Cowles Foundation

for Research in Economics
at Yale University

Cowles Foundation Discussion Paper No. 1846

EFFICIENT AUCTIONS AND INTERDEPENDENT TYPES

Dirk Bergemann, Stephen Morris, and Satoru Takahashi

January 2012

Efficient Auctions and Interdependent Types[*]

Dirk Bergemann[†]   Stephen Morris[‡]   Satoru Takahashi[§]

January 2012

Abstract

We consider the efficient allocation of a single good with interdependent values in a quasi-linear environment. We present an approach to modelling interdependent preferences distinguishing between "payoff types" and "belief types" and report a characterization of when the efficient allocation can be partially Bayesian implemented on a finite type space. The characterization can be used to unify a number of sufficient conditions for efficient partial implementation in this classical auction setting.

We report how a canonical language for discussing interdependent types - developed in a more general setting by Bergemann, Morris and Takahashi (2011) - applies in this setting and note by example that this canonical language will not allow us to distinguish some types in the payoff type - belief type language.

KEYWORDS: Mechanism Design, Robust Mechanism Design, Efficient Auctions, Interdepedent Types, Partial Implementation, Full Implementation.

JEL CLASSIFICATION: C79, D82

1

1   Introduction

Before the work of Harsanyi (1967/68) economists used to routinely argue that game theory could not be applied to economic settings because it required common knowledge of the environment. Following Harsanyi (1967/68), economists accepted that rich enough "type spaces" allowed any possible lack of common knowledge to be incorporated. But very rich type spaces would be needed and applied work remains highly sensitive to sometimes unexamined modelling choices about types. Nowhere is this more true than in mechanism design.

Two of us have written a series of papers highlighting the importance of lack of common knowledge and rich type spaces in mechanism design, now to be collected together in a book (Bergemann and Morris (2012)). Abreu and Matsushima (1992) showed that types that cannot be distinguished in this language may always pool in some equilibrium and thus the efficient allocation can only be fully implemented if it is measurable with respect to statements that can be expressed in this language.

2   Partial Implementation

There are I agents. Each agent i has a "payoff type" θ_i belonging to a finite set Θ_i. Each agent i's monetary valuation of a good depends on the profile of payoff types $\theta = (\theta_1, \ldots, \theta_I)$ and is given by a valuation function $v_i : \Theta_1 \times \cdots \times \Theta_I \to \mathbb{R}_+$. Agents are assumed to know their own payoff type θ_i but we want to allow for rich beliefs and higher order beliefs. Thus we assume each agent has a "type" t_i belonging to a finite set T_i, and write $\widehat{\theta}_i : T_i \to \Theta_i$ for a function describing an agent's payoff type and $\widehat{\pi}_i : T_i \to \Delta(T_{-i})$ for his belief type. This decomposition of an agent's type into a payoff type and a belief type is a natural one in this

2

quasi-linear setting. Two of us used this implicit description of interdependent types in our work on robust mechanism design (Bergemann and Morris (2012)). We will discuss briefly below a sense in which this decomposition can be seen as without loss of generality.

We will focus on the problem of allocating the object to the agent with the highest valuation. Write $i^*(\theta)$ for the agent who values the object the most, i.e., $i^*(\theta) \in \arg\max_i v_i(\theta)$, and in this finite type setting, we can assume without loss of generality that $i^*(\theta)$ is uniquely defined. We are interested in designing a finite mechanism involving monetary transfers which has an equilibrium where the object is always allocated to the agent who values it the most. By the revelation principle, we can restrict attention to direct mechanisms. In this context, a direct mechanism consists of a rule specifying monetary transfers to all agents, $y : T \to \mathbb{R}^I$. Now if agents report themselves to be of the type profile $t = (t_1, \ldots, t_I)$, the object will be allocated to agent $i^*\left(\widehat{\theta}(t)\right)$, where we write $\widehat{\theta}(t) = \left(\widehat{\theta}_1(t_1), \ldots, \widehat{\theta}_I(t_I)\right)$, and each agent i will receive a monetary transfer $y_i(t)$. This direct mechanism will be incentive compatible if

$$t_i \in \arg\max_{t_i' \in T_i} \sum_{t_{-i} \in T_{-i}} \left(\mathbb{I}_i\left(\widehat{\theta}(t_i', t_{-i})\right) v_i\left(\widehat{\theta}(t_i, t_{-i})\right) + y_i(t_i', t_{-i}) \right) \widehat{\pi}_i(t_{-i}|t_i), \qquad (1)$$

where the indicator function $\mathbb{I}_i(\cdot)$ of agent i is one if $i = i^*\left(\widehat{\theta}(t)\right)$, and zero otherwise. Thus we say efficient partial implementation is possible if and only if there exists $y : T \to \mathbb{R}^I$ such that (1) is satisfied.

To give a characterization of the incentive compatibility of the efficient social choice function i^*, or short, the efficient partial implementation, it is useful to first identify key properties of players' incentives to report their payoff types for beliefs about others' payoff types. Fix an agent i and fix his beliefs about the payoff types of others, $\psi_i \in \Delta(\Theta_{-i})$. Suppose that agent i expects the object to be allocated according to i^* based on the truthful reports of other

agents and his own (true or false) report. By extension of the valuation function v_i, we write $V_i\left(\theta_i, \theta_i', \psi_i\right)$ for the expected utility gain of agent i from a misreport θ_i' relative to his true payoff type θ_i, so that $V_i\left(\theta_i, \theta_i', \psi_i\right) \triangleq \sum_{\theta_{-i}}\left(\mathbb{I}_i\left(\theta_i', \theta_{-i}\right) - \mathbb{I}_i\left(\theta_i, \theta_{-i}\right)\right)\psi_i\left(\theta_{-i}\right) v_i\left(\theta_i, \theta_{-i}\right)$. Fix a subset $\overline{\Theta}_i \subseteq \Theta_i$; say that V_i is cyclically monotonic on $\left(\overline{\Theta}_i, \psi_i\right)$ if, for every sequence of types $\left(\theta_i^1, \ldots, \theta_i^K\right)$ in $\overline{\Theta}_i$,

$$V_i\left(\theta_i^1, \theta_i^2, \psi_i\right) + V_i\left(\theta_i^2, \theta_i^3, \psi_i\right) + \cdots + V_i\left(\theta_i^K, \theta_i^1, \psi_i\right) \leq 0.$$

Now, Theorem 1 in Rochet (1987) shows that there exists $\overline{y}_i : \overline{\Theta}_i \to \mathbb{R}$ such that $V_i\left(\theta_i, \theta_i, \psi_i\right) + \overline{y}_i\left(\theta_i\right) \geq V_i\left(\theta_i, \theta_i', \psi_i\right) + \overline{y}_i\left(\theta_i'\right)$ for all $\theta_i, \theta_i' \in \Theta_i$ if and only if V_i is cyclically monotonic on $\left(\overline{\Theta}_i, \psi_i\right)$. Thus cyclic monotonicity tells us that - ignoring belief types beyond the induced beliefs ψ_i over the payoff types - it is possible to choose transfers that give an agent an incentive to report his payoff type truthfully. To state this precisely, it will be useful to introduce the mapping $\widehat{\psi}_i : \Delta\left(T_{-i}\right) \to \Delta\left(\Theta_{-i}\right)$ which describes the beliefs over payoff types induced by the agents' belief types, so that $\widehat{\psi}_i\left(\pi_i\right)\left[\theta_{-i}\right] \triangleq \pi_i\left(\left\{t_{-i} : \widehat{\theta}_{-i}\left(t_{-i}\right) = \theta_{-i}\right\}\right)$ for any $\pi_i \in \Delta\left(T_{-i}\right)$. Now the incentive compatibility condition (1) can be re-written as:

$$t_i \in \underset{t_i' \in T_i}{\arg\max} \, V_i\left(\widehat{\theta}_i\left(t_i\right), \widehat{\theta}_i\left(t_i'\right), \widehat{\psi}_i\left(\widehat{\pi}_i\left(t_i\right)\right)\right) + \sum_{t_{-i} \in T_{-i}} y_i\left(t_i', t_{-i}\right)\widehat{\pi}_i\left(t_{-i}|t_i\right) \tag{2}$$

Now we consider what we can learn about agents' beliefs and how we can use what we learn. The classical environment we are considering - with quasi-linear utility and no limited liability constraints - is well-known to be very permissive in allowing "belief extraction". As observed by d'Aspremont and Gerard-Varet (1979), Myerson (1981) and Cremer and McLean (1985), it is possible to elicit agents' beliefs over other types by offering them gambles. Our general characterization of efficient partial implementation will essentially state that agents' beliefs can be elicited "for free" and what then matters is whether the set of payoff types consistent

4

with a given belief type can be distinguished. But this reduces to a cyclic monotonicity condition. Thus write $\widehat{\Theta}_i(\pi_i)$ for the set of payoff types of player i associated with belief type π_i, so that $\widehat{\Theta}_i(\pi_i) = \left\{\theta_i \in \Theta_i | \exists t_i \text{ with } \widehat{\pi}_i(t_i) = \pi_i \text{ and } \widehat{\theta}_i(t_i) = \theta_i\right\}$. Now we have:

Proposition 1 *Efficient partial implementation is possible if and only if each V_i satisfies cyclic monotonicity on $\left(\widehat{\Theta}_i(\pi_i), \widehat{\psi}_i(\pi_i)\right)$ for each π_i in the range of $\widehat{\pi}_i$.*

Proof. Suppose efficient partial implementation is possible and so (2) holds for some $y : T \to \mathbb{R}^I$. For any agent i, π_i in the range of $\widehat{\pi}_i$ and $\theta_i \in \widehat{\Theta}_i(\pi_i)$, write $\widehat{t}_i(\theta_i, \pi_i)$ for any type t_i with $\widehat{\pi}_i(t_i) = \pi_i$ and $\widehat{\theta}_i(t_i) = \theta_i$. Define $\overline{y}_i(\theta_i, \pi_i) = \sum_{t_{-i}} y_i\left(\widehat{t}_i(\theta_i, \pi_i), t_{-i}\right)\pi_i(t_{-i})$. By (2), $V_i\left(\theta_i, \theta_i, \widehat{\psi}_i(\pi_i)\right) + \overline{y}_i(\theta_i, \pi_i) \geq V_i\left(\theta_i, \theta_i', \widehat{\psi}_i(\pi_i)\right) + \overline{y}_i(\theta_i', \pi_i)$ for all $\theta_i, \theta_i' \in \widehat{\Theta}_i(\pi_i)$. Thus V_i satisfies cyclic monotonicity on $\left(\widehat{\Theta}_i(\pi_i), \widehat{\psi}_i(\pi_i)\right)$.

Suppose V_i satisfies cyclic monotonicity on $\left(\widehat{\Theta}_i(\pi_i), \widehat{\psi}_i(\pi_i)\right)$ for each π_i in the range of $\widehat{\pi}_i$. Thus there exists $\overline{y}_i(\cdot, \pi_i) : \widehat{\Theta}_i(\pi_i) \to \mathbb{R}$ such that $V_i\left(\theta_i, \theta_i, \widehat{\psi}_i(\pi_i)\right) + \overline{y}_i(\theta_i, \pi_i) \geq V_i\left(\theta_i, \theta_i', \widehat{\psi}_i(\pi_i)\right) + \overline{y}_i(\theta_i', \pi_i)$ for all $\theta_i, \theta_i' \in \widehat{\Theta}_i(\pi_i)$. Since $\sum_{t_{-i}}(\log \pi_i(t_{-i}) - \log \pi_i'(t_{-i}))\pi_i(t_{-i}) > 0$ for any $\pi_i' \neq \pi_i$, (2) holds for $y : T \to \mathbb{R}^I$ given by $y_i(t) = \overline{y}_i(\widehat{\theta}_i(t_i), \widehat{\pi}_i(t_i)) + K \log \widehat{\pi}_i(t_i)[t_{-i}]$ with sufficiently large K. ∎

We note that the present result of partial implementation is stated for the (ex post) efficient allocation rule $i^*(\theta)$. But in fact, the above result is more generally valid for every allocation rule that is measurable with respect to the payoff type profile θ. The only modification arises with respect to the utility gains $V_i(\cdot)$ from misreporting the payoff type, θ_i to θ_i', which have to be adapted to the specific allocation rule. In fact, a slightly more general version of this result was reported in Proposition 6.2 of Bergemann and Morris (2003), a working paper version of Bergemann and Morris (2005), where general allocation problems were considered.

In any case, the present result gives a sharp characterization of how payoff types and belief types matter for efficient partial implementation. In particular, differences in beliefs about others' types can be extracted for free, even if they are not payoff relevant. But once they are extracted, the non-payoff relevant content of the beliefs does not matter. All that matters is the implied belief over payoff types and the cyclic monotonicity condition given by that implied belief over the payoff types and the set of possible payoff types who could have had the original (perhaps payoff irrelevant) belief. Classic sufficient conditions for efficient partial implementation can now be seen as special cases of the above Proposition:

1. Private Values. Under private values (i.e., $v_i(\theta_i, \theta_{-i})$ does not depend on θ_{-i}), the condition becomes vacuous as the cyclic monotonicity conditions are satisfied.

2. Independent Types. If the agents' types are distributed independently ($\widehat{\pi}_i(t_i)$ does not depend on t_i), then necessary and sufficient conditions for efficient partial implementation reduce to the cyclic monotonicity conditions on the entire set Θ_i of payoff types of each agent i for his fixed beliefs over others' payoff types.

3. Linear Independence and Convex Hull. As noted by our discussant, Eric Maskin, for any set Ψ_i of linearly independent beliefs over others' payoff types, the cyclic monotonicity conditions on entire Θ_i for all beliefs in Ψ_i are sufficient for efficient partial implementation if there is common knowledge that beliefs lie in the convex hull of Ψ_i, which is less demanding than the notion of ex post equilibrium. In a similar spirit, Jehiel, Meyer-ter-Vehn and Moldovanu (2011) investigate conditions for local robust incentive compatibility, and their Lemma 1 establishes necessary monotonicity conditions for partial implementation, using the above separation of payoff types and belief types.

4. Belief Extraction. Following Neeman (2004), say that "beliefs determine preferences"

(BDP) if for every $\pi_i \in \Delta(T_{-i})$, there exists at most one $\theta_i \in \Theta_i$ such that $\widehat{\pi}_i(t_i) = \pi_i$ and $\widehat{\theta}_i(t_i) = \theta_i$ for some $t_i \in T_i$. Under BDP, the cyclic monotonicity requirements of the proposition become vacuous.[1]

5. Ex Post Incentive Compatibility. As in the work of Maskin (1992) and Dasgupta and Maskin (2000), suppose that a single crossing condition is satisfied with respect to the (ordered) payoff types of the agents, so that each v_i is strictly monotonic in θ_i and $\theta_i' > \theta_i \Rightarrow v_i(\theta_i', \theta_{-i}) - v_i(\theta_i, \theta_{-i}) > v_j(\theta_i', \theta_{-i}) - v_j(\theta_i, \theta_{-i})$ for all i, j with $i \neq j$. Then V_i satisfies cyclic monotonicity on (Θ_i, ψ_i) for every ψ_i (which implies ex post incentive compatibility) and thus again it can be shown that the cyclic monotonicity requirements of the proposition hold.[2]

6. Combining Belief Extraction and Private Values. McLean and Postlewaite (2004) analyze efficient auctions with interdependent values and multidimensional types. More precisely, each agent has a two-dimensional payoff type $\theta_i = (\theta_i^i, \theta_i^c)$, with an idiosyncratic and a common component, θ_i^i and θ_i^c respectively, as reflected by the valuation function $v_i(\theta) = v_i^i(\theta_i^i) + v_i^c(\theta^c)$. The common part of the valuation, $v_i^c(\theta^c)$, depends on the entire profile of common components $\theta^c = (\theta_1^c, \ldots, \theta_I^c)$. Now, in terms of our language, their environment in Section 4 can be described by a type space $T_i = \Theta_i$, the payoff types can be described by the identity mapping $\widehat{\theta}_i$, and the belief types can be described by $\pi_i[\theta_i]\left((\theta_j)_{j \neq i}\right) =$

---

[1]Arguments going back at least to Cremer and McLean (1985) say BDP holds generically on finite type spaces, since if we fix a finite set of types and perturb beliefs, then they will all be different and BDP will hold. Recent contributions, notably Heifetz and Neeman (2006) and Chen and Xiong (2011), examine when BDP holds on infinite type spaces, and in particular if it can be said to hold "generically" on infinite type spaces.

[2]Such sufficient conditions do not naturally arise in multidimensional payoff type space Θ_i, expect for some special cases, as noted by Maskin (1992) and shown in much more general environments by Jehiel et al. (2006).

$\pi_i^c [\theta_i^c] \left(\left(\theta_j^c \right)_{j \neq i} \right) \cdot \pi_i^i \left(\left(\theta_j^j \right)_{j \neq i} \right)$. Finally, their condition of positive informational variability implies that the common component θ_i^c can be extracted from every agent i, and the residual private information θ_i^i is independent and pertains to the private value, and hence again, our necessary and sufficient conditions are satisfied.

3 Interdependent Types and Full Implementation

Importantly, Proposition 1 only establishes partial implementation, i.e. that the efficient allocation arises in some equilibrium but there may exist other equilibria with inefficient allocations. Suppose there are three agents, $I = 3$, and each agent has two possible types, $\Theta_1 = \Theta_2 = \Theta_3 = \{0, 1\}$. Each agent's valuation of the object is given by

$$v_i \left(\theta \right) = \theta_i + \frac{2}{3} \sum_{j \neq i} \theta_j. \tag{3}$$

Each agent has only two types, and thus a single belief type for each payoff type. Suppose, in particular, that if a type has valuation θ_i, he assigns independent probability $\frac{1}{8}$ to each of the other agents having valuation $\theta_j = \theta_i$, and the remaining probability $\frac{7}{8}$ to $\theta_j = 1 - \theta_i$.

Efficient partial implementation is possible in this example (by both a belief extraction argument, i.e., sufficient condition 3 above, or an ex post incentive compatibility argument, e.g., sufficient condition 4 above). But observe that each type of each agent has an expected value of $\frac{7}{6}$ for the object. To see why, note that if agent 1 say has $\theta_1 = 1$, his expectation of θ_2 (or θ_3) is $\frac{1}{8}$ and thus his expectation of $\theta_1 + \frac{2}{3}\theta_2 + \frac{2}{3}\theta_3 = 1 + \frac{1}{12} + \frac{1}{12} = \frac{7}{6}$. But if agent 1 has $\theta_1 = 0$, his expectation of θ_2 (or θ_3) is $\frac{7}{8}$ and thus his expectation of $\theta_1 + \frac{2}{3}\theta_2 + \frac{2}{3}\theta_3 = 0 + \frac{7}{12} + \frac{7}{12} = \frac{7}{6}$. Thus, there will always be an equilibrium where both types of each player behave the same, which cannot give rise to the efficient allocation, as shown in a general payoff environment in

Bergemann and Morris (2009).

Following Bergemann, Morris and Takahashi (2011), we propose a natural description of agents' types in this setting. An agent's (unconditional) willingness to pay for the object does not depend on other agents' types and thus has a natural meaning. If it makes sense to speak of all agents' willingness to pay for the object, it makes sense to talk about an agent's beliefs about other agents' willingness to pay for the object, and his willingness to pay conditional on others' willingness to pay for the object. Call the agent's unconditional willingness to pay his first order type. Call his beliefs over others' willingness to pay and conditional willingness to pay his second order type. Now we can inductively define an agent's $(n+1)$th order type to be his beliefs over others' nth order types and his willingness to pay conditional on others' nth order types. We can thus identify an agent with a hierarchy of statements about preferences and conditional preferences.

In the simple setting of this note, such hierarchies can be defined formally as follows. An nth level type is pair $t_n = (b_n, v_n)$ consisting of a belief component and valuation component. Write T_n for the set of nth level types that can arise in a finite type space. We describe the sets T_n inductively. Let $T_1 = \{\varnothing\} \times \mathbb{R}_+$ with a typical 1st level type $t_1 = (\varnothing, v_1)$ consisting of a degenerate belief type and an unconditional valuation of the object. Now an $(n+1)$th level type $t_{n+1} = (b_{n+1}, v_{n+1})$ consists of a simple (i.e., finite support) probability distribution $b_{n+1} \in \Delta(T_n)$ with support $\text{supp}(b_{n+1})$ and a valuation function $v_{n+1} : \text{supp}(b_{n+1}) \to \mathbb{R}_+$. Now a hierarchy of types is a sequence of nth order types $(t_1, t_2, \ldots, t_n, \ldots) \in T_1 \times T_2 \times \cdots \times T_n \times \cdots$. A sequence of types is coherent if each $(n+1)$th type t_{n+1} induces beliefs over other agents' $(n-1)$th level types and willingness to pay conditional on other agents' $(n-1)$th level types that are consistent with those of t_n. (We omit the formal statement of this condition). Now if

we write T_f^* for the set of all coherent hierarchies of higher order types that can arise in finite type spaces, we have a natural language in which to discuss agents' types.

In Bergemann, Morris and Takahashi (2011), we discuss the extension to infinite types and construct a universal space of higher order preferences T^* from hierarchies. We can identify those hierarchies with beliefs over others' types and valuations conditional on their hierarchies. Thus we identify a type $t^* \in T^*$ with a probability measure over the types of others $b^* \in \Delta\left((T^*)^{I-1}\right)$ and an equivalence class of b^*-integrable valuation functions $v^* : (T^*)^{I-1} \to \mathbb{R}$ (where we identify two functions if they agree b^*-almost surely). In this sense, we have a canonical way of representing type spaces with a decomposition of types into "payoff types" and "belief types," as we did earlier. However, it is important to realize that these payoff types does not specify valuations on 0-probability events.

This language is closely related to full implementation. Abreu and Matsushima (1992) showed that a necessary condition for Bayesian full implementation of a social choice function using a finite mechanism on a finite type space was that two types which had the same hierarchies of preferences received the same allocation. While they expressed this "measurability" condition as a property of the fixed finite type space, it can be expressed without reference to a particular finite type space as we described above. Thus, if full implementation is required, we can at most achieve constrained efficiency, with the object allocated to the agent $i^{**}(t_1^*, \ldots, t_I^*)$ that maximizes $v_i^*[t_i^*](t_{-i}^*)$, where $v_i^*[t_i^*]$ denotes the payoff type component of t_i^*, and Proposition 1 must hold on the coarser type space where types with the same hierarchy of preferences are merged.

Exact full implementation is generally not possible in the allocation of a single good: even with private values, while bidding one's true value is a dominant strategy, it is often only a

weak best response and it is easy to construct inefficient equilibria in dominated strategies. For sufficient conditions for full implementation with finite mechanisms, the objective must be weakened to virtual implementation, so that $i^{**}(t_1^*, \ldots, t_I^*)$ is allocated the object with arbitrarily high probability. Abreu and Matsushima (1992) showed that virtual Bayesian full implementation is possible under Bayesian incentive compatibility and their measurability condition. In this context, this implies that constrained efficient allocation of the good is possible if the condition of Proposition 1 holds on the coarsened type space that is expressible in our language.[3]

---

[3]In Bergemann and Morris (2009), we examined the efficient allocation rule (with the single unit auction as a special case) and gave conditions under which full implementation can be achieved by a simpler mechanism, namely the direct mechanism, than the mechanism of Abreu and Matsushima (1992). It is an open question when a simpler mechanism can achieve full implemetation more generally.

References

[1] Abreu, D. and H. Matsushima (1992). "Virtual Implementation in Iteratively Undominated Strategies: Incomplete Information." mimeo.

[2] d'Aspremont, C. and L.-A. Gerard-Varet (1979). "Incentives and Incomplete Information," *Journal of Public Economics* 11, 25-45.

[3] Bergemann, D., and S. Morris (2003). "Robust Mechanism Design," Cowles Foundation Discussion Paper #1421.

[4] — (2005). "Robust Mechanism Design," *Econometrica* 73, 1771-1813.

[5] — (2009). "Robust Implementation in Direct Mechanisms" *Review of Economic Studies,* 76, 1175-1204.

[6] — (2011). "Robust Mechanism Design: An Introduction," Cowles Foundation Discussion Paper #1818.

[7] — (2012). *Robust Mechanism Design: The Role of Private Information and Higher Order Beliefs.* World Scientific Publishing, Singapore.

[8] Bergemann, D., S. Morris and S. Takahashi (2011). "Interdependent Preferences and Strategic Distinguishability," Cowles Foundation Discussion Paper #1772R.

[9] Chen, Y.-C., and S. Xiong (2011). "Genericity and the Robustness of Full Surplus Extraction," forthcoming in *Econometrica.*

[10] Cremer, J. and R. McLean (1985). "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent," *Econometrica* 53, 345-361.

[11] Dasgupta, P. and E. Maskin (2000). "Efficient Auctions," *Quarterly Journal of Economics* 115, 341-388.

[12] Harsanyi, J. (1967-68). "Games with incomplete information player by Bayesian players, I, II and III," *Management Science* 14, 159-182, 320-334 and 486-502.

[13] Heifetz, A. and Z. Neeman (2006). "On the Generic (Im)possibility of Full Surplus Extraction in Mechanism Design," *Econometrica* 74, 213-294.

[14] Jehiel, P., B. Moldovanu, and M. Meyer-ter-Vehn (2012). "Locally Robust Implementation and Its Limits," mimeo.

[15] Jehiel, P., B. Moldovanu, M. Meyer-ter-Vehn and B. Zame (2006). "The Limits of Ex Post Implementation," *Econometrica* 74, 585-610.

[16] Maskin, E. (1992). "Auctions and Privatization," in Privatization: Symposium in Honor of Herbert Giersch, ed. by H. Siebert, pp. 115-136. J.C.B. Mohr, Tuebingen.

[17] Mertens, J.-F. and S. Zamir (1985). "Formulation of Bayesian analysis for game with incomplete information," *International Journal of Game Theory* 14, 1-29.

[18] McLean, R. and A. Postlewaite (2002). "Informational Size and Efficient Auctions," *Review of Economic Studies* 71, 809-827.

[19] Myerson, R. (1981). "Optimal Auction Design," *Mathematics of Operations Research* 6, 53-73.

[20] Neeman, Z. (2004). "The Relevance of Private Information in Mechanism Design," *Journal of Economic Theory* 117, 55-77.

[21] Rochet, J.-C. (1987). "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context," *Journal of Mathematical Economics* 16, 191-200.