

Robustly Optimal Auctions with Unknown Resale Opportunities*

Gabriel Carroll Ilya Segal

Department of Economics
Stanford University, Stanford, CA 94305, USA

February 13, 2018

Abstract

The standard revenue-maximizing auction discriminates against a priori stronger bidders so as to reduce their information rents. We show that such discrimination is no longer optimal when the auction’s winner may resell to another bidder, and the auctioneer has non-Bayesian uncertainty about such resale opportunities. We identify a “worst-case” resale scenario, in which bidders’ values become publicly known after the auction and losing bidders compete Bertrand-style to buy the object from the winner. With this form of resale, misallocation no longer reduces the information rents of the high-value bidder, as he could still secure the same rents by buying the object in resale. Under regularity assumptions, we show that revenue is maximized by a version of the Vickrey auction with bidder-specific reserve prices,

*We are grateful to Adrien Auclert, Lawrence M. Ausubel, Ben Brooks, Peter Cramton, Drew Fudenberg, Simon Grant, Paul Milgrom, Isaac Sorkin, Glen Weyl, and participants of the 2016 Stony Brook Workshop on Complex Auctions and Practice, the NBER Market Design Working Group, the NYUAD Advances in Mechanism Design conference, the NSF-CEME Decentralization conference, and seminars at ANU, UNSW, UQ, HU-Berlin, Mannheim, Bonn, Northwestern, and Harvard-MIT, for helpful comments. We thank the Simons Institute at UC Berkeley for its hospitality during the Fall 2015 Economics and Computation Program. Segal acknowledges the support of the National Science Foundation (grant SES 0961693). Carroll is supported by a Sloan Research Fellowship. Alex Bloedel and Seunghwan Lim provided valuable research assistance.

first proposed by Ausubel and Cramton (2004). The proof of optimality involves constructing Lagrange multipliers on a double continuum of binding non-local incentive constraints.

1 Introduction

Standard auction theory says that when bidders are a priori asymmetric, revenue-maximizing auctions discriminate against “stronger” bidders, i.e. those who are more likely to have higher values. The optimal auction requires them to pay a premium over “weaker” bidders’ values in order to win (Myerson 1981, McAfee and McMillan 1989). This discrimination enhances revenue because it increases competition between weaker and stronger bidders and so reduces the latter bidders’ information rents.

One might suspect that the benefits of such discrimination would be viated if bidders could resell to each other after the auction, since a strong bidder might then prefer to sit back and let a weaker bidder win, in the hopes of buying from him later at a better price. In the words of Ausubel and Cramton (1999, p. 19): “The possibility of resale undermines the seller’s ability to gain by misassigning the good. The best that the seller can do is to conduct an efficient auction [i.e. selling only to the highest-value bidder], perhaps withholding ... the good.”¹ However, this is not supported by existing results on optimal auctions with resale (Zheng 2002, Calzolari and Pavan 2006), in which revenue is typically maximized by a biased auction, which induces resale in equilibrium with a positive probability. For a simple example due to Zheng (2002), suppose one bidder is known to have a zero value for the good and also to have full bargaining power in resale. Then the auctioneer would want to simply sell to this bidder, who would then resell the good using Myerson’s (1981) revenue-maximizing mechanism. The auctioneer could charge the zero-value bidder a price equal to the Myerson expected revenue. (This is the best she can do as long as bidders’ values remain private information, since the combined outcome of the auction and resale would have to satisfy the same incentive compatibility constraints as

¹The notion that price discrimination is ineffective under resale is also familiar in the context of pricing in markets. For example, Tirole (1988, p. 134) writes: “It is clear that if the transaction (arbitrage) costs between two consumers are low, any attempt to sell a given good to two consumers at different prices runs into the problem that the low-price consumer buys the good to resell it to the high-price one.”

Myerson’s mechanism.)

Ausubel and Cramton (1999, 2004) claim that it is optimal for the auctioneer not to misallocate when resale is “perfect.” Specifically, they propose a way of adding bidder-specific reserve prices to a Vickrey (second-price) auction, which we call the Ausubel-Cramton-Vickrey (ACV) auction and describe in detail below. The ACV either allocates the object efficiently among bidders, or withholds it (when no bidder beats his assigned reserve).² Ausubel and Cramton (1999) assert that the ACV auction is optimal. They do not formalize the assumption of perfect resale, and as is well known, efficient resale would generally be impossible under private information;³ but one might justify the assumption by assuming that the parties’ values exogenously become publicly known before resale.

However, even if we operationalize perfect resale in this way, it could still be optimal for the auctioneer to misallocate. Indeed, if we modify the zero-value bidder example above by letting the zero bidder observe all other bidders’ values before reselling, the optimal mechanism would extract full *first-best* surplus — by again selling the object to the zero bidder, now at a price equal to the expectation of the highest value, which is exactly what the zero bidder expects to receive in resale. Similarly, in the symmetric setting with perfect resale studied by Bulow and Klemperer (2002) and Bergemann et al. (2017a), where the auction’s winner is assumed to have full information and full bargaining power in resale, the auctioneer can command a higher price by allocating to an inefficient bidder than to the efficient one, because the inefficient bidder can later extract the efficient bidder’s surplus in resale; thus the optimal auction misallocates the good.⁴

²At least two other ways to introduce asymmetric reserves into Vickrey have been studied: “lazy” and “eager” Vickrey (Dhangwatnotai et al., 2010). “Eager” Vickrey allocates to the highest bidder among those who met their reserve prices as long as one such bidder exists, and so it misallocates in equilibrium. “Lazy” Vickrey allocates to the highest bidder provided that he met his reserve price and does not allocate otherwise. It will emerge as the solution to our relaxed problem in Subsection 4.2.

³For example, when the auction simply gives the object to bidder 1 for zero payment, so that no information is revealed, resale must be inefficient according to the theorem of Myerson and Satterthwaite (1983). The same ideas apply to resale following more general auction mechanisms.

⁴Bulow and Klemperer (2002, Example 3) find that in this setting, selling to a randomly chosen bidder at a fixed price low enough to guarantee a sale yields a higher revenue than the Vickrey auction. Bergemann et al. (2017a) show that the fixed-price mechanism is in fact the optimal mechanism among those that always sell the good. They also derive the

Our paper shows how the intuition expressed by Ausubel and Cramton can nevertheless be validated when the auction designer is uncertain about the resale procedure (including possible exogenous revelation of private information) and desires the revenue to be robust to this uncertainty. Indeed, notice that some of the optimal biased auctions described in the above examples are not robust to even a small amount of uncertainty about resale.⁵ Thus, we formulate the problem of maximizing worst-case expected revenue, where the expectation is taken over buyers' private values drawn from known distributions, and the worst case is over the possible resale procedures. We refer to this problem as the robust revenue maximization problem.

To solve this problem, we begin with a simplified model in which the auctioneer is required to always sell the object (Section 3). With no resale, the optimal auction in this model would be biased. In contrast, robust revenue maximization under resale uncertainty is achieved by the simple second-price auction (with no reserve), which allocates efficiently. To prove this, we guess a “worst-case” resale procedure, in which bidders' values exogenously become public knowledge after the auction, and then the losing bidders compete à la Bertrand to buy the object from the winner in resale. With this resale procedure, in any auction that always sells the object, each bidder can guarantee himself a payoff equal to his marginal contribution to social surplus, by sitting out to let another bidder win and then buying from the winner. Given this lower bound on the information rents captured by the bidders, the designer cannot do better for herself than the Vickrey auction. Since this auction also sustains truthful bidding as an ex post equilibrium under any other resale procedure, it solves the robust revenue maximization problem. As a side benefit, the optimal auction turns out to be independent of the bidders' value distributions, i.e., completely prior-free.

While the must-sell model cleanly illustrates the robust optimality of efficient auctions, in most real-life settings (and in the classical theory) the auctioneer has the ability to increase revenue by sometimes withholding the object, e.g. by using reserve prices. In Section 4, we turn to study such settings, and show that under appropriate regularity assumptions on the

optimal auction that can withhold the good, and show that it misallocates with a high probability. While written independently of our paper, their paper shares some technical similarities with our analysis, which we point out below.

⁵In either of the zero-bidder examples, suppose the auctioneer's guess about resale is mistaken, and there is actually an ε probability that no resale opportunity will arise. Then the zero bidder refuses to buy at the proposed price, and revenue drops to zero.

distributions, an ACV auction with appropriately chosen reserve prices is optimal. Thus, the auctioneer again never misallocates the good, but takes advantage of the known asymmetries by setting different reserve prices to different bidders. We show this result by proving that the ACV auction is optimal under the particular worst-case resale procedure described above. Since, by an argument of Ausubel and Cramton (2004) (which we include in the appendix for completeness), this auction sustains truthful bidding as an equilibrium under *any* resale procedure, it solves the robust revenue maximization problem.

Characterizing an optimal auction under our worst-case resale procedure requires new techniques that may be of separate interest. First, by folding the outcome of resale into reduced-form payoff functions, we model the auctioneer’s problem as a single-stage auction design problem, albeit one with externalities and interdependent values (since a bidder who does not win the object cares whether another bidder wins, and what the winner’s value is.)⁶ We can then apply the standard “first-order approach” to such problems, which considers only the local incentive constraints, and rewrites the objective as an appropriately-defined virtual surplus.

The solution found by this standard method reinforces our basic intuition: it never misallocates the good; it either allocates efficiently or withholds the good. However, this solution is not the right answer, because it is not incentive-compatible. It is vulnerable to non-local deviations, where a bidder underbids to lose and then buys the good in resale. To find the true optimal auction, we need to account for such non-local incentive constraints. We guess that the optimum is the ACV auction. To verify optimality, we explicitly construct supporting Lagrange multipliers on the double continuum of binding non-local incentive constraints.⁷

Our approach also yields an iterative construction of the optimal bidder-specific reserve prices in ACV auctions. Appendix F illustrates by working

⁶Other examples of externalities and interdependent values caused by post-auction interactions among buyers can be found in Jehiel and Moldovanu (2000) and Bulow and Klemperer (2002).

⁷Earlier work that constructed Lagrange multipliers as a measure over a double continuum of binding incentive constraints includes the “transport theory” approach to multidimensional screening (e.g., Daskalakis et al. (2013)). Also, Bergemann, Brooks, and Morris (2017a) independently develop a treatment of non-local incentive constraints that is the closest technically to our approach: our analysis implicitly shares with theirs the feature of considering randomized misreports that are drawn from the same distribution as the true type, but truncated.

through an example in which bidders' values are distributed uniformly with different upper limits. In this case, with bidders ordered from stronger to weaker, the optimal reserve price for the k th bidder is obtained by solving a k th-degree polynomial equation.

Some readers may find a tension between the auctioneer's use of Bayesian priors to construct optimal reserve prices in the ACV auction and her complete ignorance of the resale procedure. We do believe that there are real-life situations in which the auctioneer has knowledge that some bidders are stronger than others, which is traditionally captured by means of Bayesian priors, while being ignorant about the resale procedure.⁸ Furthermore, as we detail in Section 5, the tension can be directly ameliorated in two ways. First, we offer a model with uncertainty both about value distributions and about resale, in which our conclusion about optimality of ACV auctions carries over. Second, we offer additional results suggesting that setting the right reserve prices is quantitatively less important than ignorance of resale, so that simply using Vickrey with no reserves is a good prior-free auction choice.

Our paper joins the growing literature using the maxmin criterion to model robust mechanism design (including, in particular, Frankel (2014), Carroll (2015), and Bergemann et al. (2017b)). Conceptually, it is the closest to the work showing that a strategy-proof mechanism may emerge as optimal when the designer is uncertain about bidders' beliefs about each other's values (Chung and Ely 2007; Yamashita and Zhu 2017) or about each other's strategies (Yamashita 2015). Likewise, in our paper, a resale-proof mechanism may emerge as optimal when the designer is ignorant about the resale procedure. Also, observe that ACV auctions are also robust to bidders' beliefs and information about each other's values and their beliefs about the resale procedure, hence we obtain these additional robustness benefits "for free."

⁸To give one example, in spectrum auctions run by the U.S. Federal Communications Commission, some bidders (such as AT&T, Verizon, and T-Mobile) are believed to have strong business cases for the use of additional spectrum, while other bidders (such as Dish and Comcast) are believed to be bidding speculatively, with their values determined to a large extent by expectation of resale, even though the structure and outcomes of the resale market are hard to anticipate.

2 Setup

There are $n \geq 2$ bidders. Bidder i 's private type is his value θ_i for the object, which is distributed on $[0, 1]$ according to a c.d.f. F_i with a continuous strictly positive density f_i .⁹ 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 . We use tildes to denote random variables: $\tilde{\theta}_i$ for the random variable representing i 's value; θ_i for a specific realization.

A general auction mechanism is a triple $\Gamma = \langle M, \chi, \psi \rangle$, where

- $M = (M_1, \dots, M_n)$ is a collection of measurable message spaces for the bidders, such that $\emptyset \in M_i$ for each i , where \emptyset denotes the special non-participation message;
- $\chi : \prod_i M_i \rightarrow X$ is a measurable allocation rule and $\psi : \prod_i M_i \rightarrow \mathbb{R}^n$ is a measurable payment rule, with $\chi_i(m) = \psi_i(m) = 0$ whenever $m_i = \emptyset$.

Following allocation specified by the auction, resale may take place. We model resale in reduced form by an n -tuple of measurable functions $v = (v_1, \dots, v_n)$, where $v_i(x; \theta)$ gives bidder i 's post-resale payoff (net of payments in the auction) following allocation $x \in X$ specified by the auction when the bidders' value profile is θ . This formalism captures a setting in which all bidders' values exogenously become public after the auction, so that the outcome of resale depends only on the initial allocation and the values. (In Appendix C, we describe a more general class of resale procedures that does not assume values are revealed.)

We require that the total reduced-form payoffs not exceed the maximal total surplus available in resale:

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

We also require the resale procedure to be individually rational:

$$v_i(x; \theta) \geq \theta_i x_i \text{ for each } i. \quad (2)$$

⁹The assumptions of common support and continuous positive densities are made for expositional simplicity: in Appendix F we explain how our results can be extended to some cases in which these assumptions do not hold.

A given mechanism Γ and resale procedure v together induce a Bayesian game: the action space of player i is M_i , and his payoff is $v_i(\chi(m); \theta) - \psi_i(m)$, with corresponding revenue $\sum_i \psi_i(m)$ for the auctioneer.

Let $\text{Rev}(\Gamma, v)$ denote the supremum of expected revenue over all Bayes-Nash equilibria of this game. We state the robust revenue maximization problem as¹⁰

$$\max_{\Gamma} \left(\inf_v \text{Rev}(\Gamma, v) \right). \quad (3)$$

We will establish that a specific auction $\bar{\Gamma}$ solves the robust revenue maximization problem, by constructing a resale procedure \underline{v} and a revenue target \bar{R} such that the following two conditions are satisfied:

$$\text{Rev}(\Gamma, \underline{v}) \leq \bar{R} \text{ for all auctions } \Gamma, \quad (4)$$

$$\text{Rev}(\bar{\Gamma}, v) \geq \bar{R} \text{ for all resale procedures } v. \quad (5)$$

(4) means that given resale procedure \underline{v} , the designer could not design an auction yielding expected revenue above \bar{R} , while (5) means that given auction $\bar{\Gamma}$, an adversary could not construct a resale procedure to reduce the designer's expected revenue below \bar{R} . The logic of the Minimax Theorem then implies (the formal proof of this and all other results are in the appendix)

Lemma 1. *If (4)-(5) hold then auction $\bar{\Gamma}$ solves the robust revenue maximization problem $\max_{\Gamma} (\inf_v \text{Rev}(\Gamma, v))$, while resale procedure \underline{v} solves the worst-case resale problem $\min_v (\sup_{\Gamma} \text{Rev}(\Gamma, v))$, and the value of both problems is $\text{Rev}(\bar{\Gamma}, \underline{v}) = \bar{R}$.*

To establish (4) for a particular resale procedure \underline{v} , we can apply the Revelation Principle and restrict attention to direct mechanisms Γ , in which each bidder i 's message space is $M_i = [0, 1] \cup \{\emptyset\}$, and to the Bayes-Nash equilibrium in which all bidders participate and report truthfully, i.e., Γ satisfies

¹⁰Note, in particular, that $\text{Rev}(\Gamma, v) = -\infty$ if the induced game has no equilibrium, so the designer must guarantee equilibrium existence for all v . We view this requirement as not too onerous; for example, any mechanism Γ with finite message spaces guarantees equilibrium existence (Milgrom and Weber, 1985, Theorem 1). It is in line with the standard assumption in mechanism design that agents must play an equilibrium.

the following incentive compatibility and individual rationality constraints:¹¹

$$\mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \psi_i(\theta_i, \tilde{\theta}_{-i})] \geq \mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \psi_i(\hat{\theta}_i, \tilde{\theta}_{-i})] \quad \text{for all } \theta_i, \hat{\theta}_i; \quad (6)$$

$$\mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \psi_i(\theta_i, \tilde{\theta}_{-i})] \geq \mathbb{E}_{\tilde{\theta}_{-i}} \left[v_i(\chi(\emptyset, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) \right]. \quad \text{for all } \theta_i. \quad (7)$$

Then, we show (4) by showing that the expected revenue in any direct auction satisfying (6)-(7) cannot exceed the target \bar{R} , which we take to be the expected revenue of $\bar{\Gamma}$ when bidders behave truthfully. In Section 3, we do this for the special case in which the seller must sell the object with probability 1, and $\bar{\Gamma}$ is the standard Vickrey auction. In Section 4, we do this for the general case in which the seller can withhold the object, and $\bar{\Gamma}$ is an ACV auction (formally defined in Definition 1 ahead) with appropriately constructed bidder-specific reserve prices. These results constitute the heart of our contribution. Then, by an argument of Ausubel and Cramton (2004), truth-telling is an equilibrium of $\bar{\Gamma}$ for *any* resale procedure v , not just for \underline{v} . Thus, $\text{Rev}(\bar{\Gamma}, v)$, which is the highest expected revenue of $\bar{\Gamma}$ across all equilibria, is bounded below by \bar{R} , establishing (5). Lemma 1 then implies our main result.

3 The Must-Sell Case

We begin by considering a simpler model in which the seller must sell the object. (For example, this could be microfounded by assuming the seller has a prohibitively high cost of keeping the object.) To adapt the robust problem (3) to this setting, just redefine the objective $\text{Rev}(\Gamma, v)$ as the supremum of expected revenue over those Bayes-Nash equilibria in which the object is sold with probability 1, and $\text{Rev}(\Gamma, v) = -\infty$ if no such equilibrium exists.

In the absence of resale, if bidders' values are drawn from different distributions, the optimal must-sell auction misallocates the object towards weaker bidders, with the goal of reducing the information rents of stronger bidders

¹¹Note that the Revelation Principle would not generally apply to the robust maximization problem, since in a given auction, truth-telling may be an equilibrium for some resale procedures but not for others. This is why the robust problem is formulated using indirect mechanisms.

(Myerson (1981), McAfee and McMillan (1989)). As we shall see, this benefit of misallocation can be hampered by resale.

We guess a “worst-case” resale procedure to be the “Bertrand game” in which the auction’s losers make competing price offers to acquire the object from the winner, who then accepts one of the offers or rejects all of them and keeps the object to himself. As a result, if the auction’s winner is not the highest-value bidder, the former sells the object to the latter at price $\theta_{(2)}$, the second-highest value of all bidders.¹² Thus, bidder i ’s post-resale payoff from auction allocation x in state θ (exclusive of payments made in the auction) can be written as

$$\underline{v}_i(x; \theta) = \max\{\theta_i, \theta_{(2)}\} \cdot x_i + \max\{0, \theta_i - \theta_{(2)}\} \cdot \sum_{j \neq i} x_j. \quad (8)$$

Note that this resale procedure \underline{v} is efficient (i.e., satisfies (1) with equality) and individually rational (satisfies (2)).

Resale procedure (8) is a natural guess for the worst case because it makes the highest-value bidder a residual claimant for surplus, thus allowing him to capture information rents even if the object is allocated to another bidder. More specifically, in any must-sell auction, any bidder i can, by letting another bidder win and then buying from him if possible, assure himself an expected payoff $\mathbb{E}_{\tilde{\theta}} [\max\{\tilde{\theta}_i - \tilde{\theta}_{(2)}, 0\}]$. This payoff is bidder i ’s expected “marginal contribution” (that is, the total surplus available, minus the surplus that would be achievable if i were absent), and coincides with his expected payoff in the Vickrey auction (i.e. second-price auction) with no reserve price. If the seller cannot avoid conceding at least this much surplus to the bidders, she cannot do any better than using the Vickrey auction with no reserve price.

The above argument is not complete because it is not clear that bidder i can ensure that another bidder wins while avoiding paying anything to the auctioneer. For example, if bidder i bids exactly 0, the auctioneer might withhold the good (since she only needs to sell with probability 1), or might

¹²Specifically, this is the unique outcome arising in a subgame-perfect equilibrium in undominated strategies in the full-information Bertrand game. This is not the only worst-case resale procedure: for example, the argument below would also work if the highest-value bidder, when he fails to win the object, were to buy it back by making a take-it-or-leave-it offer to the auction’s winner. On the other hand, the results below do rely on the extreme division of bargaining power in these games. For example, if there are two bidders and bidder 1 captures a fixed share $\alpha \in (0, 1)$ of resale surplus, we can have an example where the seller can do better than the Vickrey auction.

sell it to bidder i but charge him a positive price, letting him recoup it in resale. We can address these issues by noting that since the auction must sell with probability 1, bidder i could ensure that the good is sold with probability 1 (to either another bidder or himself) by making a bid $\hat{\theta}_i$ that is arbitrarily small. Then, we can show that the individual rationality constraint of type $\hat{\theta}_i$ and resale procedure (8) guarantee bidder i an expected payoff close to his expected Vickrey payoff, which implies the result.

Theorem 1. *Under resale procedure (8), equilibrium expected revenue in any auction that sells with probability 1 does not exceed the expected truthtelling revenue in the Vickrey auction.*

This establishes condition (4) above for resale procedure (8). As for condition (5), it follows from the observation that truthful bidding is an ex post equilibrium of the Vickrey auction for *any* resale procedure satisfying (1)-(2) (and then no resale occurs, since the highest-value bidder has already won). While this will follow from Theorem 2 below, the informal argument is as follows. Since the usual arguments show that a bidder can never benefit by deviating from truthful bidding in a way that does not involve resale, we just need to check that there are no profitable deviations that involve resale either. If bidder i makes a downward deviation that causes him to lose, he cannot buy the object in resale unless he pays at least the winner's value, $\theta_{(2)}$, which is what he would have paid to win the object in the auction anyway. And if he makes an upward deviation that causes him to win, he cannot resell the object for more than the highest value, $\theta_{(1)}$, 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.

With (4) and (5) thus established, Lemma 1 implies

Corollary 1. *The Vickrey auction solves the robust revenue maximization problem (3) for a seller that must sell the object with probability 1.*

4 The Can-Keep Case

We now turn to the optimal auction when the seller can withhold the object, which is the version of the model more widely studied in the literature. As in the previous section, we conjecture that a worst-case resale procedure is given by (8), and show that given this procedure, an appropriately-designed ACV auction is optimal.

4.1 Regularity Assumptions on Distributions

The main result of this section will use the following assumptions on bidders' value distributions. (However, many of the intermediate steps will not rely on all of the assumptions, and so we will impose them only when needed.)

- A1.** For each bidder i , the hazard rate $f_i(\theta) / [1 - F_i(\theta)]$ is nondecreasing.
- A2.** For each bidder i , the reverse hazard rate $f_i(\theta) / F_i(\theta)$ is nonincreasing.
- A3.** The virtual value functions $\nu_i(\theta) = \theta - (1 - F_i(\theta)) / f_i(\theta)$ satisfy $\nu_1(\theta) \leq \dots \leq \nu_n(\theta)$ for each $\theta \in [0, 1]$.

(A1) implies, in particular, that each distribution F_i is *regular*, i.e., each function $\nu_i(\theta)$ is increasing. The value $r_i = \nu_i^{-1}(0)$ is then uniquely defined; it is the optimal price for selling to bidder i alone, and is also the optimal reserve price for bidder i in Myerson's optimal auction. Note also that both (A1) and (A2) are satisfied, in particular, when the density f_i is log-concave, which is satisfied by many standard distributions (Bagnoli and Bergstrom 2005).

Assumption (A3) is equivalent to requiring that the distributions F_i be ordered from "stronger" to "weaker" under the hazard rate ordering. In the absence of resale, it is a sufficient condition (and, for two-bidder must-sell auctions, a necessary condition) for the optimal auction to always discriminate against the stronger bidders (McAfee and McMillan 1989, Theorem 3). In particular, (A3) ensures that the optimal bidder-specific reserve prices would satisfy $r_1 \geq \dots \geq r_n$.

While assumptions (A1)-(A3) are not unreasonable and are familiar from the literature on optimal auctions, they are not the weakest possible to ensure the optimality of ACV auctions. In Section 5 we discuss the possibility of relaxing the assumptions and the challenges that arise.

4.2 Virtual Surplus and the Relaxed Problem

We begin with the standard approach of using first-order incentive compatibility constraints to substitute out the payment functions. To this end, for any proposed (direct) auction (χ, ψ) , let

$$U_i(\theta_i) = \mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \psi_i(\theta_i, \tilde{\theta}_{-i})] \quad (9)$$

be the interim expected payoff enjoyed by i when his type is θ_i , and note that by the standard envelope-theorem argument, incentive compatibility (6) implies that U_i is absolutely continuous and its derivative is given almost everywhere by¹³

$$U'_i(\theta_i) = \mathbb{E}_{\tilde{\theta}_{-i}}[v'_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})]. \quad (10)$$

Here v'_i denotes the derivative of $v_i(x; \theta)$ with respect to θ_i , which is defined when θ_i does not tie with another bidder, and then takes the form $v'_i(x; \theta) = \mathbf{1}_{\theta_i=\theta_{(2)}} \cdot x_i + \mathbf{1}_{\theta_i=\theta_{(1)}}$.

As for participation constraints, as usual we consider (7) only for type 0, for which it takes the form $U_i(0) \geq 0$. If only this participation constraint is imposed, it will be optimal to satisfy it with equality, and so integrating (10) and using (9) yields the following expression for the interim expected payment of bidder i of type θ_i :

$$\mathbb{E}_{\tilde{\theta}_{-i}}[\psi_i(\theta_i, \tilde{\theta}_{-i})] = \mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})] - \int_0^{\theta_i} \mathbb{E}_{\tilde{\theta}_{-i}}[v'_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \hat{\theta}_i, \tilde{\theta}_{-i})] d\hat{\theta}_i. \quad (11)$$

The usual integration by parts then allows us to rewrite the auction's expected revenue in terms of the allocation rule as

$$\mathbb{E}_{\tilde{\theta}} \left[\sum_i \psi_i(\tilde{\theta}) \right] = \mathbb{E}_{\tilde{\theta}} \left[\sum_i v_i(\chi(\tilde{\theta}); \tilde{\theta}) - \sum_i \frac{1 - F_i(\tilde{\theta}_i)}{f_i(\tilde{\theta}_i)} v'_i(\chi(\tilde{\theta}); \tilde{\theta}) \right]. \quad (12)$$

The standard relaxed problem maximizes (12) over all allocation rules χ , ignoring (6), as well as (7) for types other than 0. This can be done by maximizing the virtual surplus (the expression inside brackets) for each profile θ separately. At any profile θ without ties, if we allocate to the highest-value bidder i^* , the virtual surplus is

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

¹³Specifically, for any given report $\hat{\theta}_i$, $v_i(\chi(\hat{\theta}_i, \theta_{-i}); \theta_i, \theta_{-i})$ is Lipschitz continuous in θ_i and is differentiable in θ_i except when there are ties in values. This implies that $\mathbb{E}_{\tilde{\theta}_{-i}}[v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})]$ is Lipschitz continuous and differentiable in θ_i , and allows the application of Milgrom and Segal's (2002) Corollary 1 to establish absolute continuity of U_i and (10).

If we allocate to the second-highest bidder j , then the \underline{v}'_i terms are 1 for both $i = i^*$ and $i = j$, so the virtual surplus is

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

Finally, if we allocate to a bidder other than i^* or j , the virtual surplus is the same as from allocating to i^* . Thus, misallocation could only increase information rents: allocating to a bidder other than i^* still concedes information rents to higher types of i^* (who can buy the good cheaply in resale), and may leave information rents to the inefficient winner as well. So the seller cannot do better than allocating to the high-value bidder i^* , which leaves rents to him only. Furthermore, we should allocate to the highest-value bidder i^* if and only if his virtual value is positive, which is equivalent to $\theta_i > r_i = \nu_i^{-1}(0)$ provided that his virtual value function ν_i crosses zero just once (which is ensured under assumption (A1)). In particular, if there are two bidders and each ν_i crosses zero once, the allocation rule solving the relaxed problem is as shown in Figure 1, and the solution is unique up to measure-zero sets.

Given the allocation rule, transfers consistent with (11) can be obtained, in particular, by charging auction winner i^* the “threshold price” $\max\{r_{i^*}, \theta_{(2)}\}$, which is his lowest value for which he would have won the good given the values of others; and charging losers nothing. The resulting mechanism is known as the “Vickrey auction with lazy reserves” (Dhangwatnotai et al., 2010), and it is dominant-strategy incentive compatible in the absence of resale.

However, with our resale procedure \underline{v} , the solution to the relaxed problem violates global incentive compatibility constraints (6) unless all bidders have the same optimal reserve r_i . For example, consider the case of two bidders with $r_1 > r_2$, and consider bidder 1 of type $\theta_1 \in (r_2, r_1)$ deviating to report $\hat{\theta}_1 < \theta_1$, as illustrated with the horizontal arrow in Figure 1. This deviation affects the auction’s outcome when $\theta_2 \in (\max\{\hat{\theta}_1, r_2\}, \theta_1)$, and in these cases it changes the outcome from leaving the object unsold to giving it to bidder 2, allowing 1 to profitably buy it back in resale. Thus, in order to find the correct solution, we need to consider non-local incentive constraints.¹⁴

¹⁴Note that we could deter all local misreports within distance $\varepsilon > 0$ by perturbing the relaxed solution to withhold the good whenever $|\theta_2 - \theta_1| < \varepsilon$, but the perturbed mechanism would still be vulnerable to global deviations of the form described above. This is in contrast to “ironing” in the standard screening setting, where global incentive




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.

4.3 Ausubel-Cramton-Vickrey Auctions

Intuitively, to avoid the incentive to underbid to cede the object and then buy it back, we might use an allocation rule in which a lower bid never causes the object to be sold. (Ausubel and Cramton (1999) call this property *monotonicity in aggregate*.) In the two-bidder 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’s virtual value in the filled-in triangle is negative, the seller would rather shrink the size of this triangle by raising the reserve price for bidder 2 above r_2 , even though doing so also means missing out on profitable sales to bidder 2. The optimal reserve price for bidder 2 trades off these two effects. The resulting allocation rule is shown in Figure 2.

compatibility is implied by local (first- and second-order) incentive constraints (Archer and Kleinberg, 2014; Carroll, 2012).




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

This auction belongs to the following class of auctions, introduced by Ausubel and Cramton (1999, 2004):

Definition 1. An Ausubel-Cramton-Vickrey (ACV) auction with reserve prices $p_1, \dots, p_n \in [0, 1]$ is the direct revelation mechanism described as follows:

- Allocation rule:

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

where $i^*(\theta) \in \arg \max_i \theta_i$ (with arbitrary tie-breaking.)

- Payments:

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

and non-participation message \emptyset is treated the same as a report of 0.

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

A winner’s payment in the ACV auction is his “threshold price” — the minimal bid that would have allowed him to win. By standard arguments, this ensures that the auction is strategy-proof without resale.¹⁵ More important for us is that (in contrast to the solution to our relaxed problem) in an ACV auction, bidders are incentivized to bid truthfully even with resale. Specifically, truthful bidding is an ex post equilibrium, for any resale procedure and for any profile of values:

Theorem 2. *Consider an ACV auction with reserve prices p_1, \dots, p_n . For any resale payoff functions v_1, \dots, v_n satisfying (1)–(2), it is an ex post equilibrium for all bidders to participate and report their true values (after which no resale occurs). Formally, for any values $\theta_1, \dots, \theta_n$, and any possible deviation $\hat{\theta}_i \in [0, 1] \cup \{\emptyset\}$ for any bidder i ,*

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

The proof essentially follows the informal argument given in Section 3 for the Vickrey auction without reserves, but also using the fact that in an ACV auction, a downward deviation would never cause the object to be sold. While this result was proved by Ausubel and Cramton (2004), for completeness we give a proof in Appendix C. Also, in that appendix we develop additional formalism for a broader class of resale procedures that drops the assumption that values are revealed; thus, bargaining may take place under asymmetric information. The proof shows that the result holds for this broader class of resale procedures.

We note in passing that this theorem only gives us ex post equilibrium, not dominant strategies as we would have in Vickrey auctions without re-

¹⁵An ACV auction can also be indirectly implemented as a “deferred acceptance” clock auction of the kind described by Milgrom and Segal (2017). In this implementation, the auction offers the same ascending price to all bidders, letting them either accept the price or exit at any point, and stopping when both (i) there is a single bidder who is still accepting the current price, and (ii) at least one bidder (not necessarily the one who is still bidding) has ever accepted a price above his reserve price. (Note, however, that one advantage of the clock auction format – its obvious strategy-proofness (Li, 2017) – does not hold when resale is possible.)

sale.¹⁶ This is to be expected since our setting is one of interdependent values (compare Perry and Reny (2002) or Chung and Ely (2006)).

4.4 Construction of Optimal Reserve Prices

In the two-bidder ACV auction illustrated in Figure 2, bidder 1's price p_1 can be optimized without regard to bidder 2 (since it only matters when bidder 2 bids below $p_2 < p_1$, and therefore does not win), hence it is optimal to set $p_1 = r_1$. On the other hand, the optimal price for bidder 2 is raised above r_2 to increase the expected revenue on bidder 1. Extending this idea to n bidders, we construct a weakly decreasing sequence p_1, \dots, p_n of reserve prices and a corresponding sequence $\bar{R}_1, \dots, \bar{R}_n$ of revenue levels on the first k bidders recursively, initializing $p_0 = 1$ and $\bar{R}_0 = 0$ and letting for each $k \geq 1$,

$$\bar{R}_k = R_k(p_k) + F_k(p_k) [\bar{R}_{k-1} - R_{k-1}(p_k)] \quad (13)$$

$$= \max_{p \in [0, p_{k-1}]} \{R_k(p) + F_k(p) [\bar{R}_{k-1} - R_{k-1}(p)]\}, \quad (14)$$

where $R_k(p)$ is the expected revenue from the (symmetric) Vickrey auction on the first k bidders facing the same reserve price p , with $R_0(p) \equiv 0$.

Formula (13) gives an inductive construction of the expected revenue \bar{R}_k from the ACV auction on bidders $1, \dots, k$ with reserves p_1, \dots, p_k , assuming the bidders bid their true values. To see this, compare this ACV auction to the symmetric Vickrey auction with reserve p_k on the same bidders. Notice that the only states θ in which the two auctions could yield different revenues are those in which bidder k 's value θ_k is below p_k , which happens with probability $F_k(p_k)$. Conditional on any such value of θ_k , the former auction reduces to the ACV auction on the first $k - 1$ bidders with reserves p_1, \dots, p_{k-1} , yielding expected revenue \bar{R}_{k-1} , while the latter auction reduces to the Vickrey auction with reserve p_k on the first $k - 1$ bidders, yielding expected revenue $R_{k-1}(p_k)$.

Condition (14) requires that bidder k 's reserve price p_k maximize revenue on the strongest k bidders taking as given the stronger bidders' reserves

¹⁶To 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 (8), 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.

p_1, \dots, p_{k-1} and the constraint $p_k \leq p_{k-1}$. This condition does not immediately imply full optimality of the reserve prices, since we do not know whether the constraint binds, or how the choice of p_k in turn constrains the revenues on the weaker bidders. However, we will show a stronger result – that the resulting auction is indeed optimal, and not just among ACV auctions but among *all* possible auctions.¹⁷

4.5 Optimality of ACV Auction

We now come to our main theorem.

Theorem 3. *Under assumptions (A1)-(A3), the formulas (13)-(14) uniquely define reserve prices p_1, \dots, p_n . The ACV auction with these reserve prices is an optimal auction given resale procedure (8), and the resulting revenue is \bar{R}_n .*

This implies a solution to our original problem:

Corollary 2. *Under assumptions (A1)-(A3), the ACV auction with reserve prices p_1, \dots, p_n uniquely defined by (13)-(14) solves the robust revenue maximization problem (3).*

We sketch the main steps of the proof of Theorem 3. Since Theorem 2 showed that the ACV auction is feasible (i.e. satisfies the constraints (6)–(7)), and its revenue is indeed \bar{R}_n by (13), we focus on showing that \bar{R}_n is an upper bound for revenue in any auction.

As discussed in Subsection 4.2, we need to make active use of the non-local incentive constraints, which, using the formula (11) for transfers, can

¹⁷ The working paper of Ausubel and Cramton (1999) formulated the problem of optimal auction design, assuming no misallocation and monotonicity in the aggregate, motivating both properties informally by incentive compatibility under perfect resale. They also claimed (without proof) that, in the two-bidder case, the problem is solved by an ACV auction. In contrast, we *derive* no misallocation and monotonicity in aggregate as properties of a solution to the seller’s maximin problem under regularity assumptions (A1)-(A3). As argued in the Introduction, the optimal auction need not have these properties when resale is perfect but not of the worst-case form. Also, in Appendix G we show that the robust-revenue-maximizing auction need not have these properties without regularity assumptions.

be rewritten entirely in terms of the allocation rule:

$$\int_{\hat{\theta}_i}^{\theta_i} \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\tau_i, \tilde{\theta}_{-i}); \tau_i, \tilde{\theta}_{-i})] d\tau_i - \mathbb{E}_{\tilde{\theta}_{-i}} [v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - v_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \hat{\theta}_i, \tilde{\theta}_{-i})] \geq 0. \quad (15)$$

We account for these constraints by introducing Lagrange multipliers on them. Since there is a double continuum of such constraints for each bidder i , indexed by $(\theta_i, \hat{\theta}_i)$, the Lagrange multipliers (weights) on them are described with some appropriately constructed nonnegative measure \mathcal{M}_i on $[0, 1] \times [0, 1]$. The Lagrangian adds terms to the objective function (12) to penalize violations of the constraints:

$$\mathbb{E}_{\tilde{\theta}} \left[\sum_i v_i(\chi(\tilde{\theta}); \tilde{\theta}) - \sum_i \frac{1 - F_i(\tilde{\theta}_i)}{f_i(\tilde{\theta}_i)} v'_i(\chi(\tilde{\theta}); \tilde{\theta}) \right] + \sum_i \iint S_i(\theta_i, \hat{\theta}_i; \chi) d\mathcal{M}_i(\theta_i, \hat{\theta}_i), \quad (16)$$

where $S_i(\theta_i, \hat{\theta}_i; \chi)$ is the left side of (15) – the *slack* in bidder i 's incentive constraint from type θ_i to type $\hat{\theta}_i$.

We show Theorem 3 by constructing measures $\mathcal{M} = (\mathcal{M}_1, \dots, \mathcal{M}_n)$ such that $\langle \bar{\chi}, \mathcal{M} \rangle$ satisfies the following conditions, where $\bar{\chi}$ is the allocation rule from the ACV auction that is our candidate optimum.

- (a) *Optimization:* Allocation rule $\bar{\chi}$ maximizes Lagrangian (16) given measures \mathcal{M} ;
- (b) *Complementary slackness:* the support of \mathcal{M} is confined to constraints (15) that hold with equality at $\bar{\chi}$, i.e., $S_i(\theta_i, \hat{\theta}_i; \bar{\chi}) = 0$.

To see that these conditions imply the result, note that since \mathcal{M} is a nonnegative measure, the expected revenue from any incentive-compatible and individually rational direct auction with allocation rule χ cannot exceed the value of the Lagrangian (16) at $\langle \chi, \mathcal{M} \rangle$, which by (a) does not exceed the value of the Lagrangian at $\langle \bar{\chi}, \mathcal{M} \rangle$, which by (b) equals the expected revenue from $\bar{\chi}$.

We construct Lagrange multiplier measures \mathcal{M}_i to penalize deviations that would be profitable in the Vickrey auction with lazy reserves $p_1 \geq \dots \geq p_n$, of the kind that emerged as a solution to the relaxed problem. In this auction, as we argued, bidder $i < n$ with value between p_n and p_i cannot win the object outright, and so would want to underbid to increase the probability that the object is sold to another bidder, from whom he can

then buy in resale. Thus, we let the support of $\mathcal{M}_i(\theta_i, \hat{\theta}_i)$ for bidder $i < n$ be the trapezoid described by $p_n \leq \theta_i \leq p_i$ and $\hat{\theta}_i \leq \theta_i$, and set $\mathcal{M}_n \equiv 0$ for the weakest bidder. This ensures complementary slackness in the candidate-solution ACV auction: Bidder i 's deviation in the support can only affect the auction's outcome if i wins under truth-telling but his deviation cedes the good to the next-highest bidder. Then, after the deviation, i buys in resale at price $\theta_{(2)}$, which is the same price at which he would have won the auction by bidding truthfully. So, i is indifferent between truthful bidding and the deviation.

We specifically aim to construct measures \mathcal{M}_i of the following form: for appropriately chosen one-dimensional measures Λ_i and $\hat{\Lambda}_i$ with supports $[p_n, p_i]$ and $[0, p_i]$ respectively, take their product (a two-dimensional measure with support $[p_n, p_i] \times [0, p_i]$), and then restrict to the lower half-plane $\hat{\theta}_i \leq \theta_i$; this restriction is \mathcal{M}_i . We denote the two component measures' distribution functions by $\Lambda_i(\theta_i)$ and $\hat{\Lambda}_i(\hat{\theta}_i)$, respectively.

We then construct distribution functions Λ_i and $\hat{\Lambda}_i$ to satisfy the Lagrangian maximization condition (a). Once these are constructed, we verify (a) by expressing Lagrangian (16), which is a linear functional of the allocation rule χ , in the form

$$\mathbb{E}_{\tilde{\theta}} \left[\sum_i \mu_i(\tilde{\theta}) \chi_i(\tilde{\theta}) \right]. \quad (17)$$

The coefficient $\mu_i(\theta)$, which we refer to as the ‘‘modified virtual value’’ of bidder i , combines the ordinary virtual value and the terms coming from incentive compatibility constraints (15) weighted by the measures \mathcal{M} . For the measures \mathcal{M} we construct, we derive an explicit formula for $\mu_i(\theta)$ in the appendix (see (31)). Condition (a) then amounts to requiring that with probability 1, the candidate allocation rule $\bar{\chi}$ allocate the object to a bidder with the highest modified virtual value provided it is positive, and not allocate at all if all bidders have negative modified virtual values.

To illustrate how condition (a) serves to pin down measures Λ_i and $\hat{\Lambda}_i$, focus for simplicity on the two-bidder case, where we only need to construct a measure for bidder 1. In order for the proposed ACV auction to be optimal, we need two properties to be satisfied:

1. When $\theta_1 < p_2$ and $\theta_2 = p_2$, bidder 2's modified virtual value should be 0.

2. When $\theta_1 \in (p_2, p_1)$ and $\theta_2 = p_2$, bidder 1's modified virtual value should be 0.

Each property is needed for condition (a) because a slight increase in θ_2 should lead the auctioneer to sell the good to the highest-value bidder (whose modified virtual value must then be positive), while a slight decrease in θ_2 should lead the auctioneer to withhold the good (making the modified virtual value negative).

Property #1 pins down the density of $\hat{\Lambda}_1$ on $[0, p_2]$. It turns out to be proportional to $f_1(\theta_1)$ (we normalize the proportionality factor to 1). This occurs because selling to bidder 2 in any such state (θ_1, p_2) has two effects on the Lagrangian: the direct effect on revenue $\nu_2(\theta_2)$, and the effect of tightening the binding non-local incentive constraints of types of bidder 1 above p_2 , who could deviate to θ_1 and then buy the good in resale. The first effect appears with a weight proportional to the probability of state (θ_1, p_2) , while the second has weight proportional to $\hat{\Lambda}_1(\theta_1)$; these two weights must be proportional in order for these terms to cancel. Note that Property #1 does not pin down the density of $\hat{\Lambda}_1$ on the rest of its domain, $(p_2, p_1]$, but a natural guess is to take the density as $f_1(\theta_1)$ here as well.

Property #2 requires that the weighted slack in non-local incentive constraints of bidder 1's types above θ_1 induced by selling to $\theta_1 \in (p_2, p_1)$ exactly offset bidder 1's negative virtual value $\nu_1(\theta_1)$. Once $\hat{\Lambda}_1$ is fixed, this condition pins down measure Λ_1 on its domain $[p_2, p_1]$.

With more than two bidders, we need to similarly construct measures for each bidder $i < n$. We again set $\hat{\Lambda}_i$ to have the same density as i 's type distribution, f_i , restricted to the support $[0, p_i]$. As for measures Λ_i , they are pinned down by requiring the modified virtual value of the highest bidder i^* to be exactly zero so that selling to him could be made contingent on whether some other bidder j has beaten his reserve price, similarly to the argument above. The explicit formula ((30) in the appendix) depends on an auxiliary function H , constructed below. We show that under assumptions (A1)–(A3), the constructed functions Λ_i are increasing, so that each \mathcal{M}_i is in fact a nonnegative measure.

Finally, we verify that with these measures, condition (a) holds, i.e., the ACV auction maximizes the modified virtual surplus (expressed as (17)) at *every* type profile θ without ties, and not just on the regions used in the indifference conditions nailing down the measures. That is, we show that the highest-value bidder always has the highest modified virtual value, and

it is nonnegative when some bidder's value is above his reserve price and nonpositive otherwise.

Now, for the Corollary, letting $\bar{\Gamma}$ be the described ACV auction, Theorem 3 establishes bound (4) with $\bar{R} = \bar{R}_n$, while bound (5) follows from Theorem 2. Thus, by Lemma 1, $\bar{\Gamma}$ solves the robust revenue maximization problem (and Bertrand resale procedure \underline{v} solves the worst-case resale problem).

4.6 Calculation of Optimal Reserves

We now indicate briefly how one might calculate the reserve prices p_k in an application.

An interior solution p_k to maximization problem (14) must satisfy the following first-order condition:

$$\nu_k(p_k) \prod_{j < k} F_j(p_k) = \bar{R}_{k-1} - R_{k-1}(p_k). \quad (18)$$

Intuitively, increasing p_k by ε only matters when $\theta_j < p_j$ for all $j > k$, and in that case it has two first-order effects: (i) conditional on $\theta_k \in (p_k, p_k + \varepsilon)$, the change affects 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 \bar{R}_{k-1} (obtained when bidder k does not beat his reserve price); and (ii) the change also affects the expected revenue from bidder k when all stronger bidders' values are below p_k , changing it from $p_k(1 - F_k(p_k))$ to $(p_k + \varepsilon)(1 - F_k(p_k + \varepsilon))$. The first-order condition balances those two effects.

To avoid having to directly calculate \bar{R}_{k-1} and $R_{k-1}(p_k)$, we define function H on $p \in [p_n, 1)$ by

$$H(p) = \frac{\bar{R}_{k-1} - R_{k-1}(p)}{\prod_{j < k} F_j(p)} \text{ for } p \in [p_k, p_{k-1}), k = 1, \dots, n, \quad (19)$$

and thereby rewrite first-order conditions in the form

$$\nu_k(p_k) = H(p_k). \quad (20)$$

Intuitively, $H(p)$ reflects the shadow price of selling to a bidder with value p on all bidders' non-local incentive constraints, and it is used in constructing the Lagrange weights on these constraints in the proof of Theorem 3.

In Lemma 2 in Appendix D we show that function H is continuous, differentiable, and satisfies the following differential equation:

$$\frac{d}{dp} (H(p) \prod_{j < k} F_j(p)) = (\prod_{j < k} F_j(p)) \sum_{j < k} \frac{f_j(p) \nu_j(p)}{F_j(p)}, \text{ for } p \in [p_k, p_{k-1}], k \geq 2. \quad (21)$$

Furthermore, we establish that the first-order conditions (20) must be satisfied regardless of whether the solutions p_k to (14) are interior.

Differential equation (21) together with conditions (20) and the boundary condition $H(p_1) = 0$ yield the following iterative construction of function H and the reserve prices: First, $p_1 = r_1$ by the boundary condition and (20) for $k = 1$. Then for each $k \geq 2$, given p_{k-1} , integrating (21) yields an explicit expression for $H(p) \prod_{j < k} F_j(p)$ — and therefore for $H(p)$ — on the interval $p \in [p_k, p_{k-1}]$, whose unknown left endpoint p_k is then identified as the unique solution to equation (20).

In Appendix F, we apply this method to an example where each bidder i 's value is distributed uniformly on $[0, a_i]$, with the upper limits satisfying $a_1 \geq \dots \geq a_n$. We show that (20) then takes the form of a k th-degree polynomial equation to compute the reserve price p_k . For example, $p_1 = r_1 = a_1/2$, and then p_2 is given by a quadratic equation, whose positive root is

$$p_2 = \frac{1}{2} \left[\sqrt{a_1^2 + (a_1 - a_2)^2} - (a_1 - a_2) \right]. \quad (22)$$

With bidders having non-identical supports, the reserve p_k may exceed a_k , in which case bidder k , as well as all the weaker bidders, are completely excluded from the auction. In the uniform example, bidder 2 is excluded if $a_2 \leq a_1/4$. Note that as long as all bidders' supports are overlapping, such exclusion cannot occur in the Myerson optimal auction, nor even in the relaxed solution from Subsection 4.2. It is the non-local incentive constraints involving stronger bidders buying from weaker bidders in resale that sometimes make it optimal to exclude weak bidders (and more generally, make the optimal auction take the ACV form).

5 Discussion

We consider here in more detail the role of some of our modeling assumptions, and some possible alternative models.

5.1 Robustness to Information Revelation

One might be concerned with an asymmetry in our model: we have allowed information to be revealed exogenously after the auction, but assumed it is not revealed before the auction.

A common perspective on optimal auction design is that ideally it should not make assumptions on bidders' information before the auction either. Indeed, without resale, Myerson's (1981) optimal auction, though derived under the assumption that bidders share the auctioneer's prior, can actually be implemented in dominant strategies, making it robust to bidders' beliefs about each other. Similarly, in our setting with resale, we could imagine a seller who is concerned about bidders learning about each other both before *and* after the auction, and desires robustness to information arising at either stage. We have shown that an ACV auction is optimal when bidders share the designer's prior before the auction and learn each other's values after it, following the worst-case resale procedure. But since truthful bidding is an ex post equilibrium in the ACV auction, it satisfies this stronger form of robustness. Thus, we get robustness to pre-auction information revelation "for free," without needing to require it explicitly in the seller's problem.

One could alternatively try to restore symmetry to the model by assuming that exogenous information revelation cannot happen either before or after the auction. In this case, our basic conclusion — that robustness to resale implies the auctioneer should not misallocate — no longer holds. For an example, it suffices to consider the must-sell case, where we need not worry about (A1)–(A3). Suppose that there are two bidders, with bidder 1's value equal to 1 for sure, and bidder 2's value being $1/3$ or $2/3$ with probability $1/2$ each. Consider the mechanism that offers the object to bidder 1 at a fixed price $p = 2/3 - \varepsilon$, and if he rejects, gives it to bidder 2 for free. Note that if bidder 1 rejects, he does not learn anything about bidder 2's value, so the highest expected payoff he could hope to get in resale is $1/3$ (which he could obtain if he has all the bargaining power, by offering either a price of $1/3$ or $2/3$ to bidder 2). Thus, bidder 1 would prefer to buy in the mechanism. This gives a revenue of $2/3 - \varepsilon$, which exceeds the Vickrey auction's expected revenue of $1/2$. This example involves discrete distributions, but it can easily be perturbed to a continuous distribution with full support on $[0, 1]$, and then any efficient must-sell auction is revenue-equivalent to Vickrey; and still, the seller would prefer our alternative auction over Vickrey, regardless of the resale procedure. So the optimal must-sell auction must misallocate with

positive probability, contradicting the conclusion of Theorem 1.¹⁸

In summary, our main conclusion — that robustness to resale calls for no misallocation — depends on allowing information to be revealed after the auction, but does not depend on whether information can be revealed before the auction.

5.2 Robustness to Value Distributions

Another seeming asymmetry arises between the modeling assumptions of Bayesian priors over bidders’ values and complete ignorance of the resale procedure. We may instead consider a more symmetric model in which for each bidder i , there is a family of possible distributions $\mathcal{F}_i \subseteq \Delta([0, 1])$, rather than a single distribution, and the designer’s robust revenue maximization problem is

$$\max_{\Gamma} \left(\inf_{v, F_1 \in \mathcal{F}_1, \dots, F_n \in \mathcal{F}_n} \text{Rev}(\Gamma, v, F_1, \dots, F_n) \right)$$

where $\text{Rev}(\Gamma, v, F_1, \dots, F_n)$ stands for the supremum expected revenue in a Bayes-Nash equilibrium of mechanism Γ with resale procedure v and distributions F_1, \dots, F_n .¹⁹ For this problem, our result about optimality of ACV auctions carries through. Specifically:

Proposition 1. *Assume that for each i , there is a “lowest” distribution $F_i \in \mathcal{F}_i$ that is first-order stochastically dominated by every other $F'_i \in \mathcal{F}_i$. Assume that F_1, \dots, F_n have continuous densities and satisfy assumptions (A1)–(A3). Then the robust revenue maximization problem is solved by the ACV auction with reserves p_1, \dots, p_n constructed in (13)–(14).*

We prove the proposition by showing that the lowest distributions combined with resale procedure (8) constitute a worst case for the designer.

¹⁸Solving for the *optimal* auction in worst case over all resale procedures without information revelation is a difficult open question. We do know from Calzolari and Pavan (2006) that Myerson’s optimal biased allocation is generically not implementable even for a given resale procedure.

¹⁹An alternative way to avoid the asymmetry would be by interpreting our analysis as applying to a fully Bayesian auctioneer who believes that resale takes the particular worst-case form \underline{v} , joining the literature on auction design for a given known resale procedure (cited in the Introduction). However, our preferred interpretation is that the auctioneer does not presume any given specific resale procedure, instead being completely ignorant about it.

For this we need to verify that the optimal ACV auction for distributions F_1, \dots, F_n would not yield lower expected revenue when bidders' values are drawn from higher distributions. While this kind of revenue monotonicity does *not* hold for an arbitrary ACV auction, we show that it holds for the optimal ACV auction given distributions F_1, \dots, F_n .

The above result still requires the designer to know lower bounds on the possible distributions in order to choose the reserves. Can we say anything with no knowledge of the distributions at all? Our maxmin problem becomes uninteresting: the worst case is that each buyer just has value 0 with probability 1, so there is no positive revenue guarantee. However, we can give two arguments to suggest that optimization of reserve prices has little benefit when there are sufficiently many bidders, and in this case the Vickrey auction with no reserves is a reasonable prior-free choice despite possible asymmetries among bidders.²⁰

First, we show that under uncertainty about resale, setting optimal reserve prices is less valuable than adding just one bidder who is at least as strong as the existing bidders:

Proposition 2. *Suppose the value distributions of the n bidders are regular. Under our worst-case resale procedure, any auction has a lower expected revenue than the Vickrey auction with one added bidder $n + 1$, provided that the added bidder's value distribution first-order stochastically dominates the distributions of all other bidders.*

The proposition extends the classic result of Bulow and Klemperer (1996) to asymmetric settings with uncertainty about resale. In contrast, as noted by Hartline and Roughgarden (2009), the result does not extend to asymmetric settings *without resale*, in which Myerson's optimal biased auction could yield substantially higher revenue than the Vickrey auction with duplicate bidders.²¹

Second, we show that the difference in revenue between the optimal auction under worst-case resale and the Vickrey auction is very small when the number of bidders n is large: it shrinks to zero exponentially fast as n grows.

²⁰In the propositions below we do not impose assumptions (A1)-(A3), because they are proved by showing that Vickrey does well even compared to the solution of the relaxed problem, which relied only on regularity (a weakening of (A1)).

²¹Namely, in their two-bidder Example 4.6, adding a third bidder with value c.d.f. $F_3(\theta) = \min\{F_1(\theta), F_2(\theta)\}$ and using the Vickrey auction on the three bidders would only yield $\frac{3}{4}$ of the revenue of the optimal biased auction on the original two bidders.

(We require value distributions to satisfy a particular bound uniformly in n , e.g., they could be drawn from some finite family.) In contrast, if resale were impossible, Myerson’s optimal auction could attain an improvement over Vickrey that shrinks only polynomially in n (even if just two different value distributions are possible). This suggests that uncertainty about resale is quantitatively more important for auction design than optimization of reserve prices.

Proposition 3.

- (a) *If all distributions are regular and $F_i(r_j) \leq c < 1$ for all i, j , then the difference between the expected revenue from the optimal auction under worst-case resale and the expected revenue from the Vickrey auction is at most $n \cdot c^{n-1}$.*
- (b) *Let F, \hat{F} be regular distributions with respective twice-differentiable densities f, \hat{f} such that $f(1), \hat{f}(1) > 0$ and $f'(1)/f(1) < \hat{f}'(1)/\hat{f}(1)$. There is a constant $c' > 0$ with the following property: if there are $n/2$ bidders drawn from distribution F and $n/2$ from distribution \hat{F} , the expected revenue from the optimal Myerson auction (without resale) exceeds that of the Vickrey auction by at least c'/n^3 .*

5.3 Relaxing Regularity

Our regularity assumptions (A1)-(A3) are not necessary for ACV auctions to be optimal. For example, without these assumptions we can state²²

Proposition 4. *If all bidders’ virtual value functions $v_i(\theta_i)$ cross zero just once and all at the same point r , the Vickrey auction with symmetric reserve price r solves the robust revenue maximization problem.*

Indeed, this auction solves the relaxed problem in Section 4.2. From this (combined with Theorem 2), the proposition follows immediately.

²²For another example, in the two-bidder case, assumption (A3) is not needed at all (it suffices to assume that $r_1 > r_2$), while assumptions (A1)-(A2) could be replaced with the non-nested assumptions that the function $R_1(p) = p(1 - F_1(p))$ is concave, and that v_2 is increasing. (The proof for this case necessitates constructing a different Lagrange multiplier measure than the one we use.)

In particular, while we have focused on settings where the Myerson optimal auction is inefficient due to bias, it can also be inefficient due to randomization: when bidders are symmetric but have non-monotone virtual value functions, the Myerson auction calls for random allocation when multiple bidders are in an ironing region. The above proposition shows that in these cases, too, robustness to resale can make it optimal to allocate efficiently.

However, we cannot dispense with distributional assumptions entirely: Appendix G describes an example with symmetric bidders where the optimal auction is not of the ACV form. We do not know of any simple, general assumptions that nest all known cases where ACV auctions are optimal.

5.4 Further possibilities

One might argue, as Börgers (2017) does, that the maxmin criterion used in (3) is too permissive. Applied to our setting, his critique would say that the ACV mechanism we have identified as maxmin optimal is weakly dominated, in the sense that there are other mechanisms that do at least as well for every resale procedure, and strictly better for some. For example, if the resale procedure is common knowledge among bidders, the auctioneer can elicit it from them and then run an optimal auction tailored to that resale procedure. However, this elicitation mechanism would not work if the resale procedure is not common knowledge, while an ACV mechanism would still remain maxmin-optimal.²³

While we have assumed the worst case with respect to resale procedures, we have nonetheless focused on the best equilibrium, implicitly assuming that the seller gets to choose the equilibrium played by the bidders (as is usual in mechanism design). However, it is known that Vickrey auctions can have many equilibria when resale is possible, including ones that are inefficient and have lower revenue than the truthful equilibrium but are preferred by all bidders (Garratt and Tröger, 2006; Garratt, Tröger, and Zheng, 2009). We expect the same for ACV auctions. It is an open question whether the ACV auction can be modified in a way that makes its expected revenue guarantee robust to equilibrium selection.²⁴

²³As Börgers (2017) notes, an undominated mechanism may not exist in this case, since in addition to any mechanism the seller could run a side betting mechanism collecting commissions on bets when bidders have heterogenous beliefs.

²⁴In the absence of resale, it is possible to achieve robustness to equilibrium selection, as well as stronger forms of bidder collusion: Laffont and Martimort (1997) and Che and

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 own values and will learn more about them after the auction (as in Haile, 2003). Presumably the optimal auctions in such models would involve resale occurring in equilibrium.

6 Summary

We began from the classic theory of optimal auctions, which suggests that it can be optimal to bias the auction in favor of weaker bidders. A rough intuition suggests that such misallocation is no longer advantageous when it can be undone by resale among the bidders. A series of simple examples shows, however, that when the auctioneer knows the resale procedure, it typically is still optimal to misallocate, even if resale is “perfect” (contrary to the intuition of Ausubel and Cramton (1999)).

We provided a model in which resale makes it optimal not to misallocate. In our model, the auctioneer does not know the resale procedure and desires revenue to be robust to this uncertainty. In the must-sell version of the model, the optimal auction is simply a Vickrey auction. In the can-keep version, the optimum (under appropriate regularity conditions) is an ACV auction: The auctioneer boosts revenue by setting reserve prices, as in the classic model, but never benefits from misallocating the good.

References

- [1] Archer, A., and R. Kleinberg (2014), “Truthful Germs are Contagious: A Local-to-Global Characterization of Truthfulness,” *Games and Economic Behavior* 86, 340–366.
- [2] Ausubel, L. M., and P. Cramton (1999), “The Optimality of Being Efficient,” working paper, Department of Economics, University of Maryland, College Park.

Kim (2006) show how this can be done with any Bayesian incentive-compatible auction, including Myerson’s (1981) optimal auction. Che and Kim (2006, footnote 34) explain why their results on collusion do not also imply robustness to resale.

- [3] Ausubel, L. M., and P. Cramton (2004), “Vickrey Auctions with Reserve Pricing,” *Economic Theory* 23(3), 493-505.
- [4] Bagnoli, M., and T. Bergstrom (2005), “Log-Concave Probability and Its Applications,” *Economic Theory* 26(2), 445–469.
- [5] Bergemann, D., B. Brooks, and S. Morris (2017a), “Selling to Intermediaries: Optimal Auction Design in a Common Value Model,” working paper, University of Chicago.
- [6] Bergemann, D., B. Brooks, and S. Morris (2017b), “Informationally Robust Optimal Auction Design,” working paper, University of Chicago.
- [7] Börgers, T. (2017), “(No) Foundations of Dominant-Strategy Mechanisms: A Comment on Chung and Ely (2007),” *Review of Economic Design* 21, 73–82.
- [8] Bulow, J., and P. Klemperer (1996), “Auctions Versus Negotiations,” *American Economic Review* 86(1), 180-194.
- [9] Bulow, J., and P. Klemperer (2002), “Prices and the Winner’s Curse,” *RAND Journal of Economics* 33(1), 1-21.
- [10] Calzolari, G., and A. Pavan (2006), “Monopoly with Resale,” *RAND Journal of Economics* 37(2), 362–375.
- [11] Carroll, G. (2012), “When are Local Incentive Constraints Sufficient?” *Econometrica* 80(2), 661–686.
- [12] Carroll, G. (2015), “Robustness and Linear Contracts,” *American Economic Review*, 105(2), 536-563.
- [13] Che, Y.-K., and J. Kim (2006), “Robustly Collusion-Proof Implementation,” *Econometrica* 74(4), 1063–1107.
- [14] Chung, K.-S., and J.C. Ely (2006), “Ex-Post Incentive Compatible Mechanism Design,” working paper, Northwestern University.
- [15] Chung, K.-S., and J.C. Ely (2007), “Foundations of Dominant Strategy Mechanisms,” *Review of Economic Studies* 74, 447–476.

- [16] Daskalakis, C., A. Deckelbaum, and C. Tzamos (2013), “Mechanism Design via Optimal Transport,” in *Proceedings of the 14th Annual ACM Conference on Economics and Computation* (EC ’13), 269–286.
- [17] Dhangwatnotai, P., T. Roughgarden, and Q. Yan (2010), “Revenue Maximization with a Single Sample,” in *Proceedings of the 11th Annual ACM Conference on Economics and Computation* (EC ’10), 129–138.
- [18] Falk, M., J. Hüsler, and R.-D. Reiss (2010), *Laws of Small Numbers: Extremes and Rare Events* (third edition). Birkhäuser.
- [19] Frankel, A. (2014), “Aligned Delegation,” *American Economic Review*, 104(1), 66–83.
- [20] Garratt, R., and T. Tröger (2006), “Speculation in Standard Auctions with Resale,” *Econometrica* 74(3), 753–769.
- [21] Garratt, R., T. Tröger, and C. Z. Zheng (2009), “Collusion via Resale,” *Econometrica* 77(4), 1095–1136.
- [22] Haile, P. (2003), “Auctions with Private Uncertainty and Resale Opportunities,” *Journal of Economic Theory* 108 (1), 72–110.
- [23] Hartline, J., and T. Roughgarden (2009), “Simple versus Optimal Mechanisms,” in *Proceedings of the 10th ACM conference on Electronic Commerce* (EC ’09), 225–234.
- [24] Jehiel, P., and B. Moldovanu (2000), “Auctions with Downstream Interaction among Buyers,” *RAND Journal of Economics* 31 (4), 768–791.
- [25] Laffont, J.-J., and D. Martimort (1997), “Collusion under Asymmetric Information,” *Econometrica* 65, 875–911.
- [26] Li, S. (2017), “Obviously Strategy-Proof Mechanisms,” *American Economic Review* 107 (11), 3257–3287.
- [27] McAfee, R.P., and J. McMillan (1989), “Government Procurement and International Trade,” *Journal of International Economics* 26, 921–943.
- [28] Milgrom, P., and I. Segal (2002), “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica* 70(2), 583–601.

- [29] Milgrom, P., and I. Segal (2017), “Deferred-Acceptance Auctions and Radio Spectrum Reallocation,” working paper, Stanford University.
- [30] Milgrom, P., and R. Weber (1985), “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research* 10(4), 619–632.
- [31] Myerson, R. (1981), “Optimal Auction Design,” *Mathematics of Operations Research* 6(1), 58-73.
- [32] 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.
- [33] Myerson, R., and M. Satterthwaite (1983), “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory* 29, 265-281.
- [34] Perry, M., and P. Reny (2002), “An Efficient Auction,” *Econometrica* 70(3), 1199–1212.
- [35] Tirole, J. (1988), *The Theory of Industrial Organization*. MIT Press.
- [36] Yamashita, T. (2015), “Implementation in Weakly Undominated Strategies: Optimality of Second-Price Auction and Posted-Price Mechanism,” *Review of Economic Studies* 82, 1223-1246.
- [37] Yamashita, T., and S. Zhu (2017), “On the Foundations of Ex Post Incentive Compatible Mechanisms,” working paper, Toulouse School of Economics.
- [38] Zheng, C.Z. (2002), “Optimal Auctions with Resale,” *Econometrica* 70(6), 2197–2224.

A Proof of Lemma 1

(4)-(5) imply that for all auctions Γ , $\inf_v \text{Rev}(\Gamma, v) \leq \bar{R} \leq \inf_v \text{Rev}(\bar{\Gamma}, v)$, hence $\bar{\Gamma}$ solves the robust revenue maximization problem, and the problem's value is \bar{R} . Also, (4)-(5) imply that for all resale procedures v , $\sup_{\Gamma} \text{Rev}(\Gamma, v) \geq \bar{R} \geq \sup_{\Gamma} \text{Rev}(\Gamma, \underline{v})$, hence \underline{v} solves the worst-case resale problem, and the problem's value is \bar{R} . Plugging in $\Gamma = \bar{\Gamma}$ and $v = \underline{v}$ into (4)-(5) yields $\text{Rev}(\bar{\Gamma}, \underline{v}) = \bar{R}$.

B Proof of Theorem 1

Take any direct auction mechanism $\langle \chi, \psi \rangle$ satisfying (6)-(7), and define interim expected payoffs $U_i(\theta_i)$ as in (9). Rewrite incentive compatibility (6) as

$$U_i(\theta_i) - U_i(\hat{\theta}_i) \geq \mathbb{E}_{\tilde{\theta}_{-i}} [\underline{v}_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - \underline{v}_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \hat{\theta}_i, \tilde{\theta}_{-i})] \text{ for all } \theta_i, \hat{\theta}_i. \quad (23)$$

Since $\sum_j \chi_j(\tilde{\theta}) = 1$ with probability 1, there exists arbitrarily small $\hat{\theta}_i$ such that $\sum_j \chi_j(\hat{\theta}_i, \tilde{\theta}_{-i}) = 1$ with probability 1. Note that for any profile θ such that $\min_j \theta_j > \hat{\theta}_i$ and $x \equiv \chi(\hat{\theta}_i, \theta_{-i})$ satisfies $\sum_j x_j = 1$, we have $\underline{v}_i(x; \theta) - \underline{v}_i(x; \hat{\theta}_i, \theta_{-i}) \geq \theta_i - \theta_{(2)}$ if $\theta_i > \theta_{(2)}$ and $\underline{v}_i(x; \theta) \geq \underline{v}_i(x; \hat{\theta}_i, \theta_{-i})$ otherwise, hence $\underline{v}_i(x; \theta) - \underline{v}_i(x; \hat{\theta}_i, \theta_{-i}) \geq \max\{\theta_i - \theta_{(2)}, 0\}$. For any other θ , $|\underline{v}_i(x; \theta) - \underline{v}_i(x; \hat{\theta}_i, \theta_{-i}) - \max\{\theta_i - \theta_{(2)}, 0\}| \leq 2$. Hence, taking expectations, (23) implies

$$\mathbb{E}_{\tilde{\theta}_i} [U_i(\tilde{\theta}_i)] - U_i(\hat{\theta}_i) \geq \mathbb{E}_{\tilde{\theta}} \left[\max\{\tilde{\theta}_i - \tilde{\theta}_{(2)}, 0\} \right] - 2 \left(1 - \prod_j [1 - F_j(\hat{\theta}_i)] \right).$$

Since the inequality has to hold for arbitrarily small $\hat{\theta}_i$, and $U_i(\hat{\theta}_i) \geq 0$ by individual rationality (7), we must have

$$\mathbb{E}_{\tilde{\theta}_i} [U_i(\tilde{\theta}_i)] \geq \mathbb{E}_{\tilde{\theta}} \left[\max\{\tilde{\theta}_i - \tilde{\theta}_{(2)}, 0\} \right].$$

Therefore,

$$\begin{aligned} \mathbb{E}_{\tilde{\theta}} \left[\sum_i \psi_i(\tilde{\theta}) \right] &= \mathbb{E}_{\tilde{\theta}} \left[\sum_i \underline{v}_i(\chi(\tilde{\theta}); \tilde{\theta}) \right] - \sum_i \mathbb{E}_{\tilde{\theta}_i} [U_i(\tilde{\theta}_i)] \\ &\leq \mathbb{E}_{\tilde{\theta}} [\tilde{\theta}_{(1)}] - \mathbb{E}_{\tilde{\theta}} \left[\sum_i \max\{\tilde{\theta}_i - \tilde{\theta}_{(2)}, 0\} \right] = \mathbb{E}_{\tilde{\theta}} [\tilde{\theta}_{(2)}], \end{aligned}$$

which is the expected revenue in the Vickrey auction.

C Resale-Proofness of ACV Auctions

In this appendix we prove Theorem 2, and in the process we formulate a generalization that applies to a broader class of resale procedures, which may take place under asymmetric information. For this purpose, we assume the auction has the ACV form, and we allow the outcome of the resale bargaining to depend on bidders' reports $\hat{\theta}_1, \dots, \hat{\theta}_n$ in the auction, since these reports may affect bidders' beliefs about each other's values and consequently their behavior in bargaining. Therefore, we now 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.^{25,26} This formalism lets us distinguish the dependence of i 's payoff on $\hat{\theta}_j$ (which can influence the bidders' beliefs about j 's type at the time bargaining begins) from the dependence on θ_j (which can influence how j actually behaves in bargaining).²⁷

²⁵For simplicity, we assume that after the ACV auction is completed, the auctioneer discloses all bids. This avoids the need to model a bidder's inference about bids from imperfect information disclosed in the auction (e.g., his own allocation).

²⁶Note that we continue modeling resale using reduced-form payoffs rather than as an explicit noncooperative game following the auction game. The problem with the latter approach is that to ensure individual rationality in a noncooperative resale game, each bidder should have access to a "non-participation" action. However, then our formulation of the robust revenue maximization problem is too weak: *any* auction that is incentive-compatible without resale — in particular the Myerson (1981) optimal auction — has an equilibrium in which every agent always bids truthfully and then never participates in resale. We could try to rule these out by restricting to equilibria satisfying some perfection requirement in the definition of $\text{Rev}(\Gamma, v)$. But we want to ensure that a reasonably broad class of mechanisms can guarantee equilibrium existence for all v , and existence of perfect equilibria with infinite type spaces is a thorny problem (see Myerson and Reny, 2015).

²⁷We comment that our approach to modeling resale payoffs here relies on having fixed the mechanism and the proposed equilibrium. In a general mechanism, if we did not assume full information in resale, then the effect of messages on resale bargaining would depend on the off-path beliefs they induce within the equilibrium being played; thus we would have a circularity, with resale payoffs depending on the choice of equilibrium and

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:

$$\sum_i v_i(x; \hat{\theta}; \theta) \leq (\max_i \theta_i) \cdot \left(\sum_i x_i \right) \quad (24)$$

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

$$v_i(x; \hat{\theta}; \theta) \geq \theta_i x_i \quad (25)$$

for all $i, x, \hat{\theta}, \theta$.²⁸

Theorem 4. *Consider an ACV auction with reserve prices p_1, \dots, p_n . For any resale payoff functions v_1, \dots, v_n satisfying (24)–(25), it is an ex post equilibrium for all bidders to participate and report their true values (after which no resale occurs). Formally, for any values $\theta_1, \dots, \theta_n$, and any possible deviation $\hat{\theta}_i \in [0, 1] \cup \{\emptyset\}$ for any bidder i ,*

$$\theta_i \chi_i(\theta) - \psi_i(\theta) \geq v_i(\chi(\hat{\theta}_i, \theta_{-i}); \hat{\theta}_i, \theta_{-i}; \theta) - \psi_i(\hat{\theta}_i, \theta_{-i}). \quad (26)$$

We now give a proof of the theorem. Note that we need not consider the non-participation message $\hat{\theta}_i = \emptyset$ separately, since it is treated by the mechanism as equivalent to reporting 0. We consider all possible cases.

- Suppose that bidder i wins under truth-telling: $\chi_i(\theta) = 1$. Then the left-hand side of (26) 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 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 $\theta_j = \max_{k \neq i} \theta_k = \psi_i(\theta)$ (i 's threshold price under truth-telling). Then (25) 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 (24) the resale payoff of bidder i is at most $\theta_i - \theta_j$. So the right side of (26) is at most $\theta_i - \theta_j$ which equals the left side.

vice versa. Since these modeling technicalities are not our focus, we have avoided them by focusing on full-information resale for the formal statement of the robust revenue problem (3).

²⁸Note that (25) is an ex post individual rationality constraint.

- If i deviates to $\hat{\theta}_i$ such that the object goes unallocated, then no resale is possible, so the right side of (26) is zero.
- Suppose that some bidder $j \neq i$ wins under truth-telling: $\chi_j(\theta) = 1$. Then the left-hand side of (26) 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 (26) 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 (25), the resale payoffs of all other bidders are nonnegative, and therefore by (24) the resale payoff of bidder i is at most $\max_k \theta_k = \theta_j$. So the right side of (26) 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 (26) is zero.
 - If i deviates such that the object remains unsold, then no resale is possible and the right side of (26) 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:

$$\theta_j \geq \hat{\theta}_i > p_i \geq \theta_i, \text{ and } \theta_j \geq \theta_k \text{ for all } k \neq i.$$

So allocating to j is efficient for the true value profile θ , and no resale is possible, so the right side of (26) 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 (25), the resale payoffs of all other bidders are nonnegative, and therefore by (24) the resale payoff of bidder i is at most $\max_j \theta_j$. Since $\theta_i \leq p_i$ (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 (26) is at most zero.

D A Key Lemma

The following lemma is key to constructing the multipliers used to prove Theorem 3.

Lemma 2. *Under assumptions (A1)-(A3), there exist unique sequences $p_1 \geq \dots \geq p_n$ and $\bar{R}_1, \dots, \bar{R}_n$ satisfying (13)-(14). The sequence of prices p_k and the function $H : [p_n, 1) \rightarrow \mathbb{R}$ defined by (19) are jointly characterized by the following two properties:*

- (i) $p_k > 0$ and (20) holds for each $k \geq 1$,
- (ii) H is differentiable and satisfies differential equation (21) and the boundary condition $H(p) = 0$ for $p \in [p_1, 1)$.

Furthermore,

- (iii) *the function H is nonincreasing, and*
- (iv) *the reserve prices satisfy $p_k \geq r_k$ for each k (with equality for $k = 1$).*

The lemma follows from the following four claims:

Claim 1. There exist reserve prices and revenue levels satisfying (13)-(14).

Claim 2. When reserve prices and revenue levels satisfy (13)-(14) and the function H is defined by (19), properties (i)-(ii) hold.

Claim 3. When properties (i)-(ii) hold, property (iii) also holds.

Claim 4. There is a unique price sequence p_1, \dots, p_n and a corresponding function H such that properties (i)-(iii) hold. These prices satisfy (iv).

We prove these claims in turn.

D.1 Proof of Claim 1

The claim follows from the continuity of the objective function in (14).

D.2 Proof of Claim 2

First, note that H as defined by (19) is continuous at the endpoints of the intervals of its definition, i.e. the points $p = p_k$. This follows from (13). (It is immediate that H is continuous elsewhere.)

We now show properties (i)-(ii) in turn:

- (i) First we argue that the derivative of the expected revenue $R_k(p)$ in the Vickrey auction with symmetric reserve p is given by the following formula:

$$R'_k(p) = - \sum_{i=1}^k \left(\prod_{\substack{j=1 \\ j \neq i}}^k F_j(p) \right) f_i(p) \nu_i(p). \quad (27)$$

Indeed, we can express the value of this Vickrey auction as the integral of virtual surplus, as in (12):

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

Equivalently, this is the integral of $\nu_{i^*(\theta)}(\theta_{i^*(\theta)})$ over the region where bidder $i^*(\theta)$ has value at least p . As p increases marginally, the 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 $\nu_i(p)$.

Now, we show that property (i) holds, and in addition that $\bar{R}_k > R_k(0)$, by induction on $k \geq 1$. This holds for $k = 1$, as the optimal price must satisfy $p_1 \in (0, 1)$ to achieve revenue $\bar{R}_1 > R_1(0) = 0$, and must satisfy the first-order condition $\nu_1(p_1) = H(p_1) = 0$. Now, suppose that the inductive hypothesis holds for $k - 1$. Using (27), express

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

and calculate the derivative of the objective function in (14) for $p \in [0, p_{k-1}]$ as

$$\begin{aligned} D_k(p) &\equiv R'_k(p) + f_k(p) [\bar{R}_{k-1} - R_{k-1}(p)] - F_k(p) R'_{k-1}(p) \\ &= f_k(p) \left(\bar{R}_{k-1} - R_{k-1}(p) - \nu_k(p) \prod_{i=1}^{k-1} F_i(p) \right). \end{aligned}$$

In particular, $D_k(0) = f_k(0) (\bar{R}_{k-1} - R_{k-1}(0)) > 0$ by the inductive hypothesis for $k-1$, and therefore we have $\bar{R}_k > R_k(0)$ and $p_k > 0$.

Now we can express

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

where H is given by (19). Consider two remaining cases:

- If $p_k \in (0, p_{k-1})$ then the necessary first-order condition for maximization (14) is $D_k(p_k) = 0$, and therefore $H(p_k) = \nu_k(p_k)$.
- If $p_k = p_{k-1}$ then the necessary first-order condition for maximization (14) is $D_k(p_k) \geq 0$, and therefore $H(p_k) \geq \nu_k(p_k)$. On the other hand, we have

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

where the second equality uses the inductive hypothesis and the inequality uses assumption (A3). Combined with the above, we again have $H(p_k) = \nu_k(p_k)$.

- (ii) From (19) and the observation that H is continuous at the interval endpoints, $H(p) \prod_{j < k} F_j(p) = \bar{R}_{k-1} - R_{k-1}(p)$ for $p \in [p_k, p_{k-1}]$. It is then immediate from (27) that (21) holds on the interval $[p_k, p_{k-1}]$, where we take the derivative to be the right-hand derivative at p_k and the left-hand derivative at p_{k-1} . (The fact that H is zero on $[p_1, 1)$ is immediate from the definition.)

It remains to check that H is in fact differentiable at the interval endpoints, i.e. that the left-hand and right-hand derivatives are equal at each p_k . For this purpose, and for future use more generally, it is useful to rewrite (21) in the form

$$H'(p) = \sum_{j < k} \frac{f_j(p)}{F_j(p)} [\nu_j(p) - H(p)] \quad \text{for } p \in [p_k, p_{k-1}] \quad (28)$$

(where, again, we use one-sided derivatives at interval endpoints). To derive this equation, simply expand the derivative on the left side of (21) using the product rule, then rearrange.

Then the fact that the left-hand and right-hand derivatives agree at p_k is immediate, since the right side of (28) on interval $[p_{k+1}, p_k]$ differs from the formula for interval $[p_k, p_{k-1}]$ only by inclusion of the $j = k$ term, which is zero at p_k by property (i).

D.3 Proof of Claim 3

Using equation (28) above and assumption (A3), we see that for any $p \in [p_k, p_{k-1}]$ such that $H(p) \geq \nu_{k-1}(p)$ we have $H'(p) \leq 0$, and the inequality is strict if $H(p) > \nu_{k-1}(p)$.

Now, by property (i), $H(p_{k-1}) = \nu_{k-1}(p_{k-1})$, and therefore $H'(p_{k-1}) \leq 0$, hence for p in a left-neighborhood of p_{k-1} ,

$$\begin{aligned} H(p) - \nu_{k-1}(p) &\geq \nu_{k-1}(p_{k-1}) + o(p_{k-1} - p) - \nu_{k-1}(p) \\ &\geq p_{k-1} - p + o(p_{k-1} - p) > 0, \end{aligned}$$

where the second inequality uses (A1). Thus, we can choose $\hat{p} < p_{k-1}$ close enough to p_{k-1} such that $H(p) > \nu_{k-1}(p)$ for all $p \in [\hat{p}, p_{k-1})$, and therefore $H'(p) < 0$ on this interval. In particular this implies $H(\hat{p}) > H(p_{k-1})$. But then $H(p)$ cannot cross $\nu_{k-1}(p_{k-1})$ anywhere in the interval $[p_k, p_{k-1}]$: indeed, otherwise, letting $p^\circ = \max\{p \in [p_k, \hat{p}] : H(p) \leq \nu_{k-1}(p_{k-1})\}$ we would have

$$0 = H(p_{k-1}) - \nu_{k-1}(p_{k-1}) < H(\hat{p}) - \nu_{k-1}(p_{k-1}) = \int_{p^\circ}^{\hat{p}} H'(p) dp, \quad (29)$$

while on the other hand we would have $H(p) \geq \nu_{k-1}(p_{k-1}) \geq \nu_{k-1}(p)$ for all $p \in (p^\circ, \hat{p})$ and therefore $H'(p) \leq 0$ for all $p \in (p^\circ, \hat{p})$, making the right-hand side of (29) nonpositive – a contradiction.

Therefore, $H(p) \geq \nu_{k-1}(p_{k-1}) \geq \nu_{k-1}(p)$ for all $p \in [p_k, p_{k-1}]$, and therefore $H'(p) \leq 0$ on this interval. Since this holds for each $k > 1$, H is nonincreasing on $[p_n, p_1]$.

D.4 Proof of Claim 4

We prove by induction on $k \geq 1$ that price p_k is uniquely determined and satisfies $p_k \geq r_k$. First, the claim holds for $k = 1$, since the boundary condition $H(p_1) = 0$ implies that (20) takes the form $\nu_1(p_1) = 0$, and this uniquely determines $p_1 = r_1$. Now, suppose that the claim holds for all $j \leq$

$k-1$, and in particular that prices p_1, \dots, p_{k-1} are uniquely determined. Then by using property (ii) and integrating, the function H is uniquely determined on the interval $[p_k, 1)$, whatever p_k may be. Furthermore, by (i), p_k must solve (20). Observe that:

- Any value of p_k satisfying (20) must lie in $p_k \in [r_k, p_{k-1}]$. Indeed, suppose for contradiction that $p_k < r_k$. Then, using properties (i) and (iii), and strict monotonicity of ν_k (from (A1)), we have $\nu_k(p_k) < \nu_k(r_k) = 0 = H(r_1) \leq H(p_k) = \nu_k(p_k)$, a contradiction.
- (20) has at most one solution on $p_k \in [0, p_{k-1}]$. This holds because ν_k is strictly increasing by (A1) and H is nonincreasing by property (iii).

The two observations together imply that (20) has a unique solution $p_k \in [0, p_{k-1}]$, which satisfies $p_k \geq r_k$. Thus, the inductive statement also holds for k .

E Proof of Theorem 3

The proof of Theorem 3 depends on Lemma 2 in Appendix D. Henceforth we take this lemma, and the steps of its proof, as given.

We also find it useful to write $\underline{v}_i(x; \theta) = \sum_j \underline{v}_{ij}(\theta) x_j$, where $\underline{v}_{ij}(\theta) = \max\{\theta_i, \theta_{(2)}\}$ if $i = j$ and $= \max\{0, \theta_i - \theta_{(2)}\}$ otherwise. Write $\underline{v}'_{ij}(\theta)$ for the derivative with respect to θ_i (which is defined almost everywhere).

Now, to prove the theorem, 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 $\hat{\Lambda}_i(\hat{\theta}_i) = F_i(\hat{\theta}_i)$ and $\Lambda_i(\theta_i) = \bar{\Lambda}_i(p_n) - \bar{\Lambda}_i(\theta_i)$, where

$$\bar{\Lambda}_i(\theta_i) = \begin{cases} \frac{f_i(p_n)}{F_i(p_n)} [H(p_n) - \nu_i(p_n)] & \text{if } \theta_i < p_n, \\ \frac{f_i(\theta_i)}{F_i(\theta_i)} [H(\theta_i) - \nu_i(\theta_i)] & \text{if } \theta_i \in [p_n, p_i], \\ 0 & \text{if } \theta_i > p_i. \end{cases} \quad (30)$$

Note that by (A1) and Lemma 2(iii), the difference $H(\theta_i) - \nu_i(\theta_i)$ is nonincreasing, and in particular $H(\theta_i) - \nu_i(\theta_i) \geq H(p_i) - \nu_i(p_i) = 0$. Using in addition (A2), we see that $\bar{\Lambda}_i$ is nonincreasing, hence Λ_i is nondecreasing, so it indeed describes a measure. Furthermore, by construction, $\bar{\Lambda}_i(\theta_i)$, and so $\Lambda_i(\theta_i)$, is constant on $\theta_i < p_n$ and on $\theta_i > p_i$, hence Λ_i is supported on $[p_n, p_i]$.

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 $\hat{\theta}_i \leq \theta_i$. Also, define $\bar{\Lambda}_n, \mathcal{M}_n \equiv 0$. We already argued in Subsection 4.5 that the conditions (a)–(b) stated there suffice to prove the theorem, and that \mathcal{M}_i as defined here satisfies the complementary slackness condition (b), so the rest of the proof focuses on establishing the optimization condition (a).

First we rewrite the Lagrangian in terms of modified virtual values. For this purpose, note that the slack in (15) is absolutely continuous in θ_i , and

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

Then we can write, using integration by parts and the fact that $\bar{\Lambda}_i(1) = S_i(\hat{\theta}_i, \hat{\theta}_i; \chi) = 0$,

$$\int_{\hat{\theta}_i}^1 S_i(\theta_i, \hat{\theta}_i; \chi) d\Lambda_i(\theta_i) = \int_{\hat{\theta}_i}^1 \bar{\Lambda}_i(\theta_i) \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i}) - v'_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})] d\theta_i.$$

Integrating this expression with $d\hat{\Lambda}_i(\hat{\theta}_i)$, which equals $dF_i(\hat{\theta}_i)$ for $\hat{\theta}_i \leq p_i$ while elsewhere the expression is zero (since $\bar{\Lambda}_i(\theta_i) = 0$ for $\theta_i > p_i$), we obtain

$$\begin{aligned} \iint S_i(\theta_i, \hat{\theta}_i; \chi) d\mathcal{M}_i(\theta_i, \hat{\theta}_i) &= \int_0^1 \int_{\hat{\theta}_i}^1 \bar{\Lambda}_i(\theta_i) \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})] d\theta_i dF_i(\hat{\theta}_i) \\ &\quad - \int_0^1 \int_{\hat{\theta}_i}^1 \bar{\Lambda}_i(\theta_i) \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\hat{\theta}_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})] d\theta_i dF_i(\hat{\theta}_i). \end{aligned}$$

The first term can be rewritten using integration by parts as

$$\int_0^1 F_i(\theta_i) \bar{\Lambda}_i(\theta_i) \mathbb{E}_{\tilde{\theta}_{-i}} [v'_i(\chi(\theta_i, \tilde{\theta}_{-i}); \theta_i, \tilde{\theta}_{-i})] d\theta_i = \mathbb{E}_{\tilde{\theta}} \left[\frac{F_i(\tilde{\theta}_i)}{f_i(\tilde{\theta}_i)} \bar{\Lambda}_i(\tilde{\theta}_i) v'_i(\chi(\tilde{\theta}); \tilde{\theta}) \right],$$

while the second term can be rewritten as

$$-\mathbb{E}_{\tilde{\theta}} \left[\int_{\hat{\theta}_i}^1 \bar{\Lambda}_i(\tau_i) v'_i(\chi(\tilde{\theta}); \tau_i, \tilde{\theta}_{-i}) d\tau_i \right].$$

Adding these two terms across bidders i and adding the sum to the first part of Lagrangian (16), representing $v_i(x, \theta) = \sum_j v_{ij}(\theta) x_j$, and taking into

account that $\sum_i v_{ij}(\theta) = \theta_{(1)}$ for any j yields the Lagrangian in the form (17), where

$$\mu_j(\theta) = \theta_{(1)} - \sum_i \left\{ \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} - \frac{F_i(\theta_i)}{f_i(\theta_i)} \bar{\Lambda}_i(\theta_i) \right\} \underline{v}'_{ij}(\theta) - \sum_i \int_{\theta_i}^{p_i} \bar{\Lambda}_i(\tau_i) \underline{v}'_{ij}(\tau_i, \theta_{-i}) d\tau_i. \quad (31)$$

Now to show optimality condition (a), we show that the ACV auction maximizes the modified virtual surplus (17) pointwise with probability 1. In symbols, we will show that at each value profile θ that does not have ties:

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

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

To show (a-1), consider (31) for bidder $j = i^* = i^*(\theta)$. For the second term of (31), note that $\underline{v}'_{ii^*}(\theta) = 1$ when $i = i^*$ and $= 0$ otherwise. For the last term of (31), note that $\underline{v}'_{i^*i^*}(\tau_{i^*}, \theta_{-i^*}) = 1$ for all $\tau_{i^*} > \theta_{i^*}$, while for any $i \neq i^*$, $\underline{v}'_{ii^*}(\tau_i, \theta_{-i}) = 1$ when $\tau_i > \theta_{i^*}$ (since i then buys from i^* in resale) and $= 0$ otherwise. Thus

$$\mu_{i^*}(\theta) = \theta_{i^*} - \frac{1 - F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} + \frac{F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} \bar{\Lambda}_{i^*}(\theta_{i^*}) - \sum_i \int_{\theta_{i^*}}^{p_i} \bar{\Lambda}_i(\tau_i) d\tau_i. \quad (32)$$

We evaluate (32) in three different cases depending on where θ_{i^*} lies, and show that in every case, the sign of $\mu_{i^*}(\theta)$ is consistent with (a-1).

- If $\theta_{i^*} \in [p_n, p_{i^*}]$, then using (30) we see that the first three terms in (32) add up to $H(\theta_{i^*})$, and so

$$\mu_{i^*}(\theta) = H(\theta_{i^*}) - \sum_i \int_{\theta_{i^*}}^{p_i} \bar{\Lambda}_i(\tau_i) d\tau_i.$$

The derivative of this expression with respect to θ_{i^*} is

$$H'(\theta_{i^*}) + \sum_i \bar{\Lambda}_i(\theta_{i^*}) = H'(\theta_{i^*}) + \sum_{i: \theta_{i^*} < p_i} \frac{f_i(\theta_{i^*})}{F_i(\theta_{i^*})} [H(\theta_{i^*}) - \nu_i(\theta_{i^*})],$$

which equals zero by (28). Since the expression equals $H(p_1) = 0$ at $\theta_{i^*} = p_1$, it follows that $\mu_{i^*}(\theta) = 0$ 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 the third term in (32) is zero instead of $H(\theta_{i^*}) - \nu_{i^*}(\theta_{i^*})$, hence $\mu_{i^*}(\theta) = \nu_{i^*}(\theta_{i^*}) - H(\theta_{i^*})$. Since ν_{i^*} is weakly increasing, H is weakly decreasing, and $\nu_{i^*}(p_{i^*}) = H(p_{i^*})$, we have $\mu_{i^*}(\theta) \geq 0$ when $\theta_{i^*} > p_{i^*}$, which is 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 $\mu_{i^*}(\theta) \leq 0$. The first three terms in (32) add up to

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

This expression is increasing in θ_{i^*} , since ν_{i^*} is increasing, F_{i^*}/f_{i^*} is increasing (by (A2)), and $H(p_n) = \nu_n(p_n) \geq \nu_{i^*}(p_n)$ using Lemma 2. Meanwhile the last term in (32) is clearly decreasing in θ_{i^*} . Therefore, $\mu_{i^*}(\theta_{i^*}, \theta_{-i^*})$ is increasing in θ_{i^*} for $\theta_{i^*} \leq p_n$. Since we already saw that it is 0 at p_n , we get (a-1) as needed.

Now to prove (a-2), consider $\mu_j(\theta)$ for a bidder $j \neq i^*$. Note that $\underline{v}'_{ij}(\theta) = 0$ unless $i = i^*$ or $i = j$. Also, in the last term of (31), $\underline{v}'_{jj}(\tau_j, \theta_{-j}) = 1$ for $\tau_j > \theta_{(2)}$ and $= 0$ otherwise, while for $i \neq j$, $\underline{v}'_{ij}(\tau_i, \theta_{-i}) = 1$ for $\tau_i > \theta_{i^*}$ and $= 0$ otherwise. So,

$$\begin{aligned} \mu_j(\theta) &= \theta_{i^*} - \frac{1 - F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} + \frac{F_{i^*}(\theta_{i^*})}{f_{i^*}(\theta_{i^*})} \bar{\Lambda}_{i^*}(\theta_{i^*}) - \left\{ \frac{1 - F_j(\theta_j)}{f_j(\theta_j)} - \frac{F_j(\theta_j)}{f_j(\theta_j)} \bar{\Lambda}_j(\theta_j) \right\} \underline{v}'_{jj}(\theta) \\ &\quad - \sum_{i \neq j} \int_{\theta_{i^*}}^{p_i} \bar{\Lambda}_i(\tau_i) d\tau_i - \int_{\theta_{(2)}}^{p_j} \bar{\Lambda}_j(\tau_j) d\tau_j. \end{aligned}$$

Subtracting from (32), we have

$$\mu_{i^*}(\theta) - \mu_j(\theta) = \left\{ \frac{1 - F_j(\theta_j)}{f_j(\theta_j)} - \frac{F_j(\theta_j)}{f_j(\theta_j)} \bar{\Lambda}_j(\theta_j) \right\} \underline{v}'_{jj}(\theta) + \int_{\theta_{(2)}}^{\theta_{i^*}} \bar{\Lambda}_j(\tau_j) d\tau_j. \quad (33)$$

The last term in (33) is clearly nonnegative, as is \underline{v}'_{jj} . So we focus on the expression in curly braces, and consider different cases depending on where θ_j lies:

- If $\theta_j > p_j$ then $\bar{\Lambda}_j(\theta_j) = 0$ and so the term is nonnegative.

- If $\theta_j \in [p_n, p_j]$, the term equals

$$\frac{1 - F_j(\theta_j)}{f_j(\theta_j)} - \frac{F_j(\theta_j)}{f_j(\theta_j)} \cdot \frac{f_j(\theta_j)}{F_j(\theta_j)} [H(\theta_j) - \nu_j(\theta_j)] = \theta_j - H(\theta_j),$$

and

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

where we have used Lemma 2(iii) and (i).

- For $\theta_j \leq p_n$, the term equals

$$\frac{1 - F_j(\theta_j)}{f_j(\theta_j)} - \frac{F_j(\theta_j)}{f_j(\theta_j)} \bar{\Lambda}_j(p_n).$$

By (A1) and (A2), the expression is nonincreasing in θ_j , and since we already saw that it is nonnegative at $\theta_j = p_n$, it must be nonnegative for $\theta_j < p_n$.

Thus, (33) is nonnegative, establishing (a-2).

F Example with Uniform Distributions

We illustrate here the calculation technique from Subsection 4.6, by considering the case where bidder i 's value is uniformly distributed on $[0, a_i]$, with $a_1 \geq \dots \geq a_n$. While offering a convenient analytical illustration, this example also forces us to extend the main model to allow the bidders' distributions to have different supports. We argue that our formal results can be extended to cover such cases. For this purpose, assumptions (A1) and (A2) can be interpreted as stipulating the concavity of functions $\log(1 - F_i)$ and $\log F_i$, respectively, which in particular requires the supports to be intervals, while allowing function values $-\infty$ outside the supports. Assumption (A3) then implies that the supports of different bidders are ordered as we have specified.²⁹

²⁹In some places in the proofs, expressions appear that have an f_i in the denominator, which is zero outside the support. However, all such expressions are then of the form $0/0$ and the arguments remain valid if we interpret this ratio as zero. For example, bidder i 's virtual value can be extended above his support as $\nu_i(\theta_i) = \theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} = \theta_i$.

One substantive issue arising with different upper limits is that for some bidder k , the maximization problem (14) may be solved by setting $p_k \geq a_k$. According to the first-order condition (20), this occurs if and only if $H(a_k) \geq \nu_k(a_k) = a_k$.³⁰ Then, for each bidder $j > k$ we also have $H(a_k) \geq \nu_j(a_k)$, making it optimal to set $p_j \geq a_k \geq a_j$. Thus, in this case bidder k as well as all the weaker bidders are completely excluded from the auction.

In our uniform-distribution example, if bidder k is not excluded from the auction, then we can plug in $f_j(p) = 1/a_j$, $F_j(p) = p/a_j$, and $\nu_j(p) = 2p - a_j$ into formula (21), integrate the formula on the interval $[p_k, p_{k-1}]$, and multiply through by $\prod_{j < k} a_j$ to obtain

$$H(p_{k-1})p_{k-1}^{k-1} - H(p_k)p_k^{k-1} = \int_{p_k}^{p_{k-1}} p^{k-1} \left(\sum_{j < k} \left(2 - \frac{a_j}{p} \right) \right) dp.$$

Integrating, using (20) to plug in $H(p_k) = \nu_k(p_k) = 2p_k - a_k$ and likewise $H(p_{k-1})$, and rearranging yields

$$\left(\frac{p_k}{p_{k-1}} \right)^{k-1} \left[\frac{2}{k} p_k + \frac{1}{k-1} \sum_{j < k} a_j - a_k \right] = \frac{2}{k} p_{k-1} + \frac{1}{k-1} \sum_{j < k} a_j - a_{k-1}. \quad (34)$$

This k th-degree polynomial equation determines the reserve price p_k .

Note that if bidder k is excluded, (34) is not equivalent to (20) and so is not the right formula for p_k . However, it is still the case that k is excluded if and only if the solution to (34) exceeds a_k , since the left-hand side of (34) is an increasing function of p_k , and so bidder k is excluded if and only if this left-hand side falls below the right-hand side at $p_k = a_k$.

For example, for two bidders, we have $p_1 = r_1 = a_1/2$, and for $k = 2$ equation (34) reads

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

The positive root p_2 of this quadratic equation is given by formula (22) in the text. Bidder 2 is excluded if $a_2 a_1 \leq a_1^2/4$, or $a_2 \leq a_1/4$.

³⁰Note that in this case, all choices of $p_k \in [a_k, p_{k-1}]$ are optimal, but we can resolve this multiplicity by defining $\nu_k(\theta_k) = \theta_k$ for $\theta_k > a_k$, as mentioned in the previous footnote, and letting $p_k \geq a_k$ be the solution to (20), which exists by continuity since $H(1) = 0 < 1 = \nu_k(1)$. The proof of Theorem 3 then remains valid.

G Non-Optimality of ACV Auction Without Regularity

We give here an example to show that, if we make no regularity assumptions on the distributions, then the robust-revenue-maximizing auction may not be of the ACV form. Moreover, this example is symmetric. (We do not know the optimal auction in this case; we simply show that one can do better than the best ACV auction.)

Consider a distribution F on $[0, 1]$, and two values $r < s$, satisfying the following properties:

- The one-buyer revenue function R_1 has two maxima, at r and s .
- There are just three points where the virtual value $\nu(\theta)$ is zero, namely r , s , and one in between (a local minimum of the revenue function).
- $F(r) > 1/2$.

Clearly this can be constructed.

Suppose there are two bidders, with both values distributed according to F . We claim that among all ACV auctions, the unique optimal auction is the symmetric auction with $p_1 = p_2 = s$, yielding revenue $R_2(s)$. Indeed, first suppose in negation that an optimal ACV auction has $p_2 < p_1$. Then local optimality with respect to p_1 implies that either $p_1 = r$ or $p_1 = s$. In the former case, we cannot satisfy the first-order condition (18) for local optimality of $p_2 < p_1 = r$, since its left-hand side is negative while its right-hand side is positive. In the latter case, note that raising bidder 2's reserve price from p_2 to $p_1 = s$ would raise the expected revenue by

$$F(p_2) [R_1(s) - R_1(p_2)] + 2 \int_{p_2}^s [R_1(s) - R_1(\theta)] f(\theta) d\theta, \quad (35)$$

where the first term accounts for the change in expected revenue on bidder 2 when $\theta_1 < p_2$, and the second term accounts for the change in expected revenue on bidder 2 when $\theta_1 \in [p_2, s]$ and for the symmetric change in expected revenue on bidder 1 when $\theta_2 \in [p_2, s]$. Both terms are strictly positive since $R_1(\theta) < R_1(s)$ for all $\theta < s$, $\theta \neq r$. Thus, the symmetric Vickrey auction with reserve price s would be a strict improvement.

Next, note that the Vickrey auction with symmetric reserve price $p > s$ cannot be optimal, since the first-order condition for the reserve price to maximize $R_2(p)$ is $2F(p)f(p)\nu(p) = 0$ implying $\nu(p) = 0$. Finally, for the symmetric Vickrey auction with reserve price $p < s$, raising the reserve price to s would raise the expected revenue by

$$2F(p)[R_1(s) - R_1(p)] + 2 \int_p^s [R_1(s) - R_1(\theta)] f(\theta) d\theta, \quad (36)$$

with the first term accounting for the change in expected revenue on bidder 2 when $\theta_1 < p$, and symmetrically on bidder 1 when $\theta_2 < p$, and the second term accounting for the change in expected revenue on bidder 2 when $\theta_1 \in [p, s]$, and symmetrically on bidder 1 when $\theta_2 \in [p, s]$. Again, both terms are strictly positive since $R_1(\theta) < R_1(s)$ for all $\theta < s$, $\theta \neq r$. Thus, the Vickrey auction with reserve price s is the unique optimal ACV auction.

Now, we can perturb F slightly to make r be the unique maximum of R_1 , while holding fixed the stationary points of F and still having both (35) and (36) strictly positive, so that the Vickrey auction with reserve s is still the unique optimal ACV auction. From now on we consider this perturbed F .

Note that the optimal ACV auction is now not described by (13)-(14) (which requires, in particular, $p_1 = r$). But more importantly, for this distribution F , the seller can strictly improve upon the optimal ACV auction by the auction using the following allocation rule (and threshold prices):

- If $\theta_1 > \theta_2$, then buyer 1 receives the good either if $\theta_1 \geq \max\{\theta_2, s\}$, or if $\theta_2 \leq r \leq \theta_1$.
- If $\theta_2 > \theta_1$, the conditions for allocating to buyer 2 are symmetric.

Figure 3 shows the allocation rule, with the regions where the good is allocated to bidder 1, bidder 2, and not at all (0).

Note that if bidders tell the truth, the revenue from this auction differs from that of the optimal ACV auction in two cases: when $\theta_1 < r$, then the expected revenue from bidder 2 is $R_1(r)$ instead of $R_1(s)$, and similarly for the expected revenue from bidder 1 when $\theta_2 < r$. Thus, the revenue of this new auction exceeds that of the optimal ACV auction by

$$2F(r)(R_1(r) - R_1(s)) > 0.$$

So indeed this auction does strictly better.




Figure 3: Allocation rule for alternative auction that outperforms ACV.

It remains to check that truthful bidding in this auction is (Bayesian) incentive-compatible, for any resale procedure. By symmetry, it suffices to check deviations by bidder 1. Overbidding is not profitable since it can cause bidder 1 to win only in situations where the price he pays is higher than both his actual value and bidder 2's value. What about underbidding? Downward deviations to any bid above s can never be profitable, even with resale; the argument is the same as in the ACV auction. Downward deviations from a true value above s to a bid between r and s cannot be better than bidding exactly s , since the only change is that when $\theta_2 \in (r, s)$, the good now goes unsold instead of being sold to bidder 1. And downward deviations from a bid in (r, s) to a bid in (r, s) have no effect. Likewise, downward deviations from below r to below r have no effect.

Finally, consider downward deviations from a true value $\theta_1 \geq r$ to a bid below r . These deviations can potentially be beneficial ex post, because they may result in bidder 2 winning the good (and then bidder 1 maybe buying it from him) when $\theta_2 \in (r, s)$. However, the gain in any such instance is at most $\theta_1 - \theta_2 \leq \theta_1 - r$, so the total expected value of these gains is at most $(F(s) - F(r))(\theta_1 - r)$. Meanwhile, there is an expected loss of $F(r)(\theta_1 - r)$, resulting from the fact that whenever $\theta_2 < r$, the outcome changes from bidder 1 winning the good at price r to no sale. Since $F(r) > 1/2$, clearly

$F(r) > F(s) - F(r)$ which implies that the loss from deviation outweighs the gain.

Note that this counterexample also shows that Ausubel and Cramton's (1999) "monotonicity in aggregate" is not actually a necessary condition for incentive-compatibility with resale, as long as agents have no information about each other's values at the time of the auction.

H Proof of Proposition 1

We will prove the following:

Lemma 3. *Let $\bar{\Gamma} = \langle \bar{\chi}, \bar{\psi} \rangle$ be the optimal ACV auction for distributions F_1, \dots, F_n as characterized in (13)–(14). If, for each i , distribution F'_i first-order stochastically dominates distribution F_i , then $R'_n \equiv \mathbb{E}_{\tilde{\theta}}^{F'_1, \dots, F'_n} \left[\sum_i \bar{\psi}_i(\tilde{\theta}) \right] \geq \mathbb{E}_{\tilde{\theta}}^{F_1, \dots, F_n} \left[\sum_i \bar{\psi}_i(\tilde{\theta}) \right] = \bar{R}_n$.*

Together with Theorem 2, the lemma implies that $\text{Rev}(\bar{\Gamma}, v, F'_1, \dots, F'_n) \geq \bar{R}_n$ for any resale procedure v . Since Theorem 3 establishes that any mechanism Γ satisfies $\text{Rev}(\Gamma, \underline{v}, F_1, \dots, F_n) \leq \bar{R}_n$ under the Bertrand resale procedure \underline{v} , the argument of Lemma 1 establishes the result.

We prove the lemma by induction on the number n of bidders. For $n = 0$, $R'_0 = \bar{R}_0 = 0$. Now, suppose the inductive hypothesis holds for $n - 1$ bidders. Let $z(\theta_n) = \mathbb{E}_{\tilde{\theta}_{-n}}^{F_1, \dots, F_{n-1}} \left[\sum_i \bar{\psi}_i(\theta_n, \tilde{\theta}_{-n}) \right]$ and $z'(\theta_n) = \mathbb{E}_{\tilde{\theta}_{-n}}^{F'_1, \dots, F'_{n-1}} \left[\sum_i \bar{\psi}_i(\theta_n, \tilde{\theta}_{-n}) \right]$. Observe that:

- $z(\theta_n) \leq z'(\theta_n)$ for each θ_n . Indeed, for $\theta_n < p_n$, $z(\theta_n) = \bar{R}_{n-1} \leq R'_{n-1} = z'(\theta_n)$ by the inductive hypothesis, and otherwise

$$\begin{aligned} z(\theta_n) &= \mathbb{E}_{\tilde{\theta}_{-n}}^{F_1, \dots, F_{n-1}} \left[\max \left\{ p_n, (\theta_n, \tilde{\theta}_{-n})_{(2)} \right\} \right] \\ &\leq \mathbb{E}_{\tilde{\theta}_{-n}}^{F'_1, \dots, F'_{n-1}} \left[\max \left\{ p_n, (\theta_n, \tilde{\theta}_{-n})_{(2)} \right\} \right] = z'(\theta_n). \end{aligned}$$

- $z(\theta_n)$ is a nondecreasing function. For this, note that $z(\theta_n)$ is independent of θ_n when $\theta_n < p_n$ and nondecreasing in θ_n when $\theta_n > p_n$. But as θ_n moves across p_n , $z(\theta_n)$ jumps from \bar{R}_{n-1} to $R_{n-1}(p_n) + p_n \prod_{j < n} F_j(p_n)$. (The argument is the same as the argument for (18),

the first-order condition for the optimal reserve p_k .) Using (18) (which was shown to hold in Lemma 2), we get

$$R_{n-1}(p_n) + p_n \prod_{j < n} F_j(p_n) - \bar{R}_{n-1} = (p_n - \nu_n(p_n)) \prod_{j < n} F_j(p_n) \geq 0$$

which means that the jump in z is upward, as needed.

Using the two observations, we see that

$$R'_n = \mathbb{E}_{\tilde{\theta}_n}^{F'_n} [z'(\tilde{\theta}_n)] \geq \mathbb{E}_{\tilde{\theta}_n}^{F'_n} [z(\tilde{\theta}_n)] \geq \mathbb{E}_{\tilde{\theta}_n}^{F_n} [z(\tilde{\theta}_n)] = \bar{R}_n,$$

establishing the inductive hypothesis for n .

I Proof of Proposition 2

For convenience, let $\tilde{\theta}'_1, \dots, \tilde{\theta}'_n$ be n random variables, with distributions F_1, \dots, F_n , independent of $\tilde{\theta}_1, \dots, \tilde{\theta}_{n+1}$. We continue to write θ for the profile $(\theta_1, \dots, \theta_n)$ (excluding θ_{n+1}). Also, for brevity we let $i^* = i^*(\theta)$, and \tilde{i}^* for the corresponding random variable.

The expected revenue from the optimal auction on the original n bidders is at most the revenue from the relaxed problem, by maximizing (12), which is simply

$$\mathbb{E}_{\tilde{\theta}}[\max\{\nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}), 0\}]$$

So it suffices to show this latter is bounded above by the revenue from Vickrey with bidders $1, \dots, n+1$.

We have

$$\begin{aligned} \mathbb{E} \left[\max\{\nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}), 0\} \right] &\leq \mathbb{E} \left[\max\{\nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}), \nu_{\tilde{i}^*}(\tilde{\theta}'_{\tilde{i}^*})\} \right] \\ &= \mathbb{E} \left[\nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) + \mathbf{1}_{\nu_{\tilde{i}^*}(\tilde{\theta}'_{\tilde{i}^*}) \geq \nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*})} \left(\nu_{\tilde{i}^*}(\tilde{\theta}'_{\tilde{i}^*}) - \nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) \right) \right] \\ &= \mathbb{E} \left[\nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) + \mathbf{1}_{\tilde{\theta}'_{\tilde{i}^*} \geq \tilde{\theta}_{\tilde{i}^*}} \left(\nu_{\tilde{i}^*}(\tilde{\theta}'_{\tilde{i}^*}) - \nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) \right) \right] \\ &= \mathbb{E} \left[\nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) + \mathbf{1}_{\tilde{\theta}'_{\tilde{i}^*} \geq \tilde{\theta}_{\tilde{i}^*}} \left(\tilde{\theta}_{\tilde{i}^*} - \nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) \right) \right] \\ &\leq \mathbb{E} \left[\nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) + \mathbf{1}_{\tilde{\theta}_{n+1} \geq \tilde{\theta}_{\tilde{i}^*}} \left(\tilde{\theta}_{\tilde{i}^*} - \nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) \right) \right] \\ &= \mathbb{E} \left[\mathbf{1}_{\tilde{\theta}_{n+1} < \tilde{\theta}_{\tilde{i}^*}} \nu_{\tilde{i}^*}(\tilde{\theta}_{\tilde{i}^*}) + \mathbf{1}_{\tilde{\theta}_{n+1} \geq \tilde{\theta}_{\tilde{i}^*}} \tilde{\theta}_{\tilde{i}^*} \right]. \end{aligned}$$

We justify these calculations as follows:

- The first line holds because, conditional on θ , the random variable $\nu_{i^*}(\tilde{\theta}'_{i^*})$ has mean zero, and the function $\max\{\nu_{i^*}(\theta_{i^*}), \cdot\}$ is convex.
- The third line holds because each ν_i is an increasing function.
- The fourth line holds by conditioning on θ , and observing that $\int_z^1 \nu_i(\theta'_i) f_i(\theta'_i) d\theta'_i = \int_z^1 z f_i(\theta'_i) d\theta'_i$, and taking $i = i^*$, $z = \theta_{i^*}$.
- The fifth line holds because, conditional on θ , the event $\tilde{\theta}_{n+1} \geq \theta_{i^*}$ is at least as likely as $\tilde{\theta}'_{i^*} \geq \theta_{i^*}$ (by first-order stochastic dominance).

Finally, the last line is simply the expected revenue from the Vickrey auction on bidders $1, \dots, n+1$: the first term inside the expectation reflects the revenue from bidders $1, \dots, n$ (expressed in terms of virtual values as in (12)), and the last term is the revenue from bidder $n+1$.

J Proof of Proposition 3

Part (a): The revenue from the optimal auction under worst-case resale is bounded above by the relaxed solution from Subsection 4.2. So it suffices to show that the revenue from this relaxed solution exceeds the Vickrey revenue by at most $n \cdot c^{n-1}$. The former can be calculated by assuming the relaxed solution is implemented using threshold payments. Since each auction always yields revenue at most 1, the difference between the two auctions' expected revenue is at most the probability of realizing a type profile θ for which the relaxed solution yields higher revenue ex-post than Vickrey. This can happen only when there is at most one bidder i with $\theta_i > r_i$; this in turn implies there is at most one i with $\theta_i > \max_j r_j$. For each i , the event that $\theta_{i'} \leq \max_j r_j$ for all $i' \neq i$ has probability at most c^{n-1} . Since there are n possible choices of i , the bound $n \cdot c^{n-1}$ follows.

Part (b): Let $\nu(\theta_i), \hat{\nu}(\theta_i)$ be the virtual value functions corresponding to the two distributions. Using the given condition, one can check that $\nu(1) = \hat{\nu}(1)$, $\nu'(1) = \hat{\nu}'(1)$, while $\nu''(1) > \hat{\nu}''(1)$. Hence, $\nu(\theta_i) - \hat{\nu}(\theta_i)$ is quadratic in $1 - \theta_i$, for θ_i near 1. Consequently, by taking a small enough $\varepsilon > 0$, we can assume that

$$\nu(\theta_i) - \hat{\nu}(\hat{\theta}_i) > \varepsilon(1 - \theta_i)^2 \quad (37)$$

whenever

$$1 - \varepsilon < \theta_i < \hat{\theta}_i < \theta_i + \varepsilon^2(1 - \theta_i)^2.$$

Fix such an ε .

Let $i^{**}(\theta)$ denote the bidder with the highest *virtual* value at the type profile θ . Since the revenue from an auction is given by the expected virtual value of the winning bidder, in comparing the revenue from the Myerson auction (which awards the good to $i^{**}(\theta)$ whenever his virtual value is positive) against the Vickrey auction (which awards to $i^*(\theta)$), it suffices to show that the virtual value of bidder $i^{**}(\theta)$ exceeds that of $i^*(\theta)$ by at least c'/n^3 on average.

For any realization of θ , let $\theta_{i^*[F]}$ denote the highest of the $n/2$ values drawn from distribution F , and $\theta_{i^*[\hat{F}]}$ similarly. Write $\xi = 1 - \theta_{i^*[F]}$, and similarly $\hat{\xi}$.

Consider the probability that

- (i) $\xi \in (1/n - \varepsilon/n, 1/n)$; and moreover
- (ii) $\xi - \hat{\xi} \in (0, \varepsilon^2 \xi^2)$.

By standard extreme value theory, as $n \rightarrow \infty$, the distributions of $n\xi$ and $n\hat{\xi}$ converge to exponential distributions (see e.g. Falk, Hüsler, and Reiss (2010), Theorem 2.1.1(ii) with $\beta = -1$). Event (i) is equivalent to $n\xi \in (1 - \varepsilon, 1)$, hence its probability is bounded below by a positive constant as $n \rightarrow \infty$. Similarly, one can write event (ii) as $n\hat{\xi} \in (n\xi - \varepsilon^2(n\xi)^2/n, n\xi)$; so conditional on (i), its probability is bounded below by a constant times $1/n$. And when both (i) and (ii) hold, the bidder with highest value is $i^*[\hat{F}]$, while the one with highest virtual value is $i^*[F]$, and (37) implies that their virtual values differ by an amount on the order of $1/n^2$. The assertion follows.