

IFN Working Paper No. 910, 2012

Robustness to Strategic Uncertainty

Ola Andersson, Cédric Argenton and Jörgen W.
Weibull

Robustness to strategic uncertainty

OLA ANDERSSON*, CÉDRIC ARGENTON[†] AND JÖRGEN W. WEIBULL^{‡§}

March 1, 2012

ABSTRACT. In games with continuum strategy sets, we model a player's uncertainty about another player's strategy, as an atomless probability distribution over the other player's strategy set. We call a strategy profile (strictly) robust to strategic uncertainty if it is the limit, as uncertainty vanishes, of some sequence (all sequences) of strategy profiles in which every player's strategy is optimal under his or her uncertainty about the others. General properties of this robustness criterion are derived and it is shown that it is a refinement of Nash equilibrium when payoff functions are continuous. We apply the criterion to a class of Bertrand competition games. These are discontinuous games that admit a continuum of Nash equilibria. Our robustness criterion selects a unique Nash equilibrium, and this selection agrees with recent experimental findings.

Keywords: Nash equilibrium, refinement, strategic uncertainty, Bertrand competition, log-concavity.

JEL-codes: C72, D43, L13

*Research Institute of Industrial Economics (IFN)

[†]CentER & TILEC, Tilburg University.

[‡]Stockholm School of Economics, Ecole Polytechnique (Paris) and Royal Institute of Technology (Stockholm).

[§]The authors are grateful for comments from Jean-Pascal Bénassy, Axel Bernergård, Larry Blume, Hans Carlsson, Eric van Damme, David Easley, Drew Fudenberg, Sergiu Hart, David Leslie, Tom Palfrey and Philip Reny, as well as from participants in many seminar presentations. Ola Andersson gratefully acknowledges financial support from the Jan Wallander and Tom Hedelius foundation.

1. INTRODUCTION

In recent experimental studies, uncertainty about other players' actions has been found to be a clear driver of behavior, see e.g. Heinemann et al. (2009) and Cabrales et al. (2010). Following Harsanyi and Selten (1988) and Brandenburger (1996), one usually alludes to this type of uncertainty as 'strategic uncertainty', as opposed to uncertainty regarding the underlying structure of the game played, which is sometimes called 'structural uncertainty' (see e.g. Morris and Shin, 2002). Strategic uncertainty matters because in a wide range of games, many of the equilibria represent fragile situations in which players are supposed to choose a particular strategy, even though this would be optimal only if they held knife-edge beliefs about the actions taken by other players. In such situations, even the slightest uncertainty about other players' choices might lead a player to deviate from his or her equilibrium strategy. This uncertainty problem is aggravated in games with multiple Nash equilibria, and it may get particularly serious when there is a whole continuum of equilibria.

That said, in the laboratory, human subjects' behavior in games with multiple equilibria has also been found to be fairly stable and predictable in the aggregate. For instance, Abbink and Brandts (2008) produce an experimental study of Bertrand competition under strictly convex costs. Dastidar (1995) had shown that those oligopoly games admit a whole continuum of Nash equilibria, but they find that an attractor of play is the zero-monopoly-profit price. In experimental treatments with more than two firms in the market, that price is actually the modal outcome in their data. Abbink and Brandts (2008) remark that "[that] price level (...) is not predicted by any benchmark theory [they] are aware of" (p. 3).¹ We conjecture that part of this regularity may be that some equilibria are perceived as less strategically risky than others. In this paper, we introduce, and study in some generality, a robustness criterion for games with continuum strategy sets. We then proceed to show that our robustness criterion selects a unique equilibrium in the game of Bertrand competition with convex costs, and that this selection agrees with Abbink's and Brandts' (2008) empirical findings.

To be more specific about our contribution, we here formalize a notion of strategic uncertainty and propose a criterion for robustness to such uncertainty. We focus on games with continuum action spaces, with measurable and bounded, but not necessarily continuous payoff functions. Our approach is, roughly, as follows. A player's uncertainty about others' strategy choices is represented by a player-specific, atomless probability distribution, scaled with a parameter $t \geq 0$, over others' strategy sets. We do not exclude the possibility that these beliefs may be biased or that two

¹Argenton and Müller (2009) corroborate this experimental finding in a laboratory experiment with other subjects and parameter values. Van Huyck et al. (1990, 1991) (in the case of minimum effort games) and Heinemann et al. (2009) (in the case of coordination games) also identify rather predictable patterns of play in spite of the multiplicity of equilibria.

players have inter-personally inconsistent beliefs about a third one. For each value of the uncertainty parameter t , we define a t -*equilibrium* as a Nash equilibrium of the game in which each player strives to maximize her expected payoff under her strategic uncertainty so defined. For $t = 0$, this is nothing else than Nash equilibrium in the original game. We call a strategy profile *robust to strategic uncertainty* if there exists a collection of probability distributions in the admitted class, one for each player, such that some accompanying sequence of t -equilibria converges to this profile as the uncertainty parameter t tends to zero. If convergence holds for all distributions in the admitted class, we say that the strategy profile is *strictly robust* to strategic uncertainty (in this class).

We apply this definition to Bertrand competition and show that the unique, corresponding prediction in the case of strictly convex costs agrees with the findings in Abbink and Brandts (2008). Heuristically, strategic uncertainty in such discontinuous games results in uncertainty-perturbed profit functions that are continuous. The deviation incentives in some Nash equilibria may be quite asymmetric, though. At high Nash equilibrium prices, a strategically uncertain player has an incentive to slightly undercut, since she has much more to lose if others cut their prices than if they raised their prices. Conversely, for low Nash equilibrium prices, an uncertain player has an incentive to raise her price slightly, since she has a lot to lose if others raised their prices and little to lose if they cut their prices. The only Nash equilibrium price that is robust to strategic uncertainty is the price at which a monopolist would earn zero profit. In fact, that price is strictly robust to strategic uncertainty.

The proposed framework is well-suited to study games with discontinuous payoff functions, which are not uncommon in economics. However, our approach applies to a wide class of games. For games with continuous payoff functions, we show that our criterion is a refinement of Nash equilibrium, and we show that under standard (compactness and convexity) assumptions, robust equilibria exist. We also show that our definition coincides with that of Nash equilibrium in very well-behaved games, more specifically, in all games with continuously differentiable payoff functions on compact and convex strategy spaces, where each payoff function is strictly concave in the player's own strategy. We also show that, for two-player games, it implies weak perfection in the sense of Simon and Stinchcombe (1995).

The remainder of the paper is organized as follows. Section 2 provides definitions and section 3 applies them to Bertrand competition. In section 4 general results are presented for games with continuous payoff functions and in section 5 our notion of robustness is extended to higher dimensional strategy sets. Section 6 discusses related literature and section 7 concludes. Mathematical proofs are given in an appendix.

2. ROBUSTNESS TO STRATEGIC UNCERTAINTY

In this section, we introduce our definition of robustness to strategic uncertainty. Let $G = (N, S, \pi)$ be an n -player normal-form game in which the strategy set of each player $i \in N = \{1, \dots, n\}$ is $S_i = \mathbb{R}$. Thus, $S = \mathbb{R}^n$ is the set of strategy profiles, $s = (s_1, \dots, s_n)$, and $\pi : S \rightarrow \mathbb{R}^n$ is the combined payoff function, with $\pi_i(s)$ being the payoff to player i when s is played.² We assume each payoff function π_i to be Borel measurable and bounded but do not require continuity. Let \mathcal{F} be the set of continuous and everywhere positive probability density functions on \mathbb{R} . For $\phi \in \mathcal{F}$ and any $x \in \mathbb{R}$, let $\Phi(x) = \int_{-\infty}^x \phi(x) dx$. This defines the associated cumulative probability distribution on \mathbb{R} with $\Phi' = \phi$.

Definition 1. For any given $t \in [0, 1]$, a strategy profile s is a **t -equilibrium** of G if, for each player i , the strategy s_i maximizes i 's expected payoff under the probabilistic belief that all other players' strategies are random variables of the form

$$\tilde{s}_{ij} = s_j + t \cdot \varepsilon_{ij} \quad (1)$$

for some statistically independent “noise” terms ε_{ij} with densities $\phi_{ij} \in \mathcal{F}$ for all $j \neq i$.

Remark 1. For $t = 0$, this definition coincides with that of Nash equilibrium.

Remark 2. For $t > 0$, each random variable \tilde{s}_{ij} has a probability density $f_{ij}^t \in \mathcal{F}$ defined by

$$f_{ij}^t(x) = \frac{1}{t} \phi_{ij} \left(\frac{x - s_j}{t} \right) \quad \forall x \in \mathbb{R}$$

Note that we do not require that noise terms be symmetric or have expectation zero. Hence, in a t -equilibrium it may well be that some players believe that others tend more to deviate upwards than downwards.

Let \tilde{s}_{-i} denote the $(n - 1)$ -vector of random variables $(\tilde{s}_{ij})_{j \neq i}$. We note that any t -equilibrium is a Nash equilibrium of a game with perturbed payoff functions:

Remark 3. Let $t > 0$ and $\phi_{ij} \in \mathcal{F}$ for all $i \in N$ and $j \neq i$. A strategy profile $s \in S$ is a t -equilibrium of $G = (N, S, \pi)$, with $\varepsilon_{ij} \sim \phi_{ij}$, if and only if it is a Nash equilibrium of the perturbed game $G^t = (N, S, \pi^t)$, where

$$\begin{aligned} \pi_i^t(s) &= \mathbb{E}[\pi_i(s_i, \tilde{s}_{-i})] = \int \dots \int \left(\prod_{j \neq i} f_{ij}^t(x_j) \right) \pi_i(s_i, x_{-i}) dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_n \\ &= \frac{1}{t^{n-1}} \int \dots \int \left(\prod_{j \neq i} \phi_{ij} \left(\frac{x_j - s_j}{t} \right) \right) \pi_i(s_i, x_{-i}) dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_n \end{aligned}$$

²See below for how this machinery can be adapted to strategy sets with upper and/or lower bounds, and see Section 5 for a generalization to multi-dimensional strategy sets.

We are now in a position to define robustness to strategic uncertainty.

Definition 2. A strategy profile s^* in G is **robust to strategic uncertainty** if there exists a collection of probability density functions $\{\phi_{ij} \in \mathcal{F} : i \in N, j \neq i\}$ and an accompanying sequence of t -equilibria, $\langle s^{t_k} \rangle_{k \in \mathbb{N}}$ with $t_k \downarrow 0$, such that $s^{t_k} \rightarrow s^*$ as $k \rightarrow +\infty$. The strategy profile s^* is **strictly robust to strategic uncertainty** if this holds for all collections of probability density functions $\{\phi_{ij} \in \mathcal{F} : i \in N, j \neq i\}$.

Remark 4. This definition can be adapted as follows to games in which the strategy set of a player j is an interval S_j with lower and upper bounds $a_j < b_j$ (these may, but need not, belong to the interval). For any $\phi_{ij} \in \mathcal{F}$, let

$$f_{ij}^t(x) = \begin{cases} \frac{1}{t} \cdot \frac{\phi_{ij}\left(\frac{x-s_j}{t}\right)}{\Phi_{ij}\left(\frac{b_j-s_j}{t}\right) - \Phi_{ij}\left(\frac{a_j-s_j}{t}\right)} & \forall x \in S_j \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

This defines a probability density for \tilde{s}_{ij} with support S_j . Taking expectations with respect to such probability densities f_{ij}^t , one obtains a perturbed game with payoff functions

$$\begin{aligned} \pi_i^t(s) &= \mathbb{E}[\pi_i(s_i, \tilde{s}_{-i})] \\ &= \frac{1}{Y_i(t, s_{-i}) \cdot t^{n-1}} \int \dots \int \left(\prod_{j \neq i} \phi_{ij} \left(\frac{x_j - s_j}{t} \right) \right) \pi_i(s_i, x_{-i}) dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_n \end{aligned}$$

where

$$Y_i(t, s_{-i}) = \prod_{j \neq i} \left[\Phi_{ij} \left(\frac{b_j - s_j}{t} \right) - \Phi_{ij} \left(\frac{a_j - s_j}{t} \right) \right]. \quad (3)$$

We note that $Y_i(t, s_{-i}) > 0$ for all $t > 0$ and that $Y_i(t, s_{-i}) \rightarrow 1$ when $t \rightarrow 0$ and s_{-i} is interior. For $s_j = a_j$: $\Phi_{ij} \left(\frac{b_j - s_j}{t} \right) - \Phi_{ij} \left(\frac{a_j - s_j}{t} \right) \rightarrow 1 - \Phi_{ij}(0) > 0$. For $s_j = b_j$: $\Phi_{ij} \left(\frac{b_j - s_j}{t} \right) - \Phi_{ij} \left(\frac{a_j - s_j}{t} \right) \rightarrow \Phi_{ij}(0) > 0$. Hence, for all i and j , $Y_i(t, s_{-i})$ tends to a constant $\gamma_{ij} > 0$ as $t \rightarrow 0$.

Some comments are in place. First, we do not require that all players $i \neq j$ hold the same belief about a particular player j . Second, in our definition of t -equilibrium we assume statistical independence of noise terms. This is done for convenience only. Statistical dependence among the noise terms would not change our results as long as the joint distribution across other players' strategy sets has a density. (In two-player games, the issue is of course moot.) Third, as seen in (1), we assume that players' probabilistic beliefs have an additive structure. Alternatively, one could work with more general random variables representing players' subjective beliefs. However, the present formulation turns out to be rich enough to generate interesting results, and it easy to apply, as we now proceed to show.

3. BERTRAND COMPETITION

We now turn to an application of our definition of robustness to strategic uncertainty to Bertrand competition. Following Vives (1999, p.117), we take Bertrand competition to mean that (a) sellers simultaneously post their prices and (b) each firm is committed to serve all its clients at its posted price. In particular, it cannot ration its demand. As mentioned by Vives (1999), for certain utilities and auctions, provision is legally mandated and in other markets firms have a strong incentive to serve all their clients, especially in industries in which customers have an on-going relationship with suppliers (subscription, repeat purchases, etc.) or where the costs of restricting output in real time are high.³

Our definition of robustness selects a unique Nash equilibrium out of a continuum of pure-strategy equilibria in a class of Bertrand competition games with convex costs. However, before embarking on that analysis, we briefly consider the canonical Bertrand competition game with linear costs.

Example 1. Consider two identical firms, each with constant unit cost $c > 0$, in a simultaneous-move pricing game à la Bertrand in a market for a homogeneous good. Let the demand function be linear, $D(p) = a - p$, for all $p \in [0, a]$ with $a > c$.⁴ Then, the monopoly profit function, $\Pi(p) = (a - p)(p - c)$, is strictly concave with a unique maximum at $p^m = (a + c)/2 < a$ and $\Pi(p^m) > 0$. By contrast, the unique duopoly Nash equilibrium, $p_1 = p_2 = c$, results in zero profits. This Nash equilibrium is weakly dominated. Nevertheless, it is robust to strategic uncertainty. For sufficiently small degrees of strategic uncertainty, both firms will set their prices a little bit above marginal cost, and less so, the less uncertain they are. To see this, suppose that $\varepsilon_{ij} \sim \phi \in \mathcal{F}$.⁵ For each $t > 0$ and all $p_1, p_2 \in [0, a]$, and with Φ denoting the c.d.f. induced by ϕ :

$$\pi_i^t(p_i, p_j) = \left[1 - \frac{\Phi\left(\frac{p_i - p_j}{t}\right) - \Phi\left(-\frac{p_j}{t}\right)}{\Phi\left(\frac{a - p_j}{t}\right) - \Phi\left(-\frac{p_j}{t}\right)} \right] \cdot \Pi(p_i) \quad i = 1, 2, j \neq i.$$

This profit can be rewritten as

$$\pi_i^t(p_i, p_j) = \left[\Phi\left(\frac{a - p_j}{t}\right) - \Phi\left(-\frac{p_j}{t}\right) \right]^{-1} \cdot \left[\Phi\left(\frac{a - p_j}{t}\right) - \Phi\left(\frac{p_i - p_j}{t}\right) \right] \cdot \Pi(p_i)$$

³Dixon (1990) shows that modelling price competition under convex production costs and explicit penalties for turning consumers away delivers a whole continuum of Nash equilibria.

⁴To keep the intuition clear, we take a simple functional form but the argument extends to general demand curves.

⁵We focus on symmetric error distributions in this example only for expositional convenience. The Nash equilibrium is robust to strategic uncertainty under asymmetric distributions as well.

where the first factor is positive and independent of p_i . A necessary first-order condition for symmetric t -equilibrium⁶ is thus that

$$t \cdot \frac{\Pi'(p_i)}{\Pi(p_i)} = \frac{\phi(0)}{\Phi[(a - p_j)/t] - \Phi(0)} \quad i = 1, 2, j \neq i. \quad (4)$$

Consequently, in the perturbed game, it is never optimal to choose $p_i \leq c$ or $p_i \geq p^m$. On the interval (c, p^m) , the left-hand side is a continuous and strictly decreasing function of p_i , that runs from plus infinity to zero, while the right-hand side is a constant. Hence, there exists a unique symmetric t -equilibrium price, p^t , for every $t > 0$. As $t \rightarrow 0$, the denominator of the right-hand side tends to $\Phi(1) - \Phi(0) > 0$. Thus, the left-hand side has to tend to zero for (4) to hold. Consequently, $p^t \downarrow c$. This observation can be generalized from duopoly to oligopoly, with an arbitrary number $n > 1$ of firms. It is easily verified that the necessary first-order condition (4) then will have an extra factor $(n - 1)$ on its right-hand side. This does not affect the limit result $p^t \downarrow c$ but for positive t , the t -equilibrium price p^t will be lower the more firms there are in the industry.

Now, consider $n \geq 2$ identical firms $i \in N = \{1, 2, \dots, n\}$ in a market for a homogeneous good. Aggregate demand $D : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is twice differentiable and such that $D(0) = q^{\max} \in \mathbb{R}$ and $D(p^{\max}) = 0$ for some $p^{\max}, q^{\max} > 0$. Moreover, we assume that $D'(p) < 0$ for all $p \in (0, p^{\max})$. All firms i simultaneously set their prices $p_i \in \mathbb{R}_+$. Let $\mathbf{p} = (p_1, p_2, \dots, p_n)$ be the resulting strategy profile (or *price vector*). The minimal price, $p_0 := \min \{p_1, p_2, \dots, p_n\}$, will be called *the (going) market price*. Let m be the number of firms that quote the going market price, $m := |\{i : p_i = p_0\}|$. Each firm i faces the demand

$$D_i(\mathbf{p}) := \begin{cases} D(p_0)/m & \text{if } p_i = p_0 \\ 0 & \text{otherwise} \end{cases}$$

All firms have the same strictly convex cost function, $C : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which is twice differentiable with $C(0) = 0$ and $C', C'' > 0$. Each firm is required to serve all demand addressed to it at its posted price. The profit to each firm i is thus

$$\pi_i(\mathbf{p}) = \begin{cases} p_0 D(p_0)/m - C[D(p_0)/m] & \text{if } p_i = p_0 \\ 0 & \text{otherwise} \end{cases}. \quad (5)$$

This defines a simultaneous-move n -player game G in which each player i has pure-strategy set \mathbb{R}_+ and payoff function $\pi_i : \mathbb{R}_+^n \rightarrow \mathbb{R}$, defined in equation (5). A strategy profile \mathbf{p} will be called *symmetric* if $p_1 = \dots = p_n$, and we will call a price

⁶It is easily verified that there does not exist any asymmetric t -equilibrium.

$p \in \mathbb{R}_+$ a *symmetric Nash equilibrium price* if $\mathbf{p} = (p, p, \dots, p)$ is a Nash equilibrium of G . For each positive integer $m \leq n$ and non-negative price p , let

$$v_m(p) = pD(p)/m - C[D(p)/m]$$

This defines a finite collection of twice differentiable functions, $\langle v_m \rangle_{m \in \{1, 2, \dots, n\}}$, where $v_m(p)$ is the profit to each of m firms if they all quote the same price p and all other firms post higher prices (so that p is the going market price). In particular, v_1 defines the profit to a monopolist as a function of its price p .

We impose one more condition on C and D , namely, that the associated monopoly profit function, v_1 , is concave. More exactly, we assume that $v_1'' \leq 0$ and $v_1'(p^{mon}) = 0$ for some price $p^{mon} \in (0, p^{\max})$. Since the cost function is strictly convex by assumption, this concavity assumption on v_1 effectively requires the demand function to be “not too convex”. We have $v_1(p^{mon}) \geq 0$. By convexity of the cost function, there exists prices $p \in (0, p^{\max})$ at which all n firms, when quoting the same price p , make positive profits, and $v_n(p) > 0$.

The game G has a continuum of symmetric Nash equilibria.⁷ For any number of firms, $n \geq 2$, let $\check{p}_n \in (0, p^{\max})$ be the price p at which $v_n(p) = 0$ and let $\hat{p}_n \in (0, p^{\max})$ be the price p at which $v_n(p) = v_1(p)$. Dastidar (1995, Lemmas 1, 5 and 6) shows existence and uniqueness of \check{p}_n and \hat{p}_n , and that $\check{p}_n < \hat{p}_n$.⁸ All prices in the interval $P_n^{NE} = [\check{p}_n, \hat{p}_n]$ are symmetric Nash equilibrium prices in the game G , and no price outside this interval is a symmetric Nash equilibrium price (Dastidar, 1995, Proposition 1). There exists a unique price \bar{p} at which a monopolist makes zero profit, $v_1(\bar{p}) = 0$, and, moreover, $\bar{p} \in (\check{p}_n, \hat{p}_n)$ (Dastidar, 1995, Lemmas 4 and 6). Both \check{p}_n and \hat{p}_n are strictly decreasing in n (Dastidar, 1995, Lemma 7). In the present setting, it is easily verified that $\check{p}_n \downarrow 0$ and $\hat{p}_n \downarrow \bar{p}$, and hence $P_n^{NE} \rightarrow (0, \bar{p}]$, as $n \rightarrow \infty$.

Example 2. Consider a duopoly with identical firms with cost function $C(q) = cq^2$, for $c = 1/5$, facing aggregate demand $D(p) = \max\{0, 1 - p\}$. The diagram below shows the profit functions v_1 (dashed) and v_2 (solid).

⁷Dastidar (1995) and Weibull (2006) have shown existence and multiplicity of Nash equilibria under weaker conditions. In Andersson et al. (2010), we show that all equilibria are symmetric.

⁸In Andersson et al. (2010) we also show that $\hat{p}_n \leq p^{mon}$ for all $n \geq 2$.

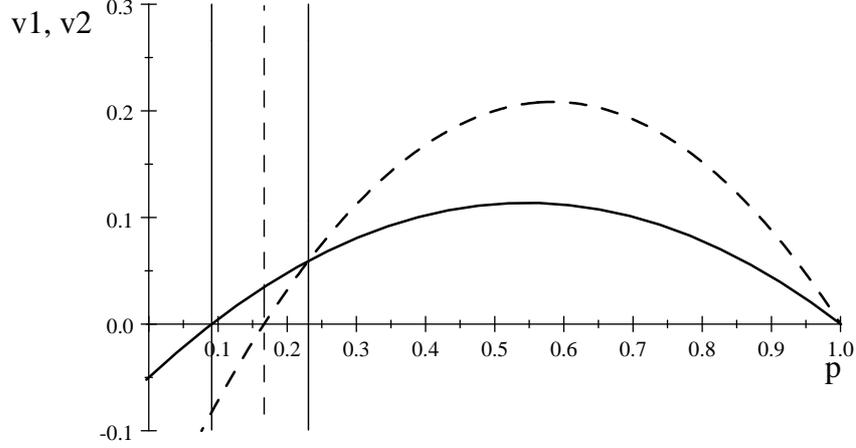


Figure 2: The monopoly (dashed) and duopoly (solid) profit functions, $v_1(p)$ and $v_2(p)$.

The set P_2^{NE} of duopoly Nash equilibrium prices is the interval $[1/11, 3/13]$, indicated by two vertical lines, and the zero monopoly-profit price $\bar{p} = 1/6$ is indicated by the dashed vertical line.

We proceed to apply our robustness definition to this class of discontinuous games. Let $t > 0$ and suppose that a firm i holds a probabilistic belief of form (1) about other firms' prices, with Φ_{ij} denoting the c.d.f. induced by each probability density $\phi_{ij} \in \mathcal{F}$. For any price p_i that firm i might contemplate to set, its subjective probability that any other firm will choose exactly the same price is zero. Hence, with probability one, its own price will either lie above the going market price or it will be the going market price and all other firms' prices will be higher, so i will then be a monopolist at its price p_i . Each firm i 's payoff function in the perturbed game $G^t = (N, S, \pi^t)$ is, for any $t > 0$, defined by

$$\pi_i^t(\mathbf{p}) = v_1(p_i) \cdot \left(\prod_{j \neq i} \left[1 - \Phi_{ij} \left(\frac{-p_j}{t} \right) \right]^{-1} \right) \cdot \left(\prod_{j \neq i} \left[1 - \Phi_{ij} \left(\frac{p_i - p_j}{t} \right) \right] \right) \quad (6)$$

The second factor is due to the restriction of the support of subjective beliefs about others' prices to \mathbb{R}_+ . This factor being positive and independent of p_i , a price profile \mathbf{p} is a Nash equilibrium of G^t if and only if

$$p_i \in \arg \max_{p \in [\bar{p}, p^{mon}]} u_i^t(p, \mathbf{p}_{-i}) \quad \forall i, \quad (7)$$

where

$$u_i^t(\mathbf{p}) = v_1(p_i) \cdot \prod_{j \neq i} \left[1 - \Phi_{ij} \left(\frac{p_i - p_j}{t} \right) \right]$$

and the restriction $p \in [\bar{p}, p^{mon}]$ is non-binding, since $v_1(p) < 0$ for all $p < \bar{p}$, $v_1(p) > 0$ for all $p \in (\bar{p}, p^{mon})$, and $v_1'(p) < 0$ for all $p > p^{mon}$. For any $t > 0$, let \bar{G}^t be the normal-form game $(N, [\bar{p}, p^{mon}]^n, u^t)$. For any $t > 0$, a price profile p is a t -equilibrium in the pricing game G if and only if it is a Nash equilibrium of the game \bar{G}^t .

Let $\mathcal{F}^* \subset \mathcal{F}$ be the class of log-concave probability distributions with finite mean. More exactly, by *log-concavity* we mean continuously differentiable probability density functions ϕ such that $\ln \phi$ is a concave function. The log-concavity assumption is common in the economics literature and has applications in mechanism design, game theory and labor economics, see Bagnoli and Bergstrom (2005). A useful feature of those distributions is that they have non-decreasing hazard rates; that is, the hazard rate

$$h(x) = \frac{\phi(x)}{1 - \Phi(x)}$$

is non-decreasing (see Corollary 2 in Bagnoli and Bergstrom, 2005). Examples of log-concave distributions are the normal, exponential and Gumbel distributions.

In the subsequent analysis we will require all density functions ϕ to belong to the subset \mathcal{F}^* . Strategy profiles that are robust to uncertainty represented by such density functions will be said to be robust to strategic uncertainty *in* \mathcal{F}^* . The next proposition uses the log-concavity of the density functions in \mathcal{F}^* to show that each player's best-reply correspondence is convex-valued. We then use Kakutani's fixed point theorem to show existence of a t -equilibrium.

Proposition 1. *Let $t > 0$ and assume that $\phi_{ij} \in \mathcal{F}^* \forall i \in N, j \neq i$. Then \bar{G}^t has at least one Nash equilibrium. Moreover, any such Nash equilibrium \mathbf{p}^t is interior.*

The following proposition uses the monotonicity of the hazard rates, which follows from our log-concavity assumption, to show symmetry of any robust equilibrium. It then follows that a t -equilibrium price, which has to lie above \bar{p} , has to lie close to \bar{p} .

Proposition 2. *The Nash equilibrium $(\bar{p}, \dots, \bar{p})$ is strictly robust to strategic uncertainty in \mathcal{F}^* . No other strategy profile of G is robust to strategic uncertainty in \mathcal{F}^* .*

Example 3. *Consider again the duopoly in Example 2. We then have $\bar{p} = 1/6 \approx 0.167$. Suppose that both firms' uncertainty takes the form of normally distributed*

noise, $\varepsilon_1, \varepsilon_2 \sim N(0, 1)$. The necessary first-order condition for interior t -equilibrium then consists of the equations

$$tv'_1(p_1) = v_1(p_1)h\left(\frac{p_1 - p_2}{t}\right)$$

and

$$tv'_1(p_2) = v_1(p_2)h\left(\frac{p_2 - p_1}{t}\right).$$

The diagram below shows these best-reply curves (dashed for player 2), for $t = 0.1$, with \bar{p} marked by thin straight lines.

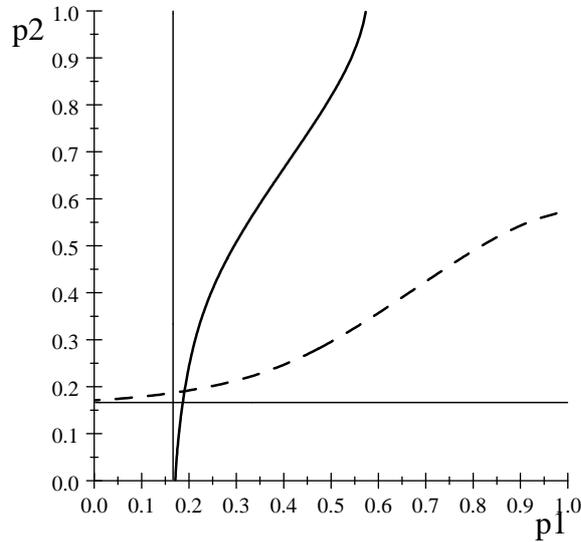


Figure 3: The best-reply curves in the perturbed duopoly pricing game.

3.1. Experimental evidence. We end this section with a short account of the experimental evidence for this class of games. In Abbink's and Brandts' (2008) experiment, the price interval was $[1, 40]$ and they conducted sessions with $n = 2, 3$ and 4 firms in a finitely repeated market interaction with fixed matching, see diagram below (their Figure 3). The associated ranges of Nash equilibrium prices were $P_2^{NE} = [12, 30]$, $P_3^{NE} = [7, 28]$ and $P_4^{NE} = [3, 27]$, respectively, with $\bar{p} = 24$ in all three cases. The diagrams show a strong tendency to this latter price.⁹ Argenton and

⁹For $n = 2$, the predominant outcome, 33, was the collusive price that maximized joint profits in that setting.

Müller (2009) replicated the experiment in Abbink and Brandts for $n = 2$, but with other cost and demand parameters. In their setting, $\bar{p} = 32$, and they also found that it was an attractor of play.

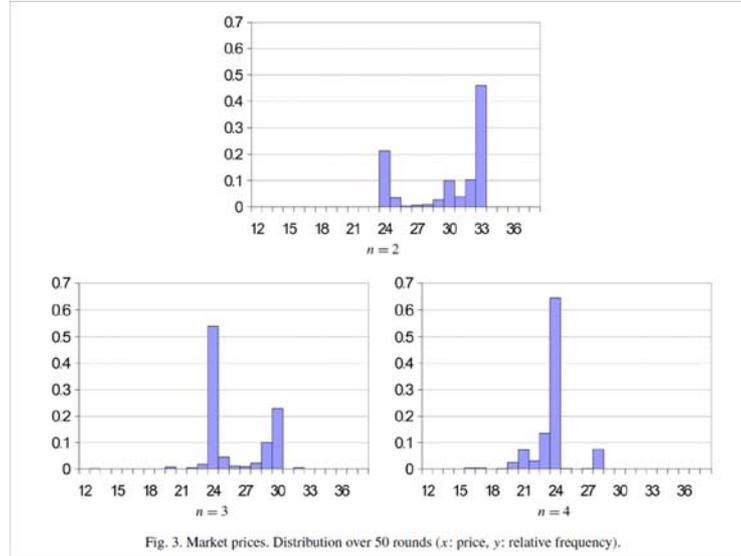


Figure 4. Empirical frequencies of prices in Abbink and Brandts (2008).

4. CONTINUOUS GAMES

A clear feature of Bertrand games with homogenous products is the discontinuity of payoffs. We now turn attention to games with continuous payoff functions to examine when our robustness criterion has some bite and how it relates to existing solution concepts. We first show that in such games, robustness to strategic uncertainty is a refinement of Nash equilibrium. Heuristically, if a strategy profile is not a Nash equilibrium in the original game, then some player has a strictly better reply to it. Such a player will then also have a strictly better reply in any slightly perturbed game, since the perturbed payoff function will be close to the original one. Formally:

Proposition 3. *If the payoff functions in G are continuous and s^* is robust to strategic uncertainty, then s^* is a Nash equilibrium.*

It is not difficult to verify existence of strategy profiles that are robust to strategic uncertainty in games with continuous payoff functions, if each strategy set is compact and convex and each payoff function π_i is concave in s_i (for every s_{-i}):

Proposition 4. *Suppose that the payoff functions in G are continuous, that the strategy sets are compact and convex, and that each payoff function π_i is concave in s_i (for every s_{-i}). Let $\phi_{ij} \in \mathcal{F}, \forall i \in N, j \neq i$. For each $t > 0$, the perturbed game G^t has at least one Nash equilibrium, and G has at least one strategy profile that is robust to strategic uncertainty.*

It follows from Proposition 3 and the proof of Proposition 4 that if a game with continuous and concave payoff functions over compact and convex strategy sets has a unique Nash equilibrium, then this equilibrium is not only robust to strategic uncertainty, but is in fact strictly robust:

Corollary 1. *Suppose that the payoff functions in G are continuous, that the strategy sets are compact and convex, and that each payoff function π_i is concave in s_i (for every s_{-i}). If s^* is the unique Nash equilibrium of G , then s^* is strictly robust to strategic uncertainty.*

The following example shows that even games that satisfy the hypotheses of Proposition 4 may have strict equilibria that are not robust to strategic uncertainty.

Example 4. *Consider the symmetric two-player game with $S_1 = S_2 = [0, 1]$ and payoff functions*

$$\pi_1(s_1, s_2) = \begin{cases} s_2 - (s_1 - s_2)^2 & \text{if } s_1 \leq s_2 \\ 2s_2 - s_1 & \text{if } s_1 > s_2 \end{cases}$$

and $\pi_2(s_1, s_2) \equiv \pi_1(s_2, s_1)$. The strategy sets are compact and convex and each player's payoff function is continuous and concave in the player's own strategy. The figure below shows the graph of player 1's payoff function when $s_2 = 1/2$.

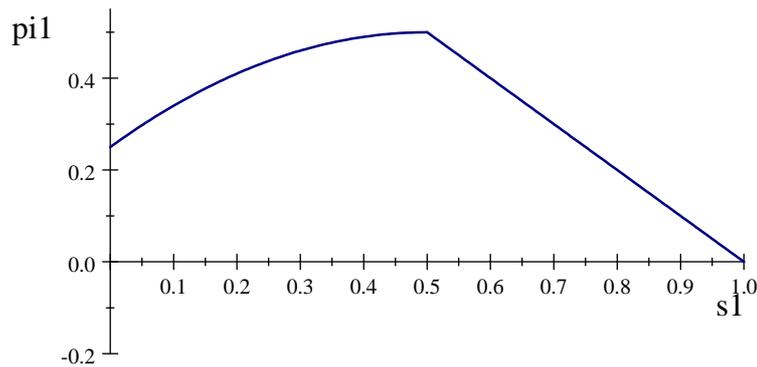


Figure 1: Player 1's payoff function in Example 1.

Any symmetric strategy profile (s, s) is a strict equilibrium. Yet, only $(0, 0)$ is robust

to strategic uncertainty. Indeed, for any density $\phi_{12} \in \mathcal{F}$ and strategy $s_2 \in (0, 1]$, the derivative of π_1^t with respect to s_1 equals

$$\frac{1}{Y_1(t, s_2)} \cdot \left[\int_{(s_1-s_2)/t}^{(1-s_2)/t} 2(s_2 + tu - s_1) \phi_{12}(u) du - \int_{-s_2/t}^{(s_1-s_2)/t} \phi_{12}(u) du \right], \quad (8)$$

where $Y_1(t, s_2) = \Phi_{12}((1-s_2)/t) - \Phi_{12}(-s_2/t) > 0$ and $\rightarrow 1$ as $t \rightarrow 0$. For $s_1 \geq s_2 > 0$ and for all $t > 0$ sufficiently small, the expression in (8) is negative. Hence, for small $t > 0$, any best reply s_1 for player 1 lies below s_2 whenever $s_2 > 0$. For small t , only $(0, 0)$ is a t -equilibrium. As the density ϕ_{12} was arbitrary, only $(0, 0)$ is robust to strategic uncertainty. (Indeed, it is strictly robust.)

Finally, we note that robustness to strategic uncertainty does not exclude all weakly dominated strategies in all continuous games. Indeed, as noted by Simon and Stinchcombe (1995), admissibility is not a property that one can generally expect from Nash equilibrium refinements in continuum-action games.

Example 5 [Example 2.1 in Simon and Stinchcombe, 1995]. Consider the symmetric two-player game with $S_1 = S_2 = [0, 1/2]$ and payoff functions

$$\pi_1(s_1, s_2) = \begin{cases} s_1 & \text{if } s_1 \leq s_2/2 \\ s_2(1-s_1)/(2-s_2) & \text{otherwise} \end{cases}$$

and $\pi_2(s_1, s_2) \equiv \pi_1(s_2, s_1)$. This game has a unique Nash equilibrium, $(0, 0)$, but for each player i , the strategy $s_i = 0$ is weakly dominated by all other strategies. Nevertheless, this is a game that meets the conditions in Corollary 1, and hence the weakly dominated Nash equilibrium $(0, 0)$ is strictly robust to strategic uncertainty.

4.1. Continuously differentiable games. Some of the above examples had continuous, but not differentiable, payoff functions. The following example shows that our robustness criterion has cutting power also in games with continuously differentiable payoff functions. The payoff functions in question are concave in the player's own strategy, and the game has a continuum of Nash equilibria only one of which is robust to strategic uncertainty.

Example 6. Consider the two-player game with $S_1 = S_2 = [0, 1]$ and payoff functions

$$\begin{cases} \pi_1(s_1, s_2) = 1 - (1-s_1)(1-s_2) \\ \pi_2(s_1, s_2) = s_2 \end{cases}$$

All strategy profiles in which $s_2 = 1$ are Nash equilibria. However, only one of these, $(1, 1)$, is undominated. It is easily verified that the latter is the unique robust equilibrium. Moreover, it is strictly robust.

However, in games where each player's payoff function is also strictly concave in the player's own strategy, all three concepts — robustness, strict robustness and Nash equilibrium — coincide. By definition, strict robustness implies robustness, and, in force of Proposition 3, robustness implies Nash equilibrium in this class of games. Hence, the point to be proved here is that every Nash equilibrium in these games is strictly robust. This follows from the observation that since the relevant partial derivatives of the payoff functions in the original game are continuous, and the strategy sets are compact, the corresponding partial derivatives in a slightly perturbed game are “close” to those in the original game. Payoff functions being strictly concave, the partial derivative is non-zero to a player who unilaterally deviates from a Nash equilibrium in the original game. Hence, this is also true in the perturbed game. Formally:

Proposition 5. *Suppose that the payoff functions in G are continuously differentiable, that the strategy sets are compact and convex, and that each payoff function π_i is strictly concave in s_i (for every s_{-i}). Then the following three statements are equivalent: A strategy profile is strictly robust to uncertainty, it is robust to strategic uncertainty, and it is a Nash equilibrium, respectively.*

We note that all hypotheses, except for differentiability, are met by the game in Example 4, and, as we showed, infinitely many of the (strict) Nash equilibria are non-robust to strategic uncertainty. Hence, the result is almost sharp (the missing link being games with payoff functions that are differentiable but not continuously differentiable).

4.2. Weak perfection. We proceed to relate our robustness criteria to the notion of weak perfection in Simon and Stinchcombe (1995). For each player $i \in N$, let $S_i = [a_i, b_i]$ for $a_i < b_i$, and let Δ_i denote the set of Borel probability measures over S_i . For any $\mu \in \square = \times_{i \in N} \Delta_i$, let $\beta_i(\mu) \in \Delta_i$ denote i 's set of mixed best replies to the mixed-strategy profile μ . Following Simon and Stinchcombe (1995), define the *weak-metric distance* between two mixed strategies, μ_i and v_i , as follows:

$$\rho^w(\mu_i, v_i) = \inf\{\delta > 0 : \mu_i(B) \leq v_i(B^\delta) + \delta \text{ and } v_i(B) \leq \mu_i(B^\delta) + \delta, \text{ for all Borel sets } B \subset S_i\},$$

where B^δ is the δ -neighborhood of B . Identify pure strategies with unit point masses.

Definition 3 [Simon and Stinchcombe, 1995]. *For any $\varepsilon > 0$, a **weak perfect ε -equilibrium** is a completely mixed-strategy profile, $\mu^\varepsilon \in \text{int}(\square)$, such that, for every player $i \in N$, $\rho_i^w(\mu_i^\varepsilon, \beta_i(\mu^\varepsilon)) < \varepsilon$. A strategy profile $\mu^* \in \square$ is **weakly perfect** if it is the limit as $\varepsilon_k \rightarrow 0$ of a sequence of weak perfect ε_k -equilibria.*

While our notion of robustness to strategic uncertainty allows two players to differ in their beliefs about a third player, such interpersonal inconsistency is, by construction, excluded by Simon and Stinchcombe, who work with mixed strategy profiles rather than with subjective beliefs. In order to establish a tight connection between their approach and ours, we call a collection of beliefs $\{\phi_{ij} \in \mathcal{F} : i \in N, j \neq i\}$ *interpersonally consistent* if, for each player $j \in N$, $\phi_{1j} = \dots = \phi_{nj}$. In other words, all players $i \neq j$ share beliefs about j 's strategy choice. We proceed to show that in games with continuous payoff functions and convex and compact strategy sets, robustness to interpersonally consistent strategic uncertainty implies weak perfection.¹⁰ The reason why interpersonal consistency matters is that a shared subjective probability distribution concerning a player's strategy choice can be viewed as a mixed strategy for that player.

Proposition 6. *If the payoff functions in G are continuous and the strategy sets convex and compact, then every strategy profile that is robust to interpersonally consistent strategic uncertainty is weakly perfect.*

Needless to say, strict robustness to strategic uncertainty requires, *inter alia*, robustness to interpersonally consistent belief systems. Hence, a strategy profile that is strictly robust is also weakly perfect. In the two-player case, only one player holds a belief about a given player, so the issue of interpersonal consistency of beliefs does not arise. Thus, in such games, robustness to strategic uncertainty implies weak perfection.

The next example shows that for the Bertrand competition game in Example 3, weak perfection in the sense of Simon and Stinchcombe (1995) admits Nash equilibria that are non-robust to strategic uncertainty:

Example 7. *Reconsider Example 3 and let the pure-strategy set for each firm be $[0, 1]$. The set of Nash equilibrium prices is the sub-interval $[1/11, 3/13]$ and we have $\bar{p} = 1/6$. Let μ_i^ε , for $i = 1, 2$ and $\varepsilon \in (0, 1)$, be a completely mixed strategy that consists of a positive uniform density on the full price interval $[0, 1]$, with total mass $\varepsilon^2 \in (0, 1)$, and a point mass on a candidate equilibrium price $p^* \in (1/11, 1/6)$ with probability mass $1 - \varepsilon^2$. We know from Proposition 2 that this Nash equilibrium price is not robust to strategic uncertainty. However, it is weakly perfect. To see this, note*

¹⁰By contrast, our notion of strict robustness to strategic uncertainty appears to be weaker than Simon's and Stinchcombe's (1995) notion of strong perfection. The reason is that they then work with point masses (attached to the pure best reply), while our probability distributions are continuous.

that the set of i 's pure best replies to the completely mixed strategy μ_j^ε (for $j \neq i$) is

$$\beta_i(\mu_j^\varepsilon) = \arg \max_{p \in [0,1]} \begin{cases} \varepsilon^2(1-p) \cdot v_1(p) + (1-\varepsilon^2) \cdot v_2(p) & \text{for } p = p^* \\ (1-\varepsilon^2)p \cdot v_1(p) & \text{for } p < p^* \\ \varepsilon^2(1-p) \cdot v_1(p) & \text{for } p > p^* \end{cases}$$

A price $p < \bar{p}$ such that $p \neq p^*$ is clearly not a best reply, since the expected profit is negative. The remaining candidates for best replies are $p = p^*$ or $p \geq \bar{p}$, so $\beta_i(\mu_j^\varepsilon) \subset \{p^*\} \cup [\bar{p}, 1]$. As $\varepsilon \rightarrow 0$, the expected profit converges to zero for each $p \geq \bar{p}$ and to $v_2(p) > 0$ for $p = p^*$. Hence, for all $\varepsilon > 0$ sufficiently small, $p = p^*$ is the unique best reply. It follows that $\mu^\varepsilon = (\mu_1^\varepsilon, \mu_2^\varepsilon)$ is a weak ε -perfect equilibrium, since $\rho^w(\mu_i^\varepsilon, \mu_i^*) = \varepsilon^2 < \varepsilon$, where μ_i^* assigns unit probability to $p_i = p^*$ (for $j \neq i$). Since $\mu_i^\varepsilon \rightarrow \mu_i^*$ as $\varepsilon \rightarrow 0$, (p^*, p^*) is weakly perfect.

5. MULTI-DIMENSIONAL STRATEGY SETS

Our criterion for robustness to strategic uncertainty is readily generalized from one-dimensional to multi-dimensional strategy sets. We here briefly outline a finite-dimensional generalization.¹¹ Let again $G = (N, S, \pi)$ be an n -player normal-form game with player set $N = \{1, \dots, n\}$, set $S = \times_{i \in N} S_i$ of strategy profiles $s = (s_1, \dots, s_n)$, and combined payoff function $\pi : S \rightarrow \mathbb{R}^n$, with $\pi_i(s)$ being the payoff to player i when strategy profile s is played.¹² Let the strategy set of each player i be a Borel set S_i , with non-empty interior, in a Euclidean space \mathbb{R}^{m_i} , for some positive integer m_i . Accordingly, S is a Borel set in \mathbb{R}^m , where $m = m_1 + \dots + m_n$. We assume each payoff function π_i to be Borel measurable and bounded. This is a large class of games. It includes all dynamic games of incomplete information with finite type spaces, a finite number of information sets and finite choice sets at non-final information sets, allowing for infinite Euclidean choice sets at final information sets.

For each positive integer k , let \mathcal{F}^k be the set of continuous and everywhere positive probability density functions on \mathbb{R}^k . For $\phi \in \mathcal{F}^k$ and any Borel set $B \subset \mathbb{R}^k$, let $\mu^\phi(B) = \int_B \phi(x) dx$. This defines a positive Borel measure on \mathbb{R}^k . Just as in the one-dimensional case, such density functions, adapted to the strategy set of a player, will represent *other* players' (subjective) beliefs about the strategy choice of that player. In particular, for two players i and j , and for any $\phi_{ij} \in \mathcal{F}^{m_j}$, let $\mu_{ij}(B) = \int_B \phi_{ij}(x) dx$ be the Borel probability measure on the Euclidean space \mathbb{R}^{m_j} that contains j 's strategy set, S_j .

Definition 4. For any given scalar $t \in [0, 1]$, a strategy profile s is a **t -equilibrium** of G if, for each player i , the strategy s_i maximizes i 's expected payoff under the

¹¹The approach can be generalized to infinite-dimensional strategy sets as well.

¹²Whether the strategies are interpreted as pure or mixed is immaterial for the subsequent analysis.

probabilistic belief that all other players' strategies are random variables of the form

$$\tilde{s}_{ij} = s_j + t \cdot \varepsilon_{ij} \quad (9)$$

for some statistically independent “noise” terms $\varepsilon_{ij} \sim \phi_{ij}^t$, for all $j \neq i$, where each noise probability density function ϕ_{ij}^t is the adaptation of some $\phi_{ij} \in \mathcal{F}^{m_j}$ to S_j :

$$\phi_{ij}^t(u) = \begin{cases} \phi_{ij}(u) / \mu_{ij}(U_j^t) & \text{if } s_j + t \cdot u \in S_j \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

where $U_j^t = \{u \in \mathbb{R}^{m_j} : s_j + t \cdot u \in S_j\}$ has non-empty interior, and μ_{ij} is the positive Borel measure induced on \mathbb{R}^{m_j} by ϕ_{ij} .

Remark 5. For $t = 0$, this definition coincides with that of Nash equilibrium.

Remark 6. For $t > 0$, the random variable \tilde{s}_{ij} has the probability density $f_{ij}^t : \mathbb{R}^{m_j} \rightarrow \mathbb{R}_+$ defined by

$$f_{ij}^t(x) = \begin{cases} \frac{1}{t} \cdot \phi_{ij}\left(\frac{x-s_j}{t}\right) / \mu_{ij}(U_j^t) & \text{if } x \in S_j \\ 0 & \text{otherwise} \end{cases}$$

Our definition of robustness to strategic uncertainty is easily extended to this large class of games:

Definition 5. A strategy profile s^* in G is **robust to strategic uncertainty** if there exists a collection of probability density functions $\{\phi_{ij} \in \mathcal{F}^{m_j} : i \in N, j \neq i\}$ and an accompanying sequence of t -equilibria, $\langle s^{t_k} \rangle_{k \in \mathbb{N}}$ with $t_k \downarrow 0$, such that $s^{t_k} \rightarrow s^*$ as $k \rightarrow +\infty$. The strategy profile s^* is **strictly robust** to strategic uncertainty if this holds for all collections of probability density functions $\{\phi_{ij} \in \mathcal{F}^{m_j} : i \in N, j \neq i\}$.

The definition in Section 2 is the special case when $m_1 = \dots = m_n = 1$. It is easily verified that all claims in Section 4 remain valid under this generalization.

6. RELATED LITERATURE

Our approach is not epistemic in the sense of being derived from a theory of how players reason. In particular, players' uncertainty is not derived from assumptions about their knowledge or beliefs about the game and/or each others' rationality (Brandenburger, 2007). Our aim is more modest; we simply study what play may arise when players are slightly strategically uncertain, whatever the reasons for their uncertainty may be. By contrast, the literature on global games (Carlsson and van Damme, 1993; Morris and Shin, 2003), derives players' uncertainty from assumptions within an incomplete-information meta-game in which the game at hand is embedded. We

view our approach as a framework within which researchers may introduce additional hypotheses about the source(s) of strategic uncertainty and thereby constrain the subjective beliefs.¹³

Simon and Stinchcombe (1995) and Méndez-Naya et al. (1995) study perfection and related ideas in games with compact and convex strategy sets, and continuous payoff functions. Our main application is to a game with discontinuous payoff functions. Carlsson and Ganslandt (1998) investigate “noisy equilibrium selection” in symmetric coordination games and derive results that agree with the experimental findings on minimal effort games in Van Huyck et al. (1990). While Carlsson’s and Ganslandt’s (1998) study is tailored to such games, our approach permits analysis of robustness to strategic uncertainty for a large class of games.

Although quantal response equilibrium (QRE), pioneered by McKelvey and Palfrey (1995), was initially developed for finite games, Anderson, Goeree and Holt (1998) extend it to infinite-action games. The QRE approach assumes probabilistic choice, driven either by idiosyncratic perturbations of individual players’ preferences (here profit functions) or by mistakes in their implementation of strategies. The choice probabilities are increasing functions of the expected payoffs. A QRE requires consistency in the sense that each player’s probability distribution follows from its presumed functional form as applied to the expected payoffs from others’ probabilistic choices. In particular, players do not best respond to their information or beliefs about each other, which occurs only in the limit as the noise level is driven down to zero. Baye and Morgan (2004) apply the continuum-action QRE to a class of Bertrand games that includes our main application (although this is implicit in their setup). More exactly, their price c is our price \bar{p} , and they study symmetric QRE for probabilistic price setting on the interval $[\bar{p}, p^{mon}]$ (in our notation), while we allow for beliefs that assign positive probability also to prices below \bar{p} . Assuming that the noise terms have a power distribution, they show the existence of a symmetric QRE for each parameter value $\lambda > 0$ (in their notation), where a higher λ means that the random choice comes closer (in probability) to a best reply. As $\lambda \rightarrow +\infty$, the probability mass moves towards a unit point mass at the price $c = \bar{p}$, that is, the same price that we show is robust under strategic uncertainty. By contrast, we do not assume symmetry. Instead, we derive both symmetry and uniqueness, and also show that \bar{p} is in fact strictly robust, that is, robust for a wide range of probabilistic beliefs about other’s strategy choices. Although our result agrees with their limit result in those games, Example 6 shows that this agreement is not general. Indeed, the unique logistic QRE (that is, with Gumbel distributed noise terms) in that example converges, as $\lambda \rightarrow +\infty$, to a limit distribution with $\mathbb{E}[s_1] = 1/2$ and $\mathbb{E}[s_2] = 1$, a weakly

¹³For instance, for some applications it may be desirable to constrain beliefs to have support on the set of rationalizable strategies.

dominated Nash equilibrium in the unperturbed game.¹⁴ By contrast, the unique undominated Nash equilibrium in this game, $(1, 1)$, is the unique robust equilibrium.

Some other selection criteria for Bertrand competition have been proposed in the literature. Spulber (1995) assumes that firms are uncertain about rivals' costs and shows that there exists a unique symmetric Nash equilibrium in pure strategies. As the number of firms grows, equilibrium pricing strategies tend to average cost pricing, which is clearly at variance with our prediction. Chowdhury and Sengupta (2004) show that, in Bertrand games with convex costs, there exists a unique coalition-proof Nash equilibrium (in the sense of Bernheim et al., 1987), which converges to the competitive outcome under free entry, again a different prediction from ours, which does not depend on the number of firms.

7. CONCLUSION

In this paper, we proposed a way to model strategic uncertainty in a straightforward fashion. After defining a robustness criterion, we investigated in detail the Bertrand competition game, which has discontinuous payoffs and continuum action spaces. We then went on to deriving a handful of general results for games with continuous payoffs. In the Bertrand competition game with convex costs, we showed that our notion of robustness to strategic uncertainty selects a unique Nash equilibrium, that, moreover, figured prominently in recent laboratory experiments.

Our criterion for robustness to strategic uncertainty comes at a relatively low analytical cost and we believe it has a wide domain of application. Indeed, in an earlier working paper (Andersson, Argenton and Weibull, 2010) we also consider the Nash demand game, and we show that robustness to symmetric strategic uncertainty singles out the Nash bargaining solution. For the case with asymmetric strategic uncertainty we find that the party who is least uncertain about the other party's bid obtains the bigger share. To name further examples that are yet to be explored, mechanisms for the provision of discrete public goods à la Bagnoli and Lipman (1989) are games with continuum action-spaces and discontinuous payoffs, admitting multiple equilibria. We think it would be worthwhile to apply our robustness criterion to such games and to check it against the experimental literature.

As mentioned earlier, our approach is agnostic regarding the source of the beliefs. However, there are several straightforward extensions involving arguably natural restrictions on beliefs. For instance, while we allow beliefs to attach positive probabilities to others' strategies that are (strictly or weakly) dominated or non-rationalizable, the notion of t -equilibrium — and hence robustness to strategic uncertainty — can be weakened by imposing requirements that the support of beliefs be restricted to others' undominated or rationalizable strategies.

¹⁴The limit QRE distribution for player 1 is the uniform distribution on 1's pure strategy set and the limit QRE distribution for player 2 is a unit point mass at $s_2 = 1$.

8. APPENDIX

Proof of Proposition 1: We note that for given $t > 0$, each player's strategy set is non-empty, convex and compact and each player's payoff function is continuous. By Weierstrass's maximum theorem, each player's best reply correspondence is non-empty and compact-valued. By Berge's maximum-theorem, it is also upper hemicontinuous. Existence of a Nash equilibrium in \bar{G}^t thus follows from Kakutani's fixed-point theorem if, in addition, each player's best-reply correspondence is convex-valued. In order to verify this, first note that no equilibrium price can lie on the boundary of the strategy set in \bar{G}^t , since $\partial u_i^t / \partial p_i$ is positive at its left boundary and negative at its right boundary. Hence, any Nash equilibrium of \bar{G}^t is in $(\bar{p}, p^{mon})^n$. It remains to show that $\arg \max_{p_i \in (\bar{p}, p^{mon})} u_i^t(\mathbf{p})$ is convex. For this purpose, note that

$$p_i^t \in \arg \max_{p_i \in (\bar{p}, p^{mon})} u_i^t(p_i, p_{-i})$$

if and only if

$$p_i^t \in \arg \max_{p_i \in (\bar{p}, p^{mon})} \left(\ln [v_1(p_i)] + \sum_{i \neq j} \ln \left[1 - \Phi_{ij} \left(\frac{p_i - p_j}{t} \right) \right] \right).$$

Since v_1 is concave by assumption, so is $\ln [v_1(\cdot)]$. By assumption, ϕ_{ij} is continuously differentiable and concave on an open interval, which, by Theorem 3 in Bagnoli and Bergstrom (2005), implies that also the survival function, $1 - \Phi_{ij}$, is log concave. Hence, each term in the above sum is a concave function of p_i (given p_j and t). Concavity is preserved under summations, so the maximand is concave, and thus i 's best-reply correspondence is convex-valued.

Proof of Proposition 2: Let $\phi_{ij} \in \mathcal{F} \forall i \in N$ and $j \neq i$. Consider any sequence $\langle t_k \rangle_{k=1}^\infty \rightarrow 0$, where each $t_k > 0$. For each $k \in N$, let p^k be a Nash equilibrium of \bar{G}^{t_k} . Since all games \bar{G}^{t_k} have the same strategy space, $[\bar{p}, p^{mon}]^n$, and this is non-empty and compact, the sequence $\langle \mathbf{p}^{t_k} \rangle_{k=1}^\infty$ contains a convergent subsequence with limit in $[\bar{p}, p^{mon}]^n$, according to the Bolzano-Weierstrass Theorem. Hence, without loss of generality we may assume that $\lim_{k \rightarrow \infty} p^k = p^* \in [\bar{p}, p^{mon}]^n$.

First, we prove that $p_i^* = p_j^*$ for all $i, j \in N$. For this purpose, first note that no price can lie on the boundary of the strategy set in \bar{G}^t , since $\partial u_i^t / \partial p_i$ is positive at its left boundary and negative at its right boundary.

$$\frac{\partial u_i^t(\mathbf{p})}{\partial p_i^k} = \prod_{j \neq i} \left[1 - \Phi_{ij} \left(\frac{p_i^k - p_j^k}{t} \right) \right] \cdot \left[v_1'(p_i^k) - \frac{v_1(p_i^k)}{t} \sum_{j \neq i} h_{ij} \left(\frac{p_i^k - p_j^k}{t} \right) \right]$$

Hence, $\bar{p} < p_i^k < p^{mon}$ for all i and k . The first-order condition gives,

$$t_k v_1'(p_i^k) = v_1(p_i^k) \sum_{j \neq i} h_{ij} \left(\frac{p_i^k - p_j^k}{t_k} \right) \quad \forall i, k \quad (11)$$

where h_{ij} is the hazard-rate function of Φ_{ij} . Consider a firm $i \in N$. Suppose that $p_j^* < p_i^*$ for some $j \neq i$, and let $\varepsilon = p_i^* - p_j^* > 0$. Then, there is a K such that $p_i^k - p_j^k > \varepsilon/2$ for all $k > K$. The hazard rate being non-decreasing, we thus have

$$h_{ij} \left(\frac{p_i^k - p_j^k}{t_k} \right) \geq h_{ij} \left(\frac{\varepsilon}{2t_k} \right)$$

for that $j \neq i$ and all $k > K$. Let $\delta = h_{ij}[\varepsilon/(2t_K)] > 0$. Then

$$h_{ij} \left(\frac{p_i^k - p_j^k}{t_k} \right) \geq \delta$$

for that $j \neq i$ and all $k > K$, and hence, since all hazard rates are positive:

$$t_k v_1'(p_i^k) > \delta \cdot v_1(p_i^k)$$

for all $k > K$. However, $t_k v_1'(p_i^k) \rightarrow 0$ and $v_1(p_i^k) \rightarrow v_1(p_i^*)$ as $k \rightarrow \infty$, since v_1 is continuous, so $v_1(p_i^*) = 0$. Hence, $p_i^* = \bar{p}$. But this contradicts the hypothesis $p_i^* > p_j^* \in [\bar{p}, p^{mon}]$. Hence, $p_j^* \geq p_i^*$. Since this holds for all i and $j \neq i$, we conclude that $p_j^* = p_i^*$ for all $i, j \in N$.

Secondly, we prove $p_i^* = \bar{p}$ for all $i \in N$. Since $v_1(p_i^k) > 0$ on (\bar{p}, p^{mon}) and all hazard rates are positive, by (11),

$$v_1(p_i^k) \cdot h_{ij} \left(\frac{p_i^k - p_j^k}{t_k} \right) \rightarrow 0 \quad \forall i, j \neq i$$

as $k \rightarrow +\infty$. Suppose that $p_i^* > \bar{p}$. Then $v_1(p_i^*) > 0$ and thus

$$h_{ij} \left(\frac{p_i^k - p_j^k}{t_k} \right) \rightarrow 0 \quad \forall j \neq i$$

implying that $p_i^k < p_j^k$ for all k sufficiently large. But, by the same token: since $p_j^* = p_i^*$, for all $j \neq i$, we also have $p_j^* > \bar{p}$ and $v_1(p_j^*) > 0$ and thus

$$h_{ji} \left(\frac{p_j^k - p_i^k}{t_k} \right) \rightarrow 0$$

implying that $p_j^k < p_i^k$ for all k sufficiently large. Both strict inequalities cannot hold. Hence, $p_i^* = \bar{p}$ for all $i \in N$. In sum: the only strategy profile that is robust to strategic uncertainty is $(\bar{p}, \dots, \bar{p})$. The strict robustness claim follows immediately from the fact that the collection $\{\Phi_{ij} \in \mathcal{F} : \forall i \in N, j \neq i\}$ was arbitrary.

Proof of Proposition 3: Let $\phi_{ij} \in \mathcal{F} \forall i \in N$ and $j \neq i$, and let $\langle s^{t_k} \rangle_{k \in \mathbb{N}}$ be a sequence of t_k -equilibria with $t_k \downarrow 0$, such that $s^{t_k} \rightarrow s^*$, where s^* is not a Nash equilibrium in G . Then, there exists a player i and pure strategy $s_i^0 \in S_i$ such that $\pi_i(s_i^0, s_{-i}^*) - \pi_i(s^*) = \delta$ for some $\delta > 0$. Now, π_i^t converges point-wise to π_i as $t \rightarrow 0$, by continuity of π_i . Hence, $\pi_i^t(s_i^0, s_{-i}^*) - \pi_i^t(s^*) > \delta/2$ for all $t > 0$ sufficiently close to zero, say, for all $t < \tau(\delta)$, where $\tau(\delta) > 0$. Moreover, since π_i^t is continuous, there exists a neighborhood $B(t)$ of s^* such that

$$\pi_i^t(s_i^0, s'_{-i}) - \pi_i^t(s') > \delta/3$$

for all $s' \in B(t)$ and $t \in (0, \tau(\delta))$. In particular, $B(t)$ contains no t -equilibrium. This holds for each $t \in (0, \tau(\delta))$, contradicting the hypothesis that all the elements of the sequence $\langle s^k \rangle_{k=1}^\infty \rightarrow s^*$ are t_k -equilibria.

Proof of Proposition 4: Let G be as stated in the proposition. Let $\phi_{ij} \in \mathcal{F}$, $\forall i \in N, j \neq i$. Consider any sequence $\langle t_k \rangle_{k=1}^\infty \rightarrow 0$, where each $t_k > 0$. For each $k \in \mathbb{N}$, suppose that s^k is a Nash equilibrium of G^{t_k} . Then, s^k is a t_k -equilibrium of G . Now $s^k \in S$ for all $k \in \mathbb{N}$, where S is a non-empty and compact set, so $\langle s^k \rangle_{k=1}^\infty$ admits a convergent subsequence, with limit $s^* \in S$, by the Bolzano-Weierstrass Theorem.

It remains to establish that for each $t > 0$, the perturbed game G^t has at least one Nash equilibrium. To apply Kakutani's fixed-point theorem, it is sufficient to show that each payoff function π_i^t is continuous and concave in s_i (for every s_{-i}). We prove this for the case of two players, but the generalization is immediate. By definition:

$$\pi_i^t(s) = \mathbb{E}[\pi_i(s_i, \tilde{s}_j)] = \int \pi_i(s_i, x) f_{ij}^t(x) dx$$

This defines π_i^t as a continuous function. Moreover, $\forall \lambda \in (0, 1)$ and $s_i, s'_i, s_j \in S$:

$$\begin{aligned} \pi_i^t[\lambda s_i + (1 - \lambda) s'_i, s_j] &\geq \int [\lambda \pi_i(s_i, x) + (1 - \lambda) \pi_i(s'_i, x)] f_{ij}^t(x) dx \\ &= \lambda \int \pi_i(s_i, x) f_{ij}^t(x) dx + (1 - \lambda) \int \pi_i(s'_i, x) f_{ij}^t(x) dx \\ &= \lambda \pi_i^t(s_i, s_j) + (1 - \lambda) \pi_i^t(s'_i, s_j). \end{aligned}$$

where the inequality holds by concavity of π . This proves that π_i^t is concave in s_i (for every s_{-i}).

Proof of Corollary 1: Let G be as stated. By Proposition 4, G admits a robust strategy profile. By Proposition 3 this is also a Nash equilibrium. By hypothesis, there is no other Nash equilibrium. Hence, since the collection $\{\phi_{ij} \subset \mathcal{F}, \forall i \in N, j \neq i\}$ in the proof of Proposition 4 was arbitrary, s^* is strictly robust.

Proof of Proposition 5: By definition, strict robustness to strategic uncertainty implies robustness to strategic uncertainty, which, by Proposition 3 implies Nash equilibrium in this class of games. It thus suffices to show that every Nash equilibrium is strictly robust. Consider a game $G = (N, S, \pi)$ in which all strategy sets are non-empty, compact and convex, and the payoff function to each player i is continuously differentiable in the sense that the mapping $s \mapsto \partial\pi_i(s)/\partial s_i$ is continuous. Let $\phi_{ij} \in \mathcal{F}, \forall i \in N, j \neq i$, and $t > 0$. By Leibniz' rule of differentiation:

$$\begin{aligned} \frac{\partial\pi_i^t(s)}{\partial s_i} &= \frac{\partial}{\partial s_i} \int \dots \int \left[\prod_{j \neq i} f_{ij}^t(x_j) \right] \pi_i(s_i, x_{-i}) dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_n \\ &= \int \dots \int \left[\prod_{j \neq i} f_{ij}^t(x_j) \right] \frac{\partial}{\partial s_i} \pi_i(s_i, x_{-i}) dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_n \end{aligned}$$

The partial derivative $\partial\pi_i(s)/\partial s_i$ being continuous:

$$\begin{aligned} \lim_{t \rightarrow 0} \frac{\partial\pi_i^t(s)}{\partial s_i} &= \lim_{t \rightarrow 0} \int \dots \int \left[\prod_{j \neq i} f_{ij}^t(x_j) \right] \frac{\partial}{\partial s_i} \pi_i(s_i, x_{-i}) dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_n \\ &= \frac{\partial\pi_i(s)}{\partial s_i} \end{aligned}$$

Arguing as in the proof of Proposition 2, we obtain that, for all $t > 0$ sufficiently small, π_i^t is strictly concave in s_i (for every s_{-i}).

Suppose that $s^* \in S$ is a Nash equilibrium of G . For any $i \in N$ and $s_i \in S_i$: $\partial\pi_i(s_i, s_{-i}^*)/\partial s_i > 0$ if $s_i < s_i^*$ and $\partial\pi_i(s_i, s_{-i}^*)/\partial s_i < 0$ if $s_i > s_i^*$. By the above observation, we also have $\partial\pi_i^t(s)/\partial s_i \rightarrow \partial\pi_i(s)/\partial s_i \forall s \in S$, for all players i . The functions being continuous on a compact set S , the convergence is uniform. For any $\varepsilon > 0$, let $B(\varepsilon) = \times_{i \in N} [s_i^* - \varepsilon, s_i^* + \varepsilon]$. For each $\varepsilon > 0$, there exists a $t_\varepsilon > 0$ such that, for each player i , $s \in S \cap B(\varepsilon)$ and $s'_i \in S_i$: $\partial\pi_i^{t_\varepsilon}(s'_i, s_{-i})/\partial s_i > 0 \forall s'_i \leq s_i^* - \varepsilon$ and $\partial\pi_i^{t_\varepsilon}(s'_i, s_{-i})/\partial s_i < 0 \forall s'_i \geq s_i^* + \varepsilon$. Hence, for each $t \in (0, t_\varepsilon)$ there exist, by standard arguments, a t -equilibrium $s^t \in S \cap B(\varepsilon)$. Moreover, by letting $\varepsilon \rightarrow 0$, we obtain a sequence of t -equilibria that converges to s^* .

Proof of Proposition 6: Suppose that s^* is robust to interpersonally consistent strategic uncertainty. Then there exist $\psi_1, \dots, \psi_n \in \mathcal{F}$ such that, for each $j \in N$,

$\phi_{ij} = \psi_j$ for all $i \neq j$, and s^* is the limit as $t_k \rightarrow 0$ of a sequence of t_k -equilibria, s^{t_k} , with these probability densities. The claim is established if we can extract a subsequence of weak perfect ε_h -equilibria with $\varepsilon_h \rightarrow 0$. For each $k \in \mathbb{N}$ and player i let

$$\hat{f}_i^k(x) = \frac{1}{t_k} \cdot \frac{\psi_i\left(\frac{x-s_i}{t_k}\right)}{\Psi_i\left(\frac{b_i-s_i}{t_k}\right) - \Psi_i\left(\frac{a_i-s_i}{t_k}\right)}$$

where Ψ_i is the c.d.f. induced by the probability density ψ_i , and where, without loss of generality, each (compact and convex) strategy set has been taken to be of the form $S_i = [a_i, b_i]$ for some $a_i < b_i$. Let \hat{F}_i^k be the c.d.f. induced by \hat{f}_i^k . For each $h \in \mathbb{N}$, let $\varepsilon_h = 1/h$, and, for each player i , let $k(i, h)$ be the minimal $k \in \mathbb{N}$ such that $\hat{F}_i^k(s_i^{t_k} + 1/h) - \hat{F}_i^k(s_i^{t_k} - 1/h) > 1 - 1/h$. Such a k clearly exists. Let $k(h) = \max_{i \in N} k(i, h)$. Then $s^{t_{k(h)}}$ is a weak perfect ε_h -equilibrium.

REFERENCES

- [1] Anderson, S.P., J. K. Goeree and C. A. Holt, (1998): Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction, *Journal of Political Economy*, 106 , 828-853.
- [2] Andersson, O., C. Argenton, and J. W. Weibull (2010): Robustness to strategic uncertainty, CentER discussion paper 2010-70.
- [3] Abbink, K., and J. Brandts (2008): 24. Pricing in Bertrand competition with increasing marginal costs, *Games and Economics Behavior*, 63, 1-31.
- [4] Argenton, C., and W. Müller (2009): Collusion in experimental Bertrand duopolies with convex costs: the role of information and cost asymmetry, CentER discussion paper 2009-87.
- [5] Bagnoli, M., and T. Bergstrom (2005): Log-concave probability and its applications, *Economic Theory*, 26, 445-469.
- [6] Bagnoli, M. and B.L. Lipman (1989): Provision of public goods: fully implementing the core through private contributions, *Review of Economic Studies*, 56, 583-601.
- [7] Baye, M. R., and J. Morgan (2004): Price Dispersion in the Lab and on the Internet: Theory and Evidence, *The RAND Journal of Economics* , 35, 449-466
- [8] Bernheim, D., B. Peleg, and M. D Whinston (1987): Coalition-proof Nash equilibria. I: Concepts, *Journal of Economic Theory*, 42, 1–12.
- [9] Brandenburger, A. (1996): Strategic and structural uncertainty in games, in R. Zeckhauser, R. Keeney and J. Sibenius (eds.), *Wise choices: games, decisions, and negotiations*, Boston: Harvard Business School Press.
- [10] Brandenburger, A. (2007): The power of paradox: some recent developments in interactive epistemology, *International Journal of Game Theory*, 35, 465-492.
- [11] Cabrales, A, R. Miniaci, M. Piovesan, and G. Ponti (2010): Social preferences and strategic uncertainty: an experiment on markets and contracts, *American Economic Review*, 100, 2261-2278.
- [12] Carlsson, H., and M. Ganslandt (1998): Noisy equilibrium selection in coordination games, *Economics Letters*, 60, 23-34.
- [13] Carlsson, H., and E. van Damme (1993): Global games and equilibrium selection, *Econometrica*, 61, 989-1018.

- [14] Chowdhury, P. R., and K. Sengupta (2004): Coalition-proof Bertrand equilibria, *Economic Theory* 24, 307-324.
- [15] Dastidar, K. G. (1995): On the existence of pure strategy Bertrand equilibrium, *Economic Theory*, 5, 19-32.
- [16] Dixon, H. D. (1990): Bertrand-Edgeworth equilibria when firms avoid turning customers away, *Journal of Industrial Economics*, 39, 131–146.
- [17] Harsanyi, J., and R. Selten (1988): *A General Theory of Equilibrium Selection in Games*, Cambridge: The MIT Press
- [18] Heinemann, F., R. Nagel, and P. Ockenfels (2009): Measuring strategic uncertainty in coordination games, *Review of Economic Studies*, 76, 181-221.
- [19] McKelvey, D., and T.R. Palfrey (1995): Quantal Response Equilibria for Normal Form Games, *Games and Economic Behavior*, 10, 6-38.
- [20] Méndez-Naya, L., I. García-Jurado and J.C. Cesco (1995): Perfection of Nash equilibria in continuous games, *Mathematical Social Sciences*, 29, 225-237.
- [21] Morris, S., and H.S. Shin (2002): Measuring strategic uncertainty, mimeo.
- [22] Morris, S., and H.S. Shin (2003): Global games: theory and applications, in M. Dewatripont, L. Hansen and S. Turnovsky (eds.), *Advances in economics and econometrics, the eighth world congress*, New York: Cambridge University Press.
- [23] Simon, C., and M. Stinchcombe (1995): Equilibrium refinement for infinite normal-form games, *Econometrica*, 63, 1421-1443.
- [24] Spulber, D. F.(1995): Bertrand competition when rival's costs are unknown, *Journal of Industrial Economics*, 43, 1–12.
- [25] van Huyck, J. B., R. C. Battalio, and R. O. Beil (1990): Tacit coordination games, strategic uncertainty, and coordination failure, *American Economic Review*, 80, 234-248.
- [26] van Huyck, J. B., R. C. Battalio, and R. O. Beil (1991): Strategic uncertainty, equilibrium selection, and coordination failure, *Quarterly Journal of Economics*, 106, 885-911.
- [27] Vives, X. (1999): *Oligopoly pricing: Old ideas and new tools*, Cambridge: The MIT Press.
- [28] Weibull, J. W. (2006): Price competition and convex costs, SSE/EFI Working Paper Series in Economics and Finance No 622.