, and the same idea can be found in many other situations (e.g., Nash demand game, location game, Bertrand duopoly, etc. See Section 5). Note that the existence of a pure Nash equilibrium in the previous game can be obtained from Reny’s theorem [26], which guarantees existence for the large class of better-reply secure games. This class encompasses many discontinuous¹ economic games, in particular the first-price auction described above. Astonishingly, there is no refinement notion of pure strategy Nash equilibrium for discontinuous games which covers this example, or more generally the class of better-reply secure games. A possible reason is that most refinement notions - like perfect equilibrium of Selten [29] - require the existence of equilibria of some auxiliary games, where the players’ strategies are perturbed by random mistakes, and payoffs are expected payoffs. If the initial game is better-reply secure, such auxiliary games are in general neither better-reply secure nor quasiconcave, thus no general existence result in pure strategies can be applied to them.²

In this paper, we introduce a new refinement of pure strategy Nash equilibrium in discontinuous games, called *prudent equilibrium*. We prove its existence (Theorem 17) for the class of *p-robust games*, which contains many discontinuous economic games. Roughly, a game is p-robust if for every strategy profile x , no player can increase largely his payoff at x by arbitrary small modifications of his strategy, even if the other players can change slightly their strategy. The first-price auction above is p-robust: for example, at every strategy profile (x, x) with $x < v_1$, if player 1 increases his strategy a little bit, player 2 can answer by the same modification, so that player 1 does not increase his payoff.

We now provide an informal definition of our main solution concept, prudent equilibrium. The main issue in the introductory example is strategic uncertainty, i.e. the uncertainty related to other players’ strategies and rationality (see Brandenburger [7]). A radical way to remove strategic uncertainty in games would be to consider extremely prudent players, who try to maximize $\tilde{u}_i(x) = \inf_{x_{-i} \in X_{-i}} u_i(x_i, x_{-i})$ with respect to their strategy x_i , u_i being the initial payoff function of player i , and X_{-i} the strategy sets of his opponents. A less extreme answer would be to assume that given player i ’s belief about the potential strategy profile x_{-i} of his opponents, he has good reasons to think that their true strategies will stay in some set $Y_{-i}(x_{-i}) \subset X_{-i}$ (for example a small ball centered at x_{-i}). Then, a prudent player would choose his strategy x_i in order to maximize $\tilde{u}_i(x) = \inf_{y_{-i} \in Y_{-i}(x_{-i})} u_i(x_i, y_{-i})$. This function can also be written as

$$\tilde{u}_i(x) = \inf_{y_{-i} \in X_{-i}} (u_i(x_i, y_{-i}) + \delta(y_{-i}, x_{-i}))$$

where

$$\delta(x_{-i}, y_{-i}) = \begin{cases} 0 & \text{if } y_{-i} \in Y_{-i}(x_{-i}), \\ +\infty & \text{otherwise.} \end{cases}$$

In the infimum above, note that only the strategies y_{-i} for which $\delta(y_{-i}, x_{-i})$ is equal to 0 are useful, and the other ones, for which $\delta(y_{-i}, x_{-i})$ is infinite, can be removed. In short, δ is a measure for player i of the importance of a potential deviation y_{-i} from the the expected strategy profile x_{-i} .

A natural generalization leads to the definition of the following auxiliary ”prudent” game

¹A game is discontinuous if some of its payoff functions are discontinuous.

²Carbonnel ([9], [8]) extends the existence of perfect equilibria in mixed strategies for some particular classes of discontinuous games. Andersson et al. [3] have recently introduced a pure strategy refinement of Nash equilibrium which has cutting power in some discontinuous games, but whose existence is guaranteed only in continuous games.

$$u_i^\lambda(x) = \inf_{y_{-i} \in X_{-i}} \left(u_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} \right) \quad (1)$$

where $c_i : X_{-i} \times X_{-i} \rightarrow \mathbb{R}$, and $\lambda > 0$ is some normalization coefficient. The function $\frac{c_i(x_{-i}, y_{-i})}{\lambda}$ could be interpreted as a (free) insurance payed to player i if the other players $-i$ play y_{-i} instead of the expected strategy x_{-i} , and $1/\lambda$ would parametrize the insurance level. By analogy with the previous case, the function $\frac{c_i(y_{-i}, x_{-i})}{\lambda}$ is also a way to parametrize strategic uncertainty: $x_{-i} \in X_{-i}$ is the strategy expected by player i , and $\frac{c_i(y_{-i}, x_{-i})}{\lambda} \leq \frac{c_i(y'_{-i}, x_{-i})}{\lambda}$ means that the possible deviation y_{-i} has more importance for player i than y'_{-i} has.

Interestingly, the prudent behaviour described above has a smoothing effect on the initial game: in general, the prudent payoff functions u_i^λ are more regular than the initial payoff functions u_i . This smoothing effect implies that for every p-robust game G , and for a large class of functions c_i (for example distances), there exists a Nash equilibrium of the prudent game associated to G (Theorem 13). This opens a route for refinement, and indeed, we prove that if the initial game G is better-reply secure, and if the level of insurance $1/\lambda$ tends to $+\infty$, then any limit point of Nash equilibria of the prudent games is a Nash equilibrium of G (Proposition 16). We call such a limit point a prudent equilibrium: for example, in the first-price auction above, the only prudent equilibrium is the intuitive solution (v_1, v_1) (see Proposition 31).

Our definition of a prudent game should be compared to variational preferences, introduced by Maccheroni, Marinacci and Rustichini [20] to model uncertainty aversion in decision theory. Recall that variational preferences on the set of acts \mathcal{F} are represented by

$$V(f) = \min_{p \in \Delta} \left(\int u(f) dp + c(p) \right),$$

where u is a utility function, $f \in \mathcal{F}$ an act, Δ the set of priors over a state space S , and $c : \Delta \rightarrow [0, +\infty]$ an index of uncertainty aversion. The interpretation by Maccheroni et al. is the following. When the decision maker contemplates choosing an act f , the malevolent Nature tries to minimize its expected utility. Any prior p can be chosen, but Nature must pay a cost $c(p)$ to do so. In their setting, this cost is also interpreted as an ambiguity index.

Our model adapts³ some of these ideas to a strategic setting. In particular, ambiguity is turned into strategic uncertainty. But a major difference is that variational preferences are valid in a probabilistic setting, although we consider only a deterministic framework. Also, the cost $c_i(x_{-i}, y_{-i})$ in our model depends on the potential strategy profile y_{-i} of $-i$, which plays the role of p in variational preferences, but also on the strategy profile x_{-i} expected by i .

Many other papers have tried to model strategic uncertainty in games. In quantal-response equilibrium models, pioneered by Kelsey and Palfrey [22], strategic uncertainty is represented by a probability distribution (some noise) added to the initial payoff of each player, which defines a perturbed game. For every mixed strategy profile σ , every player i acts optimally in the perturbed game against σ_{-i} . This induces another probability distribution over the observed actions of the players. If this probability distribution is σ , it is, by definition, a quantal-response equilibrium. In a similar vein, Andersson et

³De Marco and Romaniello [13] have adapted Maccheroni, Marinacci and Rustichini [20] to a Bayesian strategic setting. Their model and their scope is not connected with ours.

al. [3] consider that players choose pure strategies, and strategic uncertainty is now represented through probabilistic subjective beliefs about the strategies of each player’s opponents. As above, this defines an equilibrium notion in some auxiliary noisy game. Remark that both approaches are related to refinement literature (see Selten [29] or Myerson [24]), and as a matter of fact, when the level of noise tends to zero, they provide refinement concepts.

Non-additive models are an alternative to model strategic uncertainty. If strategic uncertainty is represented by a set of priors, then the preferences of each players can be defined through Choquet expected utility model (see Mukerji [23], Marinacci [21], Ryan [28], or Eichberger and Kelsey [15] who also model optimism or pessimism in strategic games), or through Gilboa-Schmeidler maximin model (Klibanoff [17], Dow-Verlang [14], Lo [19] or De Marco and Romaniello [13]). Most of the papers above differ in their definition of the support for the beliefs. Recently, Renou and Schlag [25] have proposed a dual model based on minimax behaviour: in their approach, regret guides players in forming probabilistic assessments and, ultimately, in making choices.

The main difference between our model and these models is that beliefs about the strategies of the other players are not represented by sets of priors, but by deterministic functions. It turns out to be a very tractable approach in many cases, even when the initial game is discontinuous (see Section 5), and it has several interpretations. For readers who are more interested in mixed strategies, note that our model can also be applied to the mixed extension of a game. We only give one example to illustrate this point (Example 38), and more generally, prudent equilibria in mixed strategies will be studied in a separate paper. But we think that our notion is particularly well-designed for discontinuous games, less for mixed extension of games, for which usual refinement concepts could be more adapted.

Last⁴, it should be added that prudent equilibrium shares a common feature with most previous refinement notions: it does not pretend to select the reasonable outcome in all familiar games. Indeed, in general, it is always possible to find a strategic game for which a given refinement concept is ineffective. But we think that (1) our notion is complementary to the previous ones, (2) it helps to remove non intuitive solutions for many discontinuous games, as illustrated in the paper, and (3) it is not difficult to compute for many simple games. A specific feature of prudent equilibrium in pure strategies is that the strategic uncertainty is local (contrarily to trembling-hand equilibrium concept, which considers that every strategy could be played by mistake (i.e. out of equilibrium) with some probability). According to us, it could be an interesting feature, since it only requires that agents examine local mistakes. But clearly, it implies that we do not encompass previous refinement notions for every game. This is a price to pay to be able to get an existence result in pure strategies for large and simple classes of discontinuous games. Another important feature of prudent equilibrium is that it is possible (though not automatic) that the choice of the functions c_i that parametrize strategic uncertainty influences the prudent equilibrium.⁵ This is not surprising: indeed, c_i models strategic uncertainty in a local way, and it could influence the selection of the Nash equilibrium of the initial game. In general, the choice of c_i could depend on the context of the game. Anyway, we provide examples for which the prudent equilibrium does not depend on c_i (see Examples 28 or 35).

The paper is organized as follows. Section 2 introduces p-robustness and prudent games. A measure

⁴We thank an anonymous referee for pushing us to the following discussion.

⁵This phenomenon also occurs in Anderson et al. [3], where Nash equilibria which are robust to strategic uncertainty can depend on the random variables which model strategic uncertainty.

of strategic uncertainty is defined, together with some local comparison index of strategic uncertainty between players. In addition, we characterize the class of prudent payoffs. Section 3 introduces prudent equilibrium, and proves it is a refinement of Nash equilibrium in better-reply secure games. Section 4 proposes some extensions, for example when the game possesses enough symmetry. Section 5 provides examples. Mathematical proofs are given in the last section.

2 The main solution concepts

2.1 The general framework

There are N players.⁶ The pure strategy set of each player $i \in N$, denoted by X_i , is a non-empty, compact subset of a metric topological vector space E_i . Each player i has a bounded payoff function

$$u_i : X = \prod_{i \in N} X_i \rightarrow \mathbb{R}.$$

A strategic game G is a pair $G = ((X_i)_{i \in N}, (u_i)_{i \in N})$. For every $x \in X$ and every $i \in N$, we denote $x_{-i} = (x_j)_{j \neq i}$ and $X_{-i} = \prod_{j \neq i} X_j$. Throughout this paper, a game G satisfying the above assumptions is called a compact game. Additionally, G is called a quasiconcave game if for every player i , X_i is convex, and if for every player i and every strategy $x_{-i} \in X_{-i}$, $u_i(x_i, x_{-i})$ is quasiconcave in x_i . The game G is called continuous if for every player i , u_i is continuous in x .

We shall denote by $\Gamma = \overline{\{(x, u(x)) : x \in X\}}$ the closure of the graph of G . Let us define the “secure payoff level” of player i when he plays d_i and when the other players play x_{-i} by $\underline{u}_i(d_i, x_{-i}) = \liminf_{x'_{-i} \rightarrow x_{-i}} u_i(d_i, x'_{-i})$. Following Reny [26], the game G is *better-reply secure* if whenever $(x, v) \in \Gamma$ and x is not a Nash equilibrium, some player $i \in N$ can secure a payoff strictly above v_i , i.e. there exists $d_i \in X_i$ such that $\underline{u}_i(d_i, x_{-i}) > v_i$. It is easy to check that every continuous game is better-reply secure, and we recall that every better-reply quasiconcave game admits a Nash equilibrium (Reny’s theorem [26]).

2.2 P-robustness

The following definition plays a central role in all our results.

Definition 1. *A payoff function u_i is p-robust at $x \in X$ if for every $\varepsilon > 0$ and for every neighborhood $V_{x_{-i}}$ of x_{-i} , there exists some open neighborhood V_{x_i} of x_i such that $\sup_{x'_i \in V_{x_i}} \inf_{x'_{-i} \in V_{x_{-i}}} u(x'_i, x'_{-i}) \leq u_i(x) + \varepsilon$.*

The payoff function u_i is p-robust if this holds for every $x \in X$. If for every $i \in N$, u_i is p-robust, then we say that G is p-robust.

Thus, u_i is p-robust if player i cannot largely improve his payoff by modifying slightly his strategy, if he anticipates the worst local possible modification of other players’ strategies. If u_i was not p-robust at x , a pessimistic player could have some incentive to slightly modify his strategy x_i . In particular, games with continuous payoff functions are p-robust.

⁶According to the context, N will denote the set of players or its cardinal.

A first simple class of p-robust games are games for which the following property of graph upper-semicontinuity (graph u.s.c.) is true:⁷

Definition 2. *The payoff function u_i is graph upper semicontinuous (graph u.s.c.) at $\bar{x} \in X$ if there exists V_{x_i} an open neighborhood of x_i , and a continuous mapping $f_{-i} : V_{x_i} \rightarrow X_{-i}$, such that $u_i(x_i, f_{-i}(x_i))$ is u.s.c. on V_{x_i} .*

Example 3. The first-price auction described in the introduction is p-robust: indeed, for every player $i = 1, 2$ and around any point \bar{x}_i , one can define $f_{-i}(x_i) = x_i$. Then, $u_i(x_i, f_{-i}(x_i)) = u_i(x_i, x_i)$ is continuous, thus u_i is graph u.s.c. Following the same argument, we can prove that every 2-player discontinuous game is p-robust if the payoff functions are continuous outside the diagonal and u.s.c. on the diagonal.

Example 4. Most familiar discontinuous economic games are p-robust: first-price auctions, second-price auctions, Bertrand's price competition, Cournot's model of oligopoly, Hotelling's model of spatial competition, timing game, etc. All these games are contained into the class into Diagonal games which are p-robust (see Example 7 below).

P-robustness assumption is fundamentally different from better-reply security or its generalizations. As the following example shows, there is no direct relationship between better-reply security and p-robustness.

Example 5. Consider a two-player game with $X_1 = X_2 = [0, 1]$, $u_1(x_1, x_2) = -(x_1 - x_2)^2$ for every $(x_1, x_2) \in [0, 1]^2$, $u_2(0, x_2) = x_2$ for every $x_2 \in [0, 1]$ and $u_2(x_1, x_2) = -x_2$ for every $(x_1, x_2) \in]0, 1] \times [0, 1]$. This game is p-robust, but it is not better-reply secure (because it is quasiconcave and does not possess any Nash equilibria). Conversely, define $v_1(0, 0) = 0$, $v_1(x_1, x_2) = 1$ for every $(x_1, x_2) \neq (0, 0)$, and $v_2(x_1, x_2) = x_2$. This game is better-reply secure, but v_1 is not p-robust at $(0, 0)$.

P-robustness is also fundamentally different from Baye et al. [5]: these authors provide necessary and sufficient conditions for Nash equilibrium existence, although our assumption is disconnected from Nash equilibrium existence. For a similar reason, our assumption is different from weak reciprocal upper semicontinuity, introduced by Carmona [10].

The following proposition, proved in Appendix 6.1, provides an easy way to check that a game is p-robust, and generalizes the graph u.s.c. criterium above.

Proposition 6. Consider a game $G = ((X_i)_{i \in N}, (u_i)_{i \in N})$. Assume that for every $x = (x_i, x_{-i}) \in X$, there exists $V_{x_i} \subset X_i$, an open neighborhood of x_i , and there exists a lower semicontinuous multivalued function⁸ ψ^x from V_{x_i} to X_{-i} such that $x_{-i} \in \psi^x(x_i)$, and such that the restriction of u_i so $\text{Gr}\psi^x := \{y = (y_i, y_{-i}) \in V_{x_i} \times X_{-i} : y_{-i} \in \psi^x(y_i)\}$ is upper semicontinuous at x . Then G is p-robust.

In particular, if u_i is upper semicontinuous in x_i , then it is p-robust. Indeed, for every $x = (x_i, x_{-i}) \in X$, we can apply Proposition 6 at x by defining $\psi^x(y_i) = x_{-i}$ for every $y_i \in X_i$.

The next example provides a general class of games which is p-robust.

⁷This definition is similar to graph-continuity assumption of Dasgupta and Maskin [11], except that i and $-i$ are reversed in their definition, and upper semicontinuity is replace by continuity.

⁸Let A and B be two topological spaces. A multivalued function Φ from A to B is lower semicontinuous if for every open subset V of B , the set $\{x \in A : \Phi(x) \cap V \neq \emptyset\}$ is an open subset of A .

Example 7. A game $G = ((X_i)_{i \in N}, (u_i)_{i \in N})$ is diagonal (see Bich and Laraki [6]) if for every $i \in N$, there exists:

1. f_i and g_i , some upper semicontinuous mappings from $[0, 1] \times [0, 1]$ to \mathbb{R}
2. h_i , a mapping from $[0, 1]^N$ to \mathbb{R}
3. ϕ , a continuous mapping from $[0, 1]^{N-1}$ to $[0, 1]$, such that its (multivalued) inverse $\phi^{-1}(x_i) = \{x_{-i} \in X_{-i} : x_i = \phi(x_{-i})\}$ is lower semicontinuous,

such that:

$$u_i(x_i, x_{-i}) = \begin{cases} f_i(x_i, \phi(x_{-i})) & \text{if } \phi(x_{-i}) > x_i, \\ g_i(x_i, \phi(x_{-i})) & \text{if } \phi(x_{-i}) < x_i, \\ h_i(x_i, x_{-i}) & \text{if } \phi(x_{-i}) = x_i. \end{cases}$$

The assumption on ϕ is satisfied, for example, if $\phi(y) = \{k\text{-th highest value of } \{y_1, \dots, y_{N-1}\}\}$, $k = 1, \dots, N - 1$. Such functions encompass many models of competition with complete information (e.g., auctions, wars of attrition, preemption games or Bertrand competitions).

If h_i is upper semicontinuous for every $i \in N$, then the diagonal game G is p-robust from Proposition 6. Indeed, we have to check p-robustness at every discontinuity point $x \in X$ of u_i , i.e. at $x = (\phi(x_{-i}), x_{-i})$. By assumption, $\psi^x(y_i) = \{y_{-i} \in X_{-i} : y_i = \phi(y_{-i})\}$ is a lower semicontinuous multivalued function, and Proposition 6 can be applied, since the restriction of u_i to $\text{Gr}\psi^x$ coincides with h_i .

Another example of interest is the following case⁹: assume that $N = 2$, that ϕ is equal to identity, and that for every $x \in [0, 1]$, $h_i(x, x) \geq \min\{f_i(x, x), g_i(x, x)\}$. Under these assumptions, we get a p-robust game. Indeed, consider for example the case $(x, x) \in]0, 1[\times]0, 1[$ and $h_i(x, x) \geq f_i(x, x)$. Defining $\psi^x(y_i) = (y_i, 1]$, we can apply Proposition 6. The other cases are similar.

2.3 Solution concept

To every game we can associate an auxiliary "prudent" game as follows:

Definition 8. Let \mathcal{F}_i be the set of continuous¹⁰ real-valued functions c from $X_{-i} \times X_{-i}$ to $[0, +\infty]$ such that $c(x_{-i}, y_{-i}) = 0$ if and only if $x_{-i} = y_{-i}$. Let $c = (c_i)_{i \in N} \in \prod_{i \in N} \mathcal{F}_i$. For every $\lambda > 0$, the λ -prudent game associated to $G = ((X_i)_{i \in N}, (u_i)_{i \in N})$ is $G^\lambda = ((X_i)_{i \in N}, (u_i^\lambda)_{i \in N})$, where for each player $i \in N$,

$$u_i^\lambda(x_i, x_{-i}) = \inf_{y_{-i} \in X_{-i}} \left\{ u_i(x_i, y_{-i}) + \frac{c_i(x_{-i}, y_{-i})}{\lambda} \right\}. \quad (2)$$

A λ -prudent-equilibrium (or λ -equilibrium¹¹) of G is a Nash equilibrium of G^λ .

The first important properties of the prudent game G^λ are summarized in the following proposition. The proof can be found in Appendix 6.2.

Proposition 9. For every quasiconcave game G :

⁹It can be applied, for example, to non-zero-sum, noisy games of timing (see Example 3.1. in [26]).

¹⁰The set $[0, +\infty]$ is endowed with the topology induced by the usual topology of the extended real line $[-\infty, +\infty]$.

¹¹The notion of λ -equilibrium is actually parametrized by λ and c , but in many applications, we will be interested by comparative static effects of λ for a fixed c .

1. G^λ is quasiconcave.
2. For every $x_i \in X_i$, $x_{-i} \rightarrow u_i^\lambda(x_i, x_{-i})$ is continuous.
3. If G is p-robust, then u_i^λ is upper semicontinuous with respect to x .
4. $u_i^\lambda \leq \underline{u}_i \leq u_i$.

For example, a natural choice in Definition 8 is $c_i(x_{-i}, y_{-i}) = \delta_{-i}(x_{-i}, y_{-i})^{\alpha_i}$ for some $\alpha_i > 0$, where δ_{-i} denotes any product distance on X_{-i} . The infimum in the definition of u_i^λ means that each player of G^λ is prudent (or pessimistic) with respect to the rationality¹² of opponents. The function $\frac{c_i(x_{-i}, y_{-i})}{\lambda}$ allows to weight differently opponents' actions, and could be interpreted in several ways: as discussed in the introduction, it may be seen as a functional index related to strategic uncertainty of player i about other players' strategies. Another interpretation is that it is a (free) insurance payed to player i if the other players $-i$ play y_{-i} instead of the predicted strategy x_{-i} . This "insurance" can be seen as an abstract way to modelize the degree of confidence¹³ that player i has in his belief that opponents will respect their Nash equilibrium strategies. A last interpretation is related to ambiguity: the set of strategies of opponents can be seen as a set of deterministic priors, and $\frac{c_i(x_i, y_{-i})}{\lambda}$ is a measure of ambiguity on the other players' strategies. In particular, $\frac{c_i(x_i, y_{-i})}{\lambda} = 0$ means a maximal ambiguity (which implies that player i will act as a maximin player), and $\frac{c_i(x_i, y_{-i})}{\lambda} = +\infty$ that there is no ambiguity at all.

In each interpretation above, the functions c_i parametrize some local shape (of degree of confidence, ambiguity index, etc.), and $\lambda > 0$ parametrizes the level of the ambiguity. The case $\lambda \rightarrow 0$ corresponds to perfect insurance against strategic uncertainty, or perfect confidence, or minimal ambiguity. On the opposite, when $\lambda \rightarrow +\infty$, players are getting closer to maximin agents, which corresponds to maximal ambiguity or minimal confidence level.

This could be formalized as follows:¹⁴ first recall that a maximin strategy for player $i \in N$ is a strategy $x_i^* \in X_i$ such that

$$\inf_{x_{-i} \in X_{-i}} u_i(x_i^*, x_{-i}) = \sup_{x_i \in X_i} \inf_{x_{-i} \in X_{-i}} u_i(x_i, x_{-i})$$

A maximin strategy profile of G is a profile of strategies $x = (x_i)_{i \in N}$ such that x_i is a maximin strategy for every player i .

- Proposition 10.**
1. For every p-robust game¹⁵ G , any limit of λ^n -equilibria where λ^n tends to $+\infty$ is a maximin strategy profile.
 2. For every better-reply secure game G , the limit of λ^n -equilibria where λ^n tends to 0 is a Nash equilibrium.

The proof can be found in Appendix 6.3.

¹²In this paper, rationality is considered in the following sense: we quote Aumann and Brandenburger [4]: "Suppose that each player is rational (i.e. he maximizes his utility given his beliefs), knows his own payoff function, and knows the strategy choices of the others. Then the players' choices constitute a Nash equilibrium in the game being played."

¹³In [15], Eichberger and Kelsey use capacities to modelize ambiguity and degree of confidence in a strategic game.

¹⁴A similar idea that maximin behavior in a game can be supported by beliefs that show extreme ambiguity has first been proposed in [21].

¹⁵If we remove p-robustness, we can get similarly that any limit of λ^n -equilibria where λ^n tends to $+\infty$ is a sequence of ε^n -maximin strategy profiles for some sequence ε^n converging to 0.

For two-player games G , we can introduce the following index which compares locally the prudent behaviour of the two players. Let $x_1 \in X_1$ and let $C = \{(x', f(x')) \in V_{x_1} \times X_2\}$ be a local curve parametrized by a continuous mapping $f : V_{x_1} \rightarrow X_2$ on some neighborhood V_{x_1} of x_1 . This curve connects the strategies of the two players. In applications, f is chosen to ensure that C contains the set of Nash equilibria of G (if possible), and the index below is used to refine Nash equilibria (see Example 26, for which $f(x_1) = 1 - x_1$, or Example 30, for which $f(x_1) = x_1$).

Definition 11. *The relative prudence of player 1 with respect to player 2 at x_1 along C , when it exists, is defined by*

$$p_{1|2}^C(x_1) = \lim_{(x', x'') \rightarrow (x_1, x_1)} \frac{c_2(x', x'')}{c_1(f(x'), f(x''))}.$$

If this limit is equal to 0, we say that Player 2 is infinitely more prudent than player 1 at x_1 along C .

The index $p_{1|2}^C(x_1)$ locally measures how much player 2 takes into account possible modifications of player 1 strategy in a neighborhood of x_1 , compared with player 1, who himself takes into account possible modifications of player 2's strategy in a neighborhood of $f(x_1)$. The possible modifications of both players' strategies are constrained by C .

2.4 Characterizing our class of utility functions

Given some function $c_i \in \mathcal{F}_i$, the following proposition characterizes payoff functions v_i that are "prudent payoffs", i.e. which can be written

$$v_i(x_i, x_{-i}) = \inf_{y_{-i} \in X_{-i}} \left\{ u_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} \right\} \quad (3)$$

for some other payoff function u_i .

Proposition 12. Assume that $c_i \in \mathcal{F}_i$ defines a distance on X_{-i} , and let $\lambda > 0$. Let v_i be a payoff function quasiconcave in x_i . There exists a payoff function u_i , quasiconcave in x_i , which satisfies Equation 3 if and only if

$$\forall (x, y_{-i}) \in X \times X_{-i}, |v_i(x_i, x_{-i}) - v_i(x_i, y_{-i})| \leq \frac{c_i(y_{-i}, x_{-i})}{\lambda} \quad (4)$$

The proof can be found in Appendix 6.4.

2.5 Existence of λ -equilibria in p-robust games

We now state a first important result of this paper.

Theorem 13. *Let $c = (c_i)_{i \in N} \in \prod_{i \in N} \mathcal{F}_i$ and G be a quasiconcave and p-robust game. For every $\lambda > 0$, there exists a λ -equilibrium.*

Indeed, from Proposition 9, G^λ is a compact and quasiconcave game, u_i^λ is upper semicontinuous with respect to x and continuous with respect to x_{-i} . Thus, for every $\lambda > 0$, G^λ possesses a Nash equilibrium (see, for example, Theorem 2 in [11]).

3 Prudent equilibrium and refinement of Nash equilibrium in better-reply secure games

Refinement theory refers to the selection of particular equilibria which are more plausible. Many refinement concepts exist in the literature: in his seminal paper, Selten [29] introduces trembling hand perfect equilibrium and proves its existence in finite-strategy games. In short, his idea is to select equilibria which are immune to small mistakes of the other players, a mistake being formalized by a mixed strategy close to the initial strategy. Perfect equilibria have been refined in several directions: Myerson [24] introduces proper equilibria, where players are more likely to make mistakes in directions that are least harmful to them. Kohlberg and Mertens [18] defines some stable equilibrium notion, which requires stronger requirements than perfect equilibrium; Simon and Stinchcombe [30] extend perfect and proper equilibria to infinite normal-form games.¹⁶

To the best of our knowledge, there is no general existence result of refined Nash equilibrium (1) in pure strategies (2) which allows discontinuities of the payoff functions with respect to each player's strategy (3) without using random perturbations. This section proposes such a result. Recall that Andersson et al. [3], or Carbonell-Nicolau ([9] and [8])) introduce some refinement notion in pure strategies, but in both cases, the beliefs on other mistakes are random variables, and the author requires concavity (thus continuity, except on the boundary of strategy spaces) of the payoffs with respect to each player's strategy.

Definition 14. *A strategy profile $x \in X$ is a prudent equilibrium of G if there exists $c = (c_i)_{i \in N} \in \prod_{i \in N} \mathcal{F}_i$ such that x is the limit of λ^n -equilibria for $\lambda^n \rightarrow 0$. The strategy profile x is a strictly prudent equilibrium if this holds for any $c \in \mathcal{F}$.*

Remark 15. Andersson et al. [3] were the first to propose a notion of robustness to strategic uncertainty. They consider a family \mathcal{F} of strictly positive probability density functions ϕ_{ij} (on X_j) for each pair of distinct players $i \neq j$. For every $t > 0$, they define a t -equilibrium as a Nash equilibrium of the game whose payoff of player i at $x \in X$ is $u_i(x_i, (x_j + t\varepsilon_{ij})_{j \neq i})$, where $\varepsilon_{ij} \sim \phi_{ij}$ are statistically independent random variables. Then, *robust* and *strictly robust equilibria* are defined, as in Definition 14, by considering limits of t -equilibria when t tends to zero. Existence of robust equilibria is proved under (1) continuity of the payoffs (2) concavity of each u_i with respect to x_i . Moreover, Andersson et al. [3] proves that robust equilibrium refines Nash equilibrium when the game is continuous.

From Proposition 10 we get that Prudent equilibrium refines Nash equilibrium in better-reply secure games:

Proposition 16. For every better-reply secure game, a prudent equilibrium is a Nash equilibrium.

Theorem 17. *For every compact, quasiconcave and p -robust game, there exists a prudent equilibrium.*

Proof. By definition of prudent equilibrium, for every sequence $(\lambda^n)_{n \in \mathbb{N}}$ converging to 0, any limit point of Nash equilibria of G^{λ^n} (which exists from Theorem 13 and from compactness of G) is a prudent

¹⁶All these refinement concepts use perturbations (or "mistakes") of the equilibrium strategies. Other refinement concepts consider perturbations in the payoffs of the game (essential-equilibria [31] or regular-equilibria [16]). This is not connected with our paper.

equilibrium.

From the two results above, in every better-reply secure strategic game, prudent equilibrium concept selects some particular Nash equilibria. This selection can be affected by the choice of particular functions c_i (see Example 26), or not (see example 28). This is not surprising: indeed, c_i models strategic uncertainty in a local way, and it could influence the selection of the Nash equilibrium of the initial game.

Applications are given in Section 5. For example, it is proved that in first-price sealed-bid auctions with complete information, prudent equilibrium concept selects the unique natural solution, although there is a continuum of such solutions. In case of ties, if the winner is the player with the highest value (Example 30), then the game is better-reply secure and the unique prudent equilibrium is a Nash equilibrium from Proposition 16. If we now consider an equal sharing rule (Example 32), then the game is no more better-reply secure. There is no Nash equilibrium, but a family of approximate equilibria, and the unique prudent equilibrium is the "natural" approximate Nash equilibrium of the game. Thus, Theorem 17 can refine Nash or approximate Nash equilibrium.

4 Extensions and developments

4.1 Symmetric games

The results of the previous sections can be improved upon when the game possesses enough symmetry. Following Reny [26], a game G is *symmetric*¹⁷ if:

- (1) For every players $(i, j) \in N \times N$, $X_i = X_j$. We denote $\mathbf{X} = X_1 = \dots = X_N$.
- (2) For every $(x, y) \in \mathbf{X} \times \mathbf{X}$, $u_1(x, y, \dots, y) = u_2(y, x, y, \dots, y) = \dots = u_N(y, \dots, y, x)$. We denote $v(x, y) = u_1(x, y, \dots, y) = \dots = u_N(y, \dots, y, x)$.

Thus, a symmetric game can be summarized by $G = (\mathbf{X}, v)$.

A symmetric game $G = (\mathbf{X}, v)$ is *strongly diagonally quasiconcave* (Reny [26]) if \mathbf{X} is convex, and if $v(x, y)$ is quasiconcave in x . The game G is *diagonally better-reply secure* if for every (x^*, v^*) which belongs to $\overline{\{(x, v(x, x)) : x \in \mathbf{X}\}}$, where (x^*, \dots, x^*) is not a Nash equilibrium, then there exists $d \in \mathbf{X}$ and $\varepsilon > 0$ such that $v(d, x') > v^* + \varepsilon$ for every $x' \in \mathbf{X}$ in some neighborhood of x^* .

Finally, recall that a Nash equilibrium (x_1^*, \dots, x_N^*) is symmetric if $x_1^* = \dots = x_N^*$.

Theorem 18. (Reny [26])

Every symmetric, compact, diagonally quasiconcave and diagonally better-reply secure game possesses a symmetric pure Nash equilibrium.

We now adapt Definition 1 and Definition 8 to symmetric games:

Definition 19. *The symmetric game $G = (\mathbf{X}, v)$ is symmetrically p -robust if for every $x \in \mathbf{X}$, for every $\varepsilon > 0$ and for every neighborhood W_x of x in \mathbf{X} , there exists some open neighborhood $V_x \subset \mathbf{X}$ of x such that $\sup_{x' \in V_x} \inf_{x'' \in W_x} v(x', x'') \leq v(x, x) + \varepsilon$.*

Remark 20. Adapting Example 6, it can be easily proved that every symmetric game $G = (\mathbf{X}, v)$ is symmetrically p -robust if at every $x \in \mathbf{X}$, there exists a lower semicontinuous multivalued function

¹⁷Reny uses the terminology of "quasi-symmetric" game.

ψ^x from V_x , some open neighborhood of x into \mathbf{X} , such that $x \in \psi^x(x)$ and such that v is upper semicontinuous on $\{(y, y') \in V_x \times \mathbf{X} : y' \in \psi^x(y)\}$.

Definition 21. Let \mathcal{F} be the set of continuous real-valued functions c from $\mathbf{X} \times \mathbf{X}$ to $[0, +\infty]$, such that $c(x, x') = 0$ if and only if $x = x'$. Let $c \in \mathcal{F}$. For every $\lambda > 0$, the symmetric λ -prudent game associated to the symmetric game $G = (\mathbf{X}, v)$ is the symmetric game $G_{sym}^\lambda = (\mathbf{X}, v^\lambda)$, where

$$v^\lambda(x, x') = \inf_{x'' \in X} \left\{ v(x, x'') + \frac{c(x'', x')}{\lambda} \right\} \quad (5)$$

A symmetric λ -equilibrium of G is a symmetric Nash equilibrium of G^λ . A strategy profile $x \in X$ is a symmetric prudent equilibrium of G if there exists $c \in \mathcal{F}$ such that x is the limit of symmetric λ^n -equilibria for $\lambda^n \rightarrow 0$. The strategy profile x is a symmetric strictly prudent equilibrium if this holds for any $c \in \mathcal{F}$.

Remark 22. There are two modifications with respect to the definitions of the previous sections. First, the perturbations of the other players' strategies are assumed to be symmetric. This is the price to pay to get a symmetric prudent game, and this implies that the symmetric prudent game G_{sym}^λ is, in general, different from the prudent game G^λ (except in two-player games). Second, the sequence of λ^n -equilibria considered in Definition 21 has to be symmetric, which is a strenghtening of Definition 14. In particular, in two-player symmetric games, symmetric prudent equilibrium refines prudent equilibrium.

The proof of the following theorem is similar to those of Theorem 13 and Proposition 16.

Theorem 23. Let $c \in \mathcal{F}$ and G be a compact, symmetric, diagonally quasiconcave and symmetrically p -robust game.

- (1) For every $\lambda > 0$, there exists a symmetric λ -equilibrium.
- (2) There exists a symmetric prudent equilibrium, and for every diagonally better-reply secure game, this is a Nash equilibrium.

Applications can be found in Section 5 (see Example 28 and Example 34).

4.2 Beyond p-robust games: strategic approximation.

The idea of the previous sections can be extended to any quasiconcave game G as follows. For every finite subsets $X^f = \prod_{i \in N} X_i^f$ of X and every sequence $(\lambda^k)_{k \in \mathbf{N}}$ of positive reals converging to zero, we can quasiconcavify the prudent game G^{λ^k} on X^f as follows: for every player $i \in N$ and every integer $k \geq 0$, define \tilde{u}_i^k on $co(X^f) = \prod_{i=1}^N co(X_i^f)$ by

$$\tilde{u}_i^k(x_i, x_{-i}) = \sup\{\min\{u_i^{\lambda^k}(y_i^1, x_{-i}), \dots, u_i^{\lambda^k}(y_i^n, x_{-i})\}\}$$

over all $n \in \mathbf{N}$ and all families $\{y_i^1, \dots, y_i^n\}$ of X_i^f such that $x_i \in co\{y_i^1, \dots, y_i^n\}$. Since X^f is finite and $u_i^{\lambda^k}$ is continuous with respect to the second argument, it is easy to see that \tilde{u}_i^k is upper semicontinuous with respect to x and continuous with respect to x_{-i} . Thus, for every integer $k \geq 0$, there exists a Nash equilibrium x^k of the game $\tilde{G}^k = (co(X^f), \tilde{u}_i^k)_{i \in N}$ (see, for example, Theorem 2 in [11]). The proof of the following proposition can be found in Appendix 6.5.

Proposition 24. Let G be a quasiconcave and better-reply secure game. There exists a sequence of finite approximations $X^k \subset X$ such that any limit point of Nash equilibria of $(coX^k, \tilde{u}_i^k)_{i \in N}$ is a Nash equilibrium of G .

Following Reny [27], let us define a strategic approximation of G as "a countable subset of pure strategies with the property that limits of all equilibria of all sequences of approximating games whose finite strategy sets eventually include each member of the countable set must be equilibria of the infinite game". Thus, Proposition 24 provides a pure-strategy strategic approximation of any quasiconcave and better-reply secure game, and this strategic approximation scheme is based on prudent games.

4.3 Games in mixed strategies

Denote by $M_i = \Delta(X_i)$ the set of Borel probability measures on X_i , usually called the set of mixed strategies of player i . Recall it is a compact, Hausdorff and metrizable set under the weak* topology. To every game $G = ((X_i)_{i \in N}, (u_i)_{i \in N})$, we associate its mixed strategy extension $G' = ((M_i)_{i \in N}, (\tilde{u}_i)_{i \in N})$, where \tilde{u}_i is the multi-linear extension of u_i to M . Most of the techniques introduced in the previous sections can be applied to G' . In particular:

Corollary 25. *Let G be a compact game. Then its mixed extension G' possesses a prudent equilibrium if G' is p -robust. In particular, for every finite game G , its mixed-extension admits a prudent equilibrium.*

A difficulty is to get sufficient conditions for p -robustness in mixed strategies. In general, p -robustness of G does not imply p -robustness of its mixed extension: consider the two player game defined by $X_1 = X_2 = [0, 1]$, $u_2 = 0$, $u_1(x_1, x_2) = 0$ if $x_1 = x_2$ and $u_1(x_1, x_2) = 1$ otherwise. This game is p -robust (because it is a diagonal game: see Example 7), but its mixed extension is not: indeed, consider $(\sigma_1, \sigma_2) = (0, 0)$. If player 1 plays uniformly on a small neighborhood of 0, he obtains a payoff of 1, whatever the strategy of player 2. Thus, for every $\varepsilon > 0$ small enough, for every neighborhood V_{σ_2} of σ_2 , and for every neighborhood V_{σ_1} of σ_1 , $\sup_{\sigma'_1 \in V_{\sigma_1}} \inf_{\sigma'_2 \in V_{\sigma_2}} u_1(\sigma'_1, \sigma'_2) = 1 > u_1(\sigma_1, \sigma_2) + \varepsilon = \varepsilon$, which contradicts p -robustness of G' .

Since the main objective of this paper is to study games in pure strategies, we do not push further the case of mixed strategies. At worst, Corollary 25 can be applied to continuous games or finite games. See Example 38.

5 Examples

5.1 Nash demand game

Example 26. Some amount of money can be split between two players. Each one chooses the share he demands. Then, each player receives his demand if the demand can be satisfied, and 0 otherwise. If the total amount of money is normalized at 1, the payoff of player i is

$$u_i(x_i, x_{-i}) = \begin{cases} x_i & \text{if } x_i + x_{-i} \leq 1, \\ 0 & \text{otherwise} \end{cases}$$

This game is compact, quasiconcave and better-reply secure. The set of Nash equilibria is $\{(1, 1)\} \cup \{(x, 1 - x) : x \in [0, 1]\}$. This is a p-robust game (because it is a diagonal game: see Example 7), thus it possesses a prudent Nash equilibrium. Define $C = \{(x', 1 - x') \in [0, 1] \times [0, 1]\}$

Proposition 27. Assume c_1 and c_2 are distances on $[0, 1]$. If the relative prudence $p_{1|2}^C(x)$ of player 1 with respect to player 2 at every $x \in [0, 1]$ along C is constant, equal to $\mu \in [0, +\infty]$, then the unique efficient equilibrium is $(\frac{1}{1+\mu}, \frac{\mu}{1+\mu})$.

Proof. By definition, for every $x_{-i} < 1$ and every $x_i < 1 - x_{-i}$,

$$u_i^\lambda(x_i, x_{-i}) = \inf_{y_{-i} \in X_{-i}} \left\{ u_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} \right\} = \min \left\{ u_i(x_i, x_{-i}) = x_i, \frac{c_i(1 - x_i, x_{-i})}{\lambda} \right\}.$$

Indeed, the infimum above can be reached at $y_{-i} = x_{-i}$ or for $y_{-i} \rightarrow (1 - x_i)^+$. Moreover, for every other strategy profiles, $u_i^\lambda(x_i, x_{-i}) = 0$. If $(x_1^{\lambda_n}, x_2^{\lambda_n})$ is a sequence of Nash equilibria of G^{λ_n} which converges to (x_1, x_2) when λ_n tends to zero, then either $(x_1^{\lambda_n}, x_2^{\lambda_n}) = (1, 1)$ for infinitely many $n > 0$, and at the limit, this provides the prudent equilibrium $(1, 1)$ which is not efficient. Otherwise, up to a subsequence, we can assume that $x_1^{\lambda_n} + x_2^{\lambda_n} < 1$, otherwise one player does not play a best-response in G^{λ_n} . Moreover, since the initial game is better-reply secure, (x_1, x_2) should be a Nash equilibrium, i.e. $x_1 = 1 - x_2$. In addition, we have $\lambda_n x_i^{\lambda_n} = c_i(1 - x_i^{\lambda_n}, x_{-i}^{\lambda_n})$ for $i = 1, 2$ (thus $x_i^{\lambda_n} > 0$): indeed, otherwise, $u_i^{\lambda_n}(x)$ would be either equal to x_i or to $\frac{c_i(1 - x_i, x_{-i})}{\lambda_n}$ for every x on some neighborhood of x^{λ_n} , thus player i would be able to improve its payoff in G^{λ_n} by increasing or decreasing slightly his strategy. This implies

$$p_{1|2}^C(x_1) = \mu = \lim_{n \rightarrow +\infty} \frac{c_2(1 - x_2^{\lambda_n}, x_1^{\lambda_n})}{c_1(f(1 - x_2^{\lambda_n}), f(x_1^{\lambda_n}))} = \lim_{n \rightarrow +\infty} \frac{x_2^{\lambda_n}}{x_1^{\lambda_n}} = \frac{1 - x_1}{x_1},$$

with $f(x) = 1 - x$. Thus $(x_1, x_2) = (\frac{1}{1+\mu}, \frac{\mu}{1+\mu})$. In short, the less uncertain player obtains the bigger share.¹⁸

We should add that $(1, 1)$ being a prudent equilibrium is not surprising: our strategic uncertainty modeling is local. Around the strategy profile $(1, 1)$, any local deviation does not change the payoffs of the game. This illustrates that our concept may refine partially Nash equilibria (here, we pass from a continuum of Nash equilibria to only two). Other requirements (like efficacy) can be added to select better solutions. The same remark applies, in general, for standard refinement notions.

5.2 A Location game

Example 28. Two sellers $i = 1, 2$ sell the same good at the same price. Each seller i has to find a location along some street $x_i \in [0, 1]$. Consumers are uniformly distributed on $[0, 1]$, and each consumer chooses the closest seller. In case of ties ($x_i = x_{-i}$), it is assumed that the two sellers merge. In this case, the payoff of each player is assumed to increase of $c \in]0, \frac{1}{2}[$ (economies of scale, fixed cost elimination, etc.). The payoff of seller i is

¹⁸A similar idea was first proposed in [2], in a different setting.

$$u_i(x_i, x_{-i}) = \begin{cases} \frac{x_i + x_{-i}}{2} & \text{if } x_i < x_{-i}, \\ \frac{1}{2} + c & \text{if } x_i = x_{-i} \\ 1 - \frac{x_i + x_{-i}}{2} & \text{if } x_i > x_{-i}, \end{cases}$$

Proposition 29. The set of Nash equilibria is $\{(x, x) : x \in [\frac{1}{2} - c, \frac{1}{2} + c]\}$, and $(\frac{1}{2}, \frac{1}{2})$ is the unique symmetric prudent equilibrium.

Proof. First, this game is compact, symmetric, diagonally quasiconcave and diagonally better-reply secure, second, its set of Nash equilibria is clearly $\{(x, x) : x \in [\frac{1}{2} - c, \frac{1}{2} + c]\}$. It is a symmetrically p-robust game: indeed, for every $(x, x) \in [0, 1] \times [0, 1]$, we can define the continuous mapping $\psi^x(x') = x'$, and then, recalling that u_i is continuous on $\text{Gr}\psi^x$, we can use Remark 20. Thus, it possesses a symmetric prudent Nash equilibrium. Let us prove that the only prudent Nash equilibrium is $(\frac{1}{2}, \frac{1}{2})$. Let (x, x) be a symmetric prudent Nash equilibrium, and assume, for example, $x > \frac{1}{2}$. By definition, (x, x) is the limit of symmetric λ_n -equilibria $(x^{\lambda_n}, x^{\lambda_n})$, where λ_n converges to zero. Define $v(x, y) = u_1(x, y) = u_2(y, x)$, and v^λ as in Definition 21. Remark that $v^{\lambda_n}(x^{\lambda_n}, x^{\lambda_n}) \leq 1 - x^{\lambda_n}$ (since the other player can decrease slightly its location). Moreover, for every $\varepsilon > 0$ small enough, $v^{\lambda_n}(x^{\lambda_n} - \varepsilon, x^{\lambda_n})$ converges to $v(x - \varepsilon, x) = x + \frac{\varepsilon}{2}$, which is strictly larger than $1 - x^{\lambda_n}$ for n large enough. This contradicts that $(x^{\lambda_n}, x^{\lambda_n})$ is a λ_n -equilibria. The proof is similar when $x < \frac{1}{2}$.

5.3 First-price sealed-bid auctions

Example 30. (*First-price auction with maximum value sharing rule*)

Two bidders $i = 1, 2$ submit simultaneous sealed bids $x_i \in [0, M]$ to the seller, $M > 0$. The highest bidder wins the object and pays the value of her bid. The true values of the bidders are $v_1 < v_2 < M$. The strategy spaces are $X_1 = X_2 = [0, M]$, and the payoff of player i is defined by

$$u_i(x_i, x_{-i}) = \begin{cases} v_i - x_i & \text{if } x_i > x_{-i}, \\ 0 & \text{if } x_i < x_{-i} \end{cases}$$

Assume that in case of ties ($x_i = x_{-i}$), the winner is the bidder with the highest valuation, i.e. player 2. Let $C = \{(x, x) : x \in [0, 1]\}$.

Proposition 31. 1) The set of Nash equilibria is $\{(x, x) : x \in [v_1, v_2]\}$. 2) Assume that c_1 and c_2 are distances on $[0, M]$ and that for every strategy profile (x, x) with $x \in]v_1, v_2]$, player 1 is infinitely more prudent than player 2 at $x \in X$ along C . Then (v_1, v_1) is the only prudent Nash equilibrium.

First, this game is compact, quasi-concave and better-reply secure. Second, clearly, the set of Nash equilibria is equal to $\{(x_1, x_1) : x_1 \in [v_1, v_2]\}$. Third, this game belongs to the class of diagonal games (see Example 7), thus it is p-robust. Thus, there exists a prudent Nash equilibrium (Theorem 17 and Proposition 16).

At any Nash equilibrium (x, x) , for $x > v_1$, strategic uncertainty matters a lot for player 1, because if player 2 decreases slightly its strategy, which could be expected, then player 1's payoff decreases suddenly of $x - v_1 > 0$. Thus, player 2 could anticipate that player 1 will decrease its strategy, which is harmless for player 2. Thus, strategic uncertainty should have no importance at x for him, *compared* with player

1. This explains why the assumption that player 1 is infinitely more prudent than player 2 at $x > v_1$ along C is reasonable.

If one assumes simply that relative prudence of player 1 with respect to player 2 is $p \in [0, +\infty]$, then the set of prudent equilibria is included in $\{(x, x) : x \in [v_1, \frac{pv_1}{1+p} + \frac{v_2}{1+p}]\}$. In particular, if player 1 is not prudent at all ($p = 0$), then this coincides with the set of Nash equilibria. (See the detailed proof of Proposition 31 in Appendix 6.6).¹⁹

Example 32. (*First-price auction with equal sharing rule*)

Consider the game defined in Example 30, but now assume equal sharing rule, that is $u_i(x_i, x_{-i}) = \frac{v_i - x_i}{2}$ if $x_i = x_{-i}$. There is no Nash equilibrium, because at (x, x) , one player should deviate. But the game is still p -robust and quasiconcave, thus possesses a prudent equilibrium. Moreover, there is a continuum of approximate equilibria: for every $x \in [v_1, v_2]$, (x, x) is an approximate equilibrium, meaning that it is the limit of the $\frac{1}{n}$ -Nash equilibrium profiles $(x, x + \frac{1}{n})$. Following the proof of Proposition 31, we get the following refinement result:

Proposition 33. 1) The set of approximate Nash equilibria is $\{(x, x) : x \in [v_1, v_2]\}$. 2) Assume that c_1 and c_2 are distances on $[0, M]$. Assume that for every strategy profile (x, x) with $x \in]v_1, v_2]$, player 1 is infinitely more prudent than player 2 at $x \in X$ along C . Then (v_1, v_1) is the only prudent approximate Nash equilibrium.

5.4 Bertrand duopoly with symmetric costs

Example 34. Consider²⁰ N identical firms competing for a homogenous good. Aggregate demand is a function $D : [0, +\infty[\rightarrow [0, +\infty[$, and all firms have the same cost function $C : [0, +\infty[\rightarrow [0, +\infty[$. For every integer $m \in \{1, \dots, N\}$, let $v_m(p) = \frac{pD(p)}{m} - C(\frac{D(p)}{m})$ for every $p \in [0, +\infty[$. This is the profit of each of m firms which choose a same price p , when all other firms choose higher prices. Define a symmetric game as follows: firm $i \in N$ chooses a price $p_i \in X = [0, +\infty[$. Given a strategy profile $p = (p_i)_{i \in N}$, the payoff of firm i is

$$\pi_i(p) = \begin{cases} v_{\text{card}\{j:p_j=p_i\}}(p_i) & \text{if } p_i = \min\{p_1, \dots, p_N\}, \\ 0 & \text{otherwise} \end{cases}$$

We assume the following:²¹

1. v_1 and v_N are continuous.
2. There exists $p^{\max} > 0$ such that $v_N(p^{\max}) = v_1(p^{\max}) = 0$, and $v_N(p) < 0$ for every $p > p^{\max}$.
3. There exists a unique $\check{p}_N \in]0, p^{\max}[$ and a unique $\hat{p}_N \in]\check{p}_N, p^{\max}]$, such that $v_N(\check{p}_N) = 0$ and $v_N(\hat{p}_N) = v_1(\hat{p}_N)$.

¹⁹This result is not astonishing: indeed, prudent equilibrium is an asymptotic notion: if $p < +\infty$, it implies there is some $x \in]v_1, v_2[$ such that (x, x) is a prudent equilibrium: recall this means that (x, x) is the limit of Nash equilibria $(x_1^{\lambda_n}, x_2^{\lambda_n})$ of the strategic game G^{λ_n} when λ_n tends to zero. If $x_1^{\lambda_n} < x_2^{\lambda_n}$ (which has to be true, see Appendix 6.6) and if the difference $x_1^{\lambda_n} - x_2^{\lambda_n}$ is large enough, then player 1 gets 0 when the profile $(x_1^{\lambda_n}, x_2^{\lambda_n})$ is played, and the "risk" that player 2 deviates sufficiently for some strategy $x_2 < x_1^{\lambda_n}$ could be thought to be small by Player 1, *considering his level of prudence*. This is no more the case if $p = +\infty$.

²⁰This presentation of Bertrand duopoly model follows partially [3].

²¹See [12] or [3] for natural assumptions on C and D that imply these properties.

4. There exists a unique price $\bar{p} \in]\check{p}_N, \hat{p}_N[$ such that $v_1(\bar{p}) = 0$ (price at which a monopolist makes a zero profit), where $v_1(p) > 0$ for every $p > \bar{p}$ and $v_1(p) < 0$ for every $p < \bar{p}$.
5. For every $p > \hat{p}_N$ $v_1(p) > v_N(p)$.
6. $\pi_i(p_1, p_2, \dots, p_2)$ is quasiconcave in p_1 on $\mathbf{X} = [\check{p}_N, p^{\max}]$ for every $p_2 \in \mathbf{X}$.
7. For every $p_i < \check{p}_N$ and every $m \in \{1, \dots, N\}$, $v_m(p_i) < 0$.

From these assumptions, without any loss of generality, we can restrict the strategies to $\mathbf{X} = [\check{p}_N, p^{\max}]$, since it is neither optimal for player i to play outside \mathbf{X} . Call G the game thus defined.

In the following proposition, we assume that c is a distance on $[0, +\infty[$.

Proposition 35. (1) The set of Nash equilibria of the game G above is $\{(p, \dots, p) : p \in [\check{p}_N, \hat{p}_N]\}$.
(2) $(\bar{p}, \dots, \bar{p})$ is the unique symmetric prudent equilibrium of G .

Remark that this game is symmetrically p-robust, symmetric, compact, diagonally quasiconcave and diagonally better-reply secure. See the proof in Appendix 6.7. Note that [3] gets a similar result in a different framework.

5.5 Link with the notion of strategic uncertainty of Andersson et al.

The two examples below show that prudent equilibrium concept does not coincide with Andersson et al. robust equilibrium concept.

Example 36. In this example, we provide a differentiable game where Andersson et al. robustness concept refines prudent equilibrium. Consider a two-player game with $X_1 = X_2 = [0, 1]$,

$$u_1(x_1, x_2) = \begin{cases} x_1 \cdot (2x_2 - 1) & \text{if } \frac{1}{2} \leq x_2 \leq 1, \\ 0 & \text{otherwise} \end{cases}$$

and

$$u_2(x_1, x_2) = -(x_1 - x_2)^2$$

This game is p-robust (because each u_i is u.s.c. in x_i). It is also compact and quasiconcave. The strategy profile $(0, 0)$ is not robust to uncertainty. Indeed, for every strictly positive probability density function ϕ_{12} (on X_2), $\varepsilon_{12} \sim \phi_{12}$ be a random variable, and every $(x_1^{t_n}, x_2^{t_n})$ converging to $(0, 0)$ where t_n converges to 0, we have $u_1(x_1^{t_n}, x_2^{t_n} + t_n \varepsilon_{12}) < u_1(1, x_2^{t_n} + t_n \varepsilon_{12})$, thus $(x_1^{t_n}, x_2^{t_n})$ cannot be an equilibrium of the perturbed game (in the sense of Andersson et al.). Clearly, only $(1, 1)$ is robust to uncertainty. Yet $(0, 0)$ is a prudent equilibrium: indeed, $u^\lambda(0, 0) = 0 \geq u^\lambda(x_1, 0) = 0$ for every $x_1 \in [0, 1]$ and for $\lambda > 0$ small enough.

Example 37. We provide an example of better-reply secure game where the unique Nash equilibrium is a prudent equilibrium, though there is a continuum of Nash equilibria robust to uncertainty in the sense of Andersson et al. [3]. Consider the two-player game with $X_i = [0, 1]$ for $i = 1, 2$ and

$$u_1(x_1, x_2) = \begin{cases} 0 & \text{if } x_2 = \frac{1}{2} \text{ and } x_1 \neq \frac{1}{2}, \\ 1 & \text{otherwise,} \end{cases}$$

and

$$u_2(x_1, x_2) = \begin{cases} 0 & \text{if } x_2 \neq \frac{1}{2}, \\ 1 & \text{otherwise,} \end{cases}$$

The unique Nash equilibrium is $(\frac{1}{2}, \frac{1}{2})$. This game is p-robust (because each u_i is u.s.c. in x_i). It is better-reply secure: indeed, for every strategy profiles $(x_1, \frac{1}{2})$ with $x_1 \neq \frac{1}{2}$, limit points of strategy profiles at $(x_1, \frac{1}{2})$ point can be $(1, 0)$ or $(0, 1)$. In the first case, player 2 can secure strictly more than 0 by playing $\frac{1}{2}$, in the second case, player 1 can secure strictly more than 0 by playing $\frac{1}{2}$. From Theorem 17, there exists a prudent equilibrium, and from Proposition 16, it should be the (unique) Nash equilibrium $(\frac{1}{2}, \frac{1}{2})$. Yet, for every strictly positive probability density function ϕ_{12} (on X_2), $\varepsilon_{12} \sim \phi_{12}$ a random variable, and $t > 0$, we have $u_1(x_1, x_2 + t\varepsilon_{12}) = 1$ for every x_1 and x_2 in $[0, 1]$. Similarly, the modified payoff function of player 2, when strategic uncertainty is modeled through some random variable with strictly positive probability density, is 0 if $x_2 \neq \frac{1}{2}$ and 1 otherwise. Thus, every strategy profile $(x_1, \frac{1}{2})$, $x_1 \in [0, 1]$, is robust to uncertainty in the sense of Andersson et al. (even if x is not a Nash equilibrium).

5.6 Prudent equilibrium refines Nash equilibrium in mixed strategies

Example 38. Consider the following two-player finite game:

	A	B
a	(0, 0)	(-1, 0)
b	(0, -1)	(0, 0)

Let G be the mixed extension of this game. Every strategy of player 1 (resp. of player 2) can be written $(\lambda, 1 - \lambda) \in [0, 1]^2$, where λ is the probability of a (resp. A), and $1 - \lambda$ is the probability of b (resp. B). Assume that c is quadratic, i.e. more precisely that $c((\lambda, 1 - \lambda), (\lambda', 1 - \lambda')) = (\lambda - \lambda')^2$. There are two Nash equilibria which are (a, A) and (b, B) . Only (b, B) is perfect.

Let us prove that (a, A) is not a prudent equilibrium. Indeed, by contradiction, assume that (σ_1^n, σ_2^n) is a Nash equilibrium of G^{λ^n} for $\lambda^n \rightarrow 0$, where (σ_1^n, σ_2^n) converges to $(a, A) = ((1, 0), (1, 0))$. Then, $u_1^{\lambda^n}(\sigma_1^n, \sigma_2^n) \geq u_1^{\lambda^n}(b, \sigma_2^n) = 0$ which is only possible if $\sigma_2^n = A$, and by a symmetric argument $\sigma_1^n = a$, and (a, A) is a Nash equilibrium of G^{λ^n} for n large enough. But then, by definition of $u_1^{\lambda^n}$,

$$u_1^{\lambda^n}(a, A) \leq u_1^{\lambda^n}(a, \alpha_n B + (1 - \alpha_n)A) + \frac{c(\alpha_n B + (1 - \alpha_n)A, A)}{\lambda^n} = -\alpha_n + \frac{\alpha_n^2}{\lambda^n}$$

where $\alpha_n > 0$ can be chosen small enough for every $n > 0$ such that $-\alpha_n + \frac{\alpha_n^2}{\lambda^n}$ is negative. Then we get $u_1^{\lambda^n}(a, A) < 0 = u_1^{\lambda^n}(b, A)$. But this contradicts that (a, A) is a Nash equilibrium of G^{λ^n} .

6 Appendix

6.1 Proof of Proposition 6

Let G satisfies the assumption of Proposition 6. First, we prove the following claim:

Claim 39. For every open neighborhood $V_{x_{-i}}$ of $x_{-i} \in X_{-i}$, $\inf_{x'_{-i} \in V_{x_{-i}}} u_i(x_i, x'_{-i})$ is upper semicontinuous with respect to x_i .

Proof. Let $a \in \mathbb{R}$, $V_{x_{-i}}$ be an open neighborhood of $x_{-i} \in X_{-i}$, and $(x_i^n)_{n \in \mathbb{N}}$ be a sequence of strategies converging to $x_i \in X$, and such that for every integer n ,

$$\inf_{x'_{-i} \in V_{x_{-i}}} u_i(x_i^n, x'_{-i}) \geq a.$$

By assumption, there exists $V_{x_i} \subset X_i$, an open neighborhood of x_i , and there exists a lower semicontinuous multivalued function ψ^x from V_{x_i} to X_{-i} such that $x_{-i} \in \psi^x(x_i)$, and such that the restriction of u_i so $\text{Gr}\psi^x$ is upper semicontinuous at x .

From lower semicontinuity of ψ^x , considering a subsequence of $(x_i^n)_{n \in \mathbb{N}}$ if necessary, there exists a sequence x_{-i}^n converging to x_{-i} such that $x_{-i}^n \in \psi^x(x_i^n)$ (see Aliprentis [1], Theorem 17.19). For n large enough, $x_{-i}^n \in V_{x_{-i}}$, thus $u_i(x_i^n, x_{-i}^n) \geq a$. Since $x^n \in \text{Gr}\psi^x$, and since the restriction of u_i to $\text{Gr}\psi^x$ is upper semicontinuous at x , passing to the limit, we get $u_i(x_i, x_{-i}) \geq a$. Passing to the infimum with respect to $x_{-i} \in V_{x_{-i}}$, we get

$$\inf_{x'_{-i} \in V_{x_{-i}}} u_i(x_i, x'_{-i}) \geq a,$$

which ends the proof of the claim.

Now, to prove that G is p-robust, consider $\varepsilon > 0$ and let $V_{x_{-i}}$ be an open neighborhood of x_{-i} . From the claim above, there exists V_{x_i} such that for every $x'_i \in V_{x_i}$,

$$\inf_{x'_{-i} \in V_{x_{-i}}} u_i(x'_i, x'_{-i}) \leq \inf_{x'_{-i} \in V_{x_{-i}}} u_i(x_i, x'_{-i}) + \varepsilon \leq u_i(x) + \varepsilon,$$

which proves p-robustness.

6.2 Proof of Proposition 9

(1) Remark that $u_i^\lambda(x_i, x_{-i})$ is the infimum of a family of functions which are quasiconcave in x_i , thus it is quasiconcave in x_i .

(2) First note that $u_i^\lambda(x_i, x_{-i})$ is the infimum of a family of functions which are continuous in x_{-i} . Thus it is upper semicontinuous in x_{-i} . To prove it is lower semicontinuous in x_{-i} , consider a sequence $(x_{-i}^n)_{n \in \mathbb{N}}$ converging to some x_{-i} , and such that $u_i^\lambda(x_i, x_{-i}^n) \leq \alpha$ for some real α and for every integer $n \geq 0$. By definition,

$$\inf_{y_{-i} \in X_{-i}} \left\{ u_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i}^n)}{\lambda} \right\} \leq \alpha$$

for every integer n . Given $\varepsilon > 0$, this implies that there is a sequence $y_{-i}^n \in X_{-i}$ such that

$$u_i(x_i, y_{-i}^n) + \frac{c_i(y_{-i}^n, x_{-i}^n)}{\lambda} \leq \alpha + \varepsilon.$$

Since c_i is continuous on the compact set $X_{-i} \times X_{-i}$, it is uniformly continuous, thus

$$u_i(x_i, y_{-i}^n) + \frac{c_i(y_{-i}^n, x_{-i})}{\lambda} \leq \alpha + 2\varepsilon$$

for n large enough. Passing to the infimum with respect to the second variable y_{-i}^n , we get $u_i^\lambda(x_i, x_{-i}) \leq \alpha + 2\varepsilon$. Since this is true for every $\varepsilon > 0$, this finally proves (2).

(3) Now, assume that G is p-robust, and prove that u_i^λ is u.s.c. with respect to x . Take $a \in \mathbb{R}$ and consider a sequence $(x^n)_{n \in \mathbb{N}}$ of strategy profiles converging to $x \in X$, and such that $u_i^\lambda(x^n) \geq a$ for every integer n . We

have to prove that $u_i^\lambda(x) \geq a$. By definition of $(x^n)_{n \in \mathbf{N}}$, for every integer $n \in \mathbf{N}$, we get

$$\inf_{y_{-i} \in X_{-i}} \left\{ u_i(x_i^n, y_{-i}) + \frac{c_i(y_{-i}, x_{-i}^n)}{\lambda} \right\} \geq a.$$

Let $\varepsilon > 0$ and $y_{-i} \in X_{-i}$. By p-robustness, for every integer $k \in \mathbf{N}$, choosing $V_{y_{-i}} = B(y_{-i}, \frac{1}{k})$, there is some open neighborhood $V_{x_i}^k$ of x_i such that: for every $x'_i \in V_{x_i}^k$, there exists $y_{-i}^k \in B(y_{-i}, \frac{1}{k})$ such that

$$u(x'_i, y_{-i}^k) \leq u_i(x_i, y_{-i}) + \varepsilon$$

In particular, for every k , there is n_k large enough such that for every $n \geq n_k$, there is $y_{-i}^n \in B(y_{-i}, \frac{1}{k})$ such that

$$u_i(x_i^n, y_{-i}^n) \leq u_i(x_i, y_{-i}) + \varepsilon$$

Thus

$$u_i(x_i^n, y_{-i}^n) + \frac{c_i(y_{-i}^n, x_{-i}^n)}{\lambda} \leq u_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} + \left(\frac{c_i(y_{-i}^n, x_{-i}^n)}{\lambda} - \frac{c_i(y_{-i}, x_{-i})}{\lambda} \right) + \varepsilon,$$

From continuity of c_i , for k large enough and $n \geq n_k$, we get

$$u_i(x_i^n, y_{-i}^n) + \frac{c_i(y_{-i}^n, x_{-i}^n)}{\lambda} \leq u_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} + 2\varepsilon$$

Passing to the infimum with respect to y_{-i} in the right-hand side, then to the infimum with respect to y_{-i}^n in the left-hand side, we get, for n large enough:

$$a \leq u_i^\lambda(x^n) \leq u_i^\lambda(x) + 2\varepsilon.$$

Consequently, u_i^λ is an u.s.c. function of x

(4) For the last point, take $x \in X$, and consider a sequence $(x_{-i}^n)_{n \in \mathbf{N}}$ converging to x_{-i} such that $\underline{u}_i(x) = \lim_{n \rightarrow +\infty} u_i(x_i, x_{-i}^n)$. By definition

$$u_i^\lambda(x_i, x_{-i}) = \inf_{y_{-i} \in X_{-i}} \left\{ u_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} \right\} \leq u_i(x_i, x_{-i}^n) + \frac{c_i(x_{-i}^n, x_{-i})}{\lambda}.$$

Passing to the limit, we get $u_i^\lambda(x_i, x_{-i}) \leq \underline{u}_i(x_i, x_{-i})$.

6.3 Proof of Proposition 10

(1) Let $(\lambda_n)_{n \in \mathbf{N}}$ be a sequence of non negative reals converging to $+\infty$, and $(x^n)_{n \in \mathbf{N}}$ be a sequence of λ^n -equilibria converging to $x \in X$. By definition of a Nash equilibrium in the game G^{λ_n} , we get

$$\forall i \in N, \forall d_i \in X_i, \inf_{y_{-i} \in X_{-i}} \left\{ u_i(d_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i}^n)}{\lambda_n} \right\} \leq \inf_{y_{-i} \in X_{-i}} \left\{ u_i(x_i^n, y_{-i}) + \frac{c_i(y_{-i}, x_{-i}^n)}{\lambda_n} \right\} \quad (6)$$

Since c_i is continuous and X is compact, for every $\varepsilon > 0$, there exists n large enough such that

$$\forall i \in N, \forall d_i \in X_i, \inf_{y_{-i} \in X_{-i}} \left\{ u_i(d_i, y_{-i}) \right\} \leq \inf_{y_{-i} \in X_{-i}} \left\{ u_i(x_i^n, y_{-i}) \right\} + \varepsilon \quad (7)$$

i.e. x_i^n is a ε -maximin strategy for player i for n large enough. From Claim 39, $\inf_{y_{-i} \in X_{-i}} \{u_i(x_i, y_{-i})\}$ is u.s.c. with respect to x_i . Thus, passing to the limit when $n \rightarrow +\infty$ and $\varepsilon \rightarrow 0$ in the above inequality, we get

$$\forall i \in N, \forall d_i \in X_i, \inf_{y_{-i} \in X_{-i}} \{u_i(d_i, y_{-i})\} \leq \inf_{y_{-i} \in X_{-i}} \{u_i(x_i, y_{-i})\} \quad (8)$$

that is x is a maximin strategy profile.

(2) We will use the following Claim, which lists some additional important properties of u_i^λ .

Claim 40. Let $G = ((X_i)_{i \in N}, (u_i)_{i \in N})$ be a game, and $G^\lambda = ((X_i)_{i \in N}, (u_i^\lambda)_{i \in N})$ be the λ -prudent game associated to G .

1. If $x_{-i}^n \rightarrow x_{-i}$ and $\lambda_n \rightarrow 0$ then for every $x_i \in X_i$, $\liminf_{n \rightarrow +\infty} u_i^{\lambda_n}(x_i, x_{-i}^n) \geq \underline{u}_i(x_i, x_{-i})$.
2. For every $x \in X$, $u_i^{\lambda_n}(x_i, x_{-i})$ tends to $\underline{u}_i(x_i, x_{-i})$ when λ_n tends to 0.
3. $u_i^\lambda(x_i, x_{-i}) = \inf_{y_{-i} \in X_{-i}} \{ \underline{u}_i(x_i, y_{-i}) + \frac{c_i(x_{-i}, y_{-i})}{\lambda} \}$

Proof of Claim 40. For the first point of the claim, let $x_{-i}^n \rightarrow x_{-i}$ and consider a sequence $(\lambda_n)_{n \in \mathbf{N}}$ converging to 0. By definition,

$$u_i^{\lambda_n}(x_i, x_{-i}^n) = \inf_{y_{-i} \in X_{-i}} \{ u_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i}^n)}{\lambda_n} \}.$$

Let $\varepsilon > 0$. By definition of infimum, there is a sequence $y_{-i}^n \in X_{-i}$ such that

$$u_i^{\lambda_n}(x_i, x_{-i}^n) \geq u_i(x_i, y_{-i}^n) + \frac{c_i(y_{-i}^n, x_{-i}^n)}{\lambda_n} - \varepsilon.$$

Since the sequence $u_i^{\lambda_n}(x_i, x_{-i}^n)$ is bounded by the maximum of u_i , this implies that the sequence $\frac{c_i(y_{-i}^n, x_{-i}^n)}{\lambda_n}$ is bounded, thus y_{-i}^n converges to x_{-i} (because λ_n converges to 0). Moreover, since $c_i \geq 0$, we get

$$u_i^{\lambda_n}(x_i, x_{-i}^n) \geq u_i(x_i, y_{-i}^n) - \varepsilon,$$

and passing to the infimum limit as $n \rightarrow +\infty$, then taking $\varepsilon \rightarrow 0$, we get the first point.

For the second point of the claim, use the first point proved above with a constant sequence $x_{-i}^n = x_{-i}$, and Point 4 in Proposition 9.

For the last point of the claim, first note that the inequality $u_i^\lambda(x_i, x_{-i}) \geq \inf_{y_{-i} \in X_{-i}} \{ \underline{u}_i(x_i, y_{-i}) + \frac{c_i(x_{-i}, y_{-i})}{\lambda} \}$ is obvious. To prove the converse inequality, let $\varepsilon > 0$ and $\bar{y}_{-i} \in X_{-i}$ such that $\inf_{y_{-i} \in X_{-i}} \{ \underline{u}_i(x_i, y_{-i}) + \frac{c_i(x_{-i}, y_{-i})}{\lambda} \} \geq \underline{u}_i(x_i, \bar{y}_{-i}) + \frac{c_i(x_{-i}, \bar{y}_{-i})}{\lambda} - \varepsilon$. By definition of \underline{u}_i and from the continuity of c_i , there exists a sequence $y_{-i}^n \in X_{-i}$, converging to \bar{y}_{-i} , such that $\underline{u}_i(x_i, \bar{y}_{-i}) + \frac{c_i(x_{-i}, \bar{y}_{-i})}{\lambda} - \varepsilon \geq u_i(x_i, y_{-i}^n) + \frac{c_i(x_{-i}, y_{-i}^n)}{\lambda} - 2\varepsilon$ for n large enough. Passing to the infimum with respect to the second variable in the right-hand side, we finally get $\inf_{y_{-i} \in X_{-i}} \{ \underline{u}_i(x_i, y_{-i}) + \frac{c_i(x_{-i}, y_{-i})}{\lambda} \} \geq u_i^\lambda(x_i, x_{-i}) - 2\varepsilon$, which ends the proof of the Claim.

Now, we prove the second part of Proposition 10. Assume G is a better-reply secure game, let $(\lambda_n)_{n \in \mathbf{N}}$ be a sequence of non negative reals converging to 0, and $(x^n)_{n \in \mathbf{N}}$ be a sequence of λ_n -equilibria which converges to x . From compactness of Γ , without any loss of generality, up to a subsequence, we can assume that $(x^n, u(x^n))$ converges to some $(x, v) \in \Gamma$. By definition,

$$\forall i \in N, \forall d_i \in X_i, u_i^{\lambda_n}(d_i, x_{-i}^n) \leq u_i^{\lambda_n}(x^n) \leq u_i(x^n), \quad (9)$$

the last inequality being a consequence of Proposition 9. Passing to the infimum limit as $n \rightarrow +\infty$, and using

Point 1. in the above claim, we get $\underline{u}_i(d_i, x_{-i}) \leq u_i$ for every $i \in N$ and every $d_i \in X_i$. Since G is better-reply secure, this implies that x is a Nash equilibrium.

6.4 Proof of Proposition 12

First, assume that v_i satisfies Equation 3. For every $(x, y_{-i}, z_{-i}) \in X \times X_{-i} \times X_{-i}$, from Equation 3, and since c_i is a distance, we get

$$v_i(x_i, x_{-i}) \leq u_i(x_i, z_{-i}) + \frac{c_i(z_{-i}, x_{-i})}{\lambda} \leq u_i(x_i, z_{-i}) + \frac{c_i(z_{-i}, y_{-i})}{\lambda} + \frac{c_i(y_{-i}, x_{-i})}{\lambda}. \quad (10)$$

Passing to the infimum with respect to z_{-i} in Equation 10, we get

$$v_i(x_i, x_{-i}) \leq v_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda},$$

thus v_i satisfies Equation 4.

For the converse implication, we now assume that v_i satisfies Equation 4, and we will prove that Equation 3 is true with $u_i = v_i$. By definition, we have

$$v_i(x_i, x_{-i}) \geq \inf_{y_{-i} \in X_{-i}} \left\{ v_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} \right\}$$

since we can take $y_{-i} = x_{-i}$ in the infimum and since $c(x_{-i}, x_{-i}) = 0$. For the converse inequality, remark that from Equation 4, we have

$$v_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} \geq v_i(x_i, x_{-i})$$

for every $y_{-i} \in X_{-i}$, and passing to the infimum with respect to y_{-i} we finally get

$$v_i(x_i, x_{-i}) = \inf_{y_{-i} \in X_{-i}} \left\{ v_i(x_i, y_{-i}) + \frac{c_i(y_{-i}, x_{-i})}{\lambda} \right\}$$

6.5 Proof of Proposition 24

We first prove the following lemma:

Lemma 41. If G is quasiconcave and better-reply secure, then there exists a countable set $\prod_{i=1}^N X'_i \subset X$ such that for every $(x^*, u^*) \in \Gamma$, if x^* is not a Nash equilibrium of G , then there exists $i \in N$ and $d_i \in X'_i$ such that $\underline{u}_i(d_i, x_{-i}^*) > u_i^*$.

Proof of Lemma 41. Let $\Gamma^{equ} = \{(x^*, u^*) \in \Gamma, x^* \text{ is a Nash equilibrium}\}$ and $\Gamma^{nequ} = \{(x^*, u^*) \in \Gamma, x^* \text{ is not a Nash equilibrium}\} = {}^c \Gamma^{equ}$. For every $(x^*, u^*) \in \Gamma^{nequ}$, from better-reply security, there exists some player i and some strategy $d \in X_i$ such that $u_i^* < \underline{u}_i(d, x_{-i}^*)$. Since by definition \underline{u}_i is l.s.c. with respect to x_{-i} , there exists an open neighborhood $V_{(x^*, u^*)}(d)$ of (x^*, u^*) in Γ^{nequ} such that for every $(x, u) \in V_{(x^*, u^*)}(d)$, $u_i < \underline{u}_i(d, x_{-i})$. Since Γ is a compact subset of a metric space, it is separable. Thus Γ^{nequ} , as a subset of a separable metric space, is separable. Thus, it is a Lindelöf space²², i.e. every open cover of Γ^{nequ} has a countable subcover. Consequently, there exists a countable cover \mathcal{O} of Γ^{nequ} by some open neighborhoods $V_{x^*(j), u^*(j)}(d(j))$, where $(x^*(j), u^*(j)) \in \Gamma^{nequ}$, $d(j) \in \cup_{i=1}^N X_i$ and $j \in \mathbb{N}$. Now, define $X'_i = \{d(j), j \in \mathbb{N}\} \cap X_i$ if it is nonempty, and X'_i be any singleton in X_i otherwise. By construction, it satisfies the conclusion of Lemma 41.

To prove Proposition 24, consider an increasing sequence of finite subsets $X^k = \prod_{i=1}^N X_i^k$ of X such that $\cup_k X^k = X'$ (X' being defined in Lemma 41) and take a sequence $(x^k)_{k \in \mathbb{N}}$ of Nash equilibria of the games

²²Let X be a separable metric space (which means that there exists C , a countable and dense subset of X). Then X is a Lindelöf space, i.e. every open cover of X has a countable subcover.

$(coX^k, \tilde{u}_i^k)_{i \in N}$. By compactness of X , without any loss of generality, we can suppose that x^k converges to $x^* \in X$. By definition of \tilde{u}_i^k and from Point 4. in Proposition 9, we have

$$\tilde{u}_i^k(x_i^k, x_{-i}^k) = \sup\{\min\{u_i^{\lambda^k}(y_i^1, x_{-i}^k), \dots, u_i^{\lambda^k}(y_i^n, x_{-i}^k)\}\} \leq \sup\{\min\{u_i(y_i^1, x_{-i}^k), \dots, u_i(y_i^n, x_{-i}^k)\}\},$$

the supremum being taken over all $n \in \mathbb{N}$ and all families $\{y_i^1, \dots, y_i^n\}$ of X_i^k such that $x_i^k \in co\{y_i^1, \dots, y_i^n\}$. From quasiconcavity of u_i with respect to x_i , we finally get $\tilde{u}_i^k(x_i^k, x_{-i}^k) \leq u_i(x_i^k, x_{-i}^k)$. In addition, the definition of \tilde{u}_i^k gives $u_i^{\lambda^k}(d_i, x_{-i}^k) \leq \tilde{u}_i^k(d_i, x_{-i}^k)$ for every $d_i \in X_i^k$ (since we can take, in the supremum defining $\tilde{u}_i^k(d_i, x_{-i}^k)$, $n = 1$ and $y_i^1 = d_i$).

Now, fix $d_i \in X_i'$. For $k > 0$ large enough, $d_i \in X_i^k$, and by definition of x^k we get

$$u_i^{\lambda^k}(d_i, x_{-i}^k) \leq \tilde{u}_i^k(d_i, x_{-i}^k) \leq \tilde{u}_i^k(x^k) \leq u_i(x^k).$$

Passing to the limit as $k \rightarrow +\infty$, and from Point 1. in Claim 40, we get

$$\forall d_i \in X_i', \underline{u}_i(d_i, x_{-i}^*) \leq u_i^*$$

where $(x^*, u^*) \in \Gamma$. This proves that x^* is a Nash equilibrium by construction of X' in Lemma 41.

6.6 Proof of Proposition 31

We want to prove that (v_1, v_1) is the unique prudent equilibrium. Assume by contradiction that there is a sequence of λ_n -equilibria $x^{\lambda_n} = (x_1^{\lambda_n}, x_2^{\lambda_n})$ which converges to (x, x) when $\lambda_n \rightarrow 0^+$, where $x \in]v_1, v_2]$ (recall that here, a prudent equilibrium is a Nash equilibrium, because the game is p-robust). If $x_1^{\lambda_n} \geq x_2^{\lambda_n}$ for infinitely many $n > 0$, then, up to a subsequence, we can assume $x_1^{\lambda_n} \geq x_2^{\lambda_n}$ for every $n > 0$. Then $\underline{u}_1(x^{\lambda_n}) = v_1 - x_1^{\lambda_n} < 0$ for n large enough (indeed, one can slightly decrease $x_2^{\lambda_n}$). Thus from Point 4. in Proposition 9, $u_1^{\lambda_n}(x^{\lambda_n}) \leq v_1 - x_1^{\lambda_n} < 0$ for n large enough. But, by definition, x^{λ_n} is a Nash of G^{λ_n} , thus for n large enough, we get $0 = u_1^{\lambda_n}(0, x_2^{\lambda_n}) \leq u_1^{\lambda_n}(x^{\lambda_n}) < 0$, a contradiction. Now, we can assume that $x_1^{\lambda_n} < x_2^{\lambda_n}$ for n large enough. By definition,

$$u_2^{\lambda_n}(x^{\lambda_n}) = \inf_{y_1 \in [0, M]} \left\{ u_2(y_1, x_2^{\lambda_n}) + \frac{c_2(y_1, x_1^{\lambda_n})}{\lambda_n} \right\},$$

or equivalently

$$u_2^{\lambda_n}(x^{\lambda_n}) = \min \left\{ u_2(x^{\lambda_n}) = v_2 - x_2^{\lambda_n}, \inf_{y_1 > x_2^{\lambda_n}} \left\{ u_2(y_1, x_2^{\lambda_n}) + \frac{c_2(y_1, x_1^{\lambda_n})}{\lambda_n} = \frac{c_2(y_1, x_1^{\lambda_n})}{\lambda_n} \right\} \right\}$$

because $u_2(x^{\lambda_n}) = v_2 - x_2^{\lambda_n} < u_2(y_1, x_2^{\lambda_n}) + \frac{c_2(y_1, x_1^{\lambda_n})}{\lambda_n} = v_2 - x_2^{\lambda_n} + \frac{c_2(y_1, x_1^{\lambda_n})}{\lambda_n}$ for every $0 \leq y_1 \leq x_2^{\lambda_n}$ and $y_1 \neq x_1^{\lambda_n}$. Remark also that $\inf_{y_1 > x_2^{\lambda_n}} \left\{ \frac{c_2(y_1, x_1^{\lambda_n})}{\lambda_n} \right\} = \frac{c_2(x_2^{\lambda_n}, x_1^{\lambda_n})}{\lambda_n}$ because c_2 is a distance. Thus finally,

$$u_2^{\lambda_n}(x^{\lambda_n}) = \min \left\{ v_2 - x_2^{\lambda_n}, \frac{c_2(x_2^{\lambda_n}, x_1^{\lambda_n})}{\lambda_n} \right\}.$$

By continuity, this equality is true on some open neighborhood of x^{λ_n} . If $v_2 - x_2^{\lambda_n} < \frac{c_2(x_2^{\lambda_n}, x_1^{\lambda_n})}{\lambda_n}$, then we get $u_2^{\lambda_n}(x_1^{\lambda_n}, x_2') = v_2 - x_2'$ for x_2' close enough to $x_2^{\lambda_n}$, which contradicts that $x_2^{\lambda_n}$ is a best reply of player 2 (in G^{λ_n}) to $x_1^{\lambda_n}$. A similar argument holds if $v_2 - x_2^{\lambda_n} > \frac{c_2(x_2^{\lambda_n}, x_1^{\lambda_n})}{\lambda_n}$. Thus, for n large enough, we can assume $v_2 - x_2^{\lambda_n} = \frac{c_2(x_2^{\lambda_n}, x_1^{\lambda_n})}{\lambda_n}$. In particular, the sequence $\frac{c_2(x_2^{\lambda_n}, x_1^{\lambda_n})}{\lambda_n}$ converges to $v_2 - x$.

Define

$$p_n = \frac{c_2(x_2^{\lambda_n}, x_1^{\lambda_n})}{c_1(x_2^{\lambda_n}, x_1^{\lambda_n})}.$$

By assumption, p_n converges to $p = \infty$.

Recalling that $x_1^{\lambda_n}$ is a best reply to $x_2^{\lambda_n}$ for player 1, we get for every $\varepsilon > 0$,

$$0 = u_1^{\lambda_n}(0, x_2^{\lambda_n}) \leq u_1^{\lambda_n}(x^{\lambda_n}) \leq u_1(x_2^{\lambda_n} + \varepsilon, x_2^{\lambda_n}) + \frac{c_1(x_1^{\lambda_n}, x_2^{\lambda_n} + \varepsilon)}{\lambda_n} = v_1 - x_2^{\lambda_n} - \varepsilon + \frac{c_1(x_1^{\lambda_n}, x_2^{\lambda_n} + \varepsilon)}{\lambda_n}, \quad (11)$$

the first inequality being a consequence of x^{λ_n} being a Nash of G^{λ_n} , and the second inequality being a consequence of the definition of $u_1^{\lambda_n}$. Passing to the limit when ε tends to zero, we get $\frac{c_1(x_1^{\lambda_n}, x_2^{\lambda_n})}{\lambda_n} \geq x_2^{\lambda_n} - v_1$ for every n .

Thus, $\frac{c_2(x_1^{\lambda_n}, x_2^{\lambda_n})}{\lambda_n} \geq p_n(x_2^{\lambda_n} - v_1)$. Passing at the limit, we get $v_2 - x \geq p(x - v_1)$, that is

$$x \leq \frac{pv_1}{1+p} + \frac{v_2}{1+p}.$$

In particular, since $p = +\infty$, we get only one prudent equilibrium (v_1, v_1) . If we simply assume $p \in [0, +\infty]$, we get that the set of prudent equilibria is included in $\{(x, x) : x \in [v_1, \frac{pv_1}{1+p} + \frac{v_2}{1+p}]\}$.

6.7 Proof of Proposition 35

(1) can be found in Dastidar [12].

(2) Let

$$v(p_1, p_2) = \begin{cases} v_N(p_1) & \text{if } p_1 = p_2, \\ v_1(p_1) & \text{if } p_1 < p_2, \\ 0 & \text{otherwise} \end{cases}$$

and $\mathbf{X} = [\hat{p}_N, p^{\max}]$. By definition, the game $G = (\mathbf{X}, v)$ is symmetric, since $\pi_i(p_i, p_2, \dots, p_2) = v(p_i, p_2)$ for every $(p_i, p_2) \in \mathbf{X}^2$. The game G is symmetrically p-robust: indeed, we only have to check p-robustness at (p, p) , and we can define the continuous mapping $\psi^p(p') = p'$ for every $p' \in [0, 1]$, and use the fact that the restriction of v to $\{(p', p') : p' \in [0, 1]\}$ is continuous. Then, see Remark 20. In addition, the game is strongly diagonally quasiconcave when the strategies are restricted to $[\hat{p}_N, p^{\max}]$, by Assumption 6 in the model.

To prove that G is *diagonally better-reply secure*, let (p^*, v^*) in $\overline{\{(p, v(p, p)) : p \in \mathbf{X}\}}$, where (p^*, \dots, p^*) is not a Nash equilibrium (thus $p^* > \hat{p}_N$). Since v_N is continuous, $v^* = v_N(p^*)$. But then, any $p^* - \varepsilon$ secures strictly a payoff above v^* for $\varepsilon > 0$ small enough, because from Assumption 5 in the model and continuity of v_1 , $v(p^* - \varepsilon, p^*) = v_1(p^* - \varepsilon) > v_N(p^*) = v^*$ for $\varepsilon > 0$ small enough, and from the continuity of v_1 , this inequality is robust to a small modification of the other players' strategies.

Now, applying Theorem 23, for every $c \in \mathcal{F}$, there exists a symmetric prudent equilibrium (p, \dots, p) , which is also a symmetric Nash equilibrium (which implies $p \in [\hat{p}_N, \hat{p}_N]$).

(3) We want to prove that for every $c \in \mathcal{F}$, if (p, \dots, p) is a symmetric prudent equilibrium, then $p = \bar{p}$. Assume first that $p > \bar{p}$. By definition, there exists a sequence of positive reals $(\lambda_n)_{n \in \mathbf{N}}$ converging to 0, and a sequence of symmetric equilibria $(p^{\lambda_n}, \dots, p^{\lambda_n})$ of $G_{sym}^{\lambda_n}$, which converges to (p, \dots, p) . Thus

$$\forall d \in \mathbf{X}, v^{\lambda_n}(d, p^{\lambda_n}) \leq v^{\lambda_n}(p^{\lambda_n}, p^{\lambda_n}) \quad (12)$$

Recall that by definition,

$$v^{\lambda_n}(p^{\lambda_n}, p^{\lambda_n}) = \inf_{p' \in X} \left\{ v(p^{\lambda_n}, p') + \frac{c(p', p^{\lambda_n})}{\lambda_n} \right\}.$$

For every $p' < p^{\lambda_n}$, $v(p^{\lambda_n}, p') = 0$. Thus, taking $p' \rightarrow (p^{\lambda_n})^-$, we get, from the continuity of c and from equation 12:

$$\forall d \in \mathbf{X}, v^{\lambda_n}(d, p^{\lambda_n}) \leq v^{\lambda_n}(p^{\lambda_n}, p^{\lambda_n}) \leq 0. \quad (13)$$

Now, we claim that there exists n large enough and $\varepsilon > 0$ such that $p^{\lambda_n} - \varepsilon > \bar{p}$, and such that for every $p' \leq p^{\lambda_n} - \varepsilon$, we have $\frac{c(p^{\lambda_n}, p')}{\lambda_n} > v^{\lambda_n}(p^{\lambda_n} - \varepsilon, p^{\lambda_n}) + 1$. This is possible because otherwise, up to a subsequence, we would be able to build a sequence $p'_n \leq p^{\lambda_n} - \varepsilon$ such that $\frac{c(p^{\lambda_n}, p'_n)}{\lambda_n} \leq v^{\lambda_n}(p^{\lambda_n} - \varepsilon, p^{\lambda_n}) + 1$. But $v^{\lambda_n}(p^{\lambda_n} - \varepsilon, p^{\lambda_n}) + 1$ is bounded, and $\frac{c(p^{\lambda_n}, p'_n)}{\lambda_n}$ tends to $+\infty$, a contradiction.

By definition,

$$v^{\lambda_n}(p^{\lambda_n} - \varepsilon_n, p^{\lambda_n}) = \inf_{p' \in \mathbf{X}} \left\{ v(p^{\lambda_n} - \varepsilon_n, p') + \frac{c(p', p^{\lambda_n})}{\lambda_n} \right\}.$$

Now, if $p' \leq p^{\lambda_n} - \varepsilon$ then $v(p^{\lambda_n} - \varepsilon_n, p') + \frac{c(p', p^{\lambda_n})}{\lambda_n} > v^{\lambda_n}(p^{\lambda_n} - \varepsilon, p^{\lambda_n}) + 1$ for n large enough, from our choice above, and because $v(p^{\lambda_n} - \varepsilon_n, p') \geq 0$ from the assumptions of the model.

Thus, in particular

$$v^{\lambda_n}(p^{\lambda_n} - \varepsilon_n, p^{\lambda_n}) = \inf_{p' > p^{\lambda_n} - \varepsilon} \left\{ v_1(p^{\lambda_n} - \varepsilon) + \frac{c(p^{\lambda_n}, p')}{\lambda_n} \right\}$$

and finally, since c is a distance, this minimum is reached for $p' = p^{\lambda_n}$, that is

$$v^{\lambda_n}(p^{\lambda_n} - \varepsilon_n, p^{\lambda_n}) = v_1(p^{\lambda_n} - \varepsilon).$$

But from $p^{\lambda_n} - \varepsilon > \bar{p}$ we get $v_1(p^{\lambda_n} - \varepsilon) > 0$ (because from Assumption 4 in the model, $v_1(p) > 0$ when $p > \bar{p}$). This contradicts Equation 13, since this equation implies, for $d = p^{\lambda_n} - \varepsilon$, that $v_1(p^{\lambda_n} - \varepsilon) = v^{\lambda_n}(p^{\lambda_n} - \varepsilon_n, p^{\lambda_n}) \leq 0$.

Now, assume $p < \bar{p}$. By definition, there exists a sequence of symmetric equilibria $(p^{\lambda_n}, \dots, p^{\lambda_n})$ of $G_{sym}^{\lambda_n}$, which converges to (p, \dots, p) and which satisfies Equation 12. Taking $d = p^{max}$ in this Equation, we get

$$v^{\lambda_n}(p^{max}, p^{\lambda_n}) = 0 \leq v^{\lambda_n}(p^{\lambda_n}, p^{\lambda_n}). \quad (14)$$

But by definition

$$v^{\lambda_n}(p^{\lambda_n}, p^{\lambda_n}) = \inf_{p' \in X} \left\{ v(p^{\lambda_n}, p') + \frac{c(p', p^{\lambda_n})}{\lambda_n} \right\}.$$

Taking $p' \rightarrow (p^{\lambda_n})^+$, we get in particular

$$v^{\lambda_n}(p^{\lambda_n}, p^{\lambda_n}) \leq v_1(p^{\lambda_n}) < 0 \quad (15)$$

for n large enough (because $v_1(p) < 0$ when $p < \bar{p}$, from Assumption 4 in the model). This contradicts Equation 14.

References

- [1] Charalambos D. Aliprantis and Kim Border. Infinite Dimensional Analysis: A Hitchhiker's Guide. Springer, 1999.
- [2] Ola Andersson, Cédric Argenton, and Jorgen Weibull. Robustness to strategic uncertainty in price competition. SSE/EFI Working Paper Series in Economics and Finance 0726, Stockholm School of Economics, March 2010.

- [3] Ola Andersson, Cédric Argenton, and Jorgen Weibull. Robustness to strategic uncertainty. Games and Economic Behavior, 85:272–288, 2014.
- [4] Robert Aumann and Adam Brandenburger. Epistemic Conditions for Nash Equilibrium. Econometrica, 63(5):1161–80, 1995.
- [5] Michael R. Baye, Guoqiang Tian, and Jianxin Zhou. Characterizations of the Existence of Equilibria in Games with Discontinuous and Non-quasiconcave Payoffs. The Review of Economic Studies, 60(4):935–948, 1993.
- [6] Philippe Bich and Rida Laraki. On the Existence of Approximate Equilibria and Sharing Rule Solutions in Discontinuous Games. Theoretical Economics, to appear, 2016.
- [7] Adam Brandenburger. Strategic and Structural Uncertainty in Games, volume 94. Division of Research, Harvard Business School, 1993.
- [8] Oriol Carbonell-Nicolau. On the existence of pure-strategy perfect equilibrium in discontinuous games. Games and Economic Behavior, 71(1):23–48, January 2011.
- [9] Oriol Carbonell-Nicolau. The Existence of Perfect Equilibrium in Discontinuous Games. Games, 2(3):235–256, 2011.
- [10] Guilherme Carmona. Understanding some Recent Existence Results for Discontinuous Games. Economic Theory, 48(1):31–45, 2011.
- [11] Partha Dasgupta and Eric Maskin. The Existence of Equilibrium in Discontinuous Economic Games, Part I (Theory). Review of Economic Studies, 53:1–26, 1986.
- [12] Krishnendu Dastidar. On the Existence of Pure Strategy Bertrand Equilibrium. Economic Theory, 5(1):19–32, 1995.
- [13] Giuseppe De Marco and Maria Romaniello. Games Equilibria and the Variational Representation of Preferences. Csef working papers, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy, 2013.
- [14] James Dow and Sergio Werlang. Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction. Journal of Economic Theory, 64(2):305–324, 1994.
- [15] Jorgen Eichberger and David Kelsey. Non-Additive Beliefs and Strategic Equilibria. Games and Economic Behavior, 30(2):183 – 215, 2000.
- [16] John Harsanyi and Reinhard Selten. A General Theory of Equilibrium Selection in Games, volume 1. The MIT Press, 1988.
- [17] Peter Klibanoff. Uncertainty, decision, and normal form games. Unpublished manuscript, Kellogg Graduate School of Management, Northwestern University, 1996.
- [18] Elon Kohlberg and Jean-Francois Mertens. On the Strategic Stability of Equilibria. Econometrica, 54(5):1003–37, September 1986.

- [19] Chung Kin Lo. Nash equilibrium without mutual knowledge of rationality. Economic Theory, 14(3):621–633, 1998.
- [20] Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. Econometrica, 74(6):1447–1498, 2006.
- [21] Massimo Marinacci. Ambiguous Games. Games and Economic Behavior, 31(2):191–219, May 2000.
- [22] Richard D. McKelvey and Thomas R. Palfrey. Quantal Response Equilibria for Normal Form Games. Games and Economic Behavior, 10(1):6 – 38, 1995.
- [23] Sujoy Mukerji. A theory of play for games in strategic form when rationality is not common knowledge. Discussion paper series in economics and econometrics, University of Southampton, Economics Division, School of Social Sciences, 1995.
- [24] Roger B. Myerson. Refinements of the Nash equilibrium concept. International Journal of Game Theory, 7(2):73–80, 1978.
- [25] Ludovic Renou and Karl H. Schlag. Minimax regret and strategic uncertainty. Journal of Economic Theory, 145(1):264 – 286, 2010.
- [26] Philip J. Reny. On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games. Econometrica, 67(5):1029–1056, 1999.
- [27] Philip J. Reny. Strategic Approximations of Discontinuous Games. Economic Theory, 48(1):17–29, 2011.
- [28] Matthew J. Ryan. What do uncertainty-averse decision-makers believe ? Economic Theory, 20(1):47–65, 2002.
- [29] Reinhard Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory, 4(1):25–55.
- [30] Leo Simon and Maxwell Stinchcombe. Equilibrium Refinement for Infinite Normal-Form Games. Econometrica, 63(6):1421–43, 1995.
- [31] Wen-Tsun Wu and Jia-He Jiang. Essential equilibrium points of n-person non-cooperative games. Scientia Sinica, 11:1307–1322, 1962.