Optimal Information Disclosure in Auctions and The Handicap Auction*

Péter Eső†

Northwestern University, Kellogg MEDS

Balázs Szentes‡

University of Chicago, Department of Economics

September 2003

Abstract

We analyze the situation where a monopolist is selling an indivisible good to risk neutral buyers who only have an initial estimate of their private valuations. The seller can release (but cannot observe) signals, which affect the buyers' ex-post valuations. We show that in the expected revenue maximizing mechanism, the seller discloses all the information that she can, and her expected revenue is the same as if she could observe the released signals. We also show that this mechanism can be implemented by what we call a "handicap auction" in certain interesting applications (including one where the additional signal resolves an additive risk in the buyers' valuations). In the first round of this auction, each buyer privately buys a price premium from a menu published by the seller (a smaller premium costs more), then the seller releases the additional signals. In the second round, the buyers play a second-price auction, where the winner pays the sum of his premium and the second highest non-negative bid.

JEL Classification: C72, D44, D82, D83

*This version substantially generalizes and hence supersedes our earlier paper with the same title circulated without this notice. We thank Mark Armstrong, Yeon-koo Che, Kim-Sau Chung, Eddie Dekel, Drew Fudenberg, Peter Klibanoff, Marco Ottaviani, Alessandro Pavan, Andy Postlewaite, Phil Reny, Jim Schummer, Nancy Stokey, seminar participants at Northwestern University, the University of Chicago, the University of Pennsylvania, the University of Wisconsin-Madison, the Virginia Polytechnical Institue, the 2002 Midwest Theory Meeting, and the 2003 Summer Meeting of the Econometric Society for comments, and Norbert Maier for research assistance.

†2001 Sheridan Rd, Evanston, IL 60208. Email: eso@northwestern.edu.

‡1126 E. 59th St, Chicago, IL 60637. Email: szentes@uchicago.edu.

1 Introduction

In many examples of the monopolist's selling problem (optimal auctions),[1] the seller has considerable control over the accuracy of the buyers' information concerning their own valuations. Very often, the seller can decide whether the buyers can access information that refines their valuations; however, she either cannot observe these signals, or at least, she is unaware of their significance to the buyers. For example, the seller of an oil field or a painting can determine the number and nature of the tests the buyers can carry out privately (without the seller observing the results). Another example (due to Bergemann and Pesendorfer, 2002) is where the seller of a company has detailed information regarding the company's assets (e.g., its client list), but does not know how well these assets complement the assets of the potential buyers. Here, the seller can choose the extent to which she will disclose information about the firm's assets to the buyers. Sometimes the buyers' valuations become naturally more precise over time as the uncertainty of the good's value resolves, and the seller can decide how long to wait with the sale.

When the buyers' information acquisition is controlled by the seller, that process can also be optimized by the mechanism designer. In the present paper we explore the revenue maximizing mechanism for the sale of an indivisible good in a model where the buyers initially only have an estimate (signal) of their private valuations. The seller can costlessly release, without directly observing, additional private signals to the buyers that affect their valuations.[2] This model captures the common theme of the motivating examples: the seller controls, although cannot learn, private information that the buyers care about.

Our main result is that in the revenue-maximizing mechanism the seller will fully disclose all the information she has, and that her expected revenue will be as high as if she could observe the signals whose release she controls.[3] That is, the buyers will not enjoy additional informational rents from learning more about their ex-

[1]Early seminal contributions include Myerson (1981), Harris and Raviv (1981), Riley and Samuelson (1981), and Maskin and Riley (1984).

[2]In the formal model, a given buyer's ex-post valuation will be a general function of his original type and a shock. We only impose on this function certain "non-increasing returns" conditions.

[3]In this hypothetical benchmark situation the buyers may or may not observe (but in either case can ex post verify) the additional signals.

post valuations when the access to additional information is controlled by the seller. Besides these surprising findings, an added theoretical interest of our model is that the standard revelation principle cannot be applied, yet we are able to characterize the optimal mechanism.

We also exhibit a simple mechanism, dubbed the *handicap auction*, which implements the revenue-maximizing outcome in some interesting applications of the general model. The first application is one where the signals controlled by the seller indeed only *refine* the buyers' original value-estimates. We will argue that this situation is properly modeled as one where any given buyer's ex-post valuation is the sum of his original value-estimate and an independent noise term whose realization is the signal that the seller can disclose to him. (In this application the noise must be additive and independent otherwise the buyer's original value-estimate would be informative not only about his ex-post valuation, but also the precision of his original estimate.) In the second application each buyer's ex-post valuation is the realization of a normally distributed random variable. A given buyer's initial estimate and the seller's signal(s) are normally distributed, conditionally independent noisy observations of the true ex-post valuation. Note that in this "sampling" application a buyer's private information and his shock are strictly affiliated.[4]

The handicap auction, which implements the optimal mechanism in said applications, consists of two rounds. In the first round, each buyer buys a price premium from a menu provided by the seller (a smaller premium costs more). Then the seller releases, without observing, as much information as she can. In the second round, the buyers bid in a second-price auction, where the winner is required to pay his premium over the second highest non-negative bid. We call the whole mechanism a handicap auction because buyers compete under unequal conditions in the second round: a bidder with a smaller premium has an advantage.[5] For a single buyer, the handicap auction simplifies to a menu of buy-options (a schedule consisting of option fees as a function of the strike price), where the buyer gets to observe the additional signal after paying for the option of his choice.

[4] We thank Marco Ottaviani for suggesting that we develop an application along these lines.

[5] The handicap auction can also be implemented as a mechanism where, in the first round, each bidder buys a discount (larger discounts cost more), and then participates in a second price auction with a positive reservation price, where the winner's discount is applied towards his payment.

Our model nests the classical (independent private values) auction design problem as a special case, where the additional signals are identically zero. In this case, the handicap auction implements the outcome of the optimal auction of Myerson (1981) and Riley and Samuelson (1981).

Several papers have studied issues related to how buyers learn their valuations in auctions, and what consequences that bears on the seller's revenue, both from a positive and a normative point of view. One strand of the literature, see Persico (2000), Compte and Jehiel (2001) and the references therein, focuses on the buyers' incentives to acquire information in different auction formats. Our approach is different in that we want to *design* a revenue-maximizing mechanism in which the seller has the opportunity to costlessly release (without observing) information to the buyers. In our model, it is the seller (not the buyers) who controls how much information the buyers acquire.

The issue of the informed party acquiring private information over time arises in certain dynamic principal-agent models as well. In Baron and Besanko (1984), Riordan and Sappington (1987), and Courty and Li (2000), the principal and the agent are contracting over two periods, and the agent learns payoff-relevant private information in both periods. These authors then derive the optimal two-stage revelation mechanism where the contract is signed in the first period (when the agent only knows his first-period type). In contrast, in our paper, we solve for the optimal mechanism in a multi-agent auction environment where the seller can *decide* whether or not the buyers receive additional private signals.

Information disclosure by the seller in an auction has been studied in the context of the winner's curse and the linkage principle by Milgrom and Weber (1982). They investigate whether in traditional auctions the seller should commit to disclose public signals that are affiliated with the buyers' valuations. They find that the seller gains from committing to full disclosure, because that reduces the buyers' fear of overbidding, thereby increasing their bids and hence the seller's revenue. Our problem differs from this classic one in many aspects. Most importantly, in our setting, the signals that the seller can release are private (not public) signals, in the sense that each signal affects the valuation of a single buyer and can be disclosed to that buyer only. The seller will gain from the release of information (which she does not even observe) not because of the linkage principle, but because the information can potentially improve efficiency, and she can charge for the access.

Our motivation is closer to that of Bergemann and Pesendorfer (2002). They consider the task of designing an information disclosure policy for the seller that allows to extract the most revenue in a subsequent auction. Their problem is very different from ours in that the seller is not allowed to *charge* for the release of information. Their model also differs from ours in that the buyers do not have private information at the beginning of the game. Under these assumptions, Bergemann and Pesendorfer (2002) show that the information structure that allows the seller to design the auction with the largest expected revenue is necessarily imperfect: in this structure, buyers are only allowed to learn which element of a finite partition their valuation falls into.

In contrast, in our paper, we design the expected revenue maximizing mechanism where the information structure and the rules of transaction *together* are chosen optimally. The difference may first seem subtle, it is important nevertheless. What we assume is that the seller can integrate the rules of information acquisition into the mechanism used for the sale itself. For example, in our model, the seller can charge the buyers for getting more and more accurate signals (perhaps in several rounds); the buyers could even be asked to bid for obtaining more information.

The idea that "selling" the access to information may be advantageous for the seller can be easily illustrated by an example. Suppose that there are two buyers who are both unaware of their valuation (drawn independently from the same distribution), which the seller can allow them to learn. Then consider the following mechanism. The seller charges both buyers an entry fee, which equals half of the expected difference of the maximum and minimum of two independent draws of the value-distribution. In exchange, she allows the buyers to observe their valuations (after they have paid the entry fee), and makes them play an ordinary second-price auction. The second-price auction will be efficient, and the buyers' ex-ante expected profit exactly equals the upfront entry fee. The seller ends up appropriating the entire surplus by charging the buyers for observing their valuations.[6]

[6]This example has been studied independently by Gershkov (2002), who also obtained the same result. In the context of a two-period principal-agent problem, Baron and Besanko (1984) make the same point: if the agent's type in the second period is independent of his type in the first period then a period one (constant and efficient) contract gives the whole period-two surplus to the principal.

This simple solution—the seller committing to the efficient allocation, revealing the additional signals, and charging an entry fee equal to the expected efficiency gains—only works when the buyers do not have private information to start with. Otherwise (for example, if the buyers privately observe signals, but their valuations also depend on other signals that they may see at the seller's discretion), the auctioneer, as we will show, does not want to commit to an efficient auction in the continuation, so the previously proposed mechanism does not work. We have to find a more sophisticated auction, and this is exactly what we will do in the remainder of the paper.

The paper is structured as follows. In the next section we outline the model and introduce the necessary notation. In Section 3 we characterize the revenue maximizing mechanism. In Section 4 a simple implementation (via a "handicap auction") is discussed for an important special case of the general model. Section 5 concludes. Proofs are collected in an Appendix.

2 The Model

2.1 The environment

There are n potential buyers, $i = 1, ..., n$, for an indivisible good. Each buyer i has a type, v_i, which is a random variable distributed on some interval $[\underline{v}, \bar{v}]$ with a positive density $f_i = F_i'$. The type distributions satisfy the monotone hazard rate condition, that is, $f_i/(1 - F_i)$ is weakly increasing. Buyer i's valuation for the good is $u_i(v_i, s_i)$, where s_i (called i's shock) is drawn from $(-\infty, \infty)$ according to an atomless distribution G_i. We will use v to denote the vector of types and s to denote the vector of shocks. We will also use the usual shorthand notation for the vector of types of buyers other than i, v_{-i}, and let s_{-i} denote $(s_j)_{j \neq i}$.

We assume that all shocks and types are independent random variables. If i's shock is correlated with v_i, e.g., his valuation is $\tilde{u}_i(v_i, \tilde{s}_i)$ where $\tilde{s}_i$ has cdf $G_i(\cdot|v_i)$, then we can consider a transformed problem where we re-define the shock as $s_i = G_i(\tilde{s}_i|v_i)$ and write $u_i(v_i, s_i) = \tilde{u}_i(v_i, G_i^{-1}(s_i|v_i))$. Clearly, s_i is independent of v_i, hence assuming $v_i \perp s_i$ is without loss of generality (as long as u_i can be any function). The independence of information *across buyers* is a standard assumption that rules out Crémer-McLean (1988) full rent extracting mechanisms.

The realization of v_i is observed by buyer i. Although neither the seller nor buyer i can directly observe i's shock, the seller has the ability to generate signals conditional on s_i, which only buyer i will observe. In particular, we assume that the seller can allow buyer i to observe his shock without the seller learning its value.[7]

As far as the ex-post valuation functions are concerned, for all i, we assume that u_i is strictly increasing and twice differentiable, and denote the partials by $u_{i1} = \partial u_i/\partial v_i$ and $u_{i2} = \partial u_i/\partial s_i$. (Similarly, $u_{i11} = \partial^2 u_i/\partial v_i^2$ and $u_{i12} = \partial^2 u_i/\partial v_i \partial s_i$.) Assume that $\int_{s_i} u_{i1}(v_i, s_i) dG_i(s_i) < \infty$, that is, the expected marginal value of i's type is finite.[8] Moreover, we make two assumptions regarding the marginal returns of the signals. First, $u_{i12} \leq 0$, which means that the marginal value of the shock is non-increasing in i's type. Second, $u_{i11}/u_{i1} \leq u_{i12}/u_{i2}$, which says that an increase in i's type, holding the ex-post valuation constant, weakly decreases the type's marginal value.[9] Purely for convenience, suppose that for all v_i, s_i, and $\hat{v}_i$, there exists $\hat{s}_i$ such that $u_i(v_i, s_i) = u_i(\hat{v}_i, \hat{s}_i)$.

2.2 Two applications

The leading application (or special case) of our general model is the following. Buyer i's original private information, v_i, is his estimated valuation for the good; his ex-post valuation differs from this by an additive and independent noise, s_i. That is, i's ex-post valuation is

$$u_i(v_i, s_i) = v_i + s_i,$$

where s_i is independent of v_i, and all signals are independent across i's. The seller is able to resolve the uncertainty in buyer i's valuation by disclosing without observing the realization of s_i.

Since s_i is independent of v_i, the buyer's original private information pertains only to the expected value of the good for him; the knowledge of v_i conveys no

[7]In many applications, the seller may not be able to generate just *any* random signal correlated with s. Therefore, all we assume is that the seller can "show" s_i to buyer i if she so decides.

[8]This assumption is made because the support of s_i may be unbounded.

[9]To better understand the last condition note that the total differential of u_{i1} (the change in the marginal value of i's type) is $u_{i11} dv_i + u_{i12} ds_i$. Keeping u_i constant (moving along an "iso-value" curve) means $ds_i = -u_{i1}/u_{i2} dv_i$. Substituting it into the total differential of u_{i1} yields $(u_{i11} - u_{i12} u_{i1}/u_{i2}) dv_i$. This is non-positive for $dv_i > 0$ if and only if $u_{i11}/u_{i1} \leq u_{i12}/u_{i2}$.

information about how precise (t)his estimate is. This is the cleanest model of the problem that motivated our research in the first place: the seller can release, without observing, signals that *refine* the buyers' original value-estimates. In the previous subsection we introduced the more general model (where v_i and s_i can interact in i's valuation) in order to demonstrate that the additive-independent signal structure does not play a major role in the main results.

Another special case of the general model is inspired by a practical application. Suppose that buyer i's ex-post valuation is the realization of a random variable, u_i, with prior $\mathcal{N}(\bar{u}_i, 1/\tau_{u_i})$. Let v_i and $\tilde{s}_i$ be a random sample drawn from a normal distribution with mean u_i and precisions τ_{v_i} and τ_{s_i}, respectively. (The support of v_i is the real line instead of a bounded interval, but this does not cause a problem in the analysis.) Suppose that buyer i originally observes v_i, but the seller can allow him to observe the second signal ($\tilde{s}_i$) as well. Then, the buyer's posterior on his valuation is Normal with mean $E[u_i|v_i] = (\tau_{u_i}\bar{u}_i + \tau_{v_i}v_i)/(\tau_{u_i} + \tau_{v_i})$, i.e., the precision-weighted average of the prior mean and v_i, and precision $\tau_{u_i} + \tau_{v_i}$. This implies that $\tilde{s}_i$ conditional v_i is also normally distributed, with expectation $E[u_i|v_i]$ and variance $1/(\tau_{u_i} + \tau_{v_i}) + 1/\tau_{s_i}$. Clearly, $\tilde{s}_i$ and v_i are strictly affiliated; denote the cdf corresponding to $\tilde{s}_i$'s conditional distribution by $G_i(\cdot|v_i)$.

In this application, buyer i's expected ex-post valuation given v_i and $\tilde{s}_i$ is

$$\tilde{u}_i(v_i, \tilde{s}_i) = E[u_i|v_i, \tilde{s}_i] = \frac{\tau_{u_i}\bar{u}_i + \tau_{v_i}v_i + \tau_{s_i}\tilde{s}_i}{\tau_{u_i} + \tau_{v_i} + \tau_{s_i}}.$$

Transform $\tilde{s}_i$ (in order to make the noise independent of v_i) by letting $s_i = G_i(\tilde{s}_i|v_i)$, and $u_i(v_i, s_i) = \tilde{u}_i(v_i, G_i^{-1}(s_i|v_i))$. Obviously, u_i is strictly increasing in both arguments. Moreover, it can be shown (contact either author for details) that

$$\begin{aligned}\frac{\partial}{\partial v_i}u_i(v_i, s_i) &= -\frac{\partial G_i(\tilde{s}_i|v_i)/\partial v_i}{g_i(\tilde{s}_i|v_i)}\frac{\partial}{\partial \tilde{s}_i}\tilde{u}_i(v_i, \tilde{s}_i)\bigg]_{\tilde{s}_i = G_i^{-1}(s_i|v_i)} \\ &= \frac{\tau_{v_i}}{\tau_{u_i} + \tau_{v_i}}\frac{\tau_{s_i}}{\tau_{u_i} + \tau_{v_i} + \tau_{s_i}},\end{aligned}$$

which is a constant between 0 and 1. Therefore u_i is linear in v_i, and so all of our assumptions on u_i are satisfied. (For example, if u_i's prior is improper diffuse, i.e., $\tau_{u_i} = 0$, and v_i and s_i have the same precision, then the formula yields $u_{i1} = 1/2$.)

2.3 The mechanism design problem

All parties are risk neutral. The seller's valuation for the good is zero and wants to maximize her expected revenue. Each buyer's utility is the negative of his payment to the seller, plus, in case he wins, the value of the object. The value of the buyers' outside option is zero.

The seller can design any (indirect) mechanism, which can consist of several rounds of communication between the parties (i.e., sending of messages according to rules specified by the seller); the seller may also generate and release (without observing) certain signals to each buyer i, including the realization of his shock, s_i. Transfers of the good and money may also occur as a function of the history. The set of all indirect mechanisms is rather complex, and the standard revelation principle cannot be applied in the quest for the optimal mechanism. However, this issue is avoided by the approach that we take in the next section.

3 The Optimal Mechanism

Our main result is the characterization of the expected revenue maximizing mechanism in the model introduced in Section 2. We will also show that this mechanism achieves the same expected revenue *as if* the seller could observe the realizations of the shocks. In other words, while the buyers still enjoy information rents from their types, all their rents from observing the shocks can be appropriated by the seller.

Our strategy to find the optimal selling mechanism incorporating information disclosure is the following. We will first characterize the optimal mechanism under the assumption that the seller, after having committed to a mechanism, is able to observe the realizations of the shocks. Clearly, the seller's maximum revenue in this "benchmark" case is an upper bound on her revenue when she cannot observe the shocks. Then, we will show that the same allocation rule and transfers can be implemented even if the seller cannot directly observe the shocks, but can control their release to the buyers. The constructive proof yield the optimal mechanism in our original model.

Suppose first, for benchmarking purposes only, that the seller can observe the

s_i's after having committed to a selling mechanism.[10] The Revelation Principle applies, hence we can restrict attention to mechanisms where the buyers report their types, and the seller determines the allocation and the transfers as a function of the reported types and the realization of the shocks. We will analyze truthful equilibria of direct mechanisms that consist of an allocation rule, $x_i^*(v_i, v_{-i}, s_i, s_{-i})$ for all i, and an (expected) transfer scheme, $t_i^*(v_i, v_{-i}, s_i, s_{-i})$ for all i. Here, $x_i^*(v_i, v_{-i}, s_i, s_{-i})$ is the probability that buyer i receives the good, and $t_i^*(v_i, v_{-i}, s_i, s_{-i})$ is the transfer that he expects to pay, given the reported types and the realization of the shocks.

Fix a direct mechanism in the benchmark case, and define

$$X_i^*(v_i, s_i) = \int_{s_{-i}} \int_{v_{-i}} x_i^*(v_i, v_{-i}, s_i, s_{-i})\, dF_{-i}(v_{-i}) dG_{-i}(s_{-i}), \tag{1}$$

which is buyer i's expected probability of winning in the mechanism when his type is v_i and the realization of his shock is s_i, and all buyers report their types truthfully. If buyer i with type v_i reports type $\hat{v}_i$ then his profit is,

$$\pi_i^*(v_i, \hat{v}_i) = \int_s \int_{v_{-i}} \left[x_i^*(\hat{v}_i, v_{-i}, s_i, s_{-i}) u_i(v_i, s_i) - t_i^*(\hat{v}_i, v_{-i}, s_i, s_{-i})\right] dF_{-i}(v_{-i}) dG(s). \tag{2}$$

Incentive compatibility of the mechanism is equivalent to,

$$\pi_i^*(v_i, \hat{v}_i) \leq \pi_i^*(v_i, v_i), \quad \text{for all } i \text{ and all } v_i, \hat{v}_i \in [\underline{v}, \bar{v}]. \tag{3}$$

Define $\Pi_i^*(v_i) \equiv \pi_i^*(v_i, v_i)$ as the indirect profit function of buyer i. The seller's expected revenue can be written as the difference of the social surplus and the the buyers' expected profits as,

$$\sum_{i=1}^{n} \int_s \int_v \left[u_i(v_i, s_i) x_i^*(v_i, v_{-i}, s_i, s_{-i}) - \Pi_i^*(v_i)\right] dF(v) dG(s). \tag{4}$$

The benchmark problem is to maximize (4) subject to a feasibility constraint, $\sum_{i=1}^{n} x_i(v_i, v_{-i}, s_i, s_{-i}) \leq 1$ for all v, s, the buyers' incentive compatibility con-

[10] Since the seller observes the shocks only after having committed to a mechanism, "informed principal" type problems do not arise in the benchmark. On the other hand, it does not matter whether the buyers can observe the shocks as long as the mechanism is verifiable.

straints, (3), and their participation constraints, $\Pi_i^*(v_i) \geq 0$ for all i and $v_i \in [\underline{v}, \bar{v}]$.

Using the tools of Bayesian mechanism design, we obtain the following solution to the benchmark problem. The (standard) proof is found in the Appendix.

Lemma 1 *In the revenue-maximizing mechanism of the benchmark case (when the seller can observe the s_i's after having committed to a selling mechanism), the allocation rule sets $x_i^*(v_i, v_{-i}, s_i, s_{-i}) = 1$ for the buyer with the highest non-negative "shock-adjusted virtual valuation,"*

$$W_i(v_i, s_i) = u_i(v_i, s_i) - \frac{1 - F_i(v_i)}{f_i(v_i)} u_{i1}(v_i, s_i) \tag{5}$$

(ties broken randomly), and the profit of buyer i with type v_i is

$$\Pi_i^*(v_i) = \int_{\underline{v}}^{v_i} \int_{s_i} u_{i1}(y, s_i) X_i^*(y, s_i) dG_i(s_i) dy, \tag{6}$$

where X_i^* is defined by the optimal allocation rule and (1).

From now on, we will refer to the revenue-maximizing mechanism of the benchmark case (characterized in the previous lemma) as $\{x_i^*, t_i^*\}_{i=1}^n$, and let X_i^* and Π_i^* denote i's expected probability of winning and profit functions, respectively. It is useful, for use in subsequent steps of the analysis, to further describe some properties of the allocation rule.

Corollary 1 *X_i^* induced by the optimal allocation rule and (1) is*
(i) continuous,
(ii) weakly increasing in both arguments, and
(iii) weakly increasing in v_i even if s_i is adjusted to keep $u_i(v_i, s_i)$ constant.

Continuity follows because distributions are atomless. X_i^* is weakly increasing in v_i and s_i because (5) is strictly increasing in both variables, which in turn follows from the monotone hazard rate condition, $u_{i1} > 0$, $u_{i11} \leq 0$, $u_{i2} > 0$, and $u_{i12} \leq 0$. Finally, if $v_i > \hat{v}_i$ and $s_i < \hat{s}_i$ with $u_i(v_i, s_i) = u_i(\hat{v}_i, \hat{s}_i)$ then $X_i^*(v_i, s_i) \geq X_i^*(\hat{v}_i, \hat{s}_i)$. To see this, compare (5) for (v_i, s_i) and $(\hat{v}_i, \hat{s}_i)$: the first term is the same by assumption, in the second term $(1 - F_i(v_i))/f_i(v_i) \leq (1 - F_i(\hat{v}_i))/f_i(\hat{v}_i)$, and $u_{i1}(v_i, s_i) \leq u_{i1}(\hat{v}_i, \hat{s}_i)$ because $u_{i11}/u_{i1} \leq u_{i12}/u_{i2}$ (see footnote 9).

We will now show that the (benchmark) allocation rule and indirect profit function of Lemma 1 can be implemented even if the seller cannot observe the shocks, as long as she can allow the buyers to observe them. For this implementation we will use a rather restricted class of mechanisms, in which first buyers report their v_i's, then observe their own s_i's and report them back—so that telling the truth is incentive compatible in both reporting stages—, and finally the seller allocates the good according to $x_i(v_i, v_{-i}, s_i, s_{-i})$ and the buyers pay $t_i(v_i, v_{-i}, s_i, s_{-i})$. Clearly, the class of these mechanisms does not exhaust or represent all possible mechanisms; however, since we will be able to replicate the outcome of the benchmark even in this restricted class, we need not look any further.[11]

Given such a two-stage reporting mechanism $\{x_i, t_i\}_{i=1}^n$, define

$$
\begin{aligned}
X_i(v_i, s_i) &= \int_{s_{-i}} \int_{v_{-i}} x_i(v_i, v_{-i}, s_i, s_{-i})\, dF_{-i}(v_{-i}) dG_{-i}(s_{-i}), \\
T_i(v_i, s_i) &\equiv \int_{s_{-i}} \int_{v_{-i}} t_i(v_i, v_{-i}, s_i, s_{-i})\, dF_{-i}(v_{-i}) dG_{-i}(s_{-i}).
\end{aligned}
$$

These quantities correspond to buyer i's expected probability of winning and expected transfers, respectively, when he has type v_i and shock s_i, and everybody (including i) reports truthfully in both rounds of the two-stage mechanism. We will now analyze the consequences of incentive compatibility going backwards, starting in the second stage of the mechanism.

In the second reporting stage, after truthful first round, buyer i with type v_i who observes s_i and reports $\hat{s}_i$ gets

$$\tilde{\pi}_i(s_i, \hat{s}_i; v_i) = u_i(v_i, s_i) X_i(v_i, \hat{s}_i) - T_i(v_i, \hat{s}_i). \tag{7}$$

Incentive compatibility in the second reporting stage is equivalent to

$$\tilde{\pi}_i(s_i, \hat{s}_i; v_i) \leq \tilde{\pi}_i(s_i, s_i; v_i) \quad \text{for all } i,\ v_i,\ s_i, \text{ and } \hat{s}_i. \tag{8}$$

[11] In contrast, two-stage revelation mechanisms can be (and are indeed) used to represent *all* possible indirect mechanisms in dynamic principal-agent models where contracts are signed in the first of two periods (e.g., Baron and Besanko (1984), Courty and Li (2000)) because there the Revelation Principle holds. This is not so in our model because the seller does not have to disclose (or may only partially disclose) the shocks to the buyers.

The following lemma provides conditions for a mechanism to be incentive compatible in the second reporting stage after a truthful first round.

Lemma 2 *If a two-stage mechanism is incentive compatible and $X_i(v_i, s_i)$ induced by the allocation rule is continuous in s_i then for all $s_i > \hat{s}_i$,*

$$\tilde{\pi}_i(s_i, s_i; v_i) - \tilde{\pi}_i(\hat{s}_i, \hat{s}_i; v_i) = \int_{\hat{s}_i}^{s_i} u_{i2}(v_i, z) X_i(v_i, z) dz. \tag{9}$$

Moreover, if (9) holds and X_i is weakly increasing in s_i then the two-stage mechanism is incentive compatible in the second round after truthful revelation in the first round.

In order to complete the analysis of the second round of the mechanism, we also need to know what buyer i will do if he misreports his type in the first round. The following lemma claims that he will "correct" his mistake by reporting a shock such that his ex-post valuation at the reported type and shock coincides with his true ex-post valuation.

Lemma 3 *In the second round of an incentive compatible two-stage mechanism, buyer i with type v_i who reported $\hat{v}_i$ in the first round and has observed s_i will report $\hat{s}_i = \sigma_i(v_i, \hat{v}_i, s_i)$ such that*

$$u_i(v_i, s_i) \equiv u_i(\hat{v}_i, \sigma_i(v_i, \hat{v}_i, s_i)). \tag{10}$$

Let us move back to the first round, and compute buyer i's expected profit there. When i's type is v_i but reports $\hat{v}_i$ his profit is

$$\pi_i(v_i, \hat{v}_i) = \int_{s_i} [u_i(v_i, s_i) X_i(\hat{v}_i, \sigma_i(v_i, \hat{v}_i, s_i)) - T_i(\hat{v}_i, \sigma_i(v_i, \hat{v}_i, s_i))] dG_i(s_i).$$

Using (10) we have,

$$\pi_i(v_i, \hat{v}_i) = \int_{s_i} \tilde{\pi}_i(\sigma_i(v_i, \hat{v}_i, s_i), \sigma_i(v_i, \hat{v}_i, s_i); \hat{v}_i) dG_i(s_i). \tag{11}$$

Suppose $v_i > \hat{v}_i$. Note that by (10) and the monotonicity of u_i,

$$\sigma_i(v_i, \hat{v}_i, s_i) > s_i > \sigma_i(\hat{v}_i, v_i, s_i).$$

By (9), we can rewrite $\pi_i(v_i, \hat{v}_i)$ in (11) as

$$\pi_i(v_i, \hat{v}_i) = \int_{s_i} \left[\tilde{\pi}_i(s_i, s_i; \hat{v}_i) + \int_{z=s_i}^{\sigma_i(v_i,\hat{v}_i,s_i)} u_{i2}(\hat{v}_i, z) X_i(\hat{v}_i, z) dz \right] dG_i(s_i),$$

which becomes, by $\pi_i(\hat{v}_i, \hat{v}_i) = \int_{s_i} \tilde{\pi}_i(s_i, s_i, \hat{v}_i) dG_i(s_i)$,

$$\pi_i(v_i, \hat{v}_i) = \pi_i(\hat{v}_i, \hat{v}_i) + \int_{s_i} \int_{z=s_i}^{\sigma_i(v_i,\hat{v}_i,s_i)} u_{i2}(\hat{v}_i, z) X_i(\hat{v}_i, z) dz \, dG_i(s_i). \tag{12}$$

By differentiating (10) in v_i, we have

$$u_{i1}(v_i, s_i) = u_{i2}(\hat{v}_i, \sigma_i(v_i, \hat{v}_i, s_i)) \, \sigma_{i1}(v_i, \hat{v}_i, s_i). \tag{13}$$

Note that σ_i defined by (10) is continuous and monotonic. Hence the image of $\sigma_i(y, \hat{v}_i, s_i)$ on $y \in [\hat{v}_i, v_i]$ is $[s_i, \sigma_i(v_i, \hat{v}_i, s_i)]$, and thus we can change the variable of the inside integration in (12) from $z \in [s_i, \sigma_i(v_i, \hat{v}_i, s_i)]$ to $y \in [\hat{v}_i, v_i]$ to get

$$\pi_i(v_i, \hat{v}_i) = \pi_i(\hat{v}_i, \hat{v}_i) + \int_{s_i} \int_{y=\hat{v}_i}^{v_i} u_{i2}(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) X_i(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) \sigma_{i1}(y, \hat{v}_i, s_i) dy \, dG_i(s_i).$$

Therefore, by (13),

$$\pi_i(v_i, \hat{v}_i) = \pi_i(\hat{v}_i, \hat{v}_i) + \int_{s_i} \int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) dy \, dG_i(s_i). \tag{14}$$

Similarly, when type $\hat{v}_i$ reports $v_i > \hat{v}_i$, his payoff is

$$\pi_i(\hat{v}_i, v_i) = \pi_i(v_i, v_i) - \int_{s_i} \int_{z=\sigma_i(\hat{v}_i,v_i,s_i)}^{s_i} u_{i2}(v_i, z) X_i(v_i, z) dz \, dG_i(s_i).$$

By a change of variable in the inside integration and (13),

$$\pi_i(\hat{v}_i, v_i) = \pi_i(v_i, v_i) - \int_{s_i} \int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(v_i, \sigma_i(y, v_i, s_i)) dy \, dG_i(s_i). \tag{15}$$

Incentive compatibility in the first round is equivalent to, for all $v_i > \hat{v}_i$,

$$\pi_i(v_i, \hat{v}_i) \leq \pi_i(v_i, v_i) \text{ and } \pi_i(\hat{v}_i, v_i) \leq \pi_i(\hat{v}_i, \hat{v}_i). \tag{16}$$

Using (14) and (15), incentive compatibility in the first round is equivalent to, for all $v_i > \hat{v}_i$,

$$\begin{aligned}\int_{s_i} \frac{\int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) dy}{v_i - \hat{v}_i} dG_i(s_i) \\ \leq \frac{\pi_i(v_i, v_i) - \pi_i(\hat{v}_i, \hat{v}_i)}{v_i - \hat{v}_i} \\ \leq \int_{s_i} \frac{\int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(v_i, \sigma_i(y, v_i, s_i)) dy}{v_i - \hat{v}_i} dG_i(s_i). \quad (17)\end{aligned}$$

By $X_i \leq 1$ and the concavity of u_i in v_i,

$$\frac{\int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) dy}{v_i - \hat{v}_i} \leq \frac{\int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) dy}{v_i - \hat{v}_i} \leq u_{i1}(\hat{v}_i, s_i),$$

and by assumption $u_{i1}(\hat{v}_i, s_i)$ has a finite expectation with respect to s_i. Therefore, by the Lebesgue Convergence Theorem,

$$\begin{aligned}\lim_{v_i \to \hat{v}_i} \int_{s_i} \frac{\int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) dy}{v_i - \hat{v}_i} dG_i(s_i) \\ = \int_{s_i} \lim_{v_i \to \hat{v}_i} \frac{\int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) dy}{v_i - \hat{v}_i} dG_i(s_i) \\ = \int_{s_i} u_{i1}(\hat{v}_i, s_i) X_i(\hat{v}_i, s_i) dG_i(s_i).\end{aligned}$$

By analogous reasoning,

$$\begin{aligned}\lim_{\hat{v}_i \to v_i} \int_{s_i} \frac{\int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(v_i, \sigma_i(y, v_i, s_i)) dy}{v_i - \hat{v}_i} dG_i(s_i) \\ = \int_{s_i} u_{i1}(v_i, s_i) X_i(v_i, s_i) dG_i(s_i).\end{aligned}$$

Therefore, if X_i is continuous in v_i, then $\Pi_i(v_i) \equiv \pi_i(v_i, v_i)$ is differentiable everywhere by (17), and

$$\Pi_i'(v_i) = \int_{s_i} u_{i1}(v_i, s_i) X_i(v_i, s_i) dG_i(s_i).$$

Since this derivative is finite for all v_i, Π_i is Lipschitz-continuous and hence it can be recovered from its derivative,

$$\Pi_i(v_i) = \Pi_i(\underline{v}) + \int_{\underline{v}}^{v_i} \int_{s_i} u_{i1}(y, s_i) X_i(y, s_i) dG_i(s_i) dy. \tag{18}$$

We conclude that if the optimal allocation rule (which is continuous) can be implemented with $\Pi_i(\underline{v}) = 0$ then $\Pi_i(v_i) = \Pi_i^*(v_i)$, and so the seller's revenue is the same as in the benchmark. The following result (which is the main result of the paper) shows that this is indeed the case.

Proposition 1 *In the model where the seller can disclose but cannot observe the buyers' shocks, the revenue-maximizing auction can be implemented as a mechanism where the buyers report their types, the seller discloses the shocks which the buyers report back, and the good is allocated to the buyer with the highest $W(v_i, s_i)$ defined in (5). The buyers' profits are the same as in (6), and the seller's expected revenue is the same as if she could observe the realizations of the shocks.*

In order to prove this statement we need to show is that when the allocation rule coincides with the optimal allocation rule of the benchmark case, the induced period-one profit functions, $\pi_i(v, \hat{v}_i)$ for all i, make untruthful type-reports unprofitable for the buyers. In the formal proof (relegated to the Appendix) we indeed show that when X_i set equal to X_i^* for all i in (14) and (15), the incentive compatibility constraints for reporting v_i when i's type is v_i hold, that is, (16) is satisfied.

The key feature of the optimal allocation rule that makes the proof work is that for all $v_i, \hat{v}_i \in [\underline{v}, \bar{v}]$ and $s_i, \hat{s}_i \in \mathbb{R}$ such that $v_i > \hat{v}_i$ and $u_i(v_i, s_i) = u_i(\hat{v}_i, \hat{s}_i)$, the allocation rule "favors" the pair (v_i, s_i), that is, $X_i^*(\hat{v}_i, \hat{s}_i) \leq X_i^*(v_i, s_i)$, as seen in property (iii) in Corollary 1. In words, buyer i with type v_i and a given ex-post valuation wins the object more often than he does with a lower type $\hat{v}_i$ but the same ex-post valuation. This property is a consequence of the monotone hazard rate condition and our assumptions on u_i (in particular that $u_{i1}(v_i, s_i)$ decreases when v_i increases but s_i decreases to keep $u_i(v_i, s_i)$ the same). This condition is not necessary for the incentive compatibility of the same allocation rule in the benchmark case (this can be seen from the proof of Lemma 1), but it is crucial in Proposition 1 (if it does not hold then one can easily construct an example where the two-stage mechanism is not incentive compatible).

4 The Handicap Auction

In this section we show how to implement the optimal mechanism derived in Section 3 when u_{i1} is constant (as it is in the two applications discussed in Section 2.2).

4.1 The rules of the auction

A handicap auction consists of two rounds. In the first round, each buyer i, knowing his type, chooses a price premium p_i for a fee $C_i(p_i)$, where C_i is a fee-schedule published by the seller. The buyers do not observe the premia chosen by others. The second round is a traditional auction, and the winner is required to pay his premium over the price. Between the two rounds, the seller may send messages to the buyers. In our model, the seller will allow every buyer to learn the realization of his shock between the two rounds, and the second round is a second price (or English) auction with a zero reservation price.

We call this mechanism a handicap auction because in the second round, the buyers compete under unequal conditions: a bidder with a smaller premium has a clear advantage. An interesting feature of our auction is that the bidders *buy* their premium in the initial round, which allows for some form of price discrimination. We will come back to the issue of price discrimination.

An interesting alternative way of formulating the rules of the handicap auction would be by using price discounts (or rebates) instead of price premia. In this version, each bidder first has to buy a discount from a schedule published by the seller. Then the buyers are allowed to learn the realizations of the shocks, and are invited to bid in a second price auction with a reservation price r, where the winner's discount is applied towards his payment. The reader can check that a handicap auction can be easily transformed into a mechanism like this by setting r sufficiently high (larger than the highest p_i in the original fee-schedules), and specifying that a discount $d_i = r - p_i$ is sold for a price $C(p_i)$ in the first round. In what follows, however, we will use the original form of the handicap auction.

If there is only a single buyer, then the handicap auction simplifies to a *menu of buy options*: p_i can be thought of as the strike price, and the upfront fee, $C_i(p_i)$, is the cost of the option. In the second round, the buyer can exercise his option to buy at price p_i (there is no other bidder, so the second-highest bid is zero), for which he initially paid a fee of $C_i(p_i)$. We will revisit the single-buyer case later in the context of a numerical example.

Clearly, buyer i who chose premium p_i in the first round, has type v_i and observes s_i before the second round bids $u_i(v_i, s_i) - p_i$ in the (handicapped) second-price auction. (This is so because in a second-price auction buyers bid their surplus conditional on winning in a tie.) Assuming that the buyers follow this weakly dominant strategy in the second round, the handicap auction can be represented by pairs of functions, $p_i : [\underline{v}, \bar{v}] \to \mathbb{R}$ and $c_i : [\underline{v}, \bar{v}] \to \mathbb{R}$, for $i = 1, ..., n$, where $p_i(v_i)$ is the price premium that type v_i chooses (in equilibrium) for the fee of $c_i(v_i) \equiv C_i(p_i(v_i))$.

Proposition 2 *If u_{i1} is constant then the optimal mechanism of Proposition 1 can be implemented via a handicap auction $\{c_i, p_i\}_{i=1}^n$, where*

$$p_i(v_i) = \frac{1 - F_i(v_i)}{f_i(v_i)} u_{i1}, \tag{19}$$

and $c_i(v_i)$ is defined by

$$\begin{aligned} c_i(v_i) = & \int_s \int_{v_{-i}} (W_i(v_i, s_i) - W_{-i}^0(v_{-i}, s_{-i})) \mathbf{1}_{\{W_i(v_i,s_i) \geq W_{-i}^0(v_{-i},s_{-i})\}} dF_{-i}(v_{-i}) dG(s) \\ & - \int_s \int_{v_{-i}} \int_{\underline{v}}^{v_i} u_{i1} \mathbf{1}_{\{W_i(y,s_i) \geq W_{-i}^0(v_{-i},s_{-i})\}} dy \, dF_{-i}(v_{-i}) dG(s), \end{aligned} \tag{20}$$

where $W_{-i}^0(v_{-i}, s_{-i}) = \max_{j \neq i}\{W_j(v_j, s_j), 0\}$.

4.2 Discussion of the handicap auction

From the seller's perspective, the premium–fee schedule offered in the first round of the handicap auction works as a device to discriminate among buyers with different value estimates. When a buyer decides to participate in the handicap auction, he knows his type (expected valuation), which tells him whether he is more or less likely to win. Therefore, in the first round, a buyer with a high type chooses a small price premium for a large fee in order not to pay much when he wins. Using analogous reasoning, low types choose large price premia, which are cheaper, but make winning more expensive.

It is interesting to observe that in the optimal handicap auction (and, in general, in the optimal mechanism), two buyers with the same ex-post valuation do not have the same probability of winning. The buyer with the larger v_i will choose a smaller price premium, bid higher in the second round, and will be more likely to win. This shows that the auction does not achieve full ex post efficiency, even under ex ante symmetry of the bidders and conditional on the object being sold.[12]

In order to better explain our main result (Proposition 1), consider a setup where the buyers are ex ante symmetric (the v_i's are identically distributed), and the shocks are additive, mean zero random variables (hence Proposition 2 applies). Let us compare the optimal allocation rule in the case when the seller can observe the shocks with that of the revenue maximizing auction when *nobody* (neither the seller nor the buyers) can observe them. In the latter case, the seller should allocate the good to the buyer with the largest non-negative virtual value-estimate, $v_i - (1 - F(v_i))/f(v_i)$. If the seller can observe the shocks, then, in the optimal mechanism, the good will be allocated more efficiently, as the winner will now be the buyer with the highest non-negative shock-adjusted virtual valuation, $v_i + s_i - (1 - F(v_i))/f(v_i)$.[13] According to Propositions 1 and 2, the seller, by controlling the release of the shocks and without actually observing them, can implement the same allocation, and surprisingly, can appropriate the increase in efficiency.[14]

One may suggest that the way the seller can appropriate all rents from the additional information is that in the handicap auction, she essentially charges the buyers a type-dependent up-front fee equal to the "value" of the information they are about to receive. This intuition may be appealing, but it overly simplifies the workings of the mechanism. The value of the additional information to the participants is not well-defined because it depends on the rules of the selling mechanism. This

[12] In contrast, in the classical setup with deterministic valuations, the optimal auction of Myerson (1981) and Riley and Samuelson (1981) is efficient conditional on sale, provided that the buyers are ex ante symmetric.

[13] It is easy to see that if $v_i - (1 - F(v_i))/f(v_i) < v_j - (1 - F(v_j))/f(v_j)$, but, by adding the shocks to both sides the inequality is reversed, then $v_i + s_i > v_j + s_j$. Therefore, an allocation based on the shock-adjusted virtual valuations "pointwise" improves efficiency. (This may not be true if the F_i's are not identical.)

[14] If the buyers' ex ante type-distributions are not identical then, as the seller gets to observe the signals, the efficiency of the optimal mechanism may only improve in ex ante expectation. Still, there will be some efficiency gain, which will be fully extracted by the seller even if she cannot observe the additional signals.

value could be different if the seller chose a mechanism different from the handicap auction.

4.3 Determining the Optimal Handicap Auction: A Numerical Example

It may be useful to compute a numerical example not only for illustrative purposes, but also, to see how a seller may be able to compute the parameters of the optimal handicap auction (the price premium–fee schedule) in a practical application.

We will consider the following setup. Types and shocks are additive; the types are distributed independently and uniformly on $[0, 1]$, the shocks are distributed independently according to a standard logistic distribution.[15]

First, assume that there is a single buyer, that is, $n = 1$. Then, the handicap auction can be thought of as a menu of buy options, represented by $C_1(p_1)$, where p_1 is the strike price and $C_1(p_1)$ is the fee of the option. In the first round, the buyer chooses a price p_1 and pays $C_1(p_1)$; in the second round (after having observed s_1), he has the option to buy the good at price p_1. Again, we represent this menu as a pair of functions, $c_1(v_1)$ and $p_1(v_1)$, $v_1 \in [0, 1]$.

In the uniform-logistic example with $n = 1$, the expected revenue maximizing strike price-schedule is given by (19),

$$p_1(v_1) = 1 - v_1.$$

The fee-schedule in (20) becomes,

$$c_1(v_1) = \frac{1}{2}\ln(1+e) - 1 + v_1 + \frac{1}{2}\ln\left(1 + e^{1-2v_1}\right).$$

We can also express the cost of the option as a function of the strike price,

$$c_1 = C_1(p_1) = \frac{1}{2}\ln\left[(1+e)\left(1+e^{2p_1-1}\right)\right] - p_1.$$

This (downward-sloping) schedule is depicted as the top curve in Figure 1.

[15] The cdf of the standard logistic distribution is $G_i(s_i) = e^{s_i}/(1+e^{s_i})$, $s_i \in (-\infty, +\infty)$.

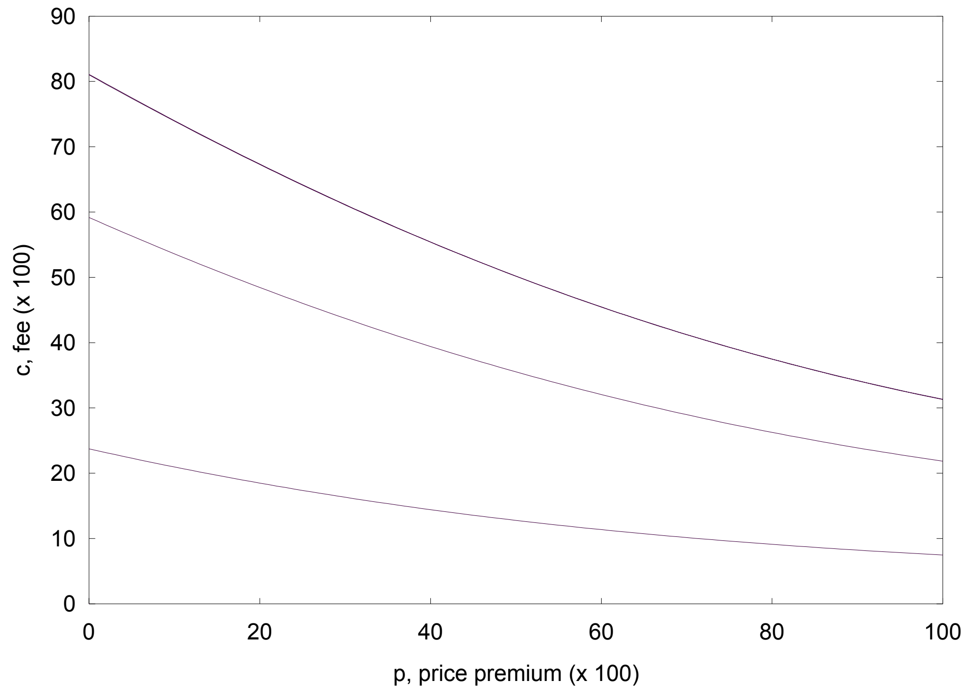

Figure 1: Fee schedules in the revenue maximizing handicap auction (uniform-logistic setup; schedules from top to bottom for n=1, 2, and 5)

If the buyer has a higher estimate then he will choose to buy an option with a lower strike price at a higher cost. For example, if the buyer has the lowest estimate, $v_1 = 0$, then he buys the option of getting the good at $p_1 = 1$, which costs $c_1 = \ln[(1 + e)/e] \approx 0.3133$ upfront, and yields zero net surplus. In contrast, the highest type, $v_1 = 1$, buys a call option with zero strike price at a cost of about 0.8133.

Now we turn to the case of many buyers, $n > 1$, in the uniform-logistic example. We will compute the optimal handicap auction represented by $\{c_i, p_i\}_{i=1}^n$. As in the case of $n = 1$, in the revenue-maximizing mechanism, $p_i(v_i) \equiv 1 - v_i$.

Instead of analytically deriving $c_i(v_i)$ for different numbers of buyers, we carry out a more practical Monte Carlo simulation. We take 100,000 random draws from the joint distribution of (s, v_{-i}), and compute $w_j = v_j + s_j - p_j(v_j)$ for all j. Then we determine $c_i(0)$ from (20), where the expectation is estimated by the sample mean. We compute $c_i(v_i)$ recursively, $c_i(v_i+\text{step}) = c_i(v_i)+\text{step}*c_i'(v_i)$, where step$= 1/100$.

The derivative is obtained from (20) as $c_i'(v_i) = \mathbf{E}_{s_i,s_{-i},v_{-i}} \left[\mathbf{1}_{\{v_i+s_i-p_i(v_i)\geq \max_{j\neq i} w_j\}}\right]$. From $c_i(v_i)$ and $p_i(v_i)$ we compute $C_i(p_i) \equiv c_i(p_i^{-1}(p_i))$.

The results of a (typical) simulation are shown in Figure 1. The top curve shows $C_i(p_i)$ for the case of $n = 1$. There are actually two (almost identical) curves superimposed on each other: one graphs the formula that we derived before, the other is the result of the Monte Carlo experiment. The curve in the middle is $C_i(p_i)$ for $n = 2$, and the one in the bottom is $C_i(p_i)$ for $n = 5$. As n increases, $C_i(p_i)$ shifts down and flattens out.

5 Conclusions

In this paper we have analyzed an auction model where the seller controls the release of private signals to the buyers that, together with the buyers' original private information, affects their valuation for the good to be sold. We have derived the expected revenue maximizing mechanism.

In the optimal mechanism, the seller discloses all the information that she can and yet obtains the same expected revenue as if she could observe the additional signals (which she can release, but cannot directly observe). The buyers do not enjoy any additional information rents from the signals whose disclosure is controlled by the seller. This result is true under a quite general signal structure: a buyer's type and signal may affect his ex-post valuation in a non-linear, non-additive way. The only restrictions imposed on the valuation functions essentially amount to requiring decreasing marginal returns.[16]

The optimal mechanism can be implemented via a "handicap auction" in important applications of the general model. Two such applications that we discussed were (i) a model where the buyer's type is interpreted as a signal solely on the expected value of his valuation, and (ii) a model where buyers obtain signals by "sampling" the distribution of their normally distributed ex-post valuations.

In the first phase of the handicap auction the seller publishes a price premium–fee schedule for each buyer; each buyer chooses a price premium and pays the

[16] Our general model can also accomodate correlation between a buyer's type and shock; we showed a simple way to transform such a setup into the independent but non-linear framework used in the paper.

corresponding fee. Then the seller allows the buyers to learn their valuations with the highest precision. In the second phase, the buyers bid for the good in a second-price sealed-bid auction with a zero reservation price, knowing that the winner will pay his premium over the price. For a single buyer, the handicap auction simplifies to a menu of buy-options.

The overall conclusion of our investigation is that under quite general conditions, the seller who controls the "flow of information" in an auction appropriates the rents of that information.

6 Appendix

Here we collect proofs omitted in the main text.

Proof of Lemma 1. Rewrite (2) as

$$\pi_i^*(v_i, \hat{v}_i) = \pi_i^*(\hat{v}_i, \hat{v}_i) + \int_{s_i} \left[u_i(v_i, s_i) - u_i(\hat{v}_i, s_i)\right] X_i(\hat{v}_i, s_i) dG_i(s_i).$$

By switching the roles of v_i and $\hat{v}_i$ we get

$$\pi_i^*(\hat{v}_i, v_i) = \pi_i^*(v_i, v_i) + \int_{s_i} \left[u_i(\hat{v}_i, s_i) - u_i(v_i, s_i)\right] X_i(v_i, s_i) dG_i(s_i).$$

Incentive compatibility of the mechanism is equivalent to, for all $v_i > \hat{v}_i$,

$$\pi_i^*(v_i, \hat{v}_i) \leq \pi_i^*(v_i, v_i) \text{ and } \pi_i^*(\hat{v}_i, v_i) \leq \pi_i^*(\hat{v}_i, \hat{v}_i).$$

Using the above expressions for the profit this becomes, for all $v_i > \hat{v}_i$,

$$\begin{aligned} &\int_{s_i} \left[u_i(v_i, s_i) - u_i(\hat{v}_i, s_i)\right] X_i(\hat{v}_i, s_i) dG_i(s_i) \\ &\qquad \leq \pi_i^*(v_i, v_i) - \pi_i^*(\hat{v}_i, \hat{v}_i) \\ &\qquad \leq \int_{s_i} \left[u_i(v_i, s_i) - u_i(\hat{v}_i, s_i)\right] X_i(v_i, s_i) dG_i(s_i). \quad (21) \end{aligned}$$

Suppose $v_i > \hat{v}_i$. Then, by $u_{i1} > 0$ and $X_i \geq 0$, the first line of the above inequality-system is non-negative, hence $\Pi_i^*(v_i) \equiv \pi_i^*(v_i, v_i)$ is weakly increasing.

Cross-divide (21) by $(v_i - \hat{v}_i)$ to get

$$\int_{s_i} \frac{u_i(v_i, s_i) - u_i(\hat{v}_i, s_i)}{v_i - \hat{v}_i} X_i(\hat{v}_i, s_i) dG_i(s_i) \leq \frac{\Pi_i^*(v_i) - \Pi_i^*(\hat{v}_i)}{v_i - \hat{v}_i} \leq \int_{s_i} \frac{u_i(v_i, s_i) - u_i(\hat{v}_i, s_i)}{v_i - \hat{v}_i} X_i(v_i, s_i) dG_i(s_i). \quad (22)$$

Consider the limit of the integral on the first line as v_i approaches $\hat{v}_i$. By concavity of u_i in v_i and $X_i \leq 1$, the integrand (pointwise, for all s_i) is weakly less than $u_{i1}(0, s_i)$, which in turn has a finite expectation with respect to s_i. Hence, by the Lebesgue Convergence Theorem (cf. Royden (1967), p.76), the order of taking the limit and the expectation can be reversed, and

$$\lim_{v_i \to \hat{v}_i} \int_{s_i} \frac{u_i(v_i, s_i) - u_i(\hat{v}_i, s_i)}{v_i - \hat{v}_i} X_i(\hat{v}_i, s_i) dG_i(s_i) = \int_{s_i} u_{i1}(\hat{v}_i, s_i) X_i(\hat{v}_i, s_i) dG_i(s_i).$$

That is, as v_i converges to $\hat{v}_i$, the first inequality of (22) implies

$$\int_{s_i} u_{i1}(\hat{v}_i, s_i) X_i(\hat{v}_i, s_i) dG_i(s_i) \leq \Pi_i^{*\prime}{}_+(\hat{v}_i)$$

whenever the right-hand derivative of Π_i^* exists at $\hat{v}_i$. By analogous reasoning, as $\hat{v}_i$ approaches v_i, the second inequality of (22) implies that

$$\int_{s_i} u_{i1}(v_i, s_i) X_i(v_i, s_i) dG_i(s_i) \geq \Pi_i^{*\prime}{}_-(v_i)$$

whenever the left-side derivative of Π_i^* exists at v_i. Since Π_i^* is weakly increasing, its derivative exists almost everywhere, therefore

$$\Pi_i^{*\prime}(v_i) = \int_{s_i} u_{i1}(v_i, s_i) X_i(v_i, s_i) dG_i(s_i) \quad (23)$$

almost everywhere.

From the second inequality of (22), concavity of u_i in v_i, and $X_i \leq 1$, it follows that

$$\frac{\Pi_i^*(v_i) - \Pi_i^*(\hat{v}_i)}{v_i - \hat{v}_i} \leq \int_{s_i} u_{i1}(0, s_i) dG_i(s_i) < \infty,$$

hence Π_i^* is Lipschitz continuous. This implies that Π_i^* is also absolutely continuous, therefore it can be recovered from its derivative (Royden (1967), p.91). We conclude that

$$\Pi_i^*(v_i) = \Pi_i^*(0) + \int_0^{v_i} \int_{s_i} u_{i1}(y, s_i) X_i(y, s_i) dG_i(s_i) dy. \tag{24}$$

Given the profit function of buyer type v_i in (24), we can write the seller's expected revenue (after the application of Fubini's Theorem) as

$$\int_s \int_v \sum_i x_i(v_i, v_{-i}, s_i, s_{-i}) \left[u_i(v_i, s_i) - \frac{1 - F_i(v_i)}{f_i(v_i)} u_{i1}(v_i, s_i) \right] dF(v) dG(s). \tag{25}$$

Our candidate optimal allocation rule will be $x_i^*(v_i, v_{-i}, s_i, s_{-i}) = 1$ for the buyer with the highest non-negative "shock-adjusted virtual valuation,"

$$u_i(v_i, s_i) - \frac{1 - F_i(v_i)}{f_i(v_i)} u_{i1}(v_i, s_i).$$

Clearly, this allocation rule pointwise maximizes the integrand in (25). The only remaining question is whether this rule is incentive compatible. Note that the candidate optimal allocation rule, $x_i^*(v_i, v_{-i}, s_i, s_{-i})$, hence $X_i^*(v_i, s_i)$, is weakly increasing in v_i because, by assumption, u_i is increasing in v_i, $(1 - F_i)/f_i$ is weakly decreasing in v_i (monotone hazard rate), and u_{i1} is weakly decreasing in v_i (u_i is concave in v_i)

We will now show that if (24) holds and $X_i(v_i, s_i)$ is weakly increasing in v_i then a mechanism inducing this allocation rule and Π_i^* profit for buyer i ($i = 1, ..., n$) is incentive compatible. Suppose $v_i > \hat{v}_i$ and use (24) to rewrite

$$\Pi_i^*(v_i) - \Pi_i^*(\hat{v}_i) = \int_{s_i} \int_{\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i(y, s_i) dy \, dG_i(s_i).$$

Decrease the integrand by replacing $X_i(y, s_i)$ with $X_i(\hat{v}_i, s_i)$,

$$\begin{aligned} \Pi_i^*(v_i) - \Pi_i^*(\hat{v}_i) &\geq \int_{s_i} \int_{\hat{v}_i}^{v_i} u_{i1}(y, s_i) dy \, X_i(\hat{v}_i, s_i) dG_i(s_i) \\ &= \int_{s_i} \left[u_i(y, s_i) - u_i(\hat{v}_i, s_i) \right] X_i(\hat{v}_i, s_i) dG_i(s_i). \end{aligned}$$

But this is exactly the incentive compatibility condition for buyer type v_i not to imitate $\hat{v}_i < v_i$, (21). A similar argument applies when $\hat{v}_i > v_i$. ■

Proof of Lemma 2. Rewrite (8) as

$$\tilde{\pi}_i(s_i, \hat{s}_i; v_i) = \tilde{\pi}_i(\hat{s}_i, \hat{s}_i; v_i) + [u_i(v_i, s_i) - u_i(v_i, \hat{s}_i)]\, X_i(v_i, \hat{s}_i) \tag{26}$$

By reversing the roles of s_i and $\hat{s}_i$, we get

$$\tilde{\pi}_i(\hat{s}_i, s_i; v_i) = \tilde{\pi}_i(s_i, s_i; v_i) - [u_i(v_i, s_i) - u_i(v_i, \hat{s}_i)]\, X_i(v_i, s_i).$$

Incentive compatibility requires, for all $s_i > \hat{s}_i$,

$$\tilde{\pi}(s_i, \hat{s}_i; v_i) \leq \tilde{\pi}_i(s_i, s_i; v_i) \text{ and } \tilde{\pi}(\hat{s}_i, s_i; v_i) \leq \tilde{\pi}_i(\hat{s}_i, \hat{s}_i; v_i).$$

Using the above expressions for $\tilde{\pi}_i(s_i, \hat{s}_i; v_i)$ and $\tilde{\pi}_i(\hat{s}_i, s_i; v_i)$ incentive compatibility becomes, for all $s_i > \hat{s}_i$,

$$\begin{aligned} [u_i(v_i, s_i) - u_i(v_i, \hat{s}_i)]\, X_i(v_i, \hat{s}_i) &\leq \tilde{\pi}_i(s_i, s_i; v_i) - \tilde{\pi}_i(\hat{s}_i, \hat{s}_i; v_i) \\ &\leq [u_i(v_i, s_i) - u_i(v_i, \hat{s}_i)]\, X_i(v_i, s_i). \end{aligned} \tag{27}$$

Cross-divide (27) by $(s_i - \hat{s}_i)$. Take $(s_i - \hat{s}_i)$ to zero: the first and third lines of (27) both converge to $u_{i2}(v_i, s_i)X_i(v_i, s_i)$. Therefore $\tilde{\pi}_i(s_i, s_i; v_i)$ is differentiable in s_i everywhere and

$$\frac{d}{ds_i}\tilde{\pi}_i(s_i, s_i; v_i) = u_{i2}(v_i, s_i)X_i(v_i, s_i).$$

This derivative is continuous in s_i (because both u_{i2} and X_i are), hence we can integrate it to get (9).

To see the second part of the claim, suppose buyer i reports $\hat{s}_i < s_i$ after seeing s_i. His deviation gain is

$$\begin{aligned} &\tilde{\pi}_i(s_i, \hat{s}_i; v_i) - \tilde{\pi}_i(s_i, s_i; v_i) \\ &\qquad = [\tilde{\pi}_i(s_i, \hat{s}_i; v_i) - \tilde{\pi}_i(\hat{s}_i, \hat{s}_i; v_i)] - [\tilde{\pi}_i(s_i, s_i; v_i) - \tilde{\pi}_i(\hat{s}_i, \hat{s}_i; v_i)]\,. \end{aligned}$$

Recall from (26) that the first bracketed difference is

$$[u_i(v_i, s_i) - u_i(v_i, \hat{s}_i)]\, X_i(v_i, \hat{s}_i),$$

and by (9) the second one can be rewritten as

$$\begin{aligned}\int_{\hat{s}_i}^{s_i} u_{i2}(v_i, z) X_i(v_i, z) dz &\geq \int_{\hat{s}_i}^{s_i} u_{i2}(v_i, z) X_i(v_i, \hat{s}_i) dz \\ &= \left[u_i(v_i, s_i) - u_i(v_i, \hat{s}_i)\right] X_i(v_i, \hat{s}_i).\end{aligned}$$

(The inequality follows from X_i being weakly increasing in s_i.) Therefore, the gain from deviation is non-positive. A similar argument can be used to show that a deviation to $\hat{s}_i > s_i$ is not profitable either. ■

Proof of Lemma 3. Had buyer i indeed have type $\hat{v}_i$ (as reported) and observed $\hat{s}_i$, incentive compatibility in the second round would require

$$u_i(\hat{v}_i, \hat{s}_i) X_i(\hat{v}_i, \hat{s}_i) - T_i(\hat{v}_i, \hat{s}_i) \geq u_i(\hat{v}_i, \hat{s}_i) X_i(\hat{v}_i, s_i') - T_i(\hat{v}_i, s_i'),$$

for all s_i'. By (10), that is $u_i(v_i, s_i) = u_i(\hat{v}_i, \hat{s}_i)$, this is equivalent to

$$u_i(v_i, s_i) X_i(\hat{v}_i, \hat{s}_i) - T_i(\hat{v}_i, \hat{s}_i) \geq u_i(v_i, s_i) X_i(\hat{v}_i, s_i') - T_i(\hat{v}_i, s_i'),$$

which means that type v_i who reported $\hat{v}_i$ in the first round and then observed s_i is indeed better off reporting $\hat{s}_i$ in the second round rather than any other s_i'. ■

Proof of Proposition 1. Set $X_i = X_i^*$ and suppose that all buyers except i report their types truthfully. Consider buyer i with type v_i contemplating to misreport to $\hat{v}_i < v_i$. Note that his deviation payoff is

$$\pi_i(v_i, \hat{v}_i) - \pi_i(v_i, v_i) = \left[\pi_i(v_i, \hat{v}_i) - \pi_i(\hat{v}_i, \hat{v}_i)\right] - \left[\pi_i(v_i, v_i) - \pi_i(\hat{v}_i, \hat{v}_i)\right].$$

By (14), and (18), the difference of the two bracketed expression can be written as

$$\int_{s_i} \int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i^*(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) dy dG_i(s_i) - \int_{s_i} \int_{y=\hat{v}_i}^{v_i} u_{i1}(y, s_i) X_i^*(y, s_i) dy dG_i(s_i) \tag{28}$$

But since for all $y \in [\hat{v}_i, v_i]$, by property (iii) of X_i^* (Corollary 1),

$$X_i^*(\hat{v}_i, \sigma_i(y, \hat{v}_i, s_i)) \leq X_i^*(y, s_i),$$

the difference in (28), and hence $\pi_i(v_i, \hat{v}_i) - \pi_i(v_i, v_i)$, is non-positive. A similar argument can be used to rule out deviation to $\hat{v}_i > v_i$. ■

Proof of Proposition 2. If, for all $j = 1, ..., n$ and $v_j \in [\underline{v}, \bar{v}]$, type v_j of buyer j purchases a premium $p_j(v_j) = (1 - F_j(v_j))/f_j(v_j)u_{i1}$ then buyer i will win in the second round if and only if, for all j,

$$u_i(v_i, s_i) - \frac{1 - F_i(v_i)}{f_i(v_i)} u_{i1} \geq \max \left\{ u_j(v_j, s_j) - \frac{1 - F_j(v_j)}{f_j(v_j)} u_{i1}, 0 \right\}.$$

This is so because in the second round, every buyer j bids $u_j(v_j, s_j) - p_j(v_j)$. Hence the allocation rule is indeed the same as in the optimal mechanism, provided that all buyers behave "truthfully," i.e., every buyer j with type v_j chooses $p_j(v_j)$ for a fee $c_j(v_j)$ defined in (20).

Finally, it is easy to check that the handicap auction defined by (19) and (20) is incentive compatible. We have seen that it defines the same allocation rule as the optimal mechanism does; moreover, the buyers' profit functions are the same as well, because by (20) and (24),

$$\Pi_i^*(v_i) = \int_s \int_{v_{-i}} (W_i(v_i, s_i) - W_{-i}^0(v_{-i}, s_{-i})) \mathbf{1}_{\{W_i(v_i, s_i) \geq W_{-i}^0(v_{-i}, s_{-i})\}} \, dF_{-i}(v_{-i}) dG(s) - c_i(v_i).$$

■

References

[1] Baron, D, and D. Besanko (1984), "Regulation and Information in a Continuing Relationship," *Information Economics and Policy,* 1:267-302.

[2] Bergemann, D., and M. Pesendorfer (2002), "Information Structures in Optimal Auctions," Cowles Foundation Discussion Paper No. 1323

[3] Compte, O., and P. Jehiel (2001), "Auctions and Information Acquisition: Sealed-bid or Dynamic Formats?," mimeo, C.E.R.A.S.

[4] Courty, P., and Li (2000), "Sequential Screening," *Review of Economic Studies*, 67:697-717.

[5] Crémer, J., and R. P. McLean (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 53:345-361.

[6] Dasgupta, P., and E. Maskin (2000), "Efficient Auctions," *The Quarterly Journal of Economics*, 115(2):341-388.

[7] Gershkov, A., (2002), "Optimal Auctions and Endogenous Information Structure," mimeo.

[8] Harris, M., and A. Raviv (1981), "Allocations Mechanisms and the Design of Auctions," *Econometrica* 49:1477-1499.

[9] Maskin, E., and J. Riley (1984), "Monopoly with Incomplete Information," *RAND Journal of Economics*, 15:171-96.

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

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

[12] Persico, N. (2000), "Information Acquisition in Auctions," *Econometrica*, 68:135-148.

[13] Riley, J., and W. Samuelson (1981), "Optimal Auctions," *American Economic Review*, 71:381-92.

[14] Riordan, M., and D. Sappington (1987), "Awarding Monopoly Franchises," *American Economic Review*, 77:375-387.

[15] Royden, H. L. (1967), "*Real Analysis*," The McMillan Company, New York.