

Stochastically Stable Implementation*

Antonio Cabrales[†] Roberto Serrano[‡]

April 2011

Abstract

Restricting attention to economic environments, we study implementation under perturbed better-response dynamics (BRD). A social choice function (SCF) is implementable in stochastically stable strategies of perturbed BRD whenever the only outcome supported by the stochastically stable strategies of the perturbed process is the outcome prescribed by the SCF. For uniform mistakes, we show that any ε -secure and strongly efficient SCF is implementable when there are at least five agents. Extensions to incomplete information environments are also obtained.

JEL Classification: C72, D70, D78.

Keywords: robust implementation; bounded rationality; evolutionary dynamics; mechanisms; stochastic stability.

*We thank Jeff Ely, Takashi Kunimoto, Eric Maskin, Stephen Morris, Olivier Tercieux, Norov Tumenasan and Fernando Vega-Redondo for helpful conversations and encouragement. The comments of the Guest Editor, Matthew Jackson, and an anonymous referee greatly helped to improve the paper. We gratefully acknowledge the financial support from Spain's Ministry of Science and Innovation under grants CONSOLIDER INGENIO 2010 CSD2006-0016 (both), and ECO2009-10531 (Cabrales). Serrano thanks CEMFI in Madrid for its hospitality.

[†]Departamento de Economía, Universidad Carlos III de Madrid, Madrid 126, 28903 Getafe, Spain, and CEPR; Email: antonio.cabrales@uc3m.es; <http://www.eco.uc3m.es/acabrales>

[‡]Department of Economics, Brown University, Providence, RI 02912, USA, and IMDEA-Social Sciences Institute, Madrid, Spain; Email: roberto_serrano@brown.edu; <http://www.econ.brown.edu/faculty/serrano>

1 Introduction

This is a companion paper to Cabrales and Serrano (2010), referred to as CS from now on.¹ As in that paper, we continue to study the implementation problem under a plausible class of learning processes, that of better-response dynamics (BRD) and perturbations thereof. Thus, we postulate a behavioral assumption by which agents (or generations of agents) interact myopically within a given mechanism, and adjust their actions in the direction of better-responses. A first criterion for successful implementation is the convergence of the better-response process to a rest point or to a set of rest points. When the outcome of a social choice function (SCF) is the only limit of the BRD in a mechanism for any allowed environment, we shall say that the SCF is *implementable in recurrent strategies* of BRD. CS provides necessary and sufficient conditions for implementability in this sense, among which the most salient condition is quasimonotonicity, a variant of Maskin monotonicity.

Those results on recurrent implementation in BRD are obtained for a general class of preferences and will stand for any perturbed process. The latter means that, if one were to perturb the BRD via mistakes (by allowing agents not to use a better response sometimes), an SCF that is implementable in recurrent strategies would also be *implementable in stochastically stable strategies* of any perturbation of BRD. That is, the outcomes prescribed by the SCF are the states of minimum stochastic potential (see, e.g., Young (1998, Chapter 3)), for any perturbed process. Therefore, quasimonotonicity is identified as the key condition to essentially characterize very robust implementation with respect to myopic BRD processes. In this way, these conclusions are immune to the Bergin and Lipman (1996) critique of uniqueness results in stochastic evolutionary implementation.

The current paper considers how to obtain implementability results in these contexts once one moves beyond quasimonotonicity. Since implementability in recurrent strategies of BRD will not be possible, given the necessity of that condition, it follows that the permissive results we describe here must rely on a different class of dynamics, such as certain perturbations of BRD. Specifically, strengthening the assumptions on preferences and mistakes processes, we show that there are mechanisms for evolutionary implementation under relatively permissive conditions on SCFs.

We present here a result that uses uniform mutations or “mistakes” in the BRD process.²

¹To avoid obvious repetitions here, we refer to CS for an extensive literature review.

²In our working paper version (available at <http://www.eco.uc3m.es/~acabrales/research/CS-stochimple-2.pdf>) we also show that, under a variant of the “more serious mistakes are less likely” assumption, any ε -secure SCF (a version of the NWA condition found in CS formulated for economic environments) is imple-

It states that, under uniform mistakes (“all mistakes are equally likely”) and an assumption on diversity of preferences, any Pareto efficient and ε -secure SCF can be reached if there are at least five agents in the environment; if the required preference diversity happens near the zero bundle, the Pareto assumption can be dispensed with altogether.

The findings in this paper, vis-à-vis those in CS, should not be interpreted as “on-the-one-hand, on-the-other-hand” type of results. We formalize a genuine tradeoff for the social planner. If the SCF he wishes to implement satisfies quasimonotonicity, he knows that he has an evolutionarily robust mechanism for implementation at his disposal. If not, there exist mechanisms that are robust under evolution, but more requirements are needed from other fundamentals of the problem. In addition, stochastically stable outcomes may require a very long time for convergence (see e.g. Ellison 2000). Hence a high degree of patience on the attainment of social goals is required for the social planner and society as a whole. Thus, unlike what some of the previous implementation literature has suggested, there is no “free lunch” in terms of implementability.

Our main insights already described are confirmed in environments with incomplete information, and some others are obtained therein. First, incentive compatibility arises as a necessary condition for stable implementation in our sense, whatever the perturbation one wishes to use, including no perturbation at all, of interim BRD. As shown in CS, if one wishes to implement in recurrent strategies, faithful to the robustness line of thinking enunciated above, the condition of Bayesian quasimonotonicity is also necessary. Moreover, that paper shows that incentive compatibility, Bayesian quasimonotonicity and ε -security are also sufficient for implementation in recurrent strategies of BRD processes when there are at least three agents. In contrast, we show here that under weak preference diversity in the environment, the condition of Bayesian quasimonotonicity can be entirely dropped. This can be done if the planner is satisfied with implementation in stochastically stable strategies under uniform mistakes, and if there are at least five agents. Thus, we find the same tradeoff described earlier: evolutionary implementation results more permissive than those relying on the quasimonotonicity conditions are possible, but they come at a cost in terms of their robustness.

mentable in stochastically stable strategies of the corresponding perturbed BRD process if there are at least three agents.

2 Preliminaries

Let $N = \{1, \dots, n\}$ be a set of agents. For simplicity, we concentrate on economic environments. Let agent i 's consumption set be a finite set, $X_i \subset \mathbb{R}_+^l$ (where we assume $0 \in X_i$, for all $i \in N$). One can specify that each agent holds initially the bundle $\omega_i \in X_i$ with $\sum_{i \in N} \omega_i = \omega$ (private ownership economies), or simply that there is an aggregate endowment of goods ω (distribution economies). The set of alternatives is the set of allocations:

$$Z = \{(x_i)_{i \in N} \in \prod_{i \in N} X_i : \sum_{i \in N} x_i \leq \omega\}.$$

Let $\theta = (\theta_i)_{i \in N}$ be a preference profile, and Θ be the set of allowable preference profiles. For now, we shall describe environments with complete information. (Section 4 will extend the analysis to incomplete information environments.) We make the following assumptions on preferences:

- (1) No consumption externalities: $\theta_i : X_i \times X_i \mapsto X_i$, that is, an agent's preference relation depends on the bundle of goods that he consumes, and not on other agents' bundles.
- (2) Strictly increasing preference: For all i and for all $x_i \in X_i$, if $y_i \geq x_i$, $y_i \succ_i^{\theta_i} x_i$.³ Note how this implies that 0 is the worst bundle for every agent.

A *social choice function* (SCF) assigns an outcome to each $\theta \in \Theta$. We shall denote an SCF by f , and thus, $f : \Theta \mapsto Z$.

A mechanism $G = ((M_i)_{i \in N}, g)$, where M_i is agent i 's (finite) message set, and $g : \prod_{i \in N} M_i \mapsto Z$ is the outcome function. A Nash equilibrium of the mechanism in state θ is a profile of messages m^* such that for every $i \in N$, $g(m^*) \succeq_i^\theta g(m_i, m_{-i}^*)$ for all $m_i \neq m_i^*$. A strict Nash equilibrium is a Nash equilibrium in which all these inequalities are strict. Given a profile $m \in \prod_{i \in N} M_i$, agent j 's (weak) better-response to m is any m'_j such that $g(m'_j, m_{-j}) \succeq_j^\theta g(m)$.

We concentrate on the following class of SCFs. An SCF f is said to be ε -secure if there exists $\varepsilon > 0$ such that for each θ , and for each $i \in N$, $f_i(\theta) \geq (\varepsilon, \dots, \varepsilon) \gg 0$.

The condition of ε -security amounts to establishing a minimum threshold of living standards in the consumption of all commodities. We shall think of ε as being a fairly small number. Then, one could easily justify the property on normative grounds.

³For vectors $x_i, y_i \in X_i$, we use the standard conventions: $x_i \geq y_i$ whenever $x_{il} \geq y_{il}$ with at least one strict inequality; and $x_i \gg y_i$ whenever $x_{il} > y_{il}$ for every commodity l .

Next, we turn to dynamics, the central approach in our paper. The mechanism will be played simultaneously each period by myopic agents. Or, in an interpretation closer to the evolutionary tradition, the mechanism will be played successively each period by generations of agents who live and care for that period only. Given a mechanism, we take the set $M = \prod_{i \in N} M_i$ of message profiles as the finite state space. We shall begin by specifying an unperturbed Markov process on this state space, i.e., a matrix listing down the transition probabilities from any state to any other in a single period.⁴ Such a process will typically have multiple long-run predictions, which we call *recurrent classes*. A recurrent class is a set of states that, if ever reached, will never be abandoned by the process, and that does not contain any other set with the same property. A singleton recurrent class is called an *absorbing state*.

The unperturbed Markov process that we shall impose on the play of the mechanism over time is the following *better-response dynamics (BRD)*. In each period t , each of the agents is given the chance, with positive, independent and fixed probability, to revise his message or strategy. Simultaneous revision opportunities for different agents are allowed. Let $m(t)$ be the strategy profile used in period t , and say agent i is chosen in period t . Then, denoting by θ_i agent i 's true preferences, agent i switches with positive probability to any m'_i such that $g(m'_i, m_{-i}(t)) \succeq_i^{\theta_i} g(m(t))$.

CS study the problem of implementability in recurrent strategies of BRD processes, and provide necessary and sufficient conditions for it. The key condition that underlies much of their analysis is quasimonotonicity, a variant of Maskin monotonicity. One way to justify the results in the current paper is the search of conditions under which implementability in terms of perturbed BRD processes may expand the set of implementable SCFs beyond quasimonotonicity. The problem, though, in trying to implement an SCF that violates quasimonotonicity is that, since it cannot be done in recurrent classes of BRD, initial conditions will matter. Thus, some paths in the BRD dynamics may lead to the SCF outcome, but others will not.

Indeed, the dependence of long-run predictions of unperturbed Markov processes on initial conditions is sometimes perceived as a drawback of this analysis. One way out is to perturb the Markov process. The class of perturbations that we are interested in specify a Markov matrix of transition probabilities that is both irreducible and aperiodic. Irreducibility means that it is always possible to transit from any state to any other in a finite number

⁴For complete formal definitions of Markov chains, and related terms (recurrent classes, absorbing states, irreducibility, aperiodicity, etc.) see e.g. Karlin and Taylor (1975), chapter 2.

of periods. Aperiodicity is implied because there is a chance that the state does not change from one period to the next. For an irreducible and aperiodic process, there is a unique *stationary distribution* with the following two properties. First, starting from any initial strategy profile, the probability distribution on period t strategy profiles is known to approach that stationary distribution as $t \rightarrow \infty$. And second, the stationary distribution also represents the proportion of time spent on each state over an infinite time horizon. If one denotes by μ^ϵ the stationary distribution of the ϵ -perturbed Markov process and takes the limit as $\epsilon \rightarrow 0$, one gets that the $\lim_{\epsilon \rightarrow 0} \mu^\epsilon = \mu^*$ exists and is one of the multiple stationary distributions of the unperturbed process. We shall refer to the states in the support of μ^* as the *stochastically stable states* of the perturbed process, which are interpreted as the only states in which the perturbed process spends a positive proportion of time in the long run when the amount of noise ϵ is positive, but negligible.

Thus, the planner, who has a long run perspective on the social choice problem, wishes to design an institution or mechanism such that, when played by myopic agents who keep adjusting their actions in the direction of better-responses most of the time, but who may also make mistakes, the socially desirable outcome as specified by the SCF, is the only *stochastically stable state* of the process. This logic suggests the following implementability notion.

An SCF f is *implementable in stochastically stable strategies* (of perturbed BRD) if there exists a mechanism G such that, for every $\theta \in \Theta$, a perturbation of the BRD process applied to its induced game when the preference profile is θ has every $f(\theta)$ as the unique outcome supported by stochastically stable strategy profiles.

Therefore, when f is implementable in stochastically stable strategies of a perturbed BRD process, in the very long run, for each θ , the proportion of time spent by the process at $a = f(\theta)$ is 1.

Before closing the section, we go over some basic concepts in perturbed Markov processes, which we will use in the sequel. In order to identify the stochastically stable strategy profiles of any perturbed BRD process, we will use the characterization of stochastic stability provided by Young (1993) and Kandori, Mailath and Rob (1993), based on the techniques developed by Freidlin and Wentzell (1984).

Call the unperturbed Markov BRD process P^0 defined on the finite state space M . We define a perturbed process of P^0 as follows: fixing $\epsilon^* > 0$, for each $\epsilon \in (0, \epsilon^*)$, the process P^ϵ is a regular perturbed Markov process if P^ϵ is irreducible for every $\epsilon \in (0, \epsilon^*)$ and for every

$m, m' \in M$, $P^\epsilon(m, m')$ approaches $P^0(m, m')$ at an exponential rate. That is,

$$\lim_{\epsilon \rightarrow 0} P^\epsilon(m, m') = P^0(m, m')$$

and

$$\exists \epsilon > 0 | P^\epsilon(m, m') > 0 \quad \text{implies} \quad \exists r(m, m') \geq 0 | 0 < \lim_{\epsilon \rightarrow 0} \frac{P^\epsilon(m, m')}{\epsilon^{r(m, m')}} < \infty.$$

The real number $r(m, m')$ is called the resistance of the transition from m to m' . Note that it is uniquely defined, i.e., there cannot be two exponents satisfying the above condition. Note also that $P^0(m, m') > 0$ if and only if $r(m, m') = 0$: transitions that can occur under P^0 have zero resistance. For convenience, we shall assume that $r(m, m') = \infty$ if $P^\epsilon(m, m') = P^0(m, m') = 0$ for every $\epsilon \in (0, \epsilon^*)$ (this way the resistance is defined for every pair of states).

Similarly, let $\xi = (z^1, \dots, z^k)$ be an (m, m') -path, i.e., a finite sequence of states in which $z^1 = m$ and $z^k = m'$. The resistance of the path ξ is the sum of the resistances of its transitions.

Let $E = \{E_0, \dots, E_k\}$ be the set of recurrent classes of the unperturbed process and consider the complete directed graph with vertex set E , which is denoted by Γ . We want to define the resistance of each one of the edges in this graph. For this, let E_i and E_j be two elements of E . The resistance of the edge (E_i, E_j) in Γ is the minimum resistance over all the resistances of the (E_i, E_j) -paths. Note that while E_i and E_j are two recurrent classes, (E_i, E_j) -paths are typically composed of any kind of states, not necessarily recurrent.

Let E_i be a recurrent class. A E_i -tree is a tree with vertex set E such that from every vertex different from E_i , there is a unique directed path in the tree to E_i . The resistance of the E_i -tree is the sum of the resistances of the edges that compose it. The *stochastic potential* of the recurrent class E_i is the minimum resistance over all the E_i -trees. Young (1993) shows that the set of stochastically stable states of the process consists of those states with minimum stochastic potential. Thus, what is key is the identification of paths of minimum resistance, and this is what the proofs of our sufficiency results in the next sections will do.

3 Complete Information

In this section we present two results for complete information environments, based on two distinct perturbed BRD processes. The first uses a perturbation of the “more serious mistakes are less likely” type, while the second uses uniform or equally-likely mistakes.

3.1 When More Serious Mistakes Are Less Likely

This subsection explores the possibilities of obtaining a permissive implementation result, by imposing a specific kind of perturbation of BRD. It is instructive to note that the institution we shall employ to this end will be essentially the same canonical mechanism used in the proof of Theorem 2 in CS. The main differences are in the fact that here we deal with social choice functions, rather than correspondences, and that we use 0 in rule (iii) rather than a modulo game. Neither of these differences is essential, in the sense that working with social choice correspondences would be easy in our approach, and using modulo games would yield the same results.

We use the following additional assumptions on preferences:

- (3) Let commodity 1 be a numeraire whose indivisible unit is $\Delta > 0$. The preference is quasilinear in the numeraire. Also, let $\Delta > 0$ be smaller than any utility gap resulting from reallocations of the non-numeraire commodities.
- (4) The preference is continuous.
- (5) For an SCF f , let $\mathbb{L}^{\phi_i}(f(\theta'))$ be the lower contour set of $f(\theta')$ under any preference ϕ_i . If for all agents i , $\mathbb{L}^{\theta_i}(f(\theta')) \subset \mathbb{L}^{\theta'_i}(f(\theta'))$, we assume that at least for one j , $\mathbb{L}^{\theta_j}(f(\theta')) \neq \mathbb{L}^{\theta'_j}(f(\theta'))$.

Assumption (3) is needed because we use the penalties in the numeraire (see rule (ii.a) in the proof below), which are smaller than any other caused by reallocations of the other goods. Assumption (4) is needed because we shall quantify the resistance of each transition through utility differences. Assumption (5) says that if there is an inclusion for all agents of the relevant lower contour sets at an outcome of the SCF, at least for one agent the inclusion is strict.

For the perturbed process used in the next theorem we shall specify a very concrete type of perturbations. The interpretation is that agents may make “mistakes” with positive, though small probability, when changing their strategies in the mechanism. When preferences are θ , each agent i is said to “make a mistake” in choosing m'_i against m_{-i} whenever $g(m) \succ_i^\theta g(m'_i, m_{-i})$.

The idea is to introduce an assumption that is a variant of “more serious mistakes are less likely.” Specifically, suppose that at the status-quo in the mechanism, agent i is receiving bundle z_i^0 . Suppose agent i takes an action in the mechanism in which he asks for bundle y_i and forces a change in outcome to bundle z_i , out of which he suffers a utility loss. In

principle, one should think of the probability of such a transition to depend on all three components: the initial and final bundles in the transition, as well as the exact messages used in the mechanism to force such a transition.

Consider a perturbation of BRD, in which one allows transitions where agent i moves and becomes worse off going from z_i^0 to z_i . We shall define the resistance of such a transition to be the following:

$$[u_i(z_i^0) - u_i(z_i)] + \lambda[u_i(y_i) - u_i(z_i)],$$

where $0 < \lambda < 1$ is small enough to ensure that this resistance is always positive, and u_i is a utility function representing agent i 's preferences.⁵

That is, the first term says that more serious mistakes are less likely (the first component of resistance is utility loss in the transition). However, this is affected by the size of the disappointment/relief of the agent inducing an outcome change when comparing the final outcome with the one proposed by him. For a given amount of disappointment/relief, the transition is all the more likely the smaller the utility loss. And, for a given utility loss, the transition is all the more likely the smaller the disappointment or the greater the relief (as if the agent exhibited disappointment aversion-relief attraction). If the term multiplying λ is positive –disappointment–, the agent is less likely to make a mistake that will imply a greater level of disappointment. If it is negative –relief–, a mistake is more likely the greater the relief. Other interpretations of the second term of the resistance are possible. For example, one could explain it in terms of how others perceive the agent that moves. For a given real utility loss suffered by i due to his action, such a transition is more likely when the others view him as “self-sacrificing” instead of “greedy.” In any event, we emphasize that these behavioral departures from the standard conventional assumptions are minimal – λ can be taken arbitrarily small.

Apart from this, any transition from any bundle other than 0 to 0 has a fixed resistance, which we will call K . (If K were large, it would be as if the planner were “extremely reluctant” to use rule (iii) in the mechanism of the proof of Theorem 1.)

THEOREM 1 *Suppose the environments satisfy Assumptions (1) through (5). Let $n \geq 3$. Then, any ϵ -secure SCF f is implementable in stochastically stable strategies of the perturbed BRD process based on the prescribed variant of “more serious mistakes are less likely.”*

Proof: Consider the mechanism $G = ((M_i)_{i \in N}, g)$, where agent i 's message set is $M_i =$

⁵Any utility function that represents the preferences will do. The existence of such a utility representation follows from Assumption (4).

$\Theta \times Z$. Denote agent i 's message $m_i = (\theta_i, a_i)$, and the agents' message profile by m . Let $\beta = (\Delta, 0, \dots, 0)$. The outcome function g be defined by the following rules:

- (i.) If for all $i \in N$ and any $z_i \in Z$, $m_i = (\theta, z_i)$, $g(m) = f(\theta)$.
- (ii.) If for all $j \neq i$, and for any $z_j \in N$, $m_j = (\theta, z_j)$ and $m_i = (\phi, z)$, $\phi \neq \theta$, one can have two cases:
 - (ii.a) If $z \succeq_i^\theta f(\theta)$, $g(m) = (f_i(\theta) - \beta, f_{-i}(\theta))$, where $f_i(\theta) \geq f_i(\theta) - \beta \in X_i$.
 - (ii.b) If $f(\theta) \succ_i^\theta z$, $g(m) = z - \lambda\beta$.
- (iii.) In all other cases, $g(m) = 0$.

We begin by arguing in the next three steps that all recurrent classes of the unperturbed better-response process must happen under rule (i). Let θ be the true preference profile.

Step 1: No message profile in rule (iii) is part of a recurrent class. Arguing by contradiction, from any profile m in (iii), one can construct a path as follows. Without loss of generality, suppose $m_1 = (\phi, z) \neq (\theta, f(\theta))$. In the path, change one by one the strategies of all agents other than 1, starting from agent n and going down to agent 2, to $(\theta, f(\theta))$. In doing this, one constructs a sequence of outcomes consisting of the zero allocation until, in the last step, when $(n - 1)$ messages are $(\theta, f(\theta))$, the outcome switches to either z or $(f_1(\theta) - \beta, f_{-1}(\theta))$, consistent with better-response dynamics. In the last step of the path, agent 1 switches from (ϕ, z) to $(\theta, f(\theta))$. This yields $f(\theta)$, from which one can never go back to the zero allocation under better-response dynamics.

Step 2: No message profile under rule (ii.a) is part of a recurrent class of better-response dynamics. We argue by contradiction. Recall that the true preference profile is θ , and, again with no loss of generality, let the message profile under rule (ii.a) in question be the following: all agents $j \neq i$ announce $m_j = (\phi, f(\phi))$, whereas agent i 's message is (ϕ', z') such that $z' \succeq_i^\phi f(\phi)$, leading to an outcome in which agent i receives $f_i(\phi) - \beta$. Because preferences are strictly increasing, one can construct a single-step path under better-response dynamics in which agent i switches to (ϕ, z) , where $z_i = f_i(\phi) - \beta'$ (for $\beta' < \beta$) and $z_j = 0$ for every $j \neq i$, which yields outcome z . But from here, each of the other agents $j \neq i$ can switch to (ϕ^j, z^j) (for some $(\phi^j, z^j) \neq (\phi, f(\phi))$). Thus, we find ourselves under rule (iii), which is a contradiction.

Step 3: No recurrent class contains profiles under rule (ii.b). Again, we argue by contradiction. As in step 2, consider a profile m such that for all $j \neq i$ $m_j = (\phi, f(\phi))$, whereas

$m_i = (\phi', z')$, satisfying that $f_i(\phi) \succ_i^\phi z'_i$. This implies that the outcome is z' . Then, construct a path in which agent i switches, if necessary, to (ϕ', z) , where $z_i = z'_i$ and for all $j \neq i$, $z_j = 0$, after which the outcome is z . But then, as before, any of the other agents can switch to yield an outcome under rule (iii), a contradiction.

Moreover, each recurrent class, containing only outcomes under rule (i), must consist exclusively of Nash equilibria of the game induced by the mechanism when the true preferences are θ . That is, a non-equilibrium strategy profile would not be a recurrent state of better-response dynamics. All the truthful strategy profiles, $((\theta, \cdot), \dots, (\theta, \cdot))$, always constitutes one of these recurrent classes, and in addition, there may exist recurrent profiles $((\phi, \cdot), \dots, (\phi, \cdot))$ with the property that for all $i \in N$, $f(\phi) \succ_i^\phi z$ implies that $f(\phi) \succ_i^\theta z$.

We classify the recurrent classes of the unperturbed process into two kinds:

Class E_0 is the truth-telling recurrent strategy profiles, i.e., for each $i \in N$, $m_i = (\theta, \cdot)$ with outcome $f(\theta)$.

Class E_j for $j = 1, \dots, J$ is the coordinated lie on profiles θ^j : for each $i \in N$, $m_i = (\theta^j, \cdot)$ with outcome $f(\theta^j)$. Such profiles are known to be Nash equilibria of the mechanism under preference profile θ . Note that, for this to be true, as we have already pointed out, the strictly lower contour set at $f(\theta^j)$ for each agent i when his preferences are θ_i^j must be contained in the strictly lower contour set of $f(\theta^j)$ when his preferences are θ_i .

Now, we show that the profiles in E_0 are the only stochastically stable profiles of the prescribed perturbed dynamics:

[a] to get out of E_0 , one can go through rule (ii.a) of the mechanism, paying Δ if the deviator i proposes an outcome indifferent to $f(\theta)$, or go through rule (ii.b) paying a cost that is exactly the smallest utility loss from $f(\theta)$ to z plus $\lambda\Delta$, which, by Assumption 3, is not smaller than $(1 + \lambda)\Delta$. After that, a mistake that takes us to rule (iii), which costs K , takes us to 0 and from there we go for free to any of the untruthful Nash equilibria in any class E_j .

[b] to get out of an arbitrary class E_j , we have those two paths as well, but the cheapest will be one under rule (ii.a') again. Indeed, let agent i deviate from the otherwise unanimous announcement (θ^j, \cdot) with outcome $f(\theta^j)$, and instead announce (ϕ, z) such that $z \succeq_i^{\theta^j} f(\theta^j)$ and $f(\theta^j) \succ_i^\theta z$ (the existence of z is guaranteed by Assumption (5)). In this case, the resistance is strictly smaller than Δ , because of the relief term. After that, we go to rule (iii) paying also K , and from there we go for free to E_0 . ■

Remark: Note the novel use of the inclusion of the lower contour sets of preferences made in the last step of the proof. Although the assumptions made on mistakes are somewhat

special, we think it is interesting that Theorem 1 dispenses with quasimonotonicity, while still making use of essentially the same mechanism as does Theorem 2 in CS.

3.2 Uniform Mistakes

We focus in this subsection on another permissive result, based on uniform mistakes. Uniform mistakes means that each “mistake” made by an agent, i.e., each revision of his strategy that goes against the better-response direction, is equally likely (say, it has a small probability $\epsilon > 0$).

To get such a result on implementability in perturbed better-responses under uniform mistakes, we use an additional assumption on the SCF, i.e., that it is (strongly Pareto) efficient:⁶ We write the definition of efficiency as we will use it:

An SCF f is (*strongly*) *Pareto efficient* if for all θ and for all alternative outcomes $z \neq f(\theta)$, there exists an individual $i(\theta, z)$ such that $f(\theta) \succ_{i(\theta, z)}^\theta z$.⁷

In addition to (1) and (2), we shall require Assumption (3) below. Before getting to it, we go over some necessary material in the next paragraphs.

First, note that since states differ because at least one of the agents’ preference varies, one has that for each pair of states θ and ϕ , there exists an agent $j(\theta, \phi)$ and alternatives $x(\theta, \phi)$ and $y(\theta, \phi)$ such that

$$x(\theta, \phi) \succ_{j(\theta, \phi)}^\theta y(\theta, \phi) \quad \text{and} \quad y(\theta, \phi) \succeq_{j(\theta, \phi)}^\phi x(\theta, \phi). \quad (1)$$

Denote by $J(\theta, \phi)$ the set of agents $j(\theta, \phi)$ for whom there exists a preference reversal between a pair of alternatives across states θ and ϕ , as specified in (1).

Also, without loss of generality, note that for all θ, ϕ , one can choose alternative $y(\theta, \phi)$ so that for all $i \neq j(\theta, \phi)$, $y_i(\theta, \phi) \neq 0$. We shall do this in the sequel.

Here is our regularity assumption on the environments:

- (3)** For each pair of states θ and ϕ , there exists $j(\theta, \phi) \in J(\theta, \phi)$ such that $j(\theta, \phi) \neq i(\theta, x(\theta, \phi))$, where $x(\theta, \phi)$ is an alternative for which agent $j(\theta, \phi)$ has a preference reversal as in (1).

This assumption is used because the mechanism in theorem 2 will move the game from some $f(\theta)$ to some other outcome $x(\theta, \phi)$ as specified in condition (1). We need that the

⁶As we shall remark after the proof of the result in this subsection, one can get rid of this by making a different assumption on the environments.

⁷Thus, we are ruling out cases such as linear indifference curves with the same slope for all agents.

identity of some agent who loses out in this move from $f(\theta)$ to $x(\theta, \phi)$ (who exists by Pareto efficiency) be different from the identity of the agent experimenting the preference reversal, and that is what condition (3) requires. For example, a “replica” economy in which the preferences in the base economy are not all identical would meet this assumption.

Now, we can prove the following result:

THEOREM 2 *Suppose the environments satisfy Assumptions (1), (2) and (3). Let $n \geq 5$. Any ε -secure and strongly Pareto efficient SCF f is implementable in stochastically stable strategies of perturbed BRD, where the perturbation consists of uniform mistakes.*

Proof: Consider the following mechanism. Let agent i 's message set be $M_i = \Theta \times Z$. Let a typical message sent by agent i be $m_i = (m_i^1, m_i^2)$ and the corresponding message profile be $m = (m^1, m^2)$. The outcome function obeys the following rules:

- (i) If for every $i \in N$, $m_i^1 = \theta$, $g(m) = f(\theta)$.
- (ii.a) If exactly $(n - 1)$ messages m_i are such that $m_i^1 = \theta$ and $m_{i(\theta, x(\theta, \phi))} = (\phi, x(\theta, \phi))$, $g(m) = (x_{i(\theta, x(\theta, \phi))}(\theta, \phi), x_{j(\theta, \phi)}(\theta, \phi), 2\varepsilon, 2\varepsilon, \dots, 2\varepsilon)$.
- (ii.b) If exactly $(n - 1)$ messages m_i are such that $m_i^1 = \theta$, but the odd man out, say agent k , does not satisfy the requirements of rule (ii.a), $g(m) = (f_k(\theta) - \beta, f_{-k}(\theta))$, where $f_k(\theta) \geq f_k(\theta) - \beta \geq (\varepsilon, \dots, \varepsilon)$.
- (iii.a) If exactly $(n - 2)$ messages m_i are such that $m_i^1 = \theta$, $m_{i(\theta, x(\theta, \phi))} = (\phi, x(\theta, \phi))$ and $m_{j(\theta, \phi)} = (\phi, y(\theta, \phi))$, $g(m) = (y_{i(\theta, x(\theta, \phi))}(\theta, \phi), y_{j(\theta, \phi)}(\theta, \phi), \varepsilon, \varepsilon, \dots, \varepsilon)$.
- (iii.b) If exactly $(n - 2)$ messages m_i are such that $m_i^1 = \theta$, but we are not under rule (iii.a), for all $k \in N$, $g_k(m) = (\varepsilon, \dots, \varepsilon)$.
- (iv) In all other cases, $g(m) = 0$.

(For rule (iii.a) to be well defined, the assumption $n \geq 5$ is needed to determine the outcome in profiles where two agents report the same state θ as part of their message and two other agents report a different state ϕ , each pair of agents involving $j(\theta, \phi)$ and $j(\phi, \theta)$, respectively.)

We begin by arguing in the next four steps that all recurrent classes of the unperturbed better-response process must happen under either rule (i), under rule (ii.a) or under rule (iii.a). But under rules (ii.a) or (iii.a) this only happens when the common announcement by $n - 1$ or $n - 2$ people is not the true preference profile.

Let θ be the true preference profile.

Step 1: No message profile in rule (iv) is part of a recurrent class. From any profile m in (iv), one can construct a path as follows. For all players it is a better response to announce $(\theta, f(\theta))$. This yields $f(\theta)$, from which one can never go back to the zero allocation under better-response dynamics.

Step 2: No message profile in rule (ii.b) or (iii.b) is part of a recurrent class. Let ϕ be the announcement of the $n - 1$ or $n - 2$ people announcing a common state. For players announcing a state $\phi' \neq \phi$ it is a better response to announce $(\phi, f(\phi))$. This yields $f(\phi)$, from which one can never go back to the allocation under (ii.b) or (iii.b) with better-response dynamics.

Step 3: No message profile in rule (ii.a) or (iii.a) is part of a recurrent class if the common announcement is θ . Let the alternative announcement be ϕ . In rule (ii.a) the agent deviating from the common announcement is agent $i(\theta, x(\theta, \phi))$, who obtains $x_{i(\theta, x(\theta, \phi))}(\theta, \phi)$, which by definition of $i(\theta, z)$ is such that $f(\theta) \succ_{i(\theta, x(\theta, \phi))}^{\theta} x_{i(\theta, x(\theta, \phi))}(\theta, \phi)$ and so for this agent announcing $(\theta, f(\theta))$ would be a best-response. From there, going back to rule (ii.a) would never happen under better-reply dynamics. Similarly, for rule (iii.a.) agent $j(\theta, \phi)$, obtains $y_{j(\theta, \phi)}(\theta, \phi)$ and by (1) $x(\theta, \phi) \succ_{j(\theta, \phi)}^{\theta} y(\theta, \phi)$, hence announcing $(\theta, f(\theta))$ would be a best-response for agent $j(\theta, \phi)$. At this point, as we argued before, announcing $(\theta, f(\theta))$ would be a best-response for agent $i(\theta, x(\theta, \phi))$.

Recall that θ is the true state. Next, we can classify the recurrent classes into three categories:

Denote by E_0 the recurrent class of BRD in which all n agents report the true state as the first part of their announcement. Note that there are multiple states within this truthful recurrent class, as agents can disagree on the allocation reported. And denote by E_j , $j = 1, \dots, k$, a typical recurrent class consisting of a profile under rule (i), where agents' unanimously reported state is θ^j , which is not θ , the true state. Finally, classes $E_{k+1}, \dots, E_{k+k'}$ comprise the possible recurrent states under rule (ii.a) or (iii.a) where the common announcement by $n - 1$ people is not the true preference profile.

For any two states m and m' , one can now define the resistance of the transition $m \rightarrow m'$ as the number of mistakes involved. We wish to show that the stochastically stable states of perturbed BRD in the game under uniform mistakes are precisely the states in the class E_0 . To show this, it will suffice to make the following observations:

[a] To get out of the class E_0 , we need some agent $i(\theta, x(\theta, \phi))$ to impose one of the reversal outcomes $x(\theta, \phi)$ – one mistake, as by definition this individual is worse off. Next,

$j(\theta, \phi)$ imposes $y(\theta, \phi)$ – second mistake, in this case by equation (1). Finally, anyone else changes and we go to rule (iv) where 0 is the outcome – third mistake. From 0, we go for free to any of the other recurrent classes. There are other paths as well, going first to (ii.b), and from there to (iii.b), and then to (iv), but all those also require three mistakes.

[b] To get out of any of the recurrent classes with untruthful profiles E_j , $j = 1, \dots, k$, (say $m^1 = \phi$ is one such profile when the true state is θ), one can take the following path: an agent $i(\phi, x(\phi, \theta))$ can impose $x(\phi, \theta)$. At this point, either $f(\phi) \succ_{i(\phi, x(\phi, \theta))}^\theta x(\phi, \theta)$, in which case this step requires a first mistake, or $x(\phi, \theta) \succeq_{i(\phi, x(\phi, \theta))}^\theta f(\phi)$, in which case this step has zero resistance. Next, agent $j(\phi, \theta)$ changes the outcome to $y(\phi, \theta)$ for free. Finally, someone else changes the outcome to 0 under rule (iv), which constitutes at most a second mistake. From there, we go for free to any of the other recurrent classes.

[c] To get out of any of the recurrent classes $E_{k+1}, \dots, E_{k+k'}$, where the common profile is $\phi \neq \theta$, announced by $n - 1$ or $n - 2$ agents, and the alternative announcement is ϕ' , let any agent who announces ϕ deviate to announcing ϕ'' . This is a mistake and leads to rule (iv). From there, we can go for free to E_0 .

Therefore, by [b] and [c] one can construct an E_0 -tree in which the resistance of each of the edges (E_j, E_0) , $j = 1, \dots, k + k'$ is at most 2. The resistance of such a tree is at most $2(k + k')$. On the other hand, any E_j -tree ($j = 1, \dots, k + k'$) must include an edge (E_0, E_m) of resistance 3 (by [a]). This fact, together with [b] and [c] for all the other edges in the tree, implies that the resistance of the E_j -tree is no less than $2(k + k') + 1$. We conclude that E_0 is the class of minimum stochastic potential, and thus, it contains all stochastically stable states. ■

Remark: If one assumes that the preference reversals specified in equation (1) occur “near enough the zero bundle,” one can show, using a similar proof, that for $n \geq 5$ any ε -secure SCF is implementable in stochastically stable strategies of a perturbed BRD based on uniform mistakes. In this sense, one can clearly interpret Theorem 2 as a very permissive result.

Remark: It appears that, to obtain meaningful implementability results using uniform mistakes, one needs to add at least a new rule to the canonical mechanism used for the result based on “more serious mistakes are less likely” of our working paper (also used in Theorem 2 of CS). Note how the proof has relied heavily on the use of the preference reversal

specified in equation (1). On the other hand, the economic environment is not essential. A mechanism very similar to the one we present but using modulo games and allowing for some punishments, based on the NWA condition of CS, would also work in non-economic environments.

4 Incomplete Information

This section tackles the extension of our results to incomplete information environments.

Each agent knows his type $\theta_i \in \Theta_i$, a finite set of possible types. Let $\Theta = \prod_{i \in N} \Theta_i$ be the set of possible states of the world, let $\Theta_{-i} = \prod_{j \neq i} \Theta_j$ of type profiles θ_{-i} of agents other than i . We shall sometimes write a state $\theta = (\theta_i, \theta_{-i})$. We assume that all states in Θ have positive ex-ante probability.⁸

Let $q_i(\theta_{-i}|\theta_i)$ be type θ_i 's interim probability distribution over the type profiles θ_{-i} of the other agents. An SCF (or state-contingent allocation) is a mapping $f : \Theta \mapsto Z$ that assigns to each state of the world a feasible allocation.

Let A denote the set of SCFs. We shall assume that uncertainty concerning the states of the world does not affect the economy's endowments, but only preferences and beliefs.

We shall write type θ_i 's interim expected utility over an SCF f as follows:

$$U_i(f|\theta_i) \equiv \sum_{\theta_{-i} \in \Theta_{-i}} q_i(\theta_{-i}|\theta_i) u_i(f(\theta_i, \theta_{-i}), (\theta_i, \theta_{-i})).$$

Note how the Bernoulli (ex-post) utility function u_i may change with the state. We shall use the obvious versions of Assumptions (1) and (2) applied to each ex-post utility function in each state.

A mechanism $G = ((M_i)_{i \in N}, g)$, played simultaneously by myopic agents, consists of agent i 's set M_i of messages (for each $i \in N$, agent i 's message is a mapping from Θ_i to M_i), and the outcome function $g : M \mapsto Z$. The direct mechanism for the SCF f is a mechanism in which for all i , $M_i = \Theta_i$ and where $g = f$. A Bayesian equilibrium is a message profile in which each type chooses an interim best-response to the other agents' messages, and a strict Bayesian equilibrium is a Bayesian equilibrium in which every type's interim best-response is a strict best-response. To prevent any kind of learning about the state, we shall assume that, after an outcome is observed, agents forget it (or, closer to the evolutionary tradition,

⁸We make this assumption for simplicity in the presentation. With some minor modifications in the arguments, one can prove similar results if $\Theta^* \neq \Theta$ is the set of states with positive probability, according to every agent's prior belief.

agents are replaced by other agents who share the same preferences and prior beliefs as their predecessors, but are not aware of their experience).⁹

Let agent i of type θ_i be allowed to revise his message in period t . He does so using the interim better-response logic, i.e., he switches with positive probability to any message that improves (weakly) his interim expected utility, given his interim beliefs $q_i(\theta_{-i}|\theta_i)$. That is, letting m^t be the message profile at the beginning of period t , type θ_i switches from $m_i^t(\theta_i)$ to any m_i' such that:

$$\sum_{\theta_{-i} \in \Theta_{-i}} q_i(\theta_{-i}|\theta_i) u_i(g(m_i', m_{-i}^t(\theta_{-i})), (\theta_i, \theta_{-i})) \geq \sum_{\theta_{-i} \in \Theta_{-i}} q_i(\theta_{-i}|\theta_i) u_i(m^t(\theta), (\theta_i, \theta_{-i})).$$

We adapt now the definitions of implementability to environments with incomplete information (the definition of implementability in recurrent strategies is borrowed from CS):

An SCF f is *implementable in recurrent strategies* (of interim BRD) if there exists a mechanism G such that the interim BRD process applied to its induced game has f as its unique outcome of the recurrent classes of the process.

An SCF f is *implementable in stochastically stable strategies* (of perturbed interim BRD) if there exists a mechanism G such that a perturbation of the interim BRD process applied to its induced game has f as the unique outcome supported by stochastically stable strategy profiles.

4.0.1 Necessity

As for the assumptions on SCFs, we still assume that it is ε -secure in each state, although this will not be a necessary condition. In contrast, we shall introduce two more properties, which will be necessary for implementability in recurrent strategies. The next one is the strict version of incentive compatibility.

An SCF f is *strictly incentive compatible* if truth-telling is a strict Bayesian equilibrium of its direct mechanism, i.e., if for all i and for all θ_i ,

$$\sum_{\theta_{-i} \in \Theta_{-i}} q_i(\theta_{-i}|\theta_i) u_i(f(\theta), (\theta_i, \theta_{-i})) > \sum_{\theta_{-i} \in \Theta_{-i}} q_i(\theta_{-i}|\theta_i) u_i(f(\theta'_i, \theta_{-i}), (\theta_i, \theta_{-i}))$$

for every $\theta'_i \neq \theta_i$.

An SCF f is *incentive compatible* if the inequalities in the preceding definition are allowed to be weak.

⁹There are a host of alternative assumptions one could make, for example, that each agent receives his type in each period as a draw from the i.i.d. underlying distribution; see Dekel, Fudenberg and Levine (2004) for an appraisal of such different modeling choices.

As it turns out, incentive compatibility is an important necessary condition for any kind of implementability in our sense.

THEOREM 3 *If f is implementable in stochastically stable strategies of an arbitrary perturbation of an unperturbed interim BRD process, f is incentive compatible. Furthermore, if at least one of the recurrent classes selected by the perturbation of the interim BRD is a singleton, f is strictly incentive compatible.*

Proof: Suppose that f is implementable in stochastically stable strategies of an arbitrary perturbation of BRD. This means that, for this perturbed process, there is a unique outcome supported by at least one of the recurrent classes of the unperturbed process, and this outcome is f . Since f is the outcome of such a recurrent set of BRD, it must be incentive compatible.

Furthermore, if one of the recurrent classes selected by the perturbation is a singleton, any deviation from the message profile that is an absorbing state of the unperturbed dynamics must worsen each type's interim expected utility, and thus, f must be strictly incentive compatible. ■

4.1 Sufficiency

Consider a strategy in a direct mechanism for agent i , i.e., a mapping $\alpha_i = (\alpha_i(\theta_i))_{\theta_i \in \Theta_i} : \Theta_i \mapsto \Theta_i$. A *deception* $\alpha = (\alpha_i)_{i \in N}$ is a collection of such mappings where at least one differs from the identity mapping.

Given an SCF f and a deception α , let $[f \circ \alpha]$ denote the following SCF: $[f \circ \alpha](\theta) = f(\alpha(\theta))$ for every $\theta \in \Theta$.

Finally, for a type $\theta'_i \in \Theta_i$, and an arbitrary SCF y , let $y_{\theta'_i}(\theta) = y(\theta'_i, \theta_{-i})$ for all $\theta \in \Theta$.

We shall make the following additional assumptions on environments:

- (6) For every deception α , there exists an agent $i \in N$, a type $\theta_i \in \Theta_i$, a strictly incentive compatible SCF x , and another SCF y such that

$$U_i(x \mid \theta_i) > U_i(y_{\theta'_i} \mid \theta_i) \forall \theta'_i \in \Theta_i \quad \text{and} \quad U_i(x \circ \alpha \mid \theta_i) \leq U_i(y \circ \alpha \mid \theta_i). \quad (2)$$

- (7) The bundles in the SCFs x and y used in (2) are componentwise no greater than ε .

In words, Assumption (4) says that the environment admits preference reversals to overcome deceptions. However, these preference reversals need not happen around f , the SCF

of interest, but around some strictly incentive compatible SCF x ; see Serrano and Vohra (2005) for an appraisal of this assumption.

For each deception α , we shall choose one test-pair x, y and one test-agent i , satisfying the conditions in (2). Denote the set of all such x by D . Finally, with very little loss of generality, choose the bundles in the SCFs y consisting of strictly positive amounts of each commodity. Then, define the SCF $[y]$ as the one that assigns in each state the componentwise minimum bundle for each agent i and each state θ : $[y]_i(\theta) \leq y_i(\theta)$ for all y .

On the other hand, Assumption (5) says that such reversals happen “near enough the zero bundle.”¹⁰ Then, one can make use of the insight in the last remark of the previous section to show our next result:

THEOREM 4 *Suppose that the environments satisfy Assumptions (1), (2), (6) and (7). Let $n \geq 5$. Let f be ε -secure in every state and strictly incentive compatible. Then, f is implementable in stochastically stable strategies of perturbed interim BRD under uniform mistakes.*

Proof: The proof follows steps similar to that of Theorem 2, but applied to the following mechanism. Let agent i 's message set be $M_i = \Theta_i \times A$. Denote a typical message sent by agent i by $m_i = (m_i^1, m_i^2)$ and the corresponding message profile by $m = (m^1, m^2)$. The outcome function obeys the following rules:

- (i) If for every $i \in N$, $m_i^2 = f$, $g(m) = f(m^1)$.
- (ii.a) If exactly $(n - 1)$ messages m_j are such that $m_j^2 = f$ and $m_i^2 = x$ for some $x \in D$, $g(m) = x(m^1)$.
- (ii.b) If exactly $(n - 1)$ messages m_j are such that $m_j^2 = f$ and $m_i^2 = x$ for some $x \notin D$, $g(m) = (f_i(m^1) - \beta, f_{-i}(m^1))$, where $f_i(\cdot) \geq f_i(\cdot) - \beta \geq (\varepsilon, \dots, \varepsilon)$.
- (iii.a) If exactly $(n - 2)$ messages m_k are such that $m_k^2 = f$, $m_i^2 = x$ for some $x \in D$ and $m_j^2 = y$ where j and y are the ones associated with x as in (**), $g(m) = y(m^1)$.
- (iii.b) If exactly $(n - 2)$ messages m_k are such that $m_k^2 = f$, but the other conditions of rule (iii.a) are not met, $g(m) = [y](m^1)$.

¹⁰In fact, if the environment allowed the use of lotteries, and making use of expected utility, one could combine the SCFs x and y in a mixture with the zero bundle, where the latter is imposed with arbitrarily high probability. This argument would allow one to take the SCFs x and y arbitrarily “near the zero bundle” without assuming it explicitly, as we do.

(iv) In all other cases, $g(m) = 0$.

We sketch the steps of the proof as follows. First, one can show that all recurrent classes of interim BRD are under rule (i). For example, to see how rule (iv) is never part of a recurrent class, use a simultaneous switch of all types to $m_i^2 = f$, and so on; similar arguments apply to rules (ii) and (iii). Within rule (i), strict incentive compatibility allows one to support truth-telling as one of these (singleton in this case) recurrent classes, but there may well be others, in which agents are using a deception α .

To finish the sketch of proof, here is a heuristic argument. One can describe the transition paths among the different recurrent states. To get out of the absorbing state in which agents are telling the truth in their first part of the announcement, one can go through rule (ii.a), which requires one mistake because any $x \in D$ is near the origin (note that any agent can be used for this mistake, by strictly increasing preferences in each state). Next, the test-agent corresponding to that x will implement rule (iii.a), where we require a second mistake. Finally, someone else makes a mistake and we go to rule (iv). A similar path can be created for each state to get to the profile of zero bundles. There are other paths one could follow: for example, through rules (ii.b) and (iii.b), but the point is that each time an agent switches to change the outcome in the direction of the zero profile, a mistake is required.

On the other hand, if one starts at an absorbing state in which a deception is being used, one gets out through any agent other than the test-agent for that deception and imposes rule (ii.a), which requires one mistake. The next step, taken by the test-agent for that deception, is free because of equation 2. From rule (iii.a), someone else changes to rule (iv), and so on. In this path, we have “saved” one mistake. Of course, from the zero profile, we go for free to any of the other absorbing states.

These arguments allow the construction of the corresponding spanning trees for each absorbing state. The result is that the truthful absorbing state is the only one of minimum stochastic potential, i.e., the only one that is stochastically stable. ■

5 Conclusion

The results presented here complement those in CS. Restricting attention to economic environments, we have studied implementation under perturbed better-response dynamics. In the working paper version of this study, we show that, for a variant of “more serious mistakes are less likely,” any ε -secure SCF is implementable when there are at least three agents. For

uniform mistakes, we have shown here that any ε -secure and strongly efficient SCF is implementable when there are at least five agents. Extensions of results to incomplete information environments have also been obtained, including the emergence of incentive compatibility as a necessary condition for any kind of robust implementation in our sense.

References

- Bergin, J. and B. Lipman (1996), "Evolution with State Dependent Mutations," *Econometrica* **64**, 943-956.
- Cabrales, A. and R. Serrano (2010), "Implementation in Adaptive Better-Response Dynamics: Towards a General Theory of Bounded Rationality in Mechanisms," *Games and Economic Behavior*, forthcoming.
- Dekel, E., D. Fudenberg and D. K. Levine (2004), "Learning to Play Bayesian Games," *Games and Economic Behavior* **46**, 282-303.
- Ellison, G. (2000), "Basins of Attraction, Long-Run Stochastic Stability, and the Speed of Step-by-Step Evolution," *Review of Economic Studies* **67**, 17-45.
- Freidlin, M. and A. Wentzell (1984), *Random Perturbations of Dynamical Systems*, Springer-Verlag, Berlin.
- Kandori, M., G. Mailath, and R. Rob (1993), "Learning, Mutations and Long Run Equilibria in Games," *Econometrica* **61**, 29-56.
- Karlin, S., H.M. Taylor (1975) *A First Course in Stochastic Processes*, NY: Academic Press.
- Serrano, R. and R. Vohra (2005), "A Characterization of Virtual Bayesian Implementation," *Games and Economic Behavior* **50**, 312-331.
- Young, H. P. (1993), "The Evolution of Conventions," *Econometrica* **61**, 57-84.
- Young, H. P. (1998), *Individual Strategy and Social Structure: an Evolutionary Theory of Institutions*, Princeton university Press, Princeton, NJ.