Mechanism Design: Recent Developments*

Sandeep Baliga
Northwestern University

Tomas Sjöström
Rutgers University

March 8, 2007

1 Possibility Results and Robustness

Game theory provides methods to predict the outcome of a given game. Mechanism design concerns the *reverse* question: given some desirable outcome, can we design a game which produces it?

Formally, the *environment* is $\langle A, N, \Theta \rangle$, where A is a set of feasible and verifiable *alternatives* or *outcomes*, $N = \{1, ..., n\}$ is a set of *agents*, and Θ is a set of possible *states of the world*. We focus on *private values* environments, where a state is $\theta = (\theta_1, ..., \theta_n) \in \times_i \Theta_i = \Theta$, each agent i knows his own "type" $\theta_i \in \Theta_i$, and his payoff $u_i(a, \theta_i)$ depends only on the chosen alternative and his own type. (This does not rule out the possibility that the agents know something about each others' types.) If values are not private, then they are said to be *interdependent*. A *mechanism* or *contract* $\Gamma = (S, h)$ specifies a set of feasible actions S_i for each agent i, and an outcome function $h : S \equiv \times_{i=1}^n S_i \to A$. An outside party (a principal or social planner), or the agents themselves, want to design a mechanism which produces optimal outcomes. These are often represented by a *social choice rule* (SCR) $F : \Theta \to A$. A *social choice function* (SCF) is a single-valued SCR. Implicitly, it is assumed that the mechanism designer does not know the true θ, and this lack of information makes it impossible for her to directly choose an outcome in $F(\theta)$. Instead, she uses the more roundabout method of designing a mechanism which produces an outcome in $F(\theta)$, *whatever the true θ may be.*

In a *revelation mechanism*, each agent simply reports what he knows (so if agent i only knows θ_i then $S_i = \Theta_i$). By definition, an *incentive compatible* revelation mechanism has a *truthful* Bayesian-Nash equilibrium, i.e., it achieves *truthful implementation*. Truthful implementation plays an important role in the theory because of the *revelation principle* (see the detailed discussion in Myerson, 2006). Myerson (2006) surveys the early literature on truthful implementation. This literature produced powerful results on optimal mechanisms for auction design, bargaining problems, and other applications. However, some

*This survey is forthcoming in *The New Palgrave Dictionary of Economics*, Second Edition, edited by Larry Blume and Steven Durlauf.

difficult issues were not resolved by the early literature. For example, it was generally assumed that the agents and the principal share a common prior over Θ, that the principal can *commit* to a mechanism, that agents cannot side-contract, that agents play equilibrium strategies even in very complicated mechanisms, etc. The recent literature which deals with these issues is surveyed here. In addition, we note that the notion of truthful implementation has a drawback: it does not rule out the possibility that non-truthful equilibria also exist, and these may produce suboptimal outcomes. Moreover, a non-truthful equilibrium may Pareto dominate the truthful equilibrium and hence may be a focal point for coordination of players' actions. To rule out these possibilities, we may require *full implementation*: for all $\theta \in \Theta$, the set of equilibrium outcomes should precisely equal $F(\theta)$.

Maskin (1999) assumed *complete information*: each agent knows the true θ. If $n \geq 3$ agents know θ, then any SCF can be truthfully implemented: let the agents report θ, and if at least $n-1$ agents announce the same θ then implement the outcome $F(\theta)$. Unilateral deviations from a consensus are disregarded, so truth-telling is a Nash equilibrium. Of course, this revelation mechanism will also have non-truthful equilibria. For full implementation, more complex mechanisms are required. (Even if $n = 2$, any SCF can be truthfully implemented if the principal can credibly threaten to "punish" both agents if they report different states; in an economic environment, this might be achieved by making each agent pay a fine.)

A necessary condition for full Nash implementation is *(Maskin) monotonicity* (Maskin, 1999). Intuitively, monotonicity requires that moving an alternative *up* in the agents' preference rankings should not make it *less* likely to be optimal. This condition can be surprisingly difficult to satisfy. For example, if the agents can have any complete and transitive preference relation on A, then any Maskin monotonic SCF must be a constant function (Saijo, 1987)! The situation is quite different if we consider *refinements* of Nash equilibrium. For example, there is a sense in which almost *any* (ordinal) SCR can be fully implemented in *undominated Nash equilibrium* when the agents have complete information (Palfrey and Srivastava, 1991; Jackson, Palfrey and Srivastava, 1994; Sjöström, 1994). Chung and Ely (2003) showed that this possibility result is not robust to small perturbations of the information structure that violate private values (there is a small chance that agent i knows more about agent j's preferences than agent j does). The violation of private values is key. For example, in Sjöström's (1994) mechanism, an agent who knows his own preferences can eliminate his dominated strategies, and a second round of elimination of *strictly* dominated strategies generates the optimal outcome. This construction is robust to small perturbations that respect private values.

A different kind of robustness was studied by McLean and Postlewaite (2002). Consider an economic environment where each agent i observes an independently drawn signal t_i which is correlated with the state θ, but no agent observes θ. The complete information structure is approximated by letting each agent's signal be very accurate. With complete information, any SCF can be truthfully implemented. McLean and Postlewaite (2002) show robustness to perturbations of the information structure: any outcome can be approximated by an incentive-compatible allocation, if the agents' signals are accurate enough. There is no need to assume private values.

The literature on *Bayesian mechanism design* typically assumes each agent i knows only his own type $\theta_i \in \Theta_i$, the agents share a common prior p over $\Theta \equiv \times_{i=1}^n \Theta_i$, and the principal knows p. In fact, for truthful implementation with $n \geq 3$, the assumption that the principal knows p is redundant. Suppose for any common prior p on Θ, there is an incentive-compatible revelation mechanism $\Gamma_p = (\times_{i=1}^n \Theta_i, h_p)$. By definition, Γ_p truthfully implements the SCF $F_p \equiv h_p$. The mechanism Γ_p is "parametric", i.e., depends on p. To be specific, consider a quasi-linear public goods environment with independent types, and suppose Γ_p is the well-known mechanism of d'Aspremont and Gérard-Varet (1979). Now consider the following non-parametric mechanism Γ: each agent i announces p and θ_i. If at least $n-1$ agents report the same p, the outcome is $h_p(\theta_1, ..., \theta_n)$. Now, if agent i thinks everyone will announce p truthfully, he may as well do so. If in addition he thinks the other agents report θ_{-i} truthfully, then he should announce θ_i truthfully by incentive compatibility of Γ_p. Therefore, for any common prior p, the non-parametric mechanism Γ truthfully implements F_p. In this sense, the principal can use Γ to extract the agents' shared information about p. Of course, this particular mechanism also has non-truthful equilibria. Choi and Kim (1999) showed how non-parametric *full* implementation can be achieved. Naturally, their mechanism is quite complex. Suppose we restrict attention to "simple" non-parametric mechanisms, where each agent i only reports θ_i, truthfully in equilibrium. Then the necessary and sufficient condition for full non-parametric Bayesian-Nash implementation for any common prior p is: incentive compatibility plus the *rectangular property* (Cason et. al., 2006).

The d'Aspremont and Gérard-Varet (1979) mechanism is budget-balanced and surplus-maximizing. The above argument shows that such outcomes can be truthfully implemented by a non-parametric mechanism in quasi-linear environments with independent types. As is well-known, this cannot be achieved by any dominant strategy mechanism. Thus, in general, non-parametric truthful implementation is easier than dominant strategy implementation. However, there are circumstances where the two concepts coincide. Bergemann and Morris (2005a) consider a model where each agent i has a *payoff type* $\theta_i \in \Theta_i$ and a *belief type* π_i. The payoff type determines the payoff function $u_i(a, \theta_i)$, while the belief type determines beliefs over other agents' types. The set of socially optimal outcomes $F(\theta)$ depends on payoff types, but not on beliefs. Bergemann and Morris (2005a) show that in quasi-linear environments with no restrictions on side-payments (hence no budget-balance requirement), truthful implementation for all possible type spaces with a common prior implies dominant strategy implementation. (For related results, see Section 4 below.)

Bergemann and Morris (2005b) consider *full* implementation of SCFs in a similar framework. The SCF $F : \Theta \to A$ is *fully robustly implemented* if there exists a mechanism which fully implements F on all possible type spaces. There is no common prior assumption. Full robust implementation turns out to be equivalent to implementation using iterated elimination of strictly dominated strategies. Although a demanding concept, there are situations where full robust implementation is possible. For example, a VCG (Vickrey-Clarke-Groves) mechanism in a public goods economy with private values and strictly concave valuation functions achieves implementation in strictly dominant strategies. However, Bergemann and Morris (2005b) show the impossibility of full robust implementation when values are sufficiently interdependent.

A generalization of Maskin monotonicity called *Bayesian monotonicity* is necessary for ("parametric") full Bayesian-Nash implementation (Postlewaite and Schmeidler, 1986; Palfrey and Srivastava, 1989a; Jackson, 1991). Again, refinements lead to possibility results (Palfrey and Srivastava, 1989b). Another way to expand the set of implementable SCRs is *virtual* implementation (Abreu and Sen, 1991; Duggan, 1997). Serrano and Vohra (2001) argue that the sufficient conditions for virtual implementation are in fact quite strong.

The work discussed so far is *consequentialist*: only the final outcome matters. The mechanisms are clearly not meant to be descriptive of real-world institutions. For example, they typically require the agents to report "all they know" before any decision is reached, an extreme form of centralized decision making hardly ever encountered in the real world. (The question of how much information must be transmitted in order to implement a given SCR is addressed by Hurwicz and Reiter, 2006, and Segal, 2004.) Delegating the power to make (verifiable) decisions to the agents would only create additional "moral hazard" constraints, as discussed by Myerson (2006). Since centralization eliminates these moral hazard constraints, it typically strictly dominates decentralization in the basic model. However, as discussed below, by introducing additional aspects such as renegotiation and collusion, we can frequently prove the optimality of more realistic decentralized mechanisms. The implicit assumption is that decentralized decision making is in itself a good thing, which is a mild form of non-consequentialism. (Other non-consequentialist arguments are discussed in Section 4 below.) We might add that there is of course no way to eliminate the moral hazard constraints if the agents take *unverifiable* decisions that cannot be contracted upon. In this case, the issue of centralization versus decentralization of decisions is moot.

2 Renegotiation and Credibility

Suppose $n = 2$ and both agents know the true θ. If a revelation mechanism is used and the agents announce different states, then we cannot identify a deviator from a "consensus", so it may be necessary to punish *both* agents in order to support a truth-telling equilibrium. But this threat is not credible if the agents can avoid punishment by *renegotiating* the outcome. Maskin and Moore (1999) capture the renegotiation process by an exogenously given function $r : A \times \Theta \to A$ which maps outcome a in state θ into an efficient outcome $r(a, \theta)$. They derive an incentive-compatibility condition which is necessary for truth-telling when $n = 2$, and show that *renegotiation monotonicity* is necessary for full Nash implementation (see also Segal and Whinston, 2002).

The idea that renegotiation may preclude the implementation of the first-best outcome, even when information is complete, has received attention in models of bilateral trade with relationship-specific investments (the hold-up problem). It is possible to implement the first-best outcome if trade is one-dimensional and investments are "selfish", in the sense that each agent's investment does not directly influence the other agent's payoff (Nöldeke and Schmidt, 1995; Edlin and Reichelstein, 1996). But if investments are not selfish then the first-best cannot always be achieved, while the second-best can be implemented without any explicit contract (Che and Hausch, 1999). Segal (1999) found a similar result in a model with k goods and selfish investments, for k large (see also Maskin and Tirole, 1999, and Hart and Moore, 1999). Adding a third party to these two-agent models can alleviate the problem of renegotiation. The third party can act as a "budget breaker" to whom fines can be paid, thereby ensuring ex post efficient outcomes. Although the third party is often believed to cause collusion, in theory collusion can often be ruled out (Baliga and Sjöström, 2006).

Credibility and renegotiation also impact trading with asymmetric information. Suppose the seller can produce goods of different quality, but the buyer's valuation is his private information. It is typically second-best optimal for the seller to offer a contract such that low valuation buyers consume less than first-best quality ("underproduction"), while high valuation buyers enjoy "information rents". Incentive-compatibility guarantees that the buyer reveals his true valuation. Now suppose trading takes place twice, and the buyer's valuation does not change. Suppose the seller cannot credibly commit to a long-run (two-period) contract. If the buyer reveals his true valuation in the first period, then in the second period the seller will leave him no rent. This is typically not the second-best outcome. The seller may prefer a "pooling" contract which does not fully reveal valuations in the first period, a commitment device which limits his ability to extract second period rents. This idea has important applications. When a regulator cannot commit to a long-run contract, a regulated firm may hide information or exert less effort to cut costs, the *ratchet effect* (Freixas, Guesnerie and Tirole, 1985). A borrower may not exert effort to improve a project knowing that a lender with deep pockets will bail him out, the *soft budget constraint* (Dewatripont and Maskin, 1995a). These problems are exacerbated if the principal is well informed and cannot commit not to use his information. Institutional or organizational design can alleviate the problems. By committing to acquire less information via "incomplete contracts", or by maintaining an "arm's length relationship", the principal can improve efficiency (Dewatripont and Maskin, 1995b; Crémer, 1995). Less frequent regulatory reviews offset the ratchet effect, and a decentralized credit market helps to cut off borrowers from future funding. Long-run contracts is a possible solution, but they may be vulnerable to renegotiation (Dewatripont, 1989). In particular, the second period outcome may be renegotiated if quality levels are known to be different from the first-best. Again, some degree of pooling may be optimal.

If the principal cannot commit even to short-run contracts, then after receiving the agents' messages, she always chooses an outcome that is optimal given her beliefs. She cannot credibly threaten punishments that she would not want to carry out. Refinements proposed in the cheap-talk literature suggest that a putative pooling equilibrium may be destroyed if an agent can reveal information by "objecting" in a credible way. This leads to a necessary condition for full implementation which is reminiscent of Maskin monotonicity (Baliga, Corchón and Sjöström, 1997). This condition involves the principal's preferences. Even if an SCR is Nash implementable when the principal can fully commit to a mechanism, there may not exist *any* preference ordering for the principal which would make implementation possible if she cannot commit.

3 Collusion

A large literature on collusion was inspired by Tirole (1986). A key contribution was made by Laffont and Martimort (1997), who assumed an uninformed third party proposes side-contracts. This circumvents the signalling problems that might arise if a privately informed agent makes collusive proposals. A side-contract for a group of colluding agents is a *collusive mechanism* which must respect incentive compatibility, individual rationality and feasibility constraints. The original mechanism Γ, designed by the principal, is called the *grand mechanism.* The objective is to design an optimal grand mechanism when collusion is possible. Typically, collusion imposes severe limits on what can be achieved.

Baliga and Sjöström (1998) study a model with moral hazard and limited liability. Two agents share information not known to the principal: agent 1's effort is observed by both agents. Agent 2's effort is known only to himself. In the absence of collusion, the optimal grand mechanism specifies a "message game": agent 2 reports agent 1's effort to the principal. Now suppose the agents can side-contract on agent 1's effort, but not on agent 2's effort (which is unobserved). Side-contracts can specify side-transfers as a function of realized output, but must respect limited liability. This collusion may destroy centralized "message games", and we obtain a theory of optimal delegation of decision-making. For some parameters, it is optimal for the principal to contract only with agent 2, and let agent 2 subcontract with agent 1. This is intuitive, since agent 2 observes agent 1's effort and can contract directly on it. More surprisingly, there are parameter values where it is better for the principal to contract only with agent 1.

Mookherjee and Tsumagari (2004) study a similar model, but with adverse selection: the agents privately observe their own production costs. In this model, delegating to a "prime supplier" creates "double marginalization of rents": the prime supplier uses underproduction to minimize the other agent's information rent. A centralized contract avoids this problem. Hence, in this model delegation is always strictly dominated by centralization, even though the agents can collude.

Mookherjee and Tsumagari (2004) assume the agents can side-contract before deciding to participate in the grand contract. Che and Kim (2005) assume side-contracting occurs only after the decision to participate in the grand mechanism has been made. They show that collusion does not limit what the principal can achieve. Hence, the timing of side-contracting is important in determining its impact. In a complete information environment, Sjöström (1999) showed that neither renegotiation nor collusion limit the possibility of undominated Nash implementation when $n \geq 3$.

4 Other Theoretical Issues

In quasi-linear environments with uncorrelated types, there exists incentive-compatible mechanisms which maximize the social surplus (e.g., d'Aspremont and Gérard-Varet, 1979). But the principal cannot extract all the surplus: the agents must get informational rents. However, Crémer and McLean (1988) showed that the principal can extract all the surplus in auctions with *correlated types*. McAfee and Reny (1992) extended this result to general quasi-linear environments.

Jehiel and Moldovanu (2001) considered a quasi-linear environment with multi-dimensional (uncorrelated) types and interdependent values. Generically, a standard revelation mechanism cannot be designed to extract information about multi-dimensional types, and no incentive-compatible and surplus-maximizing mechanism exists. There is also a literature on optimal mechanisms for a profit maximizing monopolist when consumers have multi-dimensional types and private values (Armstrong, 1996). Mezzetti (2005) presents an ingenious *two-stage* mechanism which maximizes the surplus in interdependent values environments, even when types are independent and multi-dimensional. In the first stage, the mechanism specifies an outcome decision but not transfers. Transfers are determined in the second stage by reports on payoffs realized by the outcome decision. Mezzetti (2006) shows that the principal can sometimes extract all the surplus by this method, even if types are independent.

Incentive-compatibility does not require that each agent has a dominant strategy. Nevertheless, incentive-compatible outcomes can often be replicated by dominant strategy mechanisms (Mookherjee and Reichelstein, 1992). In quasi-linear environments, incentive-compatible mechanisms that maximize the social surplus are *payoff-equivalent* to dominant strategy (VCG) mechanisms (Krishna and Perry, 1997; Williams, 1999). However, as pointed out above, dominant strategies (but not incentive-compatibility) rules out budget-balance.

Bergemann and Välimäki (2002) assume agents can update a common prior by costly information acquisition. Suppose a single-unit auction has two bidders i and j who observe statistically independent private signals θ_i and θ_j. Bidder i's valuation of the good is $u_i(\theta_i, \theta_j) = \alpha\theta_i + \beta\theta_j$, where $\alpha > \beta > 0$. Thus, values are interdependent. Efficiency requires that bidder i gets the good if and only if $\theta_i \geq \theta_j$. Suppose bidders report their signals, the good is allocated efficiently given their reports, and the winning bidder i pays the price $(\alpha + \beta)\,\theta_j$. This VCG mechanism is incentive-compatible (Maskin, 1992). If bidder i acquires negative information which causes him to lose the auction, then he imposes a negative externality on the other bidder (as $\beta > 0$). This implies the bidders have an incentive to collect too much information. Conversely, there is an incentive to collect too little information when $\beta < 0$. Bergemann and Välimäki (2002) provide a general analysis of these externalities. Similar externalities occur when members of a committee must collect information before voting. If the committee is large, each vote is unlikely to be pivotal, and free-riding occurs. Persico (2004) shows how the optimal committee is designed to encourage the members to collect information.

Consequentialism is rejected by a literature which focusses on agents' *rights*. For example, suppose a mechanism implements *envy-free outcomes*. An agent might still feel unfairly treated if his own bundle is worse than a bundle which another agent *had the right to choose* (but did not). Such agents may demand "equal rights" (Gaspart, 1995). Unfortunately, once we leave the classical exchange economy, Sen's "Paretian Liberal" paradox suggests that rights are incompatible with efficiency (Deb, Pattanaik and Razzolini, 1997). Sen originally considered rights embodied in SCRs rather than mechanisms. Peleg and Winter (2002) study *constitutional implementation* where the mechanism embodies the same rights as the SCR it implements.

5 Learning from Experiments

Cabrales, Charness and Corchón (2003) tested the so-called canonical mechanism for Nash implementation. A Nash equilibrium was played only 13% of the time (20% when monetary fines were used). Remarkably, the optimal outcome was implemented 68% of the time (80% with "fines"), because deviations from equilibrium strategies frequently did not affect the outcome. This suggests that a desirable property of a mechanism is *fault-tolerance*: it should produce optimal outcomes even if some "faulty" players deviate from the theoretical predictions. Eliaz (2002) showed that if at most $k < \frac{1}{2}n - 1$ players are "faulty" (i.e., unpredictable), then full Nash implementation is possible if *no-veto-power* and $(k + 1)$-*monotonicity* hold.

Equilibrium play can be justified by epistemic or dynamic theories. According to epistemic theories, common knowledge about various aspects of the game implies equilibrium play even in one-shot games. Experiments provide little support for this. However, there is evidence that players can reach equilibrium through a dynamic adjustment process. If a game is played repeatedly, with no player knowing any other player's payoff function, the outcome frequently converges to a Nash equilibrium of the one-shot complete information game (Smith, 1979). Dynamic theories have been applied to the mechanism design problem (e.g., Cabrales and Ponti, 2000). Chen and Tang (1998) and Chen and Gazzale (2004) argue that mechanisms which induce supermodular games produce good long-run outcomes. Unfortunately, these convergence results are irrelevant for decisions that are taken infrequently, or if the principal is too impatient to care only about the long-run outcome.

The idea of dominant strategies is less controversial than Nash equilibrium, and should be more relevant for decisions that are taken infrequently. Unfortunately, experiments on dominant strategy mechanisms have yielded negative results. Attiyeh, Franciosi and Isaac (2000) conclude pessimistically, "we do not believe that the pivot mechanism warrants further practical consideration... This is due to the fundamental failure of the mechanism, in our laboratory experiments, to induce truthful value revelation". However, VCG mechanisms (such as the pivotal mechanism) frequently have a multiplicity of Nash equilibria, some of which produce suboptimal outcomes. Cason et. al. (2006) did experiments with *secure* mechanisms, which fully implement an SCR both in dominant strategies and in Nash equilibria. The players were much more likely to use their dominant strategies in secure than in non-secure mechanisms. In the non-secure mechanisms, deviations from dominant strategies tended to correspond to Nash equilibria. However, these deviations typically did not lead to suboptimal *outcomes*. In this sense, the non-secure mechanisms were fault-tolerant. Kawagoe and Mori (2001) report experiments where deviations from dominant strategies typically corresponded to suboptimal Nash equilibria.

In experiments, subjects often violate standard axioms of rational decision making. Alternative theories, such as prospect theory, fit the experimental evidence better. But if we modify the axioms of individual behavior, the optimal mechanisms will change. Esteban and Miyagawa (2005) assume the agents have Gul-Pesendorfer preferences. They suffer from "temptation", and may prefer a smaller menu (choice set) to a larger one. Suppose each agent first chooses a menu, and then chooses an alternative from this menu. Optimal menus may contain "tempting" alternatives which are never chosen in equilibrium, because this relaxes the incentive-compatibility constraints pertaining to the choice of menu. Eliaz and Spiegler (2005) assume some agents are "sophisticated" and some are "naive". Sophisticated agents know that they are dynamically inconsistent, and would like to commit to a future decision. Naive agents are unaware that they are dynamically inconsistent. The optimal mechanism screens the agents by providing commitment devices that are chosen only by sophisticated agents.

Experiments reveal the importance of human emotions such as spite or kindness (Andreoni, 1995; Saijo, 2003). In many mechanisms in the theoretical literature, by changing his strategy an agent can have a large impact on another agent's payoff without materially changing his own. Such mechanisms may have little hope of practical success if agents are inclined to manipulate each others' payoffs due to feelings of spite or kindness.

References

[1] Abreu, D. and Sen, A. 1991. Virtual implementation in Nash equilibria. *Econometrica* 59, 997-1022.

[2] Andreoni, J. 1995. Cooperation in public goods experiments: kindness or confusion? *American Economic Review* 85, 891-904.

[3] Armstrong, M. 1996. Multiproduct nonlinear pricing. *Econometrica* 64, 51-75.

[4] Attiyeh, G., Franciosi, R. and Isaac, R.M. 2000. Experiments with the pivotal process for providing public goods. *Public Choice* 102, 95-114.

[5] Baliga, S., Corchón, L. and Sjöström, T. 1997. The theory of implementation when the planner is a player. *Journal of Economic Theory* 77, 15-33.

[6] Baliga, S. and Sjöström, T. 1998. Decentralization and collusion. *Journal of Economic Theory* 83, 196-232.

[7] Baliga, S. and Sjöström, T. 2006. Contracting with third parties. Unpublished paper, Northwestern University.

[8] Bergemann, D. and Morris, S. 2005a. Robust mechanism design. *Econometrica* 73, 1771-1813.

[9] Bergemann, D. and Morris, S. 2005b. Robust implementation: the role of large type spaces. Unpublished paper, Yale University.

[10] Bergemann, D. and Välimäki, J. 2002. Information acquisition and efficient mechanism design. *Econometrica* 70, 1007–1034.

[11] Cabrales, A., Charness, G. and Corchón, L. 2003. An experiment on Nash implementation. *Journal of Economic Behavior and Organization* 51, 161-193.

[12] Cabrales, A. and Ponti, G. 2000. Implementation, elimination of weakly dominated strategies and evolutionary dynamics. *Review of Economic Dynamics* 3, 247-282.

[13] Cason, T., Saijo, T., Sjöström, T. and Yamato, T. 2006. Secure implementation experiments: do strategy-proof mechanisms really work? *Games and Economic Behavior*, forthcoming.

[14] Che, Y. K. and Hausch, D. 1999. Cooperative investments and the value of contracting. *American Economic Review* 89, 125-147.

[15] Che, Y.K. and Kim, J. 2005. Robustly collusion-proof implementation. *Econometrica*, forthcoming.

[16] Chen, Y. and Gazzale, R. 2004. When does learning in games generate convergence to Nash equilibria? *American Economic Review* 94, 1505-1535.

[17] Chen, Y. and Tang, F.F. 1998. Learning and incentive compatible mechanisms for public goods provision: an experimental study. *Journal of Political Economy* 106, 633-662.

[18] Choi, J. and Kim, T. 1999. A nonparametric, efficient public decision mechanism: undominated Bayesian Nash implementation. *Games and Economic Behavior* 27, 64-85.

[19] Chung, K. and Ely, J. 2003. Implementation with near-complete information. *Econometrica* 71, 857-871.

[20] Crémer, J. 1995. Arm's length relationships. *Quarterly Journal of Economics* 110, 275-295.

[21] Crémer, J. and McLean, R. 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56, 1247-1257.

[22] D'Aspremont, C. and Gérard-Varet, L.A. 1979. Incentives and incomplete information. *Journal of Public Economics* 11, 25-45.

[23] Deb, R., Pattanaik, P. and Razzolini, L. 1997. Game forms, rights and the efficiency of social outcomes. *Journal of Economic Theory* 72, 74-95

[24] Duggan, J. 1997. Virtual Bayesian implementation. *Econometrica* 67, 1175-1199.

[25] Dewatripont, M. 1989. Renegotiation and information revelation over time: the case of optimal labor contracts, *Quarterly Journal of Economics* 104, 589-619.

[26] Dewatripont, M. and Maskin, E. 1995a. Credit and efficiency in centralized and decentralized economies. *Review of Economic Studies* 62, 541-555.

[27] Dewatripont, M. and Maskin, E. 1995b. Contractual contingencies and renegotiation. *RAND Journal of Economics*, 26, 704-19.

[28] Edlin, A. and Reichelstein, S. 1996. Hold-ups, standard breach remedies and optimal investment. *American Economic Review* 86, 478-501.

[29] Eliaz, K. 2002. Fault tolerant implementation. *Review of Economic Studies* 69, 589-610.

[30] Eliaz, K. and Spiegler, R. 2005. Contracting with diversely naive agents. Unpublished paper, New York University.

[31] Esteban, S. and Miyagawa, E. 2005. Optimal menu of menus with self-control preferences. Unpublished paper, Penn State University.

[32] Freixas, X., Guesnerie, R. and Tirole, J. 1985. Planning under incomplete information and the ratchet effect. *Review of Economic Studies* 52 173-192.

[33] Gaspart, F. 1995. Fair implementation in the cooperative production problem: two properties of normal form mechanisms. Unpublished paper, FUNDP Namur.

[34] Hart, O. and Moore, J. 1999. Foundations of incomplete contracts. *Review of Economic Studies* 66, 115-138.

[35] Hurwicz, L. and Reiter, S. 2006. *Designing Economic Mechanisms*. Cambridge University Press.

[36] Jackson, M. 1991. Bayesian implementation. *Econometrica* 59, 461-477.

[37] Jackson, M., Palfrey, T. and Srivastava, S. 1994. Undominated Nash implementation in bounded mechanisms. *Games and Economic Behavior* 6, 474-501

[38] Jehiel, P. and Moldovanu, B. 2001. Efficient design with interdependent valuations. *Econometrica* 69, 1237-1259.

[39] Kawagoe, T. and Mori, T. 2001. Can the pivotal mechanism induce truthtelling? *Public Choice* 108, 331-354.

[40] Krishna. V. and Perry, M. 1997. Efficient mechanism design. Unpublished paper, Penn State University.

[41] Laffont, J.J. and Martimort, D. 1997. Collusion under asymmetric information. *Econometrica* 65, 875-911.

[42] Maskin, E. 1992. Auctions and privatization. In *Privatization*, ed. H. Siebert, Tubingen: Mohr (Siebeck).

[43] Maskin, E. 1999. Nash equilibrium and welfare optimality. *Review of Economic Studies* 66, 23-38

[44] Maskin, E. and Moore, J. 1999, Implementation and renegotiation. *Review of Economic Studies* 66, 39-56.

[45] Maskin, E. and Tirole, J. 1999. Two remarks on the property rights literature. *Review of Economic Studies* 66, 139-150.

[46] McAfee, P. and Reny, P. 1992. Correlated information and mechanism design. *Econometrica* 60, 395-421.

[47] McLean, R. and Postlewaite, A. 2002. Informational size and incentive compatibility. *Econometrica* 70, 2421-2453.

[48] Mezzetti, C. 2004. Mechanism design with interdependent valuations: efficiency. *Econometrica* 72, 1617-1626.

[49] Mezzetti, C. 2004. Mechanism design with interdependent valuations: surplus extraction. *Economic Theory*, forthcoming.

[50] Mookherjee, D. and Reichelstein, S. 1992. Dominant strategy implementation of Bayesian incentive compatible allocation rules. *Journal of Economic Theory* 56, 378-399.

[51] Mookherjee, D. and Tsumagari, M. 2004. The organization of supplier networks: effects of delegation and intermediation. *Econometrica* 72, 1179-1219.

[52] Myerson, R. 2006. Mechanism design. In *The New Palgrave Dictionary of Economics*, Second Edition, New York: Norton

[53] Palfrey, T. and Srivastava, S. 1989a. Implementation with incomplete information in exchange economies. *Econometrica* 57, 115-134.

[54] Palfrey, T. and Srivastava, S. 1989b. Mechanism design with incomplete information: a solution to the implementation problem. *Journal of Political Economy* 97, 668-691.

[55] Palfrey, T. and Srivastava, S. 1991. Nash implementation using undominated strategies. *Econometrica* 59, 479-501

[56] Peleg, B. and Winter, E. 2002. Constitutional implementation. *Review of Economic Design* 7, 187-204.

[57] Persico, N. 2004. Committee design with endogenous information. *Review of Economic Studies* 71, 165-191.

[58] Postlewaite, A. and Schmeidler, D. 1986. Implementation in differential information economies. *Journal of Economic Theory* 39, 14-33

[59] Nöldeke, G. and Schmidt, K. 1995. Option contracts and renegotiation: a solution to the hold-up problem. *RAND Journal of Economics* 26, 163-79.

[60] Saijo, T. 1987. On constant Maskin monotonic social choice functions, *Journal of Economic Theory* 42, 382-386.

[61] Saijo, T. 2003. Spiteful behavior in voluntary contribution mechanism experiments. Forthcoming in *Handbook of Experimental Economics Results*, ed. C. Plott and V. Smith, Elsevier Science.

[62] Segal, I. 1999. Complexity and renegotiation: a foundation for incomplete contracts. *Review of Economic Studies* 66, 57-82.

[63] Segal, I. 2004. The communication requirements of social choice rules and supporting budget sets. Unpublished paper, Stanford.

[64] Segal, I. and Whinston, M. 2002. The Mirrlees approach to mechanism design with renegotiation. *Econometrica* 70, 1-46.

[65] Serrano, R. and Vohra, R. 2001. Some limitations of virtual Bayesian implementation. *Econometrica* 69, 785-792.

[66] Sjöström, T. 1994. Implementation in undominated Nash equilibria without using integer games. *Games and Economic Behavior* 6, 502-511.

[67] Sjöström, T. 1999. Undominated Nash implementation with collusion and renegotiation. *Games and Economic Behavior* 26, 337-352.

[68] Smith, V. 1979. Incentive compatible experimental processes for the provision of public goods. In *Research in Experimental Economics*, Vol. I, ed. V. Smith, Greenwich, Conn.: JAI Press

[69] Tirole, J. 1986. Hierarchies and Bureaucracies. *Journal of Law, Economics and Organization* 2, 181-214.

[70] Williams, S. 1999. A characterization of efficient, Bayesian incentive compatible mechanisms. *Economic Theory* 14, 155-180.