

MECHANISM DESIGN WITH INTERDEPENDENT
VALUATIONS: EFFICIENCY

CLAUDIO MEZZETTI¹

Agents' valuations are interdependent if they depend on the signals, or types, of all agents. Under the implicit assumption that agents cannot observe their outcome-decision payoffs, previous literature has shown that with interdependent valuations and independent signals efficient design is impossible. This paper shows that an efficient mechanism exists in an environment where first the final outcome (e.g., allocation of the goods) is determined, then the agents observe their own outcome-decision payoffs, and then final transfers are made.

KEYWORDS: Auctions, Efficiency, Interdependent Valuations, Mechanism Design.

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.

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 payoff from an outcome decision and a monetary transfer. An agent's outcome-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 an efficient decision (one that maximizes the sum of agent's payoffs) can be achieved by using appropriate monetary transfers. A Groves mechanism accomplishes this by using transfers that first 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)); in a mineral-rights auction bidders have private signals about the value (e.g., amount 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 Ausubel (1997), Bergemann and Välimäki (2002), and Perry and Reny (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 outcome decision *after* a decision has been made. Implicitly, the literature has ruled out the possibility of transfers after each agent has observed his own payoff from the outcome decision.

At one extreme, one might allow the mechanism designer to make a temporary outcome decision (e.g., allocation of goods), wait until the agents experience utility, and then, based on subsequent reports, determine a final outcome and transfers. The literature has focused on the other extreme, where a final outcome and transfers must be determined initially. I consider the intermediate case, where a final outcome must be determined initially, but transfers can be made after subsequent reports. More precisely, 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, an outcome is selected. After the outcome 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. It is always possible to implement an efficient decision by using the following *generalized* (or *two-stage*) *Groves mechanism*. First, the designer implements the outcome that is efficient given the signal reports of the agents in the first reporting stage. Then, each agent is given as a transfer the sum of the outcome-decision 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 Groves mechanism, additional charges that do not depend on his reports can be imposed on each agent, so as to balance the budget.

Since the transfers to all bidders depend on the realized outcome-decision payoffs, we can think of the generalized Groves 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, then the seller can raise its revenue by using contingent payments, as opposed to cash auctions. 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 Groves mechanism are contingent on the reported outcome-decision payoffs, not on publicly observable and verifiable payoffs.²

The paper is organized as follows. The next section introduces the model. Section 3 shows that it is always possible to achieve efficient decisions. Section 4 concludes by arguing that mechanisms with two reporting stages can also be used to achieve goals different from efficiency (e.g., surplus extraction).

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 $\theta = (\theta_1, \dots, \theta_n)$ is a generic element of Θ . Let F_i and F_{-i} be the cumulative probability distributions of $\theta_i \in \Theta_i$ and

² The generalized Groves 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.

$\theta_{-i} \in \Theta_{-i} = \times_{j \neq i} \Theta_j$. 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 non-monetary outcomes (also referred as decisions, or outcome decisions). For example, 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 outcome x , the state of the world ω and his monetary transfer t_i ,

$$(1) \quad U_i(x, \omega, t_i) = u_i(x, \omega) + t_i.$$

As is common in the mechanism design literature I will assume quasi-linear utility; I will refer to u_i as the outcome-decision payoff. Agent i 's expected payoff from outcome x , conditional on θ is

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

While with private values u_i depends on x and θ_i , it is standard in the literature on interdependent valuations to use a reduced-form model in which the state variable ω is suppressed and an agent's outcome-decision payoff depends directly on the types of all the agents. There are two ways of reconciling the state-of-the-world formulation with the reduced-form model. First, by interpreting θ as a vector of signals and $v_i(x, \theta)$ as the outcome-decision payoff. 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' payoffs depend directly on θ .

The reduced-form and the state-of-the-world formulation are equivalent if an agent cannot observe his own outcome-decision payoff. However, when an agent can observe his own outcome-decision payoff before final transfers are made, as I will assume, the state-of-the-world formulation is more general, because it allows for additional noise in the payoff.

In a *standard revelation mechanism* agents are not asked to report their outcome-decision payoffs. 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 outcome-decision payoffs, restricting the designer to use standard revelation schemes entails a loss of generality. In general, allowing the designer to collect any new information enlarges the set of implementable decision functions. Thus, for example, if an agent observes several signals after the outcome 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 outcome and I will study *generalized* (or *two-stage*) *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 outcome decision to be made. The second reporting stage takes place after the agents have observed their payoffs from the outcome decision; messages from both stages are used to determine the total monetary transfers to the agents.

ASSUMPTION 1: *For any $x \in X$ and any realization of the state of the world ω , each agent i observes his realized outcome-decision payoff $u_i(x, \omega)$ after the final outcome decision, but before final transfers, are made.*

This assumption is quite plausible in many important economic setting. For example, in auctions for timber and other commodities the winning bidder eventually

learns the market value of the goods being sold. In a used car market (as in Akerlof's lemon model) the buyer learns the car's quality. In a labor market (as in Spence's signaling model), employers learn the quality of the workers they employ. In a public good environment, citizens learn their payoff from a completed project. In all these cases the designer can finalize transfers after the outcome-decision payoffs have been observed. One may wonder if the chronological separation between the outcome decision and the final transfers doesn't introduce obstacles to the practical use 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 an outcome decision today; today it could also collect fees from the agents based on this information to finance the project. After the project has been built, it should collect 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.

In some instances Assumption 1 is less plausible; an example, suggested by a referee, is the following. 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 possibilities, "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 outcome-decision payoff, because he cannot observe agent 2's opinion.

3. EFFICIENCY

The deterministic decision rule $x^* : \Theta \rightarrow X$ 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

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

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

Consider a standard Groves 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 , a Groves mechanism imposes the following transfers

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

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, a standard Groves mechanism will not implement the efficient outcome. 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 outcome-decision payoff of agent $j \neq i$, as computed by the designer, does not coincide with j 's true outcome-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 outcome-decision payoff depends on the reported type θ_i^r only indirectly through the outcome decision x^* .

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 an outcome-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) :

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

The idea, as in a standard Groves 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* (or *two-stage*) *Groves mechanism*. To see that these transfers make truthtelling an equilibrium, first observe that the report of his outcome-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 outcome-decision payoffs, while agent i of type θ_i falsely reports his type to be θ_i' . Under these hypotheses, the reported outcome-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

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

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

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

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

It is well known that with private values and sufficiently rich domains, an efficient mechanism must be a standard Groves mechanism (see Green and Laffont (1977), Holmström (1979), Laffont and Maskin (1979) and, more recently, Williams (1999)). As Example 2 in the next section shows, with interdependent valuations and observ-

able payoffs there are other two-stage mechanisms, besides the generalized Groves mechanism, that yield efficient outcomes.

A drawback of the generalized Groves mechanism that needs to be stressed is that in the second stage agents are indifferent between telling the truth and lying. This implies that there might be other equilibria besides the efficient, truth-telling equilibrium. It is an open question whether more complex mechanisms can be constructed that do not have this feature.

The transfers τ_i as defined in (4) are made after agents have observed their own payoffs. If, for some reason, the designer needed to make transfers at the outcome-decision stage, he could require that the transfers γ_i , as defined in equation (3), be made at the same time the outcome is chosen, and that the transfer adjustments $\tau_i - \gamma_i$ be made after outcome-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 issues of individual rationality and budget balance. The working paper version, Mezzetti (2002), gives a necessary and sufficient condition for a generalized Groves mechanism to be budget balancing and individually rational, that is, to induce voluntary participation.³ To deal with budget balance, 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 , first introduced by D'Aspremont and Gérard-Varet (1979) and Arrow (1979),

$$(7) \quad h_i(\theta^r) = \frac{n-1}{n} \left\{ \sum_{j=1}^n v_j(x^*(\theta^r), \theta^r) - E_{-i} \left[\sum_{j=1}^n v_j(x^*(\theta_i^r, \theta_{-i}), \theta_i^r, \theta_{-i}) \right] + \right. \\ \left. + E_{-(i+1)} \left[\sum_{j=1}^n v_j(x^*(\theta_{i+1}^r, \theta_{-(i+1)}), \theta_{i+1}^r, \theta_{-(i+1)}) \right] \right\}$$

³ The condition is analogous to the one Makowski and Mezzetti (1994) provided for the case of private values.

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. In the reduced-form model, in which the outcome-decision payoffs depend directly on the type profile, these transfers balance the budget on the equilibrium path, $\sum_{i=1}^n t_i = 0$. When the state variable ω is a random function of the state profile, the budget will not be balanced for all realizations of ω , but it will be balanced (on average) for all type profiles,

$$\sum_{i=1}^n \int_{\Omega} \left[\sum_{j \neq i} u_j(x^*(\theta), \omega) - h_i(\theta) \right] d\Pi(\omega|\theta) = 0.$$

PROPOSITION 2: *It is always possible to construct an efficient perfect Bayesian two-stage revelation mechanism which balances the budget for all type profiles.*

I now present an example that shows how the results in this section can be applied. The example is purposefully simple, to highlight the main ideas.

EXAMPLE 1: A seller knows the quality θ_s of a durable good, with $\theta_s \in [-1, 1]$. The quality of the good is unknown to the buyer. A good of quality θ_s is worth θ_s to the buyer and it is worth $\alpha\theta_s$ to the seller, where $\alpha < 1$. Efficiency dictates that the buyer get the good if $\theta_s > 0$ and that the seller keep it if $\theta_s < 0$. In a standard mechanism design model the transfers and the outcome 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 $\theta'_s < 0 < \theta''_s$, incentive compatibility requires $\alpha\theta'_s + t_s(\theta'_s) \geq t_s(\theta''_s)$ and $t_s(\theta''_s) \geq \alpha\theta''_s + t_s(\theta'_s)$, which implies $\theta'_s \geq \theta''_s$, a contradiction. Thus, no efficient standard mechanism exists. On the other hand, efficiency can be obtained by using a two-stage revelation mechanism. Consider a contract between buyer, seller and a dealer (the designer) stipulating that (i) if the seller reports to the dealer that the quality of the good is $\theta^r_s > 0$, then the dealer sells the good to the buyer at a price $p = \beta\theta^r_s$, with $\alpha \leq \beta \leq 1$; (ii) after acquiring the good the buyer will publicly report her payoff u^r_b and the dealer will pay the seller an amount equal to βu^r_b .

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 outcome decision; the outcome is efficient and the buyer ends up paying a price $\beta\theta_s$ to the seller. With these transfers the dealer breaks even, that is, the budget is balanced. Furthermore, the agents will want to participate in the mechanism; individual rationality is satisfied.

If types are single dimensional and a single-crossing condition is satisfied, then standard revelation mechanisms that implement the 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 outcome-decision payoffs. This is because the generalized Groves 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. CONCLUSIONS

I have shown that when agents observe their own outcome-decision payoffs, even if these payoffs are unverifiable, payments that are contingent on payoff reports (two-stage mechanisms) 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."

While the focus here has been on efficiency, two-stage mechanisms can also be used to implement other decision functions that cannot be implemented by standard mechanisms. For example, Mezzetti (2002, 2004) shows that by using two-stage

mechanisms a seller can extract a larger surplus, and sometimes fully extract the surplus as in Crémer and McLean (1985, 1988) (see also McAfee and Reny (1992) and McLean and Postlewaite (2001)). The reason is that, even if types are independent, with interdependent valuations observing his own outcome-decision payoff $u_i(x, \omega)$ provides agent i with a signal that is correlated with the types θ_{-i} of the other agents. The following simple example gives an idea of how this correlation may be exploited.

EXAMPLE 2: An existing business is up for sale. The potential acquirers are two firms. Firm i 's payoff from acquiring the business is $u_i = \alpha\theta_i + \theta_j$, $j \neq i$, where $\alpha > 1$. Firm i privately observes signal θ_i , which is uniformly distributed on $[0, 1]$. Signals are independent. It is well known (e.g., see Myerson (1981)), that there is no standard mechanism that fully extracts the surplus. On the other hand, consider the following generalized two-stage mechanism. In the first stage each firm reports its signal θ_i . Let ℓ be the firm that reported the lower signal, θ_ℓ^r , and h be the one that reported the higher signal, θ_h^r . Firm h is assigned the business and charged $\alpha\theta_h^r + \theta_\ell^r$ (i.e., $t_h = -\alpha\theta_h^r - \theta_\ell^r$). Firm h is then asked to report its outcome-decision payoff in the second stage. If the reported decision payoff is $u_h^r = \alpha\theta_h^r + \theta_\ell^r$, then firm ℓ is charged nothing (i.e., $t_\ell = 0$). If $u_h^r \neq \alpha\theta_h^r + \theta_\ell^r$, then firm ℓ is charged a fine $F > \alpha$ (i.e., $t_\ell = -F$). It is simple to show that it is an equilibrium for firm h to truthfully report its outcome-decision payoff in the second stage and for both firms to truthfully report their signals in the first stage. Hence the proposed mechanism, which is not a generalized Groves mechanism, implements the efficient outcome decision and extracts full surplus for all type realizations. In the second stage, reporting the true outcome-decision payoff is a best reply for firm h . In the first stage, if firm j reports its true signal, then firm i could only gain (over telling the truth) by reporting a signal lower than the truth. If firm i reports $\theta_i^r < \theta_i$, it makes an expected total payoff (outcome-decision payoff plus transfer) equal to $\alpha(\theta_i - \theta_i^r)\theta_i^r - F(1 - \theta_i^r)$, since θ_i^r is the probability that firm j will have a signal lower

than θ_i^r , and the lie will be detected if firm j acquires the business. This total payoff is always negative provided that $F > \alpha$, while telling the truth guarantees firm i zero total payoff.

It is natural to ask to what extent the approach of exploiting the 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 Groves 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, 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.

Dept. of Economics, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-3305, U.S.A.; mezzetti@email.unc.edu; <http://www.unc.edu/~mezzetti/>

REFERENCES

- AKERLOF, G.A. (1970): "The Market for Lemons: Qualitative Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84, 488-500.
- ARROW, K. (1979): "The Property Rights Doctrine and Demand Revelation under Incomplete Information," in *Economics and Human Welfare*, ed. by M. Boskin, New York, Academic Press.
- AUSUBEL, L. (1997): "An Efficient Ascending-Bid Auction for Multiple Objects," <http://www.ausubel.com/auction-papers.htm>.
- BERGEMANN, D. AND J. VÄLIMÄKI (2002): "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70, 1007-33.
- CLARKE, E. (1971): "Multipart Pricing of Public Goods," *Public Choice*, 11, 17-33.
- CRÉMER, J. (1987): "Auctions with Contingent Payments: Comment," *American Economic Review*, 77, 746.
- CRÉMER, J. AND R. MCLEAN (1985): "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, 53, 345-61.
- (1988): "Full Extraction of Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247-57.
- DASGUPTA, P., AND E. MASKIN (2000): "Efficient Auctions," *Quarterly Journal of Economics*, 115, 341-389.
- D'ASPREMONT, L., AND L. GÉRARD-VARET (1979): "Incentives and Incomplete Information," *Journal of Public Economics*, 11, 25-45.
- GREEN, J. AND J.J. LAFFONT (1977): "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods," *Econometrica*, 45, 427-38.
- GROVES, T. (1973): "Incentives in Teams," *Econometrica*, 41, 617-31.
- HANSEN, R.G. (1985): "Auctions with Contingent Payments," *American Economic Review*, 75, 862-65.

HENDEL, I AND A. LIZZERI (2002): “The Role of Leasing under Adverse Selection,” *Journal of Political Economy*, 110, 113-44.

HOLMSTRÖM, B. (1979): “Groves Schemes on Restricted Domains,” *Econometrica*, 47, 1137-44.

JEHIEL, P., AND B. MOLDOVANU (2001): “Efficient Design with Interdependent Valuations,” *Econometrica*, 69, 1237-59.

LAFFONT, J.J. AND E. MASKIN (1979): “A Differential Approach to Expected Utility Maximizing Mechanisms,” in *Aggregation and Revelation of Preferences*, ed. by J.J. Laffont, Amsterdam: North Holland, 289-308.

MAKOWSKI, L. AND C. MEZZETTI (1994): “Bayesian and Weakly Robust First Best Mechanisms: Characterizations,” *Journal of Economic Theory*, 64, 500-19.

MASKIN, E. (1992): “Auctions and Privatization,” in *Privatization*, ed. by H. Siebert, Kiel: Institut fuer Weltwirtschaft der Universitaet Kiel, 115-36.

MCAFEE, P.R. AND P. RENY (1992): “Correlated Information and Mechanism Design,” *Econometrica*, 60, 395-421.

MCLEAN, R. AND A. POSTLEWAITE (2001): “Efficient Auction Mechanisms with Interdependent and Multidimensional Signals,” <http://www.ssc.upenn.edu/~apostlew/>.

MEZZETTI, C. (2002): “Auction Design with Interdependent Valuations: The Generalized Revelation Principle, Efficiency, Full Surplus Extraction and Information Acquisition,” <http://www.unc.edu/~mezzetti/>.

——— (2004): “Mechanism Design with Interdependent Valuations: Surplus Extraction,” <http://www.unc.edu/~mezzetti/>.

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

MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58-73.

PERRY, M., AND P. RENY (2002): “An Efficient Auction,” *Econometrica*, 70, 1199-1212.

SAMUELSON, W. (1987): "Auctions with Contingent Payments: Comment," *American Economic Review*, 77, 740-45.

SPENCE, A.M. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 87, 355-74.

VICKREY, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8-37.

WILLIAMS, S.R. (1999): "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms," *Economic Theory*, 14, 181-201.