Econometrica, Vol. 70, No. 6 (November, 2002), 2197–2224

OPTIMAL AUCTION WITH RESALE

BY CHARLES ZHOUCHENG ZHENG[1]

This paper investigates the design of seller-optimal auctions when winning bidders can attempt to resell the good. In that case, the optimal allocation characterized by Myerson (1981) cannot be achieved without resale. I find a sufficient and necessary condition for sincere bidding given the possibility of resale. In two-bidder cases, I prove that the Myerson allocation can be achieved under standard conditions supplemented with two assumptions. With three or more bidders, achieving the Myerson allocation is more difficult. I prove that it can be implemented in special cases. In those cases, the Myerson allocation is generated through a sequence of resale auctions, each optimally chosen by a reseller.

KEYWORDS: Auction, optimal auction, resale, mechanism design.

1. INTRODUCTION

MUCH OF THE AUCTION DESIGN LITERATURE makes the unrealistic assumption that winning bidders cannot attempt to resell the good to losing bidders. This assumption is not innocuous. When resale cannot be banned, the optimal allocation characterized by Myerson (1981) cannot be achieved without resale: since the allocation is sometimes *biased* in the sense that a winning bidder may value the good less than his rivals, resale may occur and upset the allocation (Proposition 1). In this paper, I investigate the extent to which sellers can achieve the optimal Myerson allocation despite possibility of resale. I find the incentive constraint that arises from resale possibility (Proposition 2). In two-bidder cases, I find that the Myerson allocation can be achieved under standard conditions supplemented with two assumptions (Proposition 3). With three or more bidders, there are additional complications that make achieving the Myerson allocation more difficult. I note that it can be implemented in very special circumstances (Proposition 4).

Let us examine as *Example 1* an environment initially considered by Ausubel and Cramton (1999). Two bidders, Strong and Weak, pursue an indivisible good that has zero value to the owner; its dollar value to Strong is uniformly distributed on [0, 10], and that to Weak is commonly known as \$2. Suppose that whoever buys the good can commit to a resale price. If the initial owner is restricted to

[1] This paper has incorporated a previously separate manuscript "Middlemen in Auctions." I thank Paul Milgrom for suggestions that inspired substantial revisions of the paper, an editor and the referees for comments, and many individuals for discussions, especially Colin Campbell, Peter Cramton, Tom Gresik, Phil Haile, Jakub Kaluzny, Roger Myerson, Mike Peters, Mark Satterthwaite, Jeroen Swinkels, Jean Tirole, Dan Vincent, Mark Walker, Mike Whinston, Asher Wolinsky, and the participants of the Joint CSIO/IDEI Industrial Organization Workshop in 2001 at Toulouse. The final stage of the project was supported by NSF Grant SES-0214471.

standard auctions with a uniform reserve price for both bidders, then the best he can get is \$2.5 (by offering the good for sale at price \$5). If implementable, biased allocations could be better. For example, the Myerson allocation, selling the good to Strong at price \$6 if Strong's value is above \$6 and selling it to Weak at price \$2 if otherwise, would yield \$3.6. If resale cannot be banned, this allocation obviously cannot be implemented during the initial sale. But the initial owner can sell the good exclusively to Weak at the price \$3.6. As Strong is excluded in the initial stage, Weak continues to believe that Strong's value is uniformly distributed on [0, 10] and hence asks Strong for the optimal resale price \$6, which is accepted by Strong with value above \$6; thus, Weak's expected payoff at the resale stage is \$3.6, his breakeven point. Hence this mechanism implements the biased allocation via resale. Moreover, the seller manages to obtain an expected profit that is the highest even when he could ban resale costlessly, since the Myerson allocation is the best for him in the no-resale environment and he cannot do better when he cannot ban resale.

The general two-bidder case is slightly more complicated. I examine a model in which a seller commits to a mechanism with the expectation that a winning bidder will also commit to a mechanism that offers resale to the losing bidder. Suppose that the Myerson allocation favors bidder 1. If we use the same bias as the Myerson allocation, then bidder 2 becomes the final owner too often. That is because bidder 2 may buy from bidder 1 in resale. On the other hand, if we use the extreme bias of always selling to bidder 1, then bidder 2 becomes the final owner too seldom. That is because bidder 1 with greater monopolistic power is less willing to resell the good.[2] I find that with an intermediate degree of bias we can recover the optimum. This optimal bias is carried out by inflating the favored bidder's bid t_1 into $\beta(t_1)$, so that his rival has to top the *inflated bid* $\beta(t_1)$ in order to win in the initial auction.[3] The inflated bid is calculated to ensure that the resale price optimally chosen by the winning bidder is exactly the cutoff at which the Myerson allocation is indifferent between the bidders. To obtain this result (Proposition 3), I make two assumptions in addition to monotone hazard rate. One is that the bidders' virtual utility functions can be uniformly ranked, so that a bidder is always favored against the other. The other is that the

[2] To see that formally (thanks to Paul Milgrom for suggesting the gist of the argument), let us consider transferring an incremental probability of being the final owner from bidder 1 to 2. In this transfer, the marginal revenue received from bidder 2 is the same from the viewpoints of the initial owner and the reseller bidder 1, because the belief about the excluded bidder 2 is unchanged before and after bidder 1 buys the good. The marginal costs, however, are different. For the reseller, it is equal to his own value t_1. For the initial owner, it is less than t_1, since the initial owner does not know t_1 and hence has to give up rents when that incremental probability is given to bidder 1 instead of 2. Thus, an exclusively delegated middleman is less likely to resell the good than the initial owner wishes him to do. As what the owner wishes is implementable (shown later), one can prove by the revenue equivalence theorem that exclusive delegation is suboptimal.

[3] A way to inflate bids is offering bid credits to favored bidders. In the U.S. "regional narrowband" spectrum auction in 1994, the government offered bid credits to small and women- and minority-owned firms.

function β is monotone nondecreasing, so that the incentive constraint identified by Proposition 2 is satisfied.

With three or more bidders, implementing the Myerson allocation is more difficult. When you sell to bidder 1, for example, he will want to resell to bidders 2 and 3 by resale-robust mechanisms described in the previous paragraph, based on his posterior belief. That leads to a complication: he may favor a bidder that the initial owner would like to discriminate against, because the Myerson allocation based on the winner's posterior belief may be different from the Myerson allocation based on the prior belief. Another complication is cycles: the optimal degree of bias for one bidder against another is specific to the pair; hence it is possible that bidder 1 beats bidder 2, 2 beats bidder 3, and 3 beats 1. In Section 5, I note that there are prior distributions for which these complications do not arise. Based on these distributions, I find a mechanism with which an initial owner achieves his best with n bidders (Proposition 4). On the equilibrium path, a winner in this auction faces a similar design problem and will use the same mechanism (revised by updated beliefs) for resale, and so will the winner in his resale auction. At every stage, the Myerson allocation in the corresponding updated environment is implemented via a positive probability of resale.[4] The process continues until a mechanism results in no sale, in which case the current owner, assumed to be able to commit to his reserve, becomes the final owner. Readers may think of Section 5 as having both positive and negative messages. The good news is that Myerson's characterization of optimal final allocations is resale-robust at least in a special case. The bad news is that extending the positive result beyond the special case is difficult if not impossible. To investigate the possibility of such extension, we need a full characterization of incentive feasibility with the possibility of resale. Towards that end, Proposition 2 might be useful.

This paper is closely related to the optimal auction literature; extending it to the case of resale is the present paper's main contribution. Myerson (1981) completely characterized the seller-optimal mechanisms for single-good independent-value models where a seller can ban resale costlessly. His result implies that a seller-optimal mechanism is biased when bidders are not identically distributed. That points to a fundamental conflict between seller-optimality and efficiency, a main motivation of the present paper.

The early papers on optimal auctions (Harris and Raviv (1981), Myerson (1981), Riley and Samuelson (1981), etc.) used the revelation principle to characterize optimal mechanisms in static environments. This method, however, is less useful in dynamic environments with recontracting (resale, renegotiation, etc.). Although an equilibrium of a mechanism-selection game from its outsider's viewpoint is outcome-equivalent to a Bayesian incentive feasible direct revelation mechanism (DRM), players inside the game need not be able to replicate the equilibrium by selecting this DRM. That is because the allocation of the

[4] The feature of nested post-sale auctions somewhat resembles a real episode discussed by Porter (1992).

mechanism can be altered by the exogenously available option of recontracting, making the DRM incredible.[5]

The paper is also related to the literature on auctions with resale. Please see Haile (2000) for a more complete list of that literature. Haile (2000) analyzed standard auctions followed by a resale game. In his model, new information about the good is revealed after an auction, so resale may occur even if the previous auction has achieved efficiency before the new information inflow. To focus on the tension between a seller's manipulation and the counterbalance from resale, I do not allow information to enter the economy between auction and resale. Regarding the role of middlemen in auctions, Bose and Deltas (1999) showed that, given a second-price auction with exogenous common value, it is profitable for a seller to exclude the final consumers from the initial auction in order to mitigate the winner's curse. Haile (1999) had a similar example, with the common value derived from the assumption that a winning bidder has a fixed share of the resale surplus. In a complete-information model, Milgrom (1987) showed that an auction is at least nearly optimal when a seller cannot prevent resale from himself or among bidders. Jehiel and Moldovanu (1999) analyzed several resale processes with complete information and interdependent values. Ausubel and Cramton (1999) considered an optimal multiunit auction with an efficient secondary market. Calzolari and Pavan (2001) considered an optimal auction in a two-bidder two-type model where a winning bidder need not have full bargaining power at resale.

The focus of this paper is resale among bidders. A complementary subject is auction design when the seller cannot commit to his no-sale decision. McAfee and Vincent (1997) analyzed this problem under the assumption that the seller's choice is a sequence of reserve prices subject to a fixed auction mechanism.

2. THE MODEL

I shall formulate the primitives governing the choice of mechanisms and other actions as a multistage mechanism-selection game with incomplete information.

2.1. *The Mechanism-Selection Game*

There is an indivisible good and $n+1$ players ($n = 1, 2, \dots$). Player 0 is the initial owner of the good. The game consists of at most $N+1$ stages, with N an integer much bigger than n. At stage 0, the initial owner selects a mechanism to be operated by a neutral trustworthy mediator within stage 1, which then begins. For any $k = 1, \dots, N$, if the game continues to stage k, the mediator announces and then operates the mechanism selected by the current owner at stage $k-1$. (If no selection was made, the mechanism is "no sale no matter what.") If the mechanism does not sell the good, then the current owner becomes the final

[5] To separate mechanisms viewed by an outsider from those selected by a player in a game, perhaps we should call the former *meta*-mechanisms.

owner and the mechanism-selection game ends. Otherwise, the current owner trades with the winner, who in turn becomes the new current owner. If $k = N$, the game ends; otherwise, the current owner selects a mechanism to be operated in stage $k+1$, which then begins. In the sequel, the word *stage* is reserved for the above meaning. (A mechanism within a stage can have several *rounds*.) Players other than the current owner are called (current) *bidders*.

Note that a seller is allowed to buy back the good, but he is not allowed to postpone the sale to future stages. Also note the assumption that a player can commit to his choice of mechanisms if and only if he is the current owner. This assumption facilitates comparison with the literature on optimal mechanisms, which is mostly based on a seller's commitment ability. However, a player's commitment ability does not entitle him to pick any mechanism. An exogenous constraint is "resale cannot be banned," formalized next.

Mechanisms. At each stage, the current owner is allowed to choose any *mechanism* restricted below. It is a mapping that associates a lottery to each possible profile of actions; the actions (including nonparticipation) are taken by the current bidders, and the lottery is carried out at the end of the stage. The lottery picks (i) a winner from all the players (including the current owner), (ii) a configuration of monetary transfers (carried out at the end of the stage), and (iii) an announcement (broadcasted at the end of the stage) of some actions taken during the stage. Items (i) and (ii) are standard in auction theory, and item (iii) is relevant because the mechanism-selection game is multistage.

In addition to the usual restrictions regarding nonparticipation and indivisibility, the above definition of a mechanism reflects two new restrictions pertinent to the resale context. First, a mechanism cannot result in any binding contract contingent on any event that may occur after the trade between the current owner and the winner. This is captured by the restriction that the lottery mandated by a mechanism is independent of events after the current stage. Hence a seller cannot commit to rewarding or penalizing the resale behavior of others.[6] Second, a mechanism should be *transparent:* once the mechanism for a stage is selected, its mandate at the end of the stage is independent of the current owner's messages or actions. This is captured by the restriction that the lottery is independent of the action taken by the current owner. This restriction is to avoid the unsettled issue of mechanism design by privately informed principals. Although the restriction is harmless with private values and quasilinear preferences when types are discrete (Maskin and Tirole (1990, Prop. 11)), we consider a continuum of types here. We need this restriction even if the initial owner has no private information, as future owners may have.

[6] This extreme formulation of the resale constraint is to facilitate tractability. A deeper analysis may assume that resale is an action hidden from the current seller and then derive an endogenous resale constraint from the sacrifice needed to induce confession from the participants of resale.

Preferences. If a player i's total net payment to others throughout the mechanism-selection game is p (with $-p$ denoting the net receipts), then his payoff is $t_i - p$ if he is the final owner of the good, and is $-p$ otherwise. Here the level t_i of his utility from being the final owner is a constant given to him and is called his *type*. This assumption of private values allows us to trace the impact of resale from the endogenously induced common-value component. Discounting is not considered, as it can only weaken the resale constraint.

Information Structure. At the beginning of the mechanism-selection game, a bidder's type is privately known to himself and is regarded as a random variable independently drawn from a commonly known distribution. Different bidders' distributions can be different. The initial owner's type is commonly known. A player does not know the actions of others unless the mechanism exposes them. (The current owner is not an exception, as the mechanism is not operated by himself.) After the operation of a mechanism, the identity of the winner and the public announcement constitute the new common knowledge. Hence a seller can influence the selection of future mechanisms by his disclosure policy of past information.

2.2. *The Equilibrium Concept*

For multistage mechanism-selection games with incomplete information, defining the equilibrium concept is nontrivial. We need a logically consistent treatment to an issue: What should an equilibrium say when a player has selected a mechanism that has no continuation "equilibrium" (say due to bad tie-breaking rules)? The issue is not about the existence of equilibrium, but whether the equilibrium concept is well-formed. The resolution is a modified notion of perfect Bayesian equilibrium (PBE). The modified concept is the same as PBE except for a slightly weaker condition of sequential rationality, which is still stronger than that of Bayesian Nash equilibrium (BNE). A reader willing to accept the above claim may skip to the next subsection.

In static environments, the above issue could be resolved by assuming that only equilibrium feasible mechanisms are admissible. The trouble in multistage mechanism-selection games is that usually we cannot identify equilibrium infeasible mechanisms a priori, as whether something is a continuation equilibrium depends on the belief system of the underlying equilibrium of the entire game. This Catch-22 problem is analogous to the problem in formulating Bayesian updating. As is well-known, we cannot apply Bayes' formula at surprising events and these events cannot be identified a priori. The treatment of PBE is to identify them relative to the equilibrium and impose no restriction on the belief system at such events. Analogously, we may impose no restriction on the strategy profile whenever the selected mechanism is "equilibrium infeasible" relative to the "equilibrium," provided that both notions can be defined. As these notions are intertwined, I define them by backward recursion of the mechanism-selection game. That is why the game is assumed to have a final stage N.

DEFINITION: An *equilibrium* of a continuation game at the beginning of the final stage N means a BNE of the continuation game. Hence the notion of equilibrium at this level is well-formed, since BNE is so. Pick any $k = 1, \ldots, N$. Suppose that the notion of equilibrium is already well-formed for any continuation game starting at the beginning of stage k. Thus, for any mechanism to be operated in stage k, its continuation game either has an equilibrium or not; hence we can call the mechanism *equilibrium feasible* if its continuation game has an equilibrium and call the mechanism *equilibrium infeasible* if otherwise. An *equilibrium* of a continuation game starting at the beginning of stage $k-1$ means a pair of strategy profile $\mathcal{S}$ and belief system $\mathcal{B}$ such that:

(i) $\mathcal{B}$ follows Bayes's rule relative to $\mathcal{S}$, except at any history where the decision-maker can conclude that someone else has deviated from $\mathcal{S}$ at the immediately preceding history;

(ii) $\mathcal{S}$ is sequentially rational relative to $\mathcal{B}$, except when the player who buys the good in stage $k-1$ has either

(a) selected a mechanism (to be operated in stage k) that is equilibrium infeasible, or

(b) deviated from $\mathcal{S}$ and the deviation is dominated from the standpoint when the deviation was made.

Here clause (a) gives us a logical basis for a uniform treatment when a player has selected a deviant mechanism: players abide the associated continuation equilibrium if the deviant mechanism is equilibrium feasible, and they can pick any action if otherwise. As a by-product of the definition, clause (b) allows us to not specify a current owner's optimal mechanism if in the past he made a deviant and dominated move.

2.3. *Environments, Allocations, and Myerson Allocations*

For a current owner, the *environment* consists of all the data known to him when he selects a mechanism. The environment for the initial owner is simply the exogenously given mechanism-selection game parameterized by the prior distributions of the bidders' types. For a future owner, however, the environment is endogenously determined by the posterior distributions based on observed histories and the underlying equilibrium.

Roughly speaking, allocations are type-contingent final outcomes. Selecting an equilibrium feasible mechanism, a current owner is essentially choosing an allocation. Since he is confined to transparent mechanisms (Section 2.1), we need only to consider those allocations that are independent of his type, though his choice among them is based on his type. More precisely, a *final ownership* means a lottery that picks a final owner of the good from the set of all the players. For any current owner i, an *allocation* in his environment means a function that associates a final ownership to every possible profile of the types of the players other than i. An allocation in i's environment is said to be *equilibrium feasible* if i can select an equilibrium feasible mechanism that *implements* the allocation (generates it on the equilibrium path). In contrast, an allocation in i's environment is said to be *BNE-feasible* if i can select a mechanism whose continuation game has a BNE that generates the allocation. Since our equilibrium concept is stronger than BNE, an equilibrium feasible allocation is always BNE-feasible, but not vice versa. They coincide only when all the resale stages of the mechanism-selection game are removed.

For any current owner i and any environment given to him, given any BNE-feasible allocation α in this environment, define the *surplus* for i at α as the maximum expected profit for i among all the BNEs that generate the allocation, and define the surplus for any other player j at α as the expected profit for j at that i-optimal BNE. If there is an allocation that maximizes the current owner's surplus among all BNE-feasible allocations, we call it a *Myerson allocation*. A player's surplus at a Myerson allocation is called his *Myerson surplus*.

The next lemma implies a sufficiency condition for a mechanism to be optimal for an owner: it is equilibrium feasible and gives the owner his Myerson surplus. The intuition is simply that a seller cannot do worse when he can ban resale costlessly than when he cannot.

LEMMA 2.1: *An owner's Myerson surplus, if it exists, is an upper bound of his expected profits that are equilibrium feasible. A mechanism is optimal for an owner if it implements the Myerson allocation and gives the bidders their Myerson surpluses.*

PROOF: The first sentence follows directly from the fact that an equilibrium is always a BNE. Hence the second sentence is true since the owner's surplus is equal to the entire social surplus, fixed given an equilibrium, minus the bidders' surplus. *Q.E.D.*

For each player i, let F_i denote the prior distribution of i's type at the beginning of the mechanism-selection game. The following assumption is usual in the auction literature.

ASSUMPTION 1 (Hazard Rate): *For each player i, the support T_i of F_i is convex and bounded from below and, if T_i is a nondegenerate interval, the density function f_i is positive and continuous on T_i and differentiable in its interior, and $(1-F_i(t_i))/f_i(t_i)$ is a decreasing function of t_i on T_i.*

Assumption 1 implies that the well-known *virtual utility* function V_i defined by

$$(1) \qquad V_i(t_i) := t_i - \frac{1-F_i(t_i)}{f_i(t_i)}, \qquad \forall t_i \in T_i,$$

is strictly increasing on T_i. Then one can show that all the Myerson allocations in the initial owner's environment are identical for almost all profiles of the types of the bidders, and all these allocations have the property that, for almost all profiles $(t_i)_{i=1}^n$ of bidders' types, the final owner is selected in descending order of the virtual utilities $V_i(t_i)$ across i until they drop to the owner's type. Hence *the* Myerson allocation in the sequel will refer to the allocation that satisfies this selection criterion for *every* possible profile of types.

An allocation is said to be *biased* towards a bidder if there is a positive probability with which he becomes the final owner of the good when some other bidder's value is higher than his. Clearly the Myerson allocation is biased when the virtual utility functions are different across i. This important feature creates a tension between a seller's incentive to implement a biased allocation and bidders' incentive to undo the bias via resale. The next assumption is to dramatize this tension.

ASSUMPTION 2 (Uniform Bias): *The bidders can be ranked by* $1, \ldots, n$ *so that, if bidder* i *is ranked before bidder* j $(i < j)$, *then* $T_i \subseteq T_j$ *and* $V_i(x) \geq V_j(x)$ *for all* $x \in T_i$.

Thus, the Myerson allocation is biased towards bidder 1 against bidders $2, \ldots, n$, biased towards bidder 2 against bidders $3, \ldots, n$, and so on. We shall name the players by this ranking system in the sequel.

3. PRELIMINARY ANALYSIS

3.1. *Myerson Allocations Need Resale*

Resale means trading among bidders after the current stage. Can we implement the Myerson allocation without resale? The answer is No. That is because the allocation is biased and the following fact applies.

PROPOSITION 1: *It is impossible to implement a biased allocation such that the probability of resale is zero.*

PROOF: Suppose that an allocation biased towards a bidder i is implemented via a mechanism without resale. At the associated equilibrium, this allocation is completed after stage one of the mechanism-selection game. Let us begin with the special case where the mechanism does not give a winner any news other than the winning status. In the event that i wins, he would infer that there is a positive probability with which the gain of trade between him and other bidders is at least some positive ϵ. Hence i has a strict incentive to deviate from the equilibrium and offer the good for resale at a price equal to his value plus a fraction of ϵ, a contradiction.

Next consider the general case where a winner may receive additional news. Conditional on every news ν_i and i's winning status, bidder i calculates the gain of trade, say $g_i(\nu_i)$, at resale. Integrating g_i over all possible news ν_i (with the probability measure derived from the equilibrium path), bidder i gets the expected value $\bar{g}_i$ of the gain from trade at resale from the standpoint of having seen the winning status and before receiving any additional news. As the allocation is biased, $\bar{g}_i > 0$. Hence $g_i > 0$ on a positive-probability set of news. From every such news, bidder i would infer that it is profitable to deviate from the equilibrium prescription of no resale. *Q.E.D.*

The above proposition implies that, although the Myerson allocation may be a best scenario for a seller, the seller cannot use it as a recipe. If he did, at equilibrium the winner-selection rule must be violated with a positive probability.

3.2. *Acyclic Path of Resale*

Would a bidder resell the good back to its previous owner? Would a bidder outbid a rival and then resell the good back to the losing rival? The next lemma says that the answers are No with appropriate payment schemes. The reason is that each player knows his own value of the good.

LEMMA 3.1: *If a current owner knows, before selecting a mechanism, that his type is at least as high as a player's, and if the two parties always have identical beliefs about others, then the current owner finds it unprofitable to include the player as a potential buyer. If winning in a current stage requires a bidder i to pay more than player j's type, and if i and j always have identical beliefs about others, then it is weakly dominated, from the standpoint of the current stage, for i to outbid j (or buy the good from j if j is the current owner) and then resell the good to j.*

PROOF: The first claim follows from the fact that there is zero gain of trade between the current owner and the player whose value is a priori known to be less than the owner's. Even if this player is willing to pay more than her value due to possible resale, the owner can resell the good himself. (Note that the two parties do not have different assessments of the revenue.) The second claim follows from the fact that resale back to j cannot recover i's payment, unless j is willing to pay more than her value due to possible resale, but again i can obtain that resale revenue without resale to j. *Q.E.D.*

Note: The traditional payment schemes to implement Myerson allocation do not satisfy the condition of the above lemma. For example, a winner w is supposed to pay $\max\{V_w^{-1}(V_i(t_i)): i \neq w\}$. Clearly it need not be true that $V_w^{-1}(V_i(t_i)) \geq t_i$. Hence the probable occurrence of resale after Myerson's optimal auction does not contradict Lemma 3.1.

3.3. *The Incentive Constraint against One-Shot Deviation*

Characterizing equilibrium feasible mechanisms is itself a challenging problem. Here I shall confine the investigation to the incentive condition that prevents one-shot deviations. Specifically, let m be a revelation mechanism to be operated in the current stage, and let σ be a strategy profile conditional on the operation of m. In order for (m, σ) to constitute a mechanism-equilibrium pair, three conditions must hold: (i) the allocation generated by (m, σ) is BNE-feasible; (ii) the strategy σ is sequentially rational in any possible continuation game following the operation of m (unless clause (a) or (b) in Section 2.2 applies); (iii) each player is truthful during the operation of m. The question is what (iii) exactly requires if (i) and (ii) are satisfied.

Let $q_i^*(\hat{t}_i, t_i \mid m, \sigma)$ denote the probability, calculated by player i, with which i becomes the final owner of the good if he reports $\hat{t}_i$ to mechanism m and follows the strategy σ afterwards, given his type t_i and provided that other players always follow σ. The next proposition identifies the condition making such one-shot deviations unprofitable.

PROPOSITION 2: *Consider any mechanism m, strategy profile σ, and bidder i. Suppose: (a) if σ is followed, the final allocation is BNE-feasible; (b) σ is sequentially rational for i in any possible continuation game after the operation of m (unless clause* (ii)(*a*) *or* (ii)(*b*) *in Section* 2.2 *applies*); *and* (*c*) *the distribution of i's type is supported by an interval. Bidder i always finds it optimal to be truthful during the operation of m (within the current stage) if and only if*

$$(2) \qquad \int_{t_i'}^{t_i} [q_i^*(\tau_i, \tau_i \mid m, \sigma) - q_i^*(t_i', \tau_i \mid m, \sigma)]\, d\tau_i \geq 0$$

for any possible types t_i and t_i' of i such that misrepresenting a true type t_i as t_i' is not exempted by clause (ii)(*a*) *or* (ii)(*b*).

PROOF: Notations: Without loss of generality, regard the current stage as stage one. Starting from the beginning of stage two, the player can deviate to a strategy, say $\sigma_i(\hat{s}_i, s_i)$, that σ would instruct him to follow from now on had his stage-one report to the mechanism been $\hat{s}_i$ (instead of say $\hat{t}_i$) and his type been s_i (instead of say t_i). Let $(\hat{t}_i, \hat{s}_i, s_i)$ denote the contingency plan of reporting $\hat{t}_i$ in stage one and using strategy $\sigma_i(\hat{s}_i, s_i)$ afterwards. Given this contingency plan and the expectation that others abide by σ, let $q_i(\hat{t}_i, \hat{s}_i, s_i \mid m, \sigma)$ denote the probability with which player i becomes the final owner of the good, and let $p_i(\hat{t}_i, \hat{s}_i, s_i \mid m, \sigma)$ denote the expected value of player i's total net payment to others throughout the entire mechanism-selection game, each calculated by i before the operation of mechanism m. Let $u_i(\hat{t}_i, \hat{s}_i, s_i \mid t_i, m, \sigma)$ denote a type-t_i player i's expected profit for the entire mechanism-selection game viewed before the operation of the mechanism m, given his contingency plan $(\hat{t}_i, \hat{s}_i, s_i)$ and the expected σ. Let $u_i^*(\hat{t}_i \mid t_i, m, \sigma) := u_i(\hat{t}_i, \hat{t}_i, t_i \mid t_i, m, \sigma)$.

Note that $q_i^*(\hat{t}_i, t_i \mid m, \sigma) = q_i(\hat{t}_i, \hat{t}_i, t_i \mid m, \sigma)$, and $u_i^*(\hat{t}_i \mid t_i, m, \sigma)$ is the player's expected utility from the one-shot deviation strategy that follows σ except reporting $\hat{t}_i$ to m. Also notice that both $q_i(\hat{t}_i, \hat{s}_i, s_i \mid m, \sigma)$ and $p_i(\hat{t}_i, \hat{s}_i, s_i \mid m, \sigma)$ are independent of the player's true type. That is because the stage-one outcome is determined only by the stage-one reports, and the player's future action is determined by his strategy $\sigma_i(\hat{s}_i, s_i)$ (contingent on any additional news, which in turn is determined by the stage-one reports). Since a player's payoff is quasilinear, one can calculate that

$$(3) \qquad u_i(\hat{t}_i, \hat{s}_i, s_i \mid t_i, m, \sigma) = t_i q_i(\hat{t}_i, \hat{s}_i, s_i \mid m, \sigma) - p_i(\hat{t}_i, \hat{s}_i, s_i \mid m, \sigma).$$

Pick any possible types t_i and t_i' of the player. Since the allocation generated by (m, σ) is BNE-feasible, setting $\hat{\tau}_i := \tau_i$ maximizes $u_i(\hat{\tau}_i, \hat{\tau}_i, \hat{\tau}_i \mid \tau_i, m, \sigma)$ among all possible values of $\hat{\tau}_i$. Hence the envelope theorem (Milgrom and Segal (2002)) implies

$$u_i^*(t_i \mid t_i, m, \sigma) - u_i^*(t_i' \mid t_i', m, \sigma) = \int_{t_i'}^{t_i} q_i^*(\tau_i, \tau_i \mid m, \sigma)\, d\tau_i.$$

Here we used equation (3) and the assumption that the player's type is supported by an interval. After reporting any $\hat{t}_i$ in stage one, following the instruction of σ afterwards is optimal for the player (hypothesis (b)), hence setting $s_i := \tau_i$ maximizes $u_i(\hat{t}_i, \hat{t}_i, s_i \mid \tau_i, m, \sigma)$ among all possible values of s_i (unless clause (ii)(a) or (ii)(b) applies). Again the envelope theorem, using the same hypotheses as in the previous equation, implies

$$u_i^*(\hat{t}_i \mid t_i, m, \sigma) - u_i^*(\hat{t}_i \mid t_i', m, \sigma) = \int_{t_i'}^{t_i} q_i^*(\hat{t}_i, \tau_i \mid m, \sigma)\, d\tau_i$$

for any $\hat{t}_i$ (unless clause (ii)(a) or (ii)(b) applies to $\hat{t}_i$). These two equations give

$$\begin{aligned} u_i^*(t_i \mid t_i, m, \sigma) - u_i^*(t_i' \mid t_i, m, \sigma) &= [u_i^*(t_i \mid t_i, m, \sigma) - u_i^*(t_i' \mid t_i', m, \sigma)] \\ &\quad + [u_i^*(t_i' \mid t_i', m, \sigma) - u_i^*(t_i' \mid t_i, m, \sigma)] \\ &= \int_{t_i'}^{t_i} [q_i^*(\tau_i, \tau_i \mid m, \sigma) - q_i^*(t_i', \tau_i \mid m, \sigma)]\, d\tau_i \end{aligned}$$

whenever clauses (ii)(a) and (ii)(b) do not apply to i's deviation of reporting t_i'. This proves the proposition, because the bidder's optimization would be vacuous if the deviation is covered by clause (ii)(a) or (ii)(b). *Q.E.D.*

The incentive constraint (2) comes from the multistage nature of the mechanism-selection game. If the game is one-shot, as in the mainstream framework, then one can show that this constraint is equivalent to the standard monotonicity condition, i.e., (2) holds for all t_i and t_i' if and only if $q_i^*(\tau_i, \tau_i \mid m, \sigma)$ is monotone nondecreasing in τ_i.[7] If the mechanism-selection game is multistage, however, one can construct functions q_i^* such that $q_i^*(\tau_i, \tau_i \mid m, \sigma)$ is monotone in τ_i and yet (2) does not always hold.

REMARK 1: Equation (2) holds for all t_i and t_i' if, for any τ_i, $q_i^*(\cdot, \tau_i \mid m, \sigma)$ is monotone nondecreasing. This monotonicity condition says that, given any type, a player submitting a higher bid to mechanism m has a higher expected probability of being the final owner.

[7] In the static model, the current stage is the final stage, so the probability $q_i^*(\hat{\tau}_i, \tau_i \mid m, \sigma)$ is invariant to τ_i. Hence the claim for "if" is trivial. To prove the converse, apply (2) to the case where the roles of t_i and t_i' are interchanged. Summing the obtained inequality with (2), we have

$$\int_{t_i'}^{t_i} [q_i^*(t_i, \tau_i \mid m, \sigma) - q_i^*(t_i', \tau_i \mid m, \sigma)]\, d\tau_i \geq 0.$$

As $q_i^*(\hat{\tau}_i, \tau_i \mid m, \sigma)$ is invariant to τ_i, the monotonicity condition follows.

REMARK 2: Any efficient allocation can be implemented with zero probability of resale. Coupled with sincere bidding, any standard auction, say Vickrey auction, will do. Clearly the allocation is BNE-feasible. As a winner's payment is at least as high as any rival's value (given sincere bidding from others), Lemma 3.1 implies that it is dominated to overbid now and resell the good later. Hence overbidding is dominated. The only case not exempted by clause (ii)(b) is underbidding. If a bidder wins after submitting a bid below his type, he learns that other bidders' values are less than his, hence the winner consumes the good. Thus, the probability with which a bidder becomes the final owner from underbidding is equal to the probability with which his rivals' types are less than his bid. This probability is obviously monotone nondecreasing in the bid. Hence (2) holds.[8]

CAVEAT: When the mechanism-selector at the current stage does not know what player i believes, the functional form of q_i^* is only in the mind of player i. Hence (2) need not be useful for the choice of mechanisms when beliefs are not commonly known.

4. THE TWO-BIDDER CASE

We already know that a seller wishes for a Myerson allocation but he cannot implement it without resale. The question is how to induce a reseller to do what the seller wishes. A reseller decides whether to resell and to whom to resell the good. With only two bidders, however, a reseller has only one potential buyer unless the initial seller buys it back. Thus, we need only to align a reseller's decision of whether to resell the good with the Myerson allocation. Since this allocation is biased towards bidder 1 by the uniform-bias assumption, bidder 1 is likely the reseller. Hence we shall focus on him.

Suppose that the initial owner selected a mechanism and bidder 1 has won after bidding $\hat{t}_1$ in it. Then he updates his belief about bidder 2 based on this history together with any additional news that he may have received. Suppose temporarily that the virtual utility function based on the corresponding posterior distribution can be defined by equation (1). Denote this function by $V_{2|\hat{t}_1}$.

We can think of $V_{2|\hat{t}_1}(t_2)$ as bidder 1's marginal revenue (MR) from an incremental probability of reselling the good to bidder 2 given the latter's type t_2. Bidder 1's marginal cost (MC) is his value t_1 of consuming the good himself, as long as the initial owner does not buy it back. At any interior solution for his optimal resale decision, MR is equal to MC, since a reseller is assumed to have commitment ability. Hence the reseller's optimal cutoff for resale is determined by $V_{2|t_1}(t_2) = t_1$ if he has been honest ($\hat{t}_1 = t_1$) in the initial mechanism. To implement the Myerson allocation, however, the initial owner wishes the cutoff of resale (making bidder 2 instead of 1 the final owner) to be determined by $V_2(t_2) = V_1(t_1)$. To align their interests, therefore, the mechanism needs to ensure that the two equations are equivalent.

[8] It is unknown, however, whether an efficient allocation *has to* be implemented with zero probability of resale.

Suppose that in the initial owner's mechanism bidder l's bid $\hat{t}_1$ is transformed to $\beta_{12}(\hat{t}_1)$ by some function $\beta_{12}\colon T_1 \to T_2$. Also suppose temporarily a *minimum information condition:* bidder 1 obtains no additional news other than his winning status. Then his posterior belief about bidder 2's type is the prior F_2 conditional on the event $t_2 \le \beta_{12}(\hat{t}_1)$. Hence bidder 2's virtual utility is given by the following formula with $i = 1$ and $j = 2$:

$$V_{j|\hat{t}_i}(t_j) = t_j - \frac{1 - F_j(t_j \mid t_j \le \beta_{ij}(\hat{t}_i))}{f_j(t_j \mid t_j \le \beta_{ij}(\hat{t}_i))}. \tag{4}$$

To align bidder l's resale cutoff $V_{2|t_1}(t_2) = t_1$ with the cutoff $V_2(t_2) = V_1(t_1)$ for which the initial owner wishes, we need $V_{2|t_1}(V_2^{-1}(V_1(t_1))) = t_1$. By equation (4), that means β_{12} needs to satisfy the following equation with $i = 1$ and $j = 2$:

$$F_j(\beta_{ij}(t_i)) = F_j(V_j^{-1}(V_i(t_i))) + f_j(V_j^{-1}(V_i(t_i)))[V_j^{-1}(V_i(t_i)) - t_i], \ \forall t_i \in T_i. \tag{5}$$

Note: If $i < j$ then $V_j^{-1}(V_i(t_i)) \in T_j$ for all $t_i \in T_i$. That can be proved from the uniform-bias assumption and the convexity of the range of V_j (since V_j is continuous and T_j is convex by Assumption 1). Consequently, from the hazard-rate assumption and the intermediate-value theorem, one can prove that $\beta_{ij}(t_i)$ is unique and exists between $V_j^{-1}(V_i(t_i))$ and the supremum of T_j. As $V_j^{-1}(V_i(t_i)) \ge t_i$ by the uniform-bias assumption, $\beta_{ij}(t_i) \ge t_i$, hence β_{ij} *inflates* i's bid.

With β_{12} well-defined, it is meaningful for the initial owner to select a winner by this *bid-inflation rule*: bidder 1 defeats bidder 2 if and only if $\beta_{12}(\hat{t}_1) \ge \hat{t}_2$, with $\hat{t}_i$ denoting i's bid. If a type-t_1 bidder 1 wins under this rule after bidding $\hat{t}_1$ and wants to make a take-it-or-leave offer to bidder 2, then he solves

$$\pi_1(\hat{t}_1 \mid t_1) := \max_{p \ge t_1} (p - t_1)[1 - F_2(p \mid t_2 \le \beta_{12}(\hat{t}_1))], \tag{6}$$

which assumes the aforementioned minimum information condition. By the hazard-rate assumption, one can show that the optimum of (6) exists and is $V_{2|\hat{t}_1}^{-1}(t_1)$. (The inverse exists because $V_{2|\hat{t}_1}$ is strictly increasing by the hazard-rate assumption.) If bidder 1 has bid his true type initially, this optimal resale price becomes $V_{2|t_1}^{-1}(t_1)$, which is equal to $V_2^{-1}(V_1(t_1))$ by the definition of $\beta_{12}(t_1)$. Hence the resale cutoff is exactly what the initial owner wishes ($V_1(t_1) = V_2(t_2)$) if bidder 1 is truthful initially. This leads to the following mechanism.

The Mechanism. The minimum bid for bidder i is $V_i^{-1}(t_0)$, with t_0 being the initial owner's type. Each bidder $i \in \{1, 2\}$ independently submits a bid $\hat{t}_i \in T_i$; the message "I do not participate" is taken as a bid below the bidder's minimum bid. If $\hat{t}_1$ is below bidder 1's minimum bid, skip him and make a take-it-or-leave offer to bidder 2 at the price equal to 2's minimum bid. If $\hat{t}_1$ is not below the minimum bid but $\beta_{12}(\hat{t}_1) < \hat{t}_2$, skip bidder 1 and make a take-it-or-leave offer to bidder 2 at the price equal to $V_2^{-1}(V_1(\hat{t}_1))$. (Note that this price is at least as high as bidder 2's minimum bid.) If $\hat{t}_1$ meets the minimum-bid requirement and $\beta_{12}(\hat{t})_1 \ge \hat{t}_2$, sell the good to bidder 1 at the price $p_1(\hat{t}_1)$ to be defined later.

In this mechanism, the minimum-bid requirement is to implement the bidder-specific reserve price in the Myerson allocation, which is $V_i^{-1}(t_0)$ for bidder i. The payment scheme resembles a first-price auction for bidder 1: he cannot tell bidder 2's type from the amount of payment. This guarantees the minimum information condition, which makes equation (4) valid.

To achieve the initial owner's Myerson surplus, the payment scheme needs to align each bidder's Myerson surplus with the bidder's equilibrium expected profit. Hence let us specify a strategy profile that will be proved to constitute the underlying equilibrium:

The Equilibrium Hypothesis. The initial owner selects the mechanism defined above (whose payment scheme is defined next). If he deviates to a different mechanism, clause (ii)(a) of our equilibrium concept applies. If he does not deviate, the bidders report their true types. If a type-t_1 bidder 1 has won in this mechanism after submitting a bid $\hat{t}_1$, he makes a take-it-or-leave offer to bidder 2 at the price equal to $V_{2|\hat{t}_1}^{-1}(t_1)$, and bidder 2 accepts the offer if and only if her type t_2 is not below this price; clause (ii)(a) applies if the winning bidder 1 deviates to a different resale mechanism. If bidder 2 buys the good from either the initial owner or bidder 1, she chooses not to resell the good; if bidder 2 buys the good and offers to resell it, clause (ii)(b) of our equilibrium concept applies because it will be shown in the proof of Proposition 3 that doing so is dominated for her.

The Payment Scheme. The payment scheme for bidder 2 has been defined. To define the payment scheme for bidder 1, let us calculate his expected profit for the entire game—viewed before the mechanism is operated—provided that bidder 2 follows the hypothetical equilibrium. By Lemma 3.1, it is dominated for bidder 1 to win the good and resell it to the initial owner, because the minimum bid for bidder 1 is higher than the owner's type. Thus, we consider only those bids $\hat{t}_1$ from bidder 1 such that he will find it optimal to offer resale only to bidder 2. As an owner is restricted to transparent mechanisms, any mechanism bidder 1 selects is a take-it-or-leave offer to bidder 2. Thus, if he bids $\hat{t}_1$ and wins in the initial mechanism, his expected revenue is $t_1 + \pi_1(\hat{t}_1 \mid t_1)$, with π_1 defined in (6). If bidder 1 does not win, then his expected revenue is zero: either the initial owner becomes the final owner, or bidder 2 wins it, who becomes the final owner by the equilibrium hypothesis. The probability $Q_1(\hat{t}_1)$ with which bidder 1 wins in the initial mechanism is $F_2(\beta_{12}(\hat{t}_1))$ if his bid matches the minimum-bid requirement, and is zero if otherwise. Suppose that in the payment scheme bidder 1's expected amount of payment is $\bar{p}_1(\hat{t}_1)$ if his bid is $\hat{t}_1$. Then his expected profit, viewed before the initial mechanism is operated, is equal to

$$u_i^*(\hat{t}_1 \mid t_1) = Q_1(\hat{t}_1)[t_1 + \pi_1(\hat{t}_1 \mid t_1)] - \bar{p}_1(\hat{t}_1). \tag{7}$$

At the hypothetical equilibrium, bidder 1 bids truthfully and gets expected profit $u_1^*(t_1 \mid t_1)$. Let $U_1^*(t_1)$ denote bidder 1's Myerson surplus if his type is t_1. To give the initial owner his Myerson surplus, the payment scheme needs to satisfy $u_1^*(t_1 \mid t_1) = U_1^*(t_1)$ for all t_1. A *payment scheme* satisfying this condition is: bidder 1 pays zero if he does not win and pays $p_1(\hat{t}_1)$ if he wins after bidding $\hat{t}_1$, where $p_1(\hat{t}_1)$ is defined by

$$(8) \qquad p_1(\hat{t}_1) := \hat{t}_1 + \pi_1(\hat{t}_1 \mid \hat{t}_1) - U_1^*(\hat{t}_1)/F_2(\beta_{12}(\hat{t}_1)).$$

Notice that the denominator $F_2(\beta_{12}(\hat{t}_1))$ is positive when bidder 1 wins.[9]

We have therefore designed a mechanism for the initial owner. Are bidders honest in it? Proposition 2 has identified an incentive constraint (2) coming from the possibility of resale. As Remark 1 there points out, sufficient for (2) is a monotonicity condition: given any type, a player submitting a higher bid to the mechanism has a higher expected probability of being the final owner.

Let us calculate this probability for bidder 1. At the hypothetical equilibrium, he becomes the final owner if and only if he wins in the mechanism (bidder 2's type t_2 does not exceed $\beta_{12}(\hat{t}_1)$) and resale does not occur (t_2 is less than bidder 1's optimal resale price $V_{2|\hat{t}_1}^{-1}(t_1)$). Being the optimum for problem (6), $V_{2|\hat{t}_1}^{-1}(t_1)$ cannot exceed the posterior upper bound $\beta_{12}(\hat{t}_1)$ of bidder 2's type; hence $t_2 < V_{2|\hat{t}_1}^{-1}(t_1)$ implies $t_2 < \beta_{12}(\hat{t}_1)$, so bidder 1 becomes the final owner if and only if $t_2 < V_{2|\hat{t}_1}^{-1}(t_1)$. Thus, the above monotonicity condition is translated to the condition that $F_2(V_{2|\hat{t}_1}^{-1}(t_1))$ is increasing in the bid $\hat{t}_1$. The next lemma tells us exactly when this condition is satisfied.

LEMMA 4.1: *Assume that F_2 is strictly increasing. Given the proposed mechanism and equilibrium, and given any type, bidder 1's probability of being the final owner is monotone nondecreasing in his stage-one bid if and only if β_{12} is monotone nondecreasing.*

PROOF: Since F_2 is strictly increasing by assumption, the claim is equivalent to "$V_{2|\hat{t}_1}^{-1}(t_1)$ is monotone nondecreasing in $\hat{t}_1$ if and only if β_{12} is monotone nondecreasing." Pick any bids z and z' such that $z > z'$. Let $x := V_{2|z}^{-1}(t_1)$ and $x' := V_{2|z'}^{-1}(t_1)$. Suppose that β_{12} is monotone nondecreasing. With F_2 strictly increasing, this implies $F_2(\beta_{12}(z)) \geq F_2(\beta_{12}(z'))$, so by equation (4) we have $V_{2|z'}(x) \geq V_{2|z}(x)$. Hence, since $V_{2|z}(x) = t_1 = V_{2|z'}(x')$, $V_{2|z'}(x) \geq V_{2|z'}(x')$. With

[9] A bidder's total net payment throughout the mechanism-selection game is different from what he pays the initial owner. Bidder 1's expected payment $\bar{p}_1$ to the owner is determined by equation (8). In contrast, the expected value of bidder 1's total net payment throughout the game at the hypothetical equilibrium is $\bar{p}_1^*(\hat{t}_1) = \hat{t}_1\tilde{q}_1(\hat{t}_1) - U_1^*(\hat{t}_1)$ if he bids $\hat{t}_1$, where $\tilde{q}_1(\hat{t}_1)$ is bidder 1's probability of being the final owner in the Myerson allocation. Note

$$\bar{p}_1(\hat{t}_1) - \bar{p}_1^*(\hat{t}_1) = \hat{t}_1[F_2(\beta_{12}(\hat{t}_1)) - \tilde{q}_1(\hat{t}_1)] + F_2(\beta_{12}(\hat{t}_1))\pi_1(\hat{t}_1, \hat{t}_1),$$

which is positive since our mechanism is more biased than the Myerson allocation and $\pi_1 \geq 0$ by definition. This reflects the arrangement that bidder 1 passes part of his revenues from bidder 2 on to the initial owner. Note that $\bar{p}_1(\hat{t}_1) - \bar{p}_1^*(\hat{t}_1)$ is equal to the amount by which bidder 2 indirectly pays the initial owner via bidder 1, due to the no-discounting assumption.

$V_{2|z'}$ strictly increasing (proved by the hazard-rate assumption), $x \geq x'$. This proves the "if" part. Since every step in this derivation is reversible, the "only if" part follows. *Q.E.D.*

According to the above lemma, we add the following assumption to ensure that the bid-inflation rule satisfies the incentive constraint from resale.

ASSUMPTION 3 (Resale Monotonicity): *If bidder i is ranked before bidder j $(i < j)$, the function β_{ij} defined by* (5) *is monotone nondecreasing.*

This assumption requires that a higher bid from a favored bidder be inflated to a higher amount. Given this condition, a favored bidder who has won after submitting a higher bid thinks more highly about the losers' values and hence charges them higher reserve prices for the good; that means a higher probability of no resale. If this assumption is not satisfied, the bid-inflation rule β_{ij} need not satisfy the incentive constraint for a favored bidder, and we need to consider other kinds of winner-selection rules. For simplicity, I focus on the cases where bid-inflation rules do achieve the Myerson surplus. Next is the main result in this section.

PROPOSITION 3: *Assume that there are only two bidders and their prior distributions satisfy the assumptions of hazard rate, uniform bias, and resale monotonicity. Then the mechanism designed above is optimal for the initial owner and implements the Myerson allocation.*

PROOF: Let us denote this mechanism by M^*. As step one, we show that it suffices to prove that each bidder is truthful in M^*. By Lemma 2.1, it suffices in this step to show that truth-telling in M^* implies that (i) the Myerson allocation is implemented and (ii) bidders receive exactly their Myerson surpluses. Claim (ii) follows from the construction of the payment scheme (equations (7) and (8)) and the fact that bidder 2's payment function is the same as that constructed by Myerson (1981), given bidder 1's equilibrium resale price. Let us prove claim (i). Since the reserve prices of M^* are the same as those in the Myerson allocation, "no sale" occurs exactly when it is needed. If the good is sold, each bidder would follow the hypothetical equilibrium thereafter: If bidder 2 wins in M^*, she infers that her value is higher than bidder 1's even after the latter is inflated; hence bidder 2 selects not to resell the good. If bidder 1 wins, he offers resale only to bidder 2 (Lemma 3.1); hence he makes a take-it-or-leave offer at the resale price prescribed by the hypothetical equilibrium, as calculated after equation (6). Given his truth-telling in M^*, this resale price matches precisely the Myerson allocation. Hence claim (i) follows.

We next show that bidder 2 is truthful in M^* given the equilibrium hypothesis. That is because her bid in M^* can only affect *when* she becomes the final owner and hence she is indifferent about that. First, she pays the same price $V_2^{-1}(V_1(t_1))$ for the good whether she wins it by outbidding bidder 1 in M^* or she buys it from him later. (This follows from the construction of M^* and the calculation of bidder 1's optimal resale price, done after equation (6).) Hence bidder 2 cannot profit from underbidding in M^*. Second, overbidding in M^* is not profitable by Lemma 3.1. That is because, in outbidding her rival, bidder 2 needs to pay $V_2^{-1}(V_1(t_1))$, which is more than bidder 1's value t_1 by the uniform-bias assumption.

Finally, we prove the honesty of bidder 1 by Proposition 2. With Lemma 4.1 and the resale-monotonicity assumption, equation (2) is always satisfied. Hence it suffices to verify all the conditions of that proposition. (i) As proved previously, the mechanism M^* coupled with truth-telling would generate the Myerson allocation, which we know is BNE-feasible. (ii) After the operation of M^*, the hypothetical equilibrium projected on any continuation game is clearly sequentially rational. (Here the only nontrivial point to check is that the resale offer $V_{2|\hat{t}_1}^{-1}(t_1)$ is optimal for bidder 1, already demonstrated after equation (6).) (iii) The support of bidder 1's type is convex by Assumption 1. Thus, Proposition 2 applies, as desired. *Q.E.D.*

To illustrate Proposition 3, let us consider as *Example 2* a two-bidder environment. The value of bidder i ($i = 1, 2$) is uniformly distributed on $[0, \bar{t}_i]$, with $\bar{t}_2 > \bar{t}_1$. Let

$$\Delta := \bar{t}_2 - \bar{t}_1.$$

The initial owner's value of the good is zero. This environment is depicted in Figure 1, where the rectangle IODF represents the space of possible pairs of the bidders' values. The Myerson allocation is to award the good to bidder 1 in region IJLG, award it to bidder 2 in region GLCDF, and keep the good in region JOCL. Hence the allocation picks the final owner by inflating bidder 1's bid t_1

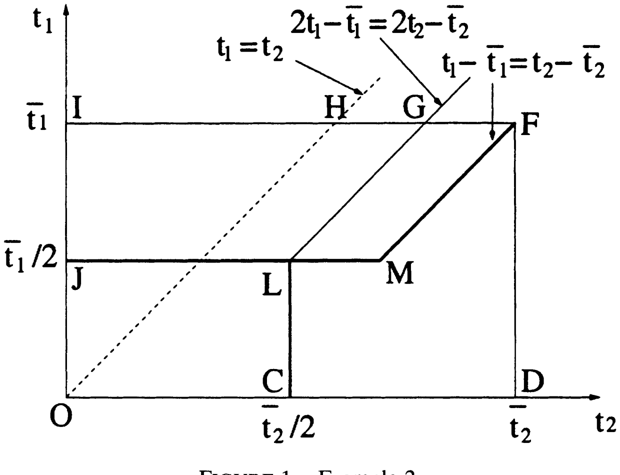

FIGURE 1.—Example 2.

to $t_1 + \Delta/2$ such that bidder 2 cannot be the final owner if his bid is below the inflated level (line GL).

The optimal mechanism M^* inflates bidder 1's bid more than the Myerson allocation does, up to the level $t_1 + \Delta$ (line FM), so that bidder 1 wins in the bigger region IJMF. Consequently, when bidder 1 wins after reporting his true type t_1, he learns that bidder 2's value is uniformly distributed on $[0, t_1 + \Delta]$. Hence bidder 1 finds it optimal to make a take-it-or-leave offer to bidder 2 at the price $t_1 + \Delta/2$, which is exactly the cutoff in the Myerson allocation. To ensure truth-telling from bidder 1, the initial owner promises to charge him a price $p_1(\hat{t}_1)$ if he bids $\hat{t}_1$ and wins, where the price calculated from (7) and (8) is

$$p_1(\hat{t}_1) = \frac{(2\hat{t}_1 + \Delta)^2 + \bar{t}_2^2}{8(\hat{t}_1 + \Delta)}.$$

We can verify that M^* gives the initial owner his Myerson surplus without using the revenue equivalence theorem. Since the payment scheme in M^* differs from that in Myerson (1981) only when the profile of types belongs to the area IJMF, it suffices to show that the expected revenues over this area are identical between the schemes. For M^*, the expected revenue is

$$\int_{\bar{t}_1/2}^{\bar{t}_1} \int_0^{t_1+\Delta} p_1(t_1)\, dt_2\, dt_1/(\bar{t}_1 \bar{t}_2) = \frac{1}{8\bar{t}_1 \bar{t}_2} \int_{\bar{t}_1/2}^{\bar{t}_1} ((2t_1 + \Delta)^2 + \bar{t}_2^2)\, dt_1$$

by the p_1 calculated above. For the scheme in Myerson (1981), the expected revenue is

$$\int_{\bar{t}_1/2}^{\bar{t}_1} \frac{1}{\bar{t}_2} \left(\int_0^{\bar{t}_2/2} \frac{\bar{t}_1}{2}\, dt_2 + \int_{\bar{t}_2/2}^{t_1+\Delta/2} (t_2 - \Delta/2)\, dt_2 + \int_{t_1+\Delta/2}^{t_1+\Delta} (t_1 + \Delta/2)\, dt_2 \right) dt_1/\bar{t}_1.$$

One can show that the two quantities are equal to each other, as desired.

5. THE n-BIDDER CASE

The main question in the two-bidder case is how to align a reseller's resale decision with the initial owner's preferences on the bidders. The answer is to inflate the favored bidder's bid to such a level that his resale decision conditional on winning is aligned with the initial owner's preferences between the bidders. Therefore, if there is a third bidder, say 3, bidder 1's inflated bid against 3 is determined by the initial owner's preferences between bidders 1 and 3; hence the inflated bid is generally different from the one against bidder 2. Thus, generalized to the n-bidder case, the new mechanism needs to conduct pairwise comparisons among the bidders so that the comparison criterion is pair-specific.

To find a winner through pairwise comparisons, we need to avoid cycles such as "1 beats 2, 2 beats 3, and yet 3 beats 1." To do so, we add an assumption. For any bidders i and j with $i < j$ and for any t_j less than or equal to the supremum of the range of β_{ij}, let $\beta_{ij}^{-1}(t_j) := \inf\{t_i \in T_i \mid \beta_{ij}(t_i) \geq t_j\}$. Hence $\beta_{ij}(t_i) \geq t_j$ implies $t_i \geq \beta_{ij}^{-1}(t_j)$.

ASSUMPTION 4 (Transitivity): *If bidder i is ranked before bidder j and j is ranked before bidder k ($i < j < k$), then for any t_j less than or equal to the supremum of the range of β_{ij}, $\beta_{ik}(\beta_{ij}^{-1}(t_j)) \geq V_k^{-1}(V_j(t_j))$.*

Since the functions β_{ij} and V_i are determined by the prior distributions, this is an assumption about the primitives. It says that, when bidders 1 and 2 tie in their pairwise comparison, if bidder 3 outbids bidder 1 and hence bidder 2, then bidder 3 would have defeated bidder 2 should the Myerson allocation be the winner-selection rule. Thus, when we conduct pairwise comparisons with the bid-inflation rules β_{ij}, if 1 defeats 2, and 3 defeats 1, then bidder 2 should not be the final owner and so we lose nothing by skipping 2 instead of comparing him with bidder 3. The next lemma demonstrates this transitivity-like property, which earns the name for the assumption.

LEMMA 5.1: *Assume that the prior distributions satisfy the assumptions of hazard rate, uniform bias, resale monotonicity, and transitivity. If $i < j < k$, $\beta_{ij}(t_i) \geq t_j$, and $\beta_{ik}(t_i) < t_k$, then $V_k(t_k) > V_j(t_j)$ and $t_k > t_j$.*

PROOF: Since $\beta_{ij}(t_i) \geq t_j$, the definition of β_{ij}^{-1} implies $t_i \geq \beta_{ij}^{-1}(t_j)$. Thus,

$$t_k > \beta_{ik}(t_i) \geq \beta_{ik}(\beta_{ij}^{-1}(t_j)) \geq V_k^{-1}(V_j(t_j)) \geq V_k^{-1}(V_k(t_j)) = t_j.$$

Here the second and third inequalities come from the assumptions of resale monotonicity and transitivity, respectively; the last inequality is due to the fact $V_j(t_j) \geq V_k(t_j)$ (the uniform-bias assumption). Note that the third inequality implies $V_k(t_k) > V_j(t_j)$. *Q.E.D.*

Suppose we sell the good to the lowest-indexed bidder who is not outbid by those ranked after him. That is, given any profile $(\hat{t}_i)_{i=1}^n$ of bids, the winner goes to the smallest $w \in \{1, \ldots, n\}$ such that $\beta_{wi}(\hat{t}_w) \geq \hat{t}_i$ for all $i > w$. Given truth-telling in this selection process, an inductive application of Lemma 5.1 implies the fact that $t_w > t_j$ and $V_w(t_w) > V_j(t_j)$ for all $j < w$. Hence the winner w finds it optimal to offer resale only to bidders $i > w$ (Lemma 3.1), and the initial owner also likes w to do so.

If he resells the good, the winner needs to select a buyer from bidders $w+1, \ldots, n$. We need to align this decision with the initial owner's preferences on them. Say the winner has only two potential buyers, and our result in the two-bidder case applies to his environment. Then the winner would implement a Myerson allocation that favors one bidder against the other. However, he may favor a bidder against whom the initial owner would like to discriminate, because the Myerson allocation based on the winner's posterior belief may be different from the Myerson allocation based on the prior belief. In general, this tension makes it difficult for the initial owner to recover his Myerson surplus. For tractability, I add the last assumption on the prior distributions so that the tension is absent when the pairwise bid-inflation rules β_{ij} are used.

ASSUMPTION 5 (Invariance): *For any* $w \in \{1, \dots, n\}$ *and any* $i, j > w$, *if* $t_i \leq \beta_{wi}(t_w)$ *and* $t_j \leq \beta_{wj}(t_w)$, *then* $V_i(t_i) >$ *(resp.* $=$*)* $V_j(t_j)$ *implies* $f_i(V_i^{-1}(V_w(t_w)))/f_i(t_i) >$ *(resp.* $=$*)* $f_j(V_j^{-1}(V_w(t_w)))/f_j(t_j)$.

Since the functions V_i and β_{wi} come from the prior distributions, this assumption is about the primitives. The assumption implies that, whenever a bidder w's β-inflated bids defeat some less favored bidders i and j, the virtual-utility ranking between i and j is invariant before and after the announcement of w's victory and bid, i.e., $V_i(t_i) >$ (resp. $=$) $V_j(t_j)$ is equivalent to $V_{i|t_w}(t_i) >$ (resp. $=$) $V_{j|t_w}(t_j)$. To prove this, one calculates $V_{i|t_w}(t_i) - V_{j|t_w}(t_j)$ by equations (4) and (5). This assumption is very restrictive, but it facilitates tractability. After proving the main result, I will present an example that satisfies all the above assumptions.

With the invariance assumption, the initial owner's task becomes implementing the pair-specific winner-selection rule based on β_{ij}. That leads to a new question: how to take care of a bidder's incentive when an auction both favors him against a rival and favors another rival against him? He should not be overly aggressive in competing with those ranked before him, but he should be aggressive enough to win at the right time so that he can probably resell the good to those ranked behind him. To accommodate both aspects, I design a multiple-round procedure where a bidder plays different roles at different rounds. In each round, only one of the bidders, called *leader* l, can possibly win and the other bidders can only affect whether l wins. We designate the most favored player, bidder 1, as the leader for the first round. In each round, the current leader l and those ranked behind him submit sealed bids, say i bids $\hat{t}_i$. If the leader's is not outbid ($\beta_{li}(\hat{t}_l) \geq \hat{t}_i$ for all $i > l$), he wins and trades with the initial owner. If the leader is outbid, we skip him and designate the most favored bidder who has outbid him as the new leader. Repeat the process until a leader wins or the least favored bidder n is reached. In every round, the leader never loses the good to those ranked before him, because they have been skipped, and the nonleaders can do nothing but help to determine whether the leader wins or not.

Since those ranked before the leader are not allowed to bid, we need to pass their competitive pressure on to the leader. Hence we update the minimum-bid requirement for the leader according to the bids submitted by them. If he bids below the required level or does not participate, skip him and designate the bidder immediately behind him as the new leader. How much should we update the minimum bid? Since a winner will not resell to a lost bidder ranked before him (shown later), the decision of letting the leader l instead of a skipped bidder i win should be consistent with the initial owner's preference between i and l. As this preference is determined by comparing the bidders' virtual utilities, the leader's minimum bid should be raised at least to the level $V_l^{-1}(V_i(\hat{t}_i))$, given the skipped bidder's $\hat{t}_i$.

The Mechanism. Based on the above ideas, let us define a mechanism for the initial owner by the following algorithm. Here l denotes the leader, w denotes the player who gets the good at the end of the current stage, and $\underline{b}_i$ denotes the minimum bid for bidder i.

$w := 0$; $\underline{b}_i := \max\{\inf T_i, V_i^{-1}(t_0)\}$ $(\forall i = 1, \ldots, n)$; $l := 1$.
While $l \leq n$ **do**
secretly inform l of $\underline{b}_l$;
each $i \in \{l, \ldots, n\}$ bids $\hat{t}_i$ independently and secretly;
bids can be different from those in previous rounds;
if i does not participate, then set $\hat{t}_i := \inf T_i$;
if $\hat{t}_l < \underline{b}_l$ or l does not participate:
then $\underline{b}_j := \max\{\underline{b}_j, V_j^{-1}(V_l(\hat{t}_l))\}$ $(\forall j > l)$,
$l := l + 1$;
else if $\beta_{li}(\hat{t}_l) \geq \hat{t}_i$ $(\forall i > l)$:
then
trade with l at price $p_l(\hat{t}_l \mid \underline{b}_l)$ determined by equation (10);
publicly announce $\hat{t}_l$;
$w := l$ and **halt**;
else
$i_1 :=$ the lowest indexed $i > l$ with $\beta_{li}(\hat{t}_l) < \hat{t}_i$,
$\underline{b}_j := \max\{\underline{b}_j, \max\{V_j^{-1}(V_i(\hat{t}_i)) : i = l, \ldots, i_1 - 1\}\}$ $(\forall j > l)$,
$l := i_1$.

Note: If all bidders are truthful in this mechanism, the winner is the lowest-indexed bidder whose inflated bids are not outbid by those ranked behind him. That is exactly the winner-selection rule motivated after Lemma 5.1. Notice the control of information: First, the winning bid is publicly announced, so the new upper bound of a loser's type is commonly known. Second, the winner's payment $p_l(\hat{t}_l \mid b_l)$ implies nothing about the bidders ranked behind him. Third, a leader's inference about the skipped bidders' types from his minimum bid is harmless because he will not resell the good to them on the equilibrium path. The updated minimum bid is told secretly just to eliminate any unnecessary information linkage. Fourth, nonparticipation from a bidder is not observed by others.

The Payment Scheme. In our mechanism, a bidder pays only if he trades with the owner. Hence we need only to design the payment function for a leader when he wins. The construction is the same as equation (8) in the two-bidder case, except that the function now is contingent on the leader's minimum bid $\underline{b}_l$. Let $U_l^*(t_l \mid \underline{b}_l)$ denote a type-t_l leader l's Myerson surplus conditional on the event that

$$\underline{b}_l = \max\left\{\inf T_l, V_l^{-1}(t_0), \max_{i<l} V_l^{-1}(V_i(t_i))\right\}. \tag{9}$$

Let $\pi_l^*(t_l \mid \underline{b}_l)$ denote the Myerson surplus for l in the posterior environment after everyone reports truthfully in the mechanism and leader l wins, conditional on the event (9). As we will see ((d) of step one in the proof of Proposition 4), $\pi_l^*(t_l \mid \underline{b}_l)$ is the leader's maximum expected profit from possible resale on the equilibrium path. Mimicking equation (8), we define a leader l's payment $p_l(\hat{t}_l \mid \underline{b}_l)$ for each $\hat{t}_l \geq \underline{b}_l$:

$$p_l(\hat{t}_l \mid \underline{b}_l) := \begin{cases} \hat{t}_l + \pi_l^*(\hat{t}_l \mid \underline{b}_l) - \dfrac{U_l^*(\hat{t}_l \mid \underline{b}_l)}{\prod_{i=l+1}^n F_i(\beta_{li}(\hat{t}_l))} & \text{if } l < n, \\ \underline{b}_l & \text{if } l = n. \end{cases} \tag{10}$$

In the first branch, the denominator $\prod_{i=l+1}^n F_i(\beta_{li}(\hat{t}_l))$ is the probability for leader l to win given $\hat{t}_l \geq \underline{b}_l$. In the second branch, the leader n bids only against his minimum bid. Thus, conditional on $\hat{t}_n \geq \underline{b}_n$, his probability of winning is one and $U_n^*(\hat{t}_n \mid \underline{b}_n) = \hat{t}_n - \underline{b}_n$. As we will demonstrate, a winner does not resell to those ranked before him, $\pi_n^*(\hat{t}_n \mid \underline{b}_n) = 0$. Hence the second branch should be $\underline{b}_n$.

LEMMA 5.2: *If the prior distributions satisfy the assumptions of hazard rate, uniform bias, and resale monotonicity, a winner in the above mechanism pays at least his minimum bid.*

PROOF: To win, a bidder l must be the leader and make some bid $\hat{t}_l \geq \underline{b}_l$. Let $\tilde{q}_l(t_i \mid \underline{b}_l)$ denote the probability for a type-t_l leader l to be the final owner in the Myerson allocation, conditional on the event (9). The revenue equivalence theorem (which uses the convexity of T_l in Assumption 1) says $U_l^*(\hat{t}_l \mid \underline{b}_l) - U_l^*(\underline{b}_l \mid \underline{b}_l) = \int_{\underline{b}_l}^{\hat{t}_l} \tilde{q}_l(z \mid \underline{b}_l)\,dz$. Since bidder l with any type below the right-hand side of (9) cannot be the final owner in the Myerson allocation, $U_l^*(\underline{b}_l \mid \underline{b}_l) = 0$ by a continuity argument. Thus, $U_l^*(\hat{t}_l \mid \underline{b}_l) = \int_{\underline{b}_l}^{\hat{t}_l} \tilde{q}_l(z \mid \underline{b}_l)\,dz$. It follows from equation (10) that $p_l(\underline{b}_l \mid \underline{b}_l) \geq \underline{b}_l$, since $\pi_l \geq 0$ by definition. Let us consider the other case, where $\hat{t}_l > \underline{b}_l$. By equation (10) and the nonnegativity of π_l,

$$\int_{\underline{b}_l}^{\hat{t}_l} \tilde{q}_l(z \mid \underline{b}_l)\,dz \geq [\hat{t}_l - p_l(\hat{t}_l \mid \underline{b}_l)] \prod_{i=l+1}^{n} F_i(\beta_{li}(\hat{t}_l)).$$

Suppose that $\hat{t}_l > \underline{b}_l$ and $p_l(\hat{t}_l \mid \underline{b}_l) < \underline{b}_l$; then for some $\xi \in (\underline{b}_l, \hat{t}_l)$,

$$(\hat{t}_l - \underline{b}_l)\tilde{q}_l(\xi \mid \underline{b}_l) > (\hat{t}_l - \underline{b}_l) \prod_{i=l+1}^{n} F_i(\beta_{li}(\hat{t}_l)).$$

By the resale-monotonicity assumption, $\prod_{i=l+1}^n F_i(\beta_{li}(\cdot))$ is nondecreasing, hence $\tilde{q}_l(\xi \mid \underline{b}_l) > \prod_{i=l+1}^n F_i(\beta_{li}(\xi))$. Thus, the leader's probability of being the final owner in the Myerson allocation is greater than his probability of winning in the mechanism. That contradicts the fact that the mechanism is more biased to a leader than the Myerson allocation ($\beta_{li}(t_l) \geq V_i^{-1}(V_l(t_l))$ if $l < i$, by the uniform-bias assumption). *Q.E.D.*

LEMMA 5.3: *If the prior distributions satisfy the assumptions of hazard rule, uniform bias, resale monotonicity, and transitivity, then for any $i < j < k$ and any $t_i \in T_i$, $\beta_{ij}(t_i) \leq \beta_{ik}(t_i)$.*

PROOF: Let $t_i \in T_i$. By the definition of β_{ij}^{-1}, $t_i \geq \beta_{ij}^{-1}(\beta_{ij}(t_i))$. Thus,

$$\beta_{ik}(t_i) \geq \beta_{ik}(\beta_{ij}^{-1}(\beta_{ij}(t_i))) \geq V_k^{-1}(V_j(\beta_{ij}(t_i))) \geq \beta_{ij}(t_i);$$

here the first inequality comes from the monotonicity of β_{ik} (resale monotonicity), the second inequality is from transitivity ($i < j < k$), and the third is from uniform-bias ($j < k$). *Q.E.D.*

Next is the main result. It says that the above mechanism generates the Myerson surplus for the initial owner, the best outcome for him even when he could ban resale costlessly. Thus, to some extent Myerson's characterization of seller-optimal allocations is robust even when resale cannot be banned. A cautionary note, however, is that the proposition is based on restrictive assumptions.

PROPOSITION 4: *Suppose that the prior distributions satisfy the assumptions of hazard rate, uniform bias, resale monotonicity, transitivity, and invariance. Then the mechanism designed above is optimal for the initial owner and implements the Myerson allocation.*

PROOF: We shall apply induction on the number of bidders. The one-bidder case is trivial. Pick any $n = 2, 3, \ldots$ and suppose that the proposition is true when the number of bidders is less than n. Let us prove the proposition when the number is n. Denote the mechanism designed above by M^*. By Lemma 2.1, it suffices to show that M^* (i) implements the Myerson allocation and (ii) gives each bidder his Myerson surplus.

First, we show that both (i) and (ii) are fulfilled if all bidders are truthful in M^*. To prove (i), let us analyze the consequence of truth-telling in M^*. If the mechanism results in no sale, then every bidder i's type is below the initial minimum bid $V_i^{-1}(t_0)$, which is exactly the case where the Myerson allocation requires no-sale. Consider the other case and let bidder w be the winner. We claim that his resale decision is consistent with the Myerson allocation. The claim is proved in five steps.

(a) The winner chooses to offer resale to only bidders ranked behind him. By an inductive application of Lemma 5.1, any bidder ranked before winner w has a lower type than w; hence the first sentence of Lemma 3.1 applies.

(b) Thus, the winner's potential buyers are the bidders in $\{w+1, \ldots, n\}$. The posterior belief about every such bidder i is commonly held and is that i's type is independently drawn from the conditional distribution $F_i(\cdot \mid t_i \leq \beta_{wi}(t_w))$. That is because the winner's identity and bid (hence type) constitute the entire new public information after the mechanism M^* is over.

(c) Assumptions from 1 to 5 are satisfied with these posterior distributions taking the role of the priors: Coming from a lower truncation of the prior support, a posterior satisfies the hazard-rate assumption. Since the priors satisfy the invariance assumption, the ranking of the virtual utilities is invariant from the prior $(V_i(t_i))_{i=w+1}^n$ to the posterior $(V_{i|t_w}(t_i))_{i=w+1}^n$. This, coupled with the fact that $\beta_{wj}(t_w) \leq \beta_{wk}(t_w)$ whenever $w < j < k$ (Lemma 5.3), implies that the posteriors satisfy the uniform-bias assumption. The invariance of the ranking also implies that $V_{j|t_w}^{-1}(V_{i|t_w}(t_j)) = V_j^{-1}(V_i(t_j))$ for all t_j if $w < i < j$. Thus, the inflation rule β_{ij} determined by the priors via (5) coincides with the β_{ij} determined by the posteriors on the posterior support. Hence the posteriors satisfy the assumptions of resale monotonicity, transitivity, and invariance.

(d) Thus, the winner's environment satisfies all the assumptions required by the proposition and consists of less than n bidders. Since the winner is restricted to transparent mechanisms for resale, his Myerson surplus in the posterior environment is the best he can achieve (Lemma 2.1). Thus, the induction hypothesis implies that the winner is willing and able to implement the posterior Myerson allocation in the posterior environment. Thus, for every profile $(t_i)_{i=w}^n$ of types, the final owner is either a bidder $w' > w$ such that $V_{w'|t_w}(t_{w'}) = \max\{V_{i|t_w}(t_i): i > w\}$ and $V_{w'|t_w}(t_{w'}) \geq t_w$, or the current owner w if such a bidder does not exist. By the construction of $\beta_{ww'}$ (equation (5)), $V_{w'|t_w}(t_{w'}) \geq t_w$ is equivalent to $V_{w'}(t_{w'}) \geq V_w(t_w)$; by the invariance assumption, $V_{i|t_w}(t_i) \geq V_{j|t_w}(t_j)$ is equivalent to $V_i(t_i) \geq V_j(t_j)$. Thus, this allocation is equivalent to the outcome where the final owner is the player $w' \geq w$ such that $V_{w'}(t_{w'}) = \max\{V_i(t_i): i \geq w\}$.

(e) An inductive application of Lemma 5.1 shows that $V_j(t_j) < V_w(t_w)$ for all bidders $j < w$. Thus, the final owner w' has the highest virtual utility among all bidders $1, \ldots, n$. The Myerson allocation in the initial environment is hence implemented.

To prove that task (ii) is fulfilled, let us calculate a type-t_i bidder i's expected profit (for the entire game) $u_i^*(\hat{t}_i \mid t_i, \underline{b}_i)$ from bidding $\hat{t}_i$ in M^*, when he is leading and secretly informed of his minimum bid $\underline{b}_i$. Conditional on $\underline{b}_i$, let $Q_i(\hat{t}_i \mid \underline{b}_i)$ denote his probability of winning in M^* and let $\pi_i(\hat{t}_i \mid t_i, \underline{b}_i)$ denote i's optimal expected profit from possible resale conditional on the history that he bids $\hat{t}_i$ and wins in M^*. If he does not win in M^*, the bidder gets zero payoff since according to (a) he will not get the good in the future. Thus,

$$u_i^*(\hat{t}_i \mid t_i, \underline{b}_i) = Q_i(\hat{t}_i \mid \underline{b}_i)[t_i - p_i(\hat{t}_i \mid \underline{b}_i) + \pi_i(\hat{t}_i \mid t_i, \underline{b}_i)]. \tag{11}$$

By the reasoning in (d), $\pi_i(t_i \mid t_i, \underline{b}_i) = \pi_i^*(t_i \mid \underline{b}_i)$ with π_i^* defined after equation (9). Hence the construction (10) of the payment function implies $u_i^*(t_i \mid t_i, \underline{b}_i) = U_i^*(t_i \mid \underline{b}_i)$. Thus, with $\underline{b}_i$ being the random variable, the expected values $E_{\underline{b}_i} u_i^*(t_i \mid t_i, \underline{b}_i)$ and $E_{\underline{b}_i} U_i^*(t_i \mid \underline{b}_i)$ are the same. The former expected value is bidder i's expected profit from truth-telling (because a bidder is engaged in no trade unless he leads) and the latter is his Myerson surplus, both calculated before M^* is operated. Thus, truth-telling in M^* induces Myerson surplus for every bidder and hence fulfills task (ii).

Second, we show that a bidder is truthful when he leads. By a trivial extension of Lemma 4.1 based on the resale-monotonicity assumption, the probability with which a leader becomes the final owner is increasing in his bid in M^*. Hence inequality (2) is always satisfied. We need only to verify all the conditions of Proposition 2. The first condition (BNE-feasibility) follows from the proved fact that the allocation is Myerson. The third condition (convex support for the leader's type) is true by assumption. The second condition is proved below: (a) If leader l wins, he pays a price above the value t_j of any bidder $j < l$. That is because the payment is at least the minimum bid (Lemma 5.2), which is at least $V_l^{-1}(V_j(t_j))$ (equation (9)), which in turn is at least t_j by the uniform-bias assumption. (b) It follows from (a) and Lemma 3.1 that we can eliminate those bids $\hat{t}_l$ such that after winning with $\hat{t}_l$ the leader would find it optimal to resell the good to those ranked before him. (c) Thus, if he wins after bidding $\hat{t}_l$, the leader's environment will effectively consist of bidders $i > l$, with commonly held posteriors $F_i(\cdot \mid \beta_{li}(\hat{t}_l) \geq t_i)$. As shown in (c) of the previous paragraph, these posteriors satisfy Assumptions from 1 to 5. Hence the induction hypothesis implies that the winning leader will find it optimal (and feasible) to implement the Myerson allocation in that posterior environment. Hence the second condition of Proposition 2 is satisfied, as desired.

Finally, we show that any bidder i is truthful when he is not leading. Let l denote the current leader. There are only two cases. Case one is that there is a bidder w such that $l < w < i$ and w's type outbids l's inflated type $\beta_{lw}(t_l)$. Then bidder i's bid cannot affect anything, since the leader in the next round will be a bidder ranked before him. Case two is that there is no such w specified in case one. Then bidder i's bid affects only (i) whether he will lead in the next round and (ii) the posterior belief about him if he is skipped. If he is skipped, the bidder gets zero payoff; hence his deviation makes no difference if it does not change whether he will lead next. Even if the deviation changes that event, it still makes no difference to his payoff, because he will be truthful when he leads (demonstrated above) and no difference is made if he does not lead. Thus, a nonleader has no strict incentive to deviate, as asserted. *Q.E.D.*

The equilibrium obtained above is interesting in its own right. It exhibits a self-similar tower of optimal auctions via optimal resale auctions. On the equilibrium path, the environment inherited by a new owner is similar to the environment of his immediate predecessor. Although an owner is allowed to resell the good to previous owners, he chooses not to do so because from the previous sales he infers that previous owners have lower values. Likewise, a winner does not resell the good to the losing bidders favored against him, because from his victory the winner knows that his value is higher than theirs even after they are inflated. The next corollary summarizes these features.

COROLLARY 5.1: *Suppose that the prior distributions satisfy Assumptions* 1 *through* 5. *Then the mechanism-selection game has an equilibrium where the initial owner selects the mechanism designed above, every subsequent owner* $w \in \{1, \ldots, n-1\}$ *selects the same mechanism, with the priors updated, that offers resale to bidders in* $\{w+1, \ldots, n\}$, *and player n selects to keep the good if he owns it. The mechanism chosen by any player* $i < n$ *results in a sale with a positive probability.*

Let us consider as *Example 3* an n-bidder environment. Bidder i's type is uniformly distributed on $[\underline{t}_i, \bar{t}_i]$, with $\underline{t}_n \leq \cdots \leq \underline{t}_1 \leq t_0 \leq \bar{t}_1 \leq \cdots \leq \bar{t}_n$. For $i < j < k$, we calculate:

$$V_i(t_i) = 2t_i - \bar{t}_i;$$

$$V_k^{-1}(V_j(t_j)) = t_j + (\bar{t}_k - \bar{t}_j)/2;$$

$$\beta_{ij}(t_i) = t_i + \bar{t}_j - \bar{t}_i;$$
$$\beta_{ik}(\beta_{ij}^{-1}(t_j)) = t_j + \bar{t}_k - \bar{t}_j;$$
$$V_{j|t_i}(t_j) = 2t_j - (t_i + \bar{t}_j - \bar{t}_i).$$

Hence the example satisfies Assumptions 1 through 5, and the Myerson allocation favors the bidders in descending order of their labels $(1, \ldots, n)$. That is, if $i < j$, the allocation wants i instead of j to be the final owner if and only if $t_i \geq t_j + (\bar{t}_j - \bar{t}_i)/2$. By Proposition 1, this allocation cannot be implemented without resale.

To implement this allocation, the initial owner uses the mechanism M^*, which offers rival-specific bid credits to each favored bidder. That is, the bid $\hat{t}_i$ from bidder i is inflated to $\hat{t}_i + \bar{t}_j - \bar{t}_i$ when he is compared to bidder $j = i+1, \ldots, n$, so that j cannot outbid i unless j's bid is higher than i's inflated bid. Conditional on a sale, the winner is the first bidder along the list $(1, \ldots, n)$ whose inflated bids are not outbid by those listed behind him. Say the winner is w. The fact that he has defeated those ranked before him implies that their values, though inflated, are lower than the winner's. Hence he offers resale only to bidders ranked behind him. The fact that no such bidder $j > w$ has outbid him implies that j's value does not exceed w's inflated bid against j. Hence the posterior belief about j is just the prior conditional on this event, and this belief is commonly held due to the control of information in the mechanism M^*. Thus, the winner's environment is similar to the initial owner's: there are $n - w$ bidders and bidder j's value is uniformly distributed on $[0, t_w + \bar{t}_j - \bar{t}_w]$; furthermore, the ranking on the bidders $(w+1, \ldots, n)$ is invariant when the virtual utilities switch from the prior V_i to posterior $V_{i|t_w}$. Thus, the winner wants to implement an allocation that coincides with the Myerson allocation from the initial owner's viewpoint. That is how a seller achieves his optimum through resale.

This example also exhibits an interesting pattern in the selection of middlemen: a good is first offered to the bidder whose value is least uncertain to the owner. The wider is the spread of one's value, the more downstream is the bidder in the distribution channel.

6. CONCLUDING REMARKS

The above results are relevant to a couple of general questions. A frequently asked question is why we should care about auction design at all given secondary markets. Behind the question is the conjecture that an item sold in an initial auction will eventually go to a bidder who values it most. The results in the paper imply that this conjecture is false even in the simplest single-item private-value environment with perfectly patient bidders. Resale among bidders is hampered by their asymmetric information and hence cannot fully undo the inefficiency of the previous sale. Hence sellers may be able to achieve their desired levels of inefficiency by mechanisms that cancel out the partial correction effect of resale.

Another general question is about the role of middlemen. A traditional view is that middlemen are to bypass exogenous restrictions in matching and trading. In this paper, we have seen that a seller achieves his optimum through and only through resale. Here the emergence of middlemen does not require the usual barriers to multilateral trades.

Two modeling aspects in this paper are worthwhile noticing. One is that the model allows players to choose how the resale game should be played. Not only does this formulation capture some interactions between sellers and resellers, it is also aligned with a fundamental appeal of mechanism design: economic predictions should be robust even when people can modify institutions for their own interests. The other aspect is that the mechanism-selection process is really decentralized: the final outcome results from the various mechanisms selected by individual players, and no single mechanism can reach it. This suggests a hope that mechanism design can be applied to various economic and political institutions, which are often shaped by different players at different times.

Department of Economics, Northwestern University, 2003 Sheridan Road, Evanston, IL 60208-2600, U.S.A.; czheng@northwestern.edu; http://faculty-web.at.northwestern.edu/economics/zheng/.

Manuscript received October, 2000; final revision received April, 2002.

REFERENCES

AUSUBEL, L. M., AND P. CRAMTON (1999): "The Optimality of Being Efficient," Mimeo, Department of Economics, University of Maryland, College Park.

BOSE, S., AND G. DELTAS (1999): "Welcoming the Middlemen: Restricting Competition in Auctions by Excluding the Consumers," in *Current Trends in Economics: Theory and Applications*, Volume 8 of *Studies in Economic Theory*, ed. by A. Alkan, C. Aliprantis, and N. Yannelis. Berlin, New York: Springer Verlag.

CALZOLARI, G., AND A. PAVAN (2001): "Monopoly with Resale," Mimeo, Department of Economics, Northwestern University.

HAILE, P. A. (1999): "Auctions with Resale," Mimeo, Department of Economics, University of Wisconsin-Madison.

——— (2000): "Auctions with Private Uncertainty and Resale Opportunities," forthcoming in *Journal of Economic Theory*.

HARRIS, M., AND A. RAVIV (1981): "Allocation Mechanisms and the Design of Auctions," *Econometrica*, 49, 1477–1499.

JEHIEL, P., AND B. MOLDOVANU (1999): "Resale Markets and the Assignment of Property Rights," *Review of Economic Studies*, 66, 971–991.

MASKIN, E., AND J. TIROLE (1990): "The Principal-Agent Relationship with an Informed Principal: The Case of Private Values," *Econometrica*, 58, 379–409.

MCAFEE, R. P., AND D. VINCENT (1997): "Sequentially Optimal Auctions," *Games and Economic Behavior*, 18, 246–276.

MILGROM, P. R. (1987): "Auction Theory," in *Advances in Economic Theory: Fifth World Congress*, ed. by T. F. Bewley. Cambridge: Cambridge University Press, 1–32.

MILGROM, P. R., AND I. SEGAL (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70, 583–601.

MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

PORTER, R. H. (1992): "Review of *Anatomy of an Auction: Rare Books at Ruxley Lodge, 1919* by Arthur Freeman and Janet Ing Freeman," *Journal of Political Economy*, 100, 433–436.

RILEY, J. G., AND W. F. SAMUELSON (1981): "Optimal Auctions," *American Economic Review*, 71, 381–392.

LINKED CITATIONS

You have printed the following article:

Optimal Auction with Resale
Charles Zhoucheng Zheng
Econometrica, Vol. 70, No. 6. (Nov., 2002), pp. 2197-2224.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28200211%2970%3A6%3C2197%3AOAWR%3E2.0.CO%3B2-W

[Footnotes]

[4] **Review: [Untitled]**
Reviewed Work(s):
Anatomy of an Auction: Rare Books at Ruxley Lodge, 1919 by Arthur Freeman; Janet Ing Freeman
Robert H. Porter
The Journal of Political Economy, Vol. 100, No. 2. (Apr., 1992), pp. 433-436.
Stable URL:
http://links.jstor.org/sici?sici=0022-3808%28199204%29100%3A2%3C433%3AAOAARB%3E2.0.CO%3B2-M

References

Allocation Mechanisms and the Design of Auctions
Milton Harris; Artur Raviv
Econometrica, Vol. 49, No. 6. (Nov., 1981), pp. 1477-1499.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198111%2949%3A6%3C1477%3AAMATDO%3E2.0.CO%3B2-5

Resale Markets and the Assignment of Property Rights
Philippe Jehiel; Benny Moldovanu
The Review of Economic Studies, Vol. 66, No. 4. (Oct., 1999), pp. 971-991.
Stable URL:
http://links.jstor.org/sici?sici=0034-6527%28199910%2966%3A4%3C971%3ARMATAO%3E2.0.CO%3B2-Q

NOTE: *The reference numbering from the original has been maintained in this citation list.*

LINKED CITATIONS

The Principal-Agent Relationship with an Informed Principal: The Case of Private Values
Eric Maskin; Jean Tirole
Econometrica, Vol. 58, No. 2. (Mar., 1990), pp. 379-409.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28199003%2958%3A2%3C379%3ATPRWAI%3E2.0.CO%3B2-C

Envelope Theorems for Arbitrary Choice Sets
Paul Milgrom; Ilya Segal
Econometrica, Vol. 70, No. 2. (Mar., 2002), pp. 583-601.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28200203%2970%3A2%3C583%3AETFACS%3E2.0.CO%3B2-J

Review: [Untitled]
Reviewed Work(s):
Anatomy of an Auction: Rare Books at Ruxley Lodge, 1919 by Arthur Freeman; Janet Ing Freeman
Robert H. Porter
The Journal of Political Economy, Vol. 100, No. 2. (Apr., 1992), pp. 433-436.
Stable URL:
http://links.jstor.org/sici?sici=0022-3808%28199204%29100%3A2%3C433%3AAOAARB%3E2.0.CO%3B2-M

Optimal Auctions
John G. Riley; William F. Samuelson
The American Economic Review, Vol. 71, No. 3. (Jun., 1981), pp. 381-392.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28198106%2971%3A3%3C381%3AOA%3E2.0.CO%3B2-T

NOTE: *The reference numbering from the original has been maintained in this citation list.*