

Optimal Quantity Caps in Discriminatory Price Auctions with Resale

Brian Baisa* and Justin Burkett†

February 10, 2017

Abstract

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 the first-price auction with resale results of Hafalir and Krishna (2008) to this setting and give a tractable characterization of equilibrium behavior. We use this to study the policy of capping the amount that may be won by large bidders in the auction. Despite being used in the US Treasury auction among others, this policy 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 revenue or efficiency has stark and contradictory implications for the choice of cap.

1 Introduction

Quantity caps place limits on the number of units a single bidder is able to win in a multi-unit auction and are a common feature of important real-world auctions. For example, in the US treasury bill market, currently no bidder may bid for more than 35% of the market supply.¹ Without caps a bidder may buy an inefficiently large quantity in the auction, thereby allowing it to exercise market power in the secondary market. Indeed, this was the case in the notable Salomon Brothers scandal in 1991. Salomon Brothers admittedly violated US Treasury auction rules and succeeded in 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 an SEC investigation (Jegadeesh, 1993; Brady et al., 1992).

Although they are used in prominent multi-unit auctions like the US Treasury auctions, the theoretical auction literature has not to our knowledge provided a detailed analysis of how and when

* Amherst College and Boston College, Department of Economics, bbaisa@amherst.edu

† Wake Forest University, Department of Economics, burketje@wfu.edu

¹This evolved over the course of the 20th century. Prior to 1990 the limitation was placed on the award amount but not the bid size. The size of the cap has also varied 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.

a quantity cap should be set. 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 use our construction to evaluate the seller's optimal choice of quantity cap for two different objectives, the expected surplus of the allocation following resale and the expected revenue generated by the auction itself.

Although ostensibly quantity caps are intended to reduce the deadweight loss caused by a large bidder exercising market power in the resale market, we are also interested in whether choosing an efficient quantity cap necessarily has a negative impact on the expected revenue from the auction. In other auction settings a trade-off exists between maximizing efficiency maximizing revenue (Myerson, 1981), but a quantity cap is an imperfect instrument for maximizing revenue, so the cap that maximizes revenue among all caps does not in general produce the revenue maximizing auction. The efficiency-revenue trade-off is important for the US Treasury, who have expressed interest in both objectives. The US Treasury's reaction to the Salomon Brothers scandal, for example, suggested the efficiency of the resale market was a major concern, but cost minimization is often cited as their primary objective, which is the analogue of revenue maximization for procurement.² Our results suggest that there is indeed an important trade-off between these objectives.³

In our model, a large bidder with downward sloping multi-unit demand competes against a "continuum" of small bidders each with "single-unit" demand in a discriminatory price auction for a divisible good. 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 monopsonist or a monopolist depending on the auction outcome. We assume that the large bidder sets the resale price and that while the auctioneer can restrict the amount won in the auction, she cannot directly restrict the amount traded in the resale market. In other words, the auctioneer's influence over the resale market — and hence the final allocation — is limited to how she may influence equilibrium behavior through a choice of quantity cap in the auction.

To derive equilibrium strategies, we extend the techniques developed in Hafalir and Krishna (2008) to our environment. Hafalir and Krishna (2008) study a two-bidder first-price auction for a single indivisible unit followed by a resale stage in which the winning bidder may resell to the losing bidder at a take-it-or-leave-it price.⁴ Their surprising result is that adding resale to an asym-

² Garbade and Ingber (2005) reports that minimizing the cost of funds is the auction objective. See Footnote 1 of that paper, for example, which references a speech made by the Under Secretary. Yet, other statements from Fed personnel seem 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".

³ Ausubel and Cramton (1999) find a contrasting result. They show that with perfect resale the efficient assignment of goods in an auction is the same as the revenue-maximizing assignment. Our results which incorporate an imperfect resale market do not agree with this conclusion and support the assertion that there are trade-offs when comparing revenue maximization to surplus maximization.

⁴They consider other resale mechanisms as well, but this model is their primary focus.

metric first-price auction yields symmetric bid distributions in equilibrium. This fact allows them to construct bidding strategies in terms of a symmetric first-price bid function, making the analysis of equilibrium bidding tractable.

The environment we study is not isomorphic to a single-unit first-price auction with two bidders, but we show that the strategic implications are essentially the same for the bidders and use this idea to extend the Hafalir and Krishna (2008) results. If one thinks of the large bidder as playing the role of one bidder and all of the small bidders as playing the role of the other, then one can think of the quantity won by the large bidder as being analogous to the probability of winning in the single-unit auction. However, this is not sufficient to make the connection between the two models, because, for example, the large bidder can submit a downward sloping bid curve in the discriminatory price auction and may have declining marginal values. We show that despite these additional complications the key insight of Hafalir and Krishna (2008) that resale makes bidding behavior symmetric extends to our environment.

After characterizing equilibrium, we analyze the influence of the cap on the total surplus of the final allocation. The quantity cap determines the interval of small bidders against whom the large bidder directly competes. These small bidders, whom we refer to as the “competitive” small bidders, are the ones with interim win probabilities in the auction between zero and one. The remaining small bidders win with probability one in equilibrium. Tightening the cap makes the large bidder compete against small bidders with lower values, and we show that this has the expected effect of reducing the quantity allocated to the large bidder in the auction.

There are two equivalent ways to describe when and how to adjust the cap. First, in equilibrium there is a monotonic relationship between the cap and the resale price chosen by any type of large bidder. If the resale price is expected to be higher than the Walrasian equilibrium price (Walrasian price), for example, tightening the cap reduces the resale price for every type of large bidder and produces a more efficient outcome. The resale price is higher than the Walrasian price whenever the large bidder purchases more than the efficient amount in the auction. The alternative way to understand the cap is that it influences the amount the large bidder wins in the auction by manipulating the distributional strength of the competitive small bidders. Tightening the cap makes the competitive small bidders relatively weaker, which causes the large bidder to bid less aggressively and win less in the auction. This is similar to how the relative strength of a bidder in a two-bidder first-price auction determines how aggressively it bids in equilibrium. (Maskin and Riley, 2000).

We relate the relative strength of the large bidder’s type distribution and the distribution of the competitive small bidders to the direction that the cap should be adjusted to improve surplus (Proposition 2). The implication of the large bidder’s strength on how much it buys relative to the efficient quantity can be determined type-by-type in the model. A given type of large bidder is locally stronger than the small bidder with the same marginal value if the large bidder’s type occurs with

smaller probability.⁵ In equilibrium, this large bidder type purchases inefficiently few units, and relaxing the cap would correct this. However, manipulating the cap affects all large bidder types at once. If a certain cap makes all of the large bidder types either stronger or weaker than the competitive small bidders, due to a first-order stochastic dominance relationship holding, then the implications for the choice of cap are clear. For example, if the large bidder is weaker in the sense of first-order stochastic dominance, a tighter cap improves total surplus.

Therefore, to justify using *any* quantity cap to increase surplus the large bidder must have a relatively weak distribution without a cap.⁶ A large bidder with a very weak distribution is one who's marginal value for the good is likely very low and is analogous to a speculator who bids in the auction, despite having little value for holding the good, because the goods can be sold at a premium in the resale market. For example, in the Salomon Brothers scandal, Salomon purchased an outsized portion of the issuance before implementing a short squeeze in the secondary market, meaning they appear to have purchased treasuries with no intention of holding them.

The conditions under which a cap maximizes revenue are similar to those under which a cap maximizes efficiency after replacing marginal values with the Myerson (1981) virtual values. Since it is not generally possible to achieve a fully efficient allocation in the auction, when efficiency is the objective the auctioneer optimally adjusts the cap to make the ex post resale price — the value of the marginal small bidder in the resale market — equal to the large bidder's marginal value of the last unit purchased in the auction “on average”.⁷ On the other hand, when revenue maximization is the objective we show that the cap should be set to on average equate the virtual valuation of the marginal small bidder with the virtual valuation of the large bidder after adjusting for the quantity purchased, where the weighting of different realizations is the same between efficiency and revenue maximization. Using the interpretation in Bulow and Roberts (1989) of virtual valuations as marginal revenues, this can be understood as adjusting the cap so that in equilibrium the marginal revenue from the marginal small bidder is equal on average to the marginal revenue from the last unit sold to the large bidder in the auction.

The difference between equating marginal values and virtual values is significant and leads to strongly contrasting policy implications for the auctioneer.⁸ When the large bidder's type distribution is weak relative to the small bidders' distribution, we have argued that a cap can be used to improve the total surplus of the final allocation. The efficient cap roughly equates the resale price and the marginal value of the large bidder. However, the weakness of the large bidder's type dis-

⁵More precisely, the type- v large bidder is locally strong given the cap if $F_L(v) < F_{S/\kappa}(p^*)$ where p^* is the Walrasian price. F_L and $F_{S/\kappa}$ are the type distributions of the large and competitive small bidders respectively.

⁶If the large bidder's distribution is strong without a cap, then using a cap magnifies the effect of large bidder's market power as a monopsonist following the auction and reduces efficiency.

⁷If it were possible to make the resale price equal to the large bidder's marginal value for every type of large bidder, this would be fully efficient and would lead to no trade in the resale market.

⁸This contrasts with the argument made in Ausubel and Cramton (1999) that in the presence of resale the two objectives are aligned (See Footnote 3).

tribution will tend to mean that at the point where the large bidder's marginal value is equal to the resale price the marginal revenue accruing to the auctioneer from that large bidder is strictly higher than the marginal revenue from the marginal small bidder. This implies that the cap should be relaxed to increase revenue. We provide an example in which this effect is extreme. In the example, if a cap less than one should be used to maximize total surplus, a cap equal to one (i.e., no cap) should be set to maximize revenue. Therefore, using a cap necessarily reduces revenues compared to no cap.

We generalize the implications of this example by using the regularity conditions introduced by Kirkegaard (2012) to rank the expected revenues of a first- and second-price auction with two bidders. Under these conditions, we show that when the surplus maximizing cap is relatively tight the revenue maximizing cap should be one (i.e., no cap should be used). Similarly when the cap should be one to maximize surplus it should be strictly less than one to maximize revenue (Propositions 3 and 4).

Related Work We are not aware of any existing theoretical work evaluating the consequences for the choice of a quantity cap on outcomes from a multi-unit auction preceding a resale stage. Back and Zender (1993) briefly consider quantity caps prior to their Theorem 3, but they do so in an auction 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. Hafalir and Kurnaz (2015) study discriminatory auctions with resale in a discrete model in which bidders only demand a single unit. Given the nature of their model, there is no way to study quantity caps in it.

Our results have parallels in the large literature on asymmetric first-price auctions. Quite a few papers develop the implications of the stochastic ordering of type distributions on the bidding strategies and outcomes in a first-price auction upon which we draw for intuition (cf. Lebrun, 1998; Arozamena and Cantillon, 2004). The literature ranking the expected revenues of the first- and second-price auction is particularly important. The conditions we use from Kirkegaard (2012) generalize the revenue rankings of Maskin and Riley (2000) and are the weakest known sufficient conditions under which a first-price auction for a single unit raises more expected revenue than a second-price auction.

While there are similarities, the analogy to a first-price auction is flawed in important ways, and results from that literature are not generally transferable. For example, our model includes small bidders who do not compete directly with the large bidder in the auction. In equilibrium, these small bidders place bids that are equal to the bid placed by the highest type of large bidder and win with probability one. The fact that revenue derived from these bidders has no analogue in the Hafalir and Krishna (2008) leads to qualitatively different results. For example, Theorem 2 of Hafalir and Krishna (2008) states that a first-price auction with resale revenue-dominates a second-

price auction with resale. One might attempt to translate this into our environment by saying that the efficient auction minimizes revenue, but this is not true precisely because the expected revenue of our auction differs from the expected revenue in the Hafalir and Krishna (2008) auction due to the presence of the sure winners.⁹

Besides the asymmetric first-price auction literature, our paper is also influenced by the literature on large multi-unit auctions. Both this paper and Baisa and Burkett (2016) augment models of large discriminatory auctions (Swinkels, 1999, 2001) to include a bidder with non-negligible demand.

One interesting application of our model is to settings in which a monopolist (the large bidder) competes against small firms composing a “competitive fringe.” Krishna (1993), for example, introduces a model in which small competitors, considered entrants, compete against a single large incumbent in a sequence of auctions for market capacity (or licenses to operate). She uses this model to critique the argument that monopolies will tend to persist because a monopolist values the additional capacity more than the entrants do (Gilbert and Newberry, 1982).

The remainder of the paper is organized as follows. The next section introduces the model. Section 3 characterizes equilibrium strategies, Section 4 uses this characterization to analyze the problem of finding the constrained efficient cap, and Section 5 relates the efficient cap to the revenue maximizing one. Section 6 concludes. An appendix contains additional material including a proof.

2 Model

A large bidder competes against a unit measure of small bidders for a unit measure of a divisible good.¹⁰ Bidders have private values and the large bidder has demand for a positive measure of the good while each small bidder demands and may purchase at most an infinitesimal unit of the good.

The large bidder has a single-dimensional type, $v_L \in [l, u_L]$ with $l \geq 0$, which is distributed according to the absolutely continuous distribution function $F_L(v_L)$. The type- v_L large bidder’s private value from buying a fraction q of the good is

$$\int_0^q v_L - g(x) dx,$$

where g is nondecreasing, $g(0) = 0$, and $l - g(1) \geq 0$.¹¹ A special case which we frequently refer to

⁹Compare the expression for expected revenue in our (12) to (14) in Hafalir and Krishna (2008). To apply the Hafalir and Krishna (2008) results to our model we would have to assume that making the quantity cap more restrictive leads to a decline the total number of small bidders who bid in the auction (e.g., to eliminate any sure winners), but this would be a strange assumption to make about a quantity cap.

¹⁰In a related paper, Baisa and Burkett (2016), we consider a very similar model (without resale) in which a large bidder competes against a measure μ_S of small bidders where μ_S is not necessarily one. We show there that in cases where $\mu_S \neq 1$ the equilibrium can be derived by first adjusting the small bidders’ distribution by either adding mass at the lower end or truncating the smallest types, and then applying the same logic as the $\mu_S = 1$. We do not analyze these cases in the current paper.

¹¹We shift the distribution up to the interval $[l, u_L]$ to avoid interpreting negative marginal values. While the functional

in the analysis is the one in which the large bidder has a constant marginal value for the good (i.e., $g(\cdot) = 0$). We refer to this as the flat demand case.

The small bidders each have value $v_S \in [l, u_S]$ for the good and are distributed according to the increasing absolutely continuous distribution function $F_S(v_S)$.¹² Our equilibrium specification requires that $u_S \geq u_L$.¹³

We will assume, as Hafalir and Krishna (2008) do, that the Myerson (1981) regularity condition holds for both F_L and F_S . This condition requires that the “virtual valuations”, $v - (1 - F_i(v))/f_i(v)$ for $i \in \{S, L\}$, are increasing functions.¹⁴

2.1 Auction Stage

We study a standard discriminatory price (or “pay-as-bid”) rule with a 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, the good is awarded to the highest bids and bidders are charged an amount equal to each of their winning bids. The rules implicitly require that bid curves be nonincreasing.¹⁵

If the type- v_L large bidder submits the nonincreasing bid schedule $b_L(v_L, q)$ for $q \leq \kappa$, then define the large bidder’s quantity demanded at b as

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

If each type- v_S small bidder bids using the increasing and continuous $b_S(v_S)$ with inverse b_S^{-1} , the quantity demanded by the small bidders at b is $1 - F_S(b_S^{-1}(b))$. The lowest winning bid, b^* , is then

$$b^* = \sup\{b | q_L(b; b_L) \geq F_S(b_S^{-1}(b))\}.$$

All of the small bidders with bids $b_S \geq b^*$ receive a payoff of $v_S - b_S$ following the auction. Given the large bidder’s bid curve, $b_L(v_L, q)$, the large bidder receives a fraction $F_S(b^*)$ of the good for a

form assumption on the demand curve is not crucial to our characterization, we use the fact that the large bidder has a single dimensional type to characterize bid behavior. We have chosen one that is simple to work with to show that downward sloping demand can be incorporated in the model.

¹²Our assumption of a common lower bound on the support of the large and small bidders is consistent with Hafalir and Krishna (2008) who assume that both type distributions have a support that includes zero.

¹³In our resale model, the large bidder decides to either buy additional units from small bidders or sell some of its units to the small bidders. Hafalir and Krishna (2008) focus on a resale model in which a winning bidder who is weak in a distributional sense offers to sell their unit to the strong bidder. They argue in an online appendix that the main results of the model hold if instead a losing strong bidder offers to buy the unit in the resale market as a monopsonist. We found that the two cases are not symmetric, and the latter requires an additional assumption on the distributions. We discuss this point further in the Appendix.

¹⁴As Hafalir and Krishna (2008) point out, this regularity condition holding implies that the conditional virtual valuations, $v - (F_i(a) - F_i(v))/f_i(v)$, are increasing for $v \leq a$.

¹⁵This is a consequence of the fact that the auctioneer treats each “incremental” bid (i.e., the level of a bid curve at some quantity q) independently in that the higher bids always receive units whenever a lower bid does.

payoff of

$$\int_0^{F_S(b^*)} (v_L - g(q) - b_L(v_L, q)) dq.$$

Note that the quantity cap implies that $F_S(b^*) \leq \kappa$.¹⁶

2.2 Resale Stage

In the resale stage, the large bidder announces a single take-it-or-leave-it price at which it is willing to buy or sell additional quantity in the resale stage. 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 efficient quantity. We do not assume that the quantity cap is enforced in the resale market, and hence allow the large bidder to finish the game with a quantity greater than κ following the resale stage. In other words, the quantity cap is solely an auction rule.

The information released following the auction is often an important policy to consider when the auction is followed by a resale market.¹⁷ However, this is not the case in this model. We simply assume that the large bidder observes the clearing price and quantity awarded in the auction stage. The large bidder in equilibrium would not benefit from acquiring any additional information about the small bidders between stages (see Footnote 19).

3 Equilibrium

We derive an equilibrium of this game with the following properties. Each type of small bidder bids according to the monotonic function, $b_S(v_S)$. As price-takers in the resale market, they sell after winning in the auction if and only if the resale price exceeds their value. They buy after losing if and only if their value exceeds the resale price. The large bidder also uses a pure strategy in the auction that is monotone in its type, while the resale price is chosen optimally conditional on the type, the quantity won, the price paid at auction and the small bidder's equilibrium bid function.

This construction allows us to extend the methodology of Hafalir and Krishna (2008) (HK) to our multi-unit setting. HK derive the resale price that should be chosen by a winning bidder conditional on the winning bid and the assumption that the opposing bidder follows an increasing bid strategy. They show that if resale takes this form and the bidders use increasing bid functions, the two equilibrium bid distributions must be symmetric. They go on to use this result to construct equilibrium bid and pricing functions in terms of model primitives.

We follow similar steps in the next two subsections. To keep the analysis similar to that of HK, we must first rule out that the potential complications introduced by the multi-unit environment do

¹⁶The cap is easily enforced by setting any positive bids from the large bidder for $q > \kappa$ to 0.

¹⁷In Hafalir and Krishna (2008) for example it is assumed that the losing bid is not announced, and this is important because it influences the inference made by the winning bidder prior to the resale stage.

not cause difficulties.

The first complication introduced in our model is that due to the quantity cap on the large bidder being arbitrary a monotone pure strategy equilibrium can involve small bidders who win with probability one. If the small bidders were to all bid according to an increasing $b_S(v_S)$, then with a cap of $\kappa < 1$ any small bidder with type $v_S > F_S^{-1}(\kappa)$ must submit a larger bid than a measure κ of the small bidders and hence must win in the auction. The remaining competitive small bidders with types $v_S \leq F_S^{-1}(\kappa)$ win in the auction if and only if their bid exceeds the clearing price. In equilibrium, the sure winners cannot bid above the largest possible clearing price, implying that b_S cannot be strictly increasing for all types and all κ . However, in the proof of Proposition 1 we argue that if the sure winners all place a bid equal to the highest possible clearing price no difficulty arises, since the quantity cap prevents the large bidder from competing with these sure winners.

The second complication is that in the discriminatory price auction the large bidder may submit a bid curve that is strictly decreasing for some or all quantities. However, doing so in this environment is never a best response to small bidders who use a nondecreasing bid function. Given an increasing bid strategy used by the competitive small bidders, with any nonincreasing bid for which $b_L(v_L, q') = b_S(F_S^{-1}(q'))$ the large bidder is assured of winning exactly q' units. That is, the large bidder faces no uncertainty about the quantity won or clearing price. The large bidder thus has no incentive to ever bid more than $b_S(F_S^{-1}(q'))$ to win the quantity q' and optimally submits a flat bid.¹⁸ We denote this flat bid by $b_L(v_L)$ for all $0 \leq q \leq \kappa$.

Finally, we allow for the large bidder to have downward sloping demand. Note that downward sloping demand does not affect the optimality of a flat bid for the large bidder. However, the assumption of downward sloping demand necessarily influences the large bidder's choice of resale price. Since the expected resale price affects the distribution of bids in the auction stage, whether the large bidder has downward demand or not also influences the bids chosen by all bidders. However, we show in the next two subsections that downward sloping demand does not change the key result that we (and HK) require to derive equilibrium bid strategies: resale implies the equilibrium bid distributions are symmetric between the large and small bidders.

3.1 Resale Stage

To formally describe equilibrium, we use backward induction evaluating first the large bidder's optimal choice of price given the clearing price, its type and the assumed increasing bid function used by competitive small bidders. In the resale stage the only information the large bidder needs to know about the small bidders is that the competitive ones bid according to an increasing bid function and that the sure winners bid an amount weakly greater than any other small bidder. This information along with the quantity won in the auction is sufficient to determine which small bidders

¹⁸This is also noted by Baisa and Burkett (2016) who discriminatory auctions without resale and quantity caps.

won in the auction. In equilibrium, the large bidder knows the small bidders' bid function, so besides the clearing price and the quantity won any additional information is superfluous to the large bidder. Furthermore, any deviation by a "single" small bidder at the auction stage cannot influence the large bidder's payoff at either stage. Given that they are always price-takers in the resale stage, the small bidders would also not benefit from additional information.¹⁹

To determine the resale price, consider the large bidder's payoff when it places a bid b in the auction and subsequently sets a resale price of p , and let $b_S(v_S)$ be increasing for $v_S < F_S^{-1}(\kappa)$ with inverse ϕ_S . With a bid $b \leq b_S(F_S^{-1}(\kappa))$ for all $q \in [0, \kappa]$ the large bidder wins the quantity $F_S(\phi_S(b))$ in the auction. The large bidder may only increase its payoff in the resale market if the auction outcome is inefficient. To find the optimal price the large bidder sets in the resale market we consider two cases. If the large bidder wins more than the efficient quantity, then there is a range of mutually agreeable prices where the large bidder sells some of her quantity to small bidder. This occurs if $\phi_S(b) > v - g(F_S(\phi_S(b)))$, meaning the large bidder's marginal value for the marginal unit is less than highest type of losing small bidder. In this case, the large bidder sets the resale price p to maximize their total payoff given their type v and b ,

$$\begin{aligned}\pi_L(v, b, p) &= (F_S(\phi_S(b)) - F_S(p))(p - b) + \int_0^{F_S(p)} (v - g(q) - b) dq \\ &= F_S(\phi_S(b))(p - b) + \int_0^{F_S(p)} (v - g(q) - p) dq.\end{aligned}\tag{1}$$

The first line divides the quantity won into the amount that is won and subsequently resold and the amount that is won and kept. The second line rearranges the first to isolate the effect of the bid.

If the large bidder wins less than the efficient quantity in the auction, then $\phi_S(b) < v - g(F_S(\phi_S(b)))$, and the large bidder can increase her payoff by offering to buy additional units at a p that maximizes

$$\begin{aligned}\pi_L(v, b, p) &= \int_{F_S(\phi_S(b))}^{F_S(p)} (v - g(q) - p) dq + \int_0^{F_S(\phi_S(b))} (v - g(q) - b) dq \\ &= F_S(\phi_S(b))(p - b) + \int_0^{F_S(p)} (v - g(q) - p) dq,\end{aligned}\tag{2}$$

which shows that the objective is the same in either case. Note that the choice of b is constrained by the quantity cap, while the choice of p is unconstrained. The optimal p given (b, v) is the same for (1) and (2),

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

¹⁹ In HK, 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 HK). We do not have a similar requirement about the information policy used by the auctioneer due to the fact that there is essentially nothing for the large bidder to learn in equilibrium in addition to the quantity won and the price paid when the competitive small bidders use an increasing bid function.

The regularity condition on F_S along with the assumption that g is weakly increasing implies that this p is unique and the condition is sufficient for optimality given (b, v) . Note that the resulting $p(b, v)$ is increasing in both arguments.

3.2 Auction Stage

The type- v large bidder's objective in the auction stage given the small bidders' strategies and its anticipated selection of a resale price is

$$\max_b \pi_L(v, b, p(b, v)) \text{ s.t. } F_S(\phi_S(b)) \leq \kappa$$

To handle the constraint we rewrite it as $F_S(\phi_S(b))/\kappa \leq 1$ and think of $F_S(v)/\kappa$ as F_S truncated to $[l, F_S^{-1}(\kappa)]$. For later use we denote this distribution by $F_{S/\kappa}$,²⁰ which can be understood as the distribution of the small bidders conditional on being in the competitive region. The constraint is then analogous to the implicit constraint in a first-price auction that a bidder cannot win with probability greater than one. To see this simply divide the large bidder's entire objective by the constant κ , the result of which may be interpreted as the large bidder's payoff per unit of available quantity. Note that the above discussion omits any mention of the small bidders with value $v_S > F_S^{-1}(\kappa)$, but as long as they place a bid that is at least as high as the largest bid they win for sure and this is their optimal bid (see the proof of Proposition 1).

The optimal bid of the large bidder can be derived using the envelope theorem and must satisfy the first-order condition

$$\phi_S'(b) f_{S/\kappa}(\phi_S(b)) (p(b, v) - b) - F_{S/\kappa}(\phi_S(b)) = 0, \quad (4)$$

regardless of whether the large bidder buys or sells in the resale market.²¹ In other words, the above argument shows that the bids of the large bidder and the competitive small bidders are analogous to the bids that would be submitted in HK's two bidder first-price auction with resale where the distributions of bidders are F_L and $F_{S/\kappa}$.

HK's symmetrization result follows from the observation in their model that the necessary conditions for optimality of each bidder's bid require that the bid distributions be everywhere equal. Symmetrization holds here as well, which can be seen by comparing the first-order condition of

²⁰Formally, this distribution is defined as

$$F_{S/\kappa}(x) \equiv \begin{cases} 0 & x < l \\ \frac{1}{\kappa} F_S(x), & x \in [l, F_S^{-1}(\kappa)] \\ 1 & x > F_S^{-1}(\kappa). \end{cases}$$

²¹The selection of b is unaffected if we use F_S instead of $F_{S/\kappa}$. We use the latter to reinforce the analogy to a first-price auction.

the large bidder to that of the small bidders. Consider a small bidder with value v_S that is in the competitive region. This bidder only wins in the auction if $b_S \geq b_L(v_L)$. If they lose and the resale price set by the large bidder is smaller than their value, $v_S > p(b_L(v_L), v_L)$, they purchase the item in the resale market. If they win and $p(b_L(v_L), v_L) > v_S$, they sell in the resale market. Using $[x]_+ \equiv \max\{0, x\}$, their payoff is therefore

$$\begin{aligned} \pi_S(v, b) = & F_L(\phi_L(b))(v - b) + \int_l^b [p(b, \phi_L(b)) - v]_+ dF_L(\phi_L(b)) \\ & + \int_b^{\bar{b}} [v - p(b, \phi_L(b))]_+ dF_L(\phi_L(b)) \end{aligned}$$

where $\bar{b} = \sup\{b | F_L(\phi_L(b)) \leq 1\}$. The first-order condition for an interior bid is

$$\begin{aligned} \phi_L'(b) f_L(\phi_L(b))(v - b + [p(b, \phi_L(b)) - v]_+ - [v - p(b, \phi_L(b))]_+) - F_L(\phi_L(b)) = \\ \phi_L'(b) f_L(\phi_L(b))(p(b, \phi_L(b)) - b) - F_L(\phi_L(b)) = 0. \end{aligned} \quad (5)$$

Notice that v always disappears from this expression even for $v > F_S^{-1}(\kappa)$. In equilibrium it turns out that all small bidders, including sure winners, are indifferent between any bid because they internalize the effect on $p(b, \phi_L(b))$.

As in HK's Proposition 1 the two first-order conditions imply that if ϕ_S and ϕ_L derive from an equilibrium with increasing bid functions for the large bidder and competitive small bidders,²² it must be that the following symmetrization identity holds.

$$F_{S/\kappa}(\phi_S(b)) = F_L(\phi_L(b)), \quad \forall b \in [l, \bar{b}] \quad (6)$$

Thus, the distribution of the large bidder's (flat) bids equals the distribution of bids placed by small bidders in the competitive region. Following the symmetrization result, HK show how to derive equilibrium strategies from primitives by constructing a new distribution function, $F(p)$, and deriving the equilibrium bid distribution of a symmetric first-price auction in which both bidders' values are distributed according to $F(p)$. The form of $F(p)$ is governed by the first-order condition for the price setter in the resale market. The analogous $F(p)$ in our environment is implicitly defined as

$$F(p) = F_L \left(p + g(F_S(p)) + \frac{F_S(p)}{f_S(p)} - \kappa \frac{F(p)}{f_S(p)} \right), \quad (7)$$

which is an increasing distribution function with support $[l, \bar{p}]$, where \bar{p} is defined as the p that makes the argument to F_L equal to u_L .

We use the symmetrized distribution, $F(p)$, to characterize the equilibrium in our setting — just

²²The bid function is only strictly increasing for the competitive small bidders, but they are the relevant small bidders for the large bidder.

as HK do in their first-price auction game. The corresponding symmetric bid function is

$$\beta(p) = \frac{1}{F(p)} \int_l^p x f(x) dx, \quad (8)$$

where $f(x)$ is the density function associated with $F(p)$. Using these definitions along with arguments similar to those in HK's Theorem 1 and Proposition 4 we prove the following.

Proposition 1. *Under the resale mechanism in which the large bidder buys or sells at a fixed resale price, symmetrization holds (i.e., Equation (6) holds), equilibrium bid strategies are*

$$b_L(v_L) = \beta(F^{-1}(F_L(v_L)))$$

$$b_S(v_S) = \begin{cases} \beta(F^{-1}(\frac{1}{\kappa}F_S(v_S))) & v_S \leq F_S^{-1}(\kappa) \\ \beta(\bar{p}) & v_S > F_S^{-1}(\kappa) \end{cases}$$

and the equilibrium resale price is determined by (3).

Proof. See Appendix. □

When the large bidder has flat demand and $\kappa = 1$ our results show that the HK model can be suitably adapted to give an equilibrium characterization for our discriminatory auction in which the large bidder has a flat demand curve. The large bidder plays the role of one bidder in the HK setting and the unit measure of small bidders plays the role of the other. Yet our analysis goes further than this in two respects. First, we show that HK's symmetrization result can accommodate downward sloping demand for the large bidder, which does not affect the large bidder's incentives to submit a flat bid curve. Second, we show how to accommodate the quantity cap by adjusting the distribution of small bidders that the large bidder competes against. In particular, with a cap the large bidder competes against a truncated distribution of small bidders.

4 Constrained Efficient Caps

Our next results use the equilibrium characterized in the previous section to evaluate the trade-offs for the auctioneer involved in the adjustment of the quantity cap. In the examples discussed in the introduction, the cap is motivated in terms of its impact on efficiency and/or auction revenue. We use our model to analyze the optimal cap choice under both criteria.

Except in special cases, a quantity cap cannot achieve full efficiency or the maximum possible revenue, and hence our results on efficient and/or revenue maximizing caps are necessarily describe second-best outcomes. The cap is a one-dimensional policy instrument, and in most cases attaining full-efficiency or revenue maximization requires a richer set of policy instruments, such as reserve prices for revenue maximization. Due to our assumption that the large bidder always has market

power at the resale stage, as long as the auction allocation is inefficient the allocation following the resale stage will be inefficient as well. Therefore, in order for full efficiency to be achieved it must be the case that the auction allocation maximizes surplus, preventing trade in the resale market. When the large bidder has flat demand the auction allocation is fully efficient if and only if there exists a cap which makes the large bidder's distribution exactly the same as the competitive small bidders' distribution (see Proposition 2).²³ That is, full efficiency is achievable in this case if the large bidder and competitive small bidders can be made symmetric.

In this section, we study the choice of a constrained efficient cap, which we hereafter refer to as an efficient cap. To determine whether a cap is optimal or not it is important to consider whether the cap makes the large bidder a buyer or a seller in the resale market. The large bidder is a seller when it bids aggressively relative to the competitive small bidders in the auction, buying more than the efficient amount. Since the large bidder has market power in the resale market the resale price exceeds the Walrasian price. The Walrasian price, p^* , solves $p^* = v_L - g(F_S(p^*))$, while the resale price would be larger, $p > p^*$. Alternatively, when the large bidder bids weakly in the auction, it becomes a buyer in the resale stage and consequently the resale price is below the Walrasian price.

Given the large bidder's type, v_L , let $p^*(v_L)$ be the unique resale price that would ensure an efficient final allocation of goods. Thus, $p^*(v_L)$ implicitly solves $p = v_L - g(F_S(p))$. In equilibrium, this large bidder chooses the price $p(b_L(v_L), v_L)$. Using these prices, the expected surplus deficit following resale is

$$\int_l^{u_L} \int_{p(b_L(v_L), v_L)}^{p^*(v_L)} (v_L - g(F_S(x)) - x) dF_S(x) dF_L(v_L).$$

Note that the inner integral measures the deadweight loss for each realization of the large bidder's type, while the outer one integrates over each possible large-bidder type. Thus, if the auctioneer's objective is to maximize the surplus of the final allocation, it chooses κ to minimize this expression.

Before discussing the general case, we first show that if the small bidders types are uniformly distributed the efficient cap can be explicitly derived.

Example 1. Suppose $g(x) = \alpha x$ with $\alpha \in [0, 1]$, $l = 1$, and F_S is the $U[1, 2]$ distribution function. We find that

$$p(b_L(v_L), v_L) = \frac{v_L + \alpha + \kappa F_L(v_L) + 1}{2 + \alpha},$$

while $p^*(v_L) = (v_L + \alpha)/(1 + \alpha)$. The first-order condition for the problem of maximizing expected

²³Formally, there exists a κ such that $u_L = F_S^{-1}(\kappa)$ and $F_L(v) = F_{S/\kappa}(v)$, $\forall v \in [l, u_L]$.

surplus reduces to

$$\begin{aligned}
& \int_1^{u_L} \frac{F_L(v)}{2+\alpha} \left(v - \alpha \left(\frac{v + \kappa F_L(v) - 1}{2+\alpha} \right) - 1 - \frac{v + \kappa F_L(v) - 1}{2+\alpha} \right) f_L(v) dv \\
&= \frac{1}{(2+\alpha)^2} \int_1^{u_L} F_L(v) (v - (1+\alpha)\kappa F_L(v) - 1) f_L(v) dv \\
&= \frac{1+\alpha}{2(2+\alpha)^2} \int_1^{u_L} v dF_L(v)^2 - \frac{\kappa(1+\alpha)}{3(2+\alpha)^2} - \frac{1}{2(2+\alpha)^2} = 0
\end{aligned}$$

which, noting that the objective is strictly concave, implies that the optimal κ is

$$\kappa_E = \min \left\{ \frac{3}{2} \int_1^{u_L} v dF_L(v)^2 - \frac{3}{2(1+\alpha)}, 1 \right\},$$

or that κ_E is determined by the expected value of the highest order statistic from two draws from the F_L distribution. With $\alpha = 0$, this can be interpreted as setting (if possible) the expected value of the higher of two draws from $F_{S/\kappa}$, which is $2\kappa/3 + 1$, equal to the expected value of the higher of two draws from F_L .

When $v_L \sim U[1, u_L]$, $\alpha = 0$, and $u_L \leq 2$, $\kappa_E = u_L - 1$. That is, κ is chosen to make the large bidder's distribution equal to the distribution of the small bidder's against whom it competes. With $\kappa = u_L - 1$, $p(b_L(v_L), v_L) = v_L$ and the outcome is fully efficient.

In the example, for arbitrary F_L the optimal cap adjusts the distribution of the competitive small bidders, $F_{S/\kappa}$, so that the mean of the first-order statistic of two draws from this distribution is equal to the mean of the first-order statistic of two draws from F_L . Given the equilibrium resale price offered by the type- v_L large bidder, $p(b_L(v_L), v_L)$, it is generally not possible to find a κ such that $p(b_L(v_L), v_L) = v_L - \alpha F_S(p(b_L(v_L), v_L))$ for all v_L , which would ensure full efficiency of the auction allocation. The optimal cap ensures this efficiency condition holds “on average”. A special case in which we can drop the “on average” qualification occurs when F_L is also a uniform distribution. In this case, the example shows that a cap can be found to make F_L and $F_{S/\kappa}$ identical, implying that the auction is symmetric and fully efficient.

Notice that in the example the resale price is increasing in κ for each type of large bidder. This is a general feature of the resale price and is one way to explain many of the results on efficient caps below. By using the symmetrization identity, (6), and the expression for the optimal resale price in (3), the resale price chosen by the type- v_L large bidder which we denote by $\hat{p}(v_L; \kappa)$ satisfies

$$\hat{p}(v_L; \kappa) = v_L - g(F_S(\hat{p}(v_L; \kappa))) + \frac{\kappa F_L(v_L) - F_S(\hat{p}(v_L; \kappa))}{f_S(\hat{p}(v_L; \kappa))}. \quad (9)$$

The function $\hat{p}(v; \kappa) : [l, u_L] \times [0, 1] \rightarrow [l, \bar{p}]$ maps large bidder types into the price offered in the resale market. Note that this is just $p(b_L(v_L), v_L)$. We refer to the following two properties of

$\hat{p}(v_L; \kappa)$ repeatedly in subsequent analysis.

Lemma 1. *The equilibrium resale price $\hat{p}(v; \kappa)$ satisfies*

- (i) *For each $v_L \in [l, u_L]$, $p(v_L; \kappa)$ is increasing in κ .*
- (ii) *$\hat{p}(v_L; \kappa) \geq p^*(v_L)$ as $F_L(v_L) \geq F_S(p^*(v_L))/\kappa$.*

Proof. Property (i) follows from (9) and the assumption that $x - (1 - F_S(x))/f_S(x)$ is increasing. From the symmetrization identity (6) it follows that the equilibrium quantity won by the type- v_L large bidder is $\kappa F_L(v_L)$. Property (ii) is then a consequence of Equation (3) being the first-order condition for the optimal choice of p . \square

In addition to explaining the monotonic relation between the cap and the resale price, (9) and Lemma 1 connect the distributional strengths of the bidders to the direction in which the resale price should be adjusted. The resale price is inefficiently high exactly when the type- v_L large bidder purchases more than the efficient quantity, which occurs when this large bidder is locally weaker than the competitive small bidder against whom it should tie in the efficient auction. By locally weaker, we mean simply that F_L evaluated at this type is larger than $F_{S/\kappa}$ evaluated at the type of small bidder against whom it should tie. Notice that the resale price is inefficiently high or low exactly when the large bidder is respectively a seller or a buyer in the resale market.

Several conclusions follow from these observations about the relation between an arbitrary cap and the efficient one. If for a given interior quantity cap $\kappa < 1$ either all types of large bidder buy or all types sell in the resale market, then this choice of cap κ cannot maximize expected post-resale surplus. When the large bidder is always a buyer, it wins inefficiently little in the auction and the resale price is always too low relative to the Walrasian price. Because the resale price is increasing in κ , relaxing the cap — which is feasible because $\kappa < 1$ — pushes the resale price closer to the Walrasian price and increases surplus. This has the effect of increasing the amount won in the auction by every type of large bidder, due to more aggressive bidding. Of course, if the large bidder is always a buyer in the resale market when the cap is one, the auctioneer is unable to increase efficiency by relaxing the cap. In this case, the constrained efficient quantity cap is one. The argument is similar if the large bidder is always a seller in the market.

A necessary and sufficient condition for the large bidder to always buy in the resale market is that F_L first-order stochastically dominates $F_{S/\kappa}$, or that the large bidder is relatively strong. Note that this characterization of strength is weaker than others used in the first-price auction literature, which define relative strength in terms of reverse hazard-rate dominance (Maskin and Riley, 2000).

We give a graphical depiction of the relationships implied by Equation (9) in Figure 1. Panel (a) compares hypothetical distributions for the large, small and competitive small bidders given κ . In this example, the small bidder distribution first-order stochastically dominates the large bidder distribution, but the competitive small bidders have a distribution that is considerably weaker than

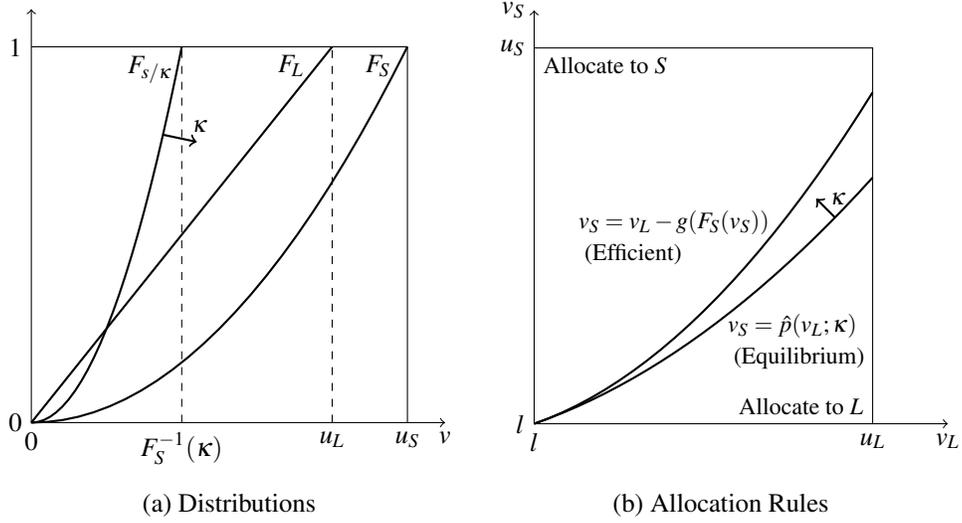


Figure 1: Effect of κ on Allocation

all the small bidders. Panel (b) depicts all possible allocations following the resale stage in (v_L, v_S) space. The upper curve represents the locus of valuations that should with tie one another for the allocation to be efficient, while the lower one depicts the tying valuations for the equilibrium allocation rule. In the example shown, the auction over-allocates to the small bidders for any v_L .

In either panel, the arrows labeled with κ show the direction of change in $F_{S/\kappa}$ and the locus $v_S = \hat{p}(v_L; \kappa)$ when κ is increased, or the cap is relaxed.

We draw the following conclusions about the efficient cap, κ_E , where we use \succ_{FO} and \succeq_{FO} to represent the strict and non-strict versions of the first-order stochastic dominance order.

Proposition 2. *Properties of the relationship between an arbitrary cap, κ , and the efficient cap, κ_E , are*

- i. *If for all $v \in [0, u_L]$, $F_L(v) = F_{S/\kappa}(v)$, $\kappa_E = \kappa$ and the allocation is fully efficient.*
- ii. *If $\kappa < 1$ and $F_L \succ_{FO} F_{S/\kappa}$, $\kappa_E > \kappa$. If $F_L \succeq_{FO} F_S$, $\kappa_E = 1$.*
- iii. *If $F_{S/\kappa} \succ_{FO} F_L$, $\kappa_E < \kappa$ (note this includes the case when $\kappa = 1$).*
- iv. *If $\kappa = \kappa_E$, the large bidder neither always buys in the resale market nor always sells.*

Proof. The objective can be written as

$$S_{PR}(\kappa) = \int_l^{u_L} \int_0^{F_S(p(v; \kappa))} (v - g(q)) dq dF_L(v) + \int_l^{u_S} v F_L(\hat{p}^{-1}(v; \kappa)) dF_S(v). \quad (10)$$

where $\hat{p}^{-1} : [l, \bar{p}] \times [0, 1] \rightarrow [l, u_L]$ is defined as the solution to $\hat{p}(\hat{p}^{-1}(x, \kappa); \kappa) = x$, and we let $F_L(\hat{p}^{-1}(v; \kappa)) = 1$ for $v > \bar{p}$. Omitting function arguments, the first-order condition for an optimal

interior choice of κ simplifies as follows.

$$\begin{aligned} 0 &= \int_l^{u_L} (v - g(F_S(\hat{p}))) f_S(\hat{p}) \hat{p}_\kappa f_L dv + \int_l^{\bar{p}} v \hat{p}_\kappa^{-1} f_L(\hat{p}^{-1}) f_S dv \\ &= \int_l^{u_L} (v - g(F_S(\hat{p})) - \hat{p}) \hat{p}_\kappa f_S(\hat{p}) f_L dv \end{aligned} \quad (11)$$

The κ subscripts indicate partial derivatives. The first equality uses the change of variables $v = \hat{p}(w; \kappa)$ in the second term and the identity $\hat{p}_v(v; \kappa) \hat{p}_\kappa^{-1}(\hat{p}(v; \kappa), \kappa) + \hat{p}_\kappa(v, \kappa) = 0$.²⁴

For (i), if there exists a κ , such that $\hat{p}(v; \kappa) = v - g(F_S(\hat{p}(v; \kappa)))$ for all $v \in [l, u_L]$, then it is fully efficient and hence optimal. For (ii), $F_L \succ_{FO} F_{S/\kappa}$ implies that $\hat{p}(v_L; \kappa) < p^*(v_L)$ for all v_L . The derivative of the objective with respect to κ is therefore positive and by Lemma 1 the objective is increasing in κ and can be improved if $\kappa < 1$. The argument is nearly the same for (iii). The final statement, (iv), follows from the observation that if the large bidder always buys (ii) applies and if it always sells (iii) does. \square

More generally our construction implies the optimal cap equates $v_L - g(F_S(\hat{p}(v_L; \kappa)))$ and $\hat{p}(v_L; \kappa)$ “on average” where the averaging is determined by Equation (11) in the proof.²⁵ In Example 1 we showed that optimality when the large bidder has a flat demand requires equating the expectation of the first of two order statistics from the F_L and F_S distribution.

The condition under which no cap is optimal is worth emphasizing. We find that for *any* binding cap to be justifiable it must be the case that the small bidders’ type distribution F_S is not first-order stochastically dominated by F_L . Therefore caps only improve the efficiency of the auction plus resale mechanism when the large bidder’s distribution of values is sufficiently weak. A very weak large bidder may be interpreted as a speculator. Such a large bidder expects its small rivals to have higher values, bids aggressively and resells a large quantity in the resale market. This is similar to the intuition that weakness breeds aggression in asymmetric first-price auctions presented by Maskin and Riley (2000). The optimal quantity cap corrects this imbalance by having the large bidder only compete in the auction against small bidders with low types.

5 Revenue Maximizing Caps

Finally, we study how changes in the cap affect the revenue raised by the auction. Our two results in this section show that under conditions guaranteeing that the surplus maximizing cap is less than one (i.e., in the interior) the revenue maximizing cap is one, and when the surplus maximizing cap is one, the revenue maximizing cap is less than one. The policy implications for choosing a cap to maximize revenue are hence opposed to those for the cap that maximizes surplus.

²⁴To get the identity, implicitly differentiate $\hat{p}(\hat{p}^{-1}(w; \kappa); \kappa) = w$ with respect to κ and set $w = \hat{p}(v; \kappa)$.

²⁵Equation (11) may be made into a proper expectation if we divide by the positive term $\int_l^{u_L} \hat{p}_\kappa f_S(\hat{p}) f_L dv$.

We are interested in the revenue maximizing cap as it relates to the surplus maximizing cap, because this relation indicates the consequences for revenue of using a quantity cap to increase surplus. The revenue maximizing cap is on its own an imperfect policy instrument for revenue maximization, and we expect that policies such as reserve prices would be more effective. Our results therefore focus on the relation between the optimal caps resulting from the two objectives and not on providing a detailed characterization of the optimal revenue cap by itself.

The arguments in this section are organized as follows. We first derive a necessary condition for the cap to maximize revenue among all caps and show that it is the same condition that characterizes the efficient cap after marginal values are replaced by virtual values. We then provide an example in which the optimal cap can be solved for explicitly. Finally, we show that the results from the example can be generalized using regularity conditions on the distributions of the large and small bidders introduced by Kirkegaard (2012) to rank the expected revenue of first- and second-price auctions.

To derive the necessary condition for the revenue maximizing cap, we need an expression for the expected revenue. It is convenient to use symmetrization of the bid distributions to view the bids by the large bidder and the competitive small bidders as being generated from a symmetric auction with distribution $F(p)$ and corresponding bids, $\beta(p)$, defined in Equation (8). The expected revenue from such an auction is equal to the expected value of the second order statistic given two draws from $F(p)$. This is similar to how HK represent expected revenue, but we must add the revenue generated by small bidders who win with certainty. The expression for expected revenue is

$$\begin{aligned} S_R(\kappa) &= \kappa \int_l^{\bar{p}} (1 - F(p))^2 dp + (1 - \kappa) \int_l^{\bar{p}} (1 - F(p)) dp \\ &= \int_l^{\bar{p}} (1 - \kappa F(p))(1 - F(p)) dp. \end{aligned} \tag{12}$$

The first line is a weighted average of the second order statistic of $F(p)$ and the mean of $F(p)$, where the weight is determined by the fraction of goods sold to the large and competitive small bidders, κ . The mean of $F(p)$ is the highest bid submitted in the auction, $\beta(\bar{p})$, and is also the bid submitted by the small bidders with types $v_S > F_S^{-1}(\kappa)$ (Proposition 1). The second line rewrites the first to isolate the direct influence of κ . It is equivalent to $E[\min\{X, Y\}]$ where X is distributed according to $F(p)$ and Y according to $\kappa F(p)$.

The first line of Equation (12) indicates that relaxing the cap has two effects on expected revenue. First, by increasing κ a larger fraction of goods are sold to the large and competitive small bidders. If $F(p)$ were constant in κ this effect would reduce revenue, because the expected value of the lower of two draws from $F(p)$ is smaller than the mean. However, relaxing the cap increases both the second order statistic and the mean — the expected auction bid and the bid of the sure winning small bidders — because $F(p)$ decreases pointwise in κ . To see this, note that the definitions of $F(p)$ and $\hat{p}(v; \kappa)$ imply the identity $F_L(v) = F(\hat{p}(v; \kappa))$ must hold for all $v \in [l, u_L]$ and κ .

Since $\hat{p}(v; \kappa)$ increases in κ , this requires that $F(p)$ be decreasing in κ for all p . Therefore, relaxing the cap increases the expected bid in the auction, but also reduces the fraction of good sold at the highest bid, and we cannot be sure at this point which effect dominates.

Using techniques similar to those used in the proof of Proposition 2, we study the precise influence of κ on expected revenue by studying its derivative. We show that the derivative of expected revenue with respect to κ is equivalent to the derivative with respect to expected total surplus, (11), after replacing types with virtual valuations. Our derivation of the next expression is lengthy and is given in the Appendix. We again omit function arguments for simplicity.

$$S'_R(\kappa) = \int_l^{\bar{p}} f_S(\hat{p}) f_L \hat{p} \kappa \left\{ v - \frac{1 - F_L}{f_L} - \hat{p} + \frac{1 - F_S(\hat{p})}{f_S(\hat{p})} - g(F_S(\hat{p})) \right\} dv. \quad (13)$$

The sign of the integrand is positive whenever the virtual valuation of type- v large bidder exceeds the virtual valuation of the small bidder with value $\hat{p}(v; \kappa)$ plus $g(F_S(\hat{p}(v; \kappa)))$. Comparing this to the analogous expression for $S'_E(\kappa)$ in (11), we find that the only distinction is between types and virtual values. In particular, it is important for our results below that the bracketed terms have the same weight, $f_S(\hat{p}) f_L \hat{p} \kappa$. This implies, for example, that if the bracketed expression is always greater than the corresponding bracketed expression for expected surplus the revenue maximizing cap must be higher than the efficient one.

Before proceeding with the general case, we illustrate the implications for the revenue maximizing cap in a linear example similar to Example 1. We adjust the distributions and assume flat demand for the large bidder to simplify the algebra.

Example 2. Let $v_S \sim U[0, u_S]$, $v_L \sim U[0, u_L]$ and $g(\cdot) = 0$ with $u_S \geq u_L \geq 1$. We find that

$$v - \hat{p}(v; \kappa) = \frac{v}{2} \left(1 - \kappa \frac{u_S}{u_L} \right),$$

which implies that the efficient cap, κ_E , is u_L/u_S . The difference in virtual valuations — the bracketed term in (13) — is

$$2(v - \hat{p}(v; \kappa)) + u_S - u_L = \left(1 - \kappa \frac{u_S}{u_L} \right) v + u_S - u_L,$$

from which it follows that to maximize revenue there should be no cap (i.e., $\kappa_R = 1$), because

$$\left(1 - \kappa \frac{u_S}{u_L} \right) v + u_S - u_L \geq \left(1 - \frac{u_S}{u_L} \right) v + u_S - u_L \geq \left(1 - \frac{u_S}{u_L} \right) u_L + u_S - u_L = 0.$$

The example presents an extreme result. The surplus maximizing cap is strictly less than one while to maximize revenue the auctioneer should not use a cap (i.e., set $\kappa_R = 1$). Evaluated at v and $\hat{p}(v; \kappa)$, the difference in virtual valuations is always positive when $\kappa < 1$, implying that revenue

can always be increased by allocating more quantity to the large bidder. This obviously decreases surplus once $\kappa > u_L/u_S$.

In the example, whenever the large bidder has a weaker distribution than the small bidders (i.e., $u_L < u_S$) the efficient cap is strictly less than one but under these conditions the virtual valuation of the large bidder is always larger than the small bidder whose value is equal to the resale price. Using Equation (13), this implies that for all $\kappa < 1$, $S'_R(\kappa) > 0$ because the integrand is always positive. In other words, the weakness of the large bidder means that the efficient cap is less than one but also that it has a relatively high virtual valuation for all types, implying that the revenue maximizing cap is higher.

We next show that this extreme result holds more generally under the regularity conditions introduced by Kirkegaard (2012). To rank the expected revenue of the first- and second-price auctions for a single-unit, Kirkegaard (2012) uses two conditions on the relative strength of the two bidders' distributions. First, he assumes that the strong bidder's distribution dominates the weak bidder's distribution in terms of the hazard-rate order. The hazard-rate of a distribution G at x is $g(x)/(1-G(x))$ where g is the density function. The distribution F_S dominates F_L in the hazard-rate order, written $F_S \succeq_{HO} F_L$, if

$$\frac{f_L(v)}{1-F_L(v)} \geq \frac{f_S(v)}{1-F_S(v)} \quad \forall v \in (l, u_L). \quad (14)$$

Krishna (2002) discusses the relationship between first-order stochastic dominance and this one in the context of auction theory, showing that hazard-rate dominance implies first-order stochastic dominance. Note that the assumption of hazard-rate dominance holds for F_S and F_L in our linear examples with $v_S \sim U[l, u_S]$, $v_L \sim U[l, u_L]$, and $u_S > u_L$.

The second condition used by Kirkegaard (2012) relates the densities of the two bidders' distributions locally. Stated using F_L and F_S this condition translates to

$$f_L(v) \geq f_S(x) \quad \forall x \in [v - g(F_S(p^*(v))), F_S^{-1}(F_L(v))]. \quad (15)$$

It requires that the large bidder's distribution be increasing faster at v than the small bidder's distribution is over a particular interval of small bidder values. Note that as written this condition is necessarily stronger than Kirkegaard's to account for the large bidder's downward sloping demand. When the large bidder has flat demand the two are equivalent.²⁶ Together with $F_S \succeq_{HR} F_L$, these conditions guarantee that F_S dominates F_L in the dispersive order (Kirkegaard, 2012).

While the hazard-rate condition simply strengthens the notion that the large bidder is weaker than the small bidder, the second condition is more difficult to interpret. The condition in (15) is only used in our proof of Proposition 3 in cases where the large bidder is a seller in the resale

²⁶Directly translating Kirkegaard's Condition (9) would yield $f_L(v) \geq f_S(x) \quad \forall x \in [v, F_S^{-1}(F_L(v))]$. Due to the downward sloping demand of the large bidder, we need to strengthen this condition by expanding the set of x over which the condition holds given v . Our condition and Kirkegaard's are equivalent with flat demand.

market. When the large bidder is a seller, it sets a price above the Walrasian price so that $\hat{p}(v_L; \kappa) > v_L - g(F_S(p^*(v_L)))$. It also sets a price no greater than $F_S^{-1}(\kappa F_L(v_L))$ which is the price at which the large bidder would sell exactly zero units. This explains the relevance of the bounds on the interval of small bidder types in (15). If it were the case that $f_L(v) < f_S(x)$ for some x in this interval then in principle the virtual valuation for the small bidder with value x could exceed the type- v large bidder's virtual valuation. The condition rules out this possibility.

The next proposition formalizes this intuition to show that a revenue maximizing auctioneer should never use a cap under these conditions.

Proposition 3. *If $F_S \succeq_{HR} F_L$ and (15) holds, $\kappa_E < \kappa_R = 1$.*

Proof. From Proposition 2, $\kappa_E < 1$. Let $\kappa < 1$, $v_L = v$, and $p = p(v; \kappa)$. If $p \geq v - g(F_S(p))$, $F_L(v) \geq F_S(p)/\kappa > F_S(p)$ and hence $f_L(v) \geq f_S(p)$ since $p \in [v - g(F_S(p)), F_S^{-1}(F_L(v))]$. In this case, $S'_R(\kappa) > 0$ because for all v

$$\begin{aligned} & v - \frac{1 - F_L(v)}{f_L(v)} - p + \frac{1 - F_S(p)}{f_S(p)} - g(F_S(p)) \\ &= - \frac{\kappa F_L(v) - F_S(p)}{f_S(p)} - \frac{1 - F_L(v)}{f_L(v)} + \frac{1 - F_S(p)}{f_S(p)} \geq \frac{F_L(v) - \kappa F_L(v)}{f_L(p)} > 0. \end{aligned}$$

If $p < v - g(F_S(p)) \leq v$ and $F_L(v) < F_S(p)/\kappa$, $S'_R(\kappa) > 0$ because

$$\begin{aligned} & v - \frac{1 - F_L(v)}{f_L(v)} - p + \frac{1 - F_S(p)}{f_S(p)} - g(F_S(p)) > 2(v - p) - \frac{1 - F_L(v)}{f_L(v)} + \frac{1 - F_S(p)}{f_S(p)} > \\ & - \frac{1 - F_L(p)}{f_L(p)} + \frac{1 - F_S(p)}{f_S(p)} \geq 0. \end{aligned}$$

The second inequality follows because the large bidder has an increase hazard rate, while the final one follows from $F_S \succeq_{HR} F_L$. \square

If one reverses the conditions, making the large bidder relatively strong, the conclusions reverse as well. From Proposition 2 no cap should be used to maximize surplus in this case. We show next that in this case a cap should definitely be used to maximize revenue. More precisely, $F_L \succ_{HR} F_S$ implies $F_L \succ_{FO} F_S$, which implies that under any cap the large bidder is always a buyer in the resale market, and hence that the efficient cap is one. The assumption that $F_L \succ_{HR} F_S$ also implies that $u_S = u_L$ and for all $v \in [l, u_L]$, $v - (1 - F_L(v))/f_L(v) < v - (1 - F_S(v))/f_S(v)$. For the purposes of revenue maximization this suggests that the auctioneer should allocate more quantity to the small bidders, which it can do by tightening the cap.

Proposition 4. *Suppose $F_L \succ_{HR} F_S$ and $f_L(v) \leq f_S(x)$ whenever $x \in [F_S^{-1}(F_L(v)), v - g(F_S(p))]$, then $\kappa_R < \kappa_E = 1$.*

Proof. Let $V_i(v)$, $i \in \{S, L\}$, be the virtual valuation of a bidder with type v . We show that $S'_R(1) < 0$, which holds if $V_L(v) - g(F_S(p(v; 1))) < V_S(p(v; 1))$ for all $v \in [l, u_L]$. Since $p < v - g(F_S(p)) \leq v$ and thus $F_S^{-1}(F_L(v)) < p$ by (9),

$$\begin{aligned} V_L(v) - V_L(p) &< v - \frac{1 - F_L(v)}{f_L(v)} - F_S^{-1}(F_L(v)) + \frac{1 - F_L(v)}{f_S(F_S^{-1}(F_L(v)))} - g(F_L(v)) \\ &< -\frac{1 - F_L(v)}{f_L(v)} + \frac{1 - F_L(v)}{f_S(F_S^{-1}(F_L(v)))} < 0. \end{aligned}$$

□

The previous two propositions justify the sense in which the goals of revenue maximization and surplus maximization have different implications, and that except for obviously bad choices of caps (e.g., one that makes a weak large bidder always a buyer in the resale market), an improvement in one objective requires a decline in the other.

For an alternative explanation of the assumptions and the results, we briefly describe the relation between this problem and a corresponding third-degree pricing problem. In a textbook third-degree pricing problem with constant marginal cost, the optimal prices offered to the different segments equate (if possible) the marginal revenue of each segment. Building on similar intuition developed in Kirkegaard (2012) and Bulow and Roberts (1989), we describe our problem in these terms and argue that the Kirkegaard conditions imply that the revenue maximizing cap should be set to one when the large bidder is weak.

Suppose that the large bidder is relatively weak, $F_S \succeq_{HO} F_L$, that its demand is flat, $g(\cdot) = 0$, and think of $q_L(p_L) = 1 - F_L(p_L)$ and $q_S(p_S) = 1 - F_S(p_S)$ as demand curves in a third-degree pricing problem. Then the large bidder's weakness implies that it has lower demand at all prices. The marginal revenue from the large bidder at $q_L(p_L)$ is $p_L - (1 - F_L(p_L))/f_L(p_L)$, which is increasing in p_L by assumption. The Kirkegaard conditions imply that the marginal revenue from the large bidder at $q_L(p)$ exceeds the marginal revenue from the small bidders at $q_S(p)$. Efficiency requires that the price be the same for both bidders, but whenever the price is the same revenue would increase by allocating more to the large bidder.

When the large bidder is a buyer in the resale market it ends up with less quantity than is efficient. In terms of the corresponding third-degree pricing problem, this is analogous to offering the large bidder a higher price than the small bidders, $p_L > p_S$. Since marginal revenue is increasing in the price, the marginal revenue for the large bidder at p_L must be greater than the large bidder marginal revenue at p_S , which from the observations above is greater than the small bidders' marginal revenue at p_S . Therefore, at these prices revenue will increase if more quantity is allocated to the large bidder as well.

Alternatively, if the large bidder is a seller in the resale market, it is as if the large bidder is charged a lower price than the small bidders in the third-degree pricing problem, $p_L > p_S$. Because

the marginal revenue from the large bidder declines between p_L and p_S more work is required to determine whether the large bidder's marginal revenue is higher. In the equilibrium of the model p_S is determined by the large bidder and is never greater than $F_S^{-1}(F_L(p_L))$ when the large bidder sells. Notice that $q_S(F_S^{-1}(F_L(p_L))) = q_L(p_L)$. The Kirkegaard conditions guarantee that for all p_S between p_L and $F_S^{-1}(F_L(p_L))$, the large bidder has a flatter demand curve and higher marginal revenue. Therefore, we find again that at such prices the large bidder should receive more units.

In each case to increase revenue the allocation of quantity should be shifted to the large bidder. The analogous implication in our model is that the cap should be relaxed.

6 Conclusion

We introduce a novel model of a discriminatory price auction with subsequent resale stage and show how to tractably characterize equilibrium behavior for any initial quantity restriction (cap) on the multi-unit bidder. In our conception of the auction and resale market, a large (multi-unit) bidder competes against many small (single-unit) bidders and then sets a take-it-or-leave-it price in the resale market. We show by extending the results of Hafalir and Krishna (2008) that our model shares key features of a two-bidder auction with resale. However, the two models are not isomorphic. In our model, the large bidder may have downward sloping demand, and depending on the choice of cap not all of the small bidders compete directly in the auction with the large bidder. Our model is therefore a strict extension of the Hafalir and Krishna (2008) model.

We use this model to study the optimal choice of quantity cap on the large bidder under the objectives of surplus and revenue maximization. Our equilibrium characterization reveals that the cap influences the distributional strength of the small bidders against whom the large bidder competes in the auction. The consequences of adjusting the strength of these bidders should be familiar to students of asymmetric first-price auctions. For example, by restricting the cap the large bidder becomes relatively strong compared to the small bidders against whom it competes in the auction. Becoming strong causes the large bidder to bid less aggressively, win less quantity in the auction, and either sell less or purchase more in the resale stage. Full efficiency attains in the resale market if the large bidder neither purchases nor sells. Constrained efficiency requires that the large bidder neither always sell nor always purchase.

To maximize surplus we show that the auctioneer adjusts the cap to align the marginal value of the large bidder with the value of the marginal small bidder in the resale market (the small bidder who's value equals the resale price). To maximize revenue the auctioneer aligns the corresponding virtual valuations of Myerson (1981). Under regularity conditions given by Kirkegaard (2012) to rank revenue in the single-unit auctions literature, we show that when surplus maximization demands an interior choice of cap, revenue maximization requires that no cap be set at all. This suggests a strong trade-off between the two objectives in this environment, meaning that an efficient

cap generates suboptimal revenue and vice versa.

References

- Arozamena, L. and E. Cantillon (2004). Investment incentives in procurement auctions. *Review of Economic Studies* 71, 1–18.
- Ausubel, L. M. and P. Cramton (1999). The optimality of being efficient. Working Paper, Department of Economics, University of Maryland.
- Back, K. and J. F. Zender (1993). Auctions of divisible goods: on the rationale for the treasury experiment. *Review of Financial Studies* 6(4), 733–764.
- Baisa, B. and J. Burkett (2016). Large multi-unit auctions with a large bidder. Working Paper.
- Bartolini, L. and C. Cottarelli (1997). *Treasury Bill Auctions: Issues and Uses*, pp. 267–336. London: Routledge.
- Brady, N., R. Breen, and A. Greenspan (1992). Joint report on the government securities market. *US Government Printing Office*.
- Bulow, J. and J. Roberts (1989). The simple economics of optimal auctions. *Journal of Political Economy* 97, 1060–1090.
- Garbade, K. and J. Ingber (2005). The treasury auction process: Objectives, structure, and recent adaptations. *Current issues in economics and finance* 11(2).
- Gilbert, R. J. and D. M. G. Newberry (1982). Preemptive patenting and the persistence of monopoly. *American Economic Review* 72(3), 514–526.
- Hafalir, I. and V. Krishna (2008). Asymmetric auctions with resale. *American Economic Review* 98(1), 87–112.
- Hafalir, I. and M. Kurnaz (2015). Discriminatory auctions with resale. Working Paper.
- Jegadeesh, N. (1993). Treasury auction bids and the salomon squeeze. *The Journal of Finance* 48(4), 1403–1419.
- Kirkegaard, R. (2012). A mechanism design approach to ranking asymmetric auctions. *Econometrica* 80(5), 2349–2364.
- Krishna, K. (1993). Auctions with endogenous valuations: The persistence of monopoly revisited. *American Economic Review* 83, 147–160.

Krishna, V. (2002). *Auction Theory*. Academic Press.

Lebrun, B. (1998). Comparative statics in first price auctions. *Games and Economic Behavior* 25, 97–110.

Maskin, E. and J. Riley (2000). Asymmetric auctions. *Review of Economic Studies* 67, 413–438.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73.

Swinkels, J. (1999). Asymptotic efficiency for discriminatory private value auctions. *Review of Economic Studies* 66, 509–528.

Swinkels, J. (2001). Efficiency of large private value auctions. *Econometrica* 69(1), 37–68.

A Appendix

Requirement that $u_S \geq u_L$. A problem arises in the model if $u_L > u_S$. Consider the case where $g(\cdot) = 0$ and $\kappa = 1$. Then to define $F(p)$ we need to find a \bar{p} such that

$$u_L = \bar{p} + \frac{F_S(\bar{p}) - 1}{f_S(\bar{p})}, \quad (16)$$

but this requires that $\bar{p} > u_S$, which would imply that in equilibrium the highest type of large bidder offers a price greater than u_S in the resale market. This is clearly not optimal and the equilibrium actually falls apart. In this situation, having a positive measure of large bidder types set $p = u_S$ breaks the equilibrium as well by encouraging small bidders to bid more aggressively than the proposed equilibrium. We found no way to fix the issues that arise here when $u_L > u_S$ for all g and κ . In their online appendix, Hafalir and Krishna (2008) argue that no problem arises in their model if they assume that the losing bidder buys from the winning bidder. One can check in their model that when the buyer has higher types than the largest of the purchasers types their equilibrium requires that the buyer offer a sub-optimal price in the resale stage, exactly as we show here. They are therefore implicitly assuming a common support in this case. \square

Proof of Proposition 1. The proposed b_L is increasing and so is b_S for competitive small bidders. Furthermore, by definition

$$F(\phi(b)) = F_L(\phi_L(b)) = F_{S/\kappa}(\phi_S(b)), \quad (17)$$

where $\phi = \beta^{-1}$, and symmetrization holds with these bid functions.

Given that for all equilibrium bids $F(\phi(b)) = F_L(\phi_L(b))$, the type- v large bidder's first-order condition for its bid (4) holds at the proposed choice of bid, $b_L(v)$, when $\phi(b) = p(b, \phi_L(b))$, which

is exactly when the resale price is chosen optimally. If the large bidder were to bid $b' \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} b_L(v)$,

$$\frac{\partial}{\partial b} \pi_L(v, b', p(b', v)) = \phi'_S(b') f_{S/\kappa}(\phi_S(b')) (p(b', v) - b') - F_{S/\kappa}(\phi_S(b')) \begin{smallmatrix} \leq \\ \geq \end{smallmatrix} 0,$$

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

For the competitive small bidders, it turns out that given $\phi(b) = p(b, \phi_L(b))$ the first-order condition (5) holds for any bid in $[0, \bar{b}]$, since they internalize the effect on $p(b, \phi_L(b))$. This implies that they are indifferent between submitting any bid.²⁷

The small bidders with values larger than $F_S^{-1}(\kappa)$ are the sure winners in the auction. Notice that (5) holds for them as well at any bid, including \bar{b} . The fact that a non-zero measure of small bidders bids \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. \square

Derivation of Equation (13).

$$\begin{aligned} S_R(\kappa) &= \int_l^{\bar{p}} (1 - \kappa F_L(p^{-1})) (1 - F_L(p^{-1})) dp. \\ S'_R(\kappa) &= \int_l^{\bar{p}} \{ -F_L(p^{-1}) (1 - F_L(p^{-1})) - \kappa f_L(p^{-1}) p_{\kappa}^{-1} (1 - F_L(p^{-1})) \\ &\quad - f_L(p^{-1}) p_{\kappa}^{-1} (1 - \kappa F_L(p^{-1})) \} dp \end{aligned}$$

²⁷HK show sufficiency of the first-order conditions describing the bids in their Proposition 4. Their argument is complicated by the fact that depending on the bid submitted in the auction the role of the bidder in the resale market as a price setter or a price taker changes, due to their assumption that the winner determines the resale price. This means that the form taken by the derivative with respect to the bid changes depending on the bid at which it is evaluated. Under our resale assumptions, these roles do not change based on the bid (e.g., the large bidder always chooses the price, although it is sometimes a buyer and sometimes a seller). The result is that the derivatives with respect to the bids take the same form across the relevant bids.

Using the definition of p , (9), $p_{\kappa}^{-1} = -F_L(p^{-1})/(f_S + \kappa f_L(p^{-1}))$, from which we get

$$\begin{aligned}
S'_R(\kappa) &= \int_l^{\bar{p}} \left\{ -F_L(p^{-1})(1 - F_L(p^{-1})) + \kappa \frac{F_L(p^{-1})f_L(p^{-1})}{f_S + \kappa f_L(p^{-1})}(1 - F_L(p^{-1})) \right. \\
&\quad \left. + \frac{F_L(p^{-1})f_L(p^{-1})}{f_S + \kappa f_L(p^{-1})}(1 - \kappa F_L(p^{-1})) \right\} dp \\
&= \int_l^{\bar{p}} \frac{f_S f_L(p^{-1}) F_L(p^{-1})}{f_S + \kappa f_L(p^{-1})} \left\{ \frac{1 - \kappa F_L(p^{-1})}{f_S} - \frac{1 - F_L(p^{-1})}{f_L(p^{-1})} \right\} dp \\
&= \int_l^{\bar{p}} \frac{f_S f_L(p^{-1}) F_L(p^{-1})}{f_S + \kappa f_L(p^{-1})} \left\{ p^{-1} - \frac{1 - F_L(p^{-1})}{f_L(p^{-1})} - p + \frac{1 - F_S}{f_S} - g(F_S) \right\} dp.
\end{aligned}$$

The last equation uses the definition of p again. The term outside of the brackets is $-p_{\kappa}^{-1}(p; \kappa) f_S(p) f_L(p^{-1}(p; \kappa))$. By differentiating the identity $p^{-1}(p(v; \kappa)) = v$ separately with respect to v and κ , we derive $-p_{\kappa}^{-1}(p(v; \kappa); \kappa) p_v(v; \kappa) = p_{\kappa}(v; \kappa)$. Therefore the change of variables $p = p(v, \kappa)$ yields

$$S'_R(\kappa) = \int_l^{\bar{p}} f_S f_L(p^{-1}) p_{\kappa} \left\{ v - \frac{1 - F_L}{f_L} - p + \frac{1 - F_S(p)}{f_S(p)} - g(F_S(p)) \right\} dv.$$

□