Implementation in Adaptive Better-Response Dynamics*

Antonio Cabrales† Roberto Serrano ‡

July 2007

Abstract

We study the classic implementation problem under the behavioral assumption that agents myopically adjust their actions in the direction of better-responses within a given institution. We offer results both under complete and incomplete information. First, we show that a necessary condition for assymptotically stable implementation is a small variation of (Maskin) monotonicity, which we call quasimonotonicity. Under standard assumptions in economic environments, we also provide a mechanism for Nash implementation which has good dynamic properties if the rule is quasimonotonic. Thus, quasimonotonicity is both necessary and almost sufficient for assymptotically stable implementation. Under incomplete information, incentive compatibility is necessary for any kind of stable implementation in our sense, while Bayesian quasimonotonicity is necessary for assymptotically stable implementation. Both conditions are also essentially sufficient for assymptotically stable implementation. We then tighten the assumptions on preferences and mutation processes and provide mechanisms for stochastically stable implementation under more permissive conditions on social choice rules.

JEL Classification: C72, D70, D78.

Keywords: implementation; bounded rationality; evolutionary dynamics; mechanisms.

---

*Cabrales gratefully acknowledges the financial support from the Spanish Ministry of Science and Technology under grants CONSOLIDER INGENIO 2010 (CSD2006-0016), and SEJ2006-11665-C02-00. Serrano thanks CEMFI in Madrid for its hospitality.

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

‡Department of Economics, Brown University, and IMDEA-Ciencias Sociales; Email: roberto_serrano@brown.edu; http://www.econ.brown.edu/faculty/serrano

1 Introduction

The correct design of institutions can be decisive for achieving economic systems with good welfare properties. But suppose that the correct design depends on the knowledge of key parameters in the environment. Then, an important problem ensues if the builder of the institutions does not have such knowledge. The theory of implementation looks in a systematic way at the design of rules for social interaction that do not assume a detailed knowledge of the fundamentals by those with power to adjudicate social outcomes.

The last decade saw impressive advances in the theory of implementation.[1] As Sjöström (1994) pointed out, 'With enough ingenuity the planner can implement "anything"'. On the other hand, several recent contributions (Cabrales (1999), Cabrales and Ponti (2000), Sandholm (2002), Eliaz (2002)) have highlighted the fact that not all mechanisms perform equally well, in terms of achieving the socially desirable outcomes. In particular, some of the mechanisms that are more permissive in terms of the span of implementable social choice functions may lead to dynamic instability or convergence to the wrong equilibrium when the players are boundedly rational. To be more precise, mechanisms for Nash implementation have been shown to be robust to boundedly rational agents. On the other hand, mechanisms for implementation under iterative deletion of dominated strategies, or for subgame-perfect implementation have bad dynamic properties (instability, convergence to the "wrong" equilibrium).

Given these findings, it is natural to ask whether the difficulty with permissive mechanisms lies in the particular mechanisms employed, or if it is a general problem. In other words, what are the necessary conditions for evolutionary implementation?

In this paper we pose the classic implementation questions for a class of evolutionary settings.[2] We postulate a behavioral assumption by which agents (or generations thereof) interact myopically within a given institution, and adjust their actions in the direction of better-responses within the mechanism. Our criterion for successful implementation will be the convergence of the better-response process to a rest point or to a set of rest points. When the outcomes of a social choice function (SCF) are the only limits of the better-

[1]See Jackson (2001), Maskin and Sjöström (2002), Palfrey (2002) or Serrano (2004) for recent surveys.

[2]References for evolutionary game theory in general are Weibull (1995), Vega-Redondo (1996), Samuelson (1997), Fudenberg and Levine (1998) and Young (1998). The stochastic stability methodology, which will be used for many of our results, is based on the techniques developed by Freidlin and Wentzell (1984), and it was first applied to evolutionary biology by Foster and Young (1990), and to game theory by Kandori, Mailath and Rob (1993) and Young (1993).

response dynamics of a mechanism for any allowed environment, we shall say that the SCF is *implementable in asymptotically stable strategies*.

The environments we shall be concerned with are *economic*. An amount of L commodities is to be allocated among n agents. Preferences are strictly increasing in all commodities, which implies that the zero bundle is the worst outcome for everyone in the economy. The typical mechanism that we shall construct will have good dynamic properties. It implements the socially desirable outcome according to the agents' reports, if there is total agreement among them. If there is an almost unanimous agreement, other outcomes will be implemented. Those outcomes are meant to elicit the "right" behavior from agents. Finally, if enough disagreement occurs in the reports, no goods will be allocated and every agent will receive the zero bundle. This form of severe punishment exploits the economic nature of the environment. Nonetheless, it allows us to avoid the use of integer or modulo games.

A necessary condition for asymptotically stable implementation is a small variation of (Maskin) monotonicity (Maskin (1999)), which we call *quasimonotonicity*. Quasimonotonicity prescribes that the social outcome not change if the *strictly* lower contour sets of preferences at the social outcome are nested for every agent across two environments. In particular, it is neither logically stronger nor weaker than monotonicity; both coincide when indifference sets are singletons, or more generally, when preferences are continuous. Furthermore, quasimonotonicity is also sufficient for asymptotically stable implementation, if there are at least three agents in the environment and the SCF is ϵ-*secure*. The condition of ϵ-security stipulates that each agent must be allocated by the rule a bundle consisting of at least ϵ units of each commodity.

Our results on asymptotic implementation are obtained for a general class of preferences and will stand for any mutation process. The latter means that, if one were to perturb the better-response dynamics via mutations, an SCF that is implementable in asymptotically stable strategies would also be *implementable in stochastically stable strategies* of any perturbation of better-response dynamics. That is, the outcomes prescribed by the SCF are the states of minimum stochastic potential, for any perturbed process. In this way, these conclusions are immune to the Bergin and Lippman (1996) critique of uniqueness results in stochastic evolutionary implementation.

Next, we strengthen the assumptions on preferences and mutation processes, and we show that there are mechanisms for evolutionary implementation under relatively permissive conditions on SCF's. Specifically, we offer two such results. The first shows that, under a variant of the "more serious mistakes are less likely" assumption, any ϵ-secure SCF is implementable in stochastically stable strategies of the corresponding perturbed better-response process if there are at least three agents. The second states that, under uniform mutations and a rather innocuous 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 zero, the Pareto assumption can be dispensed with altogether.

We do not wish to interpret our findings in this paper 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. 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. 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 better-response dynamics. If one wishes to implement in asymptotically stable strategies, faithful to the Bergin-Lipman line of thinking, the condition of Bayesian quasimonotonicity is also necessary. The comparison between this condition and Bayesian monotonicity, necessary for Bayesian implementation (e.g., Postlewaite and Schmeidler (1986), Palfrey and Srivastava (1989), Jackson (1991)), is similar to that between quasimonotonicity and Maskin's condition. Moreover, incentive compatibility, Bayesian quasimonotonicity and ϵ-security are also sufficient for implementation in asymptotically stable strategies of better-response processes when there are at least three agents. Under a weak diversity of preferences on 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 mutations, and if there are at least five agents. Thus, we find the same tradeoff enunciated 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.

Section 2 describes the model and the dynamics we use. Section 3 provides necessary and sufficient conditions for asymptotically stable implementation under complete information. Section 4 presents more permissive results for stochastically stable implementation under specific mutation processes. Section 5 collects the extensions of our results to incomplete information environments. Section 6 concludes.

2 Preliminaries

Let $N = \{1, \ldots, n\}$ be a set of agents. Let the environment be an exchange economy with a grid as its set of allocations. (For example, because of the existence of an indivisible unit for each commodity.) Let agent i's consumption set be $X_i \subset \mathbb{R}^l_+$. 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 allocations is

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

Let us now specify agents' preferences over allocations. Let θ_i denote agent i's preference ordering, on which we shall make the following assumptions:

(1) No externalities: $\theta_i : X_i \times X_i \mapsto X_i$.

(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.

Let $\theta = (\theta_i)_{i \in N}$ be a preference profile, and Θ be the set of allowable preference profiles.

A mechanism $G = ((M_i)_{i \in N}, g)$, where M_i is agent i's message set, and $g : \prod_{i \in N} M_i \mapsto Z$ is the outcome function.

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.

In this paper we shall concentrate on the following class of SCF's. An SCF f is said to be *ϵ-secure* if there exists $\epsilon > 0$ such that for each θ, and for each $i \in N$, $f_i(\theta) \geq (\epsilon, \ldots, \epsilon) \gg 0$.

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 it on normative grounds.

In addition, we shall use the following condition, which turns out to be more fundamental to the theory we develop here:

An SCF f is *quasimonotonic* if, whenever it is true that for every $i \in N$ $f(\theta) \succ_i^{\theta} z$ implies that $f(\theta) \succ_i^{\phi} z$, we have that $f(\theta) = f(\phi)$ for all $\theta, \phi \in \Theta$.

Note how quasimonotonicity resembles closely the condition of monotonicity uncovered in Maskin (1999). Indeed, the only difference is that, while Maskin's condition imposes that

[3] For 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.

the lower contour sets of preferences be nested across two environments, quasimonotonicity relies on the inclusion of the strictly lower contour sets.

Next, we turn to dynamics, the central approach in our paper.

The unperturbed Markov process that we shall impose on the play of the mechanism over time is the following better-response dynamics. In each period t one of the agents is chosen at random with positive and fixed probability so that he can revise his message or strategy. 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 $g(m'_i, m_{-i}(t)) \succeq_i^{\theta_i} g(m(t))$.

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, will eventually converge to the socially desirable outcome as specified by the SCF. This logic suggests the two notions of implementability that we shall employ in the current paper.

An SCF f is *implementable in asymptotically stable strategies* (of better-response dynamics) if there exists a mechanism G such that, for every $\theta \in \Theta$, the unique outcome of all recurrent classes of the better response process applied to its induced game when the preference profile is θ is $f(\theta)$.

After the planner solves the question of implementability in the sense just defined, he can consider "refinements" of such a notion, by allowing specific perturbations of the better-response dynamics. This will lead to the concept of implementability in stochastically stable strategies.

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

3 Necessary and Sufficient Conditions for asymptotically Stable Implementation: Complete Information

3.1 Necessity

We first seek implementation in asymptotically stable strategies, i.e., without relying on perturbations of the better response dynamics. For example, suppose one is interested in seeking robustness in stochastically stable implementation, in the sense that one would like the implementation to be successful with independence of the perturbation used – requiring implementability in this sense is desirable, given the results in Bergin and Lipman (1996). We wish to show now that quasimonotonicity of f is a necessary condition for its implementability in asymptotically stable strategies.

THEOREM **1** *If f is implementable in asymptotically stable strategies of an unperturbed better-response dynamic process, f is quasimonotonic.*

Proof 1 *Let the true preference profile be θ. Because f is implementable in asymptotically stable strategies of better-response dynamics, the only outcome that corresponds to strategy profiles in recurrent classes of the dynamics is $f(\theta)$.*

Now consider a preference profile ϕ such that for all i $f(\theta) \succ_i^\theta z$ implies that $f(\theta) \succ_i^\phi z$. Since the only outcome compatible with recurrent classes of the dynamics when preferences are θ is $f(\theta)$, this means that agent i's unilateral deviations from recurrent strategy profiles must yield either $f(\theta)$ again, or outcomes z such that $f(\theta) \succ_i^\theta z$. But this implies that $f(\theta)$ is also supported by recurrent profiles of better-response dynamics when preferences are ϕ. Since f is implementable in asymptotically stable strategies of better-response dynamics, this implies that $f(\phi) = f(\theta)$. That is, f must be quasimonotonic.

3.2 Sufficiency

We now present our next result. Together with Theorem 1, it provides almost a characterization of the SCF's that are implementable in asymptotically stable strategies, and therefore, implementable in stochastically stable strategies independently of the perturbation used.

THEOREM **2** *Suppose the environments satisfy Assumptions (1) and (2). Let $n \geq 3$. If an SCF f is ϵ-secure and quasimonotonic, it is implementable in asymptotically stable strategies of better-response dynamics.*

Proof 2 *Consider the mechanism $G = ((M_i)_{i \in N}, g)$, where agent i's message set is $M_i = \Theta \times Z$, and the outcome function g is defined by the following rules:*

(i.) *If for all $i \in N$, $m_i = (\theta, f(\theta))$, $g(m) = f(\theta)$.*

(ii.) *If for all $j \neq i$, $m_j = (\theta, f(\theta))$ and $m_i = (\phi, z) \neq (\theta, f(\theta))$, 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* $f(\theta) \succ_i^\theta z$, $g(m) = z$.

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

We begin by arguing in the next four steps that all recurrent classes of the unperturbed better-response process are singletons. 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 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. 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.*

Step 4: Thus, all recurrent classes contain only profiles under rule (i). However, one cannot abandon a profile under rule (i) to get to another under the same rule without passing through rule (ii). This establishes that all recurrent classes are singletons.

Moreover, each recurrent class, a singleton under rule (i), must consist of a Nash equilibrium of the game induced by the mechanism when the true preferences are θ*. Otherwise, one would not have an absorbing state of better-response dynamics.*

One such Nash equilibrium that always exists is the truthful profile $(\theta, f(\theta))$ reported by every agent. Note how unilateral deviations from this profile produce an outcome either under rule (ii.a) or under rule (ii.b). In either case, no such switch can happen under better-response dynamics.

In addition, one may have other (non-truthful) Nash equilibria under rule (i). Let $(\phi, f(\phi))$ be such an arbitrary Nash equilibrium. For this profile to be a Nash equilibrium, it must be true that for all $i \in N$, $f(\phi) \succ_i^\phi z$ implies that $f(\phi) \succeq_i^\theta z$, since otherwise such a z, which could be induced by i in the mechanism, would contradict that the profile is a Nash equilibrium. But in fact, we know even more. Since the profile is actually an absorbing state of the dynamics, we must have that for all $i \in N$, $f(\phi) \succ_i^\phi z$ implies that $f(\phi) \succ_i^\theta z$: that is, such allocations z that i could induce cannot be indifferent to $f(\phi)$ under θ, or the profile $(\phi, f(\phi))$ would not be absorbing.

Thus, because f is quasimonotonic, we must have that $f(\theta) = f(\phi)$ for any arbitrary absorbing state $(\phi, f(\phi))$ of the better-response dynamics. Therefore, f is implementable in asymptotically stable strategies, or equivalently, in stochastically stable strategies no matter what perturbation of better-response dynamics one takes.

4 Permissive Results

Thus far we have seen that quasimonotonicity is the key condition that essentially characterizes implementability in asymptotically stable strategies in economic environments with at least three agents. In this section, we explore the possibilities of implementing non-quasimonotonic rules under extra specific assumptions on preferences and mutations.

4.1 A Perturbed Process with Non-Uniform Mutations

This subsection explores the possibilities of obtaining a more permissive implementation result, by imposing a specific kind of perturbation of better-response dynamics. 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.

For this section we need the following additional assumptions on preferences:

(3) Let commodity 1 be a nummeraire whose indivisible unit is $\Delta > 0$. The preference is quasilinear in the nummeraire. Also, let $\Delta > 0$ be smaller than any utility gap resulting from reallocations of the non-nummeraire commodities.

(4) The preference is continuous.

Assumption (3) is needed because we use the penalties in the nummeraire (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.

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.

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 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 what happens in the mechanism.

Consider a perturbation of better-response dynamics, 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.[4]

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

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

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 "reluctant" to use rule (iii) in the mechanism of the proof of Theorem 2.)

THEOREM **3** *Suppose the environments satisfy Assumptions (1), (2), (3) and (4). Let* $n \geq 3$. *Then, any* ϵ*-secure SCF* f *is implementable in stochastically stable strategies of the prescribed perturbed better-response dynamics.*

Proof 3 *Consider the mechanism* $G = ((M_i)_{i \in N}, g)$ *proposed in the proof of Theorem 2. This proof will make extensive use of it.*

In particular, recall that, without relying on quasimonotonicity of f*, we showed in the first steps of that proof that all recurrent classes of the better-response dynamics are singletons under rule (i). Let* θ *denote the true preferences. The truthful strategy profile,* $((\theta, f(\theta)), \ldots, (\theta, f(\theta)))$*, is always one of these absorbing states, and in addition, there may exist absorbing profiles* $((\phi, f(\phi)), \ldots, (\phi, f(\phi)))$ *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 strategy profile, i.e., for each* $i \in N$*,* $m_i = (\theta, f(\theta))$*.*

Class E_j *for* $j = 1, \ldots, J$ *is the coordinated lie on profile* θ^j*: for each* $i \in N$*,* $m_i = (\theta^j, f(\theta^j))$*, known to be a Nash equilibrium 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*.*

Consider the following modification made to the outcome function of the mechanism in the proof of Theorem 2:

(ii.a'.) *Replace* β *with the vector* $(\Delta, 0, \ldots, 0)$*, i.e., the punishment takes the form of the smallest indivisible unit of the nummeraire commodity.*

Now, we show that the profile in E_0 *is the only stochastically stable profile of the perturbed dynamics:*

[a] to get out of E_0*, one can go through rule (ii.a') of the mechanism, paying* $(1 + \lambda)\Delta$ *if the deviator* i *proposes an outcome indifferent to* $f(\theta)$*, or go through rule (ii.b) paying a cost*

that is exactly $(1+\lambda)$ *times the smallest utility loss from* $f(\theta)$ *to* z*, 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* $(\theta^j, f(\theta^j))$*, and instead announce* (ϕ, z) *such that* $z \succeq_i^{\theta^j} f(\theta^j)$ *and* $f(\theta^j) \succ_i^{\theta} z$*. In this case, the resistance is strictly smaller than* $(1+\lambda)\Delta$*, 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 mutations are somewhat special, we think it is interesting that Theorem 3 dispenses with quasimonotonicity, while still making use of the same mechanism as does Theorem 2.

4.2 A Perturbed Process Based on Uniform mutations

To obtain a sufficiency result based on stochastic stability of perturbed better-responses under uniform mutations, we use an additional assumption on the SCF, i.e., that it is efficient:[5]

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 $f(\theta) \succ_{i(\theta,z)}^{\theta} z$.

In addition, we modify slightly our assumptions on the environment, as follows.

First, note that since states differ because at least one of the agent's 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

$$x(\theta,\phi) \succ_{j(\theta,\phi)}^{\theta} y(\theta,\phi) \qquad \text{and} \qquad y(\theta,\phi) \succeq_{j(\theta,\phi)}^{\phi} x(\theta,\phi). \tag{$*$}$$

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 (*).

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:

[5] As 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.

(5) 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 (*).

Now, we can prove the following result:

THEOREM **4** *Suppose the environments satisfy Assumptions (1), (2) and (5). Let $n \geq 5$. Any ϵ-secure and strongly Pareto efficient SCF f is implementable in stochastically stable strategies of perturbed better-response dynamics, where the perturbation consists of uniform mutations.*

Proof 4 *Consider the following mechanism. Let agent i's message set be $M_i = \Theta \times Z$. 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 $i \in N$, $m_i^1 = \theta$, $g(m) = f(\theta)$.*

(ii.a.) *If exactly $(n-1)$ messages m_i are such that $m_i^1 = \theta$ and $m_{i(\theta,x(\theta,\phi))} = (\phi, x(\theta,\phi))$, $g(m) = (x_{i(\theta,x(\theta,\phi))}(\theta,\phi), x_{j(\theta,\phi)}(\theta,\phi), 0, 0, \ldots, 0)$.*

(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 $f_k(\theta) \geq f_k(\theta) - \beta \geq (\epsilon, \ldots, \epsilon)$.*

(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 $m_{j(\theta,\phi)} = (\phi, y(\theta,\phi))$, $g(m) = (y_{i(\theta,x(\theta,\phi))}(\theta,\phi), y_{j(\theta,\phi)}(\theta,\phi), 0, 0, \ldots, 0)$.*

(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 $k \in N$, $g_k(m) = (\epsilon, \ldots, \epsilon)$.*

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

(For rules (iii.a) and (iii.b) to be well defined, the assumption $n \geq 5$ is needed to determine the outcome in profiles where two agents report the same state as part of their message and two other agents report a different state.)

First, we can follow steps similar to those in the proof of Theorem 2 to show that all recurrent classes of unperturbed better-response dynamics are singletons, i.e., absorbing states, and that all of them happen under rule (i).

Let θ be the true state. Next, we can classify the absorbing states into two categories:

Denote by E_0^j a typical singleton recurrent class –absorbing state– of better-response dynamics in which all n agents report the true state as the first part of their announcement. Note that there are multiple such states, as agents can disagree on the allocation reported. And denote by E_1^j a typical singleton recurrent class consisting of a uniform profile under rule (i), where agents' reported state is not θ, the true state.

We wish to show that the stochastically stable states of better-response dynamics in the game under uniform mutations are precisely the states in the classes E_0^j. To show this, it will suffice to make the following observations:

[a] *To get out of any state E_0^j, 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 mutation, in this case by equation (*). Finally, anyone else changes and we go to rule (iv) where 0 is the outcome – third mutation. From 0, we go for free to any of the other absorbing states. 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 mutations.*

[b] *On the other hand, to get out of an untruthful uniform profile, say $m^1 = \phi$ 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 $f(\phi) \succ^{\theta}_{i(\phi, x(\phi,\theta))} x(\phi, \theta)$, in which case this step requires a first mutation, or $x(\phi, \theta) \succeq^{\theta}_{i(\phi, x(\phi,\theta))} f(\phi)$, in which 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 mutation. From there, we go for free to any of the other absorbing states.*

Remark: If one assumes that the preference reversals specified in equation (*) 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 better-response dynamics based on uniform mutations. In this sense, one can clearly interpret Theorem 4 as a very permissive result.

Remark: It appears that, to obtain meaningful implementability results using uniform mutations, one needs to add at least a new rule to the canonical mechanism of Theorem 2, as we have just done. Note how the proof has relied heavily on the use of the preference reversal specified in equation (*).

5 Incomplete Information

This section tackles the extension of our results to incomplete information environments. We shall begin with the (almost) characterization of SCF's that are implementable in asymptotically stable strategies of interim better-response dynamics. Following a plan similar to that in previous sections, we shall consider more permissive results later on.

5.1 Necessary and Sufficient Conditions

We now describe an incomplete information environment. 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 shall assume that the set of states with ex-ante positive probability coincides with Θ.[6]

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 continue to make the same assumptions on ex-post preferences: no externalities, and each agent has strictly increasing utility.

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 $g : \Theta \mapsto Z$. 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, agents are replaced by other agents who share the same preferences and prior beliefs as their predecessors, but are not aware of their experience).

[6] We make this assumption for simplicity in the presentation. With some 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.

Let agent i of type θ_i be the only type that is 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:

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

The rest of the message profile $m_{-i}^t(\theta_{-i})$ is as in the previous period, and so this revision in type θ_i's message may cause a change in the induced outcome in each state where type θ_i is present, as expressed in the preceeding interim expected utility inequality.

We adapt now the definitions of implementability to environments with incomplete information:

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

5.1.1 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 asymptotically stable 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,

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

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

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

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

Theorem 5 *If f is implementable in stochastically stable strategies of any perturbation of an unperturbed interim better-response dynamic process, f is incentive compatible. Furthermore, if at least one of the recurrent classes selected by the perturbation of the interim better-response process is a singleton, f is strictly incentive compatible.*

Proof 5 *Suppose that f is implementable in stochastically stable strategies of any perturbation of better-response dynamics. 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 an absorbing set of better-response dynamics, 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.

Remark: In particular, the conclusions of Theorem 5 extend to the case of no-perturbations at all, i.e., incentive compatibility is also a necessary condition for implementability in asymptotically stable strategies of better-response dynamics. In addition, if the process admits at least one singleton recurrent class, strict incentive compatibility also becomes necessary.

The next definitions prepare the way for the other condition that we shall present in this subsection.

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 $\theta \in \Theta$.

An SCF f is *Bayesian quasimonotonic* if for all deceptions α, for all $i \in N$, and for all $\theta_i \in \Theta_i$, whenever

$$U_i(f \mid \theta_i) > U_i(y_{\theta_i'} \mid \theta_i) \forall \theta_i' \in \Theta_i \quad \text{implies} \quad U_i(f \circ \alpha \mid \theta_i) > U_i(y \circ \alpha \mid \theta_i), \qquad (**)$$

one must have that $f \circ \alpha = f$.

Note how Bayesian quasimonotonicity is to Bayesian monotonicity (e.g., Jackson (1991)) as quasimonotonicity was to Maskin monotonicity.

We move on now to our next necessity result. In it, Bayesian quasimonotonicity shows up as a necessary condition when implementability is sought in asymptotically stable strategies, or equivalently, in stochastically stable strategies for all perturbations of an unperturbed interim better-response process.

THEOREM **6** *If f is implementable in asymptotically stable strategies of an unperturbed interim better-response dynamic process, f is Bayesian quasimonotonic.*

Proof 6 *Suppose that f is implementable in asymptotically stable strategies of better-response dynamics. This means that f is the only outcome of the recurrent classes of the unperturbed dynamics. In particular, this implies that there exists a message profile m such that $g(m) = f$, in one of such recurrent classes.*

Any unilateral deviation from m made by type θ_i either results laso in f, or it changes the outcome. In the latter case, call such an outcome y. If the deviation also involves pretending to be type θ_i', the corresponding outcome imposed can be written as $y_{\theta_i'}$. But, since we are starting from a recurrent class of better-response dynamics, any such deviations that change the outcome fall in the strict lower contour set of the interim preferences for type θ_i at f. That is, $U_i(f \mid \theta_i) > U_i(y_{\theta_i'} \mid \theta_i)$. This statement holds for every θ_i and for every i.

Consider now an arbitrary deception α and suppose, following the hypothesis of Bayesian quasimonotonicity –equation (**)–, that any such $y_{\theta_i'}$, when the deception α is used, is such that $U_i(f \circ \alpha \mid \theta_i) > U_i(y \circ \alpha \mid \theta_i)$ for every i and θ_i. Consider the strategy profile $m \circ \alpha$. Its outcome is $f \circ \alpha$. Then, any unilateral deviation from it on the part of type θ_i either does not change the outcome or yields an outcome $y \circ \alpha$ such that $U_i(f \circ \alpha \mid \theta_i) > U_i(y \circ \alpha \mid \theta_i)$. To see this, note that if type $\alpha_i(\theta_i)$ found a weakly profitable deviation from $m \circ \alpha$, he would be implementing an outcome $y \circ \alpha$ that he likes at least as much as $f \circ \alpha$. But then, the same deviation made by type θ_i from m would lead to an outcome $y_{\theta_i'} \neq f$ that he would like as much as f, contradicting that m was part of a recurrent class of the better-response process.

This implies that $m \circ \alpha$ is also an element of a recurrent class of the better-response process. Thus, for f to be implementable in asymptotically stable strategies, it is required that the outcome of $m \circ \alpha$, i.e., $f \circ \alpha$, be f. That is, f must be Bayesian quasimonotonic.

Remark: Note the subtle difference between the necessity results provided in Theorems 5 and 6. Theorem 5 identifies incentive compatibility as a necessary condition for *any* kind of stochastically stable implementation. Bayesian quasimonotonicity becomes necessary only if one insists on asymptotically stable implementation.

5.1.2 Sufficiency

Our next sufficiency result follows:

THEOREM **7** *Suppose the environments satisfy Assumptions (1) and (2) in each state. Let $n \geq 3$. If an SCF f is ϵ-secure, strictly incentive compatible and Bayesian quasimonotonic, f is implementable in asymptotically stable strategies of interim better-response dynamics.*

Proof 7 *We construct the following canonical mechanism $G = ((M_i)_{i \in N}, g)$, where agent i's message set $M_i = \Theta_i \times A$. Denote $m_i = (m_i^1, m_i^2)$. The outcome function g is defined in the following rules:*

(i.) *If for every agent $i \in N$, $m_i^2 = f$, $g(m) = f(m^1)$.*

(ii.) *If for all $j \neq i$ $m_j^2 = f$ and $m_i^2 = y \neq f$, one can have two cases:*

> **(ii.a.)** *If there exist types $\theta_i, \theta_i' \in \Theta_i$ such that $U_i(y_{\theta_i'} \mid \theta_i) \geq U_i(f \mid \theta_i)$, $g(m) = (f_i(m^1) - \beta, f_{-i}(m^1))$, where $f_i(m^1) \geq f_i(m^1) - \beta \in X_i$.*
>
> **(ii.b.)** *If for all $\theta_i, \theta_i' \in \Theta_i$, $U_i(y_{\theta_i'} \mid \theta_i) < U_i(f \mid \theta_i)$, $g(m) = y(m^1)$.*

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

Following similar steps as in the proof of Theorem 2, one can show that all recurrent classes of the unperturbed dynamics are absorbing states that happen under rule (i). Moreover, their outcomes are either f or $f \circ \alpha$. But in the latter case, since f is Bayesian quasimonotonic, one can show that $f \circ \alpha = f$.

Remark: Theorems 5, 6 and 7 provide almost a characterization of the rules that are implementable in asymptotically stable strategies of better-response dynamics in economic environments with at least three agents.

5.2 More Permissive Results

In an attempt to obtain more permissive results, in this subsection we shall consider stochastically stable implementation using perturbations of interim better-responses based on uniform mutations. We shall make the following 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

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

(7) The bundles in the SCF's x and y used in (***) are componentwise no greater than ϵ.

In words, Assumption (6) 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 how weak this assumption really is.

For each deception α, we shall choose one test-pair x, y and one test-agent i, satisfying the conditions in (***). Denote the set of all such x by D. Finally, with very little loss of generality, choose the bundles in the SCF's 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 (7) says that such reversals happen "near enough the zero bundle."[7] Then, one can make use of the insight in the last remark of the previous section to show our next result:

THEOREM **8** *Suppose that the environments satisfy Assumptions (1), (2), (6) and (7). Let $n \geq 5$. Any ϵ-secure and strictly incentive compatible SCF f is implementable in stochastically stable strategies of perturbed interim better-response dynamics under uniform mutations.*

Proof 8 *The proof follows steps similar to that of Theorem 4, 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 $f_i(\cdot) \geq f_i(\cdot) - \beta \geq (\epsilon, \ldots, \epsilon)$.*

(iii.a.) *(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.) *(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, $g(m) = [y]$.*

[7] In fact, if the environment allowed the use of lotteries and making use of expected utility, one could combine the SCF's 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 SCF's 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 better-response dynamics are singletons under rule (i). Strict incentive compatibility allows one to support truth-telling as one of these absorbing states, but there may well be others, in which agents are using a deception α.

To finish the proof, the details are somewhat involved, but here is an intuitive argument. One can describe the transition paths among the different absorbing states. To get out of the one in which agents are telling the truth in their first part of the announcement, one can go through rule (ii.a), which requires one mutation because any $x \in D$ is near the origin (note that any agent can be used for this mutation, 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 mutation. Finally, someone else mutates 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 mutation 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 mutation. The next step, taken by the test-agent for that deception, is free because of equation (***). From rule (iii.a), someone else changes to rule (iv), and so on. In this path, we have "saved" one mutation. Of course, from the zero profile, we go for free to any of the other absorbing states.

6 Conclusion

This paper has studied the classic implementation problem in evolutionary settings. In particular, necessary and sufficient conditions for implementability in asymptotically stable strategies of better-response processes and stochastically stable strategies of perturbed better-response dynamics have been identified. In this exercise, variants of the well-known monotonicity conditions in implementation theory seem to be the key to capture good dynamic properties of implementation. More permissive results, beyond quasimonotonicity, are possible, but they come at a cost in terms of robustness. In the case of incomplete information, incentive compatibility shows up as the only necessary condition for evolutionary implementation in our sense, with independence of the mutation processes employed. The analysis of the implementation problem under alternative forms of bounded rationality should be the subject of further research.

References

Bergin, J. and B. Lipman (1996), "Evolution with State Dependent Mutations," *Econometrica* **64**, 943-956.

Cabrales, A. (1999), "Adaptive Dynamics and the Implementation Problem with Complete Information," *Journal of Economic Theory* **86**, 159-184.

Cabrales, A. and G. Ponti (2000), "Implementation, Elimination of Weakly Dominated Strategies and Evolutionary Dynamics," *Review of Economic Dynamics* **3**, 247-282.

Eliaz, K. (2002), "Fault-Tolerant Implementation," *Review of Economic Studies* **69**, 589-610.

Foster, D. P. and H. P. Young (1990), "Stochastic Evolutionary Game Dynamics," *Theoretical Population Biology* **38**, 219-232.

Freidlin, M. and A. Wentzell (1984), *Random Perturbations of Dynamical Systems*, Springer-Verlag, Berlin.

Fudenberg, D. and D.K. Levine (1998), *The Theory of Learning in Games*, MIT Press, Cambridge, MA.

Jackson, M. O. (1991), "Bayesian Implementation," *Econometrica* **59**, 461-477.

Jackson, M. O. (2001), "A Crash Course in Implementation Theory," *Social Choice and Welfare* **18**, 655-708.

Kandori, M., G. Mailath, and R. Rob (1993), "Learning, Mutations and Long Run Equilibria in Games," *Econometrica* **61**, 29-56.

Maskin, E. S. (1999), "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies* **66**, 23-38.

Maskin, E. S. and T. Sjöström (2002), "Implementation Theory," in *Handbook of Social Choice and Welfare (vol. I)*, ed. by K. J. Arrow, A. Sen and K. Suzumura, New York, Elsevier Science B.V.

Palfrey, T. R. (2002), "Implementation Theory," in *Handbook of Game Theory with Economic Applications (vol. III)*, ed. by R. J. Aumann and S. Hart, New York, Elsevier Science.

Palfrey, T. R. and S. Srivastava (1989), "Implementation with Incomplete Information in Exchange Economies," *Econometrica* **57**, 115-134.

Postlewaite, A. and D. Schmeidler (1986), "Implementation in Differential Information Economies," *Journal of Economic Theory* **39**, 14-33.

Samuelson, L. (1997), *Evolutionary games and equilibrium selection*, MIT Press, Cambridge, MA.

Sandholm, W. (2002), "Evolutionary Implementation and Congestion Pricing," *Review of Economic Studies* **69**, 667-689.

Serrano, R. (2004), "The Theory of Implementation of Social Choice Rules", *SIAM Review* **46**, 377-414.

Serrano, R. and R. Vohra (2005), "A Characterization of Virtual Bayesian Implementation," *Games and Economic Behavior* **50**, 312-331.

Sjóstróm, T. (1994), "Implementation in Undominated Nash Equilibria without Using Integer Games," *Games and Economic Behavior* **6**, 502-511.

Vega-Redondo, F. (1996), *Evolution, games, and economic behavior*, Oxford University Press, New York.

Weibull, J. (1995), *Evolutionary Game Theory*, MIT Press, Cambridge, MA.

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.