Provided for non-commercial research and education use. Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Games and Economic Behavior 73 (2011) 360-374



Contents lists available at ScienceDirect

Games and Economic Behavior



www.elsevier.com/locate/geb

Implementation in adaptive better-response dynamics: Towards a general theory of bounded rationality in mechanisms

Antonio Cabrales^{a,b,*}, Roberto Serrano^{c,d}

^a Departamento de Economía, Universidad Carlos III de Madrid, Madrid 126, 28903 Getafe, Spain ^b CEPR, United Kingdom

^c Department of Economics, Brown University, Providence, RI 02912, USA

^d IMDEA-Social Sciences Institute, Madrid, Spain

ARTICLE INFO

Article history: Received 18 January 2010 Available online 21 March 2011

JEL classification: C72 D70 D78

Keywords: Robust implementation Bounded rationality Evolutionary dynamics Mechanisms

ABSTRACT

We study the classic implementation problem under the behavioral assumption that agents myopically adjust their actions in the direction of better-responses or bestresponses. First, we show that a necessary condition for recurrent implementation in better-response dynamics (BRD) is a small variation of Maskin monotonicity, which we call quasimonotonicity. We also provide a mechanism for implementation in BRD if the rule is quasimonotonic and excludes worst alternatives - no worst alternative (NWA). Quasimonotonicity and NWA are both necessary and sufficient for absorbing implementation in BRD. Moreover, they characterize implementation in strict Nash equilibria. Under incomplete information, incentive compatibility is necessary for any kind of stable implementation in our sense, while Bayesian quasimonotonicity is necessary for recurrent implementation in interim BRD. Both conditions are also essentially sufficient for recurrent implementation, together with a Bayesian NWA. A characterization of implementation in strict Bayesian equilibria is also provided. Partial implementation results are also obtained.

© 2011 Elsevier Inc. All rights reserved.

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 decades have seen 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 contributions² 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, leading to a wider span of implementable social choice correspondences (SCCs) or social choice functions (SCFs), may lead to dynamic instability or convergence to the wrong

Corresponding author at: Departamento de Economía, Universidad Carlos III de Madrid, Madrid 126, 28903 Getafe, Spain. E-mail addresses: antonio.cabrales@uc3m.es (A. Cabrales), roberto_serrano@brown.edu (R. Serrano). URLs: http://www.eco.uc3m.es/acabrales (A. Cabrales), http://www.econ.brown.edu/faculty/serrano (R. Serrano).

¹ See Jackson (2001), Maskin and Sjöström (2002), Palfrey (2002), Serrano (2004) or Corchón (2009) for recent surveys.

 $^{^2}$ See e.g. Cabrales (1999) and the literature review at the end of this introduction.

equilibrium when the players are boundedly rational. This should perhaps not be surprising since those mechanisms were not designed with robustness to bounded rationality in mind, and yet, somewhat surprisingly, the canonical mechanism for Nash implementation turned out to be robust to some kinds of bounded rationality (Cabrales, 1999).

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 on SCCs for their evolutionary implementation? Are they close to being sufficient as well? That is, can one characterize the set of rules that a planner could hope to decentralize in a society populated by boundedly rational agents? The current paper answers these questions when one models bounded rationality by myopic behavior that tends to move in the direction of better replies (or best replies) to a bounded sample of past history, which conforms with a wide class of evolutionary settings.³

Thus, 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. We remark that the wide class of dynamics considered here does not require that all better-responses be played with positive probability, although all best-responses must be. Indeed, as we explain in the second remark after Theorem 2, our central results (in Section 3) and their proofs go through unchanged if we switch from better-response dynamics to (full support) best-response dynamics.⁴ 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 all the outcomes of an SCC are the only limits of the better-response dynamics (BRD) of a mechanism for any allowed environment, we shall say that the SCC is *implementable in recurrent strategies* of BRD. If the only limits of BRD lie in the SCC, we speak of *partial implementability in recurrent strategies*.

We come to describe our first main finding. A necessary condition for recurrent implementation is a small variation of (Maskin) monotonicity (Maskin, 1999), which we call *quasimonotonicity*. Quasimonotonicity prescribes that the social outcome remain in the SCC if the *strictly* lower contour sets of preferences at that social outcome are nested for every agent across two environments. In particular, it is neither logically stronger nor weaker than monotonicity, although both coincide in many settings (we discuss the relationship in Section 2). Furthermore, quasimonotonicity is also sufficient for recurrent implementation in BRD, if there are at least three agents in the environment and the SCC satisfies the "no worst alternative" (NWA) property. This property requires that, for every environment, there is always a strictly worse outcome for every agent than each outcome prescribed by the SCC. NWA becomes also necessary if the recurrent classes of the better-response dynamics are singletons (i.e., absorbing states) in the implementing mechanism.

The typical mechanism that we construct has good dynamic properties. It implements the socially desirable outcome according to the agents' reports, if there is total agreement among them. If only one agent's report disagrees with the rest, other outcomes will be implemented. Those outcomes are meant to elicit the "right" behavior from agents. Finally, if more than one agent disagrees, a "modulo" game is played. We note that the "modulo" game is not an essential part of the mechanism. It can be replaced, for example, with a unanimously bad outcome for all players. What really matters (and the "modulo" game delivers) is that outside from a situation of total (or almost total) agreement, it does not hurt to tell the "truth". This then implies that behavior can easily "drift" into unanimous truth-telling.⁵ Once agents agree, the rules for almost unanimous agreement, which make use of both quasimonotonicity and the punishments that are possible by NWA, avoid any "drifting out" of generalized agreement.

Our agents are boundedly rational – myopic better responders – yet they are able to act in an apparently complex mechanism with a modulo game. While one could see our mechanism as complex, it is finite. The argument in the previous paragraph already hints at the fact that the nature of better responses in the "complex" mechanism is not unduly complicated. It requires to understand what is your most preferred outcome, what is the true profile and not much more. More importantly, we believe that the question whether real agents are able to understand the game should be mostly empirical. In this respect, the results of Cabrales et al. (2003) are encouraging. In a laboratory experiment, the likelihood of the socially desired outcome being implemented was 0.80 in a treatment using a mechanism similar to the one we use for our sufficiency condition. Nevertheless, as already acknowledged in footnote 3, more research in this important question is needed.

Our results on recurrent implementation in BRD 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 BRD via mutations, an SCF that is implementable in recurrent strategies would also be *implementable in stochastically stable strategies* of any perturbation of BRD; the reader is referred to our companion paper for further details on this topic (Cabrales and Serrano, 2010). It follows that quasimonotonicity is identified as the key condition to essentially characterize very robust implementation with respect to myopic BRD processes.⁶

³ In doing so, we acknowledge that some of the canonical mechanisms we construct might be "too complex" for boundedly rational agents, which forces us to stress the word "towards" from our title even more strongly. References for evolutionary game theory in general are Weibull (1995), Vega-Redondo (1996), Samuelson (1997), Fudenberg and Levine (1998) and Young (1998).

⁴ Fictitious play would require some more discussion, but an assumption of limited memory would lead to the same results.

⁵ Indeed, if agents were to have some arbitrarily small cost of misrepresenting their private information, as in Kartik (2009), this drift into unanimous truth-telling could be quite fast.

⁶ It is important to stress how Maskin monotonicity has emerged in other works when robustness is required in other senses: Chung and Ely (2003) if undominated Nash implementation is to occur with near-complete information, Aghion et al. (2010) for subgame-perfect implementation also with near-complete information, and Bergemann et al. (2010) for rationalizable implementation.

Given that one can model bounded rationality in many different ways, we also explore (in Section 3.3) the robustness to other such approaches of the conditions identified for recurrent implementation under BRD. A common feature of many alternative approaches is the stability of strict Nash equilibria. Therefore we pose the question of implementation in strict Nash equilibria and provide its characterization. We find that we can dispense altogether with the standard no-veto-power sufficient condition. In general environments with at least three agents both quasimonotonicity and NWA are necessary and sufficient for implementation in strict equilibria. This provides a further strong endorsement of quasimonotonicity, and raises the stature of the NWA condition. Indeed, NWA also becomes necessary when we require implementation in absorbing strategies of BRD (see Section 3.3 for details).⁷ We view the strict Nash implementation result as a confirmation of the more general message sent by our findings for evolutionary implementation.⁸

Next, we also provide an almost characterization of partial implementability in recurrent strategies of BRD. The key condition here is a version of quasimonotonicity applied to the entire SCC. When the range of the correspondence is rich, this condition is typically trivial and thus, we learn that convergence of BRD processes to a set is much easier to obtain.

Our main insights already described are confirmed in environments with incomplete information, and some others are obtained therein (for simplicity, we present most of those results for SCFs and economic environments). 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.⁹ If one wishes to implement in recurrent strategies, 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 recurrent strategies of BRD processes when there are at least three agents.¹⁰

Finally, we provide a characterization of the SCFs that are implementable in strict Bayesian equilibria in general environments, a contribution in its own right, and also a vehicle to show the general implications of our basic results for recurrent implementation of interim BRD. That is, parallel to Section 3.3, we also study in Section 5.2 the robustness of the conditions identified for recurrent Bayesian implementation. Dropping the assumption of economic environments, we provide a full characterization of strict Bayesian implementation with at least three agents. We remark that no such characterization is available for Bayesian implementation.¹¹ The conditions for strict Bayesian implementation are strict incentive compatibility, Bayesian quasimonotonicity, and a version of the NWA condition for Bayesian environments.

1.1. Related literature

The study of implementation under bounded rationality has a long intellectual history. Muench and Walker (1984) and de Trenqualye (1988) study the stability of the Groves and Ledyard (1977) mechanism.¹² Walker (1984) describes a mechanism for which almost Walrasian allocations in large economies are stable. Vega-Redondo (1989) proposes a globally convergent mechanism (under best-response dynamics) to the Lindahl equilibrium in an economy with one private good, one public good and linear production. Along the same lines, de Trenqualye (1989) proposes a locally stable mechanism to implement Lindahl equilibria in an economy with multiple private goods, one public good, linear production and quasi-linear preferences. With respect to these early contributions, our paper does not focus on the stability of particular mechanisms, SCFs or environments, and we deal with global, rather than local convergence.

More recently, Cabrales (1999) studies sufficient conditions for global convergence of (a variation of) the canonical mechanism for Nash implementation.¹³ The sufficient conditions used there are very similar to ours. However, Cabrales (1999) does not study necessary conditions for convergence, but rather the bad dynamic properties of other (already existing) mechanisms. And unlike Cabrales and Serrano (2010), that paper does not provide either additional conditions on dynamics or preferences to deliver ergodic dynamic properties.¹⁴

Sandholm (2005) provides a sufficient concavity condition under which simple price schemes would be globally stable when implementing efficient SCFs in economic environments.¹⁵ Mathevet (2007) studies sufficient conditions for implemen-

⁷ Cabrales et al. (2003) show experimentally that in the absence of NWA, the performance of Maskin's canonical mechanism for Nash implementation is substantially worse.

⁸ The characterization of strict Nash implementation has interest in its own right. In particular, it is noteworthy how one can dispense with the "no-veto" condition in non-random settings, even outside of economic environments.

⁹ Our comment above regarding better-responses versus best-responses applies here as well.

¹⁰ Under a weak diversity of preferences 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 mutations and at least five agents; see again Cabrales and Serrano (2010).

¹¹ Theorem 2 in Jackson (1991) for general environments is a sufficiency result.

¹² Chen and Plott (1996) and Chen and Tang (1998) confirm experimentally these findings.

¹³ Cabrales et al. (2003) confirm experimentally these findings.

¹⁴ Cabrales and Ponti (2000) study the convergence and stability properties of Sjöström's (1994) mechanism assuming that the dynamics are monotonic in the sense of Samuelson and Zhang (1992). Cabrales (1999) also discusses some drawbacks, from a dynamic point of view, of the mechanisms in Abreu and Matsushima (1992, 1994).

¹⁵ Sandholm (2002) provides a concrete mechanism achieving efficiency and stability in a road congestion pricing problem.

tation in supermodular games, motivated by their good learning properties. In contrast, we do not take advantage of the good convergence properties of dynamic processes when posing our general questions and constructing our mechanisms, and therefore, the SCFs implemented in those papers must still meet our necessary conditions if the processes follow the BRD direction.

Another important line of research has studied implementation that is statically robust to bounded rationality. Eliaz (2002) shows that (a strengthening of) monotonicity and no-veto power are sufficient to implement an SCF even when agents are error-prone, provided a "majority" of players are not. Tumennasan (2008) shows that (once again) quasimonotonicity is necessary for implementation when "all" agents are error-prone in the style of Logit Quantal Response Equilibrium (LQRE), as defined in McKelvey and Palfrey (1995).¹⁶ Since the conditions used in the latter paper are closer to ours, we discuss it in more detail later on.

The papers we have discussed so far typically focus on environments with complete information, that is, the agents' know each others preferences. Notable exceptions are the previously mentioned paper of Mathevet (2007) and the impossibility result of Jordan (1986). We extend the literature by also studying incomplete information environments. In this respect, there is also a small related literature on the stability of equilibria for different auction and double auction formats (including Hon-Snir et al., 1998; Dawid, 1999; Saran and Serrano, 2010).

1.2. Plan of the paper

Section 2 describes the model and the dynamics we use. Section 3 provides necessary and sufficient conditions for recurrent implementation under complete information. Section 4 presents more permissive results – i.e., concerning nonquasimonotonic rules – if one relies on partial implementability of SCCs. Section 5 collects the extensions of our results to incomplete information environments. Section 6 concludes.

2. Preliminaries

Let $N = \{1, ..., n\}$ be a set of agents and let *Z* be a finite set of alternatives. Let θ_i denote agent *i*'s preference ordering over the set *Z*. Let $\theta = (\theta_i)_{i \in N}$ be a preference profile, and Θ be the (finite) set of allowable preference profiles.¹⁷ We assume that $|\Theta| \ge 3$, without loss of generality.¹⁸ The symbol \ge_i^{θ} stands for *i*'s weak preference when the preference profile is θ , while $>_i^{\theta}$ denotes *i*'s strict preference at θ and \sim_i^{θ} denotes indifference. We refer to the set of agents, the set of alternatives and agents' preferences over alternatives as an environment. To the extent that the sets *N* and *Z* will remain fixed throughout, two environments differ only in that the preferences of some agents change across them. (In Section 5 we shall consider also incomplete information environments, in which agents' information may also change.)

A social choice correspondence (SCC) specifies a set of alternatives for each preference profile θ . Thus, an SCC $F : \Theta \mapsto 2^Z$. A social choice function (SCF) is an SCC that always assigns a singleton 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 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'_i such that $g(m'_i, m_{-j}) \succeq_i^{\theta} g(m)$.

For our first set of results, in the next section, we begin by considering a specific class of SCCs. Most importantly, we shall use the following condition, which turns out to be central to the theory we develop here:

An SCC *F* satisfies *quasimonotonicity* whenever, for every $a \in F(\theta)$, if it is true that for every $i \in N$ and for all $z \in A$ with $a \succ_i^{\theta} z$ we have that $a \succ_i^{\phi} z$, this implies that $a \in F(\phi)$ for all $\theta, \phi \in \Theta$. Note how quasimonotonicity resembles closely the condition of monotonicity uncovered in Maskin (1999). Indeed, the

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 the lower contour sets of preferences be nested across two environments, quasimonotonicity relies on the inclusion of the strictly lower contour sets. It follows from the definitions of both conditions that quasimonotonicity is neither stronger nor weaker than Maskin monotonicity.¹⁹ The next two examples present two important correspondences that make this point:

Example 1. Consider environments in which for no pair of alternatives $z, z' \in Z$ it is the case that $z \sim_i^{\theta} z'$ for all $i \in N$. Define the strong Pareto correspondence $F^{SP} : \Theta \mapsto 2^Z$ as follows:

$$\mathsf{F}^{\mathsf{SP}}(\theta) \equiv \{ z \in Z \colon \nexists z' \in Z, \forall i \in \mathbb{N}, z' \succeq_i^{\theta} z \text{ and } \exists j \in \mathbb{N}, z' \succ_i^{\theta} z \}.$$

¹⁶ The notion of LQRE is itself closely associated with the stochastic choice models of Luce (1959) and McFadden (1973).

 $^{^{17}}$ The finiteness of Θ guarantees that we can study the dynamics for our mechanisms with finite Markov chain tools.

¹⁸ If $|\Theta| = 2$, one can always add an extra environment artificially by replicating one of the two preference profiles and use a third name for the new one. ¹⁹ This contrasts with other conditions previously found in the literature (e.g., the weaker "almost monotonicity" in Sanver, 2006), used to expand the scope of Nash implementability by means of awards.

As is well known, F^{SP} does not satisfy Maskin monotonicity: for instance, let $N = \{1, 2\}$ and $Z = \{z, z'\}$, θ be such that $z \succ_1^{\theta} z'$ and $z' \succ_2^{\theta} z$, so that $F^{SP}(\theta) = \{z, z'\}$. And let ϕ differ from θ only in that $z \sim_1^{\phi} z'$, implying that $z \notin F^{SP}(\phi)$. However, it is easy to see that F^{SP} will satisfy quasimonotonicity in any such environment. Indeed, suppose the hypoth-

However, it is easy to see that F^{SP} will satisfy quasimonotonicity in any such environment. Indeed, suppose the hypothesis of quasimonotonicity is satisfied (the inclusion of the strictly lower contour sets between θ and ϕ), and that $z \in F^{SP}(\theta)$, but $z \notin F^{SP(\phi)}$. The latter implies that there exists z' such that $z' \succeq_i^{\phi} z$ for all $i \in N$ and $z' \succ_j^{\phi} z$ for some $j \in N$. Recalling that it is not the case that $z \sim_i^{\theta} z'$ for all $i \in N$, the inclusion of the strictly lower contour sets would imply that $z' \succeq_i^{\theta} z$ for all $i \in N$ and that there exists some $j \in N$. Recalling that $i \in N$ and that there exists some $j \in N$ for whom $z' \succ_j^{\theta} z$, contradicting that $z \in F^{SP}(\theta)$.²⁰

Example 2. Consider any environment. Define the weak Pareto correspondence $F^{WP}: \Theta \mapsto 2^Z$ as follows:

$$F^{WP}(\theta) \equiv \left\{ z \in Z \colon \nexists z' \in Z, \ \forall i \in N, \ z' \succ_i^{\theta} z \right\}.$$

As is well known, F^{WP} always satisfies Maskin monotonicity. However, it may violate quasimonotonicity: for instance, let $N = \{1, 2\}$ and $Z = \{z, z'\}$, θ be such that $z \sim_1^{\theta} z'$ and $z' \succ_2^{\theta} z$, so that $F^{WP}(\theta) = \{z, z'\}$. And let ϕ differ from θ only in that $z' \succ_1^{\phi} z$, implying that $z \notin F^{WP}(\phi)$.

The above examples underscore the differences between Maskin monotonicity and quasimonotonicity. Nonetheless, there are many environments in which the two conditions coincide. For example, consider a problem of assignment of indivisible goods in which all preferences are strict (as in Shapley and Scarf, 1974; Roth and Postlewaite, 1977), or environments in which agents' preferences are continuous (in this case, the inclusion of lower contour sets and strictly lower contour sets is equivalent).

In addition to quasimonotonicity, we shall use the following condition:

An SCC *F* satisfies *no-worst-alternative* (*NWA*) whenever for every agent *i*, every preference profile θ and every $a \in F(\theta)$, there exists an outcome $z_i^{a,\theta}$ such that $a \succ_i^{\theta} z_i^{a,\theta}$.

In particular, NWA allows the possibility of punishing an agent at every outcome of the SCC.

In Section 5, in which we concentrate on economic environments and where the social choice rule allocates a bundle to each agent, we shall use a version of NWA for SCFs. An SCF f is said to be ε -secure if there exists $\varepsilon > 0$ such that for each θ , and for each $i \in N$, $f_i(\theta) \ge (\varepsilon, \dots, \varepsilon) \gg 0$.

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 it on normative grounds.

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 $\prod_{i \in N} M_i$ of message profiles as the state space (which, to avoid further measure theoretic technicalities, we assume to be countable). We shall begin by specifying a 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.²¹ 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 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 $g(m'_i, m_{-i}(t)) \geq_i^{\theta_i} g(m(t))$.²² Thus, the planner, who has a long run perspective on the social choice problem, wishes to design an institution or

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 outcomes as specified by the SCC. This logic suggests the following notion of implementability.

An SCC *F* is *implementable in recurrent strategies* (of BRD) if there exists a mechanism *G* such that, for every $\theta \in \Theta$ and every $a \in F(\theta)$, there is a recurrent class of the BRD process (applied to the induced game under preferences θ) such that the class has a unique outcome and the outcome is equal to *a*. Furthermore, the outcome of any recurrent class of that BRD process must be some $a \in F(\theta)$.

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, and this would lead to the concept of implementability in stochastically stable strategies (the interested reader should see Cabrales and Serrano, 2010).

²⁰ In environments similar to these where individual property rights are defined, other examples of correspondences that violate Maskin monotonicity, but that satisfy quasimonotonicity, include the strict individual rationality correspondence and its intersection with the strong Pareto correspondence.

²¹ For complete formal definitions of Markov chains, and related terms (recurrent classes, absorbing states, etc., see e.g. Karlin and Taylor, 1975, Chapter 2).

²² As we discuss in the second remark after Theorem 2, all the results can be extended to the case in which agents switch to play a subset of betterresponses (in particular, best-responses) to the mixed-strategy of the others given by the empirical distribution of play over some bounded sample of past history.

3. Necessary and sufficient conditions for recurrent implementation: Complete information

In this section we identify necessary and sufficient conditions for implementability in recurrent strategies of BRD. The section also discusses more general implications of these results, and in particular, the connections with implementation in strict Nash equilibria.

3.1. Necessity

We seek implementation in recurrent strategies of the BRD. We wish to show now that quasimonotonicity of F is a necessary condition for its implementability in recurrent strategies.

Theorem 1. If F is implementable in recurrent strategies of a BRD process, F is quasimonotonic.

Proof. Let the true preference profile be θ . Because F is implementable in recurrent strategies of BRD, the only outcomes that correspond to strategy profiles in recurrent classes of the dynamics are $a \in F(\theta)$.

Let $a \in F(\theta)$, and now consider a preference profile ϕ such that for all $i \ a \succ_i^{\theta} z$ implies that $a \succ_i^{\phi} z$. Since a is the only outcome compatible with a recurrent class of the dynamics when preferences are θ , this means that agent i's unilateral deviations from recurrent strategy profiles in that class must yield either a again, or outcomes z such that $a >_{\theta}^{\theta} z$. But then, the same recurrent class is also a recurrent class under ϕ , i.e., a is also supported by recurrent profiles of BRD when preferences are ϕ . Since F is implementable in recurrent strategies of BRD, this implies that $a \in F(\phi)$. 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 SCCs that are implementable in recurrent strategies of BRD.

Theorem 2. Let $n \ge 3$. If an SCC F satisfies quasimonotonicity and NWA, it is implementable in recurrent strategies of BRD.

Proof. Consider the mechanism $G = ((M_i)_{i \in N}, g)$, where agent *i*'s message set is $M_i = \Theta \times Z \times Z \times N$. Denote agent *i*'s message $m_i = (\theta_i, a_i, z_i, n_i)$, and the agents' message profile by m. Let the outcome function g be defined by the following rules:

- (i) If for all $i \in N$, $m_i = (\theta, a, z_i, n_i)$ and $a \in F(\theta)$, g(m) = a.
- (ii) If for all $j \neq i$, $m_j = (\theta, a, z_j, n_j)$ with $a \in F(\theta)$, and $m_i = (\phi, a', z, n_i)$, $(\phi, a') \neq (\theta, a)$, one can have two cases: (ii.a) If $z \succeq_i^{\theta} a$, $g(m) = z_i^{a,\theta}$ recall the definition of the NWA property. (ii.b) If $a \succ_i^{\theta} z$, g(m) = z.
- (iii) In all other cases, $g(m) = z_t$, where $t = (\sum_{i \in N} n_i) \pmod{n+1}$.

Note how, because F satisfies NWA, rule (ii.a) is well defined. Now, we begin by arguing in the next two steps that all recurrent classes of the BRD process yield outcomes 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. There are two possible cases to consider.

Case A. If fewer than n-2 agents make a common announcement of the preference profile and of the alternative in the SCC, no single agent can change his strategy and move the game to a profile outside of rule (iii). Then, it is a best-response for any agent i to choose a strategy (θ, a, a^i, n_i) with $a \in F(\theta)$, n_i suitably chosen to win the modulo game and a^i one of the best possible outcomes for i under θ . If all agents get a simultaneous revision opportunity, the game goes to rule (i) and yields $a \in F(\theta)$, from which one can never exit under BRD.

Case B. If, on the other hand, exactly n-2 agents make a common announcement of the preference profile and of the alternative in the SCC, say (ϕ, a) with $a \in F(\phi)$, there is an agent who can change his strategy and move the game to a profile in (ii). Then, let any agent announcing (ϕ, a) get a revision opportunity and change to announcing (ϕ', a', a^i, n_i) with $a' \in F(\phi')$ and ϕ' different from any preference profile announced in *m* if n > 3 (recall that $|\Theta| \ge 3$; on the other hand, if n = 3, there is no need for this change so that $\phi' = \phi$, n_i suitably chosen to win the modulo game and a^i one of the best possible outcomes for *i* under θ . This is possible under BRD, and now we are in *Case A*.

Step 2: No message profile under rule (ii) is part of a recurrent class of BRD. We argue by contradiction. Recall that the true preference profile is θ , and let the message profile under rule (ii) – (ii.a) or (ii.b) – in question be the following: all agents $j \neq i$ announce $m_j = (\phi, a, z_j, n_j)$, whereas agent *i*'s message is (ϕ', a', z', n_i) . Then, the outcome is either $z_i^{a,\phi}$ under rule (ii.a) or z' under rule (ii.b). But in either case, from here, each of the other agents $j \neq i$ can switch to $(\phi^j, a^j, a^j, n_j') \neq i$ (ϕ, a, z_j, n_j) , with n'_i suitably chosen to win the modulo game and a^j one of the best possible outcomes for j under θ . Thus, we find ourselves under rule (iii), which is a contradiction.

Therefore, from steps 1 and 2, all recurrent classes contain only profiles under rule (i). Moreover, each recurrent class must yield a unique outcome, say $a \in F(\phi)$. To see this, notice that a class which includes strategy profiles with different outcomes, must include profiles under rule (ii) or (iii), which cannot happen by steps 1 and 2.

In addition, each strategy profile in a recurrent class must be a Nash equilibrium of the game induced by the mechanism when the true preferences are θ . A non-equilibrium profile would produce with positive probability a deviation from rule (i), so the recurrent class would include profiles under rules (ii) or (iii), which cannot happen by steps 1 and 2.

There is one family of recurrent classes of Nash equilibria that exists for all games induced by the mechanism. It is composed of all the recurrent classes containing the truthful profiles (θ, a, z_i, n_i) for every $a \in F(\theta)$ reported by every agent $i \in N$. A unilateral deviation from any such profile to a profile outside the recurrent class would produce a profile either under rule (ii.a) or under rule (ii.b). In either case, no such switch can happen under BRD (recall that F satisfies NWA).

But some games induced by the mechanism may have other (non-truthful) recurrent classes of Nash equilibria under rule (i). Let us call one such class, the class (ϕ, a) and let $(\phi, a, z_i, n_i)_{i \in N}$, $a \in F(\phi)$, be one Nash equilibrium in class (ϕ, a) . Since this profile is a Nash equilibrium, it must be true that for all $i \in N$, $a \succ_i^{\phi} z$ implies that $a \succeq_i^{\theta} z$. If that were not the case (i.e. if $z \succ_i^{\theta} a$), given that *i* can obtain *z* by deviating to rule (ii.b) because $a \succ_i^{\phi} z$, he would choose to deviate and $(\phi, a, z_i, n_i)_{i \in N}$ would not be a Nash equilibrium.

But in fact, we know even more. Any $z \neq a$ that any agent i can obtain by changing his message would yield a message profile outside the class (ϕ , a). So if there is no deviation to such messages that yields $z \neq a$ in a BRD it must be because a deviation would leave the deviant strictly worse off. That is, z cannot be indifferent to a for i under θ . Hence, we must have that for all $i \in N$, $a \succ_i^{\phi} z$ implies that $a \succ_i^{\theta} z$. But, because F is quasimonotonic, this implies that $a \in F(\theta)$ for any arbitrary profile $(\phi, a, z_i, n_i)_{i \in N}$ in any recurrent

class of the BRD. Therefore, F is implementable in recurrent strategies of BRD. \Box

Remark. The reader will have noticed the similarities between the mechanism used in this proof and the one in Maskin (1999). There are some noticeable differences. Our rule (ii.a) is important because all better-responses (or best-responses) could be used with positive probability by agents who get a revision opportunity. If an agent chooses to announce some preference profile/alternative different from the consensus and he is not punished, he will thus do so with positive probability. But, from there, it is a best-response for someone else to move to the modulo game, and the mechanism would unravel. This justifies the use of the punishments in rule (ii.a). In the parallel rule of Maskin's (1999) mechanism, deviations from rule (i) are allowed and are not a problem because they do not upset a Nash equilibrium. Our rule (ii.b), on the other hand, has to be modified to complement the described change in rule (ii.a). Rule (iii) is a modulo game, instead of an integer game. Finite games are usually viewed as more natural. Finally, our strategy set allows two reported alternatives, as opposed to only one in Maskin (1999). This is done to prevent recurrent classes of the dynamics within rule (iii), creating a path back to rule (i), as explained in step 1 of the proof.

Remark. Although we define the stochastic process to be BRD, the proof of Theorem 2 uses only best-responses. This immediately shows that our results hold under best-reply dynamics (provided all best-responses are chosen with positive probability). Given that players do not switch strategies with positive probability (and hence there can be periods of arbitrary finite length without any changes), they would also apply to dynamics that best reply to the empirical distribution of play over some bounded sample of past history (a finite memory version of fictitious play²³).

3.3. Connections with other approaches

To the extent that one can model bounded rationality in many different ways, the reader may be wondering whether the conditions identified for recurrent implementation of BRD processes might have some bearing with other such approaches. This subsection answers this question in the affirmative. In particular, a common thread found by several learning, evolution and bounded rationality models in game theory is their support to certain classes of strict Nash equilibria as being "robust" to many of these exercises. This rationale perhaps justifies the question of the identification of the conditions behind strict Nash implementation, a question that the literature so far has overlooked. We turn to it next.

An SCC *F* is *implementable in strict Nash equilibrium* if there exists a mechanism *G* such that, for every $\theta \in \Theta$, the set of all strict Nash equilibrium outcomes of its induced game when the preference profile is θ is $F(\theta)$.

A mechanism is non-imposing if for each agent $i \in N$, there exist $m_i, m'_i \in M_i$, such that for some $m_{-i} \in M_{-i}$ $g(m_i, m_{-i}) \neq i$ $g(m'_i, m_{-i})$. (In particular, non-imposing mechanisms allow all agents to meaningfully participate.)

²³ Where we would again require all best-responses to be chosen with positive probability.

We first state a result that identifies necessary conditions for strict Nash implementability:

Theorem 3. If *F* is implementable in strict Nash equilibrium, *F* satisfies quasimonotonicity. Furthermore, if *F* is implementable in strict Nash equilibrium by means of a non-imposing mechanism, *F* satisfies NWA.

Proof. The simple proof is left to the reader. \Box

We move right away to the sufficiency result:

Theorem 4. Let $n \ge 3$. If an SCC F satisfies quasimonotonicity and NWA, it is implementable in strict Nash equilibrium.

Proof. We sketch the proof by specifying a canonical mechanism that will serve to prove the statement. The mechanism is a small variant of the canonical mechanism for Nash implementation, proposed by Repullo (1987) and used in Maskin (1999).

Consider the mechanism $G = ((M_i)_{i \in N}, g)$, where agent i's message set is $M_i = \Theta \times Z \times N$. Denote agent i's message $m_i = (\theta_i, z_i, n_i)$, and the agents' message profile by m. Let the outcome function g be defined by the following rules:

- (i) If for all $i \in N$, $m_i = (\theta, a, 1)$ with $a \in F(\theta)$, g(m) = a.
- (ii) If for all *j* ≠ *i*, *m_j* = (θ, *a*, 1) with *a* ∈ *F*(θ), and *m_i* = (φ, *z*, *n_i*) ≠ (θ, *a*, 1), one can have two cases:
 (ii.a) If *z* ≿^θ_{*i*} *a*, *g*(*m*) = *z*^{*a*,θ}_{*i*} recall the definition of NWA.
 (ii.b) If *a* ≻^θ_{*i*} *z*, *g*(*m*) = *z*.
 (iii) In all other cases, *g*(*m*) = *z*_{*t*}, where *t* = (∑_{*i*∈N} *n_i*) (mod *n* + 1).

Again, NWA implies that rule (ii.a) is well defined. Then, one can first establish that unanimous truthful announcements under rule (i) are strict Nash equilibria that implement each of the desired outcomes. It is easy to see that one cannot have strict equilibria either under rule (iii) or rule (ii). Thus, all strict Nash equilibria must happen under rule (i), and quasimonotonicity of F and the construction of rule (ii) imply that any such strict equilibrium must yield an outcome in $F(\theta)$.

Remark. Note how the solution concept employed – strict Nash equilibrium – allows one to dispense completely with the no-veto condition, one of the sufficient conditions for Nash implementation outside of economic environments (see Maskin, 1999).²⁴

Therefore, we have established that for *any* environment (economic or otherwise), quasimonotonicity and NWA provide a complete characterization of strict Nash implementability if there are at least three agents. The comparison with the conditions for recurrent implementability in BRD is apparent: only NWA ceases to be necessary then. Next, we investigate how one can nail down the connection even further.

An SCC *F* is *implementable in absorbing strategies* (of BRD) if it is implementable in recurrent strategies and all recurrent classes of the BRD process applied to its induced game are singletons.

Then, one can show the following connecting proposition. Its simple proof is also left to the reader:

Proposition 1. If an SCC F is implementable in absorbing strategies of BRD, F is implementable in strict Nash equilibrium. Conversely, if an SCC F is implementable in strict Nash equilibrium and there exists an implementing mechanism whose recurrent classes of BRD are all singletons, F is implementable in absorbing strategies of BRD.

Note how the second hypothesis of the converse clause in the proposition rules out non-singleton recurrent classes of BRD, such as those involved in the modulo game of the mechanism we used to prove Theorem 4. Indeed, in some environments, each agent's best alternative may still be better than the outcomes of the SCC for all agents, and in this case, we would have a recurrent class of BRD in which we cycle among each agent's top-ranked outcomes.

Following a different approach to bounded rationality, Tumennasan (2008) considers implementation in limiting logit quantal response equilibria (LLQRE), as defined in McKelvey and Palfrey (1995), where the limit is taken as noise in random payoff maximization is removed, i.e., as behavior approaches full rationality. Noting some difficulties in characterizing the structure of non-strict LLQRE, Tumennasan (2008) confines his attention to strict LLQRE and calls his implementability notion *restricted LLQRE implementation*. Interestingly, he shows that both quasimonotonicity and a stronger version of NWA are necessary for restricted LLQRE implementability, and in economic environments with at least three agents, he shows that these two conditions, together with no-veto, are also sufficient. In other words, to the extent that strict Nash equilibria

²⁴ See also Bochet (2007) and Benoit and Ok (2008), who escape no-veto within Nash implementation by using stochastic mechanisms.

Author's personal copy

A. Cabrales, R. Serrano / Games and Economic Behavior 73 (2011) 360-374

are the common thread to these robustness results, the conditions identified in the current paper are also relevant for other approaches to bounded rationality.

3.4. The examples revisited

To understand the differences between our approach and the one taken by the Nash implementation literature, it will be instructive to revisit the two examples of Section 2.

Consider the strong Pareto correspondence F^{SP} and the following two environments, somewhat enhanced with respect to Example 1 by the introduction of a uniformly worst alternative w and a third agent. Let $N = \{1, 2, 3\}$ and $Z = \{z, z', w\}$. Let state θ be such that $z \succ_1^{\theta} z' \succ_1^{\theta} w$, $z' \succ_2^{\theta} z \succ_2^{\theta} w$ and $z' \succ_3^{\theta} z \succ_3^{\theta} w$, so that $F^{SP}(\theta) = \{z, z'\}$. And let ϕ differ from θ only in that $z \sim_1^{\phi} z'$, implying that $z \notin F^{SP}(\phi)$.

There is no mechanism to Nash-implement F^{SP} if the domain includes these two environments. The reason is that, if one could find such a mechanism, it should support z by a Nash equilibrium when the state is θ . But those same messages would still constitute a Nash equilibrium of the same mechanism when the state is ϕ , and thus, the mechanism would always include an undesirable Nash equilibrium in state ϕ . On the other hand, F^{SP} satisfies quasimonotonicity and NWA over these two environments. This implies that one can apply to it the canonical mechanism of Theorem 4 to implement F^{SP} in strict Nash equilibria. Further, one can also apply to it the mechanism in the proof of Theorem 2, and make the correspondence coincide with the recurrent classes of BRD processes. Note, however, how in both cases these mechanisms fail to implement in Nash equilibrium. For example, in the mechanism of the proof of Theorem 4, the unanimous profile $(\theta, z, 1)$ is a (non-strict) Nash equilibrium when the state is ϕ : the outcome is z, but agent 1 can induce z' through rule (ii.b). Exactly the same happens in the mechanism of Theorem 2 with the unanimous profile $(\theta, z, z, 1)$ in state ϕ . In addition, both mechanisms may generally have other non-strict Nash equilibria in rules (ii) or (iii) when the modulo game, without the help of the non-veto condition, does not suffice to rule out undesired equilibria. We remark that all these additional

equilibria that prevent Nash implementation of F^{SP} are non-strict, and therefore, non-robust to our dynamic approaches. Consider now the weak Pareto correspondence F^{WP} in a similarly enhanced version of Example 2. Let $N = \{1, 2, 3\}$ and $Z = \{z, z'\}$. Let θ be such that $z \sim_{1}^{\theta} z', z' \succ_{2}^{\theta} z$ and $z' \succ_{3}^{\theta} z$, so that $F^{WP}(\theta) = \{z, z'\}$. And let ϕ differ from θ only in that $z' \succ_1^{\phi} z$, implying that $z \notin F^{WP}(\phi)$. Observe that F^{WP} satisfies monotonicity and no-veto, and thus, by Maskin's (1999) theorem, it is implementable in Nash

equilibrium. But note how, whatever the implementing mechanism is, the Nash equilibrium supporting z in state θ must be non-strict. Since F^{WP} violates quasimonotonicity in any domain that includes these two environments, it is not possible to construct a mechanism with a dynamically robust convergence to z in state θ , which explains why the correspondence fails to be implementable in recurrent strategies of BRD.

4. Non-quasimonotonic rules

Thus far we have seen that quasimonotonicity is the key condition that is necessary for (full) implementability in recurrent strategies. Moreover, with at least three agents, it is also sufficient if one imposes NWA. In this section, we explore the possibilities of implementing non-quasimonotonic rules under more permissive implementability notions. We shall investigate the conditions for partial implementability, as opposed to full.²⁵

Suppose that the planner contents herself with society eventually reaching one of the socially desirable outcomes, but is not concerned with constructing an institution that over time could get to all of them. That is, one can sometimes require that the limiting play of BRD over time lead to outcomes in the SCC, without insisting that, for each of them, there be a dynamic path that end up in it. This motivates the following notion of weak or partial implementability:

An SCC F is partially implementable in recurrent strategies (of BRD) if there exists a mechanism G such that, for every $\theta \in \Theta$, all the outcomes of all recurrent classes of the BRD process applied to its induced game when the preference profile is θ lie in $F(\theta)$.

When one considers SCCs - as opposed to just SCFs, when both notions of partial and full implementability are the same –, one would expect that the conditions that characterize partial implementability be substantially weaker than quasimonotonicity. And indeed, we next show that this is the case.

An SCC *F* is *weakly quasimonotonic* whenever for all $\theta, \phi \in \Theta$, if it is true that for every $a \in F(\theta)$ and for every $i \in N$, $a \succ_i^{\theta} z$ implies that $a \succ_i^{\phi} z$, we have that $F(\theta) \subset F(\phi)$.

For SCCs, weak guasimonotonicity is weaker than guasimonotonicity, as it imposes a restriction on the SCC only when the preferences in two environments are such that the strictly lower contour sets of each agent are nested at every outcome of the SCC. If the range of the SCC is rich, this will often be a requirement that is hard to meet, and therefore, in these cases, weak quasimonotonicity will be vacuously satisfied. Of course, for SCFs, weak quasimonotonicity reduces to quasimonotonicity.

²⁵ A second approach is pursued in Cabrales and Serrano (2010). While insisting on full implementation, in that paper we drop the requirement of converging using BRD no matter what mutation process one might use. One thus gets to implement a larger class of SCFs, but to do so, we make extra specific assumptions on environments, preferences and mutations.

We first state the necessity result:

Theorem 5. If F is partially implementable in recurrent strategies of a BRD process, F is weakly quasimonotonic.

Proof. Let the true preference profile be θ . Because F is partially implementable in recurrent strategies of BRD, the only outcomes that correspond to strategy profiles in recurrent classes of the dynamics is some $a \in F(\theta)$.

Let any $a \in F(\theta)$ be an outcome of a recurrent class of a BRD process. Let a preference profile ϕ be such that for all $b \in F(\theta)$, and for all $i, b >_i^{\theta} z$ implies that $b >_i^{\phi} z$. Since a is an outcome compatible with a recurrent class of the dynamics when preferences are θ , this means that agent *i*'s unilateral deviations from recurrent strategy profiles in that class must yield either *a* again, or outcomes *z* such that $a >_i^{\theta} z$. But this implies that *a* is also supported by recurrent profiles of BRD when preferences are ϕ . Since F is partially implementable in recurrent strategies of BRD, this implies that $a \in F(\phi)$. That is, F must be weakly quasimonotonic. \Box

To obtain the sufficiency result, we need to strengthen the NWA condition. The strong NWA condition is still compatible with many SCCs of interest, for example in standard exchange economic environments in which one can define a worst outcome - the zero bundle - for all agents:

An SCC F satisfies strong-no-worst-alternative (SNWA) whenever for every agent i, every preference profile θ there exists an outcome z_i^{θ} such that for every $a \in F(\theta)$, $a \succ_i^{\theta} z_i^{\theta}$. The sufficiency result for partial implementation in recurrent classes of BRD follows:

Theorem 6. Let $n \ge 3$. If an SCC F satisfies weak quasimonotonicity and SNWA, it is partially implementable in recurrent strategies of BRD.

Proof. We sketch the proof as follows. For the canonical mechanism we construct, we shall need the following definition: Let $\widetilde{F}(\theta) = \{a \in F(\theta) \mid \exists \phi \in \Theta, i \in N, b \in Z \text{ such that } a \succ_i^{\theta} b \text{ and } b \succeq_i^{\phi} a\}$. Note that $\widetilde{F}(\theta) = \emptyset$ implies that $F(\theta) \subset F(\phi)$ for all $\phi \in \Theta$ if F is weakly quasimonotonic.

Consider the mechanism $G = ((M_i)_{i \in N}, g)$, where agent i's message set is $M_i = \Theta \times Z \times Z \times N$. Denote agent i's message $m_i = (\theta_i, a_i, z_i, n_i)$, and the agents' message profile by m. Let the outcome function g be defined by the following rules:

- (i) If for all $i \in N$, $m_i = (\theta, a_i, z_i, n_i)$: (i.a) If $\widetilde{F}(\theta) \neq \emptyset$ and $a_i \in \widetilde{F}(\theta)$ for all *i*, then $g(m) = a_t$, where $t = (\sum_{i \in N} n_i) \pmod{n+1}$. (i.b) If $\widetilde{F}(\theta) = \emptyset$ and $a_i \in F(\theta)$ for all *i*, then $g(m) = a_t$, where $t = (\sum_{i \in N} n_i) \pmod{n+1}$.
- (ii) If for all $j \neq i$, $m_j = (\theta, a_j, z_j, n_j)$ with a_j as in rule (i), and $m_i = (\overline{\phi}, \overline{a_i}, z_i, n_i)$ with either $\phi \neq \theta$, or $a_i \notin \widetilde{F}(\theta) \neq \emptyset$ or $a_i \notin F(\theta)$, $\widetilde{F}(\theta) = \emptyset$, then one can have two cases:

(ii.a) If $z_i \geq_i^{\theta} a_j$ for some $j \in N$, $g(m) = z_i^{\theta}$ – recall the definition of the SNWA property.

(ii.b) If $a_i \succ_i^{\theta} z_i$ for all $j \in N$, $g(m) = z_i$.

(iii) In all other cases, $g(m) = z_t$, where $t = (\sum_{i \in N} n_i) \pmod{n+1}$.

Note how the mechanism makes use of two distinct modulo games, one under rule (i) and a second one under rule (iii). The former is the real novelty of this mechanism, as the recurrent classes of BRD will all make use of it. Indeed, the steps to show that no recurrent classes are compatible with either rule (ii) or rule (iii) follow similar arguments to those in the proof of Theorem 2 and we omit them.

Let θ be the true preference profile. We can have two types of recurrent classes under rule (i). One family of recurrent classes that always exists is the class that contains the truthful profiles (θ, a_i, z_i, n_i) . This can happen both when all $a_i \in F(\theta)$ and $\tilde{F}(\theta) \neq \emptyset$, or when $a_i \in F(\theta)$ and $\tilde{F}(\theta) = \emptyset$. In either case, no individual wants to move out of the class, because exiting the class involves going through rules (ii.a) or (ii.b), which strictly decrease his payoff.

But one may have other (non-truthful) recurrent classes under rule (i). Let $(\phi, a_i, z_i, n_i)_{i \in N}$, be such an arbitrary class. This can happen both when all $a_i \in \widetilde{F}(\phi)$ and $\widetilde{F}(\phi) \neq \emptyset$, or when $a_i \in F(\phi)$ and $\widetilde{F}(\phi) = \emptyset$. If $\widetilde{F}(\phi) = \emptyset$ recall that this implies that $F(\phi) \subset F(\phi')$ for all $\phi' \in \Theta$, if F is weakly quasimonotonic. In particular it is true that $F(\phi) \subset F(\theta)$, and then this agrees with our concept of implementability.

On the other hand, if $\widetilde{F}(\phi) \neq \emptyset$, we have two possible cases. One if the hypothesis of weak quasimonotonicity is met (i.e. for every $a \in F(\phi)$ and for every $i \in N$, $a \succ_i^{\phi} z$ implies that $a \succ_i^{\theta} z$). In this case we have that $F(\phi) \subset F(\theta)$, and hence every $a_i \in F(\theta)$. In the second case the hypothesis of weak quasimonotonicity does not hold. In that case, there exists an individual *i*, alternatives $a \in F(\phi)$ and *b* such that $a \succ_i^{\phi} b$ and $b \succeq_i^{\theta} a$. Note that, if the considered profile is part of a recurrent class, the same class contains a profile where all agents announce the same a (this is so because an agent can always choose an integer not to win the modulo game so as not to change the existing outcome, a weak better reply). Hence, he can weakly improve his payoff by inducing rule (ii.b) choosing (θ, a, b, n_i) . This would contradict that this outcome belongs to a recurrent class. □

A. Cabrales, R. Serrano / Games and Economic Behavior 73 (2011) 360-374

Remark. This proof shows that one can converge to a set of outcomes, where the convergence can be to more than one outcome literally: the typical recurrent class makes use of the modulo game within rule (i). This is the weak convergence notion used in Hart and Mas-Colell (2000), where regret matching dynamics is proposed as a way to converge to the set of correlated equilibrium distributions. It is interesting to note that, while the entire correspondence of correlated equilibrium distributions will satisfy the weak quasimonotonicity condition over large classes of normal form games – viewed as our environments –, certain specific correlated equilibrium distributions – viewed as SCFs – violate quasimonotonicity (see Kar et al., 2010). Taken together, these results show the difficulties of BRD converging to specific correlated equilibrium distributions, whereas convergence to the entire set is possible.

5. Incomplete information

.

This section tackles the extension of our results to incomplete information environments. For ease of exposition, we shall present our results for SCFs. We shall begin with the (almost) characterization of SCFs that are implementable in recurrent strategies of interim BRD. We consider economic environments for our results in interim BRD, although our last subsection, aiming to connect with other approaches through implementation in strict Bayesian equilibria, dispenses with this assumption.

5.1. Necessary and sufficient conditions

We now describe an incomplete information environment.

Consider for simplicity economic environments. More precisely, let agent *i*'s consumption set be a finite set, $X_i \subset \mathbb{R}^l_+$, where $0 \in X_i$. 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:

$$Z = \left\{ (x_i)_{i \in \mathbb{N}} \in \prod X_i \colon \sum_{i \in \mathbb{N}} x_i \leq \omega \right\}.$$

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.²⁶

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 also make the following assumptions on preferences:

- (1) No consumption externalities: for each state θ , $u_i : X_i \mapsto \mathbb{R}$. That is, an agent's ex-post utility function in each state depends on the bundle of goods that he consumes, and not on other agents' bundles.
- (2) Strictly increasing ex-post utility functions: For all *i*, for all θ and for all $x_i \in X_i$, if $y_i \ge x_i$, $u_i(y_i, \theta) > u_i(x_i, \theta)$.²⁷ Note how this implies that 0 is the worst bundle for every agent.

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 $g : M \mapsto Z$. 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, agents are replaced by other agents who share the same preferences and prior beliefs as their predecessors, but are not aware of their experience).²⁸

²⁶ We 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.

²⁷ For vectors $x_i, y_i \in X_i$, we use the standard conventions: $x_i \ge y_i$ whenever $x_{il} \ge y_{il}$ with at least one strict inequality; and $x_i \gg y_i$ whenever $x_{il} > y_{il}$ for every commodity *l*.

²⁸ There 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 et al. (2004) for an appraisal of such different modeling choices.

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:

$$\sum_{\theta_{-i}\in\Theta_{-i}}q_{i}(\theta_{-i}|\theta_{i})u_{i}\left(g\left(m_{i}',m_{-i}^{t}(\theta_{-i})\right),(\theta_{i},\theta_{-i})\right) \geq \sum_{\theta_{-i}\in\Theta_{-i}}q_{i}(\theta_{-i}|\theta_{i})u_{i}\left(m^{t}(\theta),(\theta_{i},\theta_{-i})\right)$$

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

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.

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 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 ,

$$\sum_{\theta_{-i}\in\Theta_{-i}}q_{i}(\theta_{-i}|\theta_{i})u_{i}\left(f(\theta),(\theta_{i},\theta_{-i})\right)>\sum_{\theta_{-i}\in\Theta_{-i}}q_{i}(\theta_{-i}|\theta_{i})u_{i}\left(f\left(\theta_{i}',\theta_{-i}\right),(\theta_{i},\theta_{-i})\right)$$

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.

Theorem 7. If f is implementable in recurrent strategies of an interim BRD process, f is incentive compatible. Furthermore, if at least one of the recurrent classes is a singleton, f is strictly incentive compatible.

Proof. Suppose that f is implementable in recurrent strategies of interim BRD. This means that, for this Markov process, there is a unique outcome supported by the recurrent classes of the process, and this outcome is f. Since f is the outcome of any such recurrent set of BRD, it must be incentive compatible.

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

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

$$U_{i}(f|\theta_{i}) > U_{i}(y_{\theta_{i}'}|\theta_{i}) \quad \forall \theta_{i}' \in \Theta_{i} \quad \text{implies} \quad U_{i}(f \circ \alpha|\theta_{i}) > U_{i}(y \circ \alpha|\theta_{i}) \tag{1}$$

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 recurrent strategies.

Theorem 8. If f is implementable in recurrent strategies of an interim BRD process, f is Bayesian quasimonotonic.

Proof. Suppose that f is implementable in recurrent strategies of interim BRD. This means that f is the only outcome of the recurrent classes of the 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 also 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 BRD, any such deviations that change the outcome fall in the strictly lower contour set of the interim preferences for type θ_i at *f*. That is, $U_i(f|\theta_i) > U_i(y_{\theta'_i}|\theta_i)$. 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 – Eq. (1) –, that any such $y_{\theta'_i}$, when the deception α is used, is such that $U_i(f \circ \alpha | \theta_i) > U_i(y \circ \alpha | \theta_i)$ for every i and θ_i . Consider the

A. Cabrales, R. Serrano / Games and Economic Behavior 73 (2011) 360-374

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 $U_i(f \circ \alpha | \theta_i) > U_i(y \circ \alpha | \theta_i)$. 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 interim BRD process.

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

5.1.2. Sufficiency

Our next sufficiency result follows:

Theorem 9. Suppose the environments satisfy Assumptions (1) and (2) in each state. Let $n \ge 3$. If an SCF f is ε -secure in each state, strictly incentive compatible and Bayesian quasimonotonic, f is implementable in recurrent strategies of interim BRD.

Proof. 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* ∈ *N*, *m_i² = f*, *g*(*m*) = *f*(*m*¹).
(ii) If for all *j* ≠ *i m_j² = f* and *m_i² = y* ≠ *f*, one can have two cases:
(ii.a) If there exist types θ_i, θ'_i ∈ Θ_i such that U_i(y_{θ'_i}|θ_i) ≥ U_i(*f*|θ_i), *g*(*m*) = (*f_i*(*m*¹) − β, *f_{-i}*(*m*¹)), where *f_i*(*m*¹) ≥ $f_i(m^1) - \beta \in X_i$.

- (ii.b) If for all $\theta_i, \theta_i' \in \Theta_i, U_i(y_{\theta_i'}|\theta_i) < U_i(f|\theta_i), g(m) = y(m^1).$
- (iii) In all other cases, g(m) = 0.

Following similar steps as in the proof of Theorem 2, properly adapted to economic environments and using the zero outcome instead of the modulo game, one can show that all recurrent classes of the dynamics 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 $f \circ \alpha = f$. \Box

Remark. Theorems 7, 8 and 9 provide almost a characterization of the rules that are implementable in recurrent strategies of BRD in economic environments with at least three agents. We shall consider non-economic environments and the more general implications of these results in the last subsection of the current section, through the notion of implementability in strict Bayesian equilibria.

5.2. Towards a more general theory: strict Bayesian implementation

We close the paper by returning to environments in which the economic assumption is not made. Paraphrasing the argument in Section 3.3, the reader may wonder how relevant the results obtained in the current incomplete information section for adaptive interim better-replies would be to alternative ways to model bounded rationality in implementation settings under incomplete information. Perhaps a common thread could be found through the use of strict Bayesian equilibrium as a solution concept - it is plausible to think that these equilibria, based on strict inequalities, will have more robust properties.29

The results on strict Bayesian implementation offered below may serve this role of connection with the previous results in the current section. Further, they have interest in their own right, Jackson (1991) offered only a sufficiency result for Bayesian implementation in general environments, and no characterization is available to date, to the best of our knowledge. We shall provide such a characterization for strict Bayesian implementation.

An SCF f is implementable in strict Bayesian equilibrium if there exists a mechanism G such that the unique outcome of all strict Bayesian equilibria of its induced game is f.

We adapt the definition of the NWA property to general environments with incomplete information as follows:

An SCF f satisfies the Bayesian no-worst-alternative (BNWA) property if for each $i \in N$ there exists an SCF y satisfying:

1. $U_i(y|\theta_i) \ge U_i(y_{\theta'_i}|\theta_i) - \varepsilon$ for all θ_i, θ'_i , and

2. if $f \circ \alpha \neq f$, $U_i(f \circ \alpha | \theta_i) \leq U_i(y \circ \alpha | \theta_i)$ for some $\theta_i \in \Theta_i$,

such that for all $\theta_i \in \Theta_i$, $U_i(f|\theta_i) > U_i(y|\theta_i)$.³⁰

²⁹ Although more work is needed in the general area of learning under incomplete information, see, e.g., Jensen et al. (2005) for a result on stochastic stability that selects among strict Bayesian equilibria.

 $^{^{30}}$ For the results below, the maximum value of ε allowed in the first requirement for y corresponds to the maximum amount of misrepresentation within the direct mechanism for y that compensates for the gap created by the strict inequality between the interim utilities of f and of y.

First, we state the following necessity result:

Theorem 10. If an SCF f is implementable in strict Bayesian equilibrium, it is strictly incentive compatible and Bayesian quasimonotonic. Furthermore, if the implementation is achieved via a non-imposing mechanism, f must also satisfy BNWA.

Proof. The proof of the first part is similar to the proofs of Theorems 7 and 8. The proof of the second part is left to the reader. 🗆

Next, we provide a sufficiency result for environments with at least three agents. The two results together essentially provide a characterization of strict Bayesian implementation for this case:

Theorem 11. Let $n \ge 3$. If an SCF f is strictly incentive compatible and Bayesian quasimonotonic, and satisfies BNWA, it is implementable in strict Bayesian equilibrium.

Proof. We construct the following canonical mechanism $G = ((M_i)_{i \in N}, g)$, where agent *i*'s message set $M_i = \Theta_i \times A \times N$. Denote $m_i = (m_i^1, m_i^2, m_i^3)$. The outcome function g is defined in the following rules:

- (i) If for every agent $i \in N$, $m_i^2 = f$ and $m_i^3 = 1$, $g(m) = f(m^1)$. (ii) If for all $j \neq i$ $m_j^2 = f$, $m_j^3 = 1$, and $(m_i^2, m_i^3) = (z, n_i) \neq (f, 1)$, one can have two cases:
 - (ii.a) If there exist types $\theta_i, \theta'_i \in \Theta_i$ such that $U_i(z_{\theta'_i}|\theta_i) \ge U_i(f|\theta_i), g(m) = y(m^1)$, where y is taken from the definition of the BNWA property.
- (ii.b) If for all $\theta_i, \theta'_i \in \Theta_i, U_i(z_{\theta'_i}|\theta_i) < U_i(f|\theta_i), g(m) = z(m^1).$ (iii) In all other cases, $g(m) = z^t(m^1)$, where $m_2^t = z^t$, $t = (\sum_{i \in N} m_i^3) \pmod{n+1}$.

Since each agent can induce one of his top-ranked SCFs with the modulo game, there cannot be strict Bayesian equilibria under either rule (ii) or (iii). Clearly, the truthful profile under rule (i) is a strict Bayesian equilibrium, by strict incentive compatibility and by the first requirement for y in the BNWA property, after unilateral deviations would induce the two versions of rule (ii). Finally, Bayesian quasimonotonicity and the second statement in the requirements for y in the BNWA property prevent the existence of a strict Bayesian equilibrium under rule (i) in which a deception α is being used and $f \circ \alpha \neq f$. \Box

6. Conclusion

This paper has studied the classic implementation problem in evolutionary settings. In particular, necessary and sufficient conditions for implementability in recurrent strategies of BRD processes 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. In the case of incomplete information, incentive compatibility and a variant of Bayesian monotonicity show up as necessary conditions for evolutionary implementation in our sense. The analysis of the implementation problem under alternative forms of bounded rationality should be the subject of further research, and to this end, we have also derived characterizations for implementability in strict equilibria, both under complete and incomplete information, which are closely related to the recurrent implementation results of the paper.

Acknowledgments

We thank Dirk Bergemann, Jeff Ely, Matt Jackson, Takashi Kunimoto, Eric Maskin, Stephen Morris, Andrea Prat, Olivier Tercieux, Norov Tumennasan and Fernando Vega-Redondo for helpful conversations and encouragement. The comments of an Associate Editor and three referees helped 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.

References

- Abreu, D., Matsushima, H., 1992. Virtual implementation in iteratively undominated strategies: complete information. Econometrica 60, 993-1008. Abreu, D., Matsushima, H., 1994. Exact implementation. J. Econ. Theory 64, 1-19.
- Aghion, P., Fudenberg, D., Holden, R., Kunimoto, T., Tercieux, O., 2010. Subgame-perfect implementation under value perturbations. Mimeo.
- Benoit, J.-P., Ok, E., 2008. Nash implementation without no-veto power. Games Econ. Behav. 64, 51-67.
- Bergemann, D., Morris, S., Tercieux, O., 2010. Rationalizable implementation. Mimeo.
- Bochet, O., 2007. Nash implementation with lottery mechanisms. Soc. Choice Welfare 28, 111-125.
- Cabrales, A., 1999. Adaptive dynamics and the implementation problem with complete information. J. Econ. Theory 86, 159-184.
- Cabrales, A., Charness, G., Corchón, L., 2003. An experiment on Nash implementation. J. Econ. Behav. Organ. 51, 161-193.

Cabrales, A., Ponti, G., 2000. Implementation, elimination of weakly dominated strategies and evolutionary dynamics. Rev. Econ. Dynam. 3, 247-282.

Cabrales, A., Serrano, R., 2010. Stochastically stable implementation. Mimeo.

Chen, Y., Plott, C.R., 1996. The Groves-Ledyard mechanism: an experimental study of institutional design. J. Public Econ. 59, 335-364.

Chen, Y., Tang, F., 1998. Learning and incentive compatible mechanisms for public goods provision: an experimental study. J. Polit. Economy 106, 633–662.

Chung, K.-S., Ely, J., 2003. Implementation with near-complete information. Econometrica 71, 857–887.

Corchón, L., 2009. The theory of implementation: what did we learn? In: Meyers, R. (Ed.), Encyclopedia of Complexity and Systems Science. Springer, New York.

Dawid, H., 1999. On the convergence of genetic learning in a double auction model. J. Econ. Dynam. Control 23, 1545–1567.

Dekel, E., Fudenberg, D., Levine, D.K., 2004. Learning to play Bayesian games. Games Econ. Behav. 46, 282-303.

de Trenqualye, P., 1988. Stability of the Groves and Ledyard mechanism. J. Econ. Theory 46, 164–171.

de Trenqualye, P., 1989. Stable implementation of Lindahl allocations. Econ. Letters 29, 291–294.

Eliaz, K., 2002. Fault-tolerant implementation. Rev. Econ. Stud. 69, 589-610.

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

Groves, T., Ledyard, J., 1977. Optimal allocation of public goods: a solution to the free rider problem. Econometrica 45, 783-809.

Hart, S., Mas-Colell, A., 2000. A simple adaptive procedure leading to correlated equilibrium. Econometrica 68, 1127-1150.

Hon-Snir, S., Monderer, D., Sela, A., 1998. A learning approach to auctions. J. Econ. Theory 82, 65–88.

Jackson, M.O., 1991. Bayesian implementation. Econometrica 59, 461–477.

Jackson, M.O., 2001. A crash course in implementation theory. Soc. Choice Welfare 18, 655-708.

Jensen, M., Sloth, B., Whitta-Jacobsen, H.J., 2005. The evolution of conventions under incomplete information. Econ. Theory 25, 171-185.

Jordan, J.S., 1986. Instability in the implementation of Walrasian allocations. J. Econ. Theory 39, 301-328.

Kar, A., Ray, I., Serrano, R., 2010. A difficulty in implementing correlated equilibrium distributions. Games Econ. Behav. 69, 189–193.

Karlin, S., Taylor, H.M., 1975. A First Course in Stochastic Processes. Academic Press, New York.

Kartik, N., 2009. Strategic communication with lying costs. Rev. Econ. Stud. 76, 1359–1395.

Luce, D., 1959. Individual Choice Behavior. Wiley, New York.

Maskin, E.S., 1999. Nash equilibrium and welfare optimality. Rev. Econ. Stud. 66, 23-38.

Maskin, E.S., Sjöström, T., 2002. Implementation theory. In: Arrow, K.J., Sen, A., Suzumura, K. (Eds.), Handbook of Social Choice and Welfare, vol. I. Elsevier Science B.V., New York.

Mathevet, L., 2007. Supermodular Bayesian implementation: learning and incentive design. Mimeo. California Institute of Technology.

McFadden, 1973. Conditional logit analysis of individual choice behavior. In: Zarembka, P. (Ed.), Frontiers of Econometrics. Academic Press, New York.

McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. Games Econ. Behav. 10, 6–38.

Muench, T., Walker, M., 1984. Are Groves-Ledyard equilibria attainable? Rev. Econ. Stud. 50, 393-396.

Palfrey, T.R., 2002. Implementation theory. In: Aumann, R.J., Hart, S. (Eds.), Handbook of Game Theory with Economic Applications, vol. III. Elsevier Science, New York.

Palfrey, T.R., Srivastava, S., 1989. Implementation with incomplete information in exchange economies. Econometrica 57, 115-134.

Postlewaite, A., Schmeidler, D., 1986. Implementation in differential information economies. J. Econ. Theory 39, 14-33.

Repullo, R., 1987. A simple proof of Maskin theorem on Nash implementation. Soc. Choice Welfare 4, 39-41.

Roth, A., Postlewaite, A., 1977. Weak versus strong domination in a market for indivisible goods. J. Math. Econ. 4, 131-137.

Samuelson, L., 1997. Evolutionary Games and Equilibrium Selection. MIT Press, Cambridge, MA.

Samuelson, L., Zhang, J., 1992. Evolutionary stability in asymmetric games. J. Econ. Theory 57, 363-392.

Sandholm, W.H., 2002. Evolutionary implementation and congestion pricing. Rev. Econ. Stud. 69, 667–689.

Sandholm, W.H., 2005. Negative externalities and evolutionary implementation. Rev. Econ. Stud. 72, 885–915.

Sanver, M., 2006. Nash implementing non-monotonic social choice rules by awards. Econ. Theory 28, 453-460.

Saran, R., Serrano, R., 2010. Ex-post regret learning in games with fixed and random matching: the case of private values. Working Paper 2010-11. Department of Economics, Brown University.

Serrano, R., 2004. The theory of implementation of social choice rules. SIAM Rev. 46, 377-414.

Shapley, L., Scarf, H., 1974. On cores and indivisibility. J. Math. Econ. 1, 23-28.

Sjöström, T., 1994. Implementation in undominated Nash equilibria without using integer games. Games Econ. Behav. 6, 502-511.

Tumennasan, N., 2008. Implementation in quantal response equilibria. Mimeo. Department of Economics, Brown University.

Vega-Redondo, F., 1989. Implementation of Lindahl equilibrium: an integration of the static and dynamic approaches. Math. Soc. Sci. 18, 211-228.

Vega-Redondo, F., 1996. Evolution, Games, and Economic Behavior. Oxford University Press, New York.

Walker, M., 1984. A simple auctioneerless mechanism with Walrasian properties. J. Econ. Theory 32, 111-127.

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

Young, H.P., 1998. Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton University Press, Princeton, NJ.

A. Cabrales, R. Serrano / Games and Economic Behavior 73 (2011) 360-374

Author's personal copy