Mechanism Design with Interdependent Valuations: Surplus Extraction

Claudio Mezzetti[1]

Department of Economics

University of North Carolina at Chapel Hill

email: mezzetti@email.unc.edu

March 23, 2004

---

[1] This research was started when I was visiting the Department of Applied Mathematics of the University of Venice, and continued while visiting the European University Institute in Florence. Their financial support is gratefully acknowledged.

Abstract

If valuations are interdependent and agents observe their own allocation payoffs, then two-stage revelation mechanisms expand the set of implementable decision functions. In a two-stage revelation mechanism, first agents report their private signals and the allocation is decided, then agents report their payoffs from the allocation and final transfers are made. Conditions are provided under which an uninformed seller can either extract full surplus, or within ε of full surplus, from a sale to privately informed buyers, in spite of the buyers' types being independent random variables.

Keywords: Auctions, Surplus Extraction, Interdependent Valuations, Mechanism Design.

JEL: D44 & D82.

1 Introduction

Consider a seller of an object facing multiple, privately informed, bidders. Can the seller extract all the surplus from the transaction? Crémer and McLean (1985, 1988) (see also McAfee and Reny (1992)) showed that full surplus extraction is generically possible when bidders' types are correlated random variables. Full surplus extraction is achieved by conditioning the payment of each bidder on the realization of the types of the other bidders. On the other hand, it is widely believed that if buyers' types are independent random variables, then full surplus extraction is not possible. This belief is indeed correct in the case in which a bidder's signal only affects his own payoff, the case of private values. A standard example of private values is when agents have private information about their own individual preferences.

In this paper I will show that, contrary to this widely held belief, full surplus extraction is possible when valuations are interdependent even if types are independent random variables. Valuations are interdependent 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. This is the case, for example, when buyers have private information about the quality of the good or service that the seller is trying to sell (e.g., in a mineral-rights auction bidders have private estimates about the quantity of minerals in the tract).

The important insight of this paper is that interdependence of valuations is a form of correlation among bidders payoffs. This correlation can be exploited to achieve full surplus extraction by using a two-stage revelation mechanism in which bidders first report their signals and then the winning bidder reports his realized allocation payoff.

In the standard mechanism design model, agents only report their types to the designer; they do not report their (pre-monetary transfer) payoffs from the allocation after an allocation has been made. Implicitly, standard mechanisms rule out the possibility of transfers after an agent has observed his own payoff from the allocation. Two-stage mechanisms, which have been introduced by Mezzetti (2004) to study efficient decisions, allow transfers to be made after a final allocation has been determined. More precisely, the designer (the seller in this paper) sets up two reporting stages. In the first stage the seller asks about the agents' types. On the basis of these reports, a winner is selected and partial transfers are made. After the winner has observed his payoff from the allocation of the object, the seller asks him to report his realized payoff in a second reporting stage. Then, final payments from all bidders are collected that are contingent on reports in both stages.

While with private values an agent cannot obtain any new information from the observation of his allocation payoff, with interdependent valuations observing his realized payoff provides the agent with new information about the types, or informational signals, of the other agents. In other words, even though types are independent random variables, with interdependent valuations the realized payoff of an agent is correlated with the types of the other agents.

I study two versions of the model. In the first version, allocation payoffs are deterministic functions of the type profile. In this case, the allocation-payoff report of the winning bidder allows the seller to detect and punish first-stage lies by losing bidders. In the second version, allocation payoffs are random functions of the type profiles. In this version, the seller can use lotteries analogous to the ones used in Crémer and McLean (1985, 1988), but based on the payoff report of the winning buyer, to extract the full surplus from the agents. An important difference with the mechanism introduced by Crémer and McLean is that, in order to fully exploit the correlation between allocation payoffs and agents' signals, the two main results of this paper (Propositions 2 and 4) require that the object be allocated randomly, and hence inefficiently, with positive probability. This inefficiency is associated with a loss of surplus to the seller, but this loss can be made arbitrarily small by making the probability of a random allocation arbitrarily small.

Some authors have criticized the full surplus extraction results of Crémer and McLean (1985, 1988) as being counterintuitive (e.g., see McAfee and Reny (1992) and Neeman (2003)). It is not the goal of this paper to contribute to the debate about the plausibility of these results. This paper simply makes the theoretical point that, by using two-stage mechanisms, the setting in which agents' types are independent random variables and valuations are interdependent is very similar to the setting in which agents' types are correlated random variables.

The paper is organized as follows. The next section introduces the model. Sections 3 and 4 contain the full surplus extraction results, while Section 5 concludes.

2 The Model

An uninformed seller of a single item, agent 0, faces n prospective buyers. Each buyer has private information about his own type, or signal, $\theta_i \in \Theta_i$. Let $\Theta = \times_{i=1}^n \Theta_i$ be the set of signal profiles, $\theta = (\theta_1, \dots, \theta_n)$ a generic element of Θ, 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.

A decision $x = (x_0, x_1, ..., x_n)$ is a probability vector; x_i is the probability that agent i gets the object. Buyer i's utility function $U_i : X \times \Omega \times \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ depends on the decision x, the state of the world $\omega \in \Omega$ and the transfer $t_i = (t_i^0, t_i^1, ..., t_i^n)$; t_i^j is buyer i's transfer to the seller when the item is allocated to agent j. The state of the world is a random variable drawn from a probability measure that depends on the agent's types.

$$U_i(x, \omega, t_i) = x_i v_i(\omega) - \sum_{j=0}^{n} x_j t_i^j, \tag{1}$$

where $v_i(\omega)$ is buyer i's allocation payoff when he gets the item and ω is the (realized) state of the world; $v_i(\cdot)$ is assumed to be a continuous and bounded function. The allocation payoff of a buyer that does not receive the object is normalized to zero (there are no allocational externalities). The seller's allocation payoff $v_0(\omega)$ is zero for all ω. Thus, the seller's payoff when the object is allocated to agent j is simply the sum of the transfers from the buyers $\sum_{i=1}^{n} t_i^j$.

In most real world examples, agents observe their own realized allocation payoffs. For example, in mineral-rights auctions the winner eventually learns its own extraction cost and the market value of the minerals found in the tract being auctioned. The buyer of a used car learns the car's quality. Acquirers of an asset learn their valuation for the asset. In many such cases, there are also no obstacles to making final transfers that are contingent on such an observation (see Mezzetti (2004) for additional discussion). Thus, it is reasonable to make the following assumption.

Assumption 1 *If the item is allocated to buyer $i = 1, ..., n$, then for any realization of the state of the world ω, buyer i observes his realized allocation payoff $v_i(\omega)$ before final transfers are made.*

Partial transfers can be made before the winning buyer observes his allocation payoff. In particular, in the mechanism I will propose the payment of the winning buyer does no depend on the reported allocation payoff and thus can be made before this is observed. Losing buyers, on the other hand, may have to make (or receive) payments that depend on the allocation report of the winner.

Under private values there is no loss of generality in assuming that the seller only uses standard revelation mechanisms in which buyers are only asked to report their signals (types). With interdependent valuations and observable allocation payoffs, allowing the seller to collect messages in two reporting stages enlarges the set of implementable decision functions. This is because the second-stage payoff reports can be used to cross check the first-stage type reports. Thus, for example, while there are no efficient standard mechanisms when valuations are interdependent and signals are multidimensional, independent, random variables (e.g., see Jehiel and Moldovanu (2001)), Mezzetti (2004) shows that efficient two-stage mechanisms always exist. In a two-stage revelation mechanism, messages about the buyers' signals are collected in the first stage and are used to determine the allocation of the item. After the agents have observed their allocation payoffs they report them in the second stage; messages from both stages are used to determine the total monetary transfers from the buyers to the seller. Thus, a *two-stage mechanism* consists of an allocation function $x : \Theta \to X$, and $n(n+1)$ transfer functions $t_i^j : \Theta \times \mathcal{V}^j \to \mathbb{R}$, where $\mathcal{V}^j \subset \mathbb{R}$ is the set of buyer j's feasible allocation payoffs, $j = 1, ..., n$, when the item is allocated to him (the only agent who obtains new information from observing his own allocation payoff is the buyer who gets the item).

In the next section, I will assume that signals are continuous variables and that the state of the world coincides with the signal profile, $\Omega = \Theta$. This version of the model, in which allocation payoffs are deterministic functions of the type profile, called Model D, permits a more transparent presentation of the effects at work with interdependent valuations. When the player receiving the object observes his own allocation payoff, he may discover without doubt that some other player misreported his type in the first reporting stage. On the basis of the allocation payoff reported by the winner in the second reporting stage, the seller can then discover and severely punish first-stage lies of the losing buyers.

Section 4 studies a second version, Model R, in which the state space does not coincide with the space of signal profiles, and allocation payoffs are random functions of the type profile.[1] In this version, by observing his own allocation payoff the winner does not detect for sure that other players lied. Full surplus extraction is still possible, however, because the interdependence of valuations implies that the (random) allocation payoff of the winner is correlated with the types of all players. The seller can then use lotteries like in Crémer and McLean (1985, 1988) to induce buyers to truthfully report their types. Roughly speaking, the winner's allocation payoff provides an informative signal about the types of the other buyers that allows the seller to severely punish first-stage lies about types. Both models have been used in the

---

[1] To keep the analysis comparable with Crémer and McLean (1988), in this version of the model the signal and state spaces are finite.

literature, (e.g., see Crémer and McLean (1985, 1988) and Gresik (1991) for Model D, and McLean and Postlewaite (2001) for Model R).

3 Surplus Extraction in Model D

In this section, I will study the case where the set of states of the world coincides with the set of type profiles.

Assumption 2 *In* Model D*: (i) The sets Θ_i are closed and bounded subset of $\mathbb{R}$; (ii) $F_i(\theta_i)$ and $F_{-i}(\theta_{-i}) = \prod_{j \neq i} F_j(\theta_j)$ are the cumulative probability distributions of $\theta_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i}$, respectively; (iii) The set of states of the world is $\Omega = \Theta$.*

In Crémer and McLean (1985, 1988), full surplus extraction occurs at the interim level. Each agent type participates in a lottery which leaves him with zero expected surplus. As we shall see, in Model D it is often possible for the seller to extract all the surplus ex-post (i.e., for all type realizations) in spite of signals being statistically independent. This also implies that on the equilibrium path the seller will only need to collect payments from the buyer who gets the item. It is useful to begin with a simple example that illustrates the main idea.

Example 1. Consider the following special case of the auction model in Myerson (1981). There is a single item for sale and two bidders (potential buyers). Bidder i observes a private signal θ_i; the other players regard θ_i as a random variable with uniform distribution over the interval $[1, 2]$. Buyer i's valuation for the object (the allocation payoff from receiving the object) is

$$v_i(\theta) = \theta_i + \alpha\theta_j \qquad i \neq j,\ i, j = 1, 2,$$

where $\alpha \in (0, 1)$ is a known parameter. Following Myerson (1981) one can show that the optimal auction is any common auction (e.g., a first-price, a second-price (Vickrey), or an ascending auction) with no reserve price. Let $\theta^{(1)} = \max\{\theta_1, \theta_2\}$ and

$\theta^{(2)} = \min\{\theta_1, \theta_2\}$. In a Vickrey or in an ascending auction, bidder i wins the object if $\theta_i = \theta^{(1)}$, and pays a price $p = (1+\alpha)\theta^{(2)}$. Thus, by using a standard optimal auction the seller does not extract the full surplus $\theta^{(1)} + \alpha\theta^{(2)}$.

I now show that the seller could exploit the interdependence of valuations and design a two-stage mechanism that extracts the full surplus. Consider the following "*shoot-the-liar*" mechanism. Bidders are first asked to report their signals (in the first reporting stage). The bidder who reports the highest signal wins the object and then is asked to report the value obtained from the object (in the second reporting stage). The payment, or transfer, to the seller from bidder i, $i \neq j$, as a function of the reports, is as follows:

$$\begin{aligned} t_i^i(\theta_1^r, \theta_2^r, v_1^r, v_2^r) &= \theta_i^r + \alpha\theta_j^r \\ t_i^j(\theta_1^r, \theta_2^r, v_1^r, v_2^r) &= \begin{cases} 0 & \text{if} \quad v_j^r = \theta_j^r + \alpha\theta_i^r \\ P & \text{if} \quad v_j^r \neq \theta_j^r + \alpha\theta_i^r, \end{cases} \end{aligned}$$

where $P > 1$ is a constant. To see that the incentive compatibility constraints for the bidders are satisfied (i.e., that bidder i wants to report truthfully), note first that t_i^i does not depend on i's reported allocation payoff v_i^r (when i obtains the item, it is common knowledge that v_j is zero and so t_i^i cannot depend on v_j^r). Hence, truthful reporting of his realized allocation payoff is optimal for the winning bidder in the second reporting stage. Suppose that bidder j truthfully reports his signal in the first reporting stage and if he wins he then truthfully reports his realized valuation in the second reporting stage. The expected payoff to player i from reporting $\theta_i^r \neq \theta_i$, while his signal is θ_i, then is

$$U_i(\theta_i^r; \theta_i) = \int_0^{\theta_i^r} (\theta_i - \theta_i^r)\, d\theta_j - \int_{\theta_i^r}^1 P d\theta_j = (\theta_i - \theta_i^r)\, \theta_i^r - P\,(1 - \theta_i^r)\,,$$

while i's expected payoff from truthfully reporting $\theta_i^r = \theta_i$ is zero. Clearly, for $P > 1$, U_i is maximized by reporting truthfully in the first stage, $\theta_i^r = \theta_i$. In this two-stage revelation mechanism each bidder obtains a zero payoff and the seller extracts the full surplus for all type realizations.

The "shoot-the-liar" mechanism contains a discrete penalty jump for being discovered lying. One could also construct continuous penalties by setting the transfers to the seller as follows

$$\begin{aligned} t_i^i(\theta_1^r, \theta_2^r, v_1^r, v_2^r) &= \theta_i^r + \alpha\theta_j^r \\ t_i^j(\theta_1^r, \theta_2^r, v_1^r, v_2^r) &= \frac{1}{\alpha^2}\gamma\left(v_j^r - \theta_j^r - \alpha\theta_i^r\right)^2 + \frac{1}{\alpha}\frac{\theta_i^r}{1-\theta_i^r}\left(v_j^r - \theta_j^r - \alpha\theta_i^r\right), \end{aligned}$$

where $\gamma > 0$ is a constant. Suppose again that bidder j truthfully reports his signal in the first reporting stage and if he wins he then truthfully reports his valuation in the second reporting stage. The expected payoff to player i from reporting θ_i^r, while his signal is θ_i, then is

$$\begin{aligned} U_i(\theta_i^r; \theta_i) &= \int_0^{\theta_i^r} (\theta_i - \theta_i^r)\, d\theta_j - \int_{\theta_i^r}^1 \left[\gamma\left(\theta_i - \theta_i^r\right)^2 + \frac{\theta_i^r}{1-\theta_i^r}\left(\theta_i - \theta_i^r\right)\right] d\theta_j \\ &= -\gamma\left(\theta_i - \theta_i^r\right)^2\left(1 - \theta_i^r\right). \end{aligned}$$

Again, U_i is maximized by reporting truthfully in the first stage, $\theta_i^r = \theta_i$.

As this example makes clear, full surplus extraction requires, roughly speaking, that potentially profitable lies in the first reporting stage be detected with positive probability if buyers truthfully report their allocation payoffs in the second stage. Before introducing a condition that captures this requirement, we need additional notation. Let $x^*(\theta_i, \theta_{-i})$ be an efficient allocation rule:

$$x_i^*(\theta_i, \theta_{-i}) > 0 \quad \text{if and only if} \quad i \in \arg\max_{j=0,\ldots,n} v_j(\theta_i, \theta_{-i})$$

(the efficient allocation is that the seller keeps the item if v_i is negative for all buyers i). Let

$$V_i(\theta_i; \theta_i') = \int_{\Theta_{-i}} x_i^*(\theta_i', \theta_{-i}) v_i(\theta_i, \theta_{-i}) dF_{-i}(\theta_{-i})$$

be the expected allocation payoff of buyer i of type θ_i when he reports θ_i', all other players report truthfully, and the seller uses an efficient allocation rule. We say that *the seller extracts full surplus for all type realizations* if he uses an efficient allocation rule and on the equilibrium path all buyers obtain zero net utility (allocation payoff minus transfer) for all type profiles. The following condition requires that it be incentive compatible for player i to truthfully report when the seller is trying to extract full surplus for all type realizations.

Condition 1 *If $V_i(\theta_i; \theta_i') > V_i(\theta_i', \theta_i')$, then there exists $j \neq i$ and $\delta > 0$ such that $x_j^*(\theta_i', \theta_{-i}) > \delta$, and $v_j(\theta_i, \theta_{-i}) \neq v_j(\theta_i', \theta_{-i})$ for all θ_{-i} that belong to a set $\Theta_{-i}^+ \subset \Theta_{-i}$ having probability measure greater than δ.*

Condition 1 requires that if type θ_i of agent i tells a profitable lie, pretending to be θ_i', then with positive probability the item will go to another agent j whose observed allocation payoffs will be inconsistent with the reported type of agent i. As Proposition 1 shows, Condition 1 is necessary and sufficient for the seller to be able to extract full surplus for all type realizations. The seller should use the following *shoot-the-liar* mechanism. First, based on the first-stage signal reports, the item is allocated efficiently (i.e., given to the buyer with the highest valuation). Second, the seller collects transfers from the buyer i that obtains the item; this transfer, which can be made before the winner observes his allocation payoff, is equal to the full value of the item to i, $v_i(\theta^r)$. Third, the winning buyer reports his realized allocation payoff in the second reporting stage. If this report is inconsistent with the type reports made by the other buyers in the first stage, then all other buyers are imposed severe (but bounded) fines.[2] On the equilibrium path no buyer will lie, and thus only the winning bidder will need to make a payment to the seller.

Condition 1 essentially requires that the shoot-the-liar mechanism be incentive compatible. A more transparent (but also much stronger) condition that guarantees that Condition 1 is satisfied is the following.

Condition 2 *(i) The valuation function of player j is strictly monotone in the type of player i. (ii) If $x_i^*(\theta_i', \theta_{-i}) = 1$ for all θ_{-i}, except possibly for a subset of Θ_{-i} with zero probability measure, then $V_i(\theta_i; \theta_i') \leq V_i(\theta_i'; \theta_i')$ for all θ_i.*

The second part of the condition rules out that a type θ_i can make a report θ_i' that both underestimates his expected payoff from owning the object and guarantees that he will always get it. It is clear that Condition 2 implies Condition 1.

Proposition 1 *In Model D, there is a perfect Bayesian equilibrium of a two-stage revelation mechanism in which the seller fully extracts the surplus from the agents for all type realizations if and only if Condition 1 holds.*

Proof. (If) Consider a two-stage revelation mechanism that uses an efficient decision rule. Let the transfer function be:

$$\begin{aligned} t_i^i(\theta^r, v_i^r) &= v_i(\theta^r) \\ t_i^j(\theta^r, v_j^r) &= \begin{cases} 0 & \text{if} \quad v_j^r = v_j(\theta^r) \\ P & \text{if} \quad v_j^r \neq v_j(\theta^r). \end{cases} \end{aligned}$$

Suppose that all the other agents always truthfully report their signals in the first stage and their allocation payoffs, if they get the item, in the second reporting stage.

---

[2] An important drawback of the shoot-the-liar mechanism is that in the second stage the winning buyer is indifferent between reporting his true allocation payoff and reporting any other payoff. As a result, there might be additional equilibria that do not yield full surplus extraction. It is an open question under which conditions more complex mechanisms that do not have this feature can be constructed. See Brusco (1998) for a study of unique implementation of the full surplus extraction outcome in the model of Crémer and McLean (1988).

Since agent i's transfer does not depend on his reported allocation payoff, he has no incentive to deviate from truthfully reporting it in the second stage. If agent i of type θ_i truthfully reports his signal in the first stage, then he gets zero total utility. If he reports a type θ_i', then, by Condition 1, he gets

$$\begin{aligned}\int_{\Theta_{-i}} x_i^*(\theta_i',\theta_{-i})\left[v_i(\theta_i,\theta_{-i}) - v_i(\theta_i',\theta_{-i})\right] dF_{-i}(\theta_{-i}) - P\sum_{j\neq i}\int_{\Theta_{-i}^j} x_j^*(\theta_i',\theta_{-i})dF_{-i}(\theta_{-i}) \\ = V_i(\theta_i;\theta_i') - V_i(\theta_i';\theta_i') - P\sum_{j\neq i}\int_{\Theta_{-i}^j} x_j^*(\theta_i',\theta_{-i})dF_{-i}(\theta_{-i}),\end{aligned}$$

where Θ_{-i}^j is the subset for which $v_j(\theta_i,\theta_{-i}) \neq v_j(\theta_i',\theta_{-i})$ for player $j \neq i$. Since, by Condition 1, for at least a $j \neq i$ the set Θ_{-i}^j has probability measure greater than δ and $x_j^*(\theta_i',\theta_{-i})$ is also greater than δ, we have $\sum_{j\neq i}\int_{\Theta_{-i}^j} x_j^*(\theta_i',\theta_{-i}) > \delta^2$. It is then possible to choose P large enough to deter any first-stage type misreport.

(Only if) Suppose Condition 1 is violated, so that there exists a buyer i and two signals θ_i and θ_i' with $V_i(\theta_i;\theta_i') > V_i(\theta_i';\theta_i')$, but for all buyers $j \neq i$ either $x_j^*(\theta_i',\theta_{-i}) \leq \delta$ for all $\delta > 0$ (i.e., $x_j^*(\theta_i',\theta_{-i}) = 0$), or $v_j(\theta_i,\theta_{-i}) = v_j(\theta_i',\theta_{-i})$ for almost all $\theta_{-i} \in \Theta_{-i}$. In any mechanism that extracts the full surplus for all type realizations, the winning buyer must be charged his valuation for the object, as computed from the type reports. Then, type θ_i profits from reporting that his signal is θ_i': the lie is never discovered and his expected utility is $V_i(\theta_i;\theta_i') - V_i(\theta_i';\theta_i') > 0$. Hence the seller cannot extract the full surplus for all type realizations. The continuity of $v_i(\theta)$ implies that expected revenue is strictly less than expected full surplus. ■

While there are important examples where it is satisfied, Condition 1 is restrictive. Even if valuations are fully interdependent (i.e., if they depend on the types of all agents) it can be violated, as shown by the following modifications of Example 1.

Example 2. As in Example 1, there is a single item for sale and two buyers with valuations $v_i(\theta) = \theta_i + \alpha\theta_j$, $\alpha < 1$, but now buyer 1's type has support $[1, 3]$, while 2's type has support $[1,2]$. For all types of buyer 1 in the set $(2,3]$, buyer 1 can report $\theta_1^r = 2$, making sure that $x_2^*(2,\theta_2) = 0$ with probability one: this lie will never be detected. Condition 1 is violated and full surplus extraction is impossible.

Example 3. As in Example 1, there is a single item for sale and two buyers with types distributed on $[1,2]$ and valuations $v_i(\theta) = \theta_i + \alpha\theta_j$, but now $\alpha \in (1,2)$. In this case it is efficient to give the item to the buyer with the lowest type. But then by reporting $\theta_i = 1$ buyer i can make sure that $x_j^*(1,\theta_j) = 0$ with probability one, and, at the same time, reduce his payment to the seller: Condition 1 is violated and full surplus cannot be extracted.

What prevents full surplus extraction in Examples 2 and 3 is that a buyer can pay less than full surplus while being sure that he will get the object. In both examples, if buyer j is given the item, then any lie by i will be discovered. The seller could then induce truthtelling by assigning the item to each buyer with positive probability, irrespective of the type reports. Knowing that with positive probability any lie will be discovered and severely punished, buyers would report their true types. Since the allocation would be inefficient with positive probability, the seller would not be able to extract full surplus. He could, however, extract all the surplus from the buyer receiving the object and, by making the probability of a random allocation arbitrarily small, the seller's expected revenue could be made arbitrarily close to the ex-ante expected surplus $S = \int_\Theta \sum_{i=1}^n x_i^*(\theta) v_i(\theta) dF(\theta)$.

Definition 1 *We say that the seller's revenue is within ε of full surplus for all type realizations if: (i) the seller extracts all the surplus from the buyer receiving the object; (ii) the seller's ex-ante expected revenue is lower than the ex-ante expected surplus by an amount smaller than ε.*

We now introduce a condition, which is satisfied by Examples 2 and 3, under which the seller can raise within ε of full surplus for all type realizations.

Condition 3 *For all i and all θ_i, θ_i' with $\theta_i \neq \theta_i'$, there exists $j \neq i$ and a $\delta > 0$ such that $v_j(\theta_i, \theta_{-i}) \neq v_j(\theta_i', \theta_{-i})$ for all $\theta_{-i} \in \Theta_{-i}^+$, where $\Theta_{-i}^+ \subset \Theta_{-i}$ has probability measure greater than δ.*

Condition 3 is much weaker and more transparent than Condition 1. It only requires that for every possible lie by agent i, there is another agent j who would detect the lie if he were allocated the object. Condition 3, unlike Condition 1, does not require that it is efficient to allocate the object to j.

For any given θ_{-i}, for a generic subset of all continuously differentiable functions $v_j(\cdot, \theta_{-i})$ with domain Θ_i, the equality $v_j(\theta_i, \theta_{-i}) = c$ admits a finite number of solutions; that is, there are a finite number of types $\theta_i, \theta_i', ...$ such that $v_j(\theta_i, \theta_{-i}) = v_j(\theta_i', \theta_{-i})$. Thus, if there are at least 3 buyers, then Condition 3 is generically satisfied, because for any given θ_{-i} and for generic, continuously differentiable functions $v_j(\cdot, \theta_{-i})$, $v_k(\cdot, \theta_{-i})$ the equalities $v_j(\theta_i, \theta_{-i}) = v_j(\theta_i', \theta_{-i})$ and $v_k(\theta_i, \theta_{-i}) = v_k(\theta_i', \theta_{-i})$, with $i \neq j \neq k$ cannot both hold.

If Condition 3 is satisfied, then the seller can raise within ε of full surplus for all type realizations, by using a modified shoot-the-liar mechanism in which (i) the efficient allocation is implemented with a probability close to, but bounded away from, one, and (ii) each player receives the object with small, but positive, probability.

Proposition 2 *In Model D, if Condition 3 holds, then for all $\varepsilon > 0$ there is a perfect Bayesian equilibrium of a two-stage revelation mechanism in which the seller raises within ε of full surplus for all type realizations.*

Proof. Let the allocation rule be

$$x_i(\theta_i^r, \theta_{-i}^r) = \frac{\delta}{n} + (1 - \delta) x_i^*(\theta_i^r, \theta_{-i}^r),$$

and the transfer function be:

$$\begin{aligned} t_i^i(\theta^r, v_i^r) &= v_i(\theta^r) \\ t_i^j(\theta^r, v_j^r) &= \begin{cases} 0 & \text{if} \quad v_j^r = v_j(\theta^r) \\ P & \text{if} \quad v_j^r \neq v_j(\theta^r). \end{cases} \end{aligned}$$

Agents have no incentive to deviate in the second stage. If all the other agents always truthfully report their signals in the first stage and their allocation payoffs in the second stage, then agent i of type θ_i gets zero total utility by truthfully reporting his signal in the first stage. If he reports a type θ_i', then by Condition 3 he gets

$$\int_{\Theta_{-i}} \left[\frac{\delta}{n} + (1-\delta)x_i^*(\theta_i', \theta_{-i})\right] \left[v_i(\theta_i, \theta_{-i}) - v_i(\theta_i', \theta_{-i})\right] dF_{-i}(\theta_{-i}) \\ - P \sum_{j \neq i} \int_{\Theta_{-i}^j} \left[\frac{\delta}{n} + (1-\delta)x_j^*(\theta_i', \theta_{-i})\right] dF_{-i}(\theta_{-i}),$$

where Θ_{-i}^j is the subset for which $v_j(\theta_i, \theta_{-i}) \neq v_j(\theta_i', \theta_{-i})$ for player $j \neq i$. By Condition 3, Θ_{-i}^j has probability measure greater than δ for at least a $j \neq i$, and the expression in square bracket in the second integral is greater than δ/n. It is then possible to choose P large enough to deter any first-stage type misreport. This shows that truthful reporting is an equilibrium of the proposed mechanism. The object is allocated efficiently with probability $(1-\delta)$ and randomly with probability δ. Since the payoff functions are bounded, it is possible to choose δ small enough so that expected revenue is within ε of full surplus, for any given ε (simply take $\delta < \varepsilon/B$, where B is an upper bound for all agents' allocation payoffs). ■

Condition 3 and Proposition 2 could be extended to the case in which θ_i is multidimensional without any modification. In such a case, for Condition 3 to hold for a generic set of valuation functions it must be that $n-1$, the number of players other than i, is greater than the dimension of θ_i.

Even though Condition 3 is satisfied by a generic set of valuation functions, there are simple functions that violate it.

Example 4. As in Example 1, there is a single item for sale and two buyers with types distributed on $[1, 2]$, but now valuations are $v_i(\theta) = \theta_i + \alpha \min\{\theta_j, 1.5\}$, with $\alpha \in (0, 1)$. For all types in the set $(1.5, 2]$, buyer i can report $\theta_i^r = 1.5$; this lie will never be detected. Condition 3 is violated and the seller cannot raise within ε of full surplus for all type realizations.

We will now see that even in cases like the one described in Example 4, by using two-stage mechanisms the seller can strictly increase his expected revenue over the revenue he would obtain in a standard revelation mechanism.

Let $\widehat{M} = \langle \widehat{x}, \widehat{t} \rangle$ be a standard single-stage revelation mechanism that maximizes the seller's expected revenue (assume such an optimal standard mechanism exists). The payoff in $\widehat{M}$ of a player of type θ_i who reports θ_i' while all other players report truthfully their types is

$$\widehat{U}_i(\theta_i; \theta_i') = \int_{\Theta_{-i}} \left[\widehat{x}_i(\theta_i', \theta_{-i}) v_i(\theta_i, \theta_{-i}) - \sum_{j=0}^{n} \widehat{x}_j(\theta_i', \theta_{-i}) \widehat{t}_i^j(\theta_i', \theta_{-i}) \right] dF_{-i}(\theta_{-i}),$$

and incentive compatibility requires $\widehat{U}_i(\theta_i) \equiv \widehat{U}_i(\theta_i; \theta_i) \geq \widehat{U}_i(\theta_i; \theta_i')$.

By using a two-stage mechanism the seller can always mimic $\widehat{M}$. All that he needs to do is ignore the second-stage report about the allocation payoff. Under the following condition, which is satisfied by Example 4, the seller can do better; that is, he can raise greater revenue.

Condition 4 *There exists i, an interval $[a, b] \subset \Theta_i$ and a $\delta > 0$ such that: (i) $\widehat{U}_i(\theta_i) > 0$ for all $\theta_i \in [a, b]$; (ii) if $\theta_i \in [a, b]$ then for all θ_i, θ_i' with $\theta_i' \neq \theta_i$ there exists a buyer $j \neq i$ such that $v_j(\theta_i, \theta_{-i}) \neq v_j(\theta_i', \theta_{-i})$ for all $\theta_{-i} \in \Theta_{-i}^+$, where $\Theta_{-i}^+ \subset \Theta_{-i}$ has probability measure greater than δ.*

Condition 4 is weaker than Condition 3. It only requires that some lies by some player, rather than all lies by all players, be detectable. In particular, it is satisfied if the valuation function of at least another player j is weakly monotone in θ_i, and strictly monotone in an interval $[a,b]$ with the property that types $\theta_i \in [a,b]$ obtain positive information rent in the optimal standard mechanism.

Proposition 3 *In Model D, if Condition* 4 *holds, then, there is a two-stage revelation mechanism that raises greater expected revenue to the seller than the optimal single-stage revelation mechanism.*

Proof. Let i^* be the player that satisfies Condition 4. Suppose the seller uses the mechanism $M = \langle x, t \rangle$ defined as follows: (i) For all i, $x_i(\theta^r) = \frac{\delta}{n} + (1-\delta)\widehat{x}_i(\theta^r)$; that is, each player is randomly assigned the object with probability $\frac{\delta}{n}$ independently of the reports, while with probability $(1-\delta)$ the probability decision function is as in mechanism $\widehat{M}$. (ii) For all $i \neq i^*$, and all j,

$$t_i^j(\theta^r, v_j^r) = \frac{(1-\delta)\widehat{x}_j(\theta^r)\widehat{t}_i^j(\theta^r)}{\frac{\delta}{n} + (1-\delta)\widehat{x}_j(\theta^r)}.$$

(iii) For $i = i^*$,

$$t_{i^*}^{i^*}(\theta^r, v_{i^*}^r) = \begin{cases} \dfrac{(1-\delta)\widehat{x}_{i^*}(\theta_{i^*}^r, \theta_{-i^*}^r)\widehat{t}_{i^*}^{i^*}(\theta_{i^*}^r, \theta_{-i^*}^r)}{\frac{\delta}{n} + (1-\delta)\widehat{x}_{i^*}(\theta_{i^*}^r, \theta_{-i^*}^r)} & \text{if} \quad \theta_{i^*}^r \notin [a,b] \\ v_{i^*}(\theta_{i^*}^r, \theta_{-i^*}^r) & \text{if} \quad \theta_{i^*}^r \in [a,b] \end{cases}$$

$$t_{i^*}^j(\theta^r, v_j^r) = \begin{cases} \dfrac{(1-\delta)\widehat{x}_j(\theta_{i^*}^r, \theta_{-i^*}^r)\widehat{t}_{i^*}^j(\theta_{i^*}^r, \theta_{-i^*}^r)}{\frac{\delta}{n} + (1-\delta)\widehat{x}_j(\theta_{i^*}^r, \theta_{-i^*}^r)} & \text{if} \quad \theta_{i^*}^r \notin [a,b] \text{ and } v_j^r = v_j(\theta_{i^*}^r, \theta_{-i^*}^r) \\ 0 & \text{if} \quad \theta_{i^*}^r \in [a,b] \text{ and } v_j^r = v_j(\theta_{i^*}^r, \theta_{-i^*}^r) \\ P & \text{if} \quad v_j^r \neq v_j(\theta_{i^*}^r, \theta_{-i^*}^r). \end{cases}$$

First note that the payoff of a player $i \neq i^*$ of type θ_i who reports θ'_i while all other players report truthfully their types and allocation payoffs is

$$U_i(\theta_i;\theta'_i) =$$
$$\int_{\Theta_{-i}} \left[\left[\frac{\delta}{n} + (1-\delta)\widehat{x}_i(\theta'_i,\theta_{-i}) \right] v_i(\theta_i,\theta_{-i}) - \sum_{j=0}^{n} (1-\delta)\widehat{x}_j(\theta'_i,\theta_{-i})\widehat{t}^{\,j}_i(\theta'_i,\theta_{-i}) \right] dF_{-i}(\theta_{-i}).$$

Thus, the incentive compatibility condition $U_i(\theta_i;\theta_i) \geq U_i(\theta_i;\theta'_i)$ is equivalent to $\widehat{U}_i(\theta_i;\theta_i) \geq \widehat{U}_i(\theta_i;\theta'_i)$. Since $\widehat{M}$ is incentive compatible, this shows that telling the truth is a best response for $i \neq i^*$ when all other players tell the truth. We now show that telling the truth is also a best response for i^*. Lies by any type $\theta_{i^*} \in [a,b]$ would be detected with probability greater than δ and thus for sufficiently large P these types do not want to deviate. The same is true for a type $\theta_{i^*} \notin [a,b]$ if he reports a type $\theta'_{i^*} \in [a,b]$. It remains to show that $\theta_{i^*} \notin [a,b]$ does not want to report another $\theta'_{i^*} \notin [a,b]$. The expected payoff from such a misreport, while all other agents report truthfully, assuming the misreport is not detected, is

$$\begin{aligned} U_{i^*}(\theta_{i^*};\theta'_{i^*}) = \int_{\Theta_{-i}} & \left[\frac{\delta}{n} + (1-\delta)\widehat{x}_{i^*}(\theta'_{i^*},\theta_{-i^*}) \right] v_{i^*}(\theta_{i^*},\theta_{-i^*}) dF_{-i^*}(\theta_{-i^*}) \\ & - \int_{\Theta_{-i}} \sum_{j=0}^{n} (1-\delta)\widehat{x}_j(\theta'_{i^*},\theta_{-i^*})\widehat{t}^{\,j}_{i^*}(\theta'_{i^*},\theta_{-i^*}) dF_{-i^*}(\theta_{-i^*}). \end{aligned}$$

As for agents $i \neq i^*$, in this case the incentive compatibility condition $U_i(\theta_{i^*};\theta_{i^*}) \geq U_i(\theta_{i^*};\theta'_{i^*})$ is equivalent to $\widehat{U}_i(\theta_{i^*};\theta_{i^*}) \geq \widehat{U}_i(\theta_{i^*};\theta'_{i^*})$. Thus, we have shown that the proposed mechanism is incentive compatible. Since it is clear that it is also individually rational, it only remains to show that it yields the seller greater expected revenue than $\widehat{M}$. Let $\widehat{R}$ be the seller's expected revenue in mechanism $\widehat{M}$ and let $\widehat{L}$ be the expected information rent in $\widehat{M}$ accruing to types $\theta_{i^*} \in [a,b]$,

$$\widehat{L} = \int_a^b \widehat{U}_i(\theta_{i^*}) dF_{i^*}(\theta_{i^*}).$$

Then the expected revenue in the proposed mechanism is bounded below by $(1-\delta)\widehat{R}+\widehat{L}$, which is greater than $\widehat{R}$ for δ sufficiently small. ■

4 Surplus Extraction in Model R

In the model of this section, observing his own allocation payoff $v_i(\omega)$ provides agent i with a signal that is imperfectly correlated with the types θ_{-i} of the other agents. It is thus possible to check the robustness of the surplus extraction results derived in the previous section to situations in which the seller cannot be certain that a player lied in the type reporting stage. For the sake of brevity, I will only derive the counterpart of Proposition 2, which I view as the main result of Section 3.

Assumption 3 *In* Model R*: (1) The sets* Θ_i *and* Ω *are finite; (2)* $f_i(\theta_i) > 0$ *and* $f_{-i}(\theta_{-i}) = \prod_{j\neq i} f_j(\theta_j) > 0$ *are the probabilities of* $\theta_i \in \Theta_i$ *and* $\theta_{-i} \in \Theta_{-i}$*; (3)* $\pi(\omega|\theta)$ *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*.*

Let $v_i(\theta) = \sum_{\omega\in\Omega} \pi(\omega|\theta) v_i(\omega)$ be the expected payoff of player i conditional on the type profile θ. Let $\Omega(v_i)$ be the set of states of the world ω for which $v_i(\omega) = v_i$, and let $\mathcal{V}^i = \{v_i \in \mathbb{R} : \exists \omega \in \Omega \text{ s.t. } v_i(\omega) = v_i\}$ be the set of feasible allocation payoffs for agent i. Since the type sets Θ_i and the set of states of the world Ω are finite, $\mathcal{V}^i$ is also finite; let k_i be its cardinality. We can identify $\mathcal{V}^i$ with the vector $(v_i^1, ..., v_i^{k_i})$.

Let $\pi(v_j, \theta_i, \theta_{-i}) = \sum_{\omega\in\Omega(v_j)} \pi(\omega|\theta_i, \theta_{-i})$ be the probability that $v_j(\omega) = v_j$, conditional on θ_{-i} and θ_i, and $\pi_j(\cdot, \theta_i, \theta_{-i}) = \left(\pi(v_j^1, \theta_i, \theta_{-i}), ..., \pi(v_j^{k_j}, \theta_i, \theta_{-i})\right)$ be the vector of probabilities of j's allocation payoffs, conditional on θ_i and θ_{-i}. Note that $\pi_j(\cdot, \theta_i, \theta_{-i})$ has dimension k_j.

Even if the winning buyer truthfully reports his allocation payoff, in Model R the seller cannot be certain that some player misreported his type in the first reporting stage. As a result, the seller has no hope of extracting within ε of full surplus for all type profiles. I will show below, however, that he can make sure that his ex-ante expected surplus (i.e., his surplus when averaging over all type profiles) is within ε of full surplus, and that all buyers obtain zero expected utility at the interim stage (i.e., when they know their own types).

Definition 2 *We say that the seller's revenue is within ε of full surplus if: (i) conditional on his true type, each buyer's expected utility is zero; (ii) the seller's ex-ante expected revenue is lower than the ex-ante expected surplus by less than ε.*

The following condition is the appropriate adaptation of Condition 3 to Model R. It is close in spirit to the condition in Theorem 2 of Crémer and McLean (1988). As will shall see, it ensures that the seller can raise within ε of full surplus.

Condition 5 *For all i and all θ_{-i}, there exists $j(\theta_{-i}) \neq i$ such that $\pi_{j(\theta_{-i})}(\cdot, \theta_i, \theta_{-i}) \neq \sum_{\theta_i' \neq \theta_i} \rho_i(\theta_i') \pi_{j(\theta_{-i})}(\cdot, \theta_i', \theta_{-i})$ for all $\theta_i \in \Theta_i$, and all $\rho_i(\theta_i') \geq 0$.*

This condition says that, for all i and θ_{-i}, there is a bidder $j \neq i$ such that j's payoff-probability vector conditional on θ_{-i} and θ_i is not a positive linear combination of all other payoff-probability vectors of bidder j conditional on θ_{-i} and $\theta_i' \neq \theta_i$.

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 (side-bet) 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 allocation rule with probability $1-\delta$; with probability δ, the object is randomly given to one of the buyers. The agent who receives the object is asked to pay his expected allocation payoff given the reported types. All other agents must pay (or be paid) an amount that depends on the reported allocation payoff of the winning buyer. If Condition 5 holds, these payments can be structured 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 allocation payoff of the winning agent, rather then the correlation among the types of all agents. Another difference is that, in order to fully exploit this correlation, with positive probability the object is assigned randomly. This introduces an inefficiency, and hence a loss of surplus to the seller, but this loss can be made arbitrarily small by lowering δ and raising the stakes in the lotteries offered to the losing buyers.[3]

Proposition 4 *In Model R, if Condition 5 holds, then for all* $\varepsilon > 0$ *there is a perfect Bayesian equilibrium of a two-stage revelation mechanism in which the seller raises within* ε *of full surplus.*

Proof. The proof is quite similar to the proof of Theorem 2 in Crémer and McLean (1988). By Condition 5 and Farkas' Lemma, for any given i and θ_{-i} there exists a $j(\theta_{-i})$ and a $k_{j(\theta_{-i})}$ dimensional vector

$$g_{j(\theta_{-i})}(\cdot, \theta_i, \theta_{-i}) = \left(g_{j(\theta_{-i})}(v^1_{j(\theta_{-i})}\cdot, \theta_i, \theta_{-i}), ..., g_{j(\theta_{-i})}(v^{k_{j(\theta_{-i})}}_{j(\theta_{-i})}, \theta_i, \theta_{-i}) \right)$$

such that the following vector-product relations hold: (i) $g_{j(\theta_{-i})}(\cdot, \theta_i, \theta_{-i})\pi_{j(\theta_{-i})}(\cdot, \theta_i, \theta_{-i}) > 0$, and (ii) $g_{j(\theta_{-i})}(\cdot, \theta_i, \theta_{-i})\pi_{j(\theta_{-i})}(\cdot, \theta_i, \theta'_{-i}) \leq 0$ for all $\theta'_i \neq \theta_i$. Then consider the two-stage revelation mechanism with allocation rule

$$x_i(\theta^r_i, \theta^r_{-i}) = \frac{\delta}{n} + (1 - \delta)x^*_i(\theta^r_i, \theta^r_{-i}),$$

---

[3] As in Crémer and McLean (1985, 1988) these lotteries may involve large prizes and penalties; see Kosmopoulou and Williams (1998) for a criticism of this feature of the mechanism.

where $x^*(\cdot)$ is an efficient allocation rule, and transfer functions:

$$\begin{aligned}
t_i^i(\theta^r, v_i^r) &= v_i(\theta^r) \\
t_i^j(\theta^r, v_j^r) &= 0 \text{ if } j \neq j(\theta_{-i}^r) \\
t_i^{j(\theta_{-i}^r)}(\theta^r, v_j^r) &= \mu \left[g_{j(\theta_{-i}^r)}(\cdot, \theta_i^r, \theta_{-i}^r) \pi_{j(\theta_{-i}^r)}(\cdot, \theta_i^r, \theta_{-i}^r) - g_{j(\theta_{-i}^r)}(v_j^r, \theta_i^r, \theta_{-i}^r) \right]
\end{aligned}$$

The winning buyer has no incentive to deviate from telling the truth in the second stage. Suppose all other bidders always truthfully report their signals in the first stage and their allocation payoffs in the second stage. Then agent i of type θ_i gets zero expected utility by reporting truthfully, because the expected value of $t_i^{j(\theta_{-i}^r)}$ is zero if i reports truthfully. Since the expected value of $t_i^{j(\theta_{-i}^r)}$ is negative if i misreports his type, and $j(\theta_{-i}^r)$ will obtain the object with probability at least δ/n, by choosing μ sufficiently large the seller can make sure that by misreporting a player obtains negative utility. This shows that truthful reporting is an equilibrium of the mechanism. Since the object is allocated efficiently with probability $(1-\delta)$ and the payoff functions are bounded, it is possible to choose δ small enough so that expected revenue is within ε of full surplus, for any given ε. ■

In the mechanism described in Proposition 4, as well as in Crémer and McLean, agents face lotteries that leave them with zero expected utility. However, while the Crémer and McLean condition is purely a restriction on the correlation among the probability distributions on Θ_{-i}, Condition 5 is a joint restriction on the conditional distribution $\pi(\omega|\theta)$ and the payoff functions of the agents. As a consequence, Condition 5 holds generically only if for each player i, the set of possible payoffs $\mathcal{V}^j$ of another player $j \neq i$ has at least as many elements as the set of types of player i. If for a player i the sets of possible payoffs of all other players j contains less elements than Θ_i, (i.e., if $k_j < m_i$) then Condition 5 cannot hold generically, because it is always possible to write one of the k_j dimensional allocation-probability vectors

$\pi_j(\cdot, \theta_i, \theta_{-i})$ as a linear combination, and in a positive-measure number of cases as a positive linear combination, of the other $m_i - 1$ vectors $\pi_j(\cdot, \theta'_i, \theta_{-i})$, with $\theta'_i \neq \theta_i$. The following example clarifies this point.

Example 5. An existing business is up for sale; there are two possible states of the world $\Omega = \{0, 1\}$. The potential acquirers are two firms. Firm 1's payoff from acquiring the business is $v_1 = 2 - \omega$, while firm 2's is $v_2 = 1 + \omega$. Suppose first that firm i privately observes one of two equally likely signals $\Theta_i \in \{1, 3\}$. Suppose that $\pi(\omega = 1|\theta_1, \theta_2) = \alpha\theta_1 + \beta\theta_2$. Then $\mathcal{V}^i = \{1, 2\}$ and the allocation-payoff probability vector of buyer 1 is $\pi_1(\cdot, \theta_2, \theta_1) = (\pi(v_1 = 1, \theta_2, \theta_1), \pi(v_1 = 2, \theta_2, \theta_1)) = (\alpha\theta_1 + \beta\theta_2, 1 - \alpha\theta_1 - \beta\theta_2)$, while the probability vector of buyer 2 is $\pi_2(\cdot, \theta_2, \theta_1) = (\pi(v_2 = 1, \theta_1, \theta_2), \pi(v_2 = 2, \theta_1, \theta_2)) = (1 - \alpha\theta_1 - \beta\theta_2, \alpha\theta_1 + \beta\theta_2)$. Condition 5 holds in this case, since $(\alpha\theta_1 + \beta, 1 - \alpha\theta_1 - \beta) \neq \rho_1(\alpha\theta_1 + 3\beta, 1 - \alpha\theta_1 - 3\beta)$ for all positive ρ_1 and $(1 - \alpha - \beta\theta_2, \alpha + \beta\theta_2) \neq \rho_2(1 - 3\alpha - 3\beta\theta_2, \alpha + 3\beta\theta_2)$ for all positive ρ_2; that is, the two payoff probability vectors of a player conditional on his own type are linearly independent. Now suppose that firm i observes one of three equally likely signals, $\Theta_i \in \{1, 2, 3\}$. If the probability density function $\pi(\omega|\theta)$ is as before, so are the allocation-probability vectors. Condition 5 does not hold in this case since, for example, $(\alpha\theta_1 + 2\beta, 1 - \alpha\theta_1 - 2\beta) = \frac{1}{2}(\alpha\theta_1 + \beta, 1 - \alpha\theta_1 - \beta) + \frac{1}{2}(\alpha\theta_1 + 3\beta, 1 - \alpha\theta_1 - 3\beta)$; the payoff probability vector of player 1 conditional on θ_1 and $\theta_2 = 2$ is a positive linear combination of the probability vector conditional on θ_1 and $\theta_2 = 1$ and the vector conditional on θ_1 and $\theta_2 = 3$.

5 Conclusions

Some authors have argued that the full surplus extraction results of Crémer and McLean (1985,1988) are counterintuitive. This paper doesn't take sides on this issue. The general point of the paper is theoretical: first, to show how two-stage revelation mechanisms can be used to expand the range of action of a seller; second, to show that the world in which agents' signals are independent random variables, but their valuations are interdependent, is very similar to a world in which agents' signals are correlated random variables. This makes intuitive sense. After all, interdependency of valuations is a form of correlation of the agents' preferences. This point has been missed in previous literature, because attention has been restricted to mechanisms in which all transfers must be made before agents observe their own allocation payoffs. Whenever it is possible to make transfers that are contingent on payoff observations, the seller can exploit the correlation between signals and allocation payoffs in a way that is very close to the way in which he can exploit any correlation among signals. In particular, the seller can use payments contingent on allocation-payoff observations to extract full surplus (or close to full surplus).

In at least two respects the full surplus extraction results when types are independent random variables and valuations are interdependent are weaker than when types are correlated random variables. First, in order to fully exploit the interdependency of valuations the seller must give the object to all buyers with small, but positive, probability. This introduces a small inefficiency and, as a result, a small loss of surplus to the seller. Second, in the case in which allocation payoffs do not perfectly reveal the signals of other players (Model R), full surplus extraction requires that the space of possible payoffs of a player has at least the same dimension as the space of signals of another player, a property which need not hold in important cases.

References

[1] Brusco, S., 1998: "Unique Implementation of the Full Surplus Extraction Outcome in Auctions with Correlated types," *Journal of Economic Theory*, 80, 185-200.

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

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

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

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

[6] Kosmopoulou, G. and S.R. Williams, 1998: "The Robustness of the Independent Private value Model in Bayesian Mechanism Design," *Economic Theory*, 12, 393-421.

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

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

[9] Mezzetti, C., 2004: "Mechanism Design with Interdependent Valuations: Efficiency," *Econometrica*, forthcoming.

[10] Myerson, R.B., 1981: "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58-73.

[11] Neeman, Z., 2003; "The Relevance of Private Information in Mechanism Design," http://econ.edu.bu/neeman/.