

INFORMATION IN MECHANISM DESIGN

By

Dirk Bergemann and Juuso Välimäki

August 2005

COWLES FOUNDATION DISCUSSION PAPER NO. 1532



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

Information in Mechanism Design

Dirk Bergemann*

Juuso Välimäki†

First Version: 06/15/2005

This Version: 08/15/2005

PRELIMINARY

COMMENTS WELCOME

Abstract

We survey the recent literature on the role of information for mechanism design. We specifically consider the role of endogeneity of and robustness to private information in mechanism design.

We view information acquisition of and robustness to private information as two distinct but related aspects of information management important in many design settings. We review the existing literature and point out directions for additional future work

KEYWORDS: Mechanism Design, Information Acquisition, Ex Post Equilibrium, Robust Mechanism Design, Interdependent Values, Information Management.

JEL CLASSIFICATION: C79, D82

*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.

Contents

1	Introduction	3
2	Setup	5
2.1	Payoff Environment	5
2.2	Information Acquisition	6
3	Motivating Examples	7
3.1	Information Acquisition in Generalized VCG auctions	7
3.2	Espionage in First-Price Auctions	9
4	Information Management	10
4.1	Information Acquisition in Committees	10
4.2	Information Acquisition in Auctions	14
4.3	Dynamic Auctions	17
4.4	Information Disclosure in Auctions	18
4.5	Information in Principal-Agent Models	20
4.6	Information and Privacy	23
5	Robustness and Type Spaces	25
5.1	Wilson Doctrine	26
5.2	Type Spaces	27
6	Robust Mechanism Design	29
6.1	Robust Implementation	31
6.2	Local Robustness	34
6.3	Rationalizability and Implementation	36
6.4	Large Types and Strategic Uncertainty	36
7	Conclusion	38

1 Introduction

The mechanism design literature of the last thirty years has been a huge success on a number of different levels. There is a beautiful theoretical literature that 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 the solution to a well defined planner's problem of achieving some objective or maximizing some utility function 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 held information in order to achieve some private or social objective. The task of the planner is then to design a game of incomplete information in which the agents have indeed an incentive to reveal the information. The optimal design of the game will commonly depend on the common prior which the principal and the agents share about the types of the agents. However, here an unfortunate disconnect between the general theory and the applications and the empirical work emerges. 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 the given common prior. Yet, the fine details of the specified environment incorporated in the common prior will often not be 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 come up with useful information for the decision making process, then it is clear that the eventual decision making process has a impact on the willingness to invest in such information generation. If this 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 such information acquisition and information disclosure as two different aspects of an information management problem that we believe is important in many mechanism design settings. In our view, it is important to recognize that 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. We also point out some directions where we think fruitful additional work should be carried out.

Second, we shall 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. He emphasized that academic mechanisms 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. 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" - i.e., they are too sensitive to fine details of the specified environment. In response to the concerns, researchers have developed many attractive and influential results by imposing (in a somewhat ad hoc way) stronger solution concepts and/or simpler mechanisms motivated by robust considerations. A natural theoretical question to ask is whether it is possible to explicitly model the robustness considerations in such a way that stronger solution concepts and/or simpler mechanisms are endogenously generated.

To the extent that the agents have or can get access to private information about their own valuation, the valuation or the beliefs of the others, the concern of the designer for the performance of the mechanism leads him naturally to adopt 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 second talk about the role of espionage 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 in terms of the Wilson doctrine and introduces the language of large type spaces and the related equilibrium notions. Section 6 discusses recent results regarding the robust mechanisms. It also emphasizes the importance of strategic uncertainty by discussing how classic auction results are modified by the introduction of large type spaces. Section 7 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 $i \in \mathcal{I} = \{1, \dots, I\}$. 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 $\theta \in \Theta = \Theta_1 \times \dots \times \Theta_I$. Each agent has utility function $u_i : Y \times \Theta \rightarrow \mathbb{R}$. An important special case is the *quasi-linear environment* where the set of outcomes Y has the product structure $Y = Y_0 \times Y_1 \times \dots \times Y_I$, where $Y_1 = Y_2 = \dots = Y_I = \mathbb{R}$, and a utility function:

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

which is linear in y_i for every agent i .

The collective choice problem is represented by a social choice correspondence $F : \Theta \rightarrow 2^Y \setminus \{\emptyset\}$, a social choice function is given by $f : \Theta \rightarrow Y$. If the true payoff type profile is θ , the planner would like the outcome to be an element of $F(\theta)$, or simply $f(\theta)$. 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' :

$$\theta_i = \theta'_i \Rightarrow u_i(y, \theta) = u_i(y, \theta'). \quad (1)$$

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 his 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 frequently be derived with ease and single crossing

conditions generating incentive compatibility are easy to identify. A conceptual advantage of ex post equilibrium is its robustness to the informational assumptions about the environment. In particular, it often seems unrealistic to allow the mechanism to depend on the designer's knowledge of the type space as Bayesian mechanisms do.¹ We shall initially focus on truthtelling in the direct mechanism and hence for equilibrium it will be sufficient to verify ex post incentive compatibility.

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

$$u_i(f(\theta), \theta) \geq u_i(f(\theta'_i, \theta_{-i}), \theta),$$

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 $f : \Theta \rightarrow Y$ is dominant strategies incentive compatible if, for all i and $\theta \in \Theta$,*

$$u_i(f(\theta_i, \theta'_{-i}), \theta) \geq u_i(f(\theta'), \theta),$$

for all $\theta' \in \Theta$.

If there are private values (i.e., each $u_i(y, \theta)$ 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 a state 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 . We denote the set of probability distributions on Ω by Θ with a generic element $\theta \in \Theta$.

¹Ex 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 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) and Jehiel, Moldovanu, Meyer-Ter-Vehn, and Zame (2004).

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 distribution $F^\alpha(\theta)$ on Θ . We also write the utility function of the agent directly in terms of the posterior and the chosen action $u(y, \theta)$. Under suitable regularity conditions, there is an optimal action $y(\theta)$ for each θ . If we denote the cost of observing experiment α by $c(\alpha)$, the information acquisition problem can be written concisely as follows:

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

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 $(\theta_1, \dots, \theta_I)$. 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 $\omega_i \in \{\omega_l, \omega_h\}$ of the good. We denote by θ_i bidder i 's probability

assessment on the event $\{\omega_i = \omega_h\}$. We assume that the player i 's payoff from obtaining the object at price y_i take the following linear form:

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

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 $y_0 \in \{1, 2\}$. Efficient allocation requires that

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

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

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

and

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

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 Dasgupta and Maskin (2000).

As we have assumed statistical independence across the two bidders' information, 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 $\beta > 0$, then the individual 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 $\theta_i \geq \theta_j$,

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

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. The divergence between the private and the social benefits to information acquisition in a single unit auction with interdependent values has first been observed in Maskin (1992).

INSERT FIGURE 1: SOCIAL GAINS FROM INFORMATION HERE

INSERT FIGURE 2: PRIVATE GAINS FROM INFORMATION HERE

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 Espionage in First-Price Auctions

Our second example demonstrates the importance of modeling information about other players' types. Consider an independent private value first price auction between two bidders for a single object. The valuations are drawn from a common distribution $F(\theta_i)$ on $[0, 1]$ and this data is common knowledge at the outset of the game. After observing one's own valuation, the players may engage in costly espionage. By paying a cost of $c > 0$, each player may observe the valuation of the other bidder. After the players have acquired the information, they bid in a first price auction. We call this game with the added opportunity for information acquisition the modified first price auction. In a second price auction such an option to acquire additional information would never be exercised as the bidders have dominant strategies. In a first price auction, however, the optimal bids depend on the bids of others and as a result, there may be scope for espionage.

For large c , it is clearly never optimal for any bidder to acquire information about her opponent. Consider hence the case of a relatively small c . It is clear that a bidder with a low valuation is not going to engage in espionage as the potential gains from a possibly lower winning bid are outweighed by the cost of espionage. Consider for the moment the bidder with the highest possible valuation $\theta_i = 1$. In the case without espionage, this bidder submits a bid of $\mathbb{E}[\theta]$ and wins the auction with probability 1. If it is possible to acquire information about the other bidder's type, this will be beneficial if

$$\mathbb{E}[\theta - b(\theta)] > c,$$

where $b(\theta)$ is the equilibrium bid of the other bidder in the standard first price auction. When this condition holds, all equilibria of the modified first price auction game display some information acquisition.

Such information about the payoff types of the other bidders turns out to have major consequences for the qualitative features of the equilibrium in the game. First of all, it should be noted that even in the case of a continuous distribution of individual valuations, the equilibria will typically be in mixed strategies. By the previous reasoning, some types of bidder i do not acquire information. Suppose these types bid according to a pure strategy. Those types of bidder j that engage in espionage can then win the auction by matching this bid.² Furthermore, in first price auctions, $b_i(v_i) < v_i$ for some types of uninformed bidders. As long as bidder j engages in espionage with a strictly positive probability, the uninformed bidder could increase her probability of winning by increasing her bid by an arbitrarily small amount.

An immediate consequence of this is that the equilibria in the modified first price auction fail to be efficient. As the uninformed bidders are randomizing, it is a positive probability occurrence that a bidder with a lower valuation for the object wins the auction. In the second part of this survey, we examine the implications of information about the other players' types more generally in mechanism design problems.

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 have assumed that the agents share a common objective function and also the role of monetary transfers has been disregarded. As a result, the mechanism design problem of eliciting information from the agents is probably at its easiest in this particular context and therefore it is easier to see what additional insights costly information acquisition brings into the model. The committee members are assumed to have common objectives or to form a team in the sense of Marschak and Radner (1972)

For concreteness and also in order to conform with most of the papers in the area, 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. The underlying idea is that in majority

²We are assuming here that all ties are broken in favor of the bidder with the higher type.

decisions, the information of several jury members is aggregated in 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 the defendant 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 $\theta \in \{0, 1\}$ where 0 stands for innocence and 1 indicates guilt and for simplicity we assume that the prior probability satisfies: $\Pr\{\theta = 1\} = \frac{1}{2}$. All jury members are assumed to have the same payoff functions $u(y_0, \theta)$ satisfying:

$$\begin{aligned} u(y_0, \theta) &= 0 \text{ if } y_0 = \theta, \\ u(0, 1) &= -d_0, \\ u(1, 0) &= -d_1. \end{aligned}$$

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

At the trial, jury members are presented with evidence on the guilt of the defendant. This is modeled 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 $\Pr\{s_i = 0 | \theta = 0\} = p > \frac{1}{2}$ and $\Pr\{s_i = 1 | \theta = 1\} = q > \frac{1}{2}$. Furthermore, we assume that the signals are independent across the jurors conditional on the state θ . Decisions in the jury are arrived at using a jury decision rule. When decisions are arrived at through a vote, jury members vote by choosing $v_i : S_i \rightarrow [0, 1]$, where $v_i(s_i)$ is understood to be the probability of voting to convict after observing signal s_i . The jury decision is then simply

$$y_0 : \{0, 1\}^J \rightarrow [0, 1],$$

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 on the guilt or innocence of the defendant based on their private signal, then the vote counts provide a better signal of θ than the individual s_i . The problem with this argument is that as pointed out by Austen-Smith and Banks (1996), it is in general not in the interest of an individual juror to vote in accordance with their private signal. When the voting stage in the jury 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. This implies that the other jury members' votes are split equally. If $p > q$, equal split together with sincere voting implies that $\theta = 1$ is much more likely than $\theta = 0$ 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.³ 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 a degree 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 depends rather 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 .

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. It is not clear that this would remain true if mixed strategies are allowed in the process of information acquisition. Mukhopadhyaya (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 for this information gathering is by approaching the committee members sequentially

³In 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.

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 when decision 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. Hence in this paper, there is no trade-off between costs of information acquisition and the accuracy of the decision.

Gerardi and Yariv (2004) 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 it may involve randomizations. In a previous version of the paper, they also showed that similar results can be obtained in a model where the jury members are allowed to communicate prior to reporting their information.

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. To our knowledge, the choice of information precision when information is privately observed remains an open question. In the notation of the current section, one could e.g. assume that there is a cost of increasing the accuracy of the signal $c(p, q)$ that is increasing in its both arguments. If this function is convex, the optimal mechanism would balance the advantages of acquiring a little information at a low marginal cost with the associated free riding costs from distributing the task of information acquisition.

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 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 focussed 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 a lot less attention. In our view, the questions relating to socially optimal information acquisition remain open for a large class of auctions models.

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 type only, coincide with the sum of payoffs to all players (up to the addition of a constant). As a result, individual incentives coincide with those of the planner.

Information acquisition in an auction 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 model in which potential bidders do not possess private information until they incur an entry cost. Upon incurring the cost, they then 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 different in general 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_{ii} on Ω_i . The utilitarian planner would like to allocate the object to bidder i such that

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

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

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

Our main finding in Bergemann and Välimäki (2002) is that if $u_j(\theta_i, \theta_{-i})$ is decreasing in θ_i for all $j \neq i$, then the VCG auction gives too low incentives for information acquisition to the individual bidders. If $u_j(\theta_i, \theta_{-i})$ is increasing in θ_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 $u_i(\theta)$ is increasing in θ_j . It is simply a local comparison of individual and social incentives for information acquisition. As such, it shows that utilitarian optimum is not achievable, but it does not tell 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 $u_i(\theta)$ is increasing in θ_j . There we

consider the case when the payoffs take a simple quasilinear form

$$u_i(y, \theta) = \alpha\theta_i + \frac{1-\alpha}{I-1} \sum_{j \neq i} \theta_j + y_i. \quad (3)$$

If $\alpha = 1$, we are in a private values case. When $\alpha = \frac{1}{I}$, we have a model with pure common values. In this case, the necessary and sufficient condition for Bayesian implementability of the efficient allocation is that

$$\alpha \geq \frac{1}{I}.$$

In the symmetric environment defined by (2), it is clear that efficient allocation of the object requires that:

$$\theta_j > \theta_i \Rightarrow y_0 \neq i.$$

In other words, the object is allocated to one of the bidders with the highest type (either realized or expected).

Incentive compatibility then requires that y_i is constant in θ_i for all announcements that induce the same allocation of the object. The allocation changes only at points where two (or more) bidders have the same type and at these points, incentive compatibility requires that all types make a zero surplus. In order to compute the generalized VCG transfers denote by θ_{-i}^* the highest type of bidders other than i . The transfers are then given by

$$y_i(\theta) = \begin{cases} -\alpha\theta_{-i}^* - \frac{1-\alpha}{I-1} \sum_{j \neq i} \theta_j, & \text{if } \theta_i \geq \theta_j \text{ for all } j \in \{1, \dots, I\}, \\ 0, & \text{if otherwise} \end{cases}$$

We contrast the optimal decisions of a planner that acquires the information for the agents with an equilibrium solution where each agent bears the cost of information acquisition privately. Denote by c_m^* the highest cost of information acquisition that is compatible with the planner acquiring the information for m agents. Similarly, denote the cost threshold in the equilibrium problem by \hat{c}_m . We show that:

1. there is excessive information acquisition in equilibrium: $\hat{c}_m \geq c_m^*$.
2. the difference $\hat{c}_m - c_m^*$ is decreasing in α .
3. the difference $\hat{c}_m - c_m^*$ is decreasing in I .

These results generalize also to auctions where multiple units are sold. As explained above, in the case that $\alpha = 1$, we have $\hat{c}_m = c_m^*$. Consider next the case where I is large. For all $c > 0$, the number of bidders that acquire information is bounded from above by $\frac{1}{c}$. As this bound is

independent of I , it is clear that the common component in the bidders' utility functions converges to $(1 - \alpha) \mathbb{E}\theta$. Hence for large I , the model reduces essentially to the private values model and the equilibrium information acquisition coincides with the efficient level.

4.3 Dynamic Auctions

Section 4.2 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 his 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. The beneficial effect of an ascending price auction on the informational decision is shown in Compte and Jehiel (2001) to extend to the case of many agents who would like to acquire information and to multiple objects. 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.⁴

The cost of acquiring information can be viewed as a specific transaction cost. Motivated by the experimental results in Lucking-Reiley (1999), Carare and Rothkopf (2001) consider the role of transaction costs in a slow Dutch auction. The Dutch auction is said to be slow as the price declines at such a rate, that the optimal bidding strategy of every agent is either to take the object at the current price or return at a later time to make a bid for the lower price. In a simple model with random arrival of the bidders, each bidder can either bid directly or at a cost return to the auction

⁴A 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 mechanism.

at a later time when the price has further depreciated. (Of course, the object might have been sold at the later time already.) Carare and Rothkopf (2001) show that as the cost of returning increases, the agents bid more aggressively and generate a higher revenue for the seller. This model provides a possible rationalization of the experimental findings of Lucking-Reiley (1999) which showed that slow Dutch auction, despite strategic equivalence, generated higher revenue than the other auction formats.

We believe that there are theoretical as well as practical reasons to keep investigating information acquisition in dynamic auctions. First of all, while the superiority of ascending price auctions to second price sealed bid auctions has been demonstrated in some settings, there is nothing to suggest that other auction formats might not perform even better. Descending price auctions also induce information disclosure and it seems to us that a combination of these two formats might perform very well. Second, many bidding processes are inherently dynamic in nature. Bidding in a takeover contest and negotiating the terms for a business proposal are obvious examples. In both of these cases, we believe that the dynamic nature reflects actual fact finding about the proposed outcomes in addition to taking strategic positions based on the information currently at hand.

4.4 Information Disclosure in Auctions

Up to this point our discussion of auctions has focussed on the case where bidder i can obtain an additional signal s_i on θ_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 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.

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 partition S_i of Θ_i and bidder i observes signal $s_i \in S_i$ with the property that $\theta_i \in s_i$. The auctioneer does not know the signal realization, but calculates its distribution from her prior distribution on Θ_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 use asymmetric partitions and not to reveal all information. This is easily seen in a two-bidder example where $\theta_i \in \{1, 3\}$ and the prior on Θ_1 is independent of the prior on Θ_2 and $\Pr\{\theta_i = 1\} = \frac{1}{2}$ for $i \in \{1, 2\}$. By choosing $S_1 = \{\{1\}, \{3\}\}$ and $S_2 = \{\{1, 3\}\}$ and running the auction where bidder 1 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 released 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 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 by the seller. A simplified version of Eso and Szentes (2004) assumes that $\theta_i = v_i + s_i$

and furthermore that both s_i and v_i are independent across bidders. The main result of the paper shows that if s_i is independent of v_i , then the seller can obtain the same revenue as she could if the signal realization 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. For the case where v_i is degenerate, the result is reminiscent of the results on optimal entry fees to auctions. Furthermore, it is shown that it is always optimal to sell the signal s_i to all bidders. Perhaps the key difference to the Bergemann and Pesendorfer (2001) model is that Eso and Szentes (2004) allow for mechanisms that are not individually rational for the bidders at the stage when the auctioneer has released her private information.

The issue of disclosure is of course also relevant in principal-agent models. Lewis and Sappington (1994) consider an 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. In contrast, 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 problem from auction models to monopoly pricing models.

In an incomplete contract setting with hold up, Lau (2004) identifies the optimal information structure. She shows that in the trade-off between ex ante efficiency and ex post efficiency, an intermediate level of asymmetry is optimal. The optimal information structure is derived in a trade-off between the information rent and the bargaining disagreement effect.

4.5 Information in Principal-Agent Models

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 his type at the moment the contract is offered. He can learn his type, say his marginal cost, either before he signs the contract or after he signs the contract. If he wishes to acquire information before signing the contract, then he has to pay a cost c , whereas after signing the contract, he will learn his type at zero cost. It is therefore socially inefficient to acquire the information before the contract is signed. The private benefit for the agent however

is that he may be able to reject contract offers which would not be profitable given his marginal cost structure. In Cremer and Khalil (1992), it is shown that the ability of the agent to acquire information will decrease the downward distortion in production in the high cost state. The optimal contract will raise the expected value of the contract, type by type, so that the agent will have no incentive to acquire the information in equilibrium. Yet, the possibility of acquiring information alone is sufficient to increase his expected value from the contract. 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 will lead to an excess in refunds.

In a recent contribution, Compte and Jehiel (2002), show that the feature that the principal does not wish the agent to acquire information before the contract depends on the presence of a single agent. If on the other hand, the principal faces many agents with unknown cost, then it is generally optimal for some agents to acquire information before they sign the contract. The idea is that faced with a choice of agents, it is now optimal to try to identify a low cost agent. The decision to acquire information before the contract is now of course socially beneficial and the principal can use competition to lower the surplus of the informed agents.

In Cremer, Khalil, and Rochet (1998b), the decision by the agent to get informed is taken before the contract is offered. The reversal in the timing of the decision now introduces strategic uncertainty for the principal as he does not know whether the agent is informed or not. The resulting equilibrium is one in which the principal offers a menu of contract, one which will be chosen by the informed and one which will be chosen by the uninformed. The two contracts will 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 contract, the production will be higher respective lower for medium and high cost types. The change in production is schedule is enacted so as to efficiently generate surplus for the informed agent and give him incentives to acquire information.

Finally, 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 contract is designed first and then the agent has to make an information to acquire information. Again,

the impact of information acquisition affects both the production schedule and the rent to the agent. For sufficiently small cost of information acquisition, the optimal contract is the standard Baron-Myerson contract. As the cost of information acquisition increases, the principal will see a decrease in the value of the contract. The optimal contract will diminish the distortion for low cost types, and increase it for high cost types. This is the most efficient way to increase the rent for the agent so that he has an incentive to acquire the information. If the cost of information increases, it will be optimal to change the contract so that the agent will receive a rent even so he does not have any privileged information. As the principal cannot receive the entire surplus, the production level will be 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 on socially efficient decision regarding information acquisition or whether it will introduce a systematic distortion in the decision of the agent.

Creswell (1988) considers a problem in which a contractor must decide whether to accept to build a house. Before accepting the contract, he spends resources investigating the disutility of production. Ex post, after he begins work he obtains better information and breaches the contract if this disutility is too high. The precontractual investigation is therefore productive as it reduces the probability of breach.

In Lewis and Sappington (1997), the decision to acquire information is embedded in a moral hazard model. In a procurement setting, the agent has to choose an optimal effort level which depends on the state of the world. The agent can acquire information about the state of the world at some cost. The cost of effort is observable and can be contracted upon. The principal then offers the agent a choice between two contract, one for the informed, the other one for the uninformed. The contract for the informed agent provides incentives to acquire information by guaranteeing him more than a dollar for every dollar in cost reduction he achieves in the favorable environment. It also guarantees him a large amount of cost sharing in the unfavorable environment to overall compensate him for the cost of information acquisition.

Shavell (1994) combines the study of information acquisition and disclosure in a simple buyer-seller setting. The study is motivated by a series of legal cases which highlight the tension between information acquisition and its return and disclosure (see Kronman (1978) for the legal analysis of this joint problem). A seller owns a single good which he 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 his cost of doing is private information and varies across types. The analysis distinguishes between two cases: (*i*) when information has no social and when it social value. In the first case, the object has the same value to all buyers which

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 value, it is socially wasteful to generate information. Yet, with voluntary disclosure, sellers which have a low cost of producing information will 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 to apply as sellers with a high cost of generating information will not produce information. In consequence, the buyer will interpret the silence of a seller as coming from two possible sources, ignorance or low quality good. But as ignorance is a possibility, due to high cost, the informed seller will be able to extract a higher value from the object than its true value, conditional on $v < v^*$. This provides the cover for the informed agent and the incentive to generate information. On the other hand, if information disclosure is mandatory, the seller will follow the efficient policy and always acquire information at the socially optimal rate, in particular acquire no information in the case of pure common value.⁵

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 his willingness

⁵Since Demski and Sappington (1987) introduced the model of *delegated expertise*, a growing literature has investigated the role of information acquisition in the optimal design of organizations. In this survey, we will not cover this research area and merely point the reader to some of the important contributions in this area. In a model of delegated expertise, as formulated by Demski and Sappington (1987), the agent has to make two decisions: first he can acquire or refine information about the nature of alternatives, second after receiving the resulting information, he can take an informed action. In the tradition of the moral hazard literature, the decision to acquire information and the received signal are unobservable, the resulting action by the agent may or may not be observable. A recent contribution by Malcolmson (2004) considerably generalizes the model of delegated expertise to a general distribution, a continuum of signals, actions and outcomes. An interesting variation is introduced in Prendergast (1993). The basic problem for worker and manager is to estimate the mean of a normal distribution. Yet, the incentive problem becomes difficult as the worker also observes a noisy signal of the managers observation. He can therefore bias his report about the signal in the direction of the (estimated) information already held by the manager. In Aghion and Tirole (1997), the cost of information acquisition determines the structure of an organization. In an incomplete contract model, they consider the allocation of decision rights among a principal and an agent. The true private and social returns of the project are unknown and information can be gathered by the agent and the principal. The focus is on the allocation of the decision right, as it is assumed to be the only instrument and monetary transfers are not used. In Dewatripont and Tirole (1999), information acquisition determines the structure of the court system. In an otherwise similar model to Aghion and Tirole (1997), the agent can be given monetary incentives based on the final decision. Finally Gromb and Martimort (2004) consider the organization of delegated expertise in setting similar to Dewatripont and Tirole (1999). Yet, in contrast to Dewatripont and Tirole (1999), Gromb and Martimort (2004) consider payments on the basis of the reports of the agents *and* the eventual outcome of the project.

to pay for the current transaction provides some information regarding his 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 (2004) 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 buyer 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 perfect correlated, and the initial supplier can sell the transaction information to future suppliers. The paper considers two different regime 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 willingness to pay by the buyer for each of 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 the information, yet if the goods are substitutes, then 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.

The previous discussion focused on the information acquired before contracting. Taylor (2004) considers a competitive market in which firms post wages, ex ante identical workers apply and information about the applicants is acquired by the firms after the applications are received. Each worker can either have a low or a high productivity and the productivity of a worker is identical across firms. If the firm were to know for sure that a worker has a low productivity, then it would

privately and socially beneficial not to hire the worker. The paper then analyzes the equilibrium incentives to acquire information and compares it to the social efficient incentives. The equilibrium incentives to acquire diverge from the social incentives depending on whether the additional information is likely to be positive or negative. The competitive market assumption guarantees that the expected surplus of the contact goes to the worker rather than the firm. The socially efficient decision compares the expected gains from an informed decision with the expected gains of an uninformed decision which simply equals the expected productivity of the worker. With a competitive market, the firm pays the worker more than the expected value of the contact. In consequence, its equilibrium is not the expected productivity but the equilibrium wage. As the equilibrium wage is larger than the expected productivity, so is difference between equilibrium wage and low productivity relative to expected productivity and low productivity. In turn, the gains from an informed decision which results in rejecting a low productivity are larger in equilibrium than in the social calculus. In consequence, the firm will overinvest in information. The argument is exactly reversed if the information is likely to be positive. Then the relative gains of the firm are lower than the social gains and he underinvests in information. In Taylor (2004), the information is acquired after the terms of trade are determined and thus the price does not reveal any information about the productivity of the worker. Yet, the productivity of the worker is identical across all firms. Hence this is a model of interdependent values, and the equilibrium incentives to acquire information have comparable efficiency properties as in Bergemann and Välimäki (2002).

5 Robustness and Type Spaces

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 gives naturally rise to type spaces where the players own payoffs do not give sufficient descriptions of the strategic environment, but 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 types.

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.⁶ The modelling strategy must be 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 his beliefs about the strategic environment, his 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 (including all possible beliefs and higher order beliefs), it is true (tautologically) 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 maintain 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*

⁶An important paper of Neeman (2004) shows how rich type spaces can be used to relax implicit common knowledge assumptions in a mechanism design context. For other approaches to formalizing robust mechanism design, see Chung and Ely (2003), Duggan and Roberts (1997), Eliaz (2002), Hagerty and Rogerson (1987), Lopomo (1998), Lopomo (2000).

agents.

We shall see shortly that the small type space assumption imposes very substantive restrictions. There has been remarkably little work since Harsanyi checking whether analysis of incomplete information games in economics is robust to the implicit common knowledge assumptions built into small type spaces. Yet towards the end of Section 6 we will discuss some recent contributions which investigate the importance of these implicit common knowledge assumptions in the context of mechanism design.

5.2 Type Spaces

While holding fixed the payoff environment, we can construct many type spaces, where an agent's type specifies both his payoff type and his belief about other agents' types. Crucially, there may be many types of an agent with the same payoff type. The larger the type space, the harder it will be to implement the social choice objective, and so the more "robust" the resulting mechanism will be. 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 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, we would like to allow for 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

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

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

$$\hat{\theta}_i : T_i \rightarrow \Delta(\Theta_i),$$

with $\hat{\theta}_i(t_i)$ being the probability distribution of agent i 's *payoff type* when his type is t_i . In particular, agent i might be uncertain about his own payoff type. A type of agent i must also

⁷Yet, 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 (2004) 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.

include a description of his beliefs about the types of the other agents. Write $\Delta(Z)$ for the space of probability measures on the Borel field of a measurable space Z . The belief of type t_i of agent i is a function

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

with $\hat{\pi}_i[t_i]$ being agent i 's *beliefs* when his type is t_i . Thus $\hat{\pi}_i(E)[t_i]$ is the probability that type t_i of agent i assigns to other agents' types, t_{-i} , being an element of a measurable set $E \subseteq T_{-i}$. In the special case where each T_j is finite, we will abuse notation slightly by writing $\hat{\pi}_i(t_{-i})[t_i]$ for the probability that type t_i of agent i assigns to other agents having types t_{-i} .

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

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

and

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

The standard approach in the mechanism design literature is to restrict attention to a common prior payoff type space (perhaps with full support). Thus it is assumed that there is common knowledge among the agents of a common prior over the payoff types. A payoff type space can be thought of the smallest type space embedding the payoff environment described above.

Fix a payoff environment and a type space \mathcal{T} . A mechanism specifies a message set for each agent and a mapping from message profiles to outcomes. Social choice function f is interim implementable if there exists a mechanism and an interim (or Bayesian) equilibrium of that mechanism such that outcomes are consistent with f . A direct mechanism is a function $f : T \rightarrow Y$.

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

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

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 $\theta = \hat{\theta}(t) = \hat{\theta}(t')$.

6 Robust Mechanism Design

In the face of a planner who does not know about agents' beliefs about others' payoff types, a recent literature has looked at mechanisms that implement the social choice correspondence in *ex post equilibrium*. Bergemann and Morris (2004) consider a situation where each player has one of a set of possible "payoff types" and the social planner wants to implement a social choice correspondence mapping payoff type profiles to sets of acceptable outcomes. They are interested in partial implementation - i.e., whether the truth-telling equilibrium in the direct mechanism 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, so that - using the solution concept of Bayesian equilibrium - partial implementability is equivalent to Bayesian incentive compatibility in the direct mechanism. Instead they ask when it is possible to implement the social choice correspondence in equilibrium, whatever the players' beliefs and higher order beliefs about other players' types. Ex post incentive compatibility is sufficient for this, but is it necessary? They provide a partial characterization for the environments where ex post incentive compatibility is equivalent to being able to implement in equilibrium independent of higher order beliefs. It is true in the economically significant cases of quasi-linear environments (without budget balance constraints) and social choice functions (i.e., when the correspondence is single valued). These results provide microfoundations for using a stronger solution concept to address robustness issues. To the extent that ex post direct mechanisms - when they exist are simpler than Bayesian direct mechanisms for arbitrary common knowledge common priors, the results also favor simpler mechanisms.

Bergemann and Morris (2004) show that the converse result is not always true. They present examples in which ex post implementation is impossible, nonetheless, interim implementation is possible on every type space. The gap arises because the planner may have the equilibrium outcome depend on the agents' higher order belief types, as well as their realized payoff type. The planner has no intrinsic interest in conditioning on non-payoff-relevant aspects of agents' types, but he is able to introduce slack in incentive constraints by doing so.

The main question in Bergemann and Morris (2004) is then to ask when the converse is true. A payoff environment is *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 his 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 we can make private value components degenerate). The other 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 we have a separable environment.

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 environment, the restriction matter. Bergemann and Morris (2004) report a simple example of 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 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.

An important paper of Neeman (2004) shows how rich type spaces can be used to relax implicit common knowledge assumptions in a mechanism design context. 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).

Chung and Ely (2004) consider 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.

6.1 Robust Implementation

Bayesian incentive compatibility analysis suffers from two important limitations. First, as discussed, the analysis typically assumes a commonly known common prior over the agents' types. This assumption may be too stringent in practice. Second, 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 outcomes. In the spirit of the "Wilson doctrine" (Wilson (1987)), 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 does seem to be a relevant one in practical mechanism design problems, particularly in the form of collusion and shill bidding, the theoretical literature is not seen as having developed 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. This problem has been analyzed under the assumption of complete information (see Maskin (1999)) and under the assumption of incomplete information (see Postlewaite and Schmeidler (1986), Palfrey and Srivastava (1989) and Jackson (1991)).

Because an ex post equilibrium is a Nash equilibrium at every type profile, there is a natural relationship between ex post and Nash implementation. In the comparison between the *complete* with the *incomplete* information settings, two important differences regarding the ability of the agents to sustain equilibrium behavior emerges. On the one hand, with complete information, the agents have the ability to coordinate their actions at every preference profile. This makes the designer's problem harder. On the other hand, with complete information the designer can detect individual deviations from the reports of the other agents. This makes the designer's problem easier. The ability of the agents to coordinate in complete information settings makes the task of implementing the social choice outcome more difficult for the designer, but it is made easier by the lack of individual incentive constraints. With incomplete information, the first problem becomes easier, but the second becomes harder. As these two effects are in conflict, they show that ex post and Maskin monotonicity are not nested notions. In particular, either one of them can hold

while the other one can fail. Interestingly, in the class of single crossing environments, ex post monotonicity is always guaranteed as is Maskin monotonicity. Even though ex post monotonicity has to include ex post incentive constraints absent in the complete information world, it turns out that the local property of single crossing indifference curves is sufficient to guarantee ex post monotonicity in the presence of strict rather than weak ex post incentive constraints.

The “augmented” mechanisms used to obtain ex post implementation results inherit some complex and unsatisfactory features from their complete information and Bayesian counterparts. Yet they also identify a number of important settings where ex post implementation is only possible when it is possible in the direct mechanism. This is true, for example, if the social choice function has a sufficiently wide range or if the environment is supermodular. In particular, they show that the direct mechanism has a unique ex post equilibrium in the problem of efficiently allocating goods when bidders have interdependent values (see Dasgupta and Maskin (2000) and Perry and Reny (2002)).

The complete information implementation literature makes the assumption of common knowledge of preferences, the Bayesian implementation literature makes the assumption that there is common knowledge of a prior on a fixed set of types; this both seems unlikely to practical market designers and is a substantive constraint when viewed as a restriction on all possible beliefs and higher order beliefs. Yet in general, robust implementation is a more stringent requirement than ex post implementation. Bergemann and Morris (2005b) analyze the problem of Bayesian implementation under the assumption that the designer has no information on the players’ types. 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 the 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 they obtain necessary and sufficient conditions for robust implementation in general environments. The necessity

⁸This follows from results in Bergemann and Morris (2004).

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. In particular, they consider a general class of interdependent preferences in which the payoff types of the agents can be linearly aggregated. In this environment they show that the social choice function can be robustly implemented if and only if the interdependence is not too large. If α is the weight of the type of agent j (relative to the type of agent i) for the utility of agent i , then the robust implementation condition can simply be stated as: $\alpha < 1/(I - 1)$, where I is the number of agents. Surprisingly, they also show that if $\alpha > 1/(I - 1)$, then not only robust implementation, but even robust virtual implementation fails.

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. In Bergemann and Morris (2005a) it is shown that in single crossing environments, the same single crossing conditions which guarantee incentive compatibility also guarantee full implementation. In contrast, in the linear aggregation environment, Bergemann and Morris (2005b) show that robust implementation imposes a strict bound on the interdependence of the preferences, which is not required by the truthtelling conditions. The contraction mapping behind the iterative argument directly points at the source of the restriction of the interaction term.

In the implementation literature, it is a standard practice to obtain the sufficiency results with augmented mechanisms. By augmenting the direct mechanism with additional messages, the designer may elicit additional information about undesirable equilibrium play by the agents. Yet, in many environments common to applied mechanism designs, such a single crossing or supermodular

preferences, the structure of the preferences may already permit direct implementation. We thus provide necessary and sufficient conditions for robust implementation in the direct mechanism. In the direct mechanism, the agents can alert the designer only by a report of their type. In consequence, the incentive compatibility conditions for the rewards are identical to the truth-telling constraints, and the necessary and sufficient conditions for robust implementation coincide.

Chung and Ely (2003) are also concerned with the robustness of implementation result. They consider complete information implementation and show that recent permissive results in the implementation literature are not robust to the introduction of some incomplete information. More precisely, they consider the notion of undominated Nash equilibrium and show that even though almost any social choice function can be implemented in undominated Nash equilibrium under complete information, the introduction of arbitrarily small incomplete information is enough to re-establish Maskin monotonicity as a necessary condition for implementation. Kunimoto (2004) complements the analysis and points out that the result by Chung and Ely (2003) depends critically on the (implicit) topology. He suggests a coarser topology under which undominated Nash equilibrium is indeed robust to the introduction of complete information.

6.2 Local Robustness

The approach of robustness in the above literature required that a mechanism could be implemented for all possible types space. This robustness criterion is therefore clearly very demanding and it is plausible to investigate weaker, in particular, local robustness criteria. In addition, it requires that the allocation problem could be defined independent of the beliefs of the designer and the agents. In contrast, revenue maximizing mechanism, such as optimal pricing and optimal auction, 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 regarding the valuations of the buyers, in our set-up 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 the 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)* he 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. As part of the analysis, they also determine how the regret varies with the amount of risk faced by the seller.

By pursuing the robust analysis with the notion of regret rather than revenue they combine the attractive features of the worst case analysis with the those of the robust analysis. In particular, for any given neighborhood, the seller uses the information contained in the model distribution and its neighborhood. The seller is minimizing expected regret and thus uses the information contained in the prior. In addition, at the worst case prior, the pricing policy which minimizes regret also maximizes revenue. Thus the regret minimization problem has a direct decision theoretic link to the original objective function of the seller, namely revenue maximization.

The robust policy in the model of Bergemann and Schlag (2005) is the result of a minmax regret problem. The seller could therefore also be interpreted as an ambiguity averse seller in the sense of Gilboa and Schmeidler (1989) if we were to maximize revenue rather than minimize regret. A recent paper by Bose, Ozdenoren, and Pape (2004) investigates the nature of the optimal auction in the presence of an ambiguity averse seller as well as ambiguity averse bidders.

Segal (2003) considers the optimal pricing mechanism with unknown demand. In his model, the seller does not know the distribution from which the buyers valuation are drawn. However, he knows that the valuation of each buyer represents a 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 he 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 achieve the same expected revenues 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, an in particular shows that the

standard optimal pricing policy of the monopolist is not robust to small model misspecifications.

6.3 Rationalizability and Implementation

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.

6.4 Large Types and Strategic Uncertainty

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 his private valuation (the valuation type) and the second element is a noisy signal about the valuation of his 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. 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 his bidding strategy without the strategic information, he now has essentially two choices. He can either increase his bid to improve his chances of winning, or he can lower his bid, and focus on winning against lower valuation type of his opponent. The optimal response to the strategic information will depend on the informativeness of 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 bidder exactly observe the valuation among 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 his 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.

Arya, Demski, and Glover (2003) consider an optimal auction environment with two bidders. The private information of each bidder consists of two elements, his own private valuation, and an improved estimate (relative to the prior *and* the posterior on the basis of his valuation alone) about the value of the competing bidders. The seller can design an optimal auction, but is a priori restricted to use as information only his prior and the report of each bidder about his own valuation. The seller is not allowed to use the additional information that bidder i has about the likelihood of valuations of bidder j . Otherwise, the environment is such that it would allow for full surplus extraction with dominant strategies as in Cremer and McLean (1988). They show that if the informational asymmetry regarding the distribution of values is small, then the full surplus extraction results still holds, but the strategies only form a Bayesian rather than a dominant strategy equilibrium. On the other hand, if the informational asymmetry is large, then full surplus extraction fails to hold, and a simple modified second-price auction is optimal.

Feinberg and Skrypacz (2005) pursue the logic of multidimensional types, in particular the separation between payoff types and belief types in the context of bargaining under incomplete information.

7 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 mechanism 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 estimate about the expected value of the competing bid is increasing over time. It is then natural whether a descending price auction might sometimes more favorable for information acquisition than an ascending price auction. The advantage of a descending price auction is that the possible bidder receives over time information that it is bid is more likely to be competitive as the competition apparently do not have very high valuation, or else they would already have stopped the clock. Interestingly, Klemperer (2002) suggest a sequential combination of English and Dutch auction to enhance entry and deter collusion. Yet a combination of English and Dutch auction could also be optimal to generate information and hence competition among the bidders.

We saw that the ex post efficient mechanisms may lead to excessive information acquisition in typical auction settings. We can then ask how we would have to modify the ex post efficient mechanism to achieve a second best solution. There are two natural modification of the ex post efficient allocation. 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 we correct the incentives to acquire information. For the given interim probability, we can then identify allocation which leads to the lowest losses in efficiency. Presumably, the solution will involve a large distortion to the allocation at the low end of the valuations as this will have the largest overall effect on transfers and incentives as they will influence all payments for the higher allocations.

The classic auctions such as auctions for art and wine, rely on a synchronization of demand and supply at a particular time and date. This has benefits for the aggregation of information and the informational efficiency of the auction. With the increased adoption of auction or auction like formats in electronic commerce, the synchronization aspect typically disappears and the arrival of supply and demand becomes decentralized and uncertain. This poses entirely new questions as to how auctions should be organized, in particular when and how long markets should be open and closed. This issue appeared at the forefront of auction theory perhaps first in the design of activity rules for the FCC auction. But in a sense, it is even more of an issue in electronic markets, where both the liquidity and the thickness of the market essentially contributes to the attractiveness of auctions and similar trading rules. Yet, until now, we do not have very good insight into dynamic mechanism design in the presence of asynchronous and decentralized trading, see Gallien (2005) and Gallien and Gupta (2005) for some recent work in this direction.

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. Yet, 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. Yet to formalize a notion of robustness, one ought to consider a set of possible beliefs.

Bergemann and Morris (2005c) 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 limit 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 with 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. The recent work by Laffont and Martimort (1997), (2000) on how collusion affects the principal in the presence of correlated information already establishes partial connections.

ACKNOWLEDGMENTS

This study was partially supported by the National Science Foundation Grant 0428422 and Yrjö Jahnsson's Foundation. We would like to thank Rahul Deb and Xianwen Shi for excellent research assistance. Bergemann would like to acknowledge joint research with Stephen Morris in the area of this survey. We would like to thank Michael Riordan and Sven DeVries for helpful conversations.

References

- ACQUISTI, A., AND H. VARIAN (2004): “Conditioning Prices on Purchase History,” *Marketing Science*, forthcoming.
- 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 Southern California.
- ARYA, A., J. DEMSKI, AND J. GLOVER (2003): “On the Optimality of Dominant Strategy Auctions under Severe Information Asymmetry,” Discussion paper, Ohio State University and University of Florida and Carnegie Mellon University.
- 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*.
- 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.
- (2004): “Robust Mechanism Design,” *Econometrica*, forthcoming.
- (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 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.
- 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.
- BOSE, S., E. OZDENOREN, AND A. PAPE (2004): “Optimal Auctions with Ambiguity,” Discussion paper, University of Texas, Austin and University of Michigan.
- 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*, forthcoming, Northwestern University and University of Bologna.
- CARARE, O., AND M. ROTHKOPF (2001): “Slow Dutch Auctions,” Discussion paper, University of Texas and Rutgers University.

- 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.
- (2001): “Auction and Information Acquisition: Sealed Bid or Dynamic Formats,” Discussion paper, ENPC and University College London.
- (2002): “Gathering Information Before Signing a Contract: A New Perspective,” Discussion paper, CERAS.
- (2004): “The Wait and See Option in Ascending Auctions,” *Journal of the European Economic Association*, 2, 494–503.
- 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.
- (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56, 1247–1258.

- CRESWELL, R. (1988): “Precontractual Investigation as an Optimal Precaution Problem,” *Journal of Legal Studies*, 17, 401–436.
- 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.
- DEWATRIPONT, M., AND J. TIROLE (1999): “Advocates,” *Journal of Political Economy*, 107, 1–39.
- 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 (2004): “Hierarchies of Belief and Interim Rationalizability,” Discussion paper, Northwestern University.
- 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.
- GALLIEN, J. (2005): “Dynamic Mechanism Design for Online Commerce,” *Operations Research*, forthcoming.
- GALLIEN, J., AND S. GUPTA (2005): “Temporary and Permanent Buyout Prices in Online Auctions,” Discussion paper, MIT.
- GANUZA, J.-J. (2004): “Ignorance Promotes Competition: An Auction Model of Endogeneous Private Valuation,” *RAND Journal of Economics*, 35.
- GERARDI, D., AND L. YARIV (2004): “Information Acquisition in Committees,” Discussion paper, Yale University.
- GERSHKOV, A., AND B. SZENTES (2004): “Optimal Voting Scheme with Costly Information Acquisition,” Discussion paper, University of Chicago.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin Expected Utility with Non-Unique Prior,” *Journal of Mathematical Economics*, 18, 141–153.
- GOLDBERG, A., J. HARTLINE, AND A. WRIGHT (2001): “Competitive Auctions and Digital Goods,” Intertrust Technical Report STAR-TR-99-01.
- GROMB, D., AND D. MARTIMORT (2004): “The Organization of Delegated Expertise,” Discussion Paper 4572, London Business School and University of Toulouse.
- 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.
- 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.

- JACKSON, M. (1991): “Bayesian Implementation,” *Econometrica*, 59, 461–477.
- 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 (2004): “The Limits of Ex Post Implementation,” Discussion paper, ENPC and University of Bonn.
- JOHNSON, R. (1979): “Auction Markets, Bid Preparation Costs and Entrance Fees,” *Land Economics*, 55, 313–318.
- KALAI, E. (2004): “Large Robust Games,” *Econometrica*, forthcoming.
- 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.
- KUNIMOTO, T. (2004): “Robust Implementation under Approximate Common Knowledge,” Discussion paper, Brown University.
- LAFFONT, J., AND D. MARTIMORT (1997): “Collusion under Asymmetric Information,” *Econometrica*, 65, 875–911.
- (2000): “Mechanism Design with Collusion and Correlation,” *Econometrica*, 65, 309–342.
- LAU, S. (2004): “Information and Bargaining in the Hold-Up Problem,” Discussion paper, Washington University.
- 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.
- (1997): “Information Management in Incentive Problems,” *Journal of Political Economy*, 105, 796–821.
- 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.
- LUCKING-REILEY, D. (1999): “Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet,” *American Economic Review*, 89, 1063–1080.
- 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.
- MARSCHAK, J., AND R. RADNER (1972): *Economic Theory of Teams*. Yale University Press, New Haven.
- 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 Acquisition 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 Acquisition 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): “Rational Expectations, Information Acquisition and Competitive Bidding,” *Econometrica*, 49, 921–943.
- 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, Boston University and Hebrew University.
- OTTAVIANI, M., AND A. PRAT (2001): “The Value of Public Information in Monopoly,” *Econometrica*, 69, 1673–1683.
- PALFREY, T., AND S. SRIVASTAVA (1989): “Implementation with Incomplete Information in Exchange Economies,” *Econometrica*, 57, 115–134.
- PARKES, D. (2004): “Auction Design with Costly Preference Elicitation,” *Annals of mathematics and AI*.
- 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.
- PRENDERGAST, C. (1993): “A Theory of ‘Yes Men’,” *American Economic Review*, 83, 757–770.
- ROGERSON, W. (1992): “Contractual Solutions to the Hold-Up Problem,” *Review of Economic Studies*, 59, 777–93.
- 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.
- TAYLOR, C. (2004): “Privacy and Information Acquisition in Competitive Markets,” Discussion paper, Duke University.
- 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.
- (1987): “Game-Theoretic Analyses of Trading Processes,” in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley, pp. 33–70, Cambridge. Cambridge University Press.
- YE, L. (2004): “Optimal Auctions with Endogenous Entry,” *Contributions to Theoretical Economics*, 4.
- 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.