Theoretical Economics 15 (2020), 1–28 1555-7561/20200001

Discriminatory price auctions with resale and optimal quantity caps

Brian Baisa
Department of Economics, Amherst College

Justin Burkett
School of Economics, Georgia Institute of Technology

We present a model of a discriminatory price auction in which a large bidder competes against many small bidders, followed by a post-auction resale stage in which the large bidder is endogenously determined to be a buyer or a seller. We extend results on first-price auctions with resale to this setting and use these results to give a tractable characterization of equilibrium behavior. We use this characterization to study the policy of capping the amount that may be won by large bidders in the auction, a policy that has received little attention in the auction literature. Our analysis shows that the trade-offs involved when adjusting these quantity caps can be understood in terms familiar to students of asymmetric first-price single-unit auctions. Furthermore, whether one seeks to maximize welfare or revenue can have contradictory implications for the choice of cap.

Keywords. Multi-unit auctions, resale, quantity caps.

JEL classification. D42, D44, D47, D82.

1. Introduction

In multi-unit auctions, bidders who demand a nonnegligible fraction of the units being auctioned may use their market power to influence the allocation and payments. When a post-auction resale market exists, the auction may be used as an instrument to obtain market power in the resale market as well. This possibility played out in the Salomon Brothers scandal in 1991. Salomon Brothers admitted to violating US Treasury auction rules and controlling almost 94% of a single issue of two-year notes. They then purportedly used their market power to implement a "short squeeze" in the secondary market, pushing the yields of these notes significantly below prevailing rates and triggering a Securities and Exchange Commission (SEC) investigation (Jegadeesh, 1993, Brady et al., 1992). Salomon Brothers subverted the Treasury's rule that restricts the amount a single bidder may win.[1]

Brian Baisa: bbaisa@amherst.edu
Justin Burkett: justin.burkett@gatech.edu

We thank the anonymous referees for their helpful comments and suggestions.

[1]This policy currently restricts bidders to winning at most 35% of the market supply, but it has evolved over the course of the 20th century, with its size varying between 25% and 35% (Garbade and Ingber, 2005). Bartolini and Cottarelli (1997) find in their survey of "treasury" auctions around the world that 23% of the countries in their sample impose a ceiling on auction awards.

 https://doi.org/10.3982/TE2939

Although they are used in prominent multi-unit auctions, the question of when and how a quantity cap should be set has received little attention in the literature.[2] In this paper we make two main contributions. First, we construct a model of a discriminatory price auction with subsequent resale market and explicitly characterize equilibrium bidding and resale behavior for any initial choice of quantity cap. We then evaluate the seller's optimal choice of quantity cap for two different objectives: the expected welfare of the allocation following resale and the expected revenue generated by the auction itself. Ostensibly quantity caps are intended to reduce the deadweight loss resulting from a large bidder holding too many units following resale; however, the US Treasury has also indicated that revenue maximization, technically cost minimization, is important in Treasury bill auctions.[3]

In our model, a single large bidder with downward sloping multi-unit demand competes against a "continuum" of small bidders in a discriminatory price auction with resale. The small bidders are heterogeneous and have private values. Each small bidder demands a single infinitesimal unit. All bidders are forward looking and anticipate that in the resale stage, the large bidder will adjust the amount it owns by acting as a single-price monopsonist or monopolist, depending on the auction outcome. The auctioneer can restrict the amount the large bidder wins in the auction by setting a quantity cap κ, but she cannot restrict holdings in the resale market. We assume the bidders' marginal values are privately known and determined by single-dimensional random variables.

While there are important differences between the large and small bidders' optimization problems, they face similar trade-offs when determining their equilibrium bids. We focus on equilibria where the small bidders use monotone pure strategies in the auction. In such equilibria, the large bidder submits a bid curve that is constant in the quantity won—a flat bid—because its residual supply curve in the auction is deterministic. Small bidders with sufficiently high values win a unit with probability 1, while the remaining "competitive" small bidders win a unit if and only if their bid exceeds the large bidder's flat bid. We construct equilibrium strategies using the observation that for both the large bidder and the competitive small bidders, the marginal value of increasing their bid is determined by the anticipated resale price and not their private value. When the large bidder increases its bid slightly, it values the additional quantity that it wins in the auction at the resale price because holding the additional quantity causes the large bidder to reduce the amount it buys (or increase the amount it sells) at the resale price following the auction. Similarly, when a small bidder increases her bid by

[2] In split award auctions, the seller decides whether to split a contract award in two and effectively restricts the ability of a firm to win the entire contract (Anton and Yao, 1992, Gong et al., 2012). Back and Zender (1993) briefly consider quantity caps prior to their Theorem 3 in an model without resale. Insofar as quantity caps are important for manipulating the resale market, it is important to explicitly model the resale market to consider the effect of quantity caps.

[3] Garbade and Ingber (2005) reports that minimizing the cost of funds is the auction objective, and cites a speech where this is stated by the Under Secretary (see footnote 1 of their paper). Other statements from Federal Reserve (Fed) personnel are more ambiguous. For example, following the Salomon Brothers scandal, Fed Vice Chairman David Mullins was quoted on page A1 of the August 26, 1991 *Wall Street Journal* as saying "We need to examine mechanisms to improve the efficiency of the market, [and] reduce the cost of Treasury finance."

a small amount and that increase is pivotal, she either purchases a unit in the auction that she would have purchased in the resale market or resells the purchased unit in the resale market. Consequently, equilibrium bids for both kinds of bidders are determined by a common distribution of resale prices. Furthermore, the large and competitive small bidders' bids are distributed symmetrically.[4] This symmetry leads to a simple characterization of equilibrium bidding.

The quantity cap's influence on expected welfare and expected auction revenue is determined by the strength of the large bidder's type distribution relative to the competitive small bidders' type distribution. Adjusting the cap changes the distributional strength of the large bidder relative to the competitive small bidders. Tightening the cap, for example, weakens the distribution of competitive small bidder values because more small bidders win with probability 1. For welfare, we show that tightening the cap increases welfare when the large bidder is weak relative to the competitive small bidders. An analogy to asymmetric first-price auctions helps to explain this result. Recall that a weak bidder in a first-price auction bids aggressively and wins more frequently than she would in an efficient auction (Maskin and Riley, 2000). Similarly in our setting, when the large bidder is weak relative to the competitive small bidders, it bids more aggressively and wins more than the welfare maximizing quantity in the auction. Although the large bidder resells some of this excess quantity, it restricts the amount it sells due to its market power. A tighter cap makes the large bidder stronger in the auction by weakening the distribution of competitive small bidders. Consequently, the tighter cap reduces both the quantity won by the large bidder in the auction and the quantity retained by the large bidder following resale.

Alternatively, tightening the cap reduces expected auction revenue when the large bidder is weak relative to the competitive small bidders and an additional regularity condition holds. Kirkegaard (2012) introduced this regularity condition to rank the revenue of first- and second-price auctions. Intuitively, a first-price auction distorts the allocation toward the aggressive weak bidder. Under Kirkegaard's condition, the weaker bidder has a higher virtual valuation at the equilibrium allocation. This distortion increases revenue since expected revenue is equal to the expected virtual valuation of the winner (Myerson, 1981). In our setting, tightening the cap distorts the allocation away from the large bidder. We show that under an adapted version of Kirkegaard's condition, the large bidder has the higher virtual valuation on the margin when it is weak. Hence, the distortion induced by the tighter cap reduces revenue when the large bidder is weak relative to the competitive small bidders.

One might expect that tightening the cap always reduces revenue, because tightening the cap always reduces the number of bids received—a measure of the level of competition in the auction. We show that this is not the case. Instead, when the large bidder is relatively strong, the logic above reverses and tightening the cap can increase revenue. Intuitively, a strong large bidder has a lower virtual valuation, so that revenue can be improved by tightening the cap and distorting the allocation away from the large

[4]This "symmetrization" effect of post-auction resale was first observed in the literature on first-price auctions with resale (Gupta and Lebrun, 1999, Hafalir and Krishna, 2008). Cheng and Tan (2010) observe bidders to behave as if they have a common value in the context of a first-price auctions with resale.

bidder. The decrease in revenue associated with tightening the cap and reducing the number of bids received by the auctioneer is offset by more aggressive bids from the competitive small bidders.

Related work

Most studies of discriminatory auctions assume resale is not possible. In the case without resale, Pycia and Woodward (2017) provides an equilibrium existence result for discriminatory auctions, building on earlier work by Wang and Zender (2002), Holmberg (2009), and Ausubel et al. (2014). Ausubel et al. (2014) compares the revenue and efficiency of discriminatory and uniform price auctions in a setting where bidders have multi-unit demands. Back and Zender (1993) and Wang and Zender (2002) study the revenue ranking of these two formats in models with common values. Swinkels (2001) studies revenue and efficiency rankings in large markets. One implication of this body is that equilibrium bidding in multi-unit auctions is inherently more complex than in single-unit settings. In uniform price auctions, each bidder strategically reduces her bids for later units to reduce the expected price paid for earlier units (Ausubel et al., 2014). In discriminatory auctions, bidders would often benefit from bidding less for their earlier units than they do for their later units, but they cannot because submitted demand must be weakly downward sloping (Woodward, 2016). In either case, these strategic effects introduce dependencies between each bidder's unit-wise bids that can be complex to disentangle.

In Baisa and Burkett (2018), we show how to characterize equilibrium bidding in a model of a large bidder competing against a continuum of small bidders by showing that the optimality conditions for marginal bids correspond to those from a first-price auction. In that paper, we compare the revenue and efficiency of discriminatory and uniform price auctions. In the current paper, we introduce resale into a discriminatory price auction and use insights from the first-price auctions with resale literature to characterize equilibrium. Our model allows for downward-sloping demand, and we use our model to study optimal quantity cap policies. Neither feature has a natural analog in a single-unit setting.

Several papers study models of discriminatory auctions with a subsequent resale stage. Hafalir and Kurnaz (2018) study a discrete model in which bidders demand only a single unit. In the market microstructure literature, Viswanathan and Wang (2004) develop a model in which a divisible good is first auctioned to a primary dealer. Following the auction, the winning dealer may resell units to other dealers. As in our model, a single large bidder controls all of the units following the auction. However, there are a couple of important distinctions with our paper. First, while Viswanathan and Wang (2004) consider a discriminatory price trading procedure with a resale market, the auction awards are all-or-nothing and are effectively single-unit auctions. Second, there is no private information among the dealers in the auction.[5]

[5]The winning dealer learns the amount of supply in the auction, which is uncertain ex ante, and this information becomes private information in the resale stage.

Coutinho (2013) studies the role of speculation in a uniform-price auction with resale where bidder preferences are common knowledge. There are two types of bidders: final investors and pure speculators. The latter group has no value for the good. Coutinho (2013) shows that the presence of speculation via resale has an ambiguous effect on auction revenue.

Results from the literature on first-price auctions with resale play an important role in our paper. Our characterization of equilibrium builds off the work on first-price auctions with resale in Gupta and Lebrun (1999), Hafalir and Krishna (2008), and Cheng and Tan (2010). Each of these papers uses the observation that resale leads to symmetric bid distributions between two bidders. Despite the similarity, the model in our paper is not equivalent to a first-price auction because a number of key features of our model are inherent to multi-unit auction models. Notably, the large bidder has downward sloping demand and we study the use of a quantity cap. Garratt and Tröger (2006) study an instance of pure speculation in single-unit auctions with resale. In their model, a speculator has zero value for the good, yet in equilibrium, the speculator purchases the good at auction and resells it to a bidder who has positive demand for the good.

We also use results from the single-unit auctions literature on the revenue ranking of first- and second-price auctions without resale (Kirkegaard, 2012, Maskin and Riley, 2000). Kirkegaard (2012) introduces a regularity condition to generalize the conditions given in Maskin and Riley (2000) under which a first-price auction raises more expected revenue than a second-price auction. We show that relaxing the quantity cap on the large bidder increases revenue under an adapted version of Kirkegaard's condition.

Organization of paper

The next section introduces the model. Section 3 characterizes equilibrium strategies, Section 4 uses this characterization to analyze the problem of finding the welfare maximizing cap, and Section 5 relates this cap to the revenue maximizing cap. Section 6 illustrates our results with an example. In Section 7, we show that allowing for additional speculation by small bidders does not change our results. Section 8 concludes. The Appendix contains proofs and supporting calculations.

2. Model

A large bidder competes against a unit measure of small bidders for a unit measure of a divisible good.[6] Bidders have private values. The large bidder has demand for a positive measure of the good, while small bidders are heterogeneous and each demands an infinitesimal unit of the good.

Payoffs

The small bidders' private values are distributed according to the increasing and absolutely continuous distribution function, $F_S(\theta_S)$, with support $[\underline{\theta}_S, \overline{\theta}_S]$, where $\overline{\theta}_S > \underline{\theta}_S \geq 0$.

[6]We extend the analysis to the case in which the measure of small bidders exceeds the measure of the good in a prior working paper (Baisa and Burkett (2019)).

Each $\theta_S \in [\underline{\theta}_S, \overline{\theta}_S]$ represents a small bidder with a private value of θ_S for holding an infinitesimal unit of the good. Thus, the type-θ_S small bidder's payoff is

$$q\theta_S - t,$$

when she finishes the game with $q \in \{0, 1\}$ infinitesimal units and makes the net payment $t \in \mathbb{R}$. She derives no additional utility from holding more than one infinitesimal unit.

We assume that F_S is regular in the sense that both $x - (1 - F_S(x))/f_S(x)$, the virtual value of a small bidder buying a unit, and $x + F_S(x)/f_S(x)$, the virtual value of a small bidder selling a unit, are increasing functions of x. These conditions ensure the uniqueness and monotonicity of the resale price chosen by the large bidder.[7]

The large bidder has a single-dimensional type, $\theta_L \in \Theta_L \equiv [\underline{\theta}_L, \overline{\theta}_L]$ with $\overline{\theta}_L > \underline{\theta}_L \geq 0$, which is distributed according to the absolutely continuous distribution function $F_L(\theta_L)$. The type-θ_L large bidder's marginal value from holding the qth increment of the good is $v(q, \theta_L)$. For all $q \geq 0$ and $\theta_L \in \Theta_L$, $v(q, \theta_L)$ is nonincreasing and continuous in q, increasing and differentiable in θ_L, and nonnegative. Without loss of generality, we assume that the large bidder's type determines the marginal value at $q = 0$ (i.e., $v(0, \theta_L) = \theta_L$ for all $\theta_L \in \Theta_L$). Thus, the type-θ_L large bidder's payoff is

$$\int_0^q v(x, \theta_L)\, dx - t$$

when it makes the net payment $t \in \mathbb{R}$ and holds the quantity q following the game. We often refer to the flat demand case of our model. In the flat demand case, the large bidder has a constant marginal value for the good (i.e., $v(q, \theta_L) = \theta_L$ for all $q \in [0, 1]$, $\theta_L \in \Theta_L$).

We assume the relation between the supports of the type distributions is such that $\underline{\theta}_S \leq \underline{\theta}_L$ and $v(1, \overline{\theta}_L) \leq \overline{\theta}_S - (\kappa - 1)/f_S(\overline{\theta}_S)$. This assumption ensures that all large bidder types choose an interior resale price. Relaxing this assumption leads to pooling of large bidder types but not to significant changes in our results as we show in a prior working paper (Baisa and Burkett (2019)).

Auction stage

We study a standard discriminatory price (or "pay-as-bid") rule with a commonly known quantity cap for the large bidder, $0 < \kappa \leq 1$, which prevents the large bidder from bidding for more than a fraction κ of the good. After bids are received, bids are awarded in

[7] The relevant conditional virtual valuations are also increasing under this assumption. Specifically,

$$x - \frac{1 - F_S(x|x < y)}{f_S(x|x < y)} = x - \frac{F_S(y) - F_S(x)}{f_S(x)}$$

is increasing in x for $x < y$, and

$$x + \frac{F_S(x|x > y)}{f_S(x|x > y)} = x + \frac{F_S(x) - F_S(y)}{f_S(x)}$$

is increasing in x for $x > y$. These facts are noted in Cheng and Tan (2010), who also observe that a unique and monotone resale price would exist under weaker conditions on the virtual values.

declining order until the supply is exhausted. The rules implicitly require that bid curves be nonincreasing. In the main model, small bidders are only allowed to bid for a single infinitesimal unit in the auction. That is, they cannot speculate with additional bids. We relax this assumption in Section 7.

The auction allocates the good by comparing bids to the lowest winning bid. If the type-θ_L large bidder submits the nonincreasing bid curve $b_L(q, \theta_L)$ for $q \leq \kappa$, then define the large bidder's quantity demanded at b as

$$q_L(b; b_L) \equiv \sup\{q | b_L(q, \theta_L) \geq b \text{ and } q \leq \kappa\}.$$

Let $G(b)$ denote the measure of small bidders who place a bid b' such that $b' \leq b$. The lowest winning bid, $b^\star$, is then

$$b^\star = \sup\{b | q_L(b; b_L) + 1 - G(b) \geq 1\}.$$

If the set $\{b | q_L(b; b_L) + 1 - G(b) \geq 1\}$ is empty, we set $b^\star = 0$. If the set $\{b | q_L(b; b_L) + 1 - G(b) = 1\}$ has nonzero measure, we assume pro rata rationing of marginal bids as in Back and Zender (1993), although the rationing rule plays no role in the results. All small bidders with bids exceeding $b^\star$ receive a unit, while the large bidder receives $q(b^\star, b_L)$.

Bidders pay the auctioneer according to their submitted bids. If the large bidder submits bid curve $b_L(q, \theta_L)$ and wins quantity q in the auction, it pays the auctioneer

$$\int_0^q b_L(x, \theta_L)\, dx.$$

If a small bidder bids b and wins an infinitesimal unit, she pays b; otherwise, she pays zero.

We assume the auctioneer reveals the lowest winning bid $b^\star$ following the auction. This assumption does not play an important role in our results (see footnote 8).

Resale stage

In the resale stage, the large bidder announces a single take-it-or-leave-it price at which it is willing to trade with the small bidders. The large bidder offers either a price at which it is only willing to buy or a price at which it is only willing to sell. There is no value to being able to offer both a buy price and a sell price in equilibrium. The large bidder's decision to buy or sell is endogenous and depends on the relation between the quantity purchased in the auction and the welfare maximizing quantity. We do not assume that the quantity cap is enforced in the resale market; hence, the large bidder can finish the game with a quantity greater than κ following the resale stage.

3. Equilibrium

We characterize equilibrium in nondecreasing strategies. This equilibrium is unique in the class of equilibria in which players use nondecreasing strategies and the small

bidders all participate in the auction. We first focus on the large bidder's problem and argue that its bid curve is constant in the quantity purchased and increasing in its type. Under the assumption that the large bidder bids in this manner, we then present the bidders' objective functions. Third, we give an expression for the resale price chosen by the large bidder. Finally, we argue that the distributions of the large bidder's flat bid and the competitive small bidders' bids must be symmetric in equilibrium. This result leads directly to our equilibrium characterization in Proposition 1.

The quantity allocated to the large bidder in the auction is determined by comparing its bid curve to the bids made by the competitive small bidders. When the size of the cap is κ and small bidders' bids are nondecreasing in their types, the $(1-\kappa)$ measure of small bidders with the highest values win with probability 1 in the auction as long as there are no ties between marginal bids. The remaining small bidders have types θ_S such that $F_S(\theta_S) \leq \kappa$. For these small bidders, the relation between their bid and the large bidder's bid determines whether they win a unit. We call these small bidders the competitive small bidders because they compete directly with the large bidder.

We illustrate the environment from the large bidder's perspective in Figure 1, in which we graph the large bidder's demand against the residual supply. The expression $q = F_S(\theta_S)$ is the residual supply curve in the sense that if the large bidder defeats small bidders with types below θ_S in the auction, it wins the quantity $F_S(\theta_S)$.

First, we observe that the large bidder's best response to an increasing bid strategy used by the competitive small bidders is to submit a flat bid curve. To see this, suppose the competitive small bidders bid according to the increasing function $b_S(\theta_S)$. If the large bidder bids a constant amount b for $q \leq \kappa$, it wins the quantity $F_S(\theta'_S)$, where $\theta'_S = \sup\{\theta_S | b_S(\theta_S) < b\}$. This quantity is deterministic given b, so the large bidder lowers her payoff by bidding more than b for any quantity $q < F_S(\theta'_S)$. Hence, a flat bid is optimal and we describe the large bidder's strategy in the auction by $b_L(\theta_L)$, which gives the value of the flat bid for all $0 \leq q \leq \kappa$.

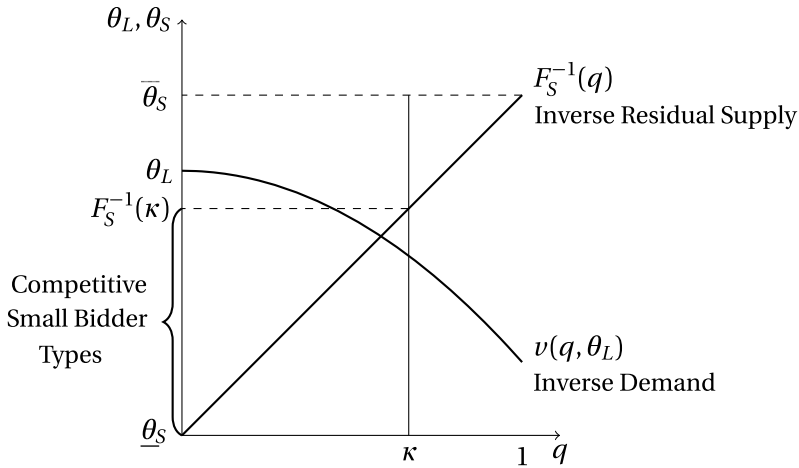

FIGURE 1. The environment from the large bidder's perspective.

Next, we determine the (interim) expected payoff $\pi_S(\theta_S, b)$ of a type-θ_S small bidder who bids b. We assume that the large bidder uses a flat bid and its behavior is captured by the functions $\phi_L(b)$ and $p(\theta_L)$. The former is the large bidder's inverse bid function or the type of the large bidder that bids b, while the latter is the resale price set by the large bidder with type θ_L. We assume both functions are increasing and we verify this in Proposition 1. The type-θ_S small bidder only wins in the auction if her bid exceeds the large bidder's flat bid. She sells (buys) in the resale market, having won (not won) a unit in the auction, if the large bidder buys (sells) at a resale price above (below) her value θ_S. Thus her expected payoff from bidding b in the auction is

$$\begin{aligned}\pi_S(\theta_S, b) = {} & \underbrace{F_L\big(\phi_L(b)\big)(\theta_S - b)}_{\text{auction}} \\ & + \underbrace{\int_{\underline{\theta}_L}^{\phi_L(b)} \big(p(t) - \theta_S\big)\mathbb{1}\big\{v\big(F_S\big(p(t)\big), t\big) \geq p(t) \geq \theta_S\big\}\, dF_L(t)}_{\text{sale post auction}} \\ & + \underbrace{\int_{\phi_L(b)}^{\overline{\theta}_L} \big(\theta_S - p(t)\big)\mathbb{1}\big\{v\big(F_S\big(p(t)\big), t\big) \leq p(t) \leq \theta_S\big\}\, dF_L(t)}_{\text{purchase post auction}}, \end{aligned} \tag{1}$$

where the braces indicate the sources of the respective terms. The indicator function $\mathbb{1}\{v(F_S(p(t)), t) \geq p(t) \geq \theta_S\}$ equals 1 if the type-t large bidder purchases additional units in the resale market at price $p(t)$ and the type-θ_S small bidder is willing to sell at this price. The indicator function $\mathbb{1}\{v(F_S(p(t)), t) \leq p(t) \leq \theta_S\}$ captures the reverse situation.

We similarly determine the large bidder's payoff $\pi_L(\theta_L, b, p)$ when its type is θ_L, it places the flat bid b, and it sets the resale price p. We let $b_S(\theta_S)$ represent the small bidders' bid strategy. We assume that $b_S(\theta_S)$ is increasing for $\theta_S < F_S^{-1}(\kappa)$ and that its inverse is $\phi_S(b)$ for $b \in (b_S(\underline{\theta}_S), b_S(F_S^{-1}(\kappa)))$. This assumption is verified in the proof of Proposition 1. Thus, the large bidder wins the quantity $F_S(\phi_S(b))$ in the auction when it bids $b < b_S(F_S^{-1}(\kappa))$ for all $q \leq \kappa$. The large bidder's payoff is

$$\pi_L(\theta_L, b, p) = \int_0^{F_S(p)} v(x, \theta_L)\, dx - \underbrace{F_S\big(\phi_S(b)\big)b}_{\text{auction-payment}} - \underbrace{\big(F_S(p) - F_S\big(\phi_S(b)\big)\big)p}_{\text{resale payment}}, \tag{2}$$

regardless of whether the large bidder buys or sells in the resale market. The large bidder retains the quantity $F_S(p)$ following resale, pays b for the quantity $F_S(\phi_S(b))$ in the auction, and pays or receives p for each unit traded in the resale market.

3.1 *Resale stage*

The large bidder chooses a price–quantity pair from the residual supply curve in the resale market, $q = F_S(p)$, based on its type θ_L and the quantity won in the auction,

$F_S(\phi_S(b))$. The price–quantity pair is determined by a standard monopoly pricing formula. Specifically, the first-order condition of the large bidder's objective with respect to p implies

$$v\big(F_S(p), \theta_L\big) = p - \frac{F_S\big(\phi_S(b)\big) - F_S(p)}{f_S(p)}. \tag{3}$$

This equation holds when an interior choice of resale price is optimal for the large bidder (i.e., the optimal resale price p is such that $\underline{\theta}_S < p < \overline{\theta}_S$). The regularity conditions on F_S (see footnote 7) along with the assumption that $v(q, \theta)$ is nonincreasing in q implies that this p is unique and that the first-order condition is sufficient for optimality given (b, θ). It is routine to show that the resulting $p(b, \theta)$ is increasing in both arguments.

Notice that the deviations of individual small bidders cannot influence the large bidder's allocation in either stage because the small bidders are infinitesimal. Furthermore, the information about bids provided to the small bidders following the auction does not influence the resale allocation, because they are price takers in the resale market.[8]

3.2 *Auction stage*

A key step toward characterizing equilibrium is to show that the large and competitive small bidders effectively face the same trade-off when choosing bids. This implies that there is an equilibrium in which the bid distributions of the large and competitive small bidders are symmetric. Gupta and Lebrun (1999) and Hafalir and Krishna (2008) study first-price auctions with resale and identify a similar symmetrization effect caused by resale. We give an intuitive argument for why symmetrization applies to our multi-unit setting as well.

First, we argue that the large bidder's value of a marginal increase in the amount won in the auction is equal to the resale price. Consider a large bidder who bids b in the auction and purchases additional quantity in the resale market at price p.[9] First, note that the resale price must be larger than the bid, $p \geq b$, because otherwise the large bidder could increase its payoff by buying less in the auction and more in the resale market. Next, suppose the large bidder increases its bid by a small amount without changing the resale price. The large bidder values the additional amount purchased in the auction at p because it would have purchased this amount in the resale market. Our conclusion does not change if we allow the resale price to increase, as it would if the large bidder buys additional quantity in the auction. The envelope theorem implies that there is no additional net effect on the large bidder's payoff resulting from the increase in the resale price.

[8] In Hafalir and Krishna (2008), the policy for revealing information following the auction is important. In fact, revealing the losing bid prevents an increasing equilibrium from existing at all (see Remark 1 in Hafalir and Krishna (2008)). We do not have a similar requirement about the information policy used by the auctioneer.

[9] The argument is easily modified if this large bidder instead sells in the resale market. The distributions of large and small bidder types as well as the value of the cap determine whether a given type of large bidder is a buyer or seller. Our analyses of the welfare and revenue maximizing caps in Sections 4 and 5 allow for the possibility that the large bidder may buy or sell in the resale market.

Similarly, each competitive small bidder's marginal value from an increase in her bid is determined by the anticipated resale price. Consider a small bidder who submits the same bid b as the large bidder in the previous paragraph. Suppose that this small bidder bids $b+\varepsilon$ instead. This bid increases the amount won by the small bidder in the event that the large bidder's flat bid falls in the interval $(b, b+\varepsilon)$. In this event, the large bidder purchases additional quantity in the resale market at a price of approximately p by the assumption of the previous paragraph. The small bidder resells her unit in this event, because she has the lowest private value among small bidders who win a unit in the auction. Therefore, for the large and small bidder types who bid b, the marginal value of an increase in either of their bids is p.

A consequence of the large and competitive small bidders having the same marginal value in the auction is that the distributions of auction bids must be symmetric. Formally, for any bid that wins with probability between 0 and 1,

$$\frac{1}{\kappa}F_S\big(\phi_S(b)\big) = F_L\big(\phi_L(b)\big), \tag{4}$$

where ϕ_S and ϕ_L are the respective inverse bid functions of the competitive small and large bidders. For a given type of large bidder, this identity determines the type of small bidder that makes the same bid. Specifically, it implies that the type-θ_L large bidder wins the quantity $F_S(\phi_S(b_L(\theta_L))) = \kappa F_L(\theta_L)$ in the auction, where we use $b_L(\theta_L)$ to represent the equilibrium choice of bid.

Observation 1. *The type-θ_L large bidder wins the quantity $\kappa F_L(\theta_L)$ in the auction.*

The quantity won by each large bidder type is independent of the shape of its demand curve and increasing in the size of the quantity cap. These properties are consequences of the symmetric dependence of bids on the resale price. Reducing the large bidder's marginal values without changing its type distribution (weakly) reduces the resale price and, consequently, the auction bids. However, the auction allocation is unchanged, because the decrease in bids is symmetric across the large and competitive small bidders. Relaxing the cap causes the large bidder to compete with a larger selection of small bidders in the auction. Intuitively, to maintain symmetry of the bid distributions across the large and competitive small bidders requires that *every* large bidder type win against more small bidders, regardless of whether the quantity cap binds. The example in Section 6 is rich enough to examine both of these comparative statics.

We use Observation 1 and the resale pricing formula (3) to determine the resale price set by a type-θ_L large bidder, $p(\theta_L)$. This function is implicitly defined by

$$v\big(F_S\big(p(\theta_L)\big), \theta_L\big) = p(\theta_L) - \frac{\kappa F_L(\theta_L) - F_S\big(p(\theta_L)\big)}{f_S\big(p(\theta_L)\big)}. \tag{5}$$

Note that (5) specifies the resale price and all bidders' final allocations in terms of primitives of the model. It is straightforward to verify that $p(\theta_L)$ is increasing in θ_L and κ. We define the distribution of resale prices using $p(\theta_L)$ as

$$F(x) \equiv \Pr\big\{p(\theta_L) \le x\big\} = F_L\big(p^{-1}(x)\big), \tag{6}$$

where $p^{-1}(x)$ represents the type of large bidder that sets the resale price x. Note that we consider $p(\theta_L)$ to be a random variable in this definition.

The distribution of resale prices determines equilibrium bidding behavior. We argue above that the marginal value of a unit won in the auction derives from the anticipated resale price. This suggests that equilibrium bids can be derived from the distribution of resale prices. In fact, the equilibrium bid of the large bidder type that sets the resale price p is equal to the equilibrium bid made by a bidder with value p in a symmetric first-price auction between two bidders, each with values distributed according to $F(x)$. One way to express this bid is

$$b(x) = \mathrm{E}_{\theta_L}\big[p(\theta_L) | p(\theta_L) \leq x\big], \tag{7}$$

where $p(\theta_L)$ represents the random resale price. Finally, if we use $\phi(b)$ for the inverse of $b(x)$, (4), (6), and (7) require that

$$F_L\big(\phi_L(b)\big) = \frac{1}{\kappa} F_S\big(\phi_S(b)\big) = F\big(\phi(b)\big), \tag{8}$$

which equates the distributions of the large and small bidders' bids to that of the hypothetical bidder in the symmetric first-price auction with value distribution $F(x)$. To see that (8) holds, observe that the large bidder type bidding b sets a resale price $\phi(b)$ and, hence, has the type $p^{-1}(\phi(b))$. Noting (6), this implies $F(\phi(b)) = F_L(\phi_L(b))$, while (4) implies the second equality in expression (8). Proposition 1 summarizes equilibrium behavior, and its proof formally verifies the formulation described above.

Proposition 1. *In equilibrium, the symmetrization identity, (8), holds and the resale price is determined by the pricing formula in (5). The lowest and highest resale prices are* $\underline{p} = p(\underline{\theta}_L)$ *and* $\overline{p} = p(\overline{\theta}_L)$*, respectively. Note that* $F(\underline{p}) = F_L(\underline{\theta}_L) = 0$ *and* $F(\overline{p}) = F_L(\overline{\theta}_L) = 1$*. Bid strategies are*

$$b_L(\theta) = b\big(F^{-1}\big(F_L(\theta)\big)\big)$$

$$b_S(\theta) = \begin{cases} b\left(F^{-1}\left(\frac{1}{\kappa} F_S(\theta)\right)\right) & F_S(\theta) \in [0, \kappa] \\ b(\overline{p}) & F_S(\theta) > \kappa. \end{cases}$$

Equilibrium strategies are unique among equilibria in which bidders use nondecreasing bidding strategies and all small bidders participate in the auction.

We construct the equilibrium bid functions using the symmetry of the bid distributions in (8) and the observation that the bid of the small bidders who win with probability 1 must equal the highest auction bid. We solve for the equilibrium bid functions for a parametric example in Section 6.

Proposition 1 extends the symmetrization result of Gupta and Lebrun (1999) and Hafalir and Krishna (2008) to our multi-unit setting. Furthermore, the proposition shows how to extend symmetrization to account for downward-sloping demand and quantity caps. Neither feature has a direct analog in a single-unit setting. We establish

uniqueness by extending an argument due to Hafalir and Krishna (2008) to our setting. Like Hafalir and Krishna (2008), we limit attention to equilibria in nondecreasing strategies. We also add the qualification that all small bidders participate in the auction. In equilibrium, a small bidder may be indifferent between participating in both the auction and the resale stages, and only participating in the resale stage. With this qualification, we ignore alternative equilibria, such as those in which a zero measure set of small bidders skips the auction.

4. Welfare maximizing caps

We next consider the influence of the cap on the welfare of the final allocation. Intuitively, if the large bidder purchases more than the welfare maximizing quantity in the auction, it resells some of its quantity in the resale market. However, due to its market power, it retains an amount above the welfare maximizing quantity causing deadweight loss. Tightening the quantity cap reduces the amount won in the auction and the amount held in the final allocation by all types of large bidder. This improves welfare whenever all types of large bidders purchase more than the welfare maximizing quantity in the auction. The large bidder wins more than the welfare maximizing quantity in the auction because it bids more aggressively than the competitive small bidders. We derive conditions that determine when the large bidder bids too aggressively and relate these conditions to well known conditions from the first-price auctions literature.

We first determine when the cap should be tightened in the flat demand case. The type-θ_L large bidder wins the quantity $\kappa F_L(\theta_L)$ in the auction (Observation 1). When the large bidder has flat demand, the welfare maximizing quantity is $F_S(\theta_L)$. This quantity is determined by the intersection of the large bidder's inverse demand, which is equal to θ_L for all q, and its inverse residual supply, which is $F_S^{-1}(q)$, in Figure 1. When the auction quantity exceeds the welfare maximizing quantity, $\kappa F_L(\theta_L) \geq F_S(\theta_L)$, the large bidder is a seller in the resale market. The large bidder's post-resale quantity, $F_S(p(\theta_L))$, is between the auction quantity and the welfare maximizing quantity because of its market power. In this case, tightening the cap reduces the amount won by this type of large bidder in the auction and pushes the final allocation closer to the welfare maximizing quantity by reducing $F_S(p(\theta_L))$. Note that $p(\theta_L)$ is increasing in κ because it is increasing in the quantity won in the auction.

Tightening the cap increases expected welfare when all large bidder types win more than the welfare maximizing quantity in the auction. In the flat demand case, this occurs when the large bidder's type distribution is weak relative to the competitive small bidder's type given κ. In such cases, we say that the large bidder is *weak at* κ, where being weak at κ is defined in terms of a first-order stochastic dominance relationship.

Definition 1. The large bidder is weak at κ if $\kappa F_L(\theta_L) \geq F_S(\theta_L)$ for all $\theta_L \in \Theta_L$ and $\kappa F_L(\theta_L) > F_S(\theta_L)$ for some $\theta_L \in \Theta_L$. The large bidder is strong at κ if $\kappa F_L(\theta_L) \leq F_S(\theta_L)$ for all $\theta_L \in \Theta_L$ and $\kappa F_L(\theta_L) < F_S(\theta_L)$ for some $\theta_L \in \Theta_L$.

The intuition for why the cap should be tightened when the large bidder is weak at κ relates to well known results from the asymmetric first-price auctions literature. Intuitively, when the large bidder is weak at κ, it bids more aggressively than the competitive

small bidders, just as a weak bidder does in a two-bidder first-price auction (Maskin and Riley, 2000). In both cases, the weak bidder wins more than she would in the welfare maximizing allocation, where "more" indicates more frequently in first-price auctions and more quantity in our model. The seller can use the cap to adjust the strength of the large bidder's distribution relative to that of the competitive small bidders. Tightening the cap makes the large bidder stronger relative to its competition in the auction, which leads the large bidder to win a smaller quantity.

In the general downward-sloping demand case, it remains true that the cap should be tightened when the large bidder's auction quantity exceeds the welfare maximizing quantity. With downward-sloping demand, the type-θ_L large bidder still wins $\kappa F_L(\theta_L)$ in the auction, but the welfare maximizing quantity is smaller. Let $q_L^W(\theta_L)$ be the welfare maximizing quantity when the large bidder type is θ_L. This quantity is implicitly defined by $v(q_L^W(\theta_L), \theta_L) = F_S^{-1}(q_L^W(\theta_L))$, which again can be seen as the intersection of demand and residual supply. Using this definition, tightening the cap (weakly) increases expected welfare if $\kappa F_L(\theta_L) \geq q_L^W(\theta_L)$ for all $\theta_L \in \Theta_L$.

Proposition 2. *Tightening the cap weakly increases expected welfare if $\kappa F_L(\theta_L) \geq q_L^W(\theta_L)$ for all $\theta_L \in \Theta_L$. If $\kappa < 1$ and $\kappa F_L(\theta_L) \leq q_L^W(\theta_L)$ for all $\theta_L \in \Theta_L$, relaxing the cap weakly increases expected welfare. If $F_L(\theta_L) \leq q_L^W(\theta_L)$ for all $\theta_L \in \Theta_L$, then $\kappa = 1$ is optimal.*

When the large bidder has flat demand, there is a close connection between the large bidder being weak or strong at κ and the inequalities in Proposition 2 because $q_L^W(\theta_L) = F_S(\theta_S)$ in this case. With downward-sloping demand, the welfare maximizing quantity is less than $F_S(\theta_S)$, i.e., $q_L^W(\theta_L) \leq F_S(\theta_S)$, because with downward-sloping demand, the large bidder has marginal values weakly below θ_L. Proposition 2 therefore implies that if the large bidder is weak at κ, the cap should be tightened in the general downward-sloping demand case as well.

Corollary 1. *If the large bidder is weak at κ, tightening the cap increases expected welfare.*

Using Corollary 1, a sufficient condition for *any* cap $\kappa < 1$ to improve expected welfare is that the large bidder is weak at κ for some $\kappa \leq 1$. This follows because being weak at some κ implies that the large bidder is weak when $\kappa = 1$.

We cannot conclude that relaxing the cap increases expected welfare when the large bidder is strong at κ, unless the large bidder has flat demand. Being strong at κ implies that the large bidder wins less than $F_S(\theta_L)$ in the auction, but being strong at κ does not imply that the quantity won is below the welfare maximizing quantity, $q_L^W(\theta_L)$, with downward-sloping demand. In other words, it is possible for the large bidder to be strong at κ and yet win more than the welfare maximizing quantity in the auction. In such a case, relaxing the cap would decrease expected welfare.

If the large bidder is neither weak nor strong at κ, the results above do not apply directly, but the analysis indicates the trade-offs involved in adjustments to the cap. Being

neither weak nor strong at κ means that the large bidder sometimes wins too much in the auction and sometimes wins too little. A tighter cap increases welfare conditional on the large bidder winning too much quantity, but decreases it conditional on the large bidder winning too little quantity. Optimally balancing these effects requires setting a cap such that a weighted average of the large bidder's marginal value in the resale market, $v(F_S(p(\theta_L)), \theta_L)$, is equal to a weighted average of the value of the marginal small bidder in the resale market, $p(\theta_L)$. We derive the first-order condition for the optimal cap in the appendix.

5. Revenue maximizing caps

In this section, we consider the influence of the cap on expected (auction) revenue, and study the relation between the revenue maximizing cap and the welfare maximizing one. Recall that Corollary 1 implies that imposing some cap increases expected welfare relative to not using a cap when the large bidder is weak at κ for some κ. In contrast, we show that when the large bidder is weak at some κ, imposing any cap reduces revenue when an additional regularity condition holds. This regularity condition is the same one used by Kirkegaard (2012) to rank first- and second-price auction revenue. Thus, the conditions under which a cap can be used to increase welfare also lead to the conclusion that a cap reduces revenue. While one might suspect that it is always the case that a cap reduces revenue, we show that imposing a cap can increase revenue. As with welfare maximization, the strength of the large bidder's types distribution relative to the competitive small bidder's is critical in determining how adjusting the cap changes expected revenue.

To determine the influence of the cap on expected revenue, we first write expected revenue as a function of the winning bidder's virtual valuations, as in Myerson (1981). From this expression it follows that the influence of a tighter cap on revenue is determined by the difference in the large and small bidders' virtual valuations at the margin. Under Kirkegaard's condition, we can sign the difference in these two virtual valuations and determine the influence of the cap on expected revenue. As in the previous section, we start with the flat demand case so as to build intuition.

We use a standard envelope theorem argument to write expected revenue in terms of virtual valuations. This argument implies that the bidders' expected surpluses in the game—and hence their expected payments—are determined by the ex post allocation. These payments include auction payments and any transfer payments made in the resale stage, so one cannot back out *each* bidder's auction payment directly using this strategy. However, the resale transfers net out when adding together all of the bidders' payments, so we can still write the *total* expected auction revenues in terms of the usual difference between total surplus and bidder surplus. Let $m_i(x) = x - (1 - F_i(x))/f_i(x)$, $i \in \{S, L\}$, be the virtual valuation of the type-x bidder. Using standard transformations, the expected revenue is

$$\mathrm{E}[\text{revenue}] = \mathrm{E}_{\theta_L}\big[F_S\big(p(\theta_L)\big)m_L(\theta_L)\big] + \mathrm{E}_{\theta_S}\big[F_L\big(p^{-1}(\theta_S)\big)m_S(\theta_S)\big], \tag{9}$$

where $p^{-1}(\theta_S)$ is the inverse of p or the type of large bidder setting the resale price θ_S. Note that θ_L and θ_S are random variables in (9). We define $p^{-1}(\theta_S) = \overline{\theta}_L$ for $\theta_S > p(\overline{\theta}_L)$

so as to capture the payments of sure winning small bidders. The first term in (9) represents the expected payment across the auction and resale stages made by the large bidder when each type, θ_L, finishes the game with the quantity $F_S(p(\theta_L))$. The second term is the corresponding expected payment from the small bidders when each type, θ_S, retains a unit following resale if and only if her type is above $p^{-1}(\theta_S)$. As noted above, the resale transfers cancel out after adding the two terms together.

Expression (9) indicates that the effect of tightening the cap on revenue can be determined by comparing the virtual value of each large bidder to that of the marginal small bidder in the resale market. Tightening the cap reduces the resale price set by each type-θ_L large bidder, $p(\theta_L)$. Intuitively, a small decrease in the cap transfers quantity from each type-θ_L large bidder to the type-$p(\theta_L)$ small bidder. Hence, tightening the cap reduces revenue if the former always has a larger virtual valuation than the latter, meaning $m_L(\theta_L) > m_S(p(\theta_L))$ for all $\theta_L \in \Theta_L$.

We use the condition introduced by Kirkegaard (2012) to sign the difference in these two virtual values. If we put the large bidder in the role of the weak bidder in Kirkegaard's paper, condition (9) from his paper is

$$f_L(\theta_L) \geq f_S(x) \quad \text{for all } x \in \left[\theta_L, F_S^{-1}\big(F_L(\theta_L)\big)\right] \text{ and all } \theta_L \in \Theta_L. \tag{10}$$

The interval of small bidder types over which (10) is assumed to hold, $[\theta_L, F_S^{-1}(F_L(\theta_L))]$, has a simple interpretation in our model. It contains all of the relevant resale prices when the large bidder is a seller. The large bidder would lose money in the resale market at any price below θ_L. No small bidder would purchase a unit at a price exceeding $F_S^{-1}(F_L(\theta_L))$, because any small bidder with type $\theta_S > F_S^{-1}(F_L(\theta_L))$ wins a unit in the auction.[10] In addition to (10), Kirkegaard assumes that the weak bidder's distribution is dominated by the strong bidder's distribution in terms of the hazard-rate order. We require only the weaker notion of first-order stochastic dominance, which is implied by (10). We also use this notion of dominance in Definition 1, where we define the phrase "weak at κ."

Lemma 1. *Kirkegaard's condition, (10), implies that the large bidder is weak at $\kappa = 1$ as long as $F_L(\theta_L)$ and $F_S(\theta_L)$ are not equal for all $\theta_L \in \Theta_L$.*

When Kirkegaard's condition holds and the large bidder has flat demand, imposing any cap reduces auction revenue. The definition of $p(\theta_L)$ implies that the difference in virtual valuations is equal to

$$m_L(\theta_L) - m_S\big(p(\theta_L)\big) = \frac{1 - \kappa F_L(\theta_L)}{f_S\big(p(\theta_L)\big)} - \frac{1 - F_L(\theta_L)}{f_L(\theta_L)}.$$

With (10) it follows that $m_L(\theta_L) > m_S(p(\theta_L))$ for all $\kappa < 1$ and all $\theta_L \in \Theta_L$. In short, Kirkegaard's condition is sufficient for the large bidder to have a higher virtual valuation

[10] More precisely, since the large bidder wins $\kappa F_L(\theta_L)$ in the auction,there are no small bidders in the resale market with types above $F_S^{-1}(\kappa F_L(\theta_L))$ who did not already purchase a unit in the auction. Thus, the large bidders sells nothing in the resale market at any price exceeding $F_S^{-1}(\kappa F_L(\theta_L))$. Using $F_S^{-1}(F_L(\theta_L))$ in (10) ensures that it holds for all κ.

than the marginal small bidder in the resale market. Under this condition,revenue falls when we impose a cap because the cap reduces the large bidder's allocation.

To account for downward-loping demand, we must correct the above analysis for the fact that the large bidder's virtual valuation depends on its value at the marginal quantity in the resale market, $F_S(p(\theta_L))$. The generalized virtual value of the large bidder with type θ_L is

$$m_L(\theta_L) = v\big(F_S\big(p(\theta_L)\big), \theta_L\big) - v_{\theta_L}\big(F_S\big(p(\theta_L)\big), \theta_L\big)\frac{1 - F_L(\theta_L)}{f_L(\theta_L)}, \tag{11}$$

where $v_{\theta_L}(F_S(p(\theta_L)), \theta_L)$ is the partial derivative of the marginal value with respect to the large bidder's type.

We can also weaken the requirements of Kirkegaard's condition by using the fact that we have an expression for the resale price in terms of model primitives, Equation (5). This removes the requirement that the condition hold over an interval of small bidder types for each large bidder type. Incorporating these two changes, the condition becomes

$$\frac{f_L(\theta_L)}{v_{\theta_L}\big(F_S\big(p(\theta_L)\big), \theta_L\big)} \geq f_S\big(p(\theta_L)\big) \quad \text{for all } \theta_L \in \Theta_L. \tag{12}$$

In our next result, we also refer to the case where the inequality is reversed in (12):

$$\frac{f_L(\theta_L)}{v_{\theta_L}\big(F_S\big(p(\theta_L)\big), \theta_L\big)} \leq f_S\big(p(\theta_L)\big) \quad \text{for all } \theta_L \in \Theta_L. \tag{13}$$

The next proposition reports the connection between (12), (13), and the revenue maximizing cap.

Proposition 3. *Auction revenue is increasing in κ for $\kappa < 1$ if (12) holds for all $\kappa < 1$. Revenue is decreasing in κ at $\kappa = 1$ if (13) holds at $\kappa = 1$ and the inequality in (13) is strict for a nonzero measure of θ_L.*

In addition to showing when tightening the cap decreases revenue, Proposition 3 indicates that tightening the cap can increase revenue. Therefore, we provide a condition under which reducing competition by reducing the number of bids received in the auction can increase revenue. The example given in the next section satisfies condition (12) or (13) depending on parameter values, and,hence,gives a case where revenue may be either increasing or decreasing in the cap.

Combined, Corollary 1 and Proposition 3 relate welfare maximization to revenue maximization. The prescription for the cap is contradictory in the following sense. Imposing some cap increases welfare if the large bidder is weak at $\kappa = 1$, but in this case any cap reduces revenue as long as (12) holds.

Proposition 4. *If the large bidder is weak at $\kappa = 1$ and (12) holds for all $\kappa < 1$, the welfare maximizing cap is strictly below* 1*, while the revenue maximizing cap is* 1.

Note that the large bidder is weak at $\kappa = 1$ if it is weak at κ for any $\kappa < 1$. In the Appendix, we derive a first-order condition for an interior choice of revenue maximizing cap. Analogous to the case with welfare maximization, choosing a cap to maximize revenue amounts to equating weighted averages of the large and small bidders' virtual valuations, $m_L(\theta_L)$ and $m_S(p(\theta_L))$.

6. An example

We use an example to illustrate how adjusting the quantity cap on the large bidder affects welfare and revenue. In the example, the distribution of small bidder values, F_S, is $U[0, 1]$, while the distribution of large bidder types, F_L, is $U[0, \overline{\theta}_L]$. Thus, the competitive small bidder types are distributed according to $U[0, \kappa]$ and the large bidder is weak at κ if $\overline{\theta}_L \leq \kappa$. The large bidder's demand curve is $v(x, \theta) = \max\{0, \theta - \alpha x\}$, where $\alpha \geq 0$. Our assumption on the supports of the type distributions require that $\overline{\theta}_L \leq 2 + \alpha - \kappa \leq 1 + \alpha$. By Observation 1, the large bidder wins the quantity $\kappa F_L(\theta_L) = \kappa\theta_L/\overline{\theta}_L$ in the auction. The pricing formula in (5) determines the resale price that the large bidder sets:

$$p(\theta_L) = \frac{\overline{\theta}_L + \kappa}{(2+\alpha)\overline{\theta}_L}\theta_L.$$

For the resale price distribution and bid functions, we calculate

$$F(p) = \frac{2+\alpha}{\overline{\theta}_L + \kappa}p \qquad b(p) = \frac{p}{2}$$

$$b_L(\theta_L) = \frac{\overline{\theta}_L + \kappa}{2\overline{\theta}_L(2+\alpha)}\theta_L \qquad b_S(\theta_S) = \frac{\overline{\theta}_L + \kappa}{2\kappa(2+\alpha)}\min\{\theta_S, \kappa\}.$$

Increases in κ or decreases in $\overline{\theta}_L$ make the large bidder's bids more aggressive and the competitive small bidders' bids less so. Intuitively, an increase in κ or a decrease in $\overline{\theta}_L$ makes the large bidder weaker relative to the competitive small bidders, either because it faces stronger competition in the auction or because its marginal values decrease. Similar to a first-price auction, a weaker bidder places more aggressive bids.

The welfare maximizing allocation equates the large bidder's demand with its residual supply. Thus, welfare is maximized if each type-θ_L large bidder retains the quantity $\theta_L/(1+\alpha)$ following resale. This occurs if

$$F_S\big(p(\theta_L)\big) = \frac{\theta_L}{1+\alpha} \quad \Longrightarrow \quad \kappa_W = \frac{\overline{\theta}_L}{1+\alpha} \leq 1,$$

where κ_W is the welfare maximizing cap. Therefore, the cap is smaller if the large bidder is weaker ($\overline{\theta}_L$ is smaller) or has a steeper demand curve (α is larger). Note that Corollary 1 implies the cap should be tightened if $\overline{\theta}_L \leq \kappa$. This is clear in the example because $\kappa_W \leq \overline{\theta}_L$. In terms of quantities, the large bidder wins $\kappa F_L(\theta_L) = \kappa\theta_L/\overline{\theta}_L$ in the auction, and this is greater than $\theta_L/(1+\alpha)$ when $\kappa > \overline{\theta}_L/(1+\alpha)$. Tightening the cap reduces the auction quantity, and the large bidder enters the resale market with a quantity closer to the welfare maximizing quantity.

We examine the virtual valuations so as to compare welfare and revenue maximizing caps. In this example, these are $m_L(\theta_L) = 2\theta_L - \alpha p(\theta_L) - \overline{\theta}_L$ (see (11)) and $m_S(\theta_S) = 2\theta_S - 1$. According to our analysis in Section 5, the sign of the difference

$$m_L(\theta_L) - m_S\big(p(\theta_L)\big) = \left(1 - \frac{\kappa}{\overline{\theta}_L}\right)\theta_L + 1 - \overline{\theta}_L$$

indicates how adjustments to the cap influence revenue. It is straightforward to show that $m_L(\theta_L) > m_S(p(\theta_L))$ for all $\kappa < 1$ if $\overline{\theta}_L < 1$. In words, relaxing the cap increases revenue when the large bidder is weak at $\kappa = 1$. The fact that the large bidder has the higher virtual valuation means that revenue increases in the cap, because relaxing the cap above some level $\kappa < 1$ increases the resale price and effectively transfers units from each type-$p(\theta_L)$ small bidder to each type-θ_L large bidder. Proposition 4 similarly indicates that the revenue maximizing cap is equal to 1 here.

In our discussion of our results, we focus on the importance of the relation between the type distributions. This example helps clarify the role of other features of the model, such as the shape of the large bidder's demand. Consider making the demand steeper by increasing α. This lowers resale prices and, consequently, lowers auction bids, because the bids are based on the resale price distribution. However, the steeper demand has no impact on the auction allocation, because the effect on bids is symmetric. Regardless of the slope of the large bidder's demand, it wins the quantity $\kappa F_L(\theta_L)$ in the auction. The seller's influence on the final allocation, which determines welfare and revenue, is limited to her influence on the auction allocation, and the auction allocation is determined by the quantity awarded to the large bidder. The shape of the large bidder's demand enters into the sellers objective only insofar as it influences the target allocation.

7. Speculation by small bidders

In the baseline model, we do not allow the small bidders to purchase more than one infinitesimal unit of the good in the auction. Since the large bidder may speculate by purchasing quantity it intends to sell in the resale market, it is reasonable to allow small bidders to do the same. An initial departure from the baseline model is to allow small bidders to submit bids for additional infinitesimal units in the auction. Specifically, in this extension we allow a small bidder to submit $k \geq 1$ additional infinitesimal bids for a total demand of $k\,dq$, and we continue to assume that the small bidder's private value for any additional infinitesimal units is zero.

Small bidders cannot profit by submitting additional bids. Importantly, the payoff of the lowest type of small bidder is zero in equilibrium. According to Proposition 1, the type-$\underline{\theta}_S$ small bidder bids $\underline{p}$ in the auction and wins with probability 0. Furthermore, the resale price exceeds $\underline{p}$ with probability 1 and $\underline{p} \geq \underline{\theta}_S$, so this small bidder has no incentive to purchase in the resale market. An additional bid by a small bidder with zero value would at most yield an additional payoff equal to the payoff of the small bidder with the lowest type in equilibrium, which is zero. Additional bids by small bidders also do not influence the large bidder's residual supply curve in the auction because the bids are infinitesimal. Since these additional bids earn at most zero and do not influence the large bidder, they cannot be profitable.

PROPOSITION 5. *A small bidder cannot increase her payoff by bidding for additional infinitesimal units if her private value for those units is zero.*

8. CONCLUSION

We study how to set a quantity cap in a discriminatory price auction between a large bidder who values a nonnegligible fraction of the divisible good and a continuum of small bidders each of whom value an infinitesimal unit. We consider the choice of quantity cap that would maximize expected welfare following a post-auction resale stage and compare it to the one that would maximize expected auction revenue.

Our results are explained by the relative strength of the type distributions of the large bidder and the small bidders who compete for the same units in the auction. Tightening the quantity cap increases expected welfare when the large bidder has a relatively weak distribution, because a weak large bidder bids too aggressively and wins more than the welfare maximizing quantity, just as a weak bidder wins more frequently in a first-price auction relative to a second-price auction. Since the large bidder has market power in the resale market, resale does not correct the auction allocation. However, tightening the cap increases the welfare of the auction allocation and this forces the ex post allocation to improve as well.

In contrast to welfare, tightening the quantity cap reduces revenue when the large bidder is weak under an additional restriction on the type distribution. Under regularity conditions on the type distributions, a weak large bidder has a higher virtual valuation than the marginal small bidder when no cap is in place. Tightening the cap distorts the allocation away from the large bidder because it causes the large bidder to bid less aggressively. This reduces revenue when the large bidder has the higher virtual valuation.

In the Salomon Brothers scandal mentioned in the Introduction, we observed a large bidder win most of the units awarded in the auction prior to entering a resale market. The large bidder in our model is weak if he is likely to have low demand for the good. Despite having low demand, the large bidder wins a large fraction of the good in the auction in the absence of a quantity cap if his type is large relative to *his* possible types, i.e., if $F_L(\theta_L)$ is large (see Observation 1). In such a case, the large bidder wins most of the units in the auction despite having low demand relative to the small bidders. The large bidder exploits his market power in the resale market and sells to the small bidders with high private values who failed to win a unit in the auction. These predictions are consistent with the Salomon Brothers story. Following the auction in the Salomon Brothers scandal, the price of the respective notes remained higher than comparable notes for weeks following the auction (Jegadeesh, 1993), suggesting that Salomon Brothers had gained significant market power in the secondary market.

APPENDIX

PROOF OF PROPOSITION 1. This proof is composed of two parts. We first show that the proposed strategies form an equilibrium. Then we argue that this equilibrium is unique among equilibria in which the large and small bidders follow nondecreasing strategies.

Verification of proposed equilibrium

The lowest and highest equilibrium bids are $\underline{b} = b(\underline{p})$ and $\overline{b} = b(\overline{p})$, respectively. The bid strategy $b : [\underline{p}, \overline{p}] \to [\underline{b}, \overline{b}]$ is increasing on the interior of its domain. Our assumptions guarantee that the same is true for b_L on $(\underline{\theta}_L, \overline{\theta}_L)$ and b_S on $(\underline{\theta}_S, F_S^{-1}(\kappa))$. Therefore, for all $b \in [\underline{b}, \overline{b})$,

$$F\big(\phi(b)\big) = F_L\big(\phi_L(b)\big) = \frac{1}{\kappa} F_S\big(\phi_S(b)\big), \tag{14}$$

where $\phi = b^{-1}$. If we define $\phi_S(\overline{b}) = \sup\{\theta_S | b_S(\theta_S) < \overline{b}\}$, symmetrization holds at $\overline{b}$ too.

Now consider a type-θ_L large bidder. Using the payoff in (2) and the envelope theorem, the first-order condition for his choice of bid is

$$\phi_S'(b) \frac{1}{\kappa} f_S\big(\phi_S(b)\big)\big(p(b, \theta_L) - b\big) - \frac{1}{\kappa} F_S\big(\phi_S(b)\big) = 0.$$

An implication of (14) is that this first-order condition holds at the proposed choice of bid, $b_L(\theta_L)$, as long as $\phi(b) = p(b, \phi_L(b))$, which is exactly when the resale price is chosen optimally. If the large bidder were to bid $b' \gtreqless b_L(\theta_L)$, then

$$\frac{\partial}{\partial b} \pi_L\big(\theta_L, b', p\big(b', \theta_L\big)\big) = \phi_S'\big(b'\big) \frac{1}{\kappa} f_S\big(\phi_S\big(b'\big)\big)\big(p\big(b', \theta_L\big) - b'\big) - \frac{1}{\kappa} F_S\big(\phi_S\big(b'\big)\big) \lesseqgtr 0,$$

since $p(b', \cdot)$ is increasing.

Given the large bidder's strategy, the type-θ_S small bidder's payoff is given by (1). Consider instead the payoff given by

$$\begin{aligned} \widehat{\pi}_S(\theta, b) = F_L\big(\phi_L(b)\big)(\theta - b) &+ \int_{\underline{\theta}_L}^{\phi_L(b)} \big(p(t) - \theta\big) \mathbb{1}\big\{p(t) \geq \theta\big\}\, dF_L(t) \\ &+ \int_{\phi_L(b)}^{\overline{\theta}_L} \big(\theta - p(t)\big) \mathbb{1}\big\{p(t) \leq \theta\big\}\, dF_L(t). \end{aligned} \tag{15}$$

The expression in (15) drops the condition from (1) that $v(F_S(p(\theta)), \theta) \geq p(\theta)$ when the large bidder purchases from the small bidder in the resale market or that $v(F_S(p(\theta)), \theta) \leq p(\theta)$ when the large bidder sells to the small bidder. Hence, $\widehat{\pi}_S(\theta, b) \geq \pi_S(\theta, b)$ for all b and $\theta_S \in \Theta_S$, but $\widehat{\pi}_S(\theta, b_S(\theta)) = \pi_S(\theta, b_S(\theta))$, meaning the two objectives are equal when the small bidder uses the equilibrium strategy. It follows that if $b_S(\theta)$ maximizes $\widehat{\pi}_S(\theta, \cdot)$, it also maximizes $\pi_S(\theta, \cdot)$. Differentiating $\widehat{\pi}_S$ with respect to the second argument, we find

$$\begin{aligned} \frac{\partial}{\partial b} \widehat{\pi}_S(\theta, b) &= \phi_L'(b) f_L\big(\phi_L(b)\big)(\theta - b) - F_L\big(\phi_L(b)\big) \\ &\quad + \phi_L'(b) f_L\big(\phi_L(b)\big)\big(p\big(\phi_L(b)\big) - \theta\big) \mathbb{1}\big\{p\big(\phi_L(b)\big) \geq \theta\big\} \\ &\quad - \phi_L'(b) f_L\big(\phi_L(b)\big)\big(\theta - p\big(\phi_L(b)\big)\big) \mathbb{1}\big\{p\big(\phi_L(b)\big) \leq \theta\big\} \\ &= \phi_L'(b) f_L\big(\phi_L(b)\big)\big(p\big(\phi_L(b)\big) - b\big) - F_L\big(\phi_L(b)\big) \\ &= \phi'(b) f\big(\phi(b)\big)\big(\phi(b) - b\big) - F\big(\phi(b)\big), \end{aligned} \tag{16}$$

where we use the facts that $F_L(\phi_L(b)) = F(\phi(b))$ and $p(\phi_L(b)) = \phi(b)$. The first-order condition holds for any bid in $[\underline{b}, \overline{b}]$, since the effect on $p(\phi_L(b))$ is internalized. This implies that $\widehat{\pi}_S(\theta, \cdot)$ is constant on $b \in [\underline{b}, \overline{b}]$, but not necessarily that $\pi_S(\theta, \cdot)$ is constant on this interval. Bidding strictly lower than $\underline{b}$ earns the same expected payoff as a bid of $\underline{b}$ because in either case, the small bidder wins with probability 0 in the auction. Bidding strictly greater than $\overline{b}$ is dominated by a bid of $\overline{b}$.

The small bidders with values larger than $F_S^{-1}(\kappa)$ are the sure winners in the auction. Notice that (16) holds for them as well at any bid, including $\bar{b}$. The fact that a nonzero measure of small bidders bid $\bar{b}$ in equilibrium does not cause any difficulty here, because once the large bidder has outbid all of the competitive small bidders with a bid of $\bar{b}$, there is no additional gain to bidding slightly higher due to the binding quantity cap.

The large and small bidders' bids are therefore optimal given the expected resale price, and as argued in the text, the resale price is optimal given the possible auction outcomes.

Uniqueness in nondecreasing strategies

To prove that the above equilibrium is unique in nondecreasing strategies, we build on the argument given by Hafalir and Krishna (2008) in their online appendix, the main elements of which are the following. Under the assumption that bid functions are nondecreasing, we first show that the large and small bidders' bid functions have a common range. Two lemmas from Hafalir and Krishna (2008), which carry over essentially unaltered to our case, imply that the equilibrium bid functions must be continuous and increasing (excluding bids placed by sure winning small bidders). We then show that symmetrization must hold in any such equilibrium. The final step is to show that the equilibrium distribution of resale prices is unique for any set of bid functions that satisfy these properties. Because bids must be derived from the resale price distribution, this implies the uniqueness result.

Suppose that $b_L : [\underline{\theta}_L, \overline{\theta}_L] \to \mathbb{R}$, $p : [\underline{\theta}_L, \overline{\theta}_L] \to \mathbb{R}$, and $b_S : [\underline{\theta}_S, \overline{\theta}_S] \to \mathbb{R}$ determine, respectively, the equilibrium choices of the large bidder's bid, the large bidder's resale price, and the small bidder's bid. Assume that all three quantities are nondecreasing and continuous. Continuity is shown later. Note that as long as b_L is nondecreasing, the resale price is increasing in the type without loss of generality, because the large bidder's optimal choice of resale price is increasing in both the quantity purchased at auction and its type.

We break the argument down into a series of claims.

Claim 1 (Common high bid). *We have* $b_S(F_S^{-1}(\kappa)) = b_S(\overline{\theta}_S) = b_L(\overline{\theta}_L)$.

Proof. Clearly $b_L(\overline{\theta}_L) \leq b_S(F_S^{-1}(\kappa))$, because otherwise large bidders with types near $\overline{\theta}_L$ win κ (i.e., the maximum allowed) and pay strictly more than the smallest amount required to win κ units, $b_S(F_S^{-1}(\kappa))$. Similarly, if $b_L(\overline{\theta}_L) < b_S(F_S^{-1}(\kappa))$, then there are small bidders who win in the auction with probability 1 who pay strictly more than the smallest amount required to win with probability 1, $b_L(\overline{\theta}_L)$. Finally, $b_S(F_S^{-1}(\kappa)) = b_S(\overline{\theta}_S)$

because small bidders with types greater than $F_S^{-1}(\kappa)$ win with probability 1 for any bid of at least $b_L(\overline{\theta}_L)$. □

CLAIM 2. *We have* $p(\underline{\theta}_L) \geq \underline{\theta}_S$.

PROOF. The large bidder always sets a price that is at least as large as the price it sets if its type is $\underline{\theta}_L$ and it wins zero units in the auction. Call this lowest price p_0. The first-order condition for the choice of p_0 is

$$H(p_0) = v\big(F_S(p_0), \underline{\theta}_L\big) - p_0 - \frac{F_S(p_0)}{f_S(p_0)} = 0.$$

Since $H(\underline{\theta}_S) = \underline{\theta}_L - \underline{\theta}_S \geq 0$ and $H(\cdot)$ is decreasing, $p_0 \geq \underline{\theta}_S$. Therefore, $p(\underline{\theta}_L) \geq p_0 \geq \underline{\theta}_S$. □

CLAIM 3. *We have* $b_L(\underline{\theta}_L) \geq b_S(\underline{\theta}_S)$.

PROOF. For a contradiction, assume $b_S(\underline{\theta}_S) > b_L(\underline{\theta}_L)$ and define θ'_L by $b_S(\underline{\theta}_S) = b_L(\theta'_L)$. Claim 1 with continuity implies such a type exists. Note that all types $\theta_L \in [\underline{\theta}_L, \theta'_L]$ win zero units in the auction and that by Claim 2, $p(\theta'_L) > p(\underline{\theta}_L) \geq \underline{\theta}_S$.

If $p(\theta'_L) > b_S(\underline{\theta}_S)$, the type-$\theta'_L$ large bidder could improve his payoff by buying units in the auction at a lower price with a bid $b' \in (b_S(\underline{\theta}_S), p(\theta'_L))$.

If $b_S(\underline{\theta}_S) \geq p(\theta'_L)(> \underline{\theta}_S)$, then the type-$\underline{\theta}_S$ small bidder wins a unit with positive probability but bids more than $\max\{\underline{\theta}_S, p(\theta_L)\}$ in the event that $\theta_L \in [\underline{\theta}_L, \theta'_L)$, which is the event in which she wins a unit, implying that the type-$\underline{\theta}_S$ small bidder has a negative payoff upon winning a unit regardless of whether it is kept or resold to the large bidder. □

CLAIM 4 (Common low bid). *We have* $b_S(\underline{\theta}_S) = b_L(\underline{\theta}_L)$.

PROOF. Using Claim 3, it is sufficient to show $b_S(\underline{\theta}_S) \geq b_L(\underline{\theta}_L)$. For a contradiction, assume $b_L(\underline{\theta}_L) > b_S(\underline{\theta}_S)$. Let θ'_S satisfy $b_S(\theta'_S) = b_L(\underline{\theta}_L)$.

If $p(\underline{\theta}_L) > b_L(\underline{\theta}_L)$, a small bidder with type $\theta_S < \theta'_S$ never wins in the auction and must pay at least $p(\underline{\theta}_L)$ in the resale market for a unit. Any such bidder can strictly improve her payoff with a bid $b' \in (b_L(\underline{\theta}_L), p(\underline{\theta}_L))$, because she may keep the unit won with positive probability for a price b' or resell it for a profit.

It cannot be that the large bidder wins units with a bid $b_L(\underline{\theta}_L)$ and optimally sets a resale price $p(\underline{\theta}_L) \leq b_L(\underline{\theta}_L)$, because by lowering the auction bid, the large bidder would reduce his auction payment and purchase instead at the weakly lower resale price. To see this, suppose that $q(b)$ is the quantity won by the large bidder in the auction for a bid b. A necessary condition for optimality at $b = b_L(\underline{\theta}_L)$ is that

$$q'(b-)\big(p(\underline{\theta}_L) - b\big) - q(b) \geq 0,$$

where we use $q'(b-)$ for the left-hand derivative of q at b in case $q(b)$ is not differentiable at b. This condition cannot hold at $b = b_L(\underline{\theta}_L)$ if $p(\underline{\theta}_L) \leq b_L(\underline{\theta}_L)$. □

Therefore, b_L and b_S have a common range. We next appeal to the online appendix of Hafalir and Krishna (2008), specifically Lemma 2(S) and Lemma 3(S), which with insignificant modifications show that b_L for all types and b_S restricted to $\theta_S \in [\underline{\theta}_S, F_S^{-1}(\kappa)]$ must be continuous and increasing, respectively. The arguments are standard from the first-price auctions literature without resale. Claim 1 implies that b_S is continuous on $[F_S^{-1}(\kappa), \overline{\theta}_S]$.

Next, we show that symmetrization must hold in equilibrium. Define ϕ_L and ϕ_S as the corresponding inverse bid functions for bids in the range $[b_L(\underline{\theta}_L), b_L(\overline{\theta}_L)]$.

CLAIM 5. *We have $F_L(\phi_L(b)) = F_S(\phi_S(b))/\kappa$ for all $b \in [b_L(\underline{\theta}_L), b_L(\overline{\theta}_L)]$.*

PROOF. The argument is very similar to that given for Proposition 1 of Hafalir and Krishna (2008). We give a brief version here. Let $\phi(b)$ be the resale price set by the large bidder type that bids b in the auction. The first-order condition for the large bidder's choice of bid implies

$$\frac{d}{db}\ln\left(\frac{1}{\kappa}F_S\big(\phi_S(b)\big)\right) = \frac{1}{\phi(b)-b},$$

while the first-order condition for each small bidder's choice of bid implies

$$\frac{d}{db}\ln\big(F_L\big(\phi_L(b)\big)\big) = \frac{1}{\phi(b)-b}.$$

Together with the common boundary condition, $F_L(\phi_L(\bar{b})) = F_S(\phi_S(\bar{b}))/\kappa = 1$, where $\bar{b} = b_L(\overline{\theta}_L)$, these differential equations imply that symmetrization holds. That is, $F_L(\phi_L(b)) = F_S(\phi_S(b))$ for all $b \in [b_L(\underline{\theta}_L), b_L(\overline{\theta}_L)]$. □

In the body of the paper, we show that symmetrization implies that in equilibrium the mapping from large bidder types to resale prices is uniquely determined independently of the bid functions used (see (5)).

CLAIM 6. *The distribution of resale prices F is uniquely determined and independent of b_S and b_L.*

Finally, symmetrization implies that $\phi(b)$ is determined as the *unique* solution to the symmetric first-price auction between two bidders with values distributed according to F. Since $F_L(\phi_L(b)) = F_S(\phi_S(b)) = F(\phi(b))$ for all $b \in [b_L(\underline{\theta}_L), b_L(\overline{\theta}_L)]$, b_S and b_L are uniquely determined for $\theta_S \in [\underline{\theta}_S, F_S^{-1}(\kappa)]$ and $\theta_L \in [\underline{\theta}_L, \overline{\theta}_L]$, respectively. Claim 1 then determines b_S for $\theta_S \in [F_S^{-1}(k), \overline{\theta}_S]$. □

FIRST-ORDER CONDITION FOR WELFARE MAXIMIZATION. The expected ex post welfare given a cap κ can be written as

$$W(\kappa) \equiv \int_{\underline{\theta}_L}^{\overline{\theta}_L}\int_0^{F_S(p(\theta))} v(x,\theta)\,dx\,dF_L(\theta) + \int_{\underline{\theta}_L}^{\overline{\theta}_L}\int_{F_S(p(\theta))}^1 F_S^{-1}(x)\,dx\,dF_L(\theta),$$

where $F_S(p(\theta))$ is the large bidder's allocation following resale. The function $p(\theta)$ is the resale price set by the type-θ large bidder from (5). To understand the second term, note that $F_S^{-1}(q)$ is the large bidder's residual supply curve. The second term integrates the residual supply curve for all quantities retained by small bidders following resale. Differentiating, we find that

$$W'(\kappa) = \int_{\underline{\theta}_L}^{\overline{\theta}_L} \big\{v\big(F_S\big(p(\theta)\big), \theta\big) - p(\theta)\big\} p_\kappa(\theta) f_S\big(p(\theta)\big)\, dF_L(\theta), \tag{17}$$

where p_κ denotes the derivative of p with respect to κ. The sign of the integrand is determined by the sign of $v(F_S(p(\theta)), \theta) - p(\theta)$ because $p_\kappa(\theta)$ is positive. Solving $W'(\kappa) = 0$ amounts to equating $v(F_S(p(\theta)), \theta)$ and $p(\theta)$ "on average," where the weighting is determined by $p_\kappa(\theta) f_S(p(\theta)) f_L(\theta)$. Note that we may divide by $\mathrm{E}_{\theta_L}[p_\kappa(\theta_L) f_S(p(\theta_L))]$ to make this a proper weighted average. □

Proof of Proposition 2. Using (5), $v(\kappa F_L(\theta), \theta) \lessgtr F_S^{-1}(\kappa F_L(\theta))$ implies

$$v\big(\kappa F_L(\theta), \theta\big) \lessgtr v\big(F_S\big(p(\theta)\big), \theta\big) \lessgtr p(\theta) \lessgtr F_S^{-1}\big(\kappa F_L(\theta)\big).$$

Notice that if $v(\kappa F_L(\theta), \theta) < F_S^{-1}(\kappa F_L(\theta))$ for all $\theta \in \Theta_L$, then $W'(\kappa') < 0$ for all $\kappa' \geq \kappa$. Hence, $\kappa_W < \kappa$. The argument is similar if $v(\kappa F_L(\theta), \theta) > F_S^{-1}(\kappa F_L(\theta))$ for all $\theta \in \Theta_L$ and $\kappa < 1$. If $v(F_L(\theta), \theta) > F_S^{-1}(F_L(\theta))$ for all θ, then it must be that $W'(\kappa) > 0$ for all $\kappa \in [0, 1]$. □

First-order condition for revenue maximization. Rewriting the expression for expected revenue in (9) and using $R(\kappa)$ to represent the expected revenue with a cap of κ, we get

$$R(\kappa) \equiv \int_{\underline{\theta}_L}^{\overline{\theta}_L} F_S\big(p(\theta)\big) m_L(\theta)\, dF_L(\theta) + \int_{\underline{\theta}_S}^{\overline{\theta}_S} F_L\big(p^{-1}(\theta)\big) m_S(\theta)\, dF_S(\theta),$$

where $m_S(x) = x - (1 - F_S(x))/f_S(x)$ and m_L is defined in (11). If we differentiate this expression, we get

$$\begin{aligned}
R'(\kappa) &= \int_{\underline{\theta}_L}^{\overline{\theta}_L} m_L(\theta_L) p_\kappa(\theta_L) f_S\big(p(\theta_L)\big)\, dF_L(\theta_L) \\
&\quad + \int_{\underline{\theta}_S}^{\overline{\theta}_S} m_S(\theta_S) p_\kappa^{-1}(\theta_S) f_L\big(p^{-1}(\theta_S)\big)\, dF_S(\theta_S) \\
&= \int_{\underline{\theta}_L}^{\overline{\theta}_L} m_L(\theta_L) p_\kappa(\theta_L) f_S\big(p(\theta_L)\big)\, dF_L(\theta_L) \\
&\quad + \int_{\underline{\theta}_S}^{F_S^{-1}(\kappa)} m_S(\theta_S) p_\kappa^{-1}(\theta_S) f_L\big(p^{-1}(\theta_S)\big)\, dF_S(\theta_S) \\
&= \int_{\underline{\theta}_L}^{\overline{\theta}_L} \big\{m_L(\theta) - m_S\big(p(\theta)\big)\big\} p_\kappa(\theta) f_S\big(p(\theta)\big)\, dF_L(\theta).
\end{aligned}$$

For the first equation, recall that we defined $p^{-1}(x) = \overline{\theta}_L$ for $x > F_S^{-1}(\kappa)$. The second equality follows from using the change of variables $\theta_S = p(\theta_L)$ in the second term. We also use the fact that $p_\kappa^{-1}(p(\theta))p'(\theta) = -p_\kappa(\theta)$, which can be derived by totally differentiating the identity $p^{-1}(p(\theta)) = \theta$ with respect to κ.

Thus, the sign of the integrand is determined by the difference $m_L(\theta) - m_S(p(\theta))$. An interior revenue maximizing κ equates these two virtual valuations "on average," where the weighting is determined by $p_\kappa(\theta)f_S(p(\theta))f_L(\theta)$, which is the same weighting function in $W'(\kappa)$ from (17). □

Proof of Lemma 1. Condition (10) immediately implies

$$F_L(\theta_L) = \int_{\underline{\theta}_L}^{\theta_L} f_L(x)\,dx \geq \int_{\underline{\theta}_L}^{\theta_L} f_S(x)\,dx = F_S(\theta_L) \quad \text{for all } \theta_L \in \Theta_L.$$

The last equality follows because (10) also implies that $\underline{\theta}_L \leq \underline{\theta}_S$. □

Proof of Proposition 3. Using the definition of $p(\theta)$ in (5), the integrand of $R'(\kappa)$ has the same sign as

$$\frac{1-\kappa F_L(\theta)}{f_S\big(p(\theta)\big)} - v_\theta\big(F_S\big(p(\theta)\big),\theta\big)\frac{1-F_L(\theta)}{f_L(\theta)}.$$

Condition (12) holding for all $\kappa \in (0,1)$ implies that for all $\theta \in \Theta_L$, this expression is strictly positive since

$$\frac{1-\kappa F_L(\theta)}{1-F_L(\theta)} > 1 \geq \frac{v_\theta\big(F_S\big(p(\theta)\big),\theta\big)f_S\big(p(\theta)\big)}{f_L(\theta)}.$$

Alternatively, when (13) holds at $\kappa = 1$, we have

$$\frac{v_\theta\big(F_S\big(p(\theta)\big),\theta\big)f_S\big(p(\theta)\big)}{f_L(\theta)} \geq 1 = \frac{1-F_L(\theta)}{1-F_L(\theta)}.$$

If, in addition, the inequality is strict for a nonzero measure of large bidder types, then $R'(1) < 0$. □

Proof of Proposition 4. The proof follows from Corollary 1 and Proposition 3. □

References

Anton, James J. and Dennis A. Yao (1992), "Coordination in split award auctions." *Quarterly Journal of Economics*, 107, 681–707. [2]

Ausubel, Lawrence M., Peter Cramton, Marek Pycia, Marzena Rostek, and Marek Weretka (2014), "Demand reduction and inefficiency in multi-unit auctions." *Review of Economic Studies*, 81, 1366–1400. [4]

Back, Kerry and Jaime F. Zender (1993), "Auctions of divisible goods: On the rationale for the treasury experiment." *Review of Financial Studies*, 6, 733–764. [2, 4, 7]

Baisa, Brian and Justin Burkett (2018), "Large multi-unit auctions with a large bidder." *Journal of Economic Theory*, 174, 1–15. [4]

Baisa, Brian and Justin Burkett (2019), "Discriminatory price auctions with resale and optimal quantity caps." (Unpublished paper, Forthcoming Theoretical Economics). [5, 6]

Bartolini, Leonardo and Carlo Cottarelli (1997), "Treasury bill auctions: Issues and uses." In *Macroeconomic Dimensions of Public Finance: Essays in Honour of Vito Tanzi* (Mario Blejer and Teresa Ter-Minassian, eds.), 258–354, Routledge, London. [1]

Brady, Nicholas F., Richard C. Breeden, and Alan Greenspan (1992), "Joint report on the government securities market." U.S. Government Printing Office. [1]

Cheng, Harrison H. and Guofu Tan (2010), "Asymmetric common-value auctions with applications to private-value auctions with resale." *Economic Theory*, 45, 253–290. [3, 5, 6]

Coutinho, Paulo Braulio (2013), "Divisible good auctions with resale." Unpublished paper, Department of Economics, University of California, Los Angeles. [4, 5]

Garbade, Kenneth D. and Jeffrey F. Ingber (2005), "The treasury auction process: Objectives, structure, and recent adaptations." Current Issues in Economics and Finance, 11. [1, 2]

Garratt, Rod and Thomas Tröger (2006), "Speculation in standard auctions with resale." *Econometrica*, 74, 753–769. [5]

Gong, Jiong, Jianpei Li, and R. Preston McAfee (2012), "Split-award contracts and investment." *Journal of Public Economics*, 96, 188–197. [2]

Gupta, Madhurima and Bernard Lebrun (1999), "First price auctions with resale." *Economics Letters*, 64, 181–185. [3, 5, 10, 12]

Hafalir, Isa and Vijay Krishna (2008), "Asymmetric auctions with resale." *American Economic Review*, 98, 87–112. [3, 5, 10, 12, 13, 22, 24]

Hafalir, Isa and Musab Kurnaz (2018), "Discriminatory auctions with resale." (Unpublished paper, forthcoming Economic Theory Bulletin). [4]

Holmberg, Pär (2009), "Supply equilibria of pay-as-bid auctions." *Journal of Regulatory Economics*, 36, 154–177. [4]

Jegadeesh, Narasimhan (1993), "Treasury auction bids and the salomon squeeze." *Journal of Finance*, 48, 1403–1419. [1, 20]

Kirkegaard, René (2012), "A mechanism design approach to ranking asymmetric auctions." *Econometrica*, 80, 2349–2364. [3, 5, 15, 16]

Maskin, Eric and John Riley (2000), "Asymmetric auctions." *Review of Economic Studies*, 67, 413–438. [3, 5, 14]

Myerson, Roger B. (1981), "Optimal auction design." *Mathematics of Operations Research*, 6, 58–73. [3, 15]

Pycia, Marek and Kyle Woodward (2017), "Pay-as-bid: Selling divisible goods." Unpublished paper, Department of Economics, University of California, Los Angeles. [4]

Swinkels, Jeroen (2001), "Efficiency of large private value auctions." *Econometrica*, 69, 37–68. [4]

Viswanathan, S. and James J. D. Wang (2004), "Inter-dealer trading in finanacial markets." *Journal of Business*, 77, 987–1040. [4]

Wang, James J. D. and Jaime F. Zender (2002), "Auctioning divisible goods." *Economic Theory*, 19, 673–705. [4]

Woodward, Kyle (2016), "Strategic ironing in pay-as-bid auctions." Unpublished paper, Department of Economics, UNC Chapel Hill. [4]

Co-editor Simon Board handled this manuscript.

Manuscript received 6 July, 2017; final version accepted 11 June, 2019; available online 14 June, 2019.