P-Stable Equilibrium : Definition and Some Properties

Gabriel Desgranges and Sayantan Ghosal

No 952

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

P-STABLE EQUILIBRIUM: DEFINITION AND SOME PROPERTIES*

GABRIEL DESGRANGES AND SAYANTAN GHOSAL
Université de Cergy-Pontoise and University of Warwick

December 17, 2010

Abstract

We define a continuous index of strategic stability, $p-$stability, which requires equilibrium to be the unique outcome compatible with common knowledge of rationality and common knowledge of $p-$beliefs (beliefs that put probability at least p on the equilibrium profile). We show that every equilibrium (within a large class) is p-stable for some $p < 1$ and justify, in smooth settings, the intuition that the slope of the best response map is related to the stability of equilibrium. We show that adding incomplete information on fundamentals could decrease the degree of strategic stability. In two applications to large markets we (i) show that a unique equilibrium globally unstable (under tâtonnement dynamics) has, nevertheless, a measure of strategic stability, (ii) characterize the conditions under which enhanced equilibrium efficiency results in decreased strategic stability.

Keywords: common knowledge, strategic uncertainty and rationalizability.

JEL Classifications C70 and D84.

1 Introduction

A common intuition relates the stability of equilibrium to the slope of best-response resulting in a "stable/unstable" typology. In this paper, we want to define a continuous index for stability that can be assigned to each equilibrium within a large class of games. Our aim is to obtain a stability index that is less coarse that the usual "stable/unstable" typology. Nevertheless, we require that our stability index, in smooth applications, justifies the intuition that the degree of equilibrium stability relates to the slope of the best response map.

Our underlying stability notion is one that requires the equilibrium to be robust to strategic uncertainty (and not uncertainty on fundamentals). Our concept is defined via an iterative elimination process. A Nash equilibrium is $p-$stable if it is the unique rationalizable outcome when all players attach at least probability $p < 1$ to the equilibrium strategy. If all players believe a Nash equilibrium action configuration to be played with probability one ($p = 1$), then, by definition, all players *only* play the Nash equilibrium. When $p = 0$, the property that only the Nash equilibrium is played is equivalent to requiring that the Nash equilibrium configuration is the unique rationalizable outcome. Our stability concept requires that *only* the Nash equilibrium configuration of actions be played when $p < 1$.

We say that a Nash equilibrium is inadmissible if the best-response map is "vertical at the equilibrium" i.e. a small change in the other players actions implies a infinitely large change in any one player's best response. Our main result shows that every admissible Nash equilibrium is $p-$stable for some $p < 1$. In a smooth setting we show that the degree of stability is related to the inverse of the slope of the best-response. Thus, our result generalizes, to models with a continuum of actions (where individual best-responses may vary continuously in the actions of other players and/or have empty basins of attraction) and a continuum of traders, the intuition that in a game with a finite number of pure strategies any strict Nash equilibrium is $p-$dominant (Morris, Rob and Shin (1995)) and is related to the idea of iterated $p-$dominance (Tercieux (2006)).

We explore the epistemic conditions for a p-stable equilibrium i.e. we exhibit (knowledge) assumptions implying that the outcome of the game is a p-stable equilibrium. We show that, for a given equilibrium, p-stability means that the equilibrium configuration of actions is the only outcome compatible with common knowledge of p-belief of equilibrium and common knowledge of rationality. This result provides an epistemic motivation for our stability concept and relate to results obtained elsewhere such as common knowledge of rationality implies that the outcome is rationalizable (Tan and Werlang (1988)), mutual knowledge of the payoff functions and of rationality, and common knowledge of the conjectures imply Nash (Aumann and Brandenburger (1995)) and common p-belief of rationality implies p-rationalizability (Hu (2007)).

Next, we develop a "contagion" argument relating strategic stability in the complete information case to incomplete information by adding payoff uncertainty (and informational asymmetry) to the underlying complete information game. We build an incomplete information game from a collection of complete

2

information games where the equilibrium is sometimes (over a non null set of fundamentals), but not always dominance solvable. In a game with a continuum of actions where the BR map is linear in the mean action, we show that adding incomplete and asymmetric information on the fundamentals decreases the degree of stability. The result holds both in an example with strategic complementarities and in an example with strategic substitutabilities. Thus the "contagion" argument developed here qualifies the point made in the "global games" literature (Carlsson and Van Damme (1993), Morris, Rob and Shin (1995), Morris and Shin (1998) amongst many others) where the introduction of incomplete information increases the degree of stability.

We, then, apply our stability concept to in two models of exchange in large markets. First, we use our concept to study the stability of equilibrium prices and allocations in the exchange economy studied by Scarf (1960). In this economy, the unique competitive equilibrium is globally unstable under tâtonnement dynamics. Modelling exchange explicitly via the Shapley Window model (Sahi and Yao (1989), Codognato and Ghosal (2000)), we show that the actions of each trader at a Nash equilibrium supporting to the unique competitive equilibrium is p-dominant and hence p-stable for some $0 < p < 1$. Therefore, although the equilibrium outcome fails to be globally stable under tâtonnement dynamics, it does have a measure of strategic stability[1]. Second, using a modified market game a la Shapley-Shubik (Shapley-Shubik (1977)), we characterize the conditions under which enhanced equilibrium efficiency, by getting rid of trading frictions, lowers the degree of strategic stability i.e. the minimum p for which equilibrium is $p-$stable increases in value.

The next section develops the stability concept, characterizes the class of equilibria which are p stable for some $p < 1$ and relates the index of stability to the slope of the best-response in the one dimensional, smooth case. Section 3 is devoted to the epistemic foundations of the stability concept while section 4 examines how strategic stability is affected by uncertainty on fundamentals. In section 5, we study two applications of our stability concept to exchange economies. The last section concludes and the appendix contains some technical material not included in the main text.

2 The stability concept

In this section, we begin by describing the underlying strategic framework. We then define $p-$stability and state the conditions under which a Nash equilibrium is p-stable.

2.1 The model

The underlying strategic framework is due to MasColell (1984).

---

[1] Although, clearly, the Nash equilibrium actions supporting the unique competitive equilibrium profile fails to be the unique rationalizable outcome (i.e. $0-$stable).

Let A be a non-empty compact metric space of actions and $\Delta(A)$ be the compact and metrizable set of Borel probability measures on A endowed with the weak convergence topology (which is metrizable using the Prohorov metric). Let U_A be the set of continuous utility functions $u : A \times \Delta(A) \to I\!R$ endowed with the supremum norm, the metric, separable and complete space of player characteristics. A game with a continuum of players is a Borel measure μ on U_A.

For any probability measure τ on $U_A \times A$, let τ_u and τ_a denote the respective marginal distribution on U_A and A respectively.

Let T denote the set of probability measures on $U_A \times A$ such that $\tau_u = \mu$: τ_u denotes the set of "strategy profiles", *i.e.*, a distribution of actions for each u.

For $\tau \in T$, let $B_\tau = \{(u,a) : u(a, \tau_a) \geq u(A, \tau_a)\}$. The best-response correspondence is a map $\phi : T \to T$ such that $\phi(\tau) = \{\tau' \in T : \tau'(B_\tau) = 1\}$, *i.e.*, ϕ is the set of "strategy profiles" putting probability one on the fact that each player plays a best-response to τ.

A Nash equilibrium is a measure $\tau^* \in T$ such that $\tau^* \in \phi(\tau^*)$.

Existence result. For a given μ, there exists a Nash equilibrium distribution τ^* (Theorem 1 in MasColell (1984)).

2.2 p-stability: definition

For a fixed equilibrium τ^* and $p \in [0, 1]$, a $p-$belief is a probability distribution $\tau_p = p\tau^* + (1-p)\tau$ for any $\tau \in T$, *i.e.*, a belief that assigns a probability p to the equilibrium τ^*. Let $T_p \subseteq T$ denote the corresponding set.

We define an equilibrium τ^* to be $p-$stable if the equilibrium distribution is the only element surviving the iterated elimination of non best-responses to a $p'-$belief for all $p' > p$. This definition relies on a "standard" definition of rationalizable outcomes in a game where the strategy set is restricted to T_p: a $p-$stable equilibrium is an equilibrium that is the only rationalizable outcome in a game with the restricted strategy set T_p.

Formally, we proceed to define $p-$stability iteratively. Let $S_p^0 = T_p$ and consider the sequence of sets $S_p^n = \left[\phi(S_p^{n-1})\right] \cap T_p$ for $n \geq 1$. This sequence is decreasing and therefore, it converges to a set S_p^∞. Note that $\tau^* \in S_p^\infty$.

Definition 1. A Nash equilibrium τ^* is $p-$stable if for all $p' > p$, $S_{p'}^\infty = \{\tau^*\}$.

Remarks:

1. We do not require that $S_p^\infty = \{\tau^*\}$ for a $p-$stable equilibrium. In some classes of games (for example in the smooth one-dimensional case below), $S_p^\infty \neq \{\tau^*\}$ at a $p-$stable equilibrium.

2. Clearly, when $p = 1$, $T_p = \{\tau^*\}$ such that every S_1^n and S_1^∞ are trivially equal to $\{\tau^*\}$. (With a slight abuse of notation, we could say that an equilibrium is always $1-$stable, but this is meaningless).

3. If τ^* is $0-$stable then τ^* is the unique rationalizable outcome.

4. For any $p < p'$, $p, p' \in [0, 1]$, we have that $S_{p'}^n \subseteq S_p^n$ as $T_{p'} \subseteq T_p$. Therefore, $S_{p'}^\infty \subseteq S_p^\infty$ and if $S_p^\infty = \{\tau^*\}$ then $S_{p'}^\infty = \{\tau^*\}$. In particular,

the set $I = \left\{ p \in [0,1] : S_p^\infty = \{\tau^*\} \right\}$ is an interval contained in $[0,1]$. Clearly, $sup.I = 1$. The interesting question is whether $inf.I < 1$.

2.3 p-stable equilibria

Define a best-response correspondence for each $u \in U_A$ to each $m \in \Delta(A)$ as $B(u,m) = \{a \in A : u(a,m) \geq u(A,m)\}$: an action in $B(u,m)$ is a best-response for $u \in U_A$ to some $m \in \Delta(A)$.

For each $m \in \Delta(A)$, consider the set

$$\tilde{U}_A(m) = \left\{ \begin{array}{c} u \in U_A : B(u,m) \text{ is not single-valued or} \\ \limsup_{m' \to m} \frac{d_A(B(u,m'),B(u,m))}{d_{\Delta(A)}(m',m)} < \infty \end{array} \right\}.$$

where d_A denotes a distance on A (recall that $d_{\Delta(A)}(.,.)$ denotes the Prohorov metric on $\Delta(A)$). Note that for each $u \in \tilde{U}_A(m)$, a small change in m induces an infinitely large change in best-responses.

Consider a given τ^*. For every u, denote $k_u = \limsup_{m \to \tau_a^*} \frac{d_A(B(u,m),B(u,\tau_a^*))}{d_{\Delta(A)}(m,\tau_a^*)}$.

Definition 2. The equilibrium τ^* of a game μ is admissible if $\mu\left(\tilde{U}_A(\tau_a^*)\right) = 0$ and

$$\sup_{u \in U_A} \mathsf{ess}\, k_u < +\infty,$$

Denote $K = \sup_{u \in U_A} \mathsf{ess}\, k_u$. K is the essential upper bound of k_u w.r.t. measure μ (that is: the set of u such that $k_u > K$ has $\mu-$measure 0).

The following lemma[2] summarizes three key properties of the Prohorov metric that will be used in the proof of the main result of the paper.

Lemma 1. (i) Consider $\tau = p\tau^* + (1-p)\tau'$. Then,

$$d_{\Delta(A)}(\tau_a, \tau_a^*) \leq (1-p)\, d_{\Delta(A)}(\tau_a', \tau_a^*).$$

(ii) Consider a Dirac measure δ_x and a distribution $\tau_a \in \Delta(A)$. Consider S the support of τ_a (the smallest closed set s.t. $\tau_a(S) = 1$) and $d = \sup_y \in S\, d_A(x,y)$ (d is the radius of the smallest ball centered on x that contains S[3]). Then,

$$d_{\Delta(A)}(\delta_x, \tau_a) \leq d.$$

(iii) Consider $\tau_a \in \Delta(A)$ defined by $\tau_a = \int \tau_\lambda f(d\lambda)$ where f is a probability distribution on a set of parameters λ. Consider another distribution $\nu \in \Delta(A)$. We have

$$d_{\Delta(A)}(\tau_a, \nu) \leq \sup_\lambda \mathsf{ess}\, d_{\Delta(A)}(\tau_\lambda, \nu).$$

---

[2] We state and prove this lemma for completeness as we are not aware of an explicit proof of the three properties of the Prohorov metric contained in the lemma and required for the proof of Proposition 1 below.

[3] Notice that x may be in S or not.

Proof. The Prohorov metric is defined by:

$$d_{\Delta(A)}\left(m,\tau_a^*\right) = \inf\left\{\varepsilon > 0 : m\left(M\right) \le \tau_a^*\left(M^\varepsilon\right) + \varepsilon \text{ for all Borel subsets } M \text{ of } A\right\},$$

where $M^\varepsilon = \{y \in A/d_A\left(x,y\right) \le \varepsilon \text{ for some } x \in M\}$.[4]

(i) For every M, for every $\varepsilon > d_{\Delta(A)}\left(\tau_a', \tau_a^*\right)$, we have

$$\tau_a^*\left(M\right) \le \tau_a'\left(M^\varepsilon\right) + \varepsilon,$$

then

$$(1-p)\,\tau_a^*\left(M\right) \le (1-p)\,\tau_a'\left(M^\varepsilon\right) + (1-p)\,\varepsilon,$$
$$\tau_a^*\left(M\right) \le p\tau_a^*\left(M\right) + (1-p)\,\tau_a'\left(M^\varepsilon\right) + (1-p)\,\varepsilon,$$
$$\tau_a^*\left(M\right) \le p\tau_a^*\left(M^\varepsilon\right) + (1-p)\,\tau_a'\left(M^\varepsilon\right) + (1-p)\,\varepsilon.$$

This implies: $d_{\Delta(A)}\left(\tau_a, \tau_a^*\right) \le (1-p)\,\varepsilon$ and

$$d_{\Delta(A)}\left(\tau_a, \tau_a^*\right) \le (1-p)\,d_{\Delta(A)}\left(\tau_a', \tau_a^*\right)$$

(ii) For a Borel set M s.t. $x \in M$, for every ε, we have $\delta_x\left(M^\varepsilon\right) = 1$ and

$$\tau_a\left(M\right) \le \delta_x\left(M^\varepsilon\right) + \varepsilon.$$

For a Borel set M that does not intersect S, for every ε, we have $\tau_a\left(M\right) = 0$ and

$$\tau_a\left(M\right) \le \delta_x\left(M^\varepsilon\right) + \varepsilon.$$

Consider now a Borel set M that does not contain x and that intersects S. For every $\varepsilon > d$, we have that $x \in M^\varepsilon$ (consider a y in $S \cap M$) and $\delta_x\left(M^\varepsilon\right) = 1$ and

$$\tau_a\left(M\right) \le \delta_x\left(M^\varepsilon\right) + \varepsilon.$$

(iii) Consider $\varepsilon > \sup_\lambda \mathbf{ess}\, d_{\Delta(A)}\left(\tau_\lambda, \nu\right)$. For $f-$almost every λ, we have: for every M

$$\tau_\lambda\left(M\right) \le \nu\left(M^\varepsilon\right) + \varepsilon.$$

Summing over λ gives:

$$\int \tau_\lambda\left(M\right) f\left(d\lambda\right) \le \int \left(\nu\left(M^\varepsilon\right) + \varepsilon\right) f\left(d\lambda\right) = \nu\left(M^\varepsilon\right) + \varepsilon,$$

as $\left(\nu\left(M^\varepsilon\right) + \varepsilon\right)$ does not depend on λ and $\int f\left(d\lambda\right) = 1$. ∎

We are now in a position to state and prove the main result of the paper.

Proposition 1. For any admissible equilibrium, there is a threshold $\hat{p} < 1$ such that the equilibrium is $p - stable$ iff $p > \hat{p}$.

---

[4] See Dudley (1989) for this definition and other properties of the Prohorov metric not explicitly mentioned in this paper.

Proof. There is a neighborhood $N \subset \Delta(A)$ of τ_a^* such that for $\mu-$almost every u in U_A,

$$\forall m \in N, d_A(B(u,m), B(u,\tau_a^*)) \leq K d_{\Delta(A)}(m, \tau_a^*). \tag{1}$$

Now for each $\tau \in T_p$ (with $\tau = p\tau^* + (1-p)\tau'$), we must have that

$$\forall M \subset A, \tau_a(M) = p\tau_a^*(M) + (1-p)\tau_a'(M),$$

It straightforwardly follows from Lemma 1(i) that

$$d_{\Delta(A)}(\tau_a, \tau_a^*) \leq (1-p) d_{\Delta(A)}(\tau_a', \tau_a^*). \tag{2}$$

As the Prohorov metric is always bounded by 1, we have $d_{\Delta(A)}(\tau_a', \tau_a^*) \leq 1$ and $d_{\Delta(A)}(\tau_a, \tau_a^*) \leq 1 - p$. Then, for p large enough, the following property holds: the inequality (1) holds for the marginal τ_a of any distribution τ in T_p. From now on, we consider p such that this property holds.

Define the set $A^n(u) \subset A$ of actions that are best responses of u to a distribution of actions τ_a that is the marginal on A of some $\tau \in S_p^{n-1}$. $\phi(S_p^{n-1})$ contains the distributions $\tau \in T$ such that, for μ-almost every u, $\tau(A^n(u)|u) = 1$.

For μ-almost every u, for every a in $A^n(u)$, a writes $B(u, \tau_a)$ for some τ in S_p^{n-1}. Inequality (1) writes:

$$d_A(a, B(u, \tau_a^*)) \leq K d_{\Delta(A)}(\tau_a, \tau_a^*).$$

As $\tau \in S_p^{n-1} = \phi(S_p^{n-2}) \cap T_p$, $\tau = p\tau^* + (1-p)\tau'$ for some $\tau' \in \phi(S_p^{n-2})$. Inequality (2) implies

$$d_A(a, B(u, \tau_a^*)) \leq K(1-p) d_{\Delta(A)}(\tau_a', \tau_a^*).$$

Denote $R_A^n(u) = \sup_{a \in A^n(u)} d_A(a, B(u, \tau_a^*))$ (this is the radius of the smallest ball containing $A^n(u)$ and centered on τ_A^*). We have

$$R_A^n(u) \leq K(1-p) d_{\Delta(A)}(\tau_a', \tau_a^*).$$

Denote $R_A^n = \sup \mathbf{ess}_{u \in U_A} R_A^n(u)$ for every n. We have

$$R_A^n \leq K(1-p) d_{\Delta(A)}(\tau_a', \tau_a^*). \tag{3}$$

Consider now that by definition, for every M, $\tau_a^*(M) = \int \tau^*(M|u)\mu(du)$. By admissibility of τ^*, for $\mu-$almost every u, the conditional distribution $\tau^*(.|u)$ is a Dirac measure $\delta_{B(\tau_a^*, u)}$ on the equilibrium action of u (denoted $B(\tau_a^*, u)$). By Lemma 1(ii), for every Dirac measure centered on $x \in A$,

$$d_{\Delta(A)}(\tau_a', \delta_x) \leq \sup_{y \in S} d_A(x, y), \tag{4}$$

where S is the support of τ_a' (the smallest closed set such that $\tau_a'(S) = 1$). Then, we have:

$$d_{\Delta(A)}(\tau_a', \delta_{B(\tau_a^*, u)}) \leq R_A^{n-2}(u)$$

7

As $\tau_a^* = \int \delta_{B(\tau_a^*, u)} \mu\, (du)$, by Lemma 1(iii),

$$d_{\Delta(A)}(\tau_a', \tau_a^*) \leq \sup_{u \in U_A} \mathsf{ess}\ d_{\Delta(A)}(\tau_a', \delta_{B(\tau_a^*, u)}) \leq \sup_{u \in U_A} \mathsf{ess}\ R_A^{n-2}\,(u) = R_A^{n-2}.$$

From Inequality (3), we have

$$R_A^n \leq K\,(1-p)\,R_A^{n-2}.$$

Hence, for p large enough, $K\,(1-p) < 1$ and the sequence of R_A^n tends to 0, which implies that S_p^n tends to $\{\tau^*\}$. We have shown $p-$stability for $p < 1$ large enough. Existence of the threshold follows from Remark 4 above stating that $p-$stability implies $p'-$stability for every $p' > p$. ∎

Heuristically, the idea underlying the proof is as follows. An equilibrium τ^* is $p-$stable if the best-response map, restricted to $p-$beliefs, generating the sequence of sets S_p^n is a contraction. For p close to one, when the equilibrium τ^* is admissible, we show that the best-response map, restricted to $p-$beliefs, cannot vary much (i.e. in the smooth case, the derivative of the best-response map, restricted to $p-$beliefs, is small). If, on the contrary, the equilibrium τ^* isn't admissible, even when restricted to $p-$beliefs, it can change dramatically around the equilibrium implying that the preceding step of the argument doesn't hold.

2.4   p-stability in the smooth one-dimensional case

To get an intuitive feel for the notion of stability being studied in this paper, consider a simple, smooth model of strategic interaction where there is a continuum of agents each whom chooses an action $a \in A$ (a compact set in $\mathbb{R}$) to maximize $u\,(a, \bar{a})$ (C^2, with $u_{aa}'' < 0$) where $\bar{a}$ is the average action. Without loss of generality, $A = [-1, 1]$. Suppose, there is a (not necessarily) unique Nash that is interior and is normalized to 0 so that $u_a'\,(0,0) = 0$. Denote $BR\,(\bar{a})$ the (unique) best response to $\bar{a}$ (characterized by $u_a'\,(BR\,(\bar{a})\,, \bar{a}) = 0$). We assume that the BR map is not vertical at equilibrium ($BR'\,(0) < +\infty$).

We are now in a position to state the following result:

Proposition 2.   There is $\hat{p} < 1$ such that the equilibrium is p-stable iff $p > \hat{p}$. If $M = \sup_{\bar{a} \in [-1,1]} |BR'\,(\bar{a})| < 1$, then $\hat{p} = 0$. Otherwise, we have:

$$1 - \frac{1}{|BR'\,(0)|} < \hat{p} < 1 - \frac{m}{M + m - 1}, \tag{5}$$

where

$$m = \inf_{a, \bar{a} \in [-1,1]} \frac{u_{a\bar{a}}''\,(a, 0)}{u_{a\bar{a}}''\,(a, \bar{a})} \in [0, 1].$$

It follows that our stability concept gives a motivation for looking at the slope of the best response map as a stability index. In particular, if $u_{a\bar{a}}''$ is

constant in $\bar{a}$, then $m = 1$ and the inequalities (5) become:

$$1 - \frac{1}{|BR'(0)|} < \hat{p} < 1 - \frac{1}{M}. \tag{6}$$

To relate $\hat{p}$ with exogenous variables, use the inequalities (5) and notice that the implicit functions theorem implies $BR'(0) = -u''_{a\bar{a}}(0,0)/u''_{aa}(0,0)$ and

$$M \leq \sup_{a,\bar{a}\in[-1,1]} \left| \frac{u''_{a\bar{a}}(a,\bar{a})}{u''_{aa}(a,\bar{a})} \right|$$

A special case of the model studied so far is the Muth model with a large number of farmers who have to commit to an output level before selling their products in a competitive market in Guesnerie (1992). Farmer i maximizes $\pi q - \frac{q^2}{2C^i}$ (π is the output price). Aggregate supply in this market is given by $S(\pi) = C\pi$ where $C = \int C^i di$. Aggregate demand in this market is:

$$D(\pi) = \begin{cases} A - B\pi \ if \ \pi \leq \frac{A}{B} \\ 0, \ otherwise \end{cases}$$

Let π^* be the competitive equilibrium price. Guesnerie (1992) shows that when the slope of the best response map $B/C < 1$, π^* is the unique rationalizable outcome.

Applying Proposition 2 immediately yields that the equilibrium in Guesnerie's model is $p-$stable iff $p > \max\{1 - C/B, 0\}$. Thus, $p-$stability describes more precisely the degree of stability of the equilibrium when it is not the unique rationalizable outcome.

The remainder of the section is devoted to the proof of Proposition 2. The proof shows that $p-$stability relies on the best response map $BR(p,\bar{a})$ (best response to beliefs "p on 0, $(1-p)$on $\bar{a}$"). When p is close to one, the slope $BR'_{\bar{a}}(p,\bar{a})$ is small enough (whatever a is). The map $BR(p,.)$ is then globally contracting and $p-$stability obtains. Intuitively, when p is close to one, the best response is not very sensible to the value $\bar{a}$ and the best response cannot deviate very much from the equilibrium value 0. This is the condition needed to get $p-$stability.

Proof. We first give some notation. For every $\bar{a}$ in $[-1,1]$, the best response $BR(p,\bar{a})$ to beliefs "p on 0, $(1-p)$on $\bar{a}$" solves:

$$\max_a pu(a,0) + (1-p)u(a,\bar{a}).$$

With the notation of the previous section, an element τ in T is such that τ_u is a Dirac measure on u. Then, τ is characterized by a distribution on A (that is τ_a). With a slight abuse of notation, we identify an element τ in T with its marginal τ_a on A.

We first check the following simple lemma:

Lemma 2. Consider an interval of actions $[a_-, a_+]$ $(0 \in [a_-, a_+])$, an action that is a best response to some beliefs on $[a_-, a_+]$ putting at least probability p on 0 is an action in the interval $[a'_-, a'_+]$ where

$$a'_- = \inf_{\bar{a} \in [a_-, a_+]} BR(p, \bar{a}) \text{ and } a'_+ = \sup_{\bar{a} \in [a_-, a_+]} BR(p, \bar{a}),$$

Proof of the lemma. The best response a of a player to belief on $[a_-, a_+]$ putting at least probability p on 0 solves a FOC

$$pu'_a(a, 0) + (1 - p) \int u'_a(a, \bar{a}) \, dP(\bar{a}) = 0,$$

where dP is some Borel measure on $[a_-, a_+]$. Notice that the LHS of this FOC is an integral over the family of functions $pu'_a(a, 0) + (1 - p) u'_a(a, \bar{a})$ (indexed by $\bar{a}$). Furthermore, $BR(p, \bar{a})$ is characterized as the solution of

$$pu'_a(a, 0) + (1 - p) u'_a(a, \bar{a}) = 0.$$

The lemma follows. ∎

We are now in a position to define the sequence of sets S_p^n. To this purpose, denote $a_-^0 = -1$ and $a_+^0 = +1$ and, for every $n \geq 1$, define iteratively the values a_-^n and a_+^n in $[-1, 1]$ by

$$\forall n \geq 1, a_-^n = \inf_{\bar{a} \in [a_-^{n-1}, a_+^{n-1}]} BR(p, \bar{a}) \text{ and } a_+^n = \sup_{\bar{a} \in [a_-^{n-1}, a_+^{n-1}]} BR(p, \bar{a}),$$

(clearly, $0 \in [a_-^n, a_+^n]$ and $[a_-^n, a_+^n] \subset [a_-^{n-1}, a_+^{n-1}]$ for every n).

- T_p (that is S_p^0) is the set of distributions on $[a_-^0, a_+^0]$ putting at least probability p on 0

- An action that is a best response to some beliefs in S_p^0 is an action in $[a_-^1, a_+^1]$ (from the Lemma above)

- As every player is rational and has beliefs in S_p^0, the aggregate action is in $[a_-^1, a_+^1]$. Hence, $\phi(S_p^0)$ is the set of distributions on $[a_-^1, a_+^1]$.

- $S_p^1 = \phi(S_p^0) \cap S_p^0$ is the set of distributions on $[a_-^1, a_+^1]$ putting at least probability p on 0.

A comment about this argument: the key point here is that p has 2 effects on the transition between S_p^{n-1} and S_p^n: the "straight" effect that S_p^n is a set of distributions on a subset X of actions putting at least probability p on one specific action (the equilibrium), and the other effect (on which the iterative contraction argument relies), that the support X on the distributions in S_p^n shrinks with p (X decreases in p, for a given size of the support of S_p^{n-1}).

We now iterate the argument:

10

- If S_p^{n-1} is the set of distributions on $\left[a_-^{n-1}, a_+^{n-1}\right]$ putting at least probability p on 0, then an action that is a best response to some beliefs in S_p^{n-1} is an action in $\left[a_-^n, a_+^n\right]$ (from the Lemma above)

- $\phi\left(S_p^{n-1}\right)$ is then the set of distributions on $\left[a_-^n, a_+^n\right]$

- $S_p^n = \phi\left(S_p^{n-1}\right) \cap S_p^0$ is the set of distributions on $\left[a_-^n, a_+^n\right]$ putting at least probability p on 0.

We now characterize the conditions implying that S_p^∞ reduces to the equilibrium. As $\left[a_-^n, a_+^n\right] = BR\left(p, \left[a_-^{n-1}, a_+^{n-1}\right]\right)$, S_p^∞ reduces to the equilibrium iff the two sequences a_-^n and a_+^n converge to 0. A necessary condition for convergence of a_-^n and a_+^n is that $BR(p,.)$ is locally contracting at 0, that is:

$$|BR_{\bar{a}}'(p,0)| < 1.$$

A sufficient condition is that

$$\forall \bar{a} \in [-1,1], |BR_{\bar{a}}'(p,\bar{a})| < 1.$$

By the implicit function theorem, we have:

$$BR_{\bar{a}}'(p,\bar{a}) = -\frac{(1-p)\,u_{a\bar{a}}''\left(BR\left(p,\bar{a}\right),\bar{a}\right)}{p u_{aa}''\left(BR\left(p,\bar{a}\right),0\right) + (1-p)\,u_{aa}''\left(BR\left(p,\bar{a}\right),\bar{a}\right)}$$

Then, $|BR_{\bar{a}}'(p,0)| < 1$ writes $(BR(p,0) = 0)$

$$1 - \left|\frac{u_{aa}''(0,0)}{u_{a\bar{a}}''(0,0)}\right| < p,$$

or, equivalently:

$$p > 1 - \frac{1}{|BR_{\bar{a}}'(0,0)|}. \tag{7}$$

On the other hand, we have:

$$BR_{\bar{a}}'(p,\bar{a}) = -\frac{u_{a\bar{a}}''\left(BR\left(p,\bar{a}\right),\bar{a}\right)}{u_{aa}''\left(BR\left(p,\bar{a}\right),\bar{a}\right)}\frac{1-p}{p\frac{u_{aa}''(BR(p,\bar{a}),0)}{u_{aa}''(BR(p,\bar{a}),\bar{a})} + 1 - p},$$

$$= BR_{\bar{a}}'(0,\bar{a})\frac{1-p}{p\frac{u_{aa}''(BR(p,\bar{a}),0)}{u_{aa}''(BR(p,\bar{a}),\bar{a})} + 1 - p}.$$

We then have:

$$\frac{1-p}{p\frac{u_{aa}''(BR(p,\bar{a}),0)}{u_{aa}''(BR(p,\bar{a}),\bar{a})} + 1 - p} \leq \frac{1-p}{pm + 1 - p},$$

$$|BR_{\bar{a}}'(p,\bar{a})| \leq |BR_{\bar{a}}'(0,\bar{a})|\frac{1-p}{pm + 1 - p},$$

11

and $\sup_{\bar{a} \in [-1,1]} |BR'_{\bar{a}}(p,a)| < 1$ is implied by:

$$M \frac{1-p}{pm+1-p} < 1.$$

If $M < 1$, then this inequality holds. Otherwise, it rewrites:

$$p > 1 - \frac{m}{m+M-1}. \tag{8}$$

The existence of $\hat{p}$ is shown in Proposition 1. Inequalities (7) and (8) imply the result.∎

3 Epistemic conditions

In this section, we give epistemic conditions for a p-stable equilibrium. We exhibit (common knowledge, CK hereafter) assumptions implying that the outcome of the game is a p-stable equilibrium. These conditions provide some epistemic motivation for our stability concept. Examples in the literature of such assumptions include CK of rationality implies that the outcome is rationalizable (Tan and Werlang 1988), mutual knowledge of the payoff functions and of rationality, and common knowledge of the conjectures imply Nash (Aumann and Brandenburger 1995) and common p-belief of rationality implies p-rationalizability (Hu 2007).

We show that, for a given equilibrium τ^*, $\phi(S_p^\infty)$ is the set of outcomes compatible with CK of p-belief of τ^* and CK of rationality. Hence, p-stability of an admissible equilibrium τ^* means that τ^* is the only outcome compatible with these CK assumptions.

In order to define events like "CK of rationality", we consider a universal beliefs space Ω (Mertens and Zamir 1985). This is a formal model where a state of the world specifies everything: the state of nature (payoffs, actions...) and players' beliefs (summarized by the type of each player). Hence, events like "everyone is rational" can be defined in Ω.

In our framework, a state of the world $\omega \in \Omega$ must specify the following two items:

- the distributions of actions played (an element in T)

- the type of each player. The type of a player specifies a probability distribution on T and others' types (a player is assumed to know his own type, not the others' ones). Hence, the type of a player with utility u is a belief on Ω, that is an element in $\Delta(\Omega)$.

So that we must have:

$$\Omega = T \times \Delta(\Omega)^{U_A(\mu)},$$

where $U_A(\mu)$ is the support of the distribution μ ($U_A(\mu)$ is a subset of the set U_A of utility functions).

The question of the existence of such a universal beliefs space Ω is not a trivial one. Mertens and Zamir (1985) give a positive answer in a framework with finitely many players. We assume existence of the universal beliefs space Ω in our framework[5].

We now define "events" (subsets of Ω) like "everyone is rational" or "CK of everyone being rational".

We first define CK of an event E. For any event $E \subseteq \Omega$, the event "everyone knows E" is defined by: μ-almost everywhere, σ_u puts probability 1 on E, that is:

$$K(E) = \{\omega = (\tau, \sigma) \in \Omega : \mu(\{u \in U_A(\mu) : \sigma_u(E) = 1\}) = 1\}.$$

where σ_u is u's (i.e. a player located at u) beliefs (a probability distribution on Ω). Higher orders of knowledge of E are then iteratively defined by:

$$\forall n \geq 1, K^{n+1}(E) = K(K^n(E)).$$

Under the assumption $K(E) \subseteq E$ ("no one knows E when E has not occurred" - this assumption is implicit in the definition of Ω), the sequence $K^n(E)$ is decreasing. Hence, it converges to a limit set $K^\infty(E) = \cap_{n \geq 1} K^n(E)$. Common knowledge of E is then defined as the event $K^\infty(E)$.

We now define CK of rationality. To this purpose, we first define the event "everyone is rational" using a BR correspondence extended to the case of heterogenous beliefs. This BR correspondence is

$$\Phi : \Delta(\Omega)^{U_A(\mu)} \rightrightarrows T,$$

where for each $\sigma = (\sigma_u)_{u \in U_A(\mu)} \in \Delta(\Omega)^{U_A(\mu)}$,

$$\tau \in \Phi(\sigma) \Leftrightarrow \mu(\{u \in U_A(\mu) : \tau(B(u, \sigma_{u,A})|u) = 1\}) = 1.$$

In this expression, $\sigma_{u,A}$ is the marginal distribution on the space A of actions ($\sigma_{u,A}$ is an element of $\Delta(A)$). Thus, $\tau \in \Phi(\sigma)$ puts probability 1 on the fact that u plays a BR to his belief $\sigma_{u,A}$ about actions.

The event that everybody is rational is described by the set

$$R_0 = \{\omega = (\tau, \sigma) \in \Omega : \tau \in \Phi(\sigma)\}.$$

Higher order knowledge of rationality is described by the sequence of sets $R_n = K^n(R_0)$ for $n \geq 1$. $R_\infty = K^\infty(R_0)$ is the set of states where rationality is common knowledge.

Next we define common knowledge of $p-$belief of τ^*. Let $B_0^p \subseteq \Omega$ be defined as follows:

$$B_0^p = \{\omega = (\tau, \sigma) \in \Omega : \mu(\{u \in U_A(\mu) : \sigma_{u,A}(\tau_a^*) \geq p\}) = 1\}.$$

B_0^p contains the states where the set of u such that $\sigma_{u,A}(\tau_a^*) \geq p$ has μ-measure 1. Higher order knowledge of $p-$belief of τ^* is described by the sequence of sets $B_n^p = K^n(B_0^p)$ for $n \geq 1$. $B_\infty^p = K^\infty(B_0^p)$ is the set of states where $p-$belief of τ^* is common knowledge.

We now state the main result of the section:

Proposition 3. $proj._T(R_\infty \cap B_\infty^p) = \phi(S_p^\infty)$ (that is: the projection on T of the set $R_\infty \cap B_\infty^p$ is $\phi(S_p^\infty)$). Further, when τ^* is admissible, $\{\tau^*\} = proj._T(R_\infty \cap B_\infty^p)$ is equivalent to $p-$stability of τ^*.

Proof. We proceed iteratively by relating the sets S_p^n to some knowledge assumptions. Recall that, for any events E and F, we have $K(E \cap F) = K(E) \cap K(F)$.

The actions profiles associated with the states in $B_0^p \cap R_0$(that is: everyone $p-$believes in τ^* and everyone is rational) are the actions profiles in $\phi(S_p^0)$.

We iterate the argument. Assume that the actions profiles associated with the states in $X_{n-1} \overset{def}{=} \cap_{0 \leq k \leq n-1} K^k(B_0^p \cap R_0)$ are the actions profiles in $\phi(S_p^{n-1})$. At a state in $B_0^p \cap K(X_{n-1})$ (that is: everyone $p-$believes in τ^* and everyone knows that the action profile is in $\phi(S_p^{n-1})$), players' beliefs are on $S_p^n = S_p^0 \cap \phi(S_p^{n-1})$. Then, the actions profiles associated with the states in $R_0 \cap B_0^p \cap K(X_{n-1})$ (that is: everyone is rational and everyone knows that the action profile is in S_p^n) are the actions profiles in $\phi(S_p^n)$. Notice that $R_0 \cap B_0^p \cap K(X_{n-1}) = X_n$.

Taking the limit on n, it follows that the actions profiles associated with the states of nature in $X_\infty = B_\infty^p \cap R_\infty$ are the actions profiles in $\phi(S_p^\infty)$. Then, $proj._T(R_\infty \cap B_\infty^p) = \phi(S_p^\infty)$.

To prove the last part of the proposition, note that $S_p^\infty = \phi(S_p^\infty) \cap S_p^0$: the outcomes in S_p^∞ are the profiles of actions that are best responses to beliefs on S_p^∞ and that put probability at least p on τ^*. If $proj._T(R_\infty \cap B_\infty^p) = \{\tau^*\}$, then $\phi(S_p^\infty) = \{\tau^*\}$ and $S_p^\infty = \{\tau^*\}$: the equilibrium is $p-$stable. Conversely, if the equilibrium is $p-$stable, then $proj._T(R_\infty \cap B_\infty^p) = \phi(S_p^\infty) = \phi(\{\tau^*\})$. By definition, when τ^* is admissible, $\phi(\{\tau^*\}) = \{\tau^*\}$. It follows that $proj._T(R_\infty \cap B_\infty^p) = \{\tau^*\}$. ∎

4 How does "fundamental" uncertainty affect strategic stability?

The example below considers a "contagion" argument to build an incomplete information game from a collection of complete information games where the equilibrium is sometimes, but not always dominant solvable. This means that we add payoff uncertainty (and informational asymmetry) to a complete information game. We show that the degree of $p-$stability is bigger in the incomplete information game than in the complete information game: *adding uncertainty on fundamentals unambiguously decreases strategic stability.* This results holds

in an example with strategic complementarities and in an example with strategic substitutabilities.

The "contagion" (or "infection") argument[6] is as follows. Consider a game (with a given set of actions for each player). There are many states of nature (affecting the payoffs of players). Under complete information, there is a dominant action in some states, and no dominant in other states. The question is: what should players do in these states with no dominant action? The "global games" answer relies on viewing the "complete information" case as a limit of the "incomplete information" case and showing that for some information structures, the equilibrium is dominant solvable in the incomplete information game. The proof of this dominant solvability property is precisely what is called the "contagion" argument: in some states, players knows "enough" about the state to know that the action that was dominant in the complete information game is still dominant in other states ("near" the states above) with this argument iterated to extend to all the states so that players' behavior is uniquely determined in every state (typically, players play in every state the action that is dominant in the states mentioned in the first step of the "contagion" argument).

The example below is an example of a "contagion" argument in a case with a continuum of actions and states. The game is the generic case where players are homogenous, players play against the mean action of others, and best responses are linear and vary continuously in the mean action. The information structure is the quite common case where each player privately observes a noisy signal of the true state (variables are normally distributed).

What makes the example more specific is that we consider a restricted strategy set in the incomplete information game. This choice is made to reduce the dimensionality of the game and has no other motivation than analytical tractability.

Lastly, the "contagion" argument is not the same as usual. In the example, there is no state where the equilibrium action is dominant under complete information. The "contagion" argument relies on iterated dominance and not on dominance only. Furthermore, we do not look at dominance solvability of the equilibria in the complete and incomplete information games, we relate $p-$stability of these equilibria.

4.1   The example

Consider a game where:

- the players are $i \in [0, 1]$,

- the state of nature is θ with mean $\bar{\theta}$ and precision τ_θ (all the stochastic variables are normally distributed on $I\!R$),

---

[6]See, for instance, Carlson and Van Damme for the seminal paper, Morris, Rob and Shin (1995), Morris and Shin (1998) amongst many others for an example in a case with a continuum of states

- the private signal of i is $s_i = \theta + \varepsilon_i$ (ε_i has mean 0 and precision τ, the ε_i are i.i.d.),

- the action of i is $x_i \in \mathbb{R}$, the mean action is denoted X (denote $X = \int x_i di$),

- a strategy of i is a function $x_i(s_i)$,

- the utility of i is

$$u(x_i, X, \theta) = e^\theta X x_i - \frac{1}{2}x_i^2,$$

The strategy profile $x_i(s_i) = 0$ is an equilibrium of this game.

For analytical tractability, we restrict attention to the strategies $x_i(s_i) = a_i e^{\frac{\tau}{\tau_\theta}s_i}$ where a_i is a real parameter in an interval A. The strategy of i is then characterized by a_i, and the strategy set is A. W.l.o.g. assume $A = [-1, 1]$.

We stress the 3 properties that allows us to solve the model: (i) the strategy set contains the equilibrium strategy $x_i(s_i) = 0$; (ii) it is self-fulfilling (as shown below): the best response to a strategy profile in this set is a strategy in this set, (iii) it follows that the equilibrium is defined as a fixed point of a one-dimensional temporary equilibrium map (the map $\mathcal{T}$ below) and stability is defined/characterized by the condition $|\mathcal{T}'| < 1$.

When the strategy profile is $(a_i)_{i\in[0,1]}$ and the private signals are $(s_i)_{i\in[0,1]}$, the mean action is defined by:

$$\int x_i(s_i)\,di = \int x_i(\theta + \varepsilon)\,dP(\varepsilon),$$

where P is the c.d.f. of a centered normal distribution with precision τ. This mean action writes then

$$X(\theta) \stackrel{def}{=} \bar{a}\int e^{\frac{\tau}{\tau_\theta}(\theta + \varepsilon)}dP(\varepsilon),$$

where $\bar{a} = \int a_i di$. Standard computations (concerning log-normal distributions)[7] imply:

$$X(\theta) = \bar{a}e^{\frac{\tau}{\tau_\theta}\left(\theta + \frac{1}{2\tau_\theta}\right)}. \tag{9}$$

We now compute the best response map. This is purely routine (because the stochastic variables are normally distributed):

Lemma 3. The best response $x_i(s_i)$ to a strategy profile $(a_i)_{i\in[0,1]}$ (with $\bar{a} = \int a_i di$) is:

$$x_i(s_i) = \mathcal{T}(\bar{a})\,e^{\frac{\tau}{\tau_\theta}s_i},$$

---

[7] Recall the mean and variance of a log-normal distribution: the stochastic variable e^x (where x is a normally distributed variable with mean μ and variance σ^2) has a mean $e^{\mu + \frac{\sigma^2}{2}}$ and a variance
$e^{2\mu + \sigma^2}(e^{\sigma^2} - 1)$.

where $\mathcal{T}(\bar{a})$ is a best response map defined by

$$\mathcal{T}(\bar{a}) \overset{def}{=} \bar{a}e^{\bar{\theta}+\frac{2\tau+\tau_\theta}{2\tau_\theta^2}}.$$

Proof. Equation (9) states that the mean action is

$$X(\theta) = \bar{a}e^{\frac{\tau}{\tau_\theta}\left(\theta+\frac{1}{2\tau_\theta}\right)}.$$

We then have:

$$x_i(s_i) = E\left(e^\theta X(\theta)\,|s_i\right) = \bar{a}e^{\frac{\tau}{\tau_\theta}\left(\frac{1}{2\tau_\theta}\right)}\int e^{\left(1+\frac{\tau}{\tau_\theta}\right)\theta}dP(\theta|s_i),$$

where $P(\theta|s_i)$ is the c.d.f. of the law of θ conditional to s_i (that is a normal distribution with mean $E = \frac{\tau_\theta\bar{\theta}+\tau s_i}{\tau+\tau_\theta}$ and precision $T = \tau + \tau_\theta$). Standard computations (concerning log-normal distributions again) imply:

$$x_i(s_i) = \bar{a}e^{\frac{\tau}{\tau_\theta}\left(\frac{1}{2\tau_\theta}\right)}e^{\left(1+\frac{\tau}{\tau_\theta}\right)E+\frac{\left(1+\frac{\tau}{\tau_\theta}\right)^2}{2T}},$$

$$x_i(s_i) = \bar{a}e^{\bar{\theta}+\frac{2\tau+\tau_\theta}{2\tau_\theta^2}+\frac{\tau}{\tau_\theta}s_i}.$$

The result follows. ∎

Notice that the equilibrium corresponds to the fixed point of $\mathcal{T}$ (that is $\bar{a} = 0$, that corresponds to a strategy profile: $a_i = 0$ for every i).

From a formal viewpoint, this game is an example of the smooth one-dimensional case studied in Section 2.4. From Section 2.4, we know that a characterization of $p-$stability relies on a best response map $BR(p, \bar{a})$ (see the proof of Proposition 2). The example under consideration here appears to be fully tractable and we get a more complete result than Proposition 2: we fully characterize the degree of $p-$stability of the equilibrium.

Proposition 4. Under the assumption that the strategy set is restricted to the strategies $x_i(s_i) = a_i e^{s_i}$, the equilibrium $x_i(s_i) = 0$ is p-stable iff $p \geq \hat{p}$, where:

$$\hat{p} = 1 - e^{-\bar{\theta}-\frac{1}{\tau_\theta}\left(\frac{1}{2}+\frac{\tau}{\tau_\theta}\right)} < 1.$$

In particular, the equilibrium is 0-stable iff

$$\bar{\theta} < -\frac{1}{\tau_\theta}\left(\frac{1}{2}+\frac{\tau}{\tau_\theta}\right).$$

The proof elaborates on the proof of Proposition 2. Notice that p^* is increasing in $\bar{\theta}$, τ and decreasing in τ_θ: stability is favored by small $\bar{\theta}$, small τ and large τ_θ. Notice that (i) small $\bar{\theta}$ means that the event $e^\theta < 1$ has a large

17

probability; (ii) a small τ means large informational asymmetries; (iii) a large τ_θ means a small uncertainty on θ; (iv) there is no uniquely defined "complete information" limit $(\lim_{(\tau_\theta,\tau)\to(+\infty,+\infty)} \frac{1}{\tau_\theta}\left(\frac{1}{2}+\frac{\tau}{\tau_\theta}\right)$ is not well defined).

Proof. $BR(p,\bar{a})$ is the best response of player i to belief "probability p on 0 and probability $(1-p)$ on $\bar{a}$". Straightforwardly, the best response of i to such belief is:

$$x_i(s_i) = pE\left(e^\theta X^*(\theta)\,|s_i\right) + (1-p)\,E\left(e^\theta X(\theta)\,|s_i\right),$$

where $X^*(\theta)$ is the equilibrium mean action (that is equal to 0) and $X(\theta)$ is defined as in Equation (9). Then, we have from Lemma 1

$$x_i(s_i) = (1-p)\,E\left(e^\theta X(\theta)\,|s_i\right) = (1-p)\,\mathcal{T}(\bar{a})\,e^{\frac{\tau}{\tau_\theta}s_i},$$

and it follows that

$$BR(p,\bar{a}) = (1-p)\,\mathcal{T}(\bar{a}).$$

As $\mathcal{T}$ is linear in $\bar{a}$, we conclude that the sequences a_-^n and a_+^n (defined in the proof of Proposition 2) converge to 0 iff $|BR'(p,\bar{a})| < 1$, that is

$$(1-p)\,e^{\bar{\theta}+\frac{2\tau+\tau_\theta}{2\tau_\theta^2}} < 1.$$

The value of $\hat{p}$ follows. $\blacksquare$

4.2 Link with the complete information game

In the complete information game associated with a state θ, the best response of i to a mean action X is:

$$x_i = e^\theta X.$$

It is straightforward that the equilibrium is $x^* = 0$ and 0-stability is equivalent to $e^\theta < 1$ (that is: $\theta < 0$). For $\theta > 0$, the equilibrium is not $0-$stable, but $p-$stability can be considered.

As this game is an example of the one-dimensional smooth case considered in Section 2.4, we know that the analysis of $p-$stability relies on the best response of i to p-beliefs assigning probability p to equilibrium $X^* = 0$ and probability $(1-p)$ to some other mean strategy X. This best response is:

$$x_i = e^\theta\left(pX^* + (1-p)\,X\right).$$

Considering the analysis made in Section 2.4 immediately shows that the equilibrium is $p-$stable iff $\left|\frac{dx_i}{dX}\right| < 1$. This condition writes $(1-p)\,e^\theta < 1$, that is: $p \geq \hat{p}(\theta)$, with:

$$\hat{p}(\theta) = 1 - e^{-\theta}.$$

$\hat{p}(\theta)$ is increasing in θ.

We have seen above that, if we introduce some asymmetric uncertainty (as defined above) on θ (this uncertainty being centered on a certain value $\bar{\theta}$), then the p-stability of the equilibrium of the incomplete information game is characterized by $p \geq \hat{p}$, with:

$$\hat{p}\left(\bar{\theta}, \tau_\theta, \tau\right) = 1 - e^{-\bar{\theta} - \frac{1}{\tau_\theta}\left(\frac{1}{2} + \frac{\tau}{\tau_\theta}\right)}.$$

Note that $e^{-\frac{1}{\tau_\theta}\left(\frac{1}{2} + \frac{\tau}{\tau_\theta}\right)} < 1$. It follows that

Proposition 5. Adding uncertainty on fundamentals unambiguously decreases strategic stability: for every θ, τ_θ and τ

$$\hat{p}\left(\bar{\theta}\right) \leq \hat{p}\left(\bar{\theta}, \tau_\theta, \tau\right).$$

Remark. No equilibrium in the incomplete information game is $0-$stable when $\bar{\theta} > 0$. In other words, it is impossible to make an unstable equilibrium (the equilibrium in the complete information game when $\theta > 0$) 0-stable by adding some uncertainty centered on this θ. To make this equilibrium 0-stable, the uncertainty on θ must be centered on a negative value (there must be some kind of a bias in the prior beliefs on θ).

4.3 The variant with strategic substitutes

The above example displays some kind of strategic complementarities. For example, in the complete information game associated with the state θ, the BR of a player i is $x_i = e^\theta X$: it is increasing in others' mean action X.

We now show that the stability result is not affected if the slope of the BR map is reversed. That is: we consider a variant of the above game where the BR to a mean strategy $X(\theta)$ of a player i observing s_i is:

$$x_i\left(s_i\right) = -E\left(e^\theta X\left(\theta\right) | s_i\right).$$

Everything is thus exactly identical to the previous game, except the sign "-" in the BR map. Straightforwardly, a best response to a mean strategy $X(\theta) = \bar{a}e^{\frac{\tau}{\tau_\theta}\left(\theta + \frac{1}{2\tau_\theta}\right)}$ (as defined in Equation (9)) is:

$$x_i\left(s_i\right) = -\mathcal{T}\left(\bar{a}\right) e^{\frac{\tau}{\tau_\theta} s_i},$$

It follows that the unique equilibrium is still 0. The game is, as the initial one, formally identical to a one-dimensional smooth example and it is easily checked that the $p-$stability condition is the same as the one in the initial game.

An intuition for this common result is the following. What we have done here is an example of an "infection" (or contagion) argument. This has nothing to do with the complements/substitutes question. The idea is that: (i) there is a subset of states where it is "CK enough" that $e^\theta < 1$ to get stability in this subset; (ii) the information structure is such that there is another subset of states where (i) is "CK enough" and where every player puts a large enough

probability on being in (i) (that is: to be in a stable state) to get stability in this other subset; (iii) there is again another subset of states where (i) and (ii) are "CK enough"... and the sequence of subsets of states defined in all these steps (i), (ii), ... covers the whole set of states of nature.

5 Applications to exchange economies

In this section, we apply our notion of stability in exchange economies. We model exchange non-cooperatively using market games whose noncooperative equilibria coincide with competitive equilibria under certain conditions. In a market game, the map that associates market prices to the actions of all individuals is well-defined both out of equilibrium and at equilibrium. We reformulate the stability problem of competitive equilibria as a coordination problem over the expectations that agents have over market variables, like prices as traders try to use structure of the game to converge back to equilibrium. The study of the resulting coordination dynamics is thus an analysis of the stability of competitive equilibria. We begin by examining the $p-$stability of a unique globally unstable equilibrium under tâtonnement dynamics. We, then, characterize the conditions under which there is a trade-off between enhanced equilibrium efficiency and stability in a simple market setting, with and without trading frictions.

5.1 p-stability of a unique globally unstable equilibrium

The underlying exchange economy is the one studied by Scarf (1960) and the unique competitive equilibrium is globally unstable under tâtonnement dynamics. Scarf's example of global instability of a unique competitive equilibrium under tâtonnement dynamics has three commodities $l = 1, 2, 3$, a continuum of individuals of measure one, and three types of agents (of equal measure), I_1, I_2, I_3 such that each $i \in I_1$ has preferences represented by a Leontieff utility function $\min\{x_1, x_2\}$ with endowments $\mathbf{w}^i = \{1, 0, 0\}$, each $j \in I_2$ has preferences represented by a Leontieff utility function $\min\{x_2, x_3\}$ with endowments $\mathbf{w}^i = \{0, 1, 0\}$, each $k \in I_3$ has preferences represented by a Leontieff utility function $\min\{x_3, x_1\}$ with endowments $\mathbf{w}^i = \{0, 0, 1\}$. The unique competitive equilibrium prices are $(1, 1, 1)$ with allocations $\{\frac{1}{2}, \frac{1}{2}, 0\}$ for each $i \in I_1$, $\{0, \frac{1}{2}, \frac{1}{2}\}$ for each $j \in I_2$ and allocations $\{\frac{1}{2}, 0, \frac{1}{2}\}$ for each $k \in I_3$.

We model non-cooperative exchange using the Shapley window model, studied by Sahi and Yao (1989) in the case with a finite number of traders and Codognato and Ghosal (2000) for the case with a continuum of agents. Applied to Scarf's example, in the Shapley window model, (1) a strategy for agent $i \in I_1$ is a vector $(b_{12}^i, b_{13}^i) \geq (0, 0)$ such that $b_{12}^i + b_{13}^i \leq 1$, (2) a strategy for agent $j \in I_2$ is a vector $(b_{21}^j, b_{23}^j) \geq (0, 0)$ such that $b_{21}^j + b_{23}^j \leq 1$, and (3) a strategy for agent $k \in I_3$ is a vector $(b_{31}^k, b_{32}^k) \geq (0, 0)$ such that $b_{31}^k + b_{32}^k \leq 1$. The interpretation is that $b_{ll'}^h$ is the amount of commodity l bid by agent h in exchange for commodity l'.

Given an assignment of bids $\mathbf{b}$, we compute the 3×3 aggregate bids matrix $\bar{\mathbf{B}}$ with $\bar{\mathbf{b}}_{ll'} = \int_I \mathbf{b}^h_{ll'} dh$ denoting the $ll'th$ component. Given $\mathbf{b}$, and an irreducible aggregate bids matrix $\bar{\mathbf{B}}$, prices π exist if

$$\pi \gg 0, \sum_{l'} \pi_{l'} \bar{\mathbf{b}}_{l'l} = \pi_l \sum_{l'} \bar{\mathbf{b}}_{ll'}, l = 1, 2, 3.$$

Intuitively, the price formation rule states that prices are formed to ensure that the value of what is bid (using commodities l') for each commodity l is equal to the value of what is bid for commodities l' using commodity l. The allocation rule is

$$x^h_l(\pi(\mathbf{s}), \mathbf{s}^n) = \begin{cases} \mathbf{w}^h_l - \sum_{l'} \mathbf{b}^h_{ll'} + \sum_{l \in L} \mathbf{b}^h_{l'l} \frac{\pi_l(\mathbf{s})}{\pi_{l'}(\mathbf{s})}, & if \ \pi \ exists, \\ \mathbf{w}^h_l, & otherwise. \end{cases}$$

for $l = 1, 2, 3$.

Note that if π exists then so does $\lambda \pi$, $\lambda > 0$ and moreover the allocation rule is unchanged if π is substituted for $\lambda \pi$. In what follows, we will adopt the normalization rule that $\pi_1 = 1$.

Codognato and Ghosal (2000) show that as long as the aggregate bids matrix $\bar{\mathbf{B}}$ is irreducible, there exists a unique (up to a scalar multiple) price $\pi(\bar{\mathbf{B}})$. It is straightforward to check that the set of Nash equilibrium prices and allocations coincide with Competitive equilibrium prices and allocations (all the assumptions of Theorem 2 Codognato and Ghosal (2000) are satisfied).

Let $\Pi(\pi^*) \subset I\!R^2_{++}$ be a compact subset of market prices containing a competitive equilibrium price π^*. Although, so far, $p-$stability has been defined directly on the strategies, motivated by the application to competitive markets, it will be convenient to define p-stability over prices. That we can do so without loss of generality is a consequence of the following useful result.

Lemma 4. Let $\Pi(\pi^*) \subset I\!R^2_{++}$ be a compact subset of market prices containing a competitive equilibrium price π^*. Let $\mathbf{b}^n$ be a sequence of assignments of strategies such that $\pi^n = \pi(\bar{\mathbf{B}}^n) \in \Pi(\pi^*)$, $n \geq 1$, with $\lim_{n \to \infty} \pi^n = \pi$. Then, there exists a limit point $\mathbf{b}^n$, $\mathbf{b}$, such that $\pi(\bar{\mathbf{B}}) = \pi$.

Proof. Notice that the sequence of assignments $\mathbf{b}^n$, $n \geq 1$, is uniformly integrable as $\mathbf{b}^n$ is bounded below by the constant 0 and above by the constant 1. By Lemma 1 in Busetto, Codognato and Ghosal (2010), there exists a limit point $\mathbf{b}$ of the sequence of the sequence of assignments $\mathbf{b}^n$, $n \geq 1$, which is itself an assignment. By proceeding to a subsequence (denoted in the same way as the original sequence to save on notation) if necessary, consider $\pi^n = \pi(\mathbf{B}^n)$ and with $\lim_{n \to \infty} \pi^n = \pi$. As $\pi \in \Pi(\pi^*) \subset I\!R^2_{++}$, by Lemma 1 in Sahi and Yao (1989), there exists $\bar{\mathbf{B}}$ the irreducible aggregate bids matrix corresponding to π. The aggregate bids matrix for each element in the sequence $\mathbf{b}^n$, $\bar{\mathbf{B}}^n$, is by construction irreducible. Then, $\bar{\mathbf{B}}$ is the aggregate bids matrix such that each element $\bar{\mathbf{b}}_{ll'} = \lim_{n \to \infty} \int_I \mathbf{b}^{h,n}_{ll'} dh$. By the generalization of Fatou's lemma in Artstein (1979), there exists a limit point $\mathbf{b}$ such that $\bar{\mathbf{b}}_{ll'} = \int_I \mathbf{b}^h_{ll'} dh = \lim_{n \to \infty} \int_I \mathbf{b}^{h,n}_{ll'} dh$, thus completing the proof.∎

Some notation. Let δ_p be the (Borel) probability distribution assigning probability at least p to π^* with $\Delta_p \Pi(\pi^*)$ the corresponding set. We have a definition by means of an iterative process. For any subset of market prices $\Pi_p^0(\pi^*) \subset I\!\!R_{++}^2$ containing a market equilibrium price π^*, for every $p \in [0,1]$, we define a sequence of price sets: for every $n \geq 0$, the set $\Pi_p^{n+1}(\pi^*)$ contains exactly the π in $\Pi_p^n(\pi^*)$ such that π_l is generated by the price formation rule of the market game where each agent is choosing a best-response to some $\delta_p \in \Pi_p^n(\pi^*)$.

The definition is consistent as every $\Pi_p^n(\pi^*)$ contains π^*. The sequence $\Pi_p^n(\pi^*)$ is decreasing and therefore converges to a limit $\Pi_p^\infty(\pi^*) = \cap_{n\geq 0}\Pi_p^n(\pi^*)$. Trivially, every $\Pi_p^\infty(\pi^*)$ contains at least π^* and $\Pi_1^\infty(\pi^*) = \{\pi^*\}$. If $\Pi_p^\infty(\pi^*) = \{\pi^*\}$ for some p, then $\Pi_{p'}^\infty(\pi^*) = \{\pi^*\}$ for every $p' \geq p$. A market equilibrium price $\pi^*$is p-stable iff there is some $p < 1$, such that $\Pi_p^\infty(\pi^*) = \{\pi^*\}$.

Note that in any best-response, for all $i \in I_1$, $\mathbf{b}_{13}^i = 0$, $j \in I_2$, $\mathbf{b}_{21}^j = 0$ and for all $k \in I_3$, $\mathbf{b}_{32}^k = 0$. The equilibrium strategy of $i \in I_1$ is $\mathbf{b}_{12}^{*i} = \frac{1}{2}$. We now show that for p large enough, each agent $i \in I_1$ plays the equilibrium strategy whatever their $p-$belief is.

Consider the case when individual $i \in I_1$ attaches a probability p to π^* and attaches a probability $1 - p$ to some other positive market price $\tilde\pi \neq \pi^* = (1,1,1)$. In this case, $i \in I_1$ solves

$$\max p \left[\min\left\{ 1 - b_{12}^i, b_{12}^i \right\} \right]$$
$$+ (1-p) \left[\min\left\{ 1 - b_{12}^i, b_{12}^i \frac{1}{\tilde\pi_2} \right\} \right]$$
$$s.t.\ 0 \leq b_{12}^i \leq 1.$$

Note that at any optimum, i must choose b_{12}^i to solve either $1 - \mathbf{b}_{12}^i = \mathbf{b}_{12}^i$ or $1 - \mathbf{b}_{12}^i = \mathbf{b}_{12}^i \frac{1}{\tilde\pi_2}$. Suppose, to begin with, i chooses b_{12}^i to solve $1 - \mathbf{b}_{12}^i = \mathbf{b}_{12}^i$ or equivalently sets $\mathbf{b}_{12}^i = \frac{1}{2}$. In this case, when $1 > \frac{1}{\tilde\pi_2}$, $i's$ payoff is $p\left(\frac{1}{2}\right) + (1-p)\left(\frac{\frac{1}{\tilde\pi_2}}{2}\right)$ while if $1 < \frac{1}{\tilde\pi_2}$, $i's$ payoff is $\frac{1}{2}$. Now suppose i chooses b_{12}^i to solve $1 - \mathbf{b}_{12}^i = \mathbf{b}_{12}^i \frac{1}{\tilde\pi_2}$ or equivalently $\mathbf{b}_{12}^i = \frac{1}{1+\frac{1}{\tilde\pi_2}}$. In this case, when $1 > \frac{1}{\tilde\pi_2}$, $i's$ payoff is $\frac{\frac{1}{\tilde\pi_2}}{1+\frac{1}{\tilde\pi_2}}$, while if $1 < \frac{1}{\tilde\pi_2}$, $i's$ payoff is $p\left(\frac{1}{1+\frac{1}{\tilde\pi_2}}\right) + (1-p)\left(\frac{\frac{1}{\tilde\pi_2}}{1+\frac{1}{\tilde\pi_2}}\right)$.

When $1 > \frac{1}{\tilde\pi_2}$, if $p = 1$, as $\frac{1}{2} > \frac{\frac{1}{\tilde\pi_2}}{1+\frac{1}{\tilde\pi_2}}$, as long as p is close to one, $\mathbf{b}_{12}^i = \frac{1}{1+\frac{1}{\pi_2^*}}$ while when $1 < \frac{1}{\tilde\pi_2}$if $p = 1$, as $\frac{1}{2} > \frac{1}{1+\frac{1}{\tilde\pi_2}}$ and therefore, as long as p is close to one, $\mathbf{b}_{12}^i = \frac{1}{2}$.

Next, we compute precise bounds on p. Consider $i \in I_1$. When $1 > \frac{1}{\tilde\pi_2}$, to ensure that $\mathbf{b}_{12}^i = \frac{1}{2}$, the inequality $p\left(\frac{1}{2}\right) + (1-p)\left(\frac{\frac{1}{\tilde\pi_2}}{2}\right) \geq \frac{\frac{1}{\tilde\pi_2}}{1+\frac{1}{\tilde\pi_2}}$ needs to be satisfied, which by computation is equivalent to $p \geq \frac{\frac{1}{\tilde\pi_2}}{1+\frac{1}{\tilde\pi_2}}$. When $1 < \frac{1}{\tilde\pi_2}$, to ensure that $\mathbf{b}_{12}^i = \frac{1}{2}$, the inequality $\frac{1}{2} \geq p\left(\frac{1}{1+\frac{1}{\tilde\pi_2}}\right) + (1-p)\left(\frac{\frac{1}{\tilde\pi_2}}{1+\frac{1}{\tilde\pi_2}}\right)$ needs to be

satisfied, which by computation is equivalent to $p \geq \frac{1}{2}$ i.e. $p \geq \max\left\{\frac{\frac{1}{\tilde{\pi}_2}}{1+\frac{1}{\tilde{\pi}_2}}, \frac{1}{2}\right\}$.

As the computations are symmetric across the three types we obtain the inequality $p \geq \max\left\{\frac{\frac{1}{\tilde{\pi}_2}}{1+\frac{1}{\tilde{\pi}_2}}, \frac{\frac{\tilde{\pi}_2}{\tilde{\pi}_3}}{1+\frac{\tilde{\pi}_2}{\tilde{\pi}_3}}, \frac{\tilde{\pi}_3}{1+\tilde{\pi}_3}, \frac{1}{2}\right\}$. Notice that each element in the right hand side of the preceding inequality is either constant or obtains a maximum in any compact set of prices $\Pi(\pi^*)$ contained in the interior of the unit simplex in $I\!\!R^3_{++}$ containing π^* The above discussion shows that, for $p > \bar{p}$, agents plays the equilibrium strategy, whatever their belief $(\tilde{\pi}_2, \tilde{\pi}_3)$ is. Straightforwardly, this statement extends to any belief in $\Delta_p \Pi(\pi^*)$ (that is: any distribution with probability p at least on π^* and with support in $\Pi(\pi^*)$). We then have shown that, for $p > \bar{p}$, $p-$belief on π^* implies that the equilibrium is played: this is $p-$dominance. This implies $p-$stability of π^* and it is even a stronger property. We summarize the above discussion as the following proposition:

Proposition 6. Suppose individuals initial expectations over prices lie in a compact set Π, a subset of the interior of the unit simplex in $I\!\!R^3_{++}$, which contains π^*. Then, there exists $0 < \bar{p} < 1$ such that whenever $p \in (\bar{p}, 1]$: the market equilibrium price π^* is $p-$dominant and hence, $p-$stable.

The above result demonstrates that the unique market equilibrium is p-stable for some whenever $p > \bar{p}$ where $0 < \bar{p} < 1$. Although the equilibrium outcome fails to be globally stable under tâtonnement dynamics, as $\bar{p} < 1$ it does have a measure of strategic stability. However, as $\bar{p} > 0$, equilibrium can never be the unique rationalizable outcome.

5.2 Equilibrium efficiency vs. the degree of strategic stability in a trading game

Here we study a simple trading game in two scenarios, one where a subset of traders are constrained and another in which no trader is constrained. We study the efficiency and stability properties of the market equilibria in two scenarios. In our trading game, a constrained market equilibrium will be *inefficient* while an unconstrained Nash equilibrium, which corresponds to a competitive equilibrium, will be *efficient*. We will show that the degree of strategic stability of the constrained market equilibrium is higher than the degree of strategic stability of the unconstrained market equilibrium although the latter yields efficient allocations while the former doesn't: enhanced equilibrium efficiency lowers the degree of strategic stability in our trading game.

We consider a market with two commodities, $l = 1, 2$, and a continuum of individuals, and two types of agents (of equal measure one each) I_1, I_2 such that each $i \in I_1$ has preferences represented by the utility function $x_1 + \frac{x_2^{1-\gamma}}{1-\gamma}$, $0 < \gamma < 1$, with endowments $\mathbf{w}^i = \{K, 0\}$, $K > 1$ each $j \in I_2$ has preferences represented by the utility function $\frac{x_1^{1-\beta}}{1-\beta} + x_2$, $0 < \beta < 1$, with endowments $\mathbf{w}^i = \{0, K\}$.

We model non-cooperative exchange using a modified Shapley-Shubik mar-

ket game (Shapley and Shubik (1977))[8] In the market game (1) a strategy for agent $i \in I_1$ is a bid $b^i \geq 0$ such that $b^i \leq 1$, (2) a strategy for agent $j \in I_2$ is an offer $q^j \geq 0$ such that $q^j \leq \alpha$, $0 < \alpha \leq K$. The interpretation is that b^i (respectively q^j) is the amount of commodity 1 (respectively, commodity 2) bid by agent i (resp. agent j) in exchange for commodity 2 (resp. commodity 1). Given an assignment of bids and offers $(\mathbf{b}, \mathbf{q})$, we compute market price π as follows:

$$\pi\left(\mathbf{B}, \mathbf{Q}\right) = \begin{cases} \frac{\mathbf{B}}{\mathbf{Q}} & \text{if } \mathbf{Q} > 0 \\ 0 & \text{otherwise} \end{cases}$$

and the allocation rule, if π exists, is $x_1^i(b^i, \pi(\mathbf{B}, \mathbf{Q})) = K - b^i, i \in I_1$, $x_2^i(b^i, \pi(\mathbf{B}, \mathbf{Q})) = \frac{b^i}{\mathbf{B}}\mathbf{Q}, i \in I_1$, $x_1^j(q^j, \pi(\mathbf{B}, \mathbf{Q})) = \frac{q^j}{\mathbf{Q}}\mathbf{B}, j \in I_2$, $x_2^j(q^j, \pi(\mathbf{B}, \mathbf{Q})) = \alpha - q^j + (K - \alpha), j \in I_2$, with the allocation rule yielding the initial endowments of traders if π doesn't exist. Given the allocation rule, α constrains the amount that type 2 traders can offer of commodity two.

In what follows we will focus on Nash equilibria with trade: we will call these market equilibria.

Note that α is a measure of the friction in the market. It is straightforward to check that when $\alpha \geq 1$ the set of market equilibrium prices and allocations coincide with competitive equilibrium prices and allocations (Dubey and Shapley (1994)). On the other hand, if $\alpha = 0$, the only Nash equilibrium involves no trade, traders consume their endowments: in this case, the Nash equilibrium is inefficient. Clearly, under the assumed boundary condition on utilities, as long as α is strictly positive but small (a precise bound is computed below), at a market equilibrium, all traders who own commodity two will be constrained and the resulting Nash equilibrium allocation will be inefficient.

We will study two distinct market equilibria:

(i) *Constrained market equilibria*: all type 2 traders who own commodity two are constrained (so that α is positive but close to zero), and

(ii) *Unconstrained market equilibria*: no type 2 trader who owns commodity two is constrained i.e. all traders are at an interior solution to their maximization problem (so that $\alpha \geq 1$).

Let $\Pi(\pi^*) \subset I\!\!R_{++}$ be a compact subset of market prices containing the competitive equilibrium price π^*. As before, we will define p−stability directly over prices (and justify this step by invoking the argument symmetric to the one used in the preceding subsection to prove Lemma 5).

p-stability of constrained market equilibrium

Fix the constrained market equilibrium price π^* (computed explicitly below). Let $\mathcal{N}(\pi^*) \subset I\!\!R_{++}$ be a compact set of prices containing π^*. From a formal viewpoint, the model under consideration is analogous to the one-dimensional smooth case considered in Section 2.4. Namely, every agent's decision is one-dimensional, the price π is a one-dimensional variable that is determined by agents' decision, and the best response of an agent is determined by his price

---

[8] In the case of two commodities, the Shapely-Shubik market game is strategically equivalent to the Shapley window model (with the normalization rule that $\pi_1 = 1$) studied in the preceding section.

24

expectation. It follows that $p-$stability of the equilibrium price is characterized using a one-dimensional temporary equilibrium map.[9]

We now define this temporary equilibrium map. To this purpose, consider the case when every individual attaches a probability p to π^* and attaches a probability $1 - p$ to some other positive market price $\tilde{\pi} \neq \pi^*$.

In this case, individual $i \in I_1$ solves

$$\max_{0 \leq b^i \leq 1} \ K - b^i + p \left[\frac{(\frac{b^i}{\pi^*})^{1-\gamma}}{1 - \gamma} \right] + (1 - p) \left[\frac{(\frac{b^i}{\tilde{\pi}})^{1-\gamma}}{1 - \gamma} \right].$$

Notice that under the assumptions made so far, the payoff function in the above maximization problem is strictly concave in b^i and moreover has a unique interior solution characterized by the first-order condition which yields the best response bid:

$$\mathbf{b}_p (\tilde{\pi}) = \left[p \left(\frac{1}{\pi^*} \right)^{1-\gamma} + (1 - p) \left(\frac{1}{\tilde{\pi}} \right)^{1-\gamma} \right]^{\frac{1}{\gamma}}$$

which is downward sloping in $\tilde{\pi}$. By computation note that the elasticity of the bid function is $\varepsilon_{b,p} (\tilde{\pi}) = \left| \frac{\mathbf{b}'_p(\tilde{\pi})}{\frac{\mathbf{b}_p(\tilde{\pi})}{\tilde{\pi}}} \right| = \frac{\frac{1-\gamma}{\gamma}(1-p)\tilde{\pi}^{\gamma-1}}{p\left(\frac{1}{\pi^*}\right)^{1-\gamma}+(1-p)\left(\frac{1}{\tilde{\pi}}\right)^{1-\gamma}}$. It turns out, given the computations reported below, the elasticity of the bid function restricted to p-beliefs and evaluated at the equilibrium price determined whether or not the equilibrium is p-stable.

Next, notice that at a constrained market equilibrium, we must have that $j \in I_2$ plays $q^* = \alpha$. This allows us to solve the market clearing equation $\mathbf{b}_p (\pi^*) = \pi^* q^*$ for the price π^*: the equilibrium market price is $\pi^* = \alpha^{-\gamma}$.

In order to ensure that at π^* it is a best-response for each $j \in I_2$ to submit a $q^* = \alpha$, observe that $j \in I_2$ solves

$$\max_{0 \leq q^j \leq \alpha} \ K - q^j + p \left[\frac{(q^j \pi^*)^{1-\beta}}{1 - \beta} \right] + (1 - p) \left[\frac{(q^j \tilde{\pi})^{1-\beta}}{1 - \beta} \right].$$

By computation, it follows that when $\alpha \leq 1$, the derivative of the payoffs of a type two trader w.r.t. q^j is positive at $q^j = \alpha$ when $\tilde{\pi}$ is in the neighborhood of π^*. We then choose $\mathcal{N} (\pi^*)$ small enough in order for $q^j = \alpha$ to be a best response when an individual $j \in I_2$ attaches a probability p to π^* and attaches a probability $1 - p$ to some other positive market price $\tilde{\pi} \neq \pi^*$ in $\mathcal{N} (\pi^*)$. The best response $q^j = \alpha$ of j is then invariant to p and $\tilde{\pi}$.

Define the temporary equilibrium map as follows

$$\Pi_p (\tilde{\pi}) = \frac{\mathbf{b}_p (\tilde{\pi})}{\alpha}.$$

$\Pi_p(\tilde{\pi})$ is the price that clears the markets whenever every agent attaches probability p on π^* and probability $(1-p)$ on $\tilde{\pi}$. The equilibrium price π^* is the unique fixed point of Π_p.

We know from the argument developed in Section 2.4 that p-stability of the equilibrium obtains iff the map Π_p is contracting on $\mathcal{N}(\pi^*)$. The condition $\left| \Pi'_p(\pi^*) \right| < 1$ is then necessary for $p-$stability. We choose $\mathcal{N}(\pi^*)$ small enough so that this condition is sufficient as well.

By computation, $\left| \Pi'_p(\pi^*) \right| = \varepsilon_{b,p}(\pi^*) = \frac{(1-\gamma)}{\gamma}(1-p)$. It follows that the equilibrium is $p-$stable iff the bid function, restricted to $p-$beliefs and evaluated at equilibrium, is *inelastic*.

Therefore, p-stability of the market equilibrium requires that $\frac{(1-\gamma)}{\gamma}(1-p) < 1$. Note as long as $1 > \gamma > \frac{1}{2}$, $\frac{(1-\gamma)}{\gamma} < 1$ and therefore the constrained market equilibrium is $0-$stable and hence $p-$stable for all $p \in [0,1]$. When $\gamma = \frac{1}{2}$, the constrained market equilibrium fails to be admissible.[10] For $0 < \gamma < \frac{1}{2}$, by computation, it is checked that the constrained market equilibrium is $p-$stable whenever $p \geq \frac{1-2\gamma}{1-\gamma}$ where $\frac{1-2\gamma}{1-\gamma} < 1$ as long as $0 < \gamma < \frac{1}{2}$.

Summing up, the minimum value of p for which the constrained Nash equilibrium is p-stable is $\hat{p}(\pi^*) = \min\left(0, \frac{1-2\gamma}{1-\gamma}\right)$.[11]

p-stability of unconstrained market equilibrium

Now suppose $\alpha = 2$ so that no trader is constrained in the trading game. Fix the unconstrained market equilibrium (and competitive equilibrium) price $\hat{\pi}^* = 1$. Let $\mathcal{N}(\hat{\pi}^*) \subset I\!R_{++}$ be a compact set of prices containing $\hat{\pi}^*$. The analysis for type 1 individuals is similar to that of the previous case and isn't repeated here. Each type 2 individual's unconstrained best-response is

$$\mathbf{q}_p(\tilde{\pi}) = \left[p(\pi^*)^{1-\beta} + (1-p)(\tilde{\pi})^{1-\beta} \right]^{\frac{1}{\beta}},$$

which is upward sloping. By computation, we obtain that the elasticity of the offer function is $\varepsilon_{q,p}(\tilde{\pi}) = \frac{\mathbf{q}'_p(\tilde{\pi})}{\frac{\mathbf{q}_p(\tilde{\pi})}{\tilde{\pi}}} = \frac{\frac{1-\beta}{\beta}(1-p)\tilde{\pi}^{1-\beta}}{p(\pi^*)^{1-\beta}+(1-p)(\tilde{\pi})^{1-\beta}}$.

Again, the characterization of $p-$stability relies on a temporary equilibrium map. Define this map as follows

$$\hat{\Pi}_p(\tilde{\pi}) = \frac{\mathbf{b}_p(\tilde{\pi})}{\mathbf{q}_p(\tilde{\pi})}.$$

$\hat{\Pi}_p(\tilde{\pi})$ is the price that clears the markets whenever every agent attaches probability p on $\hat{\pi}^*$ and probability $(1-p)$ on $\tilde{\pi}$. The equilibrium price $\hat{\pi}^*$ is the unique fixed point of Π_p.

When $\mathcal{N}(\hat{\pi}^*)$ is small enough, $p-$stability of $\hat{\pi}^*$ obtains iff $\left| \hat{\Pi}'_p(\hat{\pi}^*) \right| < 1$. By computation, $\left| \hat{\Pi}'_p(\hat{\pi}^*) \right| = \varepsilon_{b,p}(\hat{\pi}^*) + \varepsilon_{q,p}(\hat{\pi}^*)$ where $\varepsilon_{b,p}(\hat{\pi}^*) = \left| \frac{\mathbf{b}'_p(\hat{\pi}^*)}{\mathbf{b}_p(\hat{\pi}^*)/\hat{\pi}^*} \right| =$

---

[10] The slope of the best-response map is infinite.
[11] The limit case $\gamma = 1/2$ is omitted.

$\frac{(1-\gamma)}{\gamma}(1-p)$ is the elasticity of the bid function and $\varepsilon_{q,p}(\hat{\pi}^*) = \frac{\mathbf{q}'_p(\hat{\pi}^*)}{\mathbf{q}_p(\hat{\pi}^*)/\hat{\pi}^*} = \frac{(1-\hat{\beta})}{\beta}(1-p)$ is the elasticity of the offer function evaluated at $\hat{\pi}^*$. Therefore, in the unconstrained case, the market equilibrium is $p-$stable when the sum of the elasticities of the bid and the offer function, restricted to $p-$beliefs and evaluated at equilibrium, is less than one.

Hence, the condition $\left|\hat{\Pi}'_p(\hat{\pi}^*)\right| < 1$ writes $p > \hat{p}(\hat{\pi}^*)$, where

$$\hat{p}(\hat{\pi}^*) = \min\left(0, \frac{\gamma + \beta - 3\gamma\beta}{\gamma + \beta - 2\gamma\beta}\right). ^{12}$$

Lastly, it follows from the above expressions of $\left|\hat{\Pi}'_p(\hat{\pi}^*)\right|$ and $\left|\Pi'_p(\pi^*)\right|$ that $\left|\hat{\Pi}'_p(\hat{\pi}^*)\right| > \left|\Pi'_p(\pi^*)\right|$, that is: $p-$stability of the unconstrained equilibrium implies p-stability of the constrained equilibrium. In other words, we have $\hat{p}(\hat{\pi}^*) \geq \hat{p}(\pi^*)$.

We summarize the above analysis as the following proposition:

Proposition 7. When $0 < \gamma < 1$, $0 < \beta < 1$, $\gamma \neq 1/2$ and $, \beta^{-1} + \gamma^{-1} > 3$, the index of stability $\hat{p}(\pi^*)$ of the efficient unconstrained market equilibrium is larger than the index of stability $\hat{p}(\hat{\pi}^*)$ of the inefficient constrained market equilibrium. In other words, the unconstrained market equilibrium is less stable than the constrained market equilibrium.

6   Conclusion

In this paper, we have defined a continuous index for stability that can be assigned to each equilibrium within a large class of games. A Nash equilibrium is $p-$stable if it is the unique rationalizable outcome when all players attach at least probability $p < 1$ to the equilibrium strategy. In smooth applications, our stability index justifies the intuition that the degree of equilibrium stability relates to the slope of the best response map. The stability notion studied here requires the equilibrium to be robust to strategic uncertainty and is defined via an iterative elimination process. We showed that, for a given equilibrium, p-stability means that the equilibrium configuration of actions is the only outcome compatible with common knowledge of p-belief of equilibrium and common knowledge of rationality. Next, we developed a "contagion" argument relating strategic stability in the complete information case to incomplete information by adding payoff uncertainty (and informational asymmetry) to the underlying complete information game. In a game with a continuum of actions where the BR map is linear in the mean action (allowing for both strategic complementarities and substitutes), we showed that adding incomplete and asymmetric information on the fundamentals decreases the degree of stability. In two applications of our stability concept to large markets, we showed that (i) in the

---

^{12}It follows that $\hat{p}(\hat{\pi}^*) = 0$ when $\beta^{-1} + \gamma^{-1} \leq 3$ so that we will assume that $\beta^{-1} + \gamma^{-1} > 3$.

exchange economy studied by Scarf (1960) the actions of each trader at a Nash equilibrium corresponding to the unique globally unstable competitive equilibrium is p-dominant, and the Nash equilibrium configuration, p-stable for some $0 < p < 1$, (ii) enhanced equilibrium efficiency, by getting rid of trading frictions, could lower its degree of strategic stability.

Extending the stability concept studied here to large multistage games and applying the analysis to the stability of expectational coordination in intertemporal trade and financial markets are topics for future research.

References

[1] Artstein Z. (1979), "A note on Fatou's lemma in several dimensions", *Journal of Mathematical Economics*, 6(3), 277-282.

[2] Aumann R. and A. Brandenburger (1995), "Epistemic Conditions for Nash Equilibrium", *Econometrica*,63, 5, 1161-1180.

[3] Busetto, F., G. Codognato and S. Ghosal (2010), "Noncooperative oligopoly in markets with a continuum of traders", *Games and Economic Behavior*, forthcoming, doi:10.1016/j.geb.2010.06.009.

[4] Carlsson H. and E. van Damme (1993), "Global Games and Equilibrium Selection", *Econometrica*, 61,5, 998-1018.

[5] Codognato G., Ghosal S. (2000), "Cournot-Nash equilibria in limit exchange economies with complete markets and consistent prices," *Journal of Mathematical Economics* 34, 39-53.

[6] Dubey P., Shapley L.S. (1994), "Noncooperative general exchange with a continuum of traders: two models, " *Journal of Mathematical Economics*, 23, 253-293.

[7] Dudley, R.M. (1989), *Real Analysis and Probability*, Wadsworth & Brooks/Cole Advanced Books and Software, Wadsworth Inc. Belmont, California.

[8] Guesnerie R. (1992), "An Exploration of the Eductive Justifications of the Rational Expectations Hypothesis", *American Economic Review*, 82,5, 1254-1278.

[9] Hu T.W. (2007), "On p-rationalizability and approximate common certainty of rationality", *Journal of Economic Theory*, 136 (1), 379-391.

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

[11] Mertens J.F. and S. Zamir (1985), "Formulation of Bayesian analysis for games with incomplete information", *International Journal of Game Theory*, Volume 14, Number 1 / March, 1985.

[12] Morris S., R. Rob and H.S. Shin (1995), "p-Dominance and Belief Potential", *Econometrica*, 63,1, 145-157.

[13] Morris S. and H.S. Shin (1998), "Unique equilibrium in a model of self-fulfilling currency attacks", *American Economic Review*, 88, 587-597.

[14] Sahi S., Yao S. (1989), "The noncooperative equilibria of a trading economy with complete markets and consistent prices," *Journal of Mathematical Economics* 18, 325-346.

[15] Scarf H. (1960), "Some Examples of Global Instability of the Competitive Equilibrium", *International Economic Review*, 1(3), 157-172.

[16] Shapley L.S., Shubik M. (1977), "Trade using one commodity as a means of payment," *Journal of Political Economy* 85, 937-968.

[17] Tan T. and S. Werlang (1988), "On Aumann's Notion of Common Knowledge, an Alternative Approach", *Journal of Economic Theory*, 45: 370-91.

[18] Tercieux O. (2006), "p-Best response set", *Journal of Economic Theory*, 131(1), 45-70.