Mechanism Design with Interdependent Valuations: Efficiency and Full Surplus Extraction

Claudio Mezzetti[1]

Department of Economics

University of North Carolina at Chapel Hill

email: mezzetti@email.unc.edu

February 18, 2003

[1]I would like to thank the co-editor and three anonymous referees for their helpful comments. This research was started when I was visiting the Department of Applied Mathematics of the University of Venice, whose financial support is gratefully acknowledged.

Abstract

Agents' valuations are interdependent if they depend on the signals, or types, of all agents. Previous literature has implicitly assumed that agents cannot observe their payoffs from a decision and has shown that with interdependent valuations and independent signals efficient design is impossible. This paper shows that if an agent observes his own decision payoff, then it is always possible to find efficient mechanisms. The paper also provides conditions under which it is possible to extract the full surplus from the agents.

Keywords: Auctions, Efficiency, Full Surplus Extraction, Information Acquisition, Interdependent Valuations, Mechanism Design.

1 Introduction

Consider a world in which a decision affecting several agents must be made (e.g., assets must be allocated). Each agent receives private signals (has private information) about his own characteristics, or type. Utilities are quasi-linear, the sum of a decision payoff and a monetary transfer. An agent's decision payoff depends on his own type, but not the types of the other agents; that is, there are no informational externalities. The seminal contributions of Vickrey (1963) and later Clarke (1971) and Groves (1973) showed that in such a world efficient decisions (the ones that maximize the sum of agent's utilities) can be implemented by using appropriate monetary transfers. The Vickrey-Clarke-Groves (VCG) mechanism accomplishes this by using transfers that make each agent the residual claimant of the social surplus and then cover any deficit with additional charges that do not depend on his own behavior.

In many practical instances the assumption of private values, or no informational externalities, is violated. Informational externalities, or interdependent valuations, are present if the payoff of an agent depends not only on his own type, but also on the types (or informational signals) of the other agents. Among the many possible examples of interdependent valuations, consider the following three situations. A seller has private information about the quality of a good or service that he is trying to sell to a buyer (Akerlof (1970), Spence (1973)); bidders in an auction to assign the mineral rights for a tract of land have private signals about the quantity of oil in the tract (Milgrom and Weber (1982)); an existing company is either being acquired by one of several rivals, or it is going to be split among them, and each rival has different information about the many business lines of the company.

Recently, Maskin (1992), Dasgupta and Maskin (2000), and Jehiel and Moldovanu (2001) have demonstrated, in increasing generality, that if informational signals are statistically independent, multidimensional (or, if they are single dimensional but a single crossing condition is violated), and there are informational externalities, then the efficient decision rule cannot be implemented by any *standard* mechanism:

incentive compatibility and efficiency are mutually exclusive (see also Bergemann and Välimäki (2002)). In these and all previous papers in the literature, agents report their types to the designer, as in a standard mechanism design problem with private values, but they do not report their (pre-monetary transfer) payoffs from the decision *after* a decision has been made. Implicitly, the literature has assumed that an agent cannot observe his own payoff from a decision. In this paper, I will make the opposite assumption. I will assume that agents observe their own decision payoffs. I believe that this assumption is satisfied in many important economic settings (see Section 2 for a discussion of this point). While with private values an agent cannot extract any new information from the observation of his payoff, with interdependent valuations observing his realized payoff provides the agent with new information about the types, or informational signals, of the other agents. The designer, then, should collect this information and use it. Restricting attention to standard mechanisms, with only type reports, is not without loss of generality when valuations are interdependent and agents observe their own payoffs. This insight is the starting point of this paper.

I will allow the mechanism designer to set up two reporting stages. In the first stage the designer asks about the agents' types. On the basis of these reports, a decision is selected. After the decision has taken effect, the designer asks the agents to report their realized payoffs in a second reporting stage. Then transfers are finalized that depend on reports in both stages. It turns out that allowing the transfers to depend on the payoff reports completely changes the conclusions of the model.

First, it is always possible to implement an efficient decision. To achieve this, the designer should implement the decision that is efficient given the signal reports of the agents in the first reporting stage. Each agent should be given as a transfer the sum of the payoffs reported by all other agents in the second reporting stage. This is sufficient to make each agent a residual claimant, and hence gives him the incentive to truthfully report his signals in the first reporting stage. As in a VCG mechanism, additional charges that do not depend on his reports can be imposed on each agent, so as to balance the budget.

Second, even though types are independent, with interdependent valuations the realized payoff of an agent is correlated with the types of the other agents. One may then conjecture that the designer can use a mechanism analogous to the Crémer and McLean (1985, 1988) mechanism to extract the full surplus from the agents. This conjecture turns out to be only partially true. The condition that guarantees that the designer can extract the full surplus, Condition 1, is indeed analogous to the condition introduced by Crémer and McLean (1988). However, while the Crémer and McLean condition is purely a restriction on the correlation among the probability distributions of the agents' types, and it is satisfied by almost all probability distributions, Condition 1 also restricts the possible payoff functions of the agents and it is not a generic property. In many instances full surplus extraction is not possible.

The paper is organized as follows. The next section introduces the model. Section 3 shows that it is always possible to implement efficient mechanisms. Section 4 studies full surplus extraction. Section 5 concludes.

2 Generalized Revelation Mechanisms

Consider a mechanism design model with n agents. Each agent has private information about his own type $\theta_i \in \Theta_i$, where Θ_i is a closed and bounded subset of $\mathbb{R}^{m_i}$; $\Theta = \times_{i=1}^{n} \Theta_i$ is the set of type profiles and $\theta = (\theta_1, \ldots, \theta_n)$ is a generic element of Θ. Let F_i and F_{-i} be the cumulative probability distributions of $\theta_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i} = \times_{j \neq i} \Theta_j$. Types are drawn independently across agents; that is, the θ_i's are independent random variables. (We already know from Crémer and McLean (1985, 1988), McAfee and Reny (1992), and more recently McLean and Postlewaite (2001), that efficiency and full surplus extraction are possible under general conditions when there is correlation of types across agents.) Let ω be the state of the world and $\Omega \subset \mathbb{R}^k$ be the set of possible states of the world. The state of the world is a random variable that depends on the agent's types; $\Pi(\omega|\theta)$ is the conditional cumulative distribution function of ω. Let X be the set of possible decisions, or outcomes (e.g., X could be a subset of an Euclidean space and represent the set of possible allocations of private and public goods). Agent i's utility function $U_i : X \times \Omega \times \mathbb{R} \rightarrow \mathbb{R}$ depends on the decision x, the state of the world ω and his monetary transfer t_i,

$$U_i(x, \omega, t_i) = u_i(x, \omega) + t_i. \tag{1}$$

As is common in the mechanism design literature, I will assume quasi-linear utility.

While with private values u_i depends on x and θ_i, it is standard to model interdependent valuations by postulating that an agent's decision payoff depends directly on the types of all the agents. (Gresik (1991) was one of the first to study a mechanism design problem with interdependent valuations.) There are two ways of reconciling my formulation with the standard model. First, by interpreting the decision payoff in the standard model as an expected payoff: in my model agent i's expected utility from the decision x, conditional on the type profile θ is

$$v_i(x, \theta) = \int_{\Omega} u_i(x, \omega) d\Pi(\omega|\theta).$$

Second, if there exists a function $\omega(\theta)$ such that $\Pi(\omega|\theta) = 0$ for $\omega < \omega(\theta)$ and $\Pi(\omega(\theta)|\theta) = 1$, then the state of the world is a deterministic function of the type profile and agents' decision payoffs depend directly on θ.

The standard model and my formulation are equivalent if an agent cannot observe his own decision payoff. However, when an agent can observe his own payoff my formulation is more general, because it allows for additional noise in the decision payoff. As I pointed out in the Introduction, contrary to the case of private values, with interdependent valuations the observation of the decision payoff conveys information to an agent. While previous literature has implicitly assumed that an agent cannot observe his own decision payoff, I will make the opposite assumption.

Assumption 1 *Given any decision x selected by the designer, and any realization of the state of the world ω, each agent i observes his realized decision payoff $u_i(x, \omega)$.*

I share the belief, expressed by a referee, that "this assumption is very often satisfied and that it is satisfied in many important economic settings, i.e., common value auctions for offshore oil tracts, or timber; Akerlof's lemon model; Spence's signaling model; exchange economies; public goods environments." The same referee provided the following example of a situation that violates the assumption. A watch is to be allocated between two agents. Agent 1 has no private information. Agent 2's private information consists of his opinion about how the watch looks: there are two possibility, "beautiful" and "ugly." Suppose that agent 1's payoff from obtaining the watch depends on agent 2's opinion. Then, even if he gets the watch, agent 1 will not observe his decision payoff, because he cannot observe agent 2's opinion.

In a *standard revelation mechanism* agents are not asked to report their payoffs from the decision. Under private values there is no loss of generality in assuming that the designer only uses standard revelation schemes. Intuitively, in a set-up with private values, observing one's own payoff conveys no new information to an agent and thus the designer has no need to collect second-stage messages. With interdependent valuations and observable decision payoffs, restricting the designer to use standard revelation schemes is not without loss of generality. In general, allowing the designer to collect any new information enlarges the set of implementable mechanisms. Thus, for example, if an agent observes several signals after the decision (e.g., he observes his revenue and cost) then the designer should ask the agent to report all these signals. In this paper I will assume that the agent only observes the aggregate payoff from the decision and I will study *generalized revelation mechanisms* in which the designer collects messages in two stages (this is without loss of generality, see Mezzetti (2002) for more details on the appropriate version of the revelation principle in this setup). The messages collected in the first stage determine the decision to be made. The second reporting stage takes place after the agents have observed their payoffs from the decision; messages from both stages are used to determine the total monetary transfers to the agents.

One may wonder if the chronological separation between the decision and the final transfers doesn't introduce obstacles to the practical implementation of the mechanism. Thus, suppose that a community has to decide on a public project. While the project must be built and financed today, the true payoffs from the project will only be revealed later, perhaps much later. Is this an insurmountable problem? I don't think so. The authority in charge of the project should collect type information and make a decision today; it should also collect fees today from the agents to finance the project. Over time, it should collect additional information about the realized payoffs and make additional transfers among the agents that reflect the new information. In fact, as we shall see in the next section, the additional transfers can be constructed so that their expected value is zero. Thus, while the mechanism I will propose to implement efficient decisions may not always be practical, it is not much less practical than the standard Vickrey-Clarke-Groves mechanism.

3 Efficiency

The deterministic decision rule $x^* : \Theta \to X$ is efficient if, for all $\theta \in \Theta$ it is

$$x^*(\theta) \in \arg\max_{x \in X} \int_\Omega \sum_{i=1}^n u_i(x, \omega) d\Pi(\omega|\theta),$$

or, equivalently

$$x^*(\theta) \in \arg\max_{x \in X} \sum_{i=1}^n v_i(x, \theta). \tag{2}$$

I will assume that (2) is always well defined.

With private values, and for sufficiently rich domains (e.g., a simply connected or convex set of valuations) the VCG mechanism is the only one that allows the designer to implement the efficient decision rule x^* (see Green and Laffont (1977), Holmström

(1979), Laffont and Maskin (1979) and, more recently, Williams (1999)). The VCG mechanism is a standard revelation mechanism in which agents are only asked to report their types. Let θ_i^r be the type reported by agent i. Up to a function $h_i(\theta^r)$ whose expected value conditional on θ_i is independent of θ_i^r, the VCG mechanism imposes the following transfers

$$\gamma_i(\theta^r) = \sum_{j \neq i} v_j(x^*(\theta^r), \theta^r). \tag{3}$$

Then, assuming that all other agents are truthfully reporting, at the reporting stage agent i solves the following maximization problem

$$\max_{\theta_i^r \in \Theta_i} \int_{\Theta_{-i}} \left[v_i(x^*(\theta_i^r, \theta_{-i}), \theta_i, \theta_{-i}) + \sum_{j \neq i} v_j(x^*(\theta_i^r, \theta_{-i}), \theta_i^r, \theta_{-i}) \right] dF_{-i}(\theta_{-i}).$$

Contrary to the case of private values, with interdependent valuations there is no reason why $\theta_i^r = \theta_i$ should be the solution to the problem. Thus, the standard VCG mechanism will not implement the efficient decision. The problem is that with interdependent valuations the functions γ_i fail to make each agent i the residual claimant of the full surplus. This is because the expected decision payoff of agent $j \neq i$, as computed by the designer, does not coincide with j's true decision payoff. The former depends directly on i's reported type θ_i^r, while the latter depends on the true type θ_i. With private values this problem does not arise, because j's computed decision payoff depends on the reported type θ_i^r only indirectly through the decision x^*. Maskin (1992) and Dasgupta and Maskin (2000) showed that efficient standard mechanisms can still be constructed if players types are single-dimensional, $\partial v_i / \partial \theta_i > 0$ for all i, and the "single crossing" condition $\partial v_i / \partial \theta_i > \partial v_j / \partial \theta_i$ holds for all i and $j \neq i$. (See also Ausubel (1997), Jehiel and Moldovanu (2001), Bergemann and Välimäki (2002), and Perry and Reny (2002).) On the other hand, Maskin (1992), Dasgupta and Maskin (2000) and Jehiel and Moldovanu (2001) showed, at different levels of generality, that with multidimensional types it is generically impossible to construct efficient standard mechanisms.

Now suppose that the designer uses a generalized revelation mechanism. Besides reporting a type θ_i^r in the first reporting stage, agent i faces a second reporting stage in which he must report a decision payoff u_i^r. Let the designer use the efficient decision rule x^*, which only depends on first-stage type reports, but suppose that the transfer function is (again, up to a function $h_i(\theta^r)$ whose expected value conditional on θ_i is independent of θ_i^r) :

$$\tau_i(\theta^r, u^r) = \sum_{j \neq i} u_j^r. \tag{4}$$

The idea, as in the standard VCG mechanism with private values, is to make every agent the residual claimant of the full surplus. Thus, we can think of this mechanism as a *generalized VCG mechanism*. To see that these transfers make truthtelling an equilibrium, first observe that the report of his decision payoff does not affect agent i's total utility - because τ_i does not depend on it - hence it is optimal for agent i to truthfully report his payoff in the second reporting stage. Then, suppose that all agents except i truthfully report their types, $\theta_{-i}^r = \theta_{-i}$, and their decision payoffs, while agent i of type θ_i falsely reports his type to be θ_i'. Under these hypotheses, the reported decision payoff of agent j is $u_j^r = u_j(x^*(\theta_i', \theta_{-i}), \omega)$; thus, for type θ_i the expected value of u_j^r is $v_j(x^*(\theta_i', \theta_{-i}), \theta_i, \theta_{-i})$. Note that this expected value depends on the implemented decision, which is a function of the reported types, and the true type profile. As a result agent i's total expected utility when the true type profile is (θ_i, θ_{-i}) and he reports θ_i' becomes

$$v_i(x^*(\theta_i', \theta_{-i}), \theta_i, \theta_{-i}) + \sum_{j \neq i} v_j(x^*(\theta_i', \theta_{-i}), \theta_i, \theta_{-i}). \tag{5}$$

By reporting his true type, on the other hand, agent i would obtain

$$v_i(x^*(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}) + \sum_{j \neq i} v_j(x^*(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}). \tag{6}$$

Since x^* is the efficient decision, the utility in (6) is at least as great as the utility in (5). Hence, agent i will never profit from falsely reporting θ_i'; truthful reporting is a best reply to the truthful reporting of all the other agents. Thus, I have proved the following proposition.

Proposition 1 *It is always possible to construct an efficient perfect Bayesian mechanism.*

While with private values it is possible to make truthful revelation a dominant strategy for agents, with interdependent valuations the dominant strategy is lost. However, in the proposed mechanism telling the truth is a best reply for agent i independently of his beliefs about the other players. That is, telling the truth is an ex-post equilibrium: it remains a perfect Bayesian equilibrium for any prior distribution over types.

With private values the type θ_i, learned before participating in the mechanism (the interim stage), is both the only piece of private information and all that agent i needs to determine his valuation for all decisions. If valuations are interdependent, knowing θ_i is not sufficient to determine i's valuation. If the decision payoff is observed, however, an agent obtains two pieces of information: his interim type θ_i and his decision payoff type u_i, which he learns after a decision has been made (the ex-post stage). Thus, we could say that an agent's valuation for the decision that was chosen (but not necessarily all decisions) is ex-post privately known. In a sense, observing the decision payoff brings the interdependent valuations model closer to the private values model.

Telling the truth is not the unique equilibrium of the generalized VCG mechanism. This is a problem shared with the standard VCG mechanism, and it is an open question whether the tools from the literature on full Bayesian implementation could be used to construct more complex mechanisms without the untruthful equilibria.

As I mentioned earlier, with private values and sufficiently rich domains, an efficient mechanism must be a standard VCG mechanism. As it will become clear in the next section, with interdependent valuations and observable payoffs there are other mechanisms, besides the generalized VCG, that yields efficient outcomes.

Since the transfers to all bidders depend on the realized payoffs, we can think of the generalized VCG mechanism as containing contingent payments. That contingent payments are valuable tools has been pointed out before. For example, in a private values setup, Hansen (1985) and Crémer (1987) (see also Samuelson (1987)) showed that if the value of a target firm to the winning bidder in an auction becomes publicly known ex-post, then the seller can raise its revenue by using contingent payments, as opposed to cash auctions. The schemes proposed by Hansen and Crémer are not VCG mechanisms. In this paper, beside considering a much more general setup, I do not require that information becomes public, but I rely instead on the agents' reports of their own realized payoffs. Thus, the payments in the generalized VCG mechanism are contingent on the reported payoffs, not on publicly observable and verifiable payoffs.

The transfers τ_i are made after agents have observed their own payoffs. If, for some reason, the designer needed to make transfers at the decision stage, he could require that the transfers γ_i, as defined in equation (3), be made at the same time a decision is made, and that the transfer adjustments $\tau_i - \gamma_i$ are made after decision payoffs are observed. In equilibrium the expected value of the transfer adjustments is zero; in fact all transfer adjustments are exactly equal to zero if the state of the world is a deterministic function of the type profile.

So far, I have not addressed the issue of balancing the budget. However, it is simple to show, using techniques first introduced by D'Aspremont and Gérard-Varet (1979) and Arrow (1979), that the designer can balance the budget. First, note that it is always possible to make sure that the designer collects positive revenue, by selecting the following transfer functions:

$$t_i(\theta^r, u^r) = \sum_{j \neq i} u_j^r - \max_{\theta_i \in \Theta_i,\, x \in X} \sum_{j \neq i} v_j(x, \theta_i, \theta_{-i}^r). \tag{7}$$

With these transfers, the mechanism is similar to the so-called pivot scheme in the mechanism design literature with private values (a generalization of the Vickrey auction); the agent pays for the highest possible externality he causes to others (here the maximum is taken not only over the decision, as in the case of private values, but also over agent i's type).

Requiring that the budget balance (i.e., that the transfers add up to zero) is more restrictive, but it is the appropriate property if the designer is a mediator - helping agents to coordinate - and not an agent himself (e.g., an auctioneer). Let E_{-i} be the expectation operator over the random variable θ_{-i} (with $E_{-(n+1)} = E_{-1}$) and E be the expectation over θ. Consider the following additional charge h_i on agent i

$$\begin{aligned} h_i(\theta^r) \quad = \quad & \frac{n-1}{n} \left\{ \sum_{j=1}^{n} v_j(x^*(\theta^r), \theta^r) - E_{-i}\left[\sum_{j=1}^{n} v_j(x^*(\theta_i^r, \theta_{-i}), \theta_i^r, \theta_{-i}) \right] + \right. \\ & \left. + E_{-(i+1)} \left[\sum_{j=1}^{n} v_j(x^*(\theta_{i+1}^r, \theta_{-(i+1)}), \theta_{i+1}^r, \theta_{-(i+1)}) \right] \right\} \end{aligned} \tag{8}$$

With this additional charge, agent i's transfer becomes $t_i = \tau_i - h_i$. Since the expected value of h_i does not depend on the reports of agent i, truthful reporting remains an equilibrium. It is simple to check that on the equilibrium path (i.e., for $\theta = \theta^r$) the budget will be balanced, $\sum_{i=1}^{n} t_i = 0$.

Proposition 2 *It is always possible to construct an efficient, budget balancing, perfect Bayesian mechanism.*

I now present two examples that show how the results in this section can be applied. The examples are purposefully simple, to highlight the main ideas.

Example 1. A seller knows the quality θ of a durable good, with $\theta \in [-1, 1]$. The quality of the good is unknown to the buyer. A good of quality θ is worth θ to the buyer and it is worth $\alpha\theta$ to the seller, where $\alpha < 1$. Efficiency dictates that the buyer get the good if $\theta > 0$ and that the seller keep it if $\theta < 0$. In a standard mechanism design model the transfers and the decision depend on the players' reports about their types. Let t_s be the seller's transfer and suppose that the decision rule is efficient. Letting $\theta' < 0 < \theta''$, incentive compatibility requires $\alpha\theta' + t_s(\theta') \geq t_s(\theta'')$ and $t_s(\theta''_s) \geq \alpha\theta'' + t_s(\theta')$, which implies $\theta' \geq \theta''$, contradiction. Thus, no efficient standard mechanism exists. On the other hand, efficiency can be obtained by using a generalized revelation mechanism. Consider a contract between buyer, seller and a dealer (the designer) stipulating that (1) if the seller reports to the dealer that the quality of the good is $\theta^r > 0$, then the dealer sells the good to the buyer at a price $p = \beta\theta^r$, with $\alpha \leq \beta \leq 1$; (2) after acquiring the good the buyer will publicly report her payoff u^r and the dealer will pay the seller an amount equal to βu^r. Under this contract, the seller has an incentive to truthfully report the good's quality and the buyer has an incentive to report his true payoff from the decision; the outcome is efficient and the buyer ends up paying a price $\beta\theta$ to the seller; the dealer breaks even. This mechanism has some flavor in common with the dynamic model in Hendel and Lizzeri (2002). They studied the markets for new and used cars and showed that first best efficiency can be obtained with a mechanism in which higher valuation buyers lease new cars in each period and then report their quality to the dealer when returning them at the end of the lease. Lower valuation buyers purchase off-lease cars.

Example 2. An existing business is up for sale; the state of the world consists of the random variables ω_h, $h = 1, ..., 4$, representing the profitability of four different business lines. The potential acquirers are two firms. Firm i's payoff from acquiring the business is $u_i = \alpha_i\omega_i + \sum_{h=1}^{4} \omega_h$, where $\alpha_i > 0$. Firm i privately observes signal $\theta_i = (s_i, s_{i+2}) \in \mathbb{R}^2_+$, with $\omega_h = s_h + \eta_h$, where η_h is a zero-mean random variable.

Efficiency requires that firm 1 acquire the business if $\alpha_1 s_1 > \alpha_2 s_2$. As in Example 1, it is simple to see that there is no efficient standard mechanism. For example, let $s_3' - s_3'' > (1 + \alpha_1)(s_1'' - s_1') > 0$. Then, the standard incentive constraints of types $\theta_1' = (s_1', s_3')$ and $\theta_1'' = (s_1'', s_3'')$ are mutually inconsistent when the designer is trying to implement the efficient decision. The following generalized mechanism, however, implements the efficient decision and yields positive expected revenue. The designer (a merchant banker) issues a security that pays an amount equal to the payoff from the business, as reported by the winning firm after the acquisition. Each firm is asked to submit an estimate e_i. The merchant banker assigns the business to the firm i which submits the highest estimate e_i and gives the security to the other firm. Firm i pays $(1 + \alpha_j)\overline{\omega}_j + \overline{\omega}_{j+2}$, $j \neq i$, where $\overline{\omega}_h$ is the ex-ante expected value of the random variable ω_h. Since reporting the true payoff from the business is optimal after the acquisition, reporting an estimate $e_i = \alpha_i s_i$ is an equilibrium; the efficient outcome is implemented. Corporate acquisitions using security exchanges are common in practice (e.g., see, Schwert (2000)).

It is interesting to observe that in both examples the agents will want to participate in the mechanism. The working paper version, Mezzetti (2002), gives a necessary and sufficient condition for a generalized VCG mechanism to be budget balancing and individually rational, that is, to induce voluntary participation. The condition is analogous to the one Makowski and Mezzetti (1994) provided for the case of private values.

If types are single dimensional and a single-crossing condition is satisfied, then standard revelation mechanisms that implement the ex-post efficient decision rule exist even if valuations are interdependent. However, Bergemann and Välimäki (2002) showed that no such mechanism provides agents with the incentives for efficient ex-ante information acquisition (see also Maskin (1992)). It is simple to show that this inefficiency disappears if the mechanism designer is allowed to condition transfers on the players' reports of their realized payoffs. This is because the generalized VCG mechanism that I introduced in this section makes each agent the residual claimant of the full surplus and thus it also provides each agent with the incentives for the ex-ante efficient acquisition of information. See Mezzetti (2002) for details.

4 Full Surplus Extraction

Even if types are independent, with interdependent valuations observing his own decision payoff $u_i(x, \omega)$ provides agent i with a signal that is correlated with the types θ_{-i} of the other agents. Thus, it is natural to ask whether the designer (e.g., the seller in an auction) can exploit this correlation and fully extract the surplus, as in Crémer and McLean (1985, 1988) (see also McAfee and Reny (1992) and McLean and Postlewaite (2001)).

To facilitate comparison with Crémer and McLean, I will restrict attention to the special case of a single object to be assigned to one of the n agents, so that a decision $x = (x_1, ..., x_n)$ is a probability vector; x_i is the probability that agent i gets the object. Agent i's decision payoff is $x_i u_i(\omega)$. The type sets Θ_i and the set of states of the world Ω are finite; hence the set U_i of feasible decision payoffs for agent i is also finite, $U_i = \{u_i \in \mathbb{R} : \exists \omega \in \Omega \text{ s.t. } u_i(\omega) = u_i\}$. The probability of type θ_i is $f_i(\theta_i)$, while $\pi(\omega|\theta)$ is the probability of ω conditional on θ, and $\pi(\omega|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \pi(\omega|\theta_i, \theta_{-i}) f_{-i}(\theta_{-i})$ is the probability of ω conditional on θ_i, where $f_{-i}(\theta_{-i}) = \prod_{j \neq i} f_j(\theta_j)$.

Let $I^*(\theta)$ be the set of agents i such that $v_i(\theta) \geq v_j(\theta)$ for all j (i.e., the agents with the highest expected payoff, given θ), where $v_i(\theta) = \sum_{\omega \in \Omega} \pi(\omega|\theta) u_i(\omega)$, and $\#I^*(\theta)$ be the number of agents in $I^*(\theta)$. I will consider the efficient decision function $x^*(\theta)$ with $x_i^*(\theta) = 1/\#I^*(\theta)$ if $i \in I^*(\theta)$ and $x_i^*(\theta) = 0$ if $i \notin I^*(\theta)$. For all j, $\Omega(u_j)$ is the set of states of the world ω for which $u_j(\omega) = u_j$, and $\pi(u_j, \theta_{-i}, \theta_i) = \sum_{\omega \in \Omega(u_j)} \pi(\omega|\theta_i, \theta_{-i})$ is the probability that $u_j(\omega) = u_j$, conditional on θ_{-i} and θ_i.

The condition that ensures that the seller can fully extract the surplus is close in spirit to the condition in Theorem 2 of Crémer and McLean (1988).

Condition 1 *(a) For all i, all $j \neq i$ and all $\theta_{-i} \in \Theta_{-i}$, there is no type θ'_i and a family $\{\rho_i(\theta_i)\}_{\theta_i \neq \theta'_i}$ such that $\rho_i(\theta_i) \geq 0$ for all $\theta_i \neq \theta'_i$ and $\pi(u_j, \theta_{-i}, \theta'_i) = \sum_{\theta_i \neq \theta'_i} \rho_i(\theta_i)\pi(u_j, \theta_{-i}, \theta_i)$ for all $u_j \in U_j$. (b) For all i there is no type θ'_i such that $x^*_i(\theta'_i, \theta_{-i}) = 1$ for all θ_{-i}.*

In the mechanism proposed by Crémer and McLean (1985, 1988), the seller fully extracts the surplus by using an efficient auction (e.g., a first price, or a Vickrey auction), augmented with a lottery for each type of each agent. The lottery stipulates that an agent must make additional payments to the seller that depend on the types reported by the other agents. A similar approach can be followed when valuations are interdependent and decision payoffs are observable, even if the types θ_i are independent. Given the agents' reported type profile, the designer assigns the object according to the efficient decision rule. The agent who receives the object is asked to pay his expected payoff given the reported types. All other agents must pay an amount that depends on the reported decision payoff of the winning agent. If Condition 1 holds, this payments can be structured, as in Crémer and McLean, so that their expected value is zero if an agent truthfully reports his type and it is arbitrarily small if an agent lies. The important difference with Crémer and McLean is that this mechanism exploits the correlation between the types of the agents and the decision payoff of the winning agent (rather then the correlation among the types of all agents).

Proposition 3 *If Conditions 1 holds, then there is a perfect Bayesian equilibrium of a generalized revelation mechanism in the auction model in which the seller fully extracts the surplus from the agents.*

Proof. I will only sketch the proof, since it follows closely the proof of Theorem 2 in Crémer and McLean (1988). By Condition 1 and Farkas' Lemma, for all i, all $j \neq i$, and all θ_{-i} there exists a family $\{g(u_j, \theta_{-i}, \theta_i)\}_{u_j \in U_j, \theta_i \in \Theta_i}$ such that $\sum_{u_j \in U_j} g(u_j, \theta_{-i}, \theta_i)\pi(u_j, \theta_{-i}, \theta_i) > 0$ and $\sum_{u_j \in U_j} g(u_j, \theta_{-i}, \theta_i)\pi(u_j, \theta_{-i}, \theta'_i) \leq 0$ for all $\theta_i' \neq \theta_i$. Then consider the generalized revelation mechanism that uses the efficient decision x^* and the following, full surplus extracting, transfers to the agents: $t_{i^*}(\theta^r, u^r) = -v_{i^*}(\theta^r)$ for the agent i^* who gets the object (the realized draw from $I^*(\theta^r)$) and $t_i(\theta^r, u^r) = -\beta \left[\sum_{u_{i^*} \in U_{i^*}} g(u_{i^*}, \theta_{-i}^r, \theta_i^r)\pi(u_{i^*}, \theta_{-i}^r, \theta_i^r) - g(u_{i^*}^r, \theta_{-i}^r, \theta_i^r)\right]$ for an agent i who does not get the object. Truthful reporting of his payoff is optimal for the winning agent, and by choosing β sufficiently large truthful type reporting is optimal for all agents. ■

While the Crémer and McLean condition is purely a restriction on the correlation among the probability distributions on Θ_{-i}, Condition 1 also restricts the possible payoff functions of the agents. As a consequence, while the Crémer and McLean condition is generically satisfied (i.e., it is satisfied by almost all probability distributions on Θ), Condition 1 is not a generic property. First, Condition 1(b), which has no counterpart in Crémer and McLean, does not hold if there is a type θ_i' such that $x_i^*(\theta_i', \theta_{-i}) = 1$ for all $\theta_{-i} \in \Theta_{-i}$. Intuitively, if agent i could make sure to always get the object (by reporting type θ_i'), then his type report could not be statistically cross-checked (the payoff reports of the other agents, being equal to zero, would contain no information); in this case agent i may be able to obtain a positive payoff by reporting θ_i', while his true type is θ_i. Second, and most important, Condition 1(a) also imposes an implicit restriction on the payoff functions of the agents. It cannot be the case, for example, that the payoff distribution $\pi(u_j, \theta_{-i}, \theta_i)$ of agent j given θ_{-i} is the same for two different types θ_i' and θ_i'' of agent i.

Consider Example 2. First, suppose that s_3, s_4, ω_3 and ω_4 are constants, equal to zero, that $s_i = \theta_i \in \{1, 4\}$ for $i = 1, 2$, with equal probability, and that the random variables η_i, $i = 1, 2$, can also take two equally likely values, $\eta_i \in \{-1, 1\}$. Then $\omega_i = \{0, 2, 3, 5\}$, the set U_i contains 16 elements, and we can think of $\pi(u_j, \theta_{-i}, \theta_i)$ as a vector in $\mathbb{R}^{16}$. If $\alpha_i = \alpha$ for $i = 1, 2$, then Condition 1 requires that there is no non-negative constant ρ such that $\pi(u_j, \theta_{-i}, \theta_i = 1) = \rho\pi(u_j, \theta_{-i}, \theta_i = 4)$. It is easily checked that this holds and hence the designer (merchant banker) can fully extract the surplus. For example, for $j = 1$, we have $\pi(u_1 = 0, \theta_1 = 1, \theta_2 = 1) = 1/4$ and $\pi(u_1 = 7 + 2\alpha, \theta_1 = 1, \theta_2 = 1) = 0$, while $\pi(u_1 = 0, \theta_1 = 1, \theta_2 = 4) = 0$ and $\pi(u_1 = 7 + 2\alpha, \theta_1 = 1, \theta_2 = 4) = 1/4$. On the other hand, if $\alpha_1 > 4\alpha_2$, then it is efficient that firm 1 always acquire the business, Condition 1(b) fails, and full surplus extraction is not possible (firm 1 would never let the designer know that $s_1 = 4$). Now suppose that all the signals s_h, $h = 1, ...4$, can take the two values 1 and 4, with equal probability, and that the variables η_h are equally likely to be -1 and 1. In this case $\theta_i = (s_i, s_{i+2})$ and the two vectors $\pi(u_j, \theta_j, \theta_i = (1, 4))$ and $\pi(u_j, \theta_j, \theta_i = (4, 1))$ are equal. Condition 1(a) fails and full surplus extraction is impossible, because two types of player i, $\theta_i = (4, 1)$ and $\theta_i' = (1, 4)$, have different payoffs, but generate the same distribution over the payoffs of the opponent. In particular, type $\theta_i = (4, 1)$ can guarantee himself positive rents by pretending to be type $\theta_i' = (1, 4)$. It is interesting to stress the connection with the work of Neeman (2002). In a classic Crémer and McLean world of private values and correlated types, he was the first to show that if an agent's beliefs (probability distributions over the other agents' types) do not uniquely determine the agent's payoff, then full surplus extraction is not possible. With interdependent valuations and independent types, as in this paper, it is often the case that an agent's types are not in a one to one correspondence with the signals extracted from the payoff reports of the other agents.

In the mechanism described in Proposition 3, as well as in Crémer and McLean, agents face lotteries that leave them with zero expected utility. I will now show that in some cases it is possible to fully extract the surplus ex-post (i.e., for all type realizations). Consider again the general model, but with the space of states of the world Ω equal to the type space Θ, so that $u_i(x, \omega) = u_i(x, \theta)$. To fully extract the surplus the designer should use the efficient decision and collect from each agent i a payment equal to $u_i(x^*(\theta^r), \theta^r)$. If all agents with the possible exception of agent i truthfully report their types and decision payoffs, then for all $j \neq i$ the reported decision payoff will be $u_j(x^*(\theta_i^r, \theta_{-i}), \theta_i, \theta_{-i})$. On the other hand, on the basis of the type reports, and on the assumption that all agents are being truthful, the designer would have predicted a decision payoff equal to $u_j(x^*(\theta_i^r, \theta_{-i}), \theta_i^r, \theta_{-i})$. Any difference between an agent's reported and predicted payoff provides the designer with proof that some other agent lied. If the designer imposes severe penalties ("shoots the liars") on all but the agent uncovering a lie, then agents will be induced to truthfully report, provided they face a positive probability that their lies are discovered.

Condition 2 *The agents' payoff functions u_i are bounded. $\Omega = \Theta$, and there exists $\delta > 0$ such that, for all x in the range of the efficient decision function x^*, if $\theta_i \neq \theta_i'$, then there exists a $j \neq i$ such that $u_j(x, \theta_i, \theta_{-i}) \neq u_j(x, \theta_i', \theta_{-i})$ for all θ_{-i} belonging to a subset of Θ_{-i} having probability measure not smaller than δ.*

Condition 2 requires strong correlation between observed decision payoffs and agents' types: if an agent lies, then the observed payoffs will be inconsistent with the reported types with positive probability. Condition 2 could be weakened, see Mezzetti (2002).

Proposition 4 *If Conditions 2 holds, then there is a perfect Bayesian equilibrium of a generalized revelation mechanism in which the designer fully extracts the surplus from the agents ex-post (i.e., for all type realizations).*

Proof. Consider a generalized revelation mechanism that uses an efficient decision rule. Let the transfer function be: $t_i(\theta^r, u^r) = -v_i(x^*, \theta^r)$ if $u_j^r = u_j(x^*(\theta^r), \theta^r)$ for all $j \neq i$, and $t_i(\theta^r, u^r) = -L$ if $u_j^r \neq u_j(x^*(\theta^r), \theta^r)$ for at least one $j \neq i$. Suppose that all the other agents truthfully report their types and decision payoffs. Since agent i's transfer does not depend on his reported decision payoff, he has an incentive to truthfully report in the second stage. Furthermore, if agent i of type θ_i truthfully reports his type, then he gets a zero total utility. On the other hand, if he reports type $\theta_i^r \neq \theta_i$, then, by Condition 2, he will have to pay L with probability bounded away from zero. It is then possible to choose L large enough to deter any type misreport. ■

In Example 2, Condition 2 holds if (i) η_h, $h = 1, ...4$, are constants, so that $\Omega = \Theta$ and payoffs are deterministic functions of the signals (i.e., firm i has full knowledge of the profitability of business lines i and $i+2$), (ii) it is not the case that the efficient decision is for one firm to acquire the business under all circumstances, and (iii) s_3 and s_4 are constant, so that the profitability of business lines 3 and 4 is common knowledge. If this latter condition is violated, then full surplus extraction is not possible. A type $\theta_i = (s_i, s_{i+2})$ could report $\theta'_i = (s'_i, s'_{i+2})$ such that $s'_i < s_i$ and $s'_i + s'_{i+2} = s_i + s_{i+2}$. This way, if firm i wins, it pays less than its expected payoff; if it loses, the lie will go undetected.

5 Conclusions

Section 3 shows that when agents observe their own decision payoffs, even if these payoffs are unverifiable, payments that are contingent on payoff reports allow the implementation of efficient decisions. This suggests that any contractual scheme, or institutional arrangement, that facilitates the use of contingent payments may raise efficiency. It also helps us understand why, as pointed out by Samuelson (1987), "contingent pricing schemes are common in actual practice, where examples range from corporate acquisition via exchange of securities, to revenue sharing in oil lease auctions, and incentive contracts in defense procurement."

The results on surplus extraction in Section 4 show the connection, but also some important differences, between the model with correlated types and the model with independent types and interdependent valuations. The most important difference with Crémer and McLean (1988) is that full surplus extraction is not always possible, because agents with different payoff relevant types may naturally generate the same distribution of the opponents' payoffs.

It is natural to ask to what extent the approach of exploiting the ex-post observation of payoffs can be extended to general allocation problems with agents' utilities that are not quasi-linear. To do so would require to decompose the final allocation decision in two (or more) stages, with agents observing their payoffs and making reports at the end of each stage. With interdependent valuations, this certainly expands (at least weakly) the range of implementable final allocations, but the amount of extra freedom that multiple reporting stages give the designer in this general environment is not clear. On the negative side, the generalized VCG schemes defined in Section 3 rely on quasi-linearity, and with private values it is well known that efficient allocations cannot generally be implemented if utilities are not quasi-linear. On the positive side, ex-post payoff reports allow some cross-checking and it is easy to construct examples where this is enough to implement efficient allocations that could not be implemented with mechanisms that only use a single reporting stage. Further research is needed.

References

[1] Akerlof, G.A., 1970: "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84, 488-500.

[2] Arrow, K., 1979: "The Property Rights Doctrine and Demand Revelation under Incomplete Information," in M. Boskin ed., *Economics and Human Welfare*, New York, Academic Press.

[3] Ausubel, L., 1997: "An Efficient Ascending-Bid Auction for Multiple Objects," http://www.ausubel.com/auction-papers.htm.

[4] Bergemann, D. and J. Välimäki, 2002: "Information Acquisition and Efficient Mechanism Design," *Econometrica, 70, 1007-33.*

[5] Clarke, E., 1971: "Multipart Pricing of Public Goods," *Public Choice, 11, 17-33.*

[6] Crémer, J., 1987: "Auctions with Contingent Payments: Comment," *American Economic Review*, 77, 746.

[7] Crémer, J. and R. McLean, 1985: "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, 53, 345-61.

[8] Crémer, J. and R. McLean, 1988: "Full Extraction of Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247-57.

[9] Dasgupta, P., and E. Maskin, 2000: "Efficient Auctions," *Quarterly Journal of Economics*, 115, 341-389.

[10] D'Aspremont, L., and L. Gérard-Varet, 1979: "Incentives and Incomplete Information," *Journal of Public Economics*, 11, 25-45.

[11] Green, J. and J.J. Laffont, 1977: "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods," *Econometrica*, 45, 427-38.

[12] Gresik, T., 1991: "Ex ante Incentive Efficient Trading Mechanisms without the Private Valuation Restriction," *Journal of Economic Theory, 55, 41-63.*

[13] Groves, T., 1973: "Incentives in Teams," *Econometrica*, 41, 617-31.

[14] Hansen, R.G., 1985: "Auctions with Contingent Payments," *American Economic Review*, 75, 862-65.

[15] Hendel, I and A. Lizzeri, 2002: "The Role of Leasing under Adverse Selection," *Journal of Political Economy*, 110, 113-44.

[16] Holmström, B., 1979: "Groves Schemes on Restricted Domains," *Econometrica*, 47, 1137-44.

[17] Jehiel, P., and B. Moldovanu, 2001: "Efficient Design with Interdependent Valuations," *Econometrica*, 69, 1237-59.

[18] Laffont, J.J. and E. Maskin, 1979: "A Differential Approach to Expected Utility Maximizing Mechanisms," in J.J. Laffont, ed., *Aggregation and Revelation of Preferences*, Amsterdam: North Holland, 289-308.

[19] Makowski, L. and C. Mezzetti, 1994: "Bayesian and Weakly Robust First Best Mechanisms: Characterizations," *Journal of Economic Theory*, 64, 500-19.

[20] Maskin, E., 1992: "Auctions and Privatization," in H. Siebert, ed., *Privatization*, Kiel: Institut fuer Weltwirtschaft der Universitaet Kiel, 115-36.

[21] McAfee, P.R. and P. Reny, 1992: "Correlated Information and Mechanism Design," *Econometrica*, 60, 395-421.

[22] McLean, R. and A. Postlewaite, 2001: "Efficient Auction Mechanisms with Interdependent and Multidimensional Signals," http://www.ssc.upenn.edu/~apostlew/.

[23] Mezzetti, C., 2002: "Auction Design with Interdependent Valuations: The Generalized Revelation Principle, Efficiency, Full Surplus Extraction and Information Acquisition," http://www.unc.edu/˜mezzetti/.

[24] Milgrom, P.R., and R. Weber, 1982: "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089-1122.

[25] Neeman, Z., 2002; "The Relevance of Private Information in Mechanism Design," http://econ.edu.bu/neeman/.

[26] Perry, M., and P. Reny, 2002: "An Efficient Auction," *Econometrica*, 70, 1199-1212.

[27] Samuelson, W., 1987: "Auctions with Contingent Payments: Comment," *American Economic Review*, 77, 740-45.

[28] Schwert, G.W., 2000: "Hostility in Takeovers: In the Eyes of the Beholder?" *Journal of Finance*, 55, 2599-640.

[29] Spence, A.M., 1973: "Job Market Signaling," *Quarterly Journal of Economics*, 87, 355-74.

[30] Vickrey, W., 1961: "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8-37.

[31] Williams, S.R., 1999: "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms," *Economic Theory*, 14, 181-201.