Information in Mechanism Design

Dirk Bergemann[*]   Juuso Välimäki[†]

January 2006

Abstract

We survey the recent literature on the role of information in mechanism design. First, we discuss an emerging literature on the role of endogenous payoff and strategic information for the design and the efficiency of the mechanism. We specifically consider information management in the form of acquisition of new information or disclosure of existing information. Second, we argue that in the presence of endogenous information, the robustness of the mechanism to the type space and higher order beliefs becomes a natural desideratum. We discuss recent approaches to robust mechanism design and robust implementation.

[*]Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven, CT 06511, dirk.bergemann@yale.edu.

[†]Department of Economics, Helsinki School of Economics and University of Southampton, 00100 Helsinki, Finland, valimaki@hkkk.fi.

1 Introduction

The mechanism design literature of the last thirty years has been a big success on a number of different levels. A beautiful theoretical literature has shown how a wide range of institutional design questions can be formally posed as mechanism design problems with a common structure. We can understand institutions as solutions to a well defined maximization problems subject to incentive constraints. Elegant characterizations of optimal mechanisms have been obtained. Market design has become more important in many economic arenas, both because of new insights from theory and developments in information technology. A very successful econometric literature has tested auction theory in practise.

The basic issue in mechanism design is how to truthfully elicit private and decentralized information in order to achieve some private or social objective. The task of the principal is then to design a game of incomplete information in which the agents have indeed an incentive to reveal the information. The optimal design depends on the common prior which the principal and the agents share about the types of the agents. Unfortunately, the general theory, the applications and the empirical work have rather different natural starting points. The theoretical analysis begins with a given common prior, often over a small set of types, and then analyzes the optimal mechanism with respect to this prior. Yet, the fine details of the specified environment incorporated in the common prior are rarely available to the designer in practise.

In this survey, we shall pursue two distinct but closely related arguments. The first part of this survey is centered on the issue of endogenous information structures in mechanism design. In traditional mechanism design literature, the set of possible types for the participants in the design problem is exogenously given. This may be a reasonable approximation in situations such as determining Pareto efficient allocations in an exchange economy where individual preferences are private information. It is equally clear that for many applications it is not reasonable to assume that the relevant information is independent of the mechanism chosen.

To illustrate the point concretely, consider decision making in committees. If committee members have to invest privately in order to have useful information, then it is clear that their willingness to invest in such information depends on the choice of the decision making process. If additional information has little impact on the eventual decision, there is no point to acquiring it. As a second, slightly different application where the participants' information depends on the mechanism chosen, consider the optimal design of auctions. The auctioneer may have control over pieces of evidence that determine the bidders' valuation for the object on sale. Whether it is in the auctioneer's best interest to disclose this information depends on the properties of the auction to follow.

We view information acquisition and information disclosure as two different aspects of an infor-

mation management problem that we believe is important in many mechanism design settings. In our view, in many examples of great practical interest, it is not accurate to view the distribution of types as independent from the choice of the mechanism. At the most abstract level, we may think about mechanisms as institutions that coordinate societies on particular collective choices. As long as the relevant information is produced within the economies, it should be clear that this production is guided by economic incentives. Hence a good mechanism ought to provide incentives for efficient collective choices given the information collected, but at the same time a good mechanism should also provide the participants with good incentives for producing the relevant information.

We review the existing literature on information acquisition and disclosure in a number of applications. It is our intention to show that by adding an information acquisition stage a number of features arise across the spectrum of applications. Two aspects deserve mention here. With information acquisition, the extensive form, and in particular the timing, becomes more important than in standard mechanisms. Second, randomizations play a role in settings that can be analyzed using pure strategies in the basic model.

In the second part of this survey, we analyze mechanism design when the principal and the agents have little common knowledge and the type space is large. The starting point here is the influential formulation of the robustness question due to Robert Wilson. Wilson emphasized that academic mechanisms designers were tempted to assume too much common knowledge among the players, and suggested that more robust conclusions would arise if researchers were able to relax those assumptions. Practitioners have often been led to argue in favor of using simpler but apparently sub-optimal mechanisms. It is argued that the optimal mechanisms are not "robust" - that is they are too sensitive to fine details of the specified environment. In response to these concerns, attractive and influential results have been obtained by imposing (in a somewhat ad hoc way) stronger solution concepts and simpler mechanisms motivated by robustness considerations. A natural theoretical question to ask is whether it is possible to explicitly model robustness in such a way that stronger solution concepts and simpler mechanisms arise endogenously. To the extent that the agents have or can get access to private information about their own valuation, the valuations of other agents or the beliefs of the others, the designer is led to adopt a robust mechanism. Consequently, in this survey we shall study mechanism design when we relax both the *small* and the *given* type space assumptions.

The remainder of this survey is organized as follows. Section 2 provides the basic model and notation for the survey. Section 3 is meant to motivate and emphasize the perspective of this survey. We shall first discuss the role of information acquisition in generalized Vickrey Groves Clark mechanism and then talk about the role of strategic information in first price auctions. In Section 4 we survey the role of information management in mechanism design. Section 5 frames

the concern for robust mechanism by emphasizing the importance of strategic uncertainty. We discusses recent results on robust mechanisms and show how classic auction results are modified by the introduction of large type spaces expressing strategic uncertainty. Section 6 concludes the survey and discusses a number of open and note worthy research issues.

2 Setup

2.1 Payoff Environment

We consider a finite set of agents, indexed by $i \in \mathcal{I} = \{1, ..., I\}$. The agents have to make a collective choice y from a set Y of possible outcomes. The *payoff type* of agent i is $\theta_i \in \Theta_i$. We write $\theta \in \Theta = \Theta_1 \times \cdots \times \Theta_I$. Each agent has utility function $u_i : Y \times \Theta \to \mathbb{R}$. An important special case is the *quasi-linear environment* where the set of outcomes Y has the product structure $Y = Y_0 \times Y_1 \times \cdots \times Y_I$, where $Y_1 = Y_2 = ... = Y_I = \mathbb{R}$, and a utility function:

$$u_i(y, \theta) = u_i(y_0, y_1, ..., y_I, \theta) \triangleq v_i(y_0, \theta) + y_i,$$

which is linear in y_i for every agent i.

The collective choice problem is represented by a social choice correspondence $F : \Theta \to 2^Y \backslash \emptyset$, a social choice function is given by $f : \Theta \to Y$. If the true payoff type profile is θ, the planner would like the outcome to be an element of $F(\theta)$, or $f(\theta)$. This environment is fixed and informally understood to be common knowledge. We allow for interdependent types - one agent's payoff from a given outcome depends on other agents' payoff types. The model is said to be a *private value model* if for all θ, θ' :

$$\theta_i = \theta'_i \Rightarrow u_i(y, \theta) = u_i(y, \theta'). \tag{1}$$

If condition (1) is violated, then the model displays *interdependent values.*

The payoff type profile is understood to contain all information that is relevant to whether the planner achieves her objective or not. It incorporates many classic problems such as the efficient allocation of an object, the efficient provision of a public good, and arriving at a decision in a committee.

Much of the recent work on interdependent values has used the solution concept of ex post rather than Bayesian equilibrium. The analysis of ex post equilibrium is considerably more tractable, because incentive compatible transfers can often be derived with ease and single crossing conditions generating incentive compatibility are easy to identify.[1]

---

[1]Ex post incentive compatibility was discussed as "uniform incentive compatibility" by Holmstrom and Myerson (1983). Ex post equilibrium is increasingly studied in game theory (see Kalai (2004)) and is often used in mechanism

Definition 1 *A direct mechanism $f : \Theta \to Y$ is ex post incentive compatible if, for all i and $\theta \in \Theta$,*

$$u_i\left(f\left(\theta\right),\theta\right) \geq u_i\left(f\left(\theta'_i, \theta_{-i}\right), \theta\right), \quad \text{for all } \theta'_i \in \Theta_i.$$

The notion of ex post incentive compatibility requires agent i to prefer truthtelling at θ if all the other agents also report truthfully. In contrast the notion of dominant strategy implementation requires agent i to prefer truthtelling for all possible reports by the other agents, truthtelling or not.

Definition 2 *A direct mechanism $f : \Theta \to Y$ is incentive compatible in dominant strategies if, for all i and $\theta \in \Theta$,*

$$u_i\left(f\left(\theta_i, \theta'_{-i}\right), \theta\right) \geq u_i\left(f\left(\theta'\right), \theta\right), \quad \text{for all } \theta' \in \Theta.$$

If there are private values (i.e., each $u_i\left(y, \theta\right)$ depends on θ only through θ_i), then ex post incentive compatibility is equivalent to dominant strategies incentive compatibility.

2.2 Information Acquisition

In problems of choice under uncertainty, the starting point of the analysis is often the situation where an agent holds a prior probability distribution on the states of the world $\omega \in \Omega$ and must decide on an optimal action $y \in Y$. One way to model information acquisition is then to assume that the agent has access to a statistical experiment that yields additional information on ω. Each outcome in the experiment results in a posterior belief on Ω. Since the posterior belief represents the payoff relevant information of the agent, we denote (in accordance with the previous subsection) the set of probability distributions on Ω by Θ with a generic element $\theta \in \Theta$.

For the purposes of the current survey, it is easiest to formulate the information acquisition decision of the agent as a choice amongst a set of distributions on Θ. We index the experiments by $\alpha \in A$ and hence an experiment results in a distribution $F^\alpha\left(\theta\right)$ on Θ. We also write the utility function of the agent directly in terms of the posterior and the chosen action $u\left(y, \theta\right)$. Under suitable regularity conditions, there is an optimal action $y\left(\theta\right)$ for each θ. If we denote the cost of observing experiment α by $c\left(\alpha\right)$, the information acquisition problem can be written concisely as follows:

$$\max_{\alpha \in A}\left\{\int_\Theta u\left(y\left(\theta\right), \theta\right) dF^\alpha\left(\theta\right) - c\left(\alpha\right)\right\}.$$

design as a more robust solution concept (Cremer and McLean (1985)). A recent literature on interdependent value environments has obtained positive and negative results using this solution concept: Dasgupta and Maskin (2000), Bergemann and Välimäki (2002), Perry and Reny (2002), Jehiel and Moldovanu (2001), Jehiel, Moldovanu, Meyer-Ter-Vehn, and Zame (2005) and Bikhchandani (2005).

To see a concrete example that fits the framework above, consider the case where $\omega \in \{0, 1\}$. Then we may identify Θ with $[0, 1]$ where $\theta = \Pr\{\omega = 1\}$. Let θ_0 indicate the prior distribution of the agent and consider the following family of experiments:

$$F^{\alpha}(\theta) = \begin{cases} (1 - \theta_0)\,\alpha & \text{for} & \theta < \theta_0, \\ 1 - \theta_0\alpha & \text{for} & \theta_0 \leq \theta < 1, \\ 1 & \text{for} & \theta = 1. \end{cases}$$

Here α is the probability of observing a perfectly informative signal on ω. It is easy to generate richer examples of this structure.

When considering the mechanism design problem, all relevant information for the mechanism is contained in the vector of posteriors $(\theta_1, ..., \theta_I)$. It is thus possible to consider the posteriors directly as the inputs that the mechanism designer elicits from the participants in the mechanism. The choice of individual experiment α_i determines the appropriate distribution for the posteriors θ_i. Since these posteriors are in general multi-dimensional (and quite often infinite dimensional), it is clear that unless further assumptions on the payoff structures are made, the task of designing mechanisms in such settings is very complicated. We shall consider throughout the case where the ex ante investment in information is covert. As a result, the mechanism cannot be written as directly depending on α_i.

3 Motivating Examples

3.1 Information Acquisition in Generalized VCG auctions

Our first example examines the role of information acquisition in a single unit auction with interdependent values. More specifically, we are interested in the possibility of inducing the bidders to gather information in a socially efficient manner.

The auction has two bidders, each of whom has statistically independent private information on a different binary aspect $\omega_i \in \{\underline{\omega}_i, \overline{\omega}_i\} = \{0, 1\}$ of the good. We denote by θ_i bidder $i's$ probability assessment on the event $\{\omega_i = \overline{\omega}_i\}$. We assume that player $i's$ payoff from obtaining the object at price y_i takes the following linear form:

$$u_i(\theta) = \alpha\theta_i + \beta\theta_j - y_i, \tag{2}$$

where we assume that $\alpha > 0$. If $\beta = 0$, we are in the private values case. When $\alpha = \beta$, we have a model with pure common values.

Denote the allocation of the object in the auction by $y_0 \in \{1, 2\}$. Efficiency requires that

$$y_0\left(\theta_i,\theta_j\right) = i \quad \text{if} \quad \left(\alpha - \beta\right)\left(\theta_i - \theta_j\right) > 0.$$

Hence a necessary condition for incentive compatibility of the efficient allocation is that $\alpha \geq \beta$. Under this condition, it is easy to verify that the direct mechanism consisting of

$$y_i\left(\theta_i,\theta_j\right) = \begin{cases} \left(\alpha + \beta\right)\theta_j & \text{if} \quad \theta_i \geq \theta_j, \\ 0 & \text{if} \quad \theta_i < \theta_j, \end{cases}$$

and

$$y_0\left(\theta_i,\theta_j\right) = i \ \ \text{if } \theta_i \geq \theta_j,$$

is ex post incentive compatible. This mechanism is called the generalized *Vickrey-Clarke-Groves (VCG)* mechanism and its analysis in the interdependent values case is due to Maskin (1992) and Dasgupta and Maskin (2000).

With statistical independence of types, the revenue equivalence theorem implies that the expected payoffs of the two bidders in all efficient mechanisms coincide with the payoffs in the generalized VCG mechanism. As we are focusing here on socially efficient information acquisition, it is natural to ask whether an individual bidder's incentives to acquire additional information coincide with those of a utilitarian social planner.

Our main finding in Bergemann and Välimäki (2002) implies that when $\beta < 0$, the generalized VCG auction gives too low incentives for information acquisition to the individual bidders. If $\beta > 0$, the agents have an incentive to engage in excessive information acquisition.

To see the intuition for this result, notice that the generalized VCG mechanism allocates the object to i only if $\theta_i \geq \theta_j$. For $\theta_i \geq \theta_j$,

$$u_i\left(\theta_i,\theta_j\right) - u_i\left(\theta_j,\theta_j\right) = \max\{u_i\left(\theta_i,\theta_j\right), u_j\left(\theta_i,\theta_j\right)\} - u_i\left(\theta_j,\theta_j\right),$$

and hence the gains from higher θ_i are the same for bidder i and for the social planner. Bidder i's payoff is zero in the generalized VCG mechanism for all $\theta_i < \theta_j$. If $\beta > 0$ then the utilitarian planner's payoff is increasing also for $\theta_i < \theta_j$. Hence the payoff to bidder i has a sharper kink at θ_j than the planner's utility function. As a result, bidder i is locally more risk loving than the planner and hence she has stronger incentives to acquire information. It should be noted that when $\beta = 0$, bidder i's payoff equals the planner's payoff as a function of θ_i (up to a constant) and as a result, private incentives for information acquisition coincide with the planner's incentives in a private value environment.

This example shows how efficient use of information is often incompatible with efficient acquisition of information. It is clear that a second best mechanism would sacrifice some of the allocational

efficiency relative to the generalized VCG mechanism in order to achieve better alignment of private and social incentives in the information acquisition stage. Full exploration of this trade-off remains an open question at this time.

3.2 Strategic Information in First-Price Auctions

Our second example demonstrates the importance of modeling information about other players' types. In a setting with independent private value we consider a first price auction among two bidders for a single object. For simplicity, we consider a discrete space of values and bids. The valuations are given for each i by:

$$\theta_i \in \Theta_i = \{1, \quad 2, \quad 3\},$$

and the feasible bids for each i are given by:

$$b_i \in B_i = \{1/2, \quad 1, \quad 3/2, \quad 2, \quad 5/2, \quad 3\}.$$

The valuations are distributed uniformly and independently according to a common prior $p(\theta_i, \theta_j)$:

$$
\begin{array}{cccc}
 & t_j^1 & t_j^2 & t_j^3 \\
t_i^1 & \frac{1}{9} & \frac{1}{9} & \frac{1}{9} & \theta_i^1 \\
t_i^2 & \frac{1}{9} & \frac{1}{9} & \frac{1}{9} & \theta_i^2 \\
t_i^3 & \frac{1}{9} & \frac{1}{9} & \frac{1}{9} & \theta_i^3 \\
 & \theta_j^1 & \theta_j^2 & \theta_j^3
\end{array}
\tag{3}
$$

The private information of bidder i, her type t_i, consists of her true valuation, θ_i (payoff relevant type), and her belief about the valuations of the other bidders, the posterior distribution, $p(\theta_j | \theta_i)$. In the standard model of auctions, each payoff type θ_i is associated with exactly one belief type, and hence the additional notation of a type t_i may appear at first glance redundant.

We wish to consider a richer environment in which each bidder receives some additional private information about her competitor. To keep matters simple, let us suppose that every bidder with a high valuation, i.e. $\theta_i = 3$, obtains some additional information. This additional information, represented by two distinct types, t_i' and t_i'', refines her view about the strength of her competitor as follows:

	θ_j^1	θ_j^2	θ_j^3	$\Pr(t_i | \theta_i = 3)$	
$p(\theta_j | t_i')$	$\frac{3}{6}$	$\frac{2}{6}$	$\frac{1}{6}$	$\frac{2}{3}$	← "weak" competitor
$p(\theta_j | t_i'')$	0	$\frac{2}{6}$	$\frac{4}{6}$	$\frac{1}{3}$	← "strong" competitor

The last column in the above matrix represents the likelihood that a bidder with a high valuation receives either one of the two possible pieces of information, t_i' or t_i''. The posterior beliefs, $p\left(\theta_j \,|t_i\right)$ over the valuations θ_j of bidder j differ across the two types, t_i' or t_i'', indicating that bidder j is a weak or strong competitor, respectively. We also observe that the aggregate distribution over valuations θ_j given θ_i has not changed. The common prior on the new type space is now:

$$
\begin{array}{c c c c c c}
 & t_j^1 & t_j^2 & t_j^3 & t_j^4 & \\
t_i^1 & \frac{1}{9} & \frac{1}{9} & \frac{1}{9} & 0 & \theta_i^1 \\
t_i^2 & \frac{1}{9} & \frac{1}{9} & \frac{2}{27} & \frac{1}{27} & \theta_i^2 \\
t_i^3 & \frac{1}{9} & \frac{2}{27} & \frac{1}{27} & 0 & \theta_i^3 \\
t_i^4 & 0 & \frac{1}{27} & 0 & \frac{2}{27} & \theta_i^3 \\
 & \theta_j^1 & \theta_j^2 & \theta_j^3 & \theta_j^3 &
\end{array}
\tag{4}
$$

Observe that the distributions over valuations θ are identically, uniformly and independently distributed in both type spaces. Yet as we consider how the new information affects bidding in the first price auction, we observe a few important differences. In the small type space, the unique equilibrium bidding strategy, $b_i^*\left(t_i\right)$ is given by:

$$
b_i^*\left(t_i\right) = \frac{1}{2}\theta_i\left(t_i\right),
$$

which is also the bidding strategy in the continuous version of the model. However, in the larger type space, the bidding strategy changes as the bidders use their additional information to modify their bidding strategy. The unique equilibrium bidding strategy $b_i^{**}\left(t_i\right)$ is indeed given by:

$$
b_i^{**}\left(t_i\right) = \frac{1}{2}\theta_i\left(t_i\right), \text{ for } t_i = t_i^1, t_i^2
$$

but

$$
b_i^{**}\left(t_i^3\right) = 1 \neq \frac{3}{2} = b_i^*\left(t_i^4\right), \quad \text{for } t_i = t_i^3, t_i^4.
$$

The introduction of strategic uncertainty and more private information for the bidders has then a number of important implications for the equilibrium. First, even though the distribution of valuations remains identical across the two type spaces, the larger type space leads to lower bids and lower revenues for the auctioneer. Types t_i^3 and t_i^4 share the same payoff type, $\theta_i = 3$, but have different private information about their competitors and hence different equilibrium bids. In a second price auction, the bidding strategy would remain identical across the type spaces, and hence the revenues would stay constant as well. We also find that in the larger type space the first price auction does not lead to an efficient allocation. The bids of types t_i^2 and t_i^3 are identical even though $\theta_i\left(t_i^2\right) < \theta_i\left(t_i^3\right)$. In Section 5, we show more generally that type spaces richer than the

standard payoff type space lead to a failure of the revenue equivalence theorem, do not permit a revenue ranking between first and second price auction, and lead to a failure of efficiency in the first price auction.

4 Information Management

4.1 Information Acquisition in Committees

We start our survey of recent contributions to the literature on information acquisition with the problem of optimal committee design when information is costly to acquire. Most papers in this area assume that the committee members have common objectives and also that monetary transfers are not used. As a result, it is probably easiest to see what additional insights costly information acquisition brings into the model in this context.

For concreteness, we phrase our discussion of the model in terms of a jury problem. The celebrated Condorcet Jury Theorem (see e.g. Black (1958)) states in its traditional form that decision making in juries under majority rule outperforms decision making by any single individual and as the number of participating voters increases the probability of the correct social decision converges to one. The underlying idea is that in majority decisions, the information of several jury members is aggregated and therefore such decisions are superior to those arrived at by any individual jury member.

The jury chooses between two alternatives: $y_0 \in \{0, 1\}$ where 0 stands for acquitting and 1 stands for convicting the defendant. At the trial, there is uncertainty regarding the possible guilt of the defendant. We model this by a binary state $\omega \in \{0, 1\}$ where 0 stands for innocence and 1 indicates guilt and for simplicity we assume that the prior probability satisfies: $\Pr\{\omega = 1\} = \frac{1}{2}$. All jury members are assumed to have the same payoff functions $u(y_0, \omega)$ satisfying:

$$u(0,0) = u(1,1) = 0, \ u(0,1) = -d_0, \text{ and } u(1,0) = -d_1.$$

In other words, convicting guilty and acquiring innocent defendants is costless. The costs of wrongful conviction is $d_1 > 0$ and the cost of wrongful acquittal is $d_0 > 0$.

At the trial, jury members are presented with evidence on the guilt of the defendant. This is modeled through signal s_i observed by juror i. We assume that the signals are binary, i.e. $s_i \in \{0, 1\}$ and correlated with truth in the sense that $\Pr\{s_i = 0 \,|\, \omega = 0\} = p > \frac{1}{2}$ and $\Pr\{s_i = 1 \,|\, \omega = 1\} = q > \frac{1}{2}$. Furthermore, we assume that signals are independent across jurors conditional on the state ω. Decisions in the jury are reached by majority voting. The vote of juror i is denoted by $v_i : S_i \to [0, 1]$, where $v_i(s_i)$ is understood to be the probability of voting to convict after

observing signal s_i. The jury decision is given by:

$$y_0 : \{0,1\}^I \to [0,1],$$

where $y_0(v)$ gives the probability of convicting given vote profile v.

The logic behind the Condorcet Jury Theorem runs as follows. If the jury members vote based on their private signal, then the vote counts provide a better signal of ω than the individual s_i. The problem with this argument is, as pointed out by Austen-Smith and Banks (1996), that in general it is not in the interest of an individual juror to vote in accordance with their private signal. When the voting stage is seen as a Bayesian game, sincere voting, i.e. $v_i(0) = 0$, $v_i(1) = 1$ for all i is not a Bayesian equilibrium of the game. The reason for this is that at the moment of casting their votes, each jury member must condition her beliefs about the innocence of the defendant on the event that her own vote is pivotal. In a majority rule vote, this implies that the other jury members' votes are equally split. If $p > q$, equal split together with sincere voting implies that $\omega = 1$ is much more likely than $\omega = 0$ and as a result, the individual juror has an incentive to discard her own information. Feddersen and Pesendorfer (1998) compare the expected equilibrium payoff from different voting rules ranging from simple majority to unanimity as a function of the cost parameters d_0 and d_1. By concentrating on symmetric equilibria where the individual jurors' strategies are responsive to private signals, they show that a wide range of rules can be optimal.

To see how costly information acquisition changes the situation, Persico (2004) considers a simple modification to the jury problem above. The signal of each jury member is observed only with cost $c > 0$. This cost is assumed to be private and as a result, a discrepancy between social and private incentives for acquiring information arises.[2] While Feddersen and Pesendorfer (1998) obtain the result that the expected payoff from jury decisions increases in the number of members on the jury, Persico (2004) concludes that optimal jury size is bounded even if the private costs of information acquisition are not accounted for in the social welfare calculation. The reason for the difference in the results depends on the fact that information acquisition by the jurors brings in an element of moral hazard into the decision making process. In order for the jurors to be willing to pay for information, their probability of being pivotal must remain non-negligible. This is only possible in juries of bounded size. Perhaps more interestingly, Persico (2004) finds that the optimal voting rule is independent of d_0 and d_1 and instead depends on the statistical nature of evidence, i.e. on p and q. For the special case where $p = q$, he shows that for small c, the optimal supermajority in the jury decisions converges to p.

---

[2]In the literature on jury decisions, the role of monetary transfers has been ignored. This seems to be a reasonable approximation to most committee decision making processes that are observed in the real world. In addition, Persico (2004) shows that with monetary transfers the problem of inducing efficient information acquisition can be trivially solved.

A second remarkable property of jury design under costly information acquisition is that the voting rule is efficient given the information acquired by the jury members. In the setting of Persico (2004), this property arises partially from the fact that the analysis focuses on pure strategy equilibria. Under this restriction, any suboptimal decision rule would imply that some agents do not acquire information. Mukhopadhaya (2003) concentrates on the symmetric mixed strategy equilibrium and shows that for a fixed voting rule, increasing the jury size may decrease the accuracy of decisions when information acquisition is costly.

Gershkov and Szentes (2004) consider the optimal method of inducing information acquisition and eliciting it truthfully from homogenous committee members subject to the requirement that the decisions must be ex post efficient. In other words, they require that given the information collected in the committee, the decision must agree with the optimal one. They show that the optimal method of gathering information is to approach the committee members sequentially but withholding the previous record of both who has been approached and what information has been transmitted. It is also interesting to note that their optimal mechanism features randomized decisions on whether to collect additional information.

In a similar problem, Smorodinsky and Tennenholtz (2005) show that a sequential mechanism is also optimal in a class of mechanisms that arrive at the correct social decision with probability 1. In this paper, there is no trade-off between costs of information acquisition and the accuracy of the decision.

Gerardi and Yariv (2005) remove the restriction on ex post efficiency of the mechanism. They show that the optimal decision rule is not generally of the type considered in Persico (2004), but rather that it may involve randomizations and ex post inefficient decisions.

The issue of signal accuracy is addressed in Li (2001). In that paper, all jury members invest in information that is useful for determining the guilt of the defendant. In contrast to the other papers surveyed here, Li assumes that the signals are publicly observable. As a result his model is very close to traditional free-riding models of informational externalities. He shows that in order to provide good incentives for information acquisition, it may be optimal to distort the rule mapping signals to decisions. Martinelli (2006) considers a voting model in which the citizens have identical preferences but have a continuous choice regarding the precision of their private information. As the number of participating citizens increases, each individual agent decreases the precision of her private information, but with $c'(0) = c''(0) = 0$ at zero precision, in the limit the simple majority rule leads to the election of the best candidate with probability one.

Finally, Cai (2003) considers the optimal size of a committee under a fixed decision rule in a committee when the members have heterogenous payoff functions. If individual committee members have preferences different from those of the designer of the committee, they have an incentive to

distort their reports to the designer. The main observation of the paper is that preference diversity may increase the individual members' incentives for acquiring information. As a result, the optimal size of committees may be higher under preference diversity as the free rider problems are alleviated.

To summarize, the papers reviewed in this section demonstrate in a simple setting how mechanism design problems must be modified in order to take into account the costs of getting informed. When jury members have the same objectives, but bear the cost of information acquisition privately, free riding becomes an issue in models where information acquisition decisions are not observable. If it is possible to commit to decision rules at the start of the game, free riding can be fought to some extent by an appropriate choice of the decision rule. Sometimes this may involve taking decisions that are suboptimal in light of the collected information. Even when restricted to ex post optimal decision rules, the design of an appropriate extensive form for eliciting information from the jury members provides insights into the general problem.

4.2   Information in Principal-Agent Models

Throughout this paper, we assume that information acquisition is covert, in other words, the principal does not see whether the agent has acquired additional information or not. It is easy to see in the single agent setting that information acquisition adds an element of moral hazard on top of the original adverse selection model. Consider for instance the model where a principal sells an indivisible object to an initially uninformed agent. At cost c, the agent can learn privately her valuation for the object. It is clear that timing plays a crucial role in the analysis of this model. If the principal offers contracts after information acquisition, the model reduces to an adverse selection model conditional on the equilibrium level of information acquisition. If the contract is offered prior to information acquisition, we are in the traditional moral hazard world where the principal extracts all surplus from the agent. Finally, if contract offers and information acquisition decisions are simultaneous, information acquisition decisions are often in mixed strategies and the principal tries to screen the informed buyers from the uninformed.

The role of information acquisition in a principal-agent setting has been investigated in a series of papers by Cremer and Khalil (1992), Cremer, Khalil, and Rochet (1998a) and Cremer, Khalil, and Rochet (1998b). In Cremer and Khalil (1992), the basic problem is a standard adverse selection problem of regulating a monopolist with unknown cost as in Baron and Myerson (1982). The new element is that the agent does not know her type at the moment the contract is offered. She can learn her type, say her marginal cost, either before or after signing the contract. A cost c must be paid to acquire information prior to signing the contract, whereas after signing the contract the type is revealed at zero cost. Information acquisition is therefore socially inefficient. The private

benefit for the agent, however, is that she may be able to reject contract offers which would not be profitable given her marginal cost. Cremer and Khalil (1992) show that the ability of the agent to acquire information decreases the downward distortion at the production stage. The optimal contract raises the expected value of the contract, type by type, so that the agent will have no incentive to acquire the information in equilibrium.

The distinction between costly pre-contract and free post contract information is also central in a recent study by Matthews and Persico (2005) on the excess refund puzzle. They consider the optimal price and refund policy of sellers when the potential buyers can either engage in costly research to assess the value of the object or wait until delivery and inspection of the object. As the return of the object is costly, the optimal selling policy has to find a balance between returns and sales. Similar to Cremer and Khalil (1992), they show that it might be optimal for the seller to offer a refund policy sufficiently generous so as to prevent the buyer in equilibrium to acquire information. The distortion in the refund policy relative to the socially optimal policy leads to an excess in refunds.

In Cremer, Khalil, and Rochet (1998a), the setting is modified by assuming that all information about the cost structure has to be acquired at some fixed cost c. Again, the impact of information acquisition affects both the production schedule and the rent to the agent. For a sufficiently small c, the optimal contract is the standard Baron-Myerson contract. As the cost of information acquisition increases, the value of the contract decreases for the principal. The optimal contract reduces the distortion for low cost types, and increases it for high cost types. This is the most efficient way to increase the rent for the agent so that she has an incentive to acquire the information. At higher c, it is optimal to leave a rent to the uninformed agent. As the principal cannot receive the entire surplus, the production level is below the ex-ante efficient level. As information is costly, it may not be optimal to acquire information even from a social point of view. An open issue is then whether the design of the contract by the principal will lead the agent to take a socially efficient decision regarding information acquisition or whether it will introduce a systematic distortion in the decision of the agent.

Finally, in Cremer, Khalil, and Rochet (1998b), the decision by the agent to get informed is taken covertly before the contract is offered. This reversal in the timing of the decision introduces strategic uncertainty for the principal as the agent may randomize over information acquisition. The resulting equilibrium is one in which the principal offers a menu of contracts, one of which is chosen by the informed and the other is chosen by the uninformed agent. The two contracts display partial pooling, in a sense that for low marginal cost of production, informed and uninformed will produce the same quantity. For intermediate and high production cost, the informed agent will see more downward distortions, and relative to standard Baron-Myerson type contracts, the production will

be higher respective lower for medium and high cost types. The change in the production schedule is enacted so as to efficiently generate surplus for the informed agent and give him incentives to acquire information.[3]

4.3   Information Acquisition in Auctions

Within the field of mechanism design, auction theory has seen the largest number of contributions in the last decade. Surprisingly few of those papers have focused explicitly on costly information acquisition. This is somewhat puzzling given the close connections between auctions and price formation processes in competitive markets. Milgrom (1981) explores the issue of information acquisition in a model similar to the one presented in the motivating example. His main concern is on determining whether the model can be used in providing foundations for the fully revealing rational expectations equilibrium. The connections to the rational expectations equilibrium have been since worked on extensively but the issue of information acquisition has received considerably less attention. In our view, the questions relating to socially optimal information acquisition remain open for a large class of auctions models.[4]

Early contributions to the literature compared the revenue generation across different auction formats, most notably between first and second price auctions. Matthews (1977) and Matthews (1984) obtained the result that the two formats lead to the same expected revenue in a special case of an affiliated model. This result is also later found in a sequence of papers on the independent private information case. These include Hausch and Li (1991), Tan (1992) and Stegeman (1996). The most direct way of seeing why private values settings lead to same revenue rankings for different auction formats is to observe that by the revenue equivalence theorem, they are equivalent to the Vickrey auction. Hence the ex ante incentives for investing in information (or even to make more general investments) must be the same. Rogerson (1992) makes this point in a more general mechanism design setting than the current auctions model.

If the auction designer has a utilitarian welfare objective, it is again easy to see that the agents have the correct incentives to acquire information in a socially optimal manner. In the Vickrey auction, individual payoffs, when viewed as functions of own payoff type only, coincide with the sum of payoffs to all players (up to the addition of a constant). As a result, individual incentives

---

[3] A literature on *delegated expertise* that started with Demski and Sappington (1987) considers information acquisition in the moral hazard model. A recent contribution by Malcolmson (2004) reflects the state of this literature. Starting with Aghion and Tirole (1997), this model has been investigated the role of information acquisition in the optimal design of organizations from an incomplete contract point of view.

[4] A notable exception is Jackson (2003) who shows that in auction setting with a large number of bidders and costly information acquisition information aggregation may fail and not lead to the efficient allocation.

coincide with those of the planner.

Information acquisition in auctions has also been modelled as an auction with costly entry. Johnson (1979), French and McCormick (1984), McAfee and McMillan (1987), Levin and Smith (1994) formulate entry as a model in which potential bidders do not possess private information until they incur an entry cost. Upon incurring the cost, they acquire a private signal about the value of the object.

In a more general model of affiliated values, Persico (2000) shows that the incentives for information acquisition are in general different across different auction formats. In particular, he shows that the marginal incentives for acquiring additional information are higher for first price auctions than for second price auctions. This may overturn the general superiority of second price auctions as demonstrated in Milgrom and Weber (1982). In a model with affiliated values, additional information allows more accurate predictions of other players' bids. As the transfers in a first price auction depend on own bids, it is important to obtain such information in order to be able to shade own bids optimally.

In Bergemann and Välimäki (2002), we consider the possibility of maintaining the utilitarian optimal allocation in a model of interdependent but statistically independent valuations. Each bidder i acquires information on ω_i and this information is independent across the bidders. As explained above, we can view the information acquisition decision as a choice of distributions over the posterior beliefs θ_i on Ω_i. Bidder i has an expected payoff

$$u_i(\theta) = u_i(\theta_1, ..., \theta_I) = \int_\Omega \widetilde{u}_i(\omega_1, ..., \omega_I) \, d\theta_1(\omega_1)...d\theta_I(\omega_I),$$

where we recall that $d\theta_i(\omega_i)$ is the conditional distribution over ω_i given the realization of signal or posterior θ_i. The utilitarian planner would like to allocate the object to bidder i such that

$$u_i(\theta) \geq u_j(\theta) \text{ for all } j \in \{1, ..., I\}.$$

As explained in the motivating example, this can be done using the generalized VCG mechanism when the utility functions satisfy the single crossing property:

$$\frac{\partial u_i(\theta)}{\partial \theta_i} \geq \frac{\partial u_j(\theta)}{\partial \theta_i} \text{ for all } i, j \in \{1, ..., I\}.$$

We assume that ω_i and s_i satisfy monotone likelihood ratio property for all α_i. Furthermore we assume that the experiments are indexed in such a manner that $\alpha > \alpha'$ implies that experiment α is better in the sense of Lehman's order of effectiveness than α'. For this ordering, it makes sense to assume that $c(\alpha)$ is strictly increasing in α. Our main finding in Bergemann and Välimäki (2002) is that if $\widetilde{u}_j(\omega_i, \omega_{-i})$ is decreasing in ω_i for all $j \neq i$, then the VCG auction gives too low incentives

for information acquisition to the individual bidders. If $\widetilde{u}_j(\omega_i, \omega_{-i})$ is increasing in ω_i for all $j \neq i$, then the individual agents have an incentive to engage in excessive information acquisition.

It should be pointed out that this result does not guarantee that all the equilibria of the information acquisition game between the individual bidders feature excessive information acquisition in the case where $\widetilde{u}_i(\omega)$ is increasing in ω_j. It is simply a local comparison of individual and social incentives for information acquisition. As such, it shows that the utilitarian optimum is not achievable, but it does not tell us definitively whether equilibrium information acquisition is excessive or not. In any case, it is clear that the best mechanisms must trade off losses at the information acquisition stage and losses at the allocation stage. In Bergemann, Shi, and Välimäki (2005), we verify that in a model with binary information acquisition decisions equilibria of the information acquisition game feature excessive information acquisition when $\widetilde{u}_i(\omega)$ is increasing in ω_j.

4.4   Dynamic Auctions

Section 4.3 dealt with static mechanisms where the information acquisition decision is taken prior to executing the mechanism. In dynamic auctions such as the ascending price auction, information about the valuations of the opponents is disclosed as the mechanism is run. As a result, the timing of information acquisition becomes a key consideration for the bidders in such auctions. One of the main insights of the papers reviewed in this section is that the dynamic auction formats may make it easier to arrive at socially optimal decisions and they may also generate higher revenues to the seller than their static counterparts.

Compte and Jehiel (2000) compare the performance of a second price sealed bid auction and an ascending price auction in the presence of information acquisition. They consider a private value environment in which all but one agent are privately informed about the value, but the final bidder has to pay a cost to acquire and assess her valuation for the object. The ascending auction then provides the uninformed bidder with an option to acquire information should the chances of winning as expressed by bidding and drop-out behavior of the competitor be reasonably good. They show that the ascending price auction generates a higher expected welfare than the sealed bid auction. If the number of bidders is sufficiently large, then the ascending price auction also increases the expected revenue for the seller. Compte and Jehiel (2004) use the fact that the ascending price auction offers the uninformed bidder an option value to show that if some additional information is likely to arrive in the future, then the uninformed bidder will stay in the auction even when the price has reached her expected valuation. Rezende (2005) offers a dynamic auction model in which the private information of each bidder is characterized by her initial and unbiased estimate of the value of the object and a private cost to learn the true value of the object. In an ascending

price auction, each bidder observes the current price level (but not the drop-out behavior of the competing bidders) and decides if and when to acquire the additional information. It is shown that this sequential auction format is guaranteed to generate larger revenues relative to the sealed bid auction provided the number of bidders is sufficiently large. A recent paper by Cremer, Spiegel, and Zheng (2003) shows how the seller can extract the entire surplus from buyers when information is costly to acquire.[5]

The cost of acquiring information motivates the analysis of indicative bidding in Ye (2005). Commonly, the sale of assets or entire companies is conducted through a two-stage auction process. In the first stage, a large group of bidders is invited to make indicative, but non-binding offers, and in the second stage a subset of the first stage bidders is invited to make final bids for the object of sale. Ye (2005) derives the optimal auction in the presence of information or due diligence costs between the first and second stage bids.

4.5   Information Disclosure in Auctions

Up to this point our discussion of auctions has focused on the case where bidder i can obtain an additional signal s_i on ω_i. In the previous section, we allowed for the possibility of learning about other bidders' valuations during the auction. In some circumstances, it is natural to consider also the case where other players may provide additional information to a bidder. In this section, we concentrate on the case where the auctioneer has access to signals that she may reveal to the bidders. Examples of such information disclosures include allowing the bidders to inspect the object prior to the auction and providing an independent evaluation of the authenticity of a painting etc.

While the focus in the previous sections was on the case where information is costly to acquire, a natural starting point for this section is the case where information is free. The reason for this difference is that in contrast to the previous setting, it may now be in the best interest of the auctioneer not to provide the bidders with full information even when there is no charge associated with this information release. Once the form of optimal information release has been determined, we can address the question of optimal information production by the auctioneer.

Since the discovery of the 'linkage principle' in Milgrom and Weber (1982), a lot of attention has been devoted to the question of information disclosure by an informed auctioneer. As shown by Milgrom and Weber in an affiliated values models, it is revenue enhancing for the auctioneer to disclose information publicly to the participants in a wide range of auction formats.

In the last few years, the issue of information disclosure in auctions has received a lot of

---

[5]A complementary literature in theoretical computer science investigates mechanism design when it is costly to elicit the preference profile, see e.g. Parkes (2004). This literature emphasizes the role of proxy bidding and the use of indirect mechanisms.

attention. If the affiliated values model is asymmetric in the sense that the public information affects the bidders' valuations in a differential manner, Ganuza (2004) shows that linkage principle may fail and it may be optimal for the auctioneer to reveal her private information partially. Furthermore, Perry and Reny (1999) and Foucault and Lovo (2003) show that linkage principle does not necessarily hold in auctions with multi-dimensional signals. With independent information, Board (2005) shows that releasing information is in general revenue decreasing for second price auctions when there are only two bidders.[6]

Starting with Mares and Harstad (2003), more general ways of communicating information to the bidders have been considered. Mares and Harstad assume that the auctioneer can commit to revealing the information to only one of the bidders. They give examples where this type of proprietary disclosure of information dominates public disclosure in terms of generating higher revenues. They also show that it may be particularly useful for the seller to release the proprietary information to bidders that are initially disadvantaged.

Information disclosure has also been studied in models with private information. For such models, the effects behind the original linkage principle are absent and the incentives for disclosing information must have a different origin. Bergemann and Pesendorfer (2001) study a model where an auctioneer chooses the form of a signal s_i to show to each bidder i. More specifically, the auctioneer chooses a general information structure S_i of Ω_i and bidder i observes signal $s_i \in S_i$ associated with a conditional probability $p(\omega_i | s_i)$. The auctioneer does not know the signal realization, but calculates its distribution from her prior distribution on Ω_i. Once bidders have their information, an optimal auction in the sense of Myerson (1981) is run. The main result of the paper is that it is in general optimal for the auctioneer to provide each bidder with a coarse partition, which reveals information only partially, and, if feasible assign asymmetric partitions. This is easily seen in a two-bidder example example where $\omega_i \in \{1, 3\}$ and the prior on Ω_1 is independent of the prior on Ω_2 and $\Pr\{\omega_i = 1\} = \frac{1}{2}$ for $i \in \{1, 2\}$. By choosing $S_1 = \{\{1\}, \{3\}\}$ and $S_2 = \{\{1, 3\}\}$ and running the auction where bidder i wins if $s_1 = \{3\}$ and pays 3 and bidder 2 wins if $s_1 = \{1\}$ and pays 2. The expected revenue from this auction is $\frac{5}{2}$ which is more than the optimal revenue of 2 when no information is release or $\frac{9}{4}$ when all information is released.

In Eso and Szentes (2004), a different approach to information disclosure is adopted. Rather than giving the information for free to the potential bidders, the auctioneer sells additional information to possibly privately informed bidders. The starting point for this paper is that bidders

---

[6]Ivanov (2005) considers information disclosure in the model of strategic information transmission of Crawford and Sobel (1982). As in Bergemann and Pesendorfer (2001), the principal controls the information structure, but only the agent can observe the realization of the signal. In consequence, the optimal information structure is again coarse, yet improves the ex-ante welfare of the principal.

may have some initial private information relating to their valuation for the object. In addition to this, the auctioneer possesses information that determines the total valuation. To model this, let v_i be a random variable representing the private information of bidder i and let s_i denote the signal controlled (but not observed) by the seller. The main result of the paper shows that when s_i is independent of v_i, the seller can obtain the same revenue as she could if s_i was observable to her. The mechanism that allows for this is one where the bidders pay for the right to participate in an auction whose payment and allocation rules are determined by the initial bids. Furthermore, the paper shows that it is optimal to disclose s_i to bidder i (at a cost). For the case where v_i is degenerate, the result is reminiscent of the results on optimal entry fees to auctions. The key difference to the model in Bergemann and Pesendorfer (2001) is that here the participation decision of the bidders takes place prior to observing s_i and hence the individual rationality constraints for the bidders differ across the two papers.

Cremer, Spiegel, and Zheng (2004) considers a sequentially optimal auction in which the seller incurs a cost to disclose the information to each individual bidder. They show that the optimal sequencing is similar to a symmetric information search problem after replacing true by virtual utilities. Shavell (1994) combines the study of information acquisition and disclosure in a simple auction setting. The study is motivated by a series of legal cases highlighting the tension between information acquisition and disclosure (see Kronman (1978) for the legal analysis of this joint problem). A seller owns a single good which she offers to competing buyers. The buyers value the object identically but are uncertain about its true value. The seller can generate information about the true value of the object, but her cost of doing is private information. The analysis distinguishes between two cases: when information has no social and when it has social value. In the first case, the object has the same value to all buyers who value it higher than the seller, whereas in the second case, the optimal use for (or investment in) the object by the buyer will depend on its value. In the case of pure common values, it is socially wasteful to generate information. Yet, with voluntary disclosure, sellers with a low cost of producing information generate the information and disclose the value if it is above a critical value v^* and are silent if the true value is below v^*. The typical unravelling result fails as sellers with a high cost do not produce information. In consequence, the buyer interprets silence as resulting either from ignorance or from low quality. As ignorance is a possibility, an informed seller may be able to extract a value higher than v, conditional on $v < v^*$. This provides cover for the informed type and the incentive to generate information. On the other hand, if information disclosure is mandatory, the seller follows the efficient policy and always acquires information at the socially optimal rate, and therefore acquires no information in the case of pure common values.

The issue of disclosure is of course also relevant in principal-agent models. Lewis and Sappington

(1994) consider a optimal monopoly pricing model with incomplete information. The seller can choose how much information, which improves their estimate about their taste for the products, to disclose to the buyers. They show that typically the optimal release of information is either not to release any information or to release the maximal amount of information. In Lewis and Sappington (1994), the informative signal is private information to the buyer and not observable by the seller. Johnson and Myatt (2006) model advertising as the disclosure of information and analyze the optimal level of advertising in the context of an optimal monopoly pricing problem. Ottaviani and Prat (2001) show in an affiliated value model of monopoly pricing and public disclosure of the signal, that the principal is always better off by committing to disclose any affiliated signal publicly. This result is an extension of the linkage principle from auction models to monopoly pricing models.

4.6   Information and Privacy

A more implicit source of information acquisition arises in repeated interactions with private information. Consider the relationship of a customer with one or more suppliers. If her willingness to pay for the current transaction provides some information regarding her future purchases, then the optimal selling policy today may be affected by considerations about the future value of the relationship. A series of recent papers analyzes these issues, partly motivated by discussion about the role of privacy in electronic retailing. Acquisti and Varian (2005) suggest a two period model in which a single customer purchases repeatedly from a single seller and analyze the optimal pricing policy of the seller. With forward looking buyers and perfectly correlated willingness to pay across the two periods the optimal pricing policy is a sequence of static prices, reminiscent of the analysis of the ratchet effect (see Freixas, Guesnerie, and Tirole (1985)). However, if the buyer displays some myopia, then dynamic pricing, taking into account past purchase decision is optimal even under full commitment. Taylor (2002) also considers a two period model but with different suppliers in every period. The willingness to pay of the customer is positively, but not perfectly correlated, and the initial supplier can sell the transaction information to future suppliers. The paper considers two different regimes regarding the transmission of information, an anonymity and a recognition regime. In line with the ratchet effect, it is shown that forward looking buyers prefer the anonymity regime, but with some myopia, the customer recognition regime and the resulting dynamic pricing may be preferred by customers and sellers. Calzolari and Pavan (2005) consider a two period model, in which a single customer interacts sequentially with two different sellers. The buyer's willingness to pay for the two goods is perfectly correlated. The focus of the paper is on the optimal disclosure policy of the firms, in particular whether the first firm should be allowed to sell the transaction information to the second firm. Calzolari and Pavan (2005) show that if the goods

are complements then the optimal disclosure policy is to provide full information. If the goods are substitutes, then the optimal information policy is non-disclosure.

In an earlier paper, Rothkopf, Teisberg, and Kahn (1991) argued that the advantage of privacy protection conferred by the English auction is one reason why the Vickrey auction is adopted less frequently in practice than might have been expected from its multitude of theoretical advantages. If the true valuation of the winning bidder is revealed in the bidding process, this may open the door for opportunistic behavior by the seller or by third parties. If bidders have such a fear, it may no longer be in their best interest to bid their valuation in the Vickrey auction. In the English auction, only the valuation of the losing bidders can be inferred. As the winning bidders maintain (at least partially) their private information, there is less reason to distort bidding behavior.

5 Robustness

In the first part of the survey, we emphasized the role of endogenous information for the design and the performance of mechanisms. In the second part of the survey, we report when and how mechanisms can achieve their objective even if the planner has little information about the agents' beliefs about each other. As we have seen in the second motivating example, acquiring information about other bidders naturally gives rise to type spaces where the players own payoffs do not give a sufficient description of the strategic environment, but where one must account for higher order beliefs as well. The main task here is to identify which properties of the mechanism guarantee that the mechanism is robust to strategic uncertainty and hence large type spaces.

The discussion of robustness is an old theme in the mechanism design literature. Hurwicz (1972) discussed the need for "nonparametric" mechanisms (independent of parameters of the model). Wilson (1985) states that a desirable property of a trading rule is that it "does not rely on features of the agents' common knowledge, such as their probability assessments." Dasgupta and Maskin (2000) "seek auction rules that are independent of the details - such as functional forms or distribution of signals - of any particular application and that work well in a broad range of circumstances".

5.1 Wilson Doctrine

"Game theory has a great advantage in explicitly analyzing the consequences of trading rules that presumably are really common knowledge; it is deficient to the extent it assumes other features to be common knowledge, such as one player's probability assessment about another's preferences or information.

I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality." Robert Wilson (1987)

Our starting point is the influential formulation of robustness due to Robert Wilson. Wilson emphasized that academic mechanism designers were tempted to assume too much common knowledge information among the players, and suggested that more robust conclusions would arise as researchers were able to relax those common knowledge assumptions. He suggested that the problem is that we make too many implicit common knowledge assumptions in our description of the planner's problem. A possible modelling strategy therefore is to first make *explicit* the *implicit common knowledge* assumptions, and then weaken them. The approach to modelling incomplete information introduced by Harsanyi (1967-68) and formalized by Mertens and Zamir (1985) is ideally suited to this task. Harsanyi argued that by allowing an agent's type to include her beliefs about the strategic environment, her beliefs about other agents' beliefs, and so on, any environment of incomplete information could be captured by a type space. With this sufficiently large type space, the universal type space, it is true that there is common knowledge among the agents of each agent's set of possible types and each type's beliefs over the types of other agents.

However, as a practical matter, applied economic analysis tends to assume much smaller type spaces than the universal type space, and yet *maintains* the assumption that there is common knowledge among the agents of each agent's type space and each type's beliefs over the types of other agents. An important early paper by Neeman (2004) showed how rich type spaces can be used to relax implicit common knowledge assumptions in a mechanism design context. In particular, he considered a model of surplus extraction as Cremer and McLean (1985) and showed how rich types space may lead to a failure of the surplus extraction result. Heifetz and Neeman (2006) strengthen this insight and show that generic priors do not permit full surplus extraction. We shall shortly see further instances in which the small type space assumption imposes very substantive restrictions.

5.2   Robust Mechanism Design

In order to accommodate a planner who knows little about the agents' beliefs about other agents' types, a recent literature has looked at mechanisms that implement the social choice correspondence in *ex post equilibrium*. Bergemann and Morris (2005c) consider a situation where each player has one of a set of possible payoff types and the social planner seeks to implement a social choice objective mapping payoff type profiles to sets of acceptable outcomes. They are interested in partial implementation - i.e., whether truthtelling in the direct mechanism is consistent with the

social choice correspondence? The usual approach to this question would be to assume a commonly known common prior on the payoff types. Partial implementability is then equivalent to Bayesian incentive compatibility in the direct mechanism. Bergemann and Morris (2005c) ask instead when it is possible to implement the social choice correspondence in equilibrium, whatever the players' beliefs and higher order beliefs about other players' types.

Holding fixed the payoff environment, one can construct many type spaces where an agent's type specifies both her payoff type and her belief about other agents' types, as we illustrated in the introductory example. Crucially, there may be many types of an agent with the same payoff type. Intuitively, the larger the type space, the harder it is to implement the social choice objective, as there are more incentive constraints to be satisfied, and so the more "robust" the resulting mechanism is. The smallest type space is the *payoff type space* where the possible types of each agent are equal to the set of payoff types and a common knowledge prior over this type space is assumed. This is the canonical type space in the mechanism design literature. The largest type space is the union of all possible type spaces that could have arisen from the payoff environment. This is in many circumstances equivalent to working with a *universal type space* in the sense of Mertens and Zamir (1985).[7] There are many type spaces in between the payoff type space and the universal type space that are also of interest. While maintaining that the above payoff environment is common knowledge, one would like to allow the agents to have all possible beliefs and higher order beliefs about their types. A flexible framework for modelling such beliefs and higher order beliefs are "type spaces". A type space is a collection

$$\mathcal{T} = \left(T_i, \widehat{\theta}_i, \widehat{\pi}_i \right)_{i=1}^{I}.$$

Agent i's *type* is $t_i \in T_i$. The type of agent i must include a description of her payoff type. Thus there is a function

$$\widehat{\theta}_i : T_i \to \Delta\left(\Theta_i\right),$$

with $\widehat{\theta}_i\left(t_i\right)$ being the probability distribution of agent i's *payoff type* when her type is t_i. In particular, agent i might be uncertain about her own payoff type. A type of agent i must also include a description of her beliefs about the types of the other agents. Write $\Delta\left(Z\right)$ for the space of probability measures on the Borel field of a measurable space Z. The belief of type t_i of agent i

---

[7]Yet, Bergemann and Morris (2001) and Battigalli and Siniscalchi (2003b) emphasize that type spaces may allow for more correlation than is captured in the belief hierarchies of types as in Mertens and Zamir (1985). More precisely, identifying types that have identical hierarchies may lead to a loss of information. Dekel, Fudenberg, and Morris (2005) and Ely and Peski (2006) propose interim rationalizability as a solution concept under which all type spaces that have the same hierarchies of beliefs also have the same interim rationalizable outcomes.

is a function

$$\widehat{\pi}_i : T_i \rightarrow \Delta\left(T_{-i}\right),$$

with $\widehat{\pi}_i\left[t_i\right]$ being agent i's *beliefs* when her type is t_i. Thus $\widehat{\pi}_i\left(E\right)\left[t_i\right]$ is the probability that type t_i of agent i assigns to other agents' types, t_{-i}, being an element of a measurable set $E \subseteq T_{-i}$.

A type space $\mathcal{T}$ is a *payoff type space* if each $T_i = \Theta_i$ and each $\widehat{\theta}_i$ is the identity map. Type space $\mathcal{T}$ is *finite* if each T_i is finite. Finite type space $\mathcal{T}$ has *full support* if $\widehat{\pi}_i\left(t_i\right)\left[t_{-i}\right] > 0$ for all i and t. Finite type space $\mathcal{T}$ satisfies the *common prior assumption* (with prior p) if there exists $p \in \Delta\left(T\right)$ such that

$$\sum_{t_{-i} \in T_{-i}} p\left(t_i, t_{-i}\right) > 0 \text{ for all } i \text{ and } t_i$$

and

$$\widehat{\pi}_i\left(t_{-i}\right)\left[t_i\right] = \frac{p\left(t_i, t_{-i}\right)}{\displaystyle\sum_{t'_{-i} \in T_{-i}} p\left(t_i, t'_{-i}\right)}.$$

Definition 3 *A direct mechanism $f : T \rightarrow Y$ is interim incentive compatible on type space $\mathcal{T}$ if*

$$\int_{t_{-i} \in T_{-i}} u_i\left(f\left(t_i, t_{-i}\right), \widehat{\theta}\left(t_i, t_{-i}\right)\right) d\widehat{\pi}_i\left(t_i\right) \geq \int_{t_{-i} \in T_{-i}} u_i\left(f\left(t'_i, t_{-i}\right), \widehat{\theta}\left(t_i, t_{-i}\right)\right) d\widehat{\pi}_i\left(t_i\right)$$

for all i, $t \in T$ and $t'_i \in T_i$.

The notion of interim incentive compatibility is often referred to as Bayesian incentive compatibility. We use the former terminology as there need not be a common prior on the type space. It should be emphasized that a direct mechanism f can prescribe varying allocations for a given payoff profile θ as different types, t and t', may have an identical payoff profile $\theta = \widehat{\theta}\left(t\right) = \widehat{\theta}\left(t'\right)$. By inspection of the ex post incentive constraints in Definition 1, ex post incentive compatibility is sufficient for interim incentive compatibility, but is it necessary?

Bergemann and Morris (2005c) show that interim incentive compatibility on all common prior payoff type spaces is equivalent to ex post incentive compatibility in separable environments. An environment is called *separable* if the outcome space has a common component and a private value component for each agent. Each agent cares only about the common component and her own private component. The social choice correspondence picks a unique element from the common component and has a product structure over all components. In separable environments, interim implementation on all common prior payoff type spaces implies ex post implementation. Whenever the social choice correspondence is a function, the environment has a separable representation since the private value components can be made degenerate. A second leading example of a separable environment is the problem of choosing an allocation when arbitrary transfers are allowed and

agents have quasi-linear utility. If the allocation choice is a function but the planner does not care about the level and distribution of transfers, then the environment is separable.

This result provides a strong foundation for using ex post equilibrium as a solution concept in separable environments. Since ex post implementation implies interim implementation on all type spaces (with or without the common prior or the payoff type restrictions), it also shows the equivalence between ex post implementation and interim implementation on all type spaces. To the extent that the mechanisms required for ex post implementation are simpler than the mechanisms required for Bayesian implementation, these results contribute to the literature on detail free implementation and the "Wilson doctrine".

For separable environments, the restriction to payoff type spaces is not important. But interestingly, outside of separable environments, the restriction matters. Bergemann and Morris (2005c) report a simple example of a two agent quasi-linear environment where the balanced budget requirement holds: transfers must add up to zero. In this example, ex post implementation and interim implementation on all type spaces are both impossible, but interim implementation on all common prior payoff type spaces is possible. The quasi-linear environments with budget balance is a leading example of an economic non-separable environment. With two agents, there is an equivalence between ex post implementation and interim implementation on all type spaces. With at most two payoff types for each agent, there is the stronger equivalence between ex post implementation and interim implementation on all payoff type spaces. But with three or more agents with three or more types, equivalence between ex post implementation and interim implementation on all type spaces breaks down.

For other approaches to formalizing robust mechanism design, see Chung and Ely (2003), Duggan and Roberts (1997), Eliaz (2002), Hagerty and Rogerson (1987), and Lopomo (1998), (2000) and Auriol and Gary-Bobo (2005).

Chung and Ely (2004) consider the optimal auction with private values in large type spaces. They show that a dominant strategy mechanism may achieve a higher payoff than any Bayesian equilibrium mechanism provided that the type space is large. The intuition is that for any given mechanism, there may exist a type space which exposes weaknesses in the incentive constraints and leads to an inferior expected revenue result in comparison to a dominant strategy mechanism in which the agent are only asked to report their payoff type, but not to report any belief type.

5.3 Robust Implementation

The revelation principle only establishes that the direct mechanism has *an* equilibrium that achieves the social choice function. In general, there may be other equilibria that deliver undesirable out-

comes. In the spirit of the "Wilson doctrine", it is then natural to look for implementation results that are *robust* to different assumptions about what players do or do not know about other agents' types. While the possibility of multiple equilibria seems relevant for practical mechanism design problems the theoretical literature has not resulted in many practical insights (with a few recent exceptions such as Ausubel and Milgrom (2005) and Yokoo, Sakurai, and Matsubara (2004)).

In light of the earlier results on robust incentive compatibility, it is natural to ask whether implementation in Bayesian equilibrium for all possible higher order beliefs is equivalent to ex post implementation in the payoff type space. Bergemann and Morris (2005a) investigate the conditions required for ex post implementation i.e. they ask whether it is the case that all ex post equilibria deliver outcomes in the social choice correspondence. The task for the designer, who does not know the agents' types, is to choose a mechanism such that in *every* equilibrium of the mechanism, agents' play of the game results in the outcome specified by the social choice objective at every type profile.

The complete information implementation literature (see Maskin (1999)) makes the assumption of common knowledge of preferences, the Bayesian implementation literature (see Postlewaite and Schmeidler (1986), Palfrey and Srivastava (1989), and Jackson (1991)) makes the assumption that there is common knowledge of a prior on a fixed set of types. This assumption is unlikely to be valid for practical market designers and it imposes a substantive constraint when viewed as a restriction on all possible beliefs and higher order beliefs. Bergemann and Morris (2005b) show that robust implementation is a more stringent requirement than ex post implementation. While the incentive compatibility constraints for this problem are the same as for the ex post implementation problem,[8] the resulting "robust monotonicity" condition (equivalent to Bayesian monotonicity on all type spaces) is strictly stronger than ex post monotonicity (and Maskin monotonicity). The resulting robust monotonicity notions provide full implementation counterparts to the robust mechanism design (i.e. partial implementation) questions discussed earlier. In particular, they show that interim implementation on all type spaces is possible if and only if it is possible to implement the social choice function using an iterative deletion procedure. The observation about iterative deletion illustrates a general point well-known from the literature on epistemic foundations of game theory (e.g., Brandenburger and Dekel (1987), Battigalli and Siniscalchi (2003b)): equilibrium solution concepts only have bite if we make strong assumptions about type spaces, i.e., we assume small type spaces where the common prior assumption holds.

By exploiting the equivalence between robust and iterative implementation, Bergemann and Morris (2005b) obtain necessary and sufficient conditions for robust implementation in general en-

---

[8]This follows from results in Bergemann and Morris (2005c).

vironments. The necessity argument is conceptually novel, exploiting the iterative characterization. The necessary conditions for robust implementation are ex post incentive compatibility of the social choice function and a condition - *robust monotonicity* - that is equivalent to requiring interim monotonicity on every type space. The robust monotonicity condition is very strong and implies both Maskin monotonicity and ex post monotonicity conditions (but is strictly weaker than dominant strategies). As an added benefit, the robust implementation analysis removes the frequent gap between pure and mixed strategy implementation in the literature. The iterative characterization comes with the additional benefit that tight implementation results can be proved via a fixed point of a contraction mapping.

An important paper of Chung and Ely (2001) analyzes the single (and multi-unit) auction with interdependent valuations with dominance solvability (elimination of weakly rather than strictly dominated actions). In a linear and symmetric setting, they reported sufficient conditions for direct implementation that coincide with the ones derived in Bergemann and Morris (2005b). In the environment with linear aggregation, under strict incentive compatibility, the basic insight extends from the single unit auction model to general allocations models, with elimination of strictly dominated actions only (thus Chung and Ely (2001) require deletion of weakly dominated strategies only because incentive constraints are weak). By comparing the conditions for ex post and robust implementation, it becomes apparent that robust implementation typically imposes additional constraints on the allocation problem.

5.4 Local Robustness

The approach of robustness in the above literature requires that a mechanism could be implemented for all possible types spaces. This robustness criterion is therefore clearly very demanding and it is plausible to investigate weaker local robustness criteria. In addition, the approach above requires that the allocation problem can be defined independent of the beliefs of the designer and the agents. Yet there are cases such as revenue maximizing mechanism, (e.g. optimal pricing and optimal auction), that depend on the beliefs of the designer.

Bergemann and Schlag (2005) investigate a robust version of the classic problem of optimal monopoly pricing with incomplete information. The robust version of the problem is distinct in two aspects. First, instead of a given true distribution of valuations, the seller only knows that the true distribution is in a neighborhood of a given model distribution. The enlargement of the set of possible priors represents model misspecification. Second, the objective function of the seller is formulated as a regret minimization rather than a revenue maximization problem. The regret is the difference between the actual valuation of the buyer for the object and the actual revenue

obtained by the seller. The regret of the seller can be positive for two reasons: (*i*) the buyer has a low valuation relative to the price and hence does not purchase the object, or (*ii*) she has a high valuation relative to the price and hence the seller could have obtained a higher revenue. For a given neighborhood of possible distributions, they then characterize the pricing policy which minimizes maximal regret. They describe how the robust policies depend on the model distribution and the size of the risk as represented by the size of the neighborhood.

Segal (2003) also considers optimal pricing with unknown demand. In his model, the seller does not know the distribution from which the buyers' valuations are drawn. However, she knows that the valuation of each buyer represents an independent draw from the same distribution. He then suggest an optimal pricing mechanism in which the seller offers individualized prices. The price of individual i however only depends on the information she received from all customers but i. By making the price independent of the report of agent i, the equilibrium strategy of each bidder is an ex post equilibrium strategy. Similarly, Baliga and Vohra (2003) consider trading models when buyers and sellers do not know the distribution of valuations. They consider dynamic and adaptive mechanism with and without intermediaries. They show that as the number of traders becomes large, the adaptive mechanism achieves the same expected revenue as if the seller were to know the true distribution of the demand. Goldberg, Hartline, and Wright (2001) consider a similar problem but in contrast do not even make the i.i.d. assumption about the valuations of the customers. Without any Bayesian information, they derive the optimal selling mechanism under the competitive ratio. In other words, they maximize the worst case revenue relative to the optimal revenue which could be obtained if the seller were to know the true valuations of the buyers. The worst case analysis and the notion of competitiveness is central in many optimal design problems analyzed in computer science (see the recent survey to online design problems by Borodin and El-Yaniv (1998)). In auction theory, Neeman (2003) analyzes the competitiveness of the second price auction. A recent article by Prasad (2003) presents negative result, and in particular shows that the standard optimal pricing policy of the monopolist is not robust to small model misspecifications.

5.5   Rationalizability and Robustness

An alternative approach of allowing richer beliefs and strategic uncertainty into standard mechanism design is to relax the solution concept from equilibrium to rationalizability, an approach pursued by Battigalli and Siniscalchi (2003a) and Dekel and Wolinsky (2003). Battigalli and Siniscalchi (2003a) consider the standard private value auction with a continuum of valuations and bids. They show that any positive bid up to some level above the Nash equilibrium is rationalizable. In contrast, Dekel and Wolinsky (2003) consider a set-up with a finite number of valuations and

bids, but allow for some degree of affiliation. They show that as the number of bidders increases, the set of rationalizable bids converges to the bid closest to the true valuation. Similarly Cho (2005) considers the first price auction in a model with affiliated values, and analyzes rationalizable strategies after imposing the additional restriction that all feasible bidding strategies have to be monotone. He shows that the winning bid in the set of rationalizable bidding strategies converges to the competitive equilibrium price as the number of bidder increases. Cho (2004) extends the rationalizability analysis to large uniform and double price auctions.

5.6   Strategic Uncertainty in Auction Theory

We finally discuss how rich type spaces and strategic uncertainty modify and change central results in auction theory. Fang and Morris (2005) illustrate the role of large type spaces for the revenue equivalence theorem. They analyze a model of independent private values with two bidders. However each bidder receives a two-dimensional signal, the first element is her private valuation (the valuation type) and the second element is a noisy signal about the valuation of her competitor (the information type). The addition of the second signal enriches the strategic information of each bidder but obviously reduces common knowledge among bidders and auctioneer. The model is thus a natural generalization of the discrete type framework offered in the motivating example. In this simple setting, they compare first and second price auctions and conclude that the revenue equivalence theorem fails and that no definite revenue ranking exists with multidimensional signals, even though the setting remains a private value model. Naturally, the additional strategic information does not change the bidding strategy in the second price auction, but affects the bidding strategy in the first price auction. The additional information can have two distinct effects on the bidding strategy. Suppose that bidder 1 receives a signal that bidder 2 is likely to have a similar valuation. Relative to her bidding strategy without the strategic information, she now has essentially two choices. She can either increase her bid to improve her chances of winning, or she can lower her bid, and focus on winning against lower valuation types of her opponent. The optimal response to the strategic information will depend on the informativeness of the signal and may go either way. In consequence, bidding may become more fierce or more subdued, leaving the revenue ranking open to go in either direction. The multi-dimensional private value model is closely related to the affiliated value model of Wilson (1977) and Milgrom and Weber (1982). Yet, in Fang and Morris (2005), the belief of bidder 1 about bidder 2 depends directly on the value type of bidder 2 rather than the value type of bidder 1 as in the affiliated value model.

Kim and Che (2004) analyze the role of strategic information in a similar setting. In an independent private value setting with I bidders, a subset of bidders observe the valuation of each agent

in its subset but no additional information about the agents in the complementary set. They also find that the revenue equivalence theorem fails and establish that a second price auction generates a higher expected value than the first price auction. Andreoni, Che, and Kim (2005) pursue an experimental study of this set-up and largely confirm the theoretical predictions. Ye (2004) considers an auction with entry. Each bidder has to incur a cost before learning her own valuation. Yet, in contrast to earlier work, each bidder will also receive some noisy information about the value of the competing bidders. If the information potentially available to the bidders after entry is sufficiently rich, then he shows that the Vickrey auction is the only optimal sealed bid auction. Finally a recent paper by Feinberg and Skrypacz (2005) pursues the logic of multidimensional types, in particular the separation between payoff types and belief types in the context of bargaining under incomplete information.

6 Conclusion

In this survey we emphasized the role of information for mechanism design. First, we discussed an emerging literature on the role of endogenous information for the design and the efficiency of the relevant mechanism. Second, we argued that in the presence of endogenous information, the robustness of the mechanism of the type space becomes a natural desideratum. We then discussed some recent approaches to robust mechanism design and implementation.

During our discussion of the recent contributions, we have indicated that many questions remain wide open, and in fact the current research poses and creates many new questions. We end this survey by collecting a few of them.

As we consider the role of information acquisition, it is natural to consider dynamic and in particular mechanisms in which information is acquired sequentially. Recent work by Compte and Jehiel (2000) showed that the ascending price auction improves upon the static second price auction by allowing for contingent information acquisition. Yet in the ascending price auction information arrives in a particular way. The estimated expected value of the competing bidder is increasing over time. It is then natural to ask whether a descending price auction might sometimes be more favorable for information acquisition than an ascending price auction. The advantage of a descending price auction is that bidders receive over time information that their bids are more likely to be competitive, otherwise the clock would have been stopped by a competitor. Interestingly, Klemperer (2002) suggests a sequential combination of English and Dutch auction to enhance entry and deter collusion. A combination of English and Dutch auction could also be optimal to generate information and hence competition among the bidders. As many bidding processes are inherently dynamical in nature, we believe that there are further theoretical as well as

practical reasons to investigate information acquisition in dynamic settings. Bidding in a takeover contest and negotiating the terms for a business proposal are obvious examples. The dynamic nature of bidding process here reflects the actual fact finding about the proposed outcomes and in addition determines the strategic positions based on the information currently at hand.

We saw that the ex post efficient mechanisms may lead to excessive information acquisition in typical auction settings. We can then ask how the ex post efficient mechanism should be modified to achieve a second best solution. There are two natural modification. The slope of the probability that an agent gets the object could be reduced until information acquisition in equilibrium coincides with the social equilibrium. With a completely randomized decision to allocate the object, the agent will not have any incentives to acquire information. Thus if we change the probability from efficient to completely inefficient we eventually correct the incentives to acquire information. For the given interim probability distribution, we can then identify the allocation which leads to the lowest loss in terms of efficiency.

In the area of robustness, much of the recent work focused on testing the robustness of a social choice function or mechanism which can be identified independent of the beliefs of the agents and the designer, the problem of finding an efficient allocation is a classical example. In many relevant design problems, the beliefs of the designer and the agents enter into the determination of the mechanism, the leading example here is seller maximizing revenue from an optimal auction. Formulating the robust mechanism design problem for this class of problems becomes conceptually more difficult. In order to maximize revenue, the designer must be endowed with some beliefs over the agents' types. To formalize a notion of robustness, one ought to consider a set of possible beliefs.

Bergemann and Morris (2005d) suggest one possible way to proceed by maintaining the assumption that the principal is certain about the true distribution over payoff types, but allow the principal to be uncertain about agents' beliefs and higher order beliefs about other agents' types. For a given prior distribution over payoff types, they try to find (i) the optimal mechanism for a given type space, and (ii) the worst case type space which minimizes the revenue of the designer. Even though the distribution over payoff types is kept constant at a given prior, the strategic uncertainty severely limits the designer to extract the surplus. They show that in many instances, the revenue of the auctioneer can be reduced to the level which could be obtained in the ex post equilibrium of the game.

We discussed in some detail the role of large type spaces for implementation. If the agents possess large amounts of private information relative to the designer, then their ability to coordinate actions ought to increase and hence the equilibrium multiplicity problem may become severe. If the agents succeed in coordinating their actions on equilibrium play which is undesirable from the

principal's point of view, then the issue of multiplicity is essential an issue of collusion among the agents. It is thus conceivable that a common framework and characterization techniques to understand robustness, equilibrium multiplicity and collusion in the context of mechanism design might emerge as one result of this research on large type spaces.

ACKNOWLEDGMENTS

References

ACQUISTI, A., AND H. VARIAN (2005): "Conditioning Prices on Purchase History," *Marketing Science*, 24, 367–381.

AGHION, P., AND J. TIROLE (1997): "Formal and Real Authority in Organizations," *Journal of Political Economy*, pp. 1–29.

ANDREONI, J., Y.-K. CHE, AND J. KIM (2005): "Asymmetric Information About Rivals' Types in Standard Auctions: An Experiment," Discussion paper, University of Wisconsin and University of Soutnern California.

AURIOL, E., AND R. GARY-BOBO (2005): "On Robust Constitution Design," Discussion paper, University of Poulouse and University of Paris 1.

AUSTEN-SMITH, D., AND J. BANKS (1996): "Information Aggregation, Rationality, and the Condorcet Jury Theorem," *American Political Science Review*, 90, 34–45.

AUSUBEL, L. M., AND P. MILGROM (2005): "The Lovely But Lonely Vickrey Auction," in *Combinatorial Auctions*, ed. by P. Cramton, R. Steinberg, and Y. Shoham.

BALIGA, S., AND R. VOHRA (2003): "Market Research and Market Design," *Advances in Theoretical Economics*, 3.

BARON, D., AND R. MYERSON (1982): "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50, 911–930.

BATTIGALLI, P., AND M. SINISCALCHI (2003a): "Rationalizable Bidding in First Price Auctions," *Games and Economic Behavior*, 45, 38–72.

BATTIGALLI, P., AND M. SINISCALCHI (2003b): "Rationalization and Incomplete Information," *Advances in Theoretical Economics*, 3, Article 3.

BERGEMANN, D., AND S. MORRIS (2001): "Robust Mechanism Design," Cowles Foundation, Yale University.

——— (2005a): "Ex Post Implementation," Discussion Paper 1502, Cowles Foundation for Research in Economics, Yale University.

——— (2005b): "Robust Implementation: The Role of Large Type Spaces," Discussion Paper 1519, Cowles Foundation, Yale University.

——— (2005c): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

——— (2005d): "Robust Revenue Maximization," Discussion paper, Yale University and Princeton University.

BERGEMANN, D., AND M. PESENDORFER (2001): "Information Structures in Optimal Auctions," Discussion Paper Cowles Foundation Discussion Paper 1323, Yale University.

BERGEMANN, D., AND K. SCHLAG (2005): "Robust Monopoly Pricing: The Case of Regret," Discussion Paper 1527, Yale University and European University Institute.

BERGEMANN, D., X. SHI, AND J. VÄLIMÄKI (2005): "Information Acquisition in Large Markets," Discussion paper, Yale University and University of Helsinki.

BERGEMANN, D., AND J. VÄLIMÄKI (2002): "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70, 1007–1033.

BIKHCHANDANI, S. (2005): "The Limits of Ex Post Implementation Revisited," Discussion paper, University of California at Los Angeles.

BLACK, D. (1958): *The Theory of Committees and Elections*. Cambridge University Press, Cambridge.

BOARD, S. (2005): "Revealing Information in Auctions: The Efficiency Effect," Discussion paper, University of Toronto.

BORODIN, A., AND R. EL-YANIV (1998): *Online Computation and Competitive Analysis*. Cambridge University Press, Cambridge.

BRANDENBURGER, A., AND E. DEKEL (1987): "Rationalizability and Correlated Equilibria," *Econometrica*, 55, 1391–1402.

CAI, H. (2003): "Optimal Committee Design with Heterogeneous Preferences," Discussion paper, UCLA.

CALZOLARI, G., AND A. PAVAN (2005): "On the Optimality of Privacy in Sequential Contracting," *Journal of Economic Theory*.

CHO, I. (2004): "Rationalizability and Monotonicity in Large Uniform Price and Double Auctions," Discussion paper, University of Illinois at Urbana-Champaign.

——— (2005): "Monotonicity and Rationalizability in Large First Price Auction," *Review of Economic Studies*, forthcoming.

CHUNG, K.-S., AND J. ELY (2003): "Implementation with Near-Complete Information," *Econometrica*, 71, 857–871.

CHUNG, K.-S., AND J. C. ELY (2001): "Efficient and Dominance Solvable Auctions with Interdependent Valuations," Discussion paper, Northwestern University.

——— (2004): "Foundations of Dominant Strategy Mechanisms," Discussion paper, Northwestern University.

COMPTE, O., AND P. JEHIEL (2000): "On the Virtues of the Ascending Price Auction: New Insights in the Private Value Setting," Discussion paper, CERAS-ENPC.

——— (2004): "The Wait and See Option in Ascending Auctions," *Journal of the European Economic Association*, 2, 494–503.

CRAWFORD, V., AND J. SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1452.

CREMER, J., AND F. KHALIL (1992): "Gathering Information Before Signing a Contract," *American Economic Review*, 82, 566–578.

CREMER, J., F. KHALIL, AND J.-C. ROCHET (1998a): "Contracts and Productive Information Gathering," *Games and Economic Behavior*, 25, 174–193.

——— (1998b): "Strategic Information Gathering Before a Contract is Offered," *Journal of Economic Theory*, 81, 163–200.

CREMER, J., AND R. MCLEAN (1985): "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent," *Econometrica*, 53, 345–361.

CREMER, J., Y. SPIEGEL, AND C. ZHENG (2003): "Optimal Selling Mechanisms with Costly Information Aquisition," Discussion paper, Northwestern University.

——— (2004): "Optimal Search Auctions," Discussion paper, Northwestern University.

DASGUPTA, P., AND E. MASKIN (2000): "Efficient Auctions," *Quarterly Journal of Economics*, 115, 341–388.

DEKEL, E., D. FUDENBERG, AND S. MORRIS (2005): "Interim Rationalizability," Discussion paper, Tel-Aviv University, Harvard University and Yale University.

DEKEL, E., AND A. WOLINSKY (2003): "Rationalizable Outcomes of Large Independent Private-Value First Price Discrete Auctions," *Games and Economic Behavior*, 43, 175–188.

DEMSKI, J., AND D. SAPPINGTON (1987): "Delegated Expertise," *Journal of Accounting Research*, 25, 288–310.

DUGGAN, J., AND J. ROBERTS (1997): "Robust Implementation," .

ELIAZ, K. (2002): "Fault Tolerant Implementation," *Review of Economic Studies*, 69, 589–610.

ELY, J. C., AND M. PESKI (2006): "Hierarchies of Belief and Interim Rationalizability," *Theoretical Economics*, 1.

ESO, P., AND B. SZENTES (2004): "Optimal Information Disclosure in Auctions and the Handicap Auction," Discussion paper, Northwestern University and University of Chicago.

FANG, H., AND S. MORRIS (2005): "Multidimensional Private Value Auctions," *Journal of Economic Theory*, forthcoming.

FEDDERSEN, T., AND W. PESENDORFER (1998): "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting," *American Political Science Review*, 92, 23–35.

FEINBERG, Y., AND A. SKRYPACZ (2005): "Uncertainty About Uncertainty and Delay in Bargaining," *Econometrica*, 73, 69–91.

FOUCAULT, T., AND S. LOVO (2003): "Linkage Principle, Multi-Dimensional Signals and Blind Auctions," Discussion paper, HEC School of Management.

FREIXAS, X., R. GUESNERIE, AND J. TIROLE (1985): "Planning under Incomplete Information and the Ratchet Effect," *Review of Economic Studies*, 52, 173–191.

FRENCH, K., AND R. MCCORMICK (1984): "Selaed Bids, Sunk Costs, and the Process of Competition," *Journal of Business*, 57, 417–441.

GANUZA, J.-J. (2004): "Ignorance Promotes Competition: An Auction Model of Endogeneous Private Valuation," *RAND Journal of Economics*, 35.

GERARDI, D., AND L. YARIV (2005): "Information Acquisition in Committees," *Games and Economic Behavior*, forthcoming.

GERSHKOV, A., AND B. SZENTES (2004): "Optimal Voting Scheme with Costly Information Aquisition," Discussion paper, University of Chicago.

GOLDBERG, A., J. HARTLINE, AND A. WRIGHT (2001): "Competitive Auctions and Digital Goods," in *12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 735–744, Washington, D.C. ACM-SIAM.

HAGERTY, K., AND W. ROGERSON (1987): "Robust Trading Mechanisms," *Journal of Economic Theory*, 42, 94–107.

HARSANYI, J. (1967-68): "Games with Incomplete Information Played by 'Bayesian' Players," *Management Science*, 14, 159–189, 320–334, 485–502.

HAUSCH, D., AND L. LI (1991): "Private Values Auctions with Endogeneous Information: Revenue Equivalence and Non-Equivalence," University of Wisconsin-Madison.

HEIFETZ, A., AND Z. NEEMAN (2006): "On the Generic (Im)Possibility of Full Surplus Extraction in Mechanism Design," *Econometrica*.

HOLMSTROM, B., AND R. MYERSON (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51, 1799–1819.

HURWICZ, L. (1972): "On Informationally Decentralized Systems," in *Decisions and Organizations*, ed. by C. McGuire, and R. Radner, pp. 297–336. North-Holland, Amsterdam.

IVANOV, M. (2005): "Optimal Strategic Communication: Can a Less Informed Expert Be More Informative," Discussion paper, Pennsylvania State University.

JACKSON, M. (1991): "Bayesian Implementation," *Econometrica*, 59, 461–477.

JACKSON, M. (2003): "Efficiency and Information Aggregation in Auctions with Costly Information," *Review of Economic Design*, 8, 121–41.

JEHIEL, P., AND B. MOLDOVANU (2001): "Efficient Design with Interdependent Valuations," *Econometrica*, 69, 1237–1259.

JEHIEL, P., B. MOLDOVANU, M. MEYER-TER-VEHN, AND B. ZAME (2005): "The Limits of Ex Post Implementation," Discussion paper, ENPC and University of Bonn.

JOHNSON, J., AND D. MYATT (2006): "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, forthcoming.

JOHNSON, R. (1979): "Auction Markets, Bid Preparation Costs and Entrance Fees," *Land Economics*, 55, 313–318.

KALAI, E. (2004): "Large Robust Games," *Econometrica*, 72, 1631–1666.

KIM, J., AND Y.-K. CHE (2004): "Asymmetric Information About Rivals' Types in Standard Auctions," *Games and Economic Behavior*, 46, 383–397.

KLEMPERER, P. (2002): "What Really Matters in Auction Design," *Journal of Economic Perspectives*, 16, 169–189.

KRONMAN, A. (1978): "Mistake, Disclosure, Information and the Law of Contracts," *Journal of Legal Studies*, 7, 1–34.

LEVIN, D., AND J. SMITH (1994): "Equilibrium in Auctions with Entry," *American Economic Review*, 84, 585–599.

LEWIS, T., AND D. SAPPINGTON (1994): "Supplying Information to Facilitate Price Discrimination," *International Economic Review*, 35, 309–327.

LI, H. (2001): "Theory of Conservatism," *Journal of Political Economy*, 109, 617–636.

LOPOMO, G. (1998): "The English Auction is Optimal Among Simple Sequential Auctions," *Journal of Economic Theory*, 82, 144–166.

——— (2000): "Optimality and Robustness of the English Auction," *Games and Economic Behavior*, 36, 219–240.

MALCOLMSON, J. (2004): "Principal and Expert Agent," Discussion Paper 193, Oxford University.

MARES, V., AND R. HARSTAD (2003): "Private Information Revelation in Common Value Auctions," *Journal of Economic Theory*, 109, 264–282.

MARTINELLI, C. (2006): "Would Rational Voters Acquire Costly Information?," *Journal of Economic Theory*, forthcoming.

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.

——— (1999): "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66, 23–38.

MATTHEWS, S. (1977): "Information Acqusition in Competitive Bidding Process," California Institute of Technology.

——— (1984): "Information Acquisition in Discriminatory Auctions," in *Bayesian Models in Economic Theory*, ed. by M. Boyer, and R. Kihlstrom, pp. 181–207. North-Holland, Amsterdam.

MATTHEWS, S. A., AND N. PERSICO (2005): "Information Aquisition and the Excess Refund Puzzle," Discussion paper, University of Pennsylvania.

MCAFEE, P., AND J. MCMILLAN (1987): "Auctions with Entry," *Economic Letters*, 23, 343–347.

MERTENS, J., AND S. ZAMIR (1985): "Formalization of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1–29.

MILGROM, P. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12, 380–391.

MILGROM, P., AND R. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

MUKHOPADHAYA, K. (2003): "Jury Size and the Free Rider Problem," *Journal of Law, Economics and Organization*, 19, 24–44.

MYERSON, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

NEEMAN, Z. (2003): "The Effectiveness of English Auctions," *Games and Economic Behavior*, 43, 214–238.

NEEMAN, Z. (2004): "The Relevance of Private Information in Mechanism Design," *Journal of Economic Theory*, 117, 55–77.

OTTAVIANI, M., AND A. PRAT (2001): "The Value of Public Information in Monopoly," *Econometrica*, 69, 1673–1683.

PALFREY, T., AND S. SRIVASTAVA (1989): "Mechanism Design with Incomplete Information: A Solution to the Implementation Problem," *Journal of Political Economy*, 97, 668–691.

PARKES, D. (2004): "Auction Design with Costly Preference Elicitation," *Annals of Mathematics and Artificial Intelligence*, pp. 1–38.

PERRY, M., AND P. RENY (2002): "An Ex Post Efficient Auction," *Econometrica*, 70, 1199–1212.

PERRY, M., AND P. J. RENY (1999): "On the Failure of the Linkage Principle in Multi-Unit Auctions," *Econometrica*, 67, 895–900.

PERSICO, N. (2000): "Information Acquisition in Auctions," *Econometrica*, 68, 135–148.

——— (2004): "Committee Design with Endogenous Information," *The Review of Economic Studies*, 71(1), 165–94.

POSTLEWAITE, A., AND D. SCHMEIDLER (1986): "Implementation in Differential Information Economies," *Journal of Economic Theory*, 39, 14–33.

PRASAD, K. (2003): "Non-Robustness of some Economic Models," *Topics in Theoretical Economics*, 3, 1–7.

REZENDE, L. (2005): "Mid-Auction Information Acquisition," Discussion paper, University of Illinois.

ROGERSON, W. (1992): "Contractual Solutions to the Hold-Up Problem," *Review of Economic Studies*, 59, 777–793.

ROTHKOPF, M., T. TEISBERG, AND E. KAHN (1991): "Why Are Vickrey Auctions Rare," *Journal of Political Economy*, 98, 94–109.

SEGAL, I. (2003): "Optimal Pricing Mechanism with Unknown Demand," *American Economic Review*, 93, 509–529.

SHAVELL, S. (1994): "Acquisition and Disclosure of Information Prior to Sale," *RAND Journal of Economics*, 25, 20–36.

SMORODINSKY, R., AND M. TENNENHOLTZ (2005): "Overcoming Free Riding in Multi-Party Computations - The Anonymous Case," Discussion paper, Technion.

STEGEMAN, M. (1996): "Participation Costs and Efficient Auctions," *Journal of Economic Theory*, 71, 228–259.

TAN, G. (1992): "Entry and R and D in Procurement Contracting," *Journal of Economic Theory*, 58, 41–60.

TAYLOR, C. (2002): "Private Demands and Demands for Privacy: Dynamic Pricing and the Market for Customer Information," *RAND Journal of Economics*, forthcoming.

WILSON, R. (1977): "A Bidding Model of Perfect Competition," *Review of Economic Studies*, 44, 511–518.

——— (1985): "Incentive Efficiency of Double Auctions," *Econometrica*, 53, 1101–16.

YE, L. (2004): "Optimal Auctions with Endogenous Entry," *Contributions to Theoretical Economics*, 4, 1–27.

——— (2005): "Indicative Bidding and a Theory of Two-Stage Auctions," Discussion paper, Ohio State University.

YOKOO, M., Y. SAKURAI, AND S. MATSUBARA (2004): "The Effect of False-Name Bids in Combinatorial Auctions: New Fraud in Internet Auctions," *Games and Economic Behavior*, 46, 174–188.