

Robustly Optimal Auctions with Unknown Resale Opportunities*

Gabriel Carroll
Ilya Segal

Department of Economics, Stanford University, Stanford, CA 94305

August 22, 2016

Abstract

We study robust revenue maximization by the designer of a single-object auction who has Bayesian beliefs about bidders' independent private values but is ignorant about post-auction resale opportunities (including possible leakage of private information). We show the optimality of a “Vickrey auction with bidder-specific reserve prices” proposed by Ausubel and Cramton (2004), which allocates the object efficiently provided that at least one of the bidders has bid above his reserve price. In this auction, truthful bidding and no resale is an ex post equilibrium for any individually rational resale procedure. We show optimality of this auction for a “worst-case” resale procedure in which the highest-value bidder learns all other bidders' values and has full bargaining power. The proof involves construction of Lagrange multipliers on the incentive constraints representing non-local deviations in which a bidder underbids to lose and then purchases from the auction's winner.

*We are grateful to Peter Cramton, Paul Milgrom, Glen Weyl, and participants of the 2016 Stony Brook Workshop on Complex Auctions and Practice for helpful comments. The authors thank the Simons Institute for its hospitality during the Fall 2015 Economics and Computation Program. Segal acknowledges the support of the National Science Foundation grant (SES 0961693). Alex Bloedel provided valuable research assistance.

1 Introduction

It is well known that revenue-maximizing auctions for settings with a priori asymmetric bidders implement inefficient allocations that are biased in favor of “weaker” bidders (Myerson 1981, McAfee and McMillan 1989). On the other hand, real-life auctions are often followed by resale. One might wonder whether the advantage of biased auctions is undermined due to a strong bidder’s ability to buy the object from another bidder in resale at a possibly lower price than what he would have to pay in the auction. This paper examines what auctions are optimal when resale is possible but the designer is ignorant about the resale procedure.¹

There exists a substantial literature on the design of optimal auctions when resale procedures are known, including Zheng (2002), Calzolari and Pavan (2006), and Dworzak (2015). The usual assumption in this literature is that no private information is leaked before resale (other than information about auction bids that the designer may choose to disclose). It is easy to see that under this assumption, the possibility of resale can only reduce the auctioneer’s optimal revenues: indeed, the equilibrium allocation of any auction mechanism followed by resale must be incentive compatible and therefore would be feasible for the designer if resale were impossible. Nevertheless, in spite of the above observation, it has been found that the optimal auction in the presence of known resale typically implements a biased allocation followed by resale.²

On the other hand, Ausubel and Cramton (2004) proposed a class of auctions, called “Vickrey auctions with reserve prices,” which induce an ex-post Nash equilibrium with truthful bidding and no resale. These auctions allocate the object efficiently provided that at least one of the bidders has beaten his reserve price. While these auctions have nice properties, in light of the literature mentioned above, it has been unclear why they would be optimal for the auctioneer.

Ausubel and Cramton (1999, Theorem 4) attempted to derive the op-

¹For expositional simplicity we focus on revenue maximization, but the analysis extends to the auctioneer’s profit maximization when she has a non-zero value for keeping the object, by defining bidders’ values to be net of the auctioneer’s value.

²In a simple example due to Zheng (2002), if a bidder is known to have a zero value for the good and the ability to design a revenue-maximizing resale mechanism, the designer would optimally sell to this bidder at the price equal to the optimal expected revenue without resale.

tinality of Vickrey auctions with reserves from the assumption of efficient resale. However, the foundations for this assumption are unclear. By the theorem of Myerson and Satterthwaite (1983), we would generally expect resale to be inefficient unless the parties’ private values are revealed before the resale. However, if resale takes place under symmetric information, then typically, an auctioneer who can anticipate the resale procedure would *not* want to allocate the object efficiently; instead she would again want to run an auction that induces resale in equilibrium. For an extreme example, suppose that the auctioneer knows that one particular bidder has a zero value, but will have full information and full bargaining power in resale. Then the optimal auction would sell to that bidder at the price equal to the expected maximal bidder value for the object, letting the auctioneer extract full surplus.

This paper provides the missing foundations. We derive the robust optimality of single-object Vickrey auctions with reserves when the designer is ignorant about the resale procedure (including possible private information revelation before resale). Namely, we show that such auctions maximize the designer’s worst-case expected revenue, where the expectation is taken over buyers’ independent private values and the worst case is over the possible resale procedures. Our conclusion is conceptually similar to the robust optimality of strategy-proof auctions when the designer is ignorant about bidders’ beliefs about each other’s values (Chung and Ely 2007) or about each other’s strategies (Yamashita 2015). Namely, while in those cases revenue maximization that is robust to bidder strategizing makes it optimal to use strategy-proof mechanisms, in our case revenue maximization that is robust to resale makes it optimal to use resale-proof mechanisms. However, the proof techniques are quite different. Also, observe that Vickrey auctions with reserve prices are robust to bidders’ beliefs about each other’s values and their beliefs about the resale procedure, hence we obtain those additional robustness benefits “for free.”

We begin with the simple case in which the auctioneer is restricted to always sell the object (Section 3). For this case, we show that the simple Vickrey auction (second-price sealed-bid auction) with no reserve price is optimal. To do this, we guess a “worst-case” resale procedure, in which the highest-value bidder learns other bidders’ values and has full bargaining power in resale. With this resale procedure, in any auction that always sells the object, the highest-value bidder would be able to extract at least his marginal contribution to the total surplus by bidding low to let another bidder win and then buying from the winner. Given that the designer is

unable to reduce bidders' information rents below their expected marginal contributions to the total surplus, she could do no better than the simple Vickrey auction with no reserve. Since this auction sustains truthful bidding as an ex post equilibrium under any resale procedure, it is robustly optimal with unknown resale.

We proceed to the more complex setting in which the designer can withhold the object, for simplicity starting with the two-bidder case (Section 4). We continue with the worst-case resale procedure guessed in Section 3, and derive bidders' reduced-form utilities from auction allocations under this resale procedure. Note that these reduced-form utilities exhibit both externalities and interdependent values, since a bidder who does not win cares whether the other bidder wins and, if so, what the other bidder's value is. Nevertheless, the usual envelope-theorem approach to local (first-order) incentive compatibility constraints yields a simple expression for bidders' information rents, which allows us to express the expected revenue as the expectation of an appropriately defined virtual surplus. If we were to ignore all other incentive constraints and solve the resulting relaxed problem by maximizing the virtual surplus state-by-state, the solution would always allocate the object between the bidders efficiently. Intuitively, selling to the inefficient bidder would yield information rents both to him and to the efficient bidder (who would buy it in resale), which is dominated by selling to the efficient bidder, eliminating the inefficient bidder's information rents. The solution to the relaxed problem sells to the efficient bidder if and only if his value exceeds the optimal reserve price for him.

Unfortunately, the solution to the relaxed problem violates non-local incentive constraints: a reduction in the bid of the "strong" bidder would sometimes give the object to the "weak" bidder rather than leave it unsold, giving the strong bidder an incentive to underbid and then buy in resale. We guess that the solution to the auctioneer's full problem is the Vickrey auction with reserves described by Ausubel and Cramton (1999, 2004). To establish that this is indeed a solution, we construct Lagrange multipliers on non-local downward incentive constraints such that maximization of the Lagrangian yields the solution. Since there is an incentive constraint for each type and each possible misreport, and a double continuum of such incentive constraints is binding, our Lagrange multipliers are defined by a measure over this double continuum. In the two-bidder case, we construct a product measure that works for our purpose: maximization of the Lagrangian (which can be written as the expectation of a "modified virtual surplus," which can

be maximized state-by-state) yields the solution.

Finally, in Section 5 we extend the approach to the case of many bidders. Some additional complications arise because it becomes necessary to consider binding non-local downward incentive constraints both *from* a given type (when this type is the highest-value bidder but may underreport to buy from another bidder in resale) and *to* the same type (when this type is reported by some higher type so as to concede the object to another bidder and buy it from him in resale). This necessitates a somewhat more complex construction of Lagrange multipliers. Under traditional regularity assumptions on value distributions we can construct nonnegative Lagrange multipliers that yield a Vickrey auction with reserves as an optimal auction. Since this auction sustains truthful bidding as an ex post equilibrium under any resale procedure, it is robustly optimal with unknown resale.

Our approach also yields an iterative construction of the optimal bidder-specific reserve prices in n -bidder Vickrey auctions with reserves. We illustrate this (in Section 6) for the case where bidders' values are distributed uniformly with different upper limits. In this case, the optimal reserve price to the k th bidder is obtained by solving a k th-degree equation, which cannot be done analytically for $k > 4$ but is easily done numerically.

2 Setup

There are n bidders. Bidder i 's value for the object is θ_i , which is distributed according to a c.d.f. F_i with a continuous strictly positive density f_i on support $[0, 1]$.³ We write $\nu_i(\theta_i) = \theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)}$ for the traditional virtual value of type θ_i . Values are independent across bidders. We write $\theta = (\theta_1, \dots, \theta_n)$ for the profile of values. The space of (possibly randomized) allocations is $X = \{x \in [0, 1]^n : \sum_i x_i \leq 1\}$, where $x_i \in [0, 1]$ is the probability of allocating the object to bidder i .

To describe the general space of mechanisms with resale, in general we need to think about how the mechanism influences resale. For example, the information disclosed by the mechanism will in general affect the outcome of resale. However, for simplicity we now restrict attention to resale proce-

³The assumption that the distributions have a common support is made for expositional simplicity: in Section 6 below we argue that the results extend to cases in which the supports' upper limits differ. The assumption of continuous density is also made for expositional simplicity, but it will follow from assumption (A1) made in Section 5 below.

dures in which all private information is revealed before resale (but after the auction), and the resale outcome depends only on the allocation specified by the mechanism but not on any other features of the mechanism. For such resale procedures, the expected post-resale payoff of bidder i net of payment specified in the mechanism can be written as a reduced-form function $v_i(x, \theta)$ of the allocation $x \in X$ specified by the mechanism and the bidders' value profile θ , with the total reduced-form payoffs not exceeding the maximal total surplus available in resale:

$$\sum_i v_i(x; \theta) \leq \max_i \theta_i \cdot \sum_i x_i. \quad (1)$$

We also require the resale procedure to be individually rational:

$$v_i(x; \theta) \geq \theta_i x_i \quad (2)$$

for each i .

We will specify a worst-case resale procedure of this form, but the optimal mechanisms we identify will prove to be robust to a broader class of resale procedures (described in Section 7 below).

If the designer knows the resale procedure and it is described by reduced-form resale payoffs, then we can appeal to the revelation principle and focus on mechanisms where each bidder directly reports θ_i . Thus an *auction* is a pair of measurable functions (χ, ψ) , where

- $\chi : [0, 1]^n \rightarrow X$ specifies the (possibly probabilistic) allocation rule;
- $\psi : [0, 1]^n \rightarrow \mathbb{R}^n$ specifies the payments.

An auction must satisfy the usual incentive compatibility and individual rationality constraints:

$$\mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \psi_i(\theta_i, \tilde{\theta}_{-i})] \geq \mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \psi_i(\hat{\theta}_i, \tilde{\theta}_{-i})] \quad \text{for all } \theta_i, \hat{\theta}_i; \quad (3)$$

$$\mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \psi_i(\theta_i, \tilde{\theta}_{-i})] \geq 0 \quad \text{for all } \theta_i. \quad (4)$$

The seller's expected revenue maximization problem is then

$$\max_{\substack{\chi: [0, 1]^n \rightarrow X \\ \psi: [0, 1]^n \rightarrow \mathbb{R}^n}} \mathbb{E}_{\tilde{\theta}} \left[\sum_i \psi_i(\tilde{\theta}) \right] \quad \text{subject to (3)-(4)}. \quad (5)$$

3 The Must-Sell Case

In this section, we restrict attention to auctions that must sell the object with probability 1. (For example, this could be motivated by the seller’s prohibitively high cost of keeping the object.) Recall from Myerson (1981) and McAfee and McMillan (1989) that if there is no resale, and if each bidder’s virtual value function is increasing, the optimal auction allocates the object to the bidder with the highest virtual value $\nu_i(\theta_i)$. Thus, if bidders have different distributions and therefore different virtual value functions, the optimal auction misallocates the object. In particular, the auction is biased towards “weaker” bidders, which are those with higher virtual value functions. Intuitively, optimal misallocation trades off reduction of bidders’ information rents against the reduction of the expected total surplus.

Now we turn to the analysis of auctions with resale, and make a guess about a worst-case resale procedure: that the highest-value bidder learns the values of everybody else and has full bargaining power. Let $i^*(\theta)$ denote the bidder with the highest value at profile θ , breaking ties in favor of earlier-numbered bidders (the choice of tiebreak is inconsequential). Then, each bidder i ’s reduced-form payoff is

$$v_i(x; \theta) = \begin{cases} \theta_i x_i + \sum_{j \neq i} (\theta_i - \theta_j) x_j & \text{if } i = i^*(\theta); \\ \theta_i x_i & \text{otherwise.} \end{cases} \quad (6)$$

Intuitively, this is a worst-case resale procedure because it gives bidders maximal information rents: by letting another bidder win and then buying from the winner whenever this is efficient, bidder i would get an expected information rent equal to at least his expected marginal contribution to the total surplus, $\mathbb{E}_{\tilde{\theta}} \left[\max \left\{ \tilde{\theta}_i - \max_{j \neq i} \tilde{\theta}_j, 0 \right\} \right]$.⁴ Given this, the expected revenue is maximized by the Vickrey auction with no reserve, which achieves the lower bound on bidders’ information rents and at the same time maximizes total surplus without any resale.

To complete the argument we only need to establish that a bidder can always submit a bid that lets another bidder win with probability 1. Note

⁴ Note that the same argument would work if the highest-value bidder $i = i^*(\theta)$, when he fails to win the object, would buy it back from the auction’s winner at the price equal to the second-highest value $\max_{j \neq i} \theta_j$. This would be the appropriate model if the resale procedure involved Bertrand competition among the auction’s losers to buy the object from the winner.

that bidding 0 may not accomplish this because the seller only needs to sell with probability 1, and so can withhold the object if any bidder bids 0. Yet the argument can be modified to obtain the desired result:

Proposition 1. *If the highest bidder learns all other values and the resale payoffs are given by (6) then the Vickrey auction with no reserve price is optimal among all auctions that sell with probability 1.*

Proof. Take any auction mechanism (χ, ψ) satisfying (3)-(4), and let

$$U_i(\theta_i) = \mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \psi_i(\theta_i, \tilde{\theta}_{-i})] \quad (7)$$

be the interim expected payoff enjoyed by i when his type is θ_i . Since $\sum_j \chi_j(\tilde{\theta}) = 1$ with probability 1, there exists arbitrarily small $\hat{\theta}_i > 0$ such that $\sum_j \chi_j(\hat{\theta}_i, \tilde{\theta}_{-i}) = 1$ with probability 1. By incentive compatibility (3), for all θ_i ,

$$\begin{aligned} U_i(\theta_i) &\geq \mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta) - \psi_i(\hat{\theta}_i, \tilde{\theta}_{-i})] \\ &\geq \mathbb{E}_{\tilde{\theta}_{-i}} \left[\max \left\{ \theta_i - \max_{j \neq i} \tilde{\theta}_j, 0 \right\} \right] - \hat{\theta}_i, \end{aligned}$$

since $v_i(x; \theta) \geq \max \{ \theta_i - \max_{j \neq i} \theta_j, 0 \}$ whenever $\sum_j x_j(\theta) = 1$ and $\mathbb{E}_{\tilde{\theta}_{-i}}[\psi_i(\hat{\theta}_i, \tilde{\theta}_{-i})] \leq \hat{\theta}_i$ by individual rationality (4). Therefore, we can write

$$\begin{aligned} \mathbb{E}_{\tilde{\theta}} \left[\sum_i \psi_i(\tilde{\theta}) \right] &= \mathbb{E}_{\tilde{\theta}} \left[\sum_i v_i(\chi(\tilde{\theta}); \tilde{\theta}) \right] - \sum_i \mathbb{E}_{\tilde{\theta}_i} [U_i(\tilde{\theta}_i)] \\ &\leq \mathbb{E}_{\tilde{\theta}} [\tilde{\theta}_{i^*(\tilde{\theta})}] - \mathbb{E}_{\tilde{\theta}} \sum_i \left[\max \left\{ \tilde{\theta}_i - \max_{j \neq i} \tilde{\theta}_j, 0 \right\} \right] + \sum_i \hat{\theta}_i. \end{aligned}$$

Since the first two terms describe the expected revenue in the Vickrey auction, and $\hat{\theta}_i > 0$ can be chosen to be arbitrarily small, the result obtains. \square

Next, note that the Vickrey auction sustains truthful bidding and no resale as an ex post equilibrium for any resale procedure. Indeed, note when bidder i deviates downward and loses, he cannot buy the object in resale for below the winner's value, $\max_{j \neq i} \theta_j$, which is what he would pay for winning the object in an auction, and conversely when bidder i deviates upward and wins, he cannot resell the object for more than the highest loser's value,

$\max_{j \neq i} \theta_j$, which is the price he would have to pay to win the auction. In both cases he does no better than bidding his true value. (A more general version of this result, inspired by Ausubel and Cramton (2004), is stated in Section 7 below.)

This shows that the Vickrey auction attains the same expected revenue in equilibrium (namely, the expectation of the second-highest value) regardless of the resale procedure; and no other auction can guarantee higher expected revenue, by Proposition 1. Thus, the Vickrey auction solves the seller's maxmin problem: it is a robustly optimal auction with unknown resale.⁵

4 Can-keep case: Two bidders

We now turn to our main focus: the optimal auction when the seller can withhold the object. In this section we mainly study the problem with two bidders, although some of the intermediate ideas we develop are useful for any number of bidders, and will be stated in this generality. In the next section, we fully extend the result to any number of bidders.

As in the previous section, we conjecture that the worst-case resale procedure has the following form: after the auction, nature reveals all the bidders' values to each other; and the highest-value bidder gets all the bargaining power in resale, which yields reduced-form resale payoffs given by (6). So we can set up the mechanism design problem assuming this specific resale procedure, and solve for the optimal mechanism. It will then turn out that, for this mechanism, truth-telling (followed by no resale) is an ex-post equilibrium regardless of the resale procedure. Therefore, the mechanism we derive actually solves the robust optimization problem.

4.1 Analysis: the relaxed problem

We begin by studying the relaxed problem for program (5) subject to (3)–(4), which replaces (3) with the local first-order incentive compatibility constraints. For this purpose, note that by the envelope theorem of Milgrom and Segal (2002, Corollary 1), incentive compatibility (3) implies that U_i is

⁵The argument above can be extended to multi-unit auctions. It can also be extended to auctions with correlated values, provided that we also require robustness to information disclosure *before* the auction. Thus, the proposition above can be stated for all such cases.

absolutely continuous and its derivative is given almost everywhere by

$$U'_i(\theta_i) = \mathbb{E}_{\tilde{\theta}_{-i}}[v'_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})].$$

Here v'_i denotes the derivative of v_i with respect to θ_i , which is defined except when there are ties in values (a probability-zero case), and takes the form $v'_i(x; \theta) = x_i + \mathbf{1}_{i=i^*(\theta)} \cdot \sum_{j \neq i} x_j$.

Since $U_i(0) = 0$ at the optimal mechanism (type 0's participation constraints are binding), the interim expected payment of bidder i given type θ_i is

$$\mathbb{E}_{\tilde{\theta}_{-i}}[\psi_i(\theta_i, \tilde{\theta}_{-i})] = \mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})] - \int_0^{\theta_i} \mathbb{E}_{\tilde{\theta}_{-i}} \left[v'_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \hat{\theta}_i, \tilde{\theta}_{-i}) \right] d\hat{\theta}_i. \quad (8)$$

The usual integration by parts allows us to rewrite the objective (5) as

$$\mathbb{E}_{\tilde{\theta}} \left[\sum_i \psi_i(\tilde{\theta}) \right] = \mathbb{E}_{\tilde{\theta}} \left[\sum_i v_i(\chi(\tilde{\theta}), \tilde{\theta}) - \sum_i \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} v'_i(\chi(\tilde{\theta}), \tilde{\theta}) \right]. \quad (9)$$

The standard next step is to solve a relaxed problem: maximize (9) over all allocation rules χ , without worrying about (3). This can be done by maximizing the virtual surplus (the expression inside brackets) pointwise. The solution must be a deterministic mechanism (with probability 1). At any profile θ , if we allocate to the high-value bidder $i^* = i^*(\theta)$, the virtual surplus is

$$\theta_{i^*} - \frac{1 - F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} = \nu_{i^*}(\theta_{i^*}).$$

If we allocate to some other bidder $j \neq i^*$, then the v'_i terms are nonzero both for $i = i^*$ and for $i = j$, so we get

$$\theta_{i^*} - \frac{1 - F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} - \frac{1 - F_j(\theta_j)}{f_j(\theta_j)}.$$

This is less than the virtual surplus from allocating to i^* . Hence, it is never optimal to allocate inefficiently: we either allocate to the highest-value bidder i^* or not at all. Intuitively, allocating to some non-high-value bidder j at a type profile θ concedes informational rents to higher types of j (who can acquire the good by pretending to be θ_j , and then possibly resell it) as well as for higher types of i^* (who can buy the good back from j); whereas allocating




Figure 1: Allocation rule from relaxed problem. 1 means allocate to bidder 1; 2 means allocate to bidder 2. In the remaining regions, the good is not sold.

to the high-value bidder i^* leaves rents to him only, and so is preferable for the seller. Furthermore, we should allocate to the highest-value bidder i^* if and only if his virtual value is positive, i.e. his value is above r_{i^*} , where $r_i = \nu_i^{-1}(0)$ is the optimal price for a monopolist selling to bidder i only.

The allocation rule is shown in Figure 1 for the case of two bidders. Assume that $r_1 > r_2$. Transfers consistent with (8) can be achieved, for example, using threshold prices: if i^* is allocated the good, he is charged $\max\{r_{i^*}, \max_{j \neq i^*}(\theta_j)\}$; losers pay nothing.

With these payments, the auction would be incentive-compatible if there were no resale. However, with resale, it violates incentive compatibility constraints (3). To see this, consider bidder 1's type $\theta_1 \in (r_2, r_1)$. By telling the truth, bidder 1 gets a payoff of 0, since he never gets the object (and cannot get it from 2 in resale). However, he could report a lower type $\hat{\theta}_1$, in which case 2 wins the object and then 1 can profitably buy it back in

resale if $\theta_2 \in (\hat{\theta}_1, \theta_1)$. (This deviation is illustrated with a horizontal arrow in Figure 1.) Thus, in our model, the solution to the relaxed problem is not incentive-compatible: we must consider non-local incentive constraints to find the correct solution.

4.2 Vickrey auction with reserves

Intuitively, to avoid this incentive to deviate and buy back, we might use an allocation rule in which a lower bid never causes the auctioneer to sell the object. (Ausubel and Cramton (1999) call this property *monotonicity in aggregate*.) For example, we might try to fix the allocation rule by “filling in” the triangular region $r_2 < \theta_2 < \theta_1 < r_1$ in Figure 1, allocating to bidder 1 in this region (based on the above intuition that we prefer to allocate to the high-value bidder or to nobody). Note, however that this solution can be improved: since bidder 1 has a negative virtual value in the filled-in triangle, the seller would rather not sell to him. By raising the reserve price for bidder 2 above r_2 , she can shrink the size of this triangle, although doing so also means missing out on profitable sales to bidder 2. The optimal reserve price trades off these two effects. This leads to an allocation rule of the form shown in Figure 2.

This auction belongs to a class of auctions that we formally define now.

Definition 1 (Ausubel and Cramton 1999, 2004). A Vickrey auction with reserves is the auction mechanism parameterized by reserve prices $p_1, \dots, p_n \in [0, 1]$, defined as follows:

- *Allocation rule:*

$$\chi_i(\theta) = \begin{cases} 1 & \text{if } i = i^*(\theta) \text{ and } \theta_j > p_j \text{ for some } j, \\ 0 & \text{otherwise.} \end{cases}$$

- *Payments:*

$$\psi_i(\theta) = \begin{cases} \max\{p_i, \max_{j \neq i} \theta_j\} & \text{if } \chi_i(\theta) = 1 \text{ and } \theta_j \leq p_j \text{ for all } j \neq i, \\ \max_{j \neq i} \theta_j & \text{if } \chi_i(\theta) = 1 \text{ and } \theta_j > p_j \text{ for some } j \neq i, \\ 0 & \text{otherwise.} \end{cases}$$

In words, if at least one bidder i beats his reserve price p_i , then the good is allocated to the high-value bidder, otherwise, the good is left unsold.




Figure 2: Allocation rule for Vickrey auction with reserves. 1 means allocate to bidder 1; 2 means allocate to bidder 2. In the remaining regions, the good is not sold.

Importantly, since the reserve prices are asymmetric, a bidder i can win the good without meeting his reserve price p_i — if another bidder j with a lower reserve has met his reserve p_j .

The winner’s payments in the auction is constructed to be his “threshold price” – the minimal bid that would have allowed him to win. By the standard argument, this ensures that the auction is strategy-proof without resale. More importantly for us, as noted by Ausubel and Cramton (2004), in Vickrey auctions with reserves it is an *ex post* equilibrium for bidders to bid truthfully even if bidders believe resale will occur, for any beliefs about the (individually rational) resale procedure and for any profile of values. To see that the possibility of resale does not create any advantageous deviations, note that when a deviation causes another bidder to win, the deviator would need to pay at least the winner’s value to buy the object back, but he could have instead acquired the object at this price by submitting a high bid in the auction. On the other hand, if a bidder’s deviation causes *him* to win, he would pay at least the others’ highest value for the object, and would not be able to resell it for a higher price.

It turns out that a Vickrey auction with reserves is the correct solution to our optimization problem. In the remainder of this section we will sketch the argument for the case of two bidders, leaving the general formal result and proof for the next section.⁶

4.3 Optimal Reserve Prices

We now discuss how to identify the optimal reserve prices. Since the object is always allocated to the higher-value bidder, the formula (9), expressing

⁶ The working paper by Ausubel and Cramton (1999) considered a standard auction with the constraint that the good should only be sold to the highest-value bidder or not at all (but without explicitly modeling resale), and derived the solution shown in Figure 1. They also observed that this auction invites deviations if there is resale. Their paper then considered the problem of the optimal auction subject to monotonicity in aggregate, as well as selling only to the highest bidder, and claimed (without proof) the Vickrey auction with reserves as the solution in a special case with two bidders. Here, in contrast, we *derive* the properties of monotonicity in aggregate and highest-bidder-only from the maxmin problem of a seller who is concerned with revenue, rather than assuming those properties. We also prove that this auction format is optimal, both with two bidders and more generally.

revenue as the integral of virtual surplus, can be simplified to

$$\mathbb{E}_{\tilde{\theta}} \left[\mathbf{1} \left\{ \tilde{\theta}_i \geq p_i \text{ for some } i \right\} \cdot \nu_{i^*(\tilde{\theta})}(\tilde{\theta}_{i^*(\tilde{\theta})}) \right]. \quad (10)$$

Assume the optimal reserves satisfy $p_1 > p_2$ (this will turn out to be true, given our assumption that $r_1 > r_2$). The effect of bidder 1's reserve price p_1 on the expected revenue occurs only when $\theta_2 < p_2$. Conditional on θ_2 in this range, the expected revenue is given by $R_1(p_1) \equiv p_1(1 - F_1(p_1))$ (the expected revenue on bidder 1 as a function of the price charged to him). By assumption, this is maximized by setting $p_1 = r_1$.

As for the optimal reserve price for bidder 2, consider the effects of raising bidder 2's reserve price p_2 slightly to $p_2 + \varepsilon$. This change would have two first-order effects, both in the sliver $\theta_2 \in (p_2, p_2 + \varepsilon)$, the probability of which is $f_2(p_2)\varepsilon$. The first effect is to increase the expected revenue on bidder 1 from $R_1(p_2)$ to its maximal value $R_1(r_1)$. The second effect is to reduce the expected revenue on bidder 2 when $\theta_1 < p_2$ (thus with probability $F_1(p_2)$), by not selling to him when $\theta_2 \in (p_2, p_2 + \varepsilon)$, by an amount equal to his virtual value $\nu_2(p_2)$. The first-order condition for p_2 equalizes these two effects:

$$F_1(p_2)\nu_2(p_2) = R_1(r_1) - R_1(p_2). \quad (11)$$

Note that (11) has a solution $p_2 \in [r_2, r_1]$ provided that ν_2 is nondecreasing. Indeed, at $p_2 = r_2$, the left-hand side is zero and the right-hand side is nonnegative, while at $p_2 = r_1$, the left-hand side is nonnegative and the right-hand side is zero. Hence, existence obtains by the Intermediate Value Theorem. We can also see that the solution is unique when ν_2 is strictly increasing (which makes the left-hand side strictly increasing) and R_1 is concave (which makes the right-hand side nonincreasing on $[r_2, r_1]$).

4.4 Optimality of Vickrey with Reserves

To show that Vickrey with reserves is optimal, we need to make active use of the non-local incentive constraints, since using only the local incentive constraints gave us the incorrect solution in Figure 1. For the two-bidder case we will only need to use the non-local incentive constraints (3) for bidder 1. Using the formula (8) for transfers, the constraints can be rewritten entirely

in terms of the allocation rule:

$$\begin{aligned}
& S(\theta_1, \hat{\theta}_1; \chi) \\
\equiv & \int_{\hat{\theta}_1}^{\theta_1} \mathbb{E}_{\tilde{\theta}_2} [v_1'(\chi(\tau_1, \tilde{\theta}_2); \tau_1, \tilde{\theta}_2)] d\tau_1 - \mathbb{E}_{\tilde{\theta}_2} [v_1(\chi(\hat{\theta}_1, \tilde{\theta}_2); \theta_1, \tilde{\theta}_2) - v_1(\chi(\hat{\theta}_1, \tilde{\theta}_2); \hat{\theta}_1, \tilde{\theta}_2)] \geq 0.
\end{aligned} \tag{12}$$

We account for these constraints using a Lagrangian approach, by adding extra terms to the objective function (9) to penalize violations of the constraints. Since there is a continuum of such constraints, the Lagrange multipliers (weights) on the constraints must be described with some appropriately constructed nonnegative measure \mathcal{M} on $[0, 1] \times [0, 1]$. Our Lagrangian then takes the form

$$\mathbb{E}_{\tilde{\theta}} \left[\sum_i v_i(\chi(\tilde{\theta}), \tilde{\theta}) - \sum_i \frac{1 - F_i}{f_i(\theta_i)} v_i'(\chi(\tilde{\theta}), \tilde{\theta}) \right] + \iint S(\theta_1, \hat{\theta}_1; \chi) d\mathcal{M}(\theta_1, \hat{\theta}_1). \tag{13}$$

Because this Lagrangian is a bounded linear functional of the allocation rule χ , by the Riesz Representation Theorem (e.g., Royden 1988) it can be written in the form

$$\mathbb{E}_{\tilde{\theta}} \left[\sum_i \mu_i(\tilde{\theta}) \chi_i(\tilde{\theta}) \right]. \tag{14}$$

We will refer to $\mu_i(\theta)$ as the “modified virtual value” of bidder i . It combines the ordinary virtual value and the terms coming from (12). Of course, $\mu_i(\theta)$ depends on the choice of measure \mathcal{M} . (For the types of measures proposed below, an explicit formula for modified virtual values appears in display (34) in the appendix.)

Now, we show optimality of allocation rule χ in the candidate solution by constructing a measure \mathcal{M} such that (χ, \mathcal{M}) is a saddle point of the Lagrangian (13) (see Luenberger (1969), Theorem 2 on p. 221), i.e., the following conditions hold:

- (a) The allocation rule χ maximizes the Lagrangian given measure \mathcal{M} ;
- (b) The candidate solution with allocation rule χ satisfies all incentive compatibility constraints;
- (c) *Complementary slackness*: \mathcal{M} puts zero measure on all constraints (12) that are slack (hold with strict inequality) at χ , i.e., the second term in the Lagrangian (13) is zero. (Together with (b), this ensures that given χ , (13) is minimized by measure \mathcal{M} .)

To see the sufficiency of conditions (a)-(c) for optimality of the candidate solution, note that the revenue from any alternative incentive-compatible auction with allocation rule χ' will be at most the value of the Lagrangian (13) at (χ', \mathcal{M}) , which by (a) is at most the value of the Lagrangian at (χ, \mathcal{M}) , which by (c) equals the revenue at χ .

Regarding condition (b), the feasibility of Vickrey auctions with reserves was already argued informally, and will be stated formally in Section 7 below.

Regarding condition (c), note that the constraints (12) that hold with equality at the candidate solution are those with $\hat{\theta}_1 \leq \theta_1$ and either $\theta_1, \hat{\theta}_1 \leq p_1$ or $\theta_1, \hat{\theta}_1 \geq p_1$. (The constraints involving upward deviations are slack because they carry the risk of getting the good at a price above value, and the constraints from above p_1 to below p_1 are slack because those deviations may cause the good to be unsold.) To ensure complementary slackness, the support of \mathcal{M} must be restricted to those constraints. But we can restrict the support of \mathcal{M} further by observing that our Vickrey auction differs from the relaxed allocation rule in two ways: it sells to bidder 1 more often when $\theta_1 \in (p_2, p_1)$, and it sells to bidder 2 less often when $\theta_1 < p_2$. This suggests that we should focus on the binding constraints with $\theta_1 \in (p_2, p_1)$ and $\hat{\theta}_1 < p_2$. Thus we will look for a measure \mathcal{M} with support $[p_2, p_1] \times [0, p_2]$ (and any such measure will automatically satisfy complementary slackness).

We construct a measure $\mathcal{M}(\theta_1, \hat{\theta}_1)$ that takes the form of a product of a marginal measure over θ_1 and another one over $\hat{\theta}_1$. We denote the two measures' distribution functions by $\Lambda(\theta_1)$ and $\hat{\Lambda}(\hat{\theta}_1)$, respectively. The two measures will be nailed down by two indifference conditions in maximizing the Lagrangian (13):

1. Indifference between selling to bidder 2 and withholding the good when $\theta_1 < \theta_2 = p_2$, which dictates that bidder 2's modified virtual value at these profiles should be 0;
2. Indifference between selling to bidder 1 and withholding the good when $\theta_2 = p_2 < \theta_1$, which dictates that bidder 1's modified virtual value at these profiles should be 0.

To see the necessity of these indifferences, note that the indifference should be broken in either direction with an arbitrarily small change in θ_2 .

For condition #1, note that selling to bidder 2 when $\theta_1 < \theta_2 = p_2$ affects the Lagrangian in two ways: adding the virtual value of agent 2, weighted by

the density of $\theta = (\theta_1, p_2)$, to the first part of (13), and inducing information rents $\tau_1 - p_2$ for all types $\tau_1 > p_2$ of bidder 1, integrated over τ_1 with measure Λ and weighted by $\hat{\Lambda}'(\theta_1)$ (the density of $\hat{\Lambda}$), to be subtracted from the second part of (13). Thus, indifference condition #1 takes the form

$$f_1(\theta_1) \nu_2(p_2) = \hat{\Lambda}'(\theta_1) \int_{p_2}^{p_1} (\tau_1 - p_2) d\Lambda(\tau_1) \text{ for all } \theta_1 \in [0, p_2]. \quad (15)$$

Note, in particular, that the condition requires the density of $\hat{\Lambda}(\theta_1)$ to be proportional to the density $f_1(\theta_1)$. Intuitively, this proportionality is necessary to ensure that the sale's effect on tightening bidder 1's non-local incentive constraints binding to θ_1 (whose weight is proportional to the density of $\hat{\Lambda}(\theta_1)$) exactly offset the effect of adding bidder 2's virtual value (whose weight is proportional to the density $f_1(\theta_1)$ of bidder 1's type distribution). By rescaling Λ and $\hat{\Lambda}$ by constants, without loss of generality we can let the proportionality factor be 1, so $\hat{\Lambda}(\hat{\theta}_1) = F_1(\hat{\theta}_1)$.

Now we turn to indifference condition #2. One way of understanding this condition is to condition on $\theta_2 = p_2$, so that the allocation is a function of θ_1 only. A seller trying to maximize this conditional Lagrangian should then be indifferent across many such allocation rules. In particular, setting any reserve price $p \in [p_2, p_1]$ creates an alternative allocation rule, which allocates to 1 when $\theta_1 > p$, and does not allocate at all when $\theta_1 \leq p$. The conditional Lagrangian maximizer should be indifferent between any such allocation rule and the actual allocation rule used in the Vickrey auction. Note that this observation depends on indifference condition #1, to ensure the seller is willing to withhold the good when $\theta_1 < p_2$.

Now, imagine the conditional-Lagrangian-maximizing seller starting from a reserve of p_1 and changing to any other reserve $p \in [p_2, p_1]$. The indifference means that the loss in the expected profit on bidder 1 (represented by the first term of the conditional Lagrangian), $R_1(p_1) - R_1(p)$, should be exactly offset by the newly created weighted slack of the non-local incentive constraints (the second term of the conditional Lagrangian). The slack incentive constraints are now those from $\theta_1 \in [p, p_1]$ to $\hat{\theta}_1 < p$: indeed, any such incentive constraint is now slack by $\theta_1 - p$, since the deviation would cause the good to be unsold. Integrating over those constraints with measures Λ and $\hat{\Lambda}$ yields the equation

$$\hat{\Lambda}(p) \int_p^{p_1} (\theta_1 - p) d\Lambda(\theta_1) = R_1(p_1) - R_1(p). \quad (16)$$

Recall from the above that $\hat{\Lambda}$ is supported on $[0, p_2]$, and $\hat{\Lambda}(p) = \hat{\Lambda}(p_2) = F_1(p_2)$ for $p > p_2$. Differentiating both sides of (16) yields

$$-F_1(p_2) \int_p^{p_1} d\Lambda(\theta_1) = -R_1'(p).$$

Taking into account that the integral in the above display equals $\Lambda(p_1) - \Lambda(p)$, we obtain

$$\Lambda(p_1) - \Lambda(\theta_1) = R_1'(\theta_1)/F_1(p_2) \text{ for } \theta_1 \in [p_2, p_1].$$

This equation describes the measure Λ . Note that Λ is nondecreasing, and so the constructed weight measure really is nonnegative, when function R_1 is concave.

Also, plugging in this Λ into (15) and integrating by parts yields the first-order condition (11) for p_2 .

To complete the proof of optimality of the Vickrey auction with reserves, it remains to check condition (a) – that the Lagrangian is maximized by the candidate solution. For this, we can use the Lagrangian expressed by means of modified virtual values, (14). We need to check that the proposed Vickrey auction with reserves maximizes the modified virtual value for every profile (θ_1, θ_2) , and not just when $\theta_2 = p_2$ and $\theta_1 \leq p_1$, which we have already checked. This is some work but is reasonably straightforward.

5 Can-Keep Case with Many Bidders

We now proceed to the general setting with n bidders. In this section we make the following assumptions on bidders' value distributions (in addition to those made in Section 2):

A1. Both F_i and $1 - F_i$ are log-concave for each i .

A2. $\nu_1(\theta) \leq \dots \leq \nu_n(\theta)$ for all θ .

Note that the latter part of (A1) is the usual monotone hazard rate assumption. Both parts of (A1) are satisfied, in particular, when the density f_i is log-concave, which is satisfied by many standard distributions (Bagnoli and Bergstrom 2005). Note also that (A1) is not nested with the assumption used in the previous section that the revenue function $R_i(\theta) = \theta(1 - F_i(\theta))$

is concave. This leads us to believe that (A1) is not the weakest possible assumption, but it yields the result in a simple way.

Assumption (A2) means that the bidders are uniformly ranked from “stronger” to “weaker”; for example, it ensures that in the absence of resale, the optimal auction would always discriminate against the stronger bidders (McAfee and McMillan 1989), and the optimal bidder-specific reserve prices would satisfy $r_1 \geq \dots \geq r_n$.

We begin with a characterization of the candidate optimal auction, which is a Vickrey auction with reserve prices. We characterize the optimal reserve prices p_1, \dots, p_n . Then we establish the optimality of this auction, not just among Vickrey auctions with reserves, but among *all* possible auctions.

5.1 Characterization of Candidate Optimal Auction

As a first step toward defining the optimal reserve prices, define $R_k(p)$, for any $k = 1, \dots, n$ and any price p , to be the expected revenue from a Vickrey auction in which only the first k bidders participate and all participants face the same reserve price p . Also define $R_0(p) = 0$.

Now, recursively define a weakly decreasing sequence of reserve prices p_k , and a sequence of revenue levels R_k^* for $k = 1, \dots, n$, by initializing $p_0 = 1$ and $R_0^* = 0$ and letting for all $k \geq 1$,

$$R_k^* = R_k(p_k) + F_k(p_k) [R_{k-1}^* - R_{k-1}(p_k)] \quad (17)$$

$$= \max_{p \in [0, p_{k-1}]} \{ R_k(p) + F_k(p) [R_{k-1}^* - R_{k-1}(p)] \}. \quad (18)$$

Inductively, R_k^* for $k \geq 1$ is the revenue from a Vickrey auction on bidders $1, \dots, k$ with reserves p_1, \dots, p_k . To see that it satisfies formula (17), compare this asymmetric Vickrey auction to the symmetric Vickrey with reserve p_k . Notice that the two auctions differ only when bidder k 's value is below p_k , which happens with probability $F_k(p_k)$. In this case, the former auction reduces to an auction on the first $k - 1$ bidders with reserves p_1, \dots, p_{k-1} , yielding expected revenue R_{k-1}^* ; and the latter is an auction with symmetric reserve p_k on the first $k - 1$ bidders, yielding revenue $R_{k-1}(p_k)$. (18) requires that $p_k \in [0, p_{k-1}]$ is chosen to maximize expected revenues.

First we make the following observation (proven in the Appendix):

Lemma 1. $p_k > 0$ for all $k = 1, \dots, n$.

If the maximization problem (18) has an interior solution p_k , it has to satisfy the first-order condition

$$\nu_k(p_k) \prod_{j < k} F_j(p_k) = R_{k-1}^* - R_{k-1}(p_k).$$

Intuitively, increasing p_k by ε only matters when $\theta_k \in (p_k, p_k + \varepsilon)$ and $\theta_j < p_j$ for all $j > k$, and in that case it has two first-order effects: (i) on the expected revenue from bidder k , when all other bidders' values are below p_k , and (ii) on the expected revenue from the $k - 1$ strongest bidders, changing it from $R_{k-1}(p_k)$ (obtained when bidder k beats his reserve price) to R_{k-1}^* (obtained when bidder k does not beat his reserve price). The first-order condition equalizes those two effects.

The first-order condition can be rewritten in the form

$$\nu_k(p_k) = H(p_k), \tag{19}$$

where the function H is defined on $p \in [p_n, 1)$ by

$$H(p) = \frac{R_{k-1}^* - R_{k-1}(p)}{\prod_{j < k} F_j(p)} \text{ for } p \in [p_k, p_{k-1}), k = 1, \dots, n. \tag{20}$$

Furthermore, the following lemma (proven in the appendix) shows that the first-order condition (19) has to hold even if (18) has a corner solution, and yields an approach to constructing the function H and the reserve prices:

Lemma 2. *Under assumptions (A1)-(A2), the function H defined by (20) has the following properties:*

- (i) H is continuous,
- (ii) H satisfies (19) for each $k \geq 1$,
- (iii) For $p \in [p_k, p_{k-1}]$, for each $k \geq 2$,

$$H(p) = \frac{1}{\prod_{j < k} F_j(p)} \left[(\prod_{j < k} F_j(p_{k-1})) \nu_{k-1}(p_{k-1}) - \int_p^{p_{k-1}} (\prod_{j < k} F_j(\tau)) \sum_{j < k} \frac{f_j(\tau) \nu_j(\tau)}{F_j(\tau)} d\tau \right], \tag{21}$$

- (iv) H is nonincreasing.

Formula (21) can be used to construct the sequence p_1, \dots, p_n as follows: First let $p_1 = r_1$, since this value solves (18) for $k = 1$. Then we construct p_k for $k \geq 2$ iteratively as follows: Given p_{k-1} , (21) describes $H(p)$ on the interval $p \in [p_k, p_{k-1}]$, and then p_k can be obtained by solving equation (19). We can see that (19) has a solution $p_k \in [r_k, p_{k-1}]$, by noting that on the one hand, $\nu_k(r_k) = 0 = H(r_1) \leq H(r_k)$ (using part (iv) of the lemma), and on the other hand, $\nu_k(p_{k-1}) \leq \nu_{k-1}(p_{k-1}) = H(p_{k-1})$ (using (A2) and (19) for $k-1$). Thus, a solution to (19) exists by the Intermediate Value Theorem (using continuity of H and ν_k). Furthermore, the solution to (19) is unique since ν_k is strictly increasing due to assumption (A1) and H is nonincreasing by part (iv) of the lemma. In Section 6 below we use this approach to calculate optimal reserve prices when bidders' value distributions are uniform with different supports.

5.2 Optimality of the Auction

The maximization in (18) suggests that p_k is the optimal reserve for bidder k , taking as given the reserves p_1, \dots, p_{k-1} of the stronger bidders and given the constraint $p_k \leq p_{k-1}$. Now we show that the constraint does not bind and this is indeed the optimal auction — and not just among Vickrey auctions with reserves, but among *all* possible auctions:

Theorem 1. *Under assumptions (A1)-(A2) and the resale procedure described in (6), the Vickrey auction with reserves p_1, \dots, p_n characterized above is an optimal auction for the seller (and the optimal revenue is R_n^*).⁷*

A full proof of the theorem is given in the appendix, but here we describe the key steps of the proof and develop some intuition. Just as in the two-bidder case, the theorem is proven by constructing Lagrange multipliers on non-local incentive constraints in such a way that the saddle-point conditions (a)-(c) listed in Section 4 are satisfied. With n bidders, we use Lagrange multipliers $\mathcal{M}_i(\theta_i, \hat{\theta}_i)$ for the incentive constraints of bidders $i = 1, \dots, n-1$. Similarly to (12), the incentive constraints can be written as

$$\begin{aligned} & S_i(\theta_i, \hat{\theta}_i; \chi) \\ \equiv & \int_{\hat{\theta}_i}^{\theta_i} \mathbb{E}_{\tilde{\theta}_{-i}}[v'_i(\chi(\tau_i, \tilde{\theta}_{-i}); \tau_i, \tilde{\theta}_{-i})] d\tau_i - \mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \hat{\theta}_i, \tilde{\theta}_{-i})] \geq 0, \end{aligned}$$

⁷This auction is also optimal under the alternative resale assumption in footnote 4.

and the Lagrangian takes the form

$$\mathbb{E}_{\tilde{\theta}} \left[\sum_i v_i(\chi(\tilde{\theta}), \tilde{\theta}) - \sum_i \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} v'_i(\chi(\tilde{\theta}), \tilde{\theta}) \right] + \sum_{i=1}^{n-1} \iint S_i(\theta_i, \hat{\theta}_i; \chi) d\mathcal{M}_i(\theta_i, \hat{\theta}_i). \quad (22)$$

An important difference from the two-bidder case is that it is no longer possible to restrict attention to multipliers that take a product form $\Lambda_i(\theta_i) \times \hat{\Lambda}_i(\hat{\theta}_i)$. The reason is that for each bidder $i < n$ there will be a range of types θ_i where the non-local incentive constraints *both* for higher types to imitate θ_i , and for θ_i to imitate lower types, are binding. This means that in a product form, the supports of the marginal distributions of θ_i and $\hat{\theta}_i$ would have to overlap, meaning that some upward incentive constraints are binding, which is inconsistent with complementary slackness because all upward non-local incentive constraints are slack in our candidate optimal auction.

It turns out that the correct weights have just a slightly more general form: a product measure $\Lambda_i \times \hat{\Lambda}_i$ on $[0, 1] \times [0, 1]$, *restricted* to the half-plane $\hat{\theta}_i \leq \theta_i$. The support of Λ_i will be $[p_n, p_i]$ and the support of $\hat{\Lambda}_i$ will be $[0, p_i]$. As before, we will use Λ_i and $\hat{\Lambda}_i$ to denote the distribution functions of the respective measures. (The weights constructed here would also work for bidder $i = 1$ in two-bidder case, establishing the result under assumptions (A1)-(A2) instead of concavity of R_1 .)

Now we describe informally the construction of the distributions Λ_i and $\hat{\Lambda}_i$, using modified versions of the arguments developed for the two-bidder case. First, we argue that the distribution $\hat{\Lambda}_i$ should be proportional to the bidder's value distribution F_i restricted to $[0, p_i]$, and so without loss we normalize $\hat{\Lambda}_i = F_i$. Just as in the two-bidder case, proportionality is necessary in order to exactly balance the effect of bidder i 's non-local incentive constraints into all $\theta_i < p_i$ on allocating the object to other bidders against those bidders' virtual values, so that the other bidders' allocation in the region where bidder i does not win and does not beat his reserve price is not contingent on θ_i .

Next, we can derive the distributions Λ_i by a calculation that is similar to that used in the two-bidder case, but a bit trickier. Namely, to derive $\Lambda_i(\theta_i)$, where $\theta_i \in (p_k, p_{k-1})$ for some $k > i$, we use the observation that the maximizer of the Lagrangian (22) should be indifferent between allocating to the highest-value bidder and not allocating at all at any type profile $(\theta_1, \dots, \theta_n)$ at which $\theta_j = p_j$ for bidders $j \geq k$ and $\theta_j \in (p_n, p_j)$ for bidders $j < k$. (This indifference must hold, because if any bidder $j < k$ has his bid

perturbed slightly upward, the candidate optimum must allocate to the high bidder; if slightly downward, the good must be left unsold.) In particular, this implies that once we condition on $\theta_j = p_j$ for $j \geq k$, the conditional Lagrangian maximizer is indifferent to changing the reserve prices for the first $k - 1$ bidders, as long as the new reserve price for each bidder j lies in the interval $[p_n, p_j]$.

Specifically, we consider setting a reserve $q \in (p_k, p_{k-1})$ for bidder i and a reserve $p \in [q, p_{k-1})$ for all other bidders $j < k$, $j \neq i$. Let $\bar{R}_{k-1}^i(p, q)$ denote the expected revenue from the resulting Vickrey auction for the $k - 1$ strongest bidders. The Lagrangian indifference means that the loss in the expected revenue – the first term of the Lagrangian (22), $R_{k-1}^* - \bar{R}_{k-1}^i(p, q)$, should be exactly offset by the second term in (22), which is the weighted slack created in agents’ non-local incentive constraints.

The created slack in the non-local incentive constraints of a bidder $j < k$ involves constraints from every θ_j above his new reserve price to every $\hat{\theta}_j$ below his new reserve price when bidding θ_j would give the object to this bidder while bidding $\hat{\theta}_j$ would cause it to be unsold. For bidder i , the slack is the bidder’s expected payoff loss from a deviation from $\theta_i > q$ to any $\hat{\theta}_i < q$. This payoff loss arises only when the highest bid of the other $k - 1$ strongest bidders, $\tilde{\theta}^* \equiv \max_{j < k, j \neq i} \theta_j$, is below both θ_i and p , so bidder i ’s bid of θ_i makes him win while his bid of $\hat{\theta}_i < q$ causes the good to be unsold. This payoff loss can be written by accounting both for states in which bidding θ_i makes him win and pay q and for states in which bidding θ_i makes him win and pay $\tilde{\theta}^*$:

$$s_i(\theta_i, q, p) = F^*(q) \cdot (\theta_i - q) + \mathbb{E}_{\tilde{\theta}^*} \left[\mathbf{1}_{q < \tilde{\theta}^* < \min\{p, \theta_i\}} \cdot (\theta_i - \tilde{\theta}^*) \right],$$

where F^* denotes the c.d.f. of $\tilde{\theta}^*$.

For bidders $j \neq i$, $j < k$, we will not derive the slack exactly, but note that it occurs only when $\theta_i < q$, and in this event it is independent of q , hence the weighted expected slack takes the form $F_i(q) Q_j(p)$ for some function Q_j .

Adding up yields the equation

$$\hat{\Lambda}_i(q) \int_0^{p_i} s_i(\theta_i, q, p) d\Lambda_i(\theta_i) + \sum_{j < k, j \neq i} F_i(q) Q_j(p) = R_{k-1}^* - \bar{R}_{k-1}^i(p, q). \quad (23)$$

Now we divide both sides of (23) by $\hat{\Lambda}_i(q) = F_i(q)$ and differentiate with

respect to q at $q = p$. On the left-hand side of (23) we obtain

$$\int_0^{p_i} \frac{\partial s_i(\theta_i, q, p)}{\partial q} d\Lambda_i(\theta_i) \Big|_{q=p} = -F^*(p) \int_p^{p_i} d\Lambda_i(\theta_i) = -F^*(p) [\Lambda_i(p_i) - \Lambda_i(p)].$$

On the right-hand side of (23) we obtain

$$\frac{\partial}{\partial q} \left(\frac{R_{k-1}^* - \bar{R}_{k-1}^i(p, q)}{F_i(q)} \right) \Big|_{q=p} = -\frac{f_i(p) (R_{k-1}^* - \bar{R}_{k-1}^i(p, p))}{F_i^2(p)} - \frac{(\partial \bar{R}_{k-1}^i(p, q) / \partial q) \Big|_{q=p}}{F_i(p)}.$$

Now note that $\bar{R}_{k-1}^i(p, p) = R_{k-1}(p)$ and that $\partial \bar{R}_{k-1}^i(p, q) / \partial q \Big|_{q=p} = -F^*(p) f_i(p) \nu_i(p)$ (this is the first-order effect on the expected virtual surplus of agent i , while the effect on the expected virtual surplus of the other agents is second-order since it only occurs when $q < \tilde{\theta}_i < \tilde{\theta}^* < p$, the probability of which is $O((p - q)^2)$). Thus we obtain the equation

$$F^*(p) [\Lambda_i(p_i) - \Lambda_i(p)] = \frac{f_i(p)}{F_i(p)} \left[\frac{R_{k-1}^* - R_{k-1}(p)}{F_i(p)} - F^*(p) \nu_i(p) \right].$$

Dividing both sides by $F^*(p) = \prod_{j < k, j \neq i} F_j(p)$ and using (20) yields

$$\Lambda_i(p_i) - \Lambda_i(p) = \frac{f_i(p)}{F_i(p)} [H(p) - \nu_i(p)]. \quad (24)$$

Intuitively, the reason function H appears in the construction of Λ_i is that, as suggested by the first-order condition (19) for optimal reserve pricing, $H(p)$ describes the total shadow weight of all bidders' non-local incentive constraints involving deviations to just below p and then buying from the winner, which would be tightened if we sell to a bidder i when his value is p . One way to interpret (24) is by writing it as

$$\nu_i(p) - H(p) + \frac{F_i(p)}{f_i(p)} [\Lambda_i(p_i) - \Lambda_i(p)] = 0.$$

In the proof, we show that the left-hand-side expression is the modified virtual value of selling to agent i (recall equation (14)) when his value is p and he is the highest bidder. The first term is the usual virtual value of agent i . The second term, $H(p)$, is the total shadow weight of all agents' non-local incentive constraints that would be tightened by selling to bidder i when

his value is p . The third term is the shadow weight of agent i 's non-local incentive constraint from values above p into values below p , which are all relaxed by selling to agent i when his value is p (since this raises the utility of all types above p without changing the utility of any types below p). This shadow weight is therefore the measure of the event $\hat{\theta}_i < p < \theta_i$, which is $\hat{\Lambda}_i(p) [\Lambda_i(p_i) - \Lambda_i(p)]$ (recall that $\hat{\Lambda}_i(p) = F_i(p)$), divided by $f_i(p)$ so that upon taking take expectation over bidder i 's values we obtain the integral with the Lagrange multiplier measure. The equation says that the modified virtual value should be zero, so that the sale to agent i can be conditioned on whether another bidder has met his reserve price.

We can show that (24) indeed describes a nonnegative measure:

Lemma 3. *Under assumptions (A1)-(A2), the function Λ_i defined by*

$$\Lambda_i(\theta_i) = \begin{cases} 0 & \text{if } \theta_i < p_n, \\ \frac{f_i(p_n)}{F_i(p_n)}[H(p_n) - \nu_i(p_n)] - \frac{f_i(\theta_i)}{F_i(\theta_i)}[H(\theta_i) - \nu_i(\theta_i)] & \text{if } \theta_i \in [p_n, p_i], \\ \frac{f_i(p_n)}{F_i(p_n)}[H(p_n) - \nu_i(p_n)] & \text{if } \theta_i > p_i \end{cases}$$

is nondecreasing.

Proof. Consider $\theta_i \in [p_n, p_i]$. Observe that ν_i is nondecreasing by log-concavity of $1 - F_i$ and H is nonincreasing by Lemma 2(iv), hence $H(\theta_i) - \nu_i(\theta_i)$ is nonincreasing and in particular $H(\theta_i) - \nu_i(\theta_i) \geq H(p_i) - \nu_i(p_i) = 0$. Also, log-concavity of F_i means that $\frac{f_i(\theta_i)}{F_i(\theta_i)}$ is nonincreasing, and putting together, we see that $\Lambda_i(\theta_i)$ is nondecreasing on $\theta_i \in [p_n, p_i]$. Finally, by construction $\Lambda_i(\theta_i) = \Lambda_i(p_n)$ for $\theta_i < p_n$ and $\Lambda_i(\theta_i) = \Lambda_i(p_i)$ for $\theta_i > p_i$. \square

The last step of the proof of Theorem 1, just as in the two-bidder case, is to show that with the constructed Lagrange multipliers, our Vickrey auction with reserves maximizes the modified virtual value (recall (14)) at *every* type profile θ : that is, that the modified virtual value of a non-highest-value bidder is always lower than that of the highest-value bidder, and that when every bidder's value is below his reserve, all bidders' modified virtual values are negative.

6 Example: Uniform Distributions

In this section we apply the results to the setting in which the value distribution of each bidder i is uniform on $[0, a_i]$, with different upper limits a_i . At

first glance the analysis is inapplicable to this setting because it has assumed that all distributions have the same support. However, the analysis can be extended to supports with different upper limits, by formally defining bidder i 's virtual value for $\theta_i > a_i$ to be $\nu_i(\theta_i) = \theta_i$ to avoid the $0/0$ division in that region while ensuring continuity. Note that with different upper limits, assumption (A2) requires that $a_1 \geq \dots \geq a_n$.

With different upper limits it is possible that the reserve price p_k of bidder k satisfies $p_k \geq a_k$. Note that in this case by Lemma 2(ii) we have

$$H(p_k) = \nu_k(p_k) = p_k = \nu_j(p_k) \text{ for all } j \geq k,$$

and therefore $p_j = p_k \geq a_j$ for all $j \geq k$. Therefore, in this case bidder k as well as all the weaker bidders are excluded from the auction.

For the case of uniform distributions, formula (21) can be calculated as

$$H_k(p) = \left(\frac{p_{k-1}}{p}\right)^{k-1} \left[\frac{2}{k} p_{k-1} + \frac{1}{k-1} \sum_{j < k} a_j - a_{k-1} \right] + \frac{2(k-1)}{k} p - \frac{1}{k-1} \sum_{j < k} a_j$$

for $p \in [p_k, p_{k-1}]$.

Then p_k is given by equation (19). If bidder k is not excluded from the auction, then $\nu_k(p_k) = 2p_k - a_k$ and (19) can be expressed as

$$\left(\frac{p_k}{p_{k-1}}\right)^{k-1} \left[\frac{2}{k} p_k + \frac{1}{k-1} \sum_{j < k} a_j - a_k \right] = \frac{2}{k} p_{k-1} + \frac{1}{k-1} \sum_{j < k} a_j - a_{k-1}. \quad (25)$$

This is a k th-degree equation for p_k , so for $k \geq 5$ it needs to be solved numerically. To check whether bidder k is excluded, it suffices to evaluate the left-hand side of (25) at $p_k = a_k$: since it is an increasing function of p_k , if at $p_k = a_k$ it exceeds the right-hand side of (25) then the equation's solution is below a_k and bidder k is not excluded, otherwise bidder k is excluded (and so are all the weaker bidders).

For example, for two bidders, we have $p_1 = a_1/2$ and equation (25) for $k = 2$ takes the form

$$p_2 [p_2 + a_1 - a_2] = p_1^2 = a_1^2/4.$$

Bidder 2 is excluded if $a_2 a_1 \leq a_1^2/4$, or $a_2 \leq a_1/4$. Otherwise, the optimal p_2 is given by the positive root of the above quadratic equation:

$$p_2 = \frac{1}{2} \left[\sqrt{a_1^2 + (a_1 - a_2)^2} - (a_1 - a_2) \right].$$

7 Resale-Proofness of Vickrey with Reserves

We now wrap up by giving the complementary result to Theorem 1, showing that any Vickrey auction with reserves is robust to resale: truthful bidding is an ex-post equilibrium under any resale procedure. Together with Theorem 1, this means that a Vickrey auction with reserves yields the best possible revenue guarantee that is robust to the resale procedure.

The analysis of this section largely follows Ausubel and Cramton (2004), who also considered a more general setting of multi-unit auctions with interdependent values. We offer this material in part for the sake of completeness, and in part to be clear about the general class of resale payoffs that we model.

Specifically, we allow that bidders' values may not be publicly revealed before resale takes place (relaxing the implicit assumption of full-information resale made in Section 2). In that case, the continuation equilibrium of the resale game may depend on bidders' beliefs about each other's values, which in turn are affected by their reports $\hat{\theta}_1, \dots, \hat{\theta}_n$ in the auction. Therefore, we describe the payoff that bidder i receives in post-auction bargaining by a function

$$v_i(x; \hat{\theta}_1, \dots, \hat{\theta}_n; \theta_1, \dots, \theta_n),$$

where $\theta_1, \dots, \theta_n$ are the bidders' true values, $\hat{\theta}_1, \dots, \hat{\theta}_n$ are the values that were reported in the auction, and x was the allocation chosen by the auction.⁸ This formalism lets us distinguish the dependence of i 's payoff on $\hat{\theta}_j$

⁸Note that instead of explicitly modeling the post-auction bargaining game, we continue using reduced-form payoffs to represent what would happen in resale. The ideal modeling approach would be to have post-auction bargaining explicitly modeled as a noncooperative game. The ideal theorem would then say that for any such game, the composite game, consisting of the auction followed by bargaining, has an equilibrium where the bidders are always truthful. However, this theorem will only be true if we impose some individual rationality restrictions on bargaining off the equilibrium path (i.e. after some bidder has misreported his value). One way to impose such a restriction would be to specify that in the bargaining game, each bidder should have access to a “non-participation” strategy. However, then our robustness requirement is too weak: *any* auction that is incentive-compatible without resale — in particular the Myerson (1981) optimal auction — is “robust to resale” in this sense, because there is an equilibrium in which every agent always bids truthfully in the auction and then never participates in resale.

A way to strengthen the requirement would be to impose some kind of perfection or sequential rationality requirement on the resale procedure. However, we then need to limit the resale procedures to ensure that such an equilibrium always exists, and we run into thorny problems of guaranteeing perfect equilibrium existence in games with infinite type and action spaces (see Myerson and Reny, 2015).

(which can influence the bidders' beliefs about j 's type at the time bargaining begins⁹) from the dependence on θ_j (which can influence how j actually behaves in bargaining). We will show that Vickrey auctions with reserve prices are robust even if post-auction bargaining can take place under asymmetric information. Relatedly, we also allow a player j with value θ_j to behave differently in the bargaining continuation after he has deviated than after he has truthfully reported his value.

So a *resale procedure* is described by a profile of functions $v_1, \dots, v_n : X \times [0, 1]^{2n} \rightarrow \mathbb{R}$, satisfying the conditions that total resale payoffs do not exceed the total surplus available among the bidders:

$$\sum_i v_i(x; \hat{\theta}; \theta) \leq (\max_i \theta_i) \cdot \left(\sum_i x_i \right) \quad (26)$$

for all $x, \hat{\theta}, \theta$, and are *individually rational*:

$$v_i(x; \hat{\theta}; \theta) \geq \theta_i x_i \quad (27)$$

for all $x, \hat{\theta}, \theta$.

Note that (27) is an ex post individual rationality constraint: the interpretation is that at the end of the bargaining, a deal (possibly probabilistic) is revealed, and then each player still has the option of walking away and keeping the allocation he received in the auction.

We will say that “no resale occurs” if $v_i(x; \hat{\theta}; \theta) = \theta_i x_i$ for all i (note that it allows resale that does not affect payoffs, i.e., between bidders who value the object equally). Note in particular that (26) and (27) imply that if either (a) x is efficient with respect to the true values, or (b) the good is not allocated ($x_i = 0$ for all i) then no resale occurs. Thus, if bidders have been truthful in a Vickrey auction with reserves then no resale will occur. However, this in itself does not guarantee that bidders who anticipate resale do not have advantageous deviations from truthtelling (recall the solution to the relaxed problem in Subsection 4.1). Yet, this turns out to be true for Vickrey auctions with reserves (the proof is in the Appendix):

Theorem 2. *Consider a Vickrey auction with reserves p_1, \dots, p_n . Then, for any individually rational resale payoffs satisfying (26) and (27) it is an ex*

⁹To be explicit about these beliefs we would want to also specify the disclosure rule used in the auction. We avoid modeling this, simply allowing the effects of disclosure to be implicitly captured by the v_i 's.

post equilibrium for all bidders to reveal their values truthfully (followed by no resale in equilibrium).

Formally, for any values $\theta_1, \dots, \theta_n$, and any possible deviation $\hat{\theta}_i$ for any bidder i ,

$$\theta_i \chi_i(\theta) - \psi_i(\theta) \geq v_i(\chi(\hat{\theta}_i, \theta_{-i}); \hat{\theta}_i, \theta_{-i}; \theta) - \psi_i(\hat{\theta}_i, \theta_{-i}). \quad (28)$$

Note that we obtain only ex post equilibrium here, not dominant strategies as we would in Vickrey auctions without resale. To see this concretely, imagine that there are two bidders, no reserve prices, and bidder 1 expects 2 to bid higher than his true value. Then, under our resale procedure (6), 1 has an incentive to underbid and make 2 win, since if 1 wins the object in the auction he has to pay 2's (exaggerated) bid, whereas to buy it in resale he only has to pay 2's true value. Thus, truthful bidding is not a dominant strategy. This is to be expected since our setting is one of interdependent values (compare Perry and Reny (2002) or Chung and Ely (2006)).

8 Conclusion

We have given here a model of an asymmetric, single-object auction in which the possibility of resale eliminates any benefit to the seller from conducting a biased auction. Under our robust revenue criterion, the optimal auction is a Vickrey auction with asymmetric reserves, which either allocates the good efficiently, or does not allocate it at all (when all bidders are below their reserve prices). We have also shown how to recursively compute the optimal reserves. We conclude by briefly outlining some possible criticisms and future directions.

One might take issue with our worst-case model, in which the bidders are uninformed about each other's values before the auction yet exogenously become informed after the auction. One possible defense of our approach is that additional information may be revealed after the auction (for example, through bargaining procedure itself). Our worst-case model is simply an extreme case, in which enough information is revealed to allow the high-value buyer to perfectly infer the auction winner's value. Another defense is that we could have instead written down a model in which the seller has even broader robustness concerns, allowing for bidders to learn about each other both before *and* after the auction, as well as allowing for resale. Since in a Vickrey auction with reserves, truthful bidding is an ex post equilibrium,

these auctions satisfy this additional form of robustness. Thus, the robustness to pre-auction information comes “for free,” without needing to require it explicitly in the seller’s problem.

One may also note that Vickrey auctions with reserves are susceptible to collusion, which, unlike resale, involves interim or even ex ante agreements among bidders. Of course, collusion is a problem even in the usual Vickrey auction, in which bidder i may bribe bidder j to underbid so as to win the object at a lower price. However, in addition to this kind of collusion, the Vickrey auction with reserves is susceptible to bidder i bribing bidder j to *overbid* so as to meet his reserve price, triggering the object’s sale to bidder i . Intuitively, the latter kind of collusion may be more incentive compatible than the former kind.¹⁰ One response to this criticism is that, unlike resale, collusion in auctions is usually illegal. Another response is that if the seller wants an auction that is robust to the latter kind of collusion as well as to resale, then, based on our analysis, she probably could do no better than using a Vickrey auction with the same reserve price to all bidders.

It may also be interesting to ask what the best robust auction is if the seller knows that bidders never learn each other’s values, so any resale procedure must be Bayesian incentive-compatible. For this case, Laffont and Martimort (1997) and Che and Kim (2006) show that any Bayesian incentive-compatible auction, including Myerson’s (1981) optimal auction, can be implemented in a way that is robust to collusion, and is therefore robust to resale as well (which may be viewed as a weaker form of collusion). Thus, it is essential to our result that some private information may be revealed before resale.

Finally, we have presented here a model in which resale can occur only because the outcome of the auction was inefficient. In practice, however, resale can occur for many reasons, and it would be natural to consider robust auction design for situations where resale is socially desirable: for example, buyers are uncertain about their values and will learn more about them after the auction; or, buyers anticipate flow benefits from owning the object, and the buyer with the highest value immediately after the auction may not be the one with highest value farther in the future. Presumably the optimal auctions in such models would involve resale occurring in equilibrium.

¹⁰Indeed, bidder i would not have an incentive to bribe bidder j to overbid unless his value is high and he believes that he is likely to win in any case, making bidder j more inclined to follow through. We thank Glen Weyl for this observation.

9 Appendix: Proofs

9.1 Proof of Lemma 1

The statement is implied by the claim that $R_k^* > R_k(0)$ for each $k = 1, \dots, n$, which we show inductively. For $k = 1$, the claim holds because $R_1^* = \max_{p \in [0,1]} R_1(p) > 0 = R_1(0)$. Now take $k > 1$ and suppose the inductive statement holds for $k - 1$. The derivative of the objective function in (18) is

$$R'_k(p) + f_k(p) [R_{k-1}^* - R_{k-1}(p)] - F_k(p) R'_{k-1}(p).$$

We can calculate $R'_k(p)$: as in (9), the revenue of any auction is the integral of virtual surplus over the region where the object is sold; in this case that is simply the region $\{\theta_i \geq p \text{ for some } i\}$. As p increases marginally, this region shrinks by losing the surface where one bidder's value is p and all other bidders' values are below p . Therefore, the derivative of R_k is the negative of the integral of virtual surplus over this surface. For each bidder i , there is one portion of the surface where i 's value is p (and other bidders are below p), and the virtual surplus is $\nu_i(p)$. Therefore,

$$R'_k(p) = - \sum_{i=1}^k \left(\prod_{\substack{j=1 \\ j \neq i}}^k F_j(p) \right) f_i(p) \nu_i(p). \quad (29)$$

Thus

$$R'_k(p) = F_k(p) R'_{k-1}(p) - \left(\prod_{i=1}^{k-1} F_i(p) \right) \nu_k(p) f_k(p),$$

and the derivative of the objective function in (18) can be written for $p \in [p_k, p_{k-1}]$ as

$$f_k(p) \left(R_{k-1}^* - R_{k-1}(p) - \nu_k(p) \prod_{i=1}^{k-1} F_i(p) \right). \quad (30)$$

The inductive statement for $k - 1$ implies that (30) is strictly positive at $p = 0$, and therefore we must have $p_k > 0$ and $R_k^* > R_k(0)$.

9.2 Proof of Lemma 2

Before proving the lemma, we establish the following important equation:

$$H'(p) = \sum_{j < k} \frac{f_j(p)}{F_j(p)} [\nu_j(p) - H(p)] \text{ for } p \in (p_k, p_{k-1}). \quad (31)$$

This equation can be derived by writing (20) on this interval as

$$H(p) \cdot \prod_{j < k} F_j(p) = R_{k-1}^* - R_{k-1}(p),$$

differentiating both sides, using (29) to express $R'_{k-1}(p)$, then dividing both sides by $\prod_{j < k} F_j(p)$ and rearranging.

Now we proceed to prove Lemma 2.

- (i) At $p \in (p_k, p_{k-1})$, $H(p)$ is continuous due to continuity of $R_{k-1}(p)$. At $p = p_k$, H is continuous by (17).
- (ii) The derivative of the objective function in (18) is given by (30), and can be expressed as

$$f_k(p) (H(p) - \nu_k(p)) \prod_{i=1}^{k-1} F_i(p)$$

where H is given by (20). It is necessary for maximization in (18) that this derivative evaluated at $p = p_k > 0$ must be zero if $p_k \in (0, p_{k-1})$ and nonnegative if $p_k = p_{k-1}$, so we must have

$$H(p_k) \geq \nu_k(p_k), \text{ and } H(p_k) = \nu_k(p_k) \text{ if } p_k < p_{k-1}. \quad (32)$$

Furthermore, we show that we actually have $H(p_k) = \nu_k(p_k)$ for each $k = 1, 2, \dots, n$, by induction on k . It is true for $k = 1$ (it is clearly optimal to set $p_1 < p_0$ which is the upper limit of the support, so that the revenue is positive). Suppose we know that $H(p_{k-1}) = \nu_{k-1}(p_{k-1})$. Consider two cases:

- If $p_k = p_{k-1}$ then using the inductive hypothesis and ordered virtual values we have

$$H(p_k) = H(p_{k-1}) = \nu_{k-1}(p_{k-1}) \leq \nu_k(p_{k-1}) = \nu_k(p_k),$$

and combined with (32) we have $H(p_k) = \nu_k(p_k)$.

– If $p_k < p_{k-1}$ then by (32) we have $H(p_k) = \nu_k(p_k)$.

(iii) For $p \in [p_k, p_{k-1}]$,

$$\begin{aligned} R_{k-1}^* - R_{k-1}(p) &= R_{k-1}^* - R_{k-1}(p_{k-1}) + \int_p^{p_{k-1}} R'_{k-1}(\tau) d\tau = \\ &= \nu_{k-1}(p_{k-1}) \prod_{j < k} F_j(p_{k-1}) - \int_p^{p_{k-1}} \left(\prod_{j < k} F_j(\tau) \right) \sum_{j < k} \frac{f_j(\tau) \nu_j(\tau)}{F_j(\tau)} d\tau \end{aligned}$$

where we have used part (ii) and the formula (29) for $R'_{k-1}(p)$.

(iv) Using (31) and assumption (A2), we see that for any $p \in [p_{i+1}, p_i]$ such that $H(p) \geq \nu_i(p)$ we have $H'(p) \leq 0$, and the inequality is strict if $H(p) > \nu_i(p)$. (Here we take $H'(p)$ to be the left-derivative at $p = p_i$ and right-derivative at $p = p_{i+1}$. Notice that part (i) above implies that the formula used to define H on the interval $[p_{i+1}, p_i)$ is also valid at the endpoint p_i , and the reasoning used to obtain (31) also applies at these endpoints.)

Now, recall from part (ii) that $H(p_i) = \nu_i(p_i)$, and therefore $H'(p_i) \leq 0$, hence for p in a left-neighborhood of p_i

$$\begin{aligned} H(p) - \nu_i(p) &\geq \nu_i(p_i) + o(p_i - p) - \nu_i(p) \\ &\geq p_i - p + o(p_i - p) > 0, \end{aligned}$$

where the second inequality uses log-concavity of $1 - F_i$ (monotone hazard rate). Thus, we can choose $\hat{p} < p_i$ close enough to p_i such that $H(p) > \nu_i(p)$ for all $p \in [\hat{p}, p_i)$, and therefore $H'(p) < 0$ strictly on this interval. In particular this implies $H(\hat{p}) > H(p_i)$. But then $H(p)$ cannot cross $\nu_i(p_i)$ anywhere in the interval $[p_{i+1}, p_i]$: indeed, otherwise, letting $p^0 = \max\{p \in [p_{i+1}, p_i] : H(p) \leq \nu_i(p_i)\}$ we would have

$$0 = H(p_i) - \nu_i(p_i) < H(\hat{p}) - \nu_i(p_i) = \int_{p^0}^{\hat{p}} H'(p) dp,$$

while on the other hand we would have $H(p) \geq \nu_i(p_i) \geq \nu_i(p)$ for all $p \in (p_0, \hat{p})$ and therefore $H'(p) \leq 0$ for all $p \in (p_0, \hat{p})$, making the right-hand side nonpositive – a contradiction.

Therefore, $H(p) \geq \nu_i(p_i) \geq \nu_i(p)$ for all $p \in [p_{i+1}, p_i]$, and therefore $H'(p) \leq 0$ on this interval. Since this holds for each i , H is nonincreasing on $[p_n, p_1]$.

9.3 Proof of Theorem 1

First we record a helpful reformulation of equation (31):

Corollary 1. *For any $k = 2, \dots, n$ and any $p \in [p_k, p_{k-1}]$,*

$$H(p) = \sum_{j < k} \int_{\theta}^{p_j} \frac{f_j(p)}{F_j(p)} [H(p) - \nu_j(p)] dp.$$

Proof. We check that the equality holds on the interval $[p_k, p_{k-1}]$ by induction on k . We may assume it holds at p_{k-1} , by the induction hypothesis and the fact that (by Lemma 2 (ii)) the $j = k - 1$ term of the sum is zero there. (The only exception is the base case $k = 2$ and $p = p_1$; in this case both sides of the desired equation are zero.) Then, given that H is continuous (Lemma 2 (i)), (31) ensures that the desired equation holds throughout the interval $[p_k, p_{k-1}]$. \square

Now, to prove the theorem we write out our Lagrangian (22) explicitly: For each $i = 1, \dots, n - 1$, define one-dimensional measures $\Lambda_i, \hat{\Lambda}_i$, with supports $[p_n, p_i]$ and $[0, p_i]$ respectively, whose distribution functions are $\hat{\Lambda}_i(\theta_i)$ as in Lemma 3 and $\hat{\Lambda}_i(\hat{\theta}_i) = F_i(\hat{\theta}_i)$. Lemma 3 ensures that Λ_i is a nondecreasing function, so this makes sense. Define $\mathcal{M}_i(\theta_i, \hat{\theta}_i)$ to be the restriction of the product measure $\Lambda_i \times \hat{\Lambda}_i$ to the half-plane $\hat{\theta}_i \leq \theta_i$. For notational convenience, also define $\Lambda_n, \hat{\Lambda}_n, \mathcal{M}_n$ to be zero everywhere.

Then the Lagrangian is

$$\begin{aligned} \mathbb{E}_{\tilde{\theta}} \left[\sum_i v_i(\chi(\tilde{\theta}); \tilde{\theta}) - \sum_i \frac{1 - F_i(\tilde{\theta}_i)}{f_i(\tilde{\theta}_i)} v'_i(\chi(\tilde{\theta}); \tilde{\theta}) \right] + \\ \sum_i \iint \left(\int_{\hat{\theta}_i}^{\theta_i} \mathbb{E}_{\tilde{\theta}_{-i}} \left[v'_i(\chi(\tau_i, \tilde{\theta}_{-i}); \tau_i, \tilde{\theta}_{-i}) \right] d\tau_i - \right. \\ \left. \mathbb{E}_{\tilde{\theta}_{-i}} \left[v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \hat{\theta}_i, \tilde{\theta}_{-i}) \right] \right) d\mathcal{M}_i(\theta_i, \hat{\theta}_i). \end{aligned} \quad (33)$$

We need to show the Lagrangian saddle-point conditions (a)-(c) stated in Section 4 are satisfied. Part (b) was already argued informally in the text,

and will also follow from Theorem 2 below, so we will not give a separate proof here.

For (c), consider any $\theta_i, \hat{\theta}_i$ with $\hat{\theta}_i \leq \theta_i < p_i$. We check that, in the Vickrey auction with reserves, type θ_i of bidder i gets the same expected payoff from telling the truth as by reporting type $\hat{\theta}_i$. In fact we can check that this is true ex post, i.e. for any fixed realizations of the other bidders' types, bidder i gets the same payoff from truth-telling as from misreporting.

Indeed: If $\theta_k < p_k$ for all $k \neq i$, then either when i tells the truth or when he lies, the object is not sold and his payoff is zero. Suppose $\theta_k \geq p_k$ for some k , so that the object is sold in both cases. If i wins the auction in both cases, then either way he has to pay a price $\max_{k \neq i} \theta_k$, and his payoff is the same. If he loses in both cases, then the winner of the auction (in both cases) has a value at least θ_i , so i does not buy the object in resale even after misreporting, and his payoff is zero in both cases. The only interesting possibility is that i wins the auction by truthfully reporting θ_i but loses by reporting $\hat{\theta}_i$. In this case, if i tells the truth, he wins at a price $\max_{k \neq i} \theta_k$. If i lies, the object is sold to the highest of the other bidders, and then i again pays $\max_{k \neq i} \theta_k$ to buy it back in resale. So either way, i 's payoff is $\theta_i - \max_{k \neq i} \theta_k$. This proves (c).

The rest of the proof will focus on (a). We rewrite the Lagrangian as

$$\mathbb{E}_{\tilde{\theta}} \left[\mu(\chi(\tilde{\theta}); \tilde{\theta}) \right],$$

where $\mu(x; \theta)$ is the “modified virtual surplus” from allocation x when the type profile is θ :

$$\begin{aligned} \mu(x; \theta) = & \sum_i \left[v_i(x; \theta) - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} v_i'(x; \theta) \right. \\ & + \frac{1}{f_i(\theta_i)} \hat{\Lambda}_i(\theta_i) (\Lambda_i(p_i) - \Lambda_i(\theta_i)) v_i'(x; \theta) \\ & \left. - \frac{1}{f_i(\theta_i)} \hat{\Lambda}_i'(\theta_i) \int_{\theta_i}^{p_i} [v_i(x; \tau_i, \theta_{-i}) - v_i(x; \theta)] d\Lambda_i(\tau_i) \right]. \end{aligned} \quad (34)$$

To see this, write the double-integral in the second and third lines of (33)

as

$$\begin{aligned} & \iint \mathbf{1}_{\hat{\theta}_i \leq \theta_i} \left[\int_{\hat{\theta}_i}^{\theta_i} \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\tau_i, \tilde{\theta}_{-i}); \tau_i, \tilde{\theta}_{-i})] d\tau_i \right] d\Lambda_i(\theta_i) d\hat{\Lambda}_i(\hat{\theta}_i) \\ & - \iint \mathbf{1}_{\hat{\theta}_i \leq \theta_i} \mathbb{E}_{\tilde{\theta}_{-i}} \left[v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \hat{\theta}_i, \tilde{\theta}_{-i}) \right] d\Lambda_i(\theta_i) d\hat{\Lambda}_i(\hat{\theta}_i). \end{aligned} \quad (35)$$

The first line rewrites as

$$\begin{aligned} & \iiint \mathbf{1}_{\hat{\theta}_i \leq \tau_i \leq \theta_i} \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\tau_i, \tilde{\theta}_{-i}); \tau_i, \tilde{\theta}_{-i})] d\tau_i d\Lambda_i(\theta_i) d\hat{\Lambda}_i(\hat{\theta}_i) \\ & = \int \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\tau_i, \tilde{\theta}_{-i}); \tau_i, \tilde{\theta}_{-i})] \left(\int \mathbf{1}_{\tau_i \leq \theta_i} d\Lambda_i(\theta_i) \right) \left(\int \mathbf{1}_{\hat{\theta}_i \leq \tau_i} d\hat{\Lambda}_i(\hat{\theta}_i) \right) d\tau_i \\ & = \int (\Lambda_i(p_i) - \Lambda_i(\tau_i)) \hat{\Lambda}_i(\tau_i) \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\tau_i, \tilde{\theta}_{-i}); \tau_i, \tilde{\theta}_{-i})] d\tau_i \\ & = \mathbb{E}_{\tilde{\theta}} \left[\frac{1}{f_i(\tilde{\theta}_i)} (\Lambda_i(p_i) - \Lambda_i(\tilde{\theta}_i)) \hat{\Lambda}_i(\tilde{\theta}_i) v'_i(\chi(\tilde{\theta}); \tilde{\theta}) \right] \end{aligned}$$

where we have made the change of variables $\tau_i \rightarrow \tilde{\theta}_i$ in the last line.

Similarly, the second line of (35) rewrites as

$$\begin{aligned} & - \int \left[\int_{\hat{\theta}_i}^{p_i} \mathbb{E}_{\tilde{\theta}_{-i}} \left[v_i(\chi(\tilde{\theta}_i, \tilde{\theta}_{-i}); \tau_i, \tilde{\theta}_{-i}) - v_i(\chi(\tilde{\theta}_i, \tilde{\theta}_{-i}); \tilde{\theta}_i, \tilde{\theta}_{-i}) \right] d\Lambda_i(\tau_i) \right] d\hat{\Lambda}_i(\tilde{\theta}_i) \\ & = -\mathbb{E}_{\tilde{\theta}} \left[\frac{\hat{\Lambda}'_i(\tilde{\theta}_i)}{f_i(\tilde{\theta}_i)} \int_{\hat{\theta}_i}^{p_i} \left[v_i(\chi(\tilde{\theta}); \tau_i, \tilde{\theta}_{-i}) - v_i(\chi(\tilde{\theta}); \tilde{\theta}) \right] d\Lambda_i(\tau_i) \right] \end{aligned}$$

(as long as $\hat{\Lambda}_i$ is absolutely continuous, which it is in our case.)

Putting these terms together, we see that the Lagrangian (33) is indeed the expectation of the modified virtual surplus spelled out in (34).

Note that for each type profile θ , the modified virtual surplus $\mu(x; \theta)$ is a linear function of the allocation x . We can write $\mu_i(\theta)$ for the x_i -coefficient, that is, the virtual value of allocating to bidder i ; thus the Lagrangian rewrites as

$$\mathbb{E}_{\tilde{\theta}} \left[\sum_i \mu_i(\tilde{\theta}) \chi_i(\tilde{\theta}) \right]$$

as in (14).

To show that our Vickrey auction with reserves is optimal, we show that it maximizes the modified virtual surplus pointwise: at every profile θ , the

virtual surplus is maximized by selling the good to the highest-value bidder if at all, and allocating to the high-value bidder is desirable if $\theta_i \geq p_i$ for some i and undesirable if $\theta_i < p_i$ for all i .

In symbols, we need to show:

(a-1) $\mu_{i^*(\theta)}(\theta) \geq 0$ if $\theta_i \geq p_i$ for some i , and ≤ 0 otherwise;

(a-2) $\mu_{i^*(\theta)}(\theta) \geq \mu_j(\theta)$ for all other bidders j .

We first show (a-1). In this case, writing $i^* = i^*(\theta)$, the allocation x that allocates to bidder i^* satisfies $v_{i^*}(x, \theta) = \theta_{i^*}$, $v'_{i^*}(x, \theta) = 1$ (aside from the measure-zero case of ties), and $v_i(x, \theta) = v'_i(x, \theta) = 0$ for every other bidder i . So we have

$$\begin{aligned} \mu_{i^*}(\theta) &= \theta_{i^*} - \frac{1 - F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} + \frac{\hat{\Lambda}_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} (\Lambda_{i^*}(p_{i^*}) - \Lambda_{i^*}(\theta_{i^*})) \\ &\quad - \sum_i \frac{\hat{\Lambda}'_i(\theta_i)}{f_i(\theta_i)} \int_{\theta_i}^{p_i} [v_i(x; \tau_i, \theta_{-i}) - v_i(x; \theta)] d\Lambda_i(\tau_i). \end{aligned} \quad (36)$$

To evaluate this modified virtual surplus, we consider cases depending on where θ_{i^*} lies.

- If $\theta_{i^*} \in [p_n, p_{i^*}]$, then we can plug in the formulas for $\hat{\Lambda}_{i^*}$ and Λ_{i^*} , and we see that the first line on the right side of (36) becomes

$$\nu_{i^*}(\theta_{i^*}) + \frac{F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} \left(\frac{f_{i^*}(\theta_{i^*})}{F_{i^*}(\theta_{i^*})} [H(\theta_{i^*}) - \nu_{i^*}(\theta_{i^*})] \right) = H(\theta_{i^*}).$$

Meanwhile, for the second line, we consider the expression in the integral, $v_i(x; \tau_i, \theta_{-i}) - v_i(x; \theta)$. For $i = i^*$, this is simply $\tau_{i^*} - \theta_{i^*}$. On the other hand, for any other bidder i , when $\tau_i < \theta_{i^*}$, both v_i terms are zero; and when $\tau_i \geq \theta_{i^*}$, the first term becomes $\tau_i - \theta_{i^*}$ since now i can buy back from i^* in resale. So for *each* bidder i , the integrand is equal to $\max\{0, \tau_i - \theta_{i^*}\}$. Thus (36) becomes

$$\begin{aligned} H(\theta_{i^*}) - \sum_i \int_{\theta_i}^{p_i} \max\{0, \tau_i - \theta_{i^*}\} d\Lambda_i(\tau_i) \\ &= H(\theta_{i^*}) - \sum_i \int_{\theta_{i^*}}^{p_i} (\tau_i - \theta_{i^*}) d\Lambda_i(\tau_i) \\ &= H(\theta_{i^*}) - \sum_i \int_{\theta_{i^*}}^{p_i} (\Lambda_i(p_i) - \Lambda_i(\tau_i)) d\tau_i \end{aligned}$$

where the last equation is integration by parts. Note that the sum is only over i such that $\theta_{i^*} \leq p_i$, since for other i 's the integral is zero. Plugging in

$$\Lambda_i(p_i) - \Lambda_i(\tau_i) = \frac{f_i(\tau_i)}{F_i(\tau_i)} [H(\tau_i) - \nu_i(\tau_i)],$$

we see that the sum of integrals is exactly $H(\theta_{i^*})$, by Corollary 1. So the modified virtual surplus is zero, whenever $\theta_{i^*} \in [p_n, p_{i^*}]$.

- If $\theta_{i^*} > p_{i^*}$, the analysis is the same as in the previous case, except that in the first line of (36), the third term is now zero. So this line equals $\nu_{i^*}(\theta_{i^*})$ instead of $H(\theta_{i^*})$. That is, the modified virtual surplus evaluates to $\nu_{i^*}(\theta_{i^*}) - H(\theta_{i^*})$.

Since ν_{i^*} is weakly increasing, H is decreasing, and they are equal at p_{i^*} , the modified virtual surplus is nonnegative — as it should be in order to be consistent with (a-1).

- If $\theta_{i^*} < p_n$, then $\theta_i < p_n \leq p_i$ for all i , so (a-1) requires that the modified virtual value should be (weakly) negative.

Hold fixed the identity of i^* , and let the value of θ_{i^*} on the interval $[0, p_n]$. The first line of (36) is

$$\nu_{i^*}(\theta_{i^*}) + \frac{F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} \left(\frac{f_{i^*}(p_n)}{F_{i^*}(p_n)} [H(p_n) - \nu_{i^*}(p_n)] \right).$$

This is increasing in θ_{i^*} , since ν_{i^*} is increasing, F_{i^*}/f_{i^*} is increasing (by the log-concavity assumption), and $H(p_n) = \nu_n(p_n) \geq \nu_{i^*}(p_{i^*})$ using Lemma 2. Meanwhile the sum in the second line of (36) is still equal to $\sum_i \int_{\theta_{i^*}}^{p_i} (\Lambda_i(p_i) - \Lambda_i(\tau_i)) d\tau_i$. This is clearly decreasing in θ_{i^*} .

Thus, the modified virtual value $\mu_{i^*}(\theta)$ is increasing in θ_{i^*} on the interval $[0, p_n]$. Since we already saw that it equals zero at the upper endpoint p_n , it is negative on this interval, as needed.

Thus in every case, the modified virtual value $\mu_{i^*}(\theta)$ has a sign consistent with (a-1).

Now we prove (a-2). The above calculations show that

$$\begin{aligned}\mu_{i^*}(\theta) &= \theta_{i^*} - \frac{1 - F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} + \frac{\hat{\Lambda}_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} (\Lambda_{i^*}(p_{i^*}) - \Lambda_{i^*}(\theta_{i^*})) \\ &\quad - \sum_i \int_{\theta_i}^{p_i} \max\{0, \tau_i - \theta_{i^*}\} d\Lambda_i(\tau_i).\end{aligned}$$

Now consider allocating to any $j \neq i^*$; we need to compute the modified virtual value from allocating to j , $\mu_j(\theta)$. In the first line of formula (34), the v_i terms sum to θ_{i^*} , and the v'_i terms are 1 for $i = i^*, j$ and 0 otherwise. In the second line, again the v'_i terms are 1 for $i = i^*, j$ and 0 otherwise. In the third line, there are three cases for the bracketed expression $v_i(x; \tau_i, \theta_{-i}) - v_i(x; \theta)$:

- for $i = i^*$, it equals $(\tau_{i^*} - \theta_j) - (\theta_{i^*} - \theta_j) = \tau_{i^*} - \theta_{i^*}$;
- for $i = j$, it equals $\tau_j - \theta_j$;
- for any $i \neq i^*, j$, then when $\tau_i < \theta_{i^*}$ it equals $0 - 0 = 0$, and when $\tau_i > \theta_{i^*}$ it equals $(\tau_i - \theta_j) - 0 = \tau_i - \theta_j$.

Notice that also in the third line of (34), the fraction $\hat{\Lambda}'_i(\theta_i)/f_i(\theta_i)$ can be simplified to 1 (this is true unless $\theta_i > p_i$ but in that case the integral is zero anyway).

So, the modified virtual value is

$$\begin{aligned}\mu_j(\theta) &= \theta_{i^*} - \frac{1 - F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} - \frac{1 - F_j(\theta_j)}{f_j(\theta_j)} \\ &\quad + \frac{\hat{\Lambda}_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} (\Lambda_{i^*}(p_{i^*}) - \Lambda_{i^*}(\theta_{i^*})) + \frac{\hat{\Lambda}_j(\theta_j)}{f_j(\theta_j)} (\Lambda_j(p_j) - \Lambda_j(\theta_j)) \\ &\quad - \sum_{i \neq i^*, j} \int_{\theta_{i^*}}^{p_i} [\tau_i - \theta_j] d\Lambda_i(\tau_i) - \int_{\theta_{i^*}}^{p_{i^*}} [\tau_{i^*} - \theta_{i^*}] d\Lambda_{i^*}(\tau_{i^*}) - \int_{\theta_j}^{p_j} [\tau_j - \theta_j] d\Lambda_j(\tau_j).\end{aligned}$$

Subtracting, we have

$$\begin{aligned}\mu_{i^*}(\theta) - \mu_j(\theta) &= \frac{1 - F_j(\theta_j)}{f_j(\theta_j)} - \frac{\hat{\Lambda}_j(\theta_j)}{f_j(\theta_j)} (\Lambda_j(p_j) - \Lambda_j(\theta_j)) \\ &\quad + \sum_{i \neq i^*, j} \int_{\theta_{i^*}}^{p_i} [\theta_{i^*} - \theta_j] d\Lambda_i(\tau_i) \\ &\quad - \int_{\theta_{i^*}}^{p_j} [\tau_j - \theta_{i^*}] d\Lambda_j(\tau_j) + \int_{\theta_j}^{p_j} [\tau_j - \theta_j] d\Lambda_j(\tau_j).\end{aligned}\tag{37}$$

We show that each of the three lines of (37) is nonnegative.

For the first line, if $\theta_j > p_j$ then the second term is zero, so we have just $(1 - F_j(\theta_j))/f_j(\theta_j)$. If $\theta_j \in [p_n, p_j]$, the first line is

$$\begin{aligned} \frac{1 - F_j(\theta_j)}{f_j(\theta_j)} - \frac{F_j(\theta_j)}{f_j(\theta_j)} \cdot \frac{f_j(\theta_j)}{F_j(\theta_j)} [H(\theta_j) - \nu_j(\theta_j)] &= (\theta_j - \nu_j(\theta_j)) - [H(\theta_j) - \nu_j(\theta_j)] \\ &= \theta_j - H(\theta_j). \end{aligned}$$

But

$$\theta_j - H(\theta_j) \geq p_n - H(p_n) \geq \nu_n(p_n) - H(p_n) = 0$$

where we have used Lemma 2(iv) and (ii). Finally, for $\theta_j \leq p_n$, the first line of (37) is

$$\frac{1 - F_j(\theta_j)}{f_j(\theta_j)} - \frac{F_j(\theta_j)}{f_j(\theta_j)} [\Lambda_j(p_j) - \Lambda_j(p_n)].$$

The difference in brackets is a positive constant as θ_j varies; the log-concavity assumptions ensure $(1 - F_j(\theta_j))/f_j(\theta_j)$ is decreasing and $F_j(\theta_j)/f_j(\theta_j)$ is increasing, so the whole expression is decreasing in θ_j . Since we already saw that it is nonnegative at $\theta_j = p_n$, it must be nonnegative for $\theta_j < p_n$.

The second line of (37) is clearly nonnegative. And the third line is

$$\int_{\theta_j}^{p_j} [\min\{\tau_j, \theta_{i^*}\} - \theta_j] d\Lambda_j(\tau_j)$$

which is also clearly nonnegative.

Thus $\mu_{i^*}(\theta) - \mu_j(\theta) \geq 0$, proving (a-2).

At this point we have shown (a-1) and (a-2). Together, these imply that at every possible θ , our Vickrey auction with reserves p_1, \dots, p_n picks out an allocation that maximizes the modified virtual surplus, which shows that this auction maximizes the Lagrangian (33). This completes the proof.

9.4 Proof of Theorem 2

We consider all possible cases.

- Suppose that bidder i wins under truth-telling: $\chi_i(\theta) = 1$. Then the left-hand side of (28) is $\theta_i - \psi_i(\theta) \geq 0$.
 - If i deviates to $\hat{\theta}_i$ such that he still wins, then he pays the same price (and by the note preceding the theorem, no resale occurs).

- If i deviates to $\hat{\theta}_i$ such that some bidder $j \neq i$ wins (and i then pays zero), this is only possible if $\theta_j = \max_{k \neq j} \theta_k = \psi_i(\theta)$ (i 's threshold price under truth-telling). Then (27) implies that the resale payoff of bidder j is at least θ_j and the resale payoffs of all bidders $k \notin \{i, j\}$ are nonnegative, and therefore by (26) the resale payoff of bidder i is at most $\theta_i - \theta_j$. So the right side of (28) is at most $\theta_i - \theta_j$ which equals the left side.
- If i deviates to $\hat{\theta}_i$ such that the object goes unallocated, then no resale is possible, so the right side of (28) is zero.
- Suppose that some bidder $j \neq i$ wins under truth-telling: $\chi_j(\theta) = 1$. Then the left-hand side of (28) is zero.
 - If i deviates to a $\hat{\theta}_i$ that does not win the object, then either j still wins, or the object is unsold. Either way, no resale occurs, and the right side of (28) is zero.
 - If i deviates to win the object, he pays the threshold price $\psi_i(\hat{\theta}_i, \theta_{-i}) = \theta_j = \max_k \theta_k$. By (27), the resale payoffs of all other bidders are nonnegative, and therefore by (26) the resale payoff of bidder i is at most $\max_k \theta_k = \theta_j$. So the right side of (28) is at most $\theta_j - \theta_j = 0$.
- Suppose that the object is left unsold under truth-telling: $\chi(\theta) = 0$. Then the left-hand side of (28) is zero.
 - If i deviates such that the object remains unsold, then no resale is possible and the right side of (28) is zero.
 - If i deviates such that some bidder $j \neq i$ wins the object, then allocating to j is efficient for the reported value profile $(\hat{\theta}_i, \theta_{-i})$, and $\hat{\theta}_i$ is above the reserve while the true value θ_i is not:
$$\theta_j \geq \hat{\theta}_i \geq p_i \geq \theta_i, \text{ and } \theta_j \geq \theta_k \text{ for all } k \neq i.$$
So allocating to j is efficient for the true value profile θ , and no resale is possible, so the right side of (28) is zero.
 - If i deviates to win the object, he pays the threshold price $\psi_i(\hat{\theta}_i, \theta_{-i}) = \max\{p_i, \max_{j \neq i} \theta_j\}$. By (27), the resale payoffs of all other bidders are nonnegative, and therefore by (26) the resale payoff of

bidder i is at most $\max_j \theta_j$. Since $\theta_i < p_i$ (because the object is unsold under truth-telling), this expression does not exceed $\max\{p_i, \max_{j \neq i} \theta_j\} = \psi_i(\hat{\theta}_i, \theta_{-i})$, so the right side of (28) is at most zero.

References

- [1] Ausubel, L. M., and P. Cramton (1999): “The Optimality of Being Efficient,” working paper, Department of Economics, University of Maryland, College Park.
- [2] Ausubel, Lawrence M. and Peter Cramton (2004), “Vickrey Auctions with Reserve Pricing,” *Economic Theory*, Volume 23, Issue 3, pp. 493–505
- [3] Bagnoli, M., and T. Bergstrom (2005), “Log-concave probability and its applications,” *Economic Theory* 26(2): 445–469
- [4] Calzolari, G., A. Pavan (2006), “Monopoly with Resale,” *Rand Journal of Economics* 37(2): 362–375.
- [5] Che, Y.-K., and J. Kim (2006), “Robustly Collusion-Proof Implementation,” *Econometrica* 74(4): 1063–1107
- [6] Chung, K.-S., and J.C. Ely (2006), “Ex-Post Incentive Compatible Mechanism Design,” working paper, Northwestern University.
- [7] Chung, K.-S., and J.C. Ely (2007), “Foundations of Dominant Strategy Mechanisms,” *Review of Economic Studies* 74: 447–476
- [8] Dworzak, P. (2015), “The Effects of Post-Auction Bargaining between Bidders,” working paper, Stanford University
- [9] Laffont, J.-J., and D. Martimort (1997), “Collusion under Asymmetric Information,” *Econometrica*, 65, pp. 875–911.
- [10] Luenberger, D. G. (1969), *Optimization by Vector Space Methods*. John Wiley & Sons.

- [11] McAfee, R.P., and J. McMillan (1989), “Government Procurement and International Trade,” *Journal of International Economics* 26, pp. 921–943
- [12] Milgrom, P., and I. Segal (2002), “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica* 70(2), 583-601
- [13] Myerson, R. (1981), “Optimal Auction Design,” *Mathematics of Operations Research* 6(1), pp. 58-73
- [14] Myerson, R., and M. Satterthwaite (1983), “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory* 29: 265-281
- [15] Myerson, R., and P. Reny (2015), “Open Sequential Equilibria of Multi-Stage Games with Infinite Sets of Types and Actions,” working paper, University of Chicago.
- [16] Perry, M., and P. Reny (2002), “An Efficient Auction,” *Econometrica* 70(3): 1199–1212.
- [17] Royden, H.L. (1988), *Real Analysis*, Prentice-Hall.
- [18] Yamashita, T. (2015), “Implementation in Weakly Undominated Strategies: Optimality of Second-Price Auction and Posted-Price Mechanism,” *Review of Economic Studies* 82, pp. 1223-1246
- [19] Zheng, Charles Zhoucheng (2002), “Optimal Auctions with Resale,” *Econometrica*, Vol. 70, No. 6, 2197–2224