Contests for Priority Access

Brennan C. Platt*†

May 13, 2011

Abstract

This paper analyzes two contests in which customers make non-refundable bids to determine the order in which they are served. In the auction for priority access, customers are served in descending order of their bids, where each in turn is allowed to purchase any amount of a divisible good at an exogenous per-unit price, until the fixed supply is exhausted. This model is immediately applicable to kickbacks paid to circumvent price controls, allocation by queuing, and other forms of rent seeking.

This auction implements a two-part tariff, allowing the seller to capture nearly all rents. Revenue is greatest when bidders have similar valuations, and when exactly one customer is unable to make a purchase. Unlike all-pay auctions of an indivisible prize, the exclusion principle — that revenues increase when the strictly-highest-valuation bidder is excluded from participation — does not always hold.

The auction is also compared to lotteries for priority access, an extension of single prize contests. The auction raises more revenue than the lottery when agents' valuations are not too different. Also, the auction is typically more efficient, resulting in higher expected total welfare.

Keywords: contests, all-pay auctions, priority access, rationing, exclusion principle, two-part tariffs

JEL Classification: D44, D45, D72

*Department of Economics, Brigham Young University, 135 FOB, Provo, UT 84602, (801) 422-8904, brennan_platt@byu.edu

†The author thanks Beth Allen, Val Lambson, Nuray Akin, and seminar participants at the University of Miami, the 2010 Stony Brook Game Theory Festival, and the 2011 Tournaments, Contests and Relative Performance Evaluation Conference for fruitful discussions on this work. Partial funding came from NSF grants DMI-0070257 and DMI-0217974. All errors remain my own.

1 Introduction

When a price ceiling is imposed in a market, sellers often find creative ways to circumvent its effects. One such method is to have customers compete for the order in which they are served, each offering (under-the-table) payments to obtain higher priority. For instance, would-be tenants of rent-controlled apartments may offer "key money" to promote their application to the top of the waiting list. Similarly, patrons might pay a store owner for first access to the next shipment of goods in short supply.

These sorts of payments clearly violate the spirit (and, often, the letter) of price control laws, since customers effectively pay more than the legally prescribed price. However, it would likely be easier to keep such payments hidden than if the seller simply charged a higher price per unit. Access payments would be fewer in number than per-unit payments, and thus easier to keep off-the-books and unadvertised. Moreover, once access is granted, the contracts that follow (such as leases or purchase orders) use the official price, while any agreement to higher per-unit prices (such as a monthly rent supplement) would be unenforceable in the legal system.

Access payments also offer a more subtle benefit: the seller can effectively use these as a form of second degree price discrimination to capture much (if not all) of the consumer surplus. Indeed, this method can generate greater profits than selling all units at the market clearing per-unit price. Ironically, this price discrimination scheme is only possible because of the price-ceiling-induced shortage, which forces buyers to compete for access.

This paper investigates access payments by developing a model of contests for priority access.[1] We consider a single seller with a fixed supply of a perfectly divisible good, to be sold among several buyers. We characterize two contest mechanisms for allocating the limited resource: an *all-pay auction* and a *lottery*.

In the auction for priority access, each buyer simultaneously submits a nonrefundable bid (access payment), which determines the priority order of service. The agent with the highest bid is given the first opportunity to procure some of the good, paying an additional fixed amount per unit. If any supply remains, access to the good is then granted to the second highest bidder, and so forth. Any ties are randomly broken.

[1]In addition to kickbacks during price controls, described above, our model is readily applicable to rationing by queuing (such as for concert tickets or Black Friday deals) and other forms of rent seeking. We reinterpret our results for these environments in the Conclusion.

In the lottery for priority access, buyers again submit non-refundable bids, but each buyer's probability of winning first priority is equal to that buyer's bid divided by the sum of all other bids. Similarly, second priority is determined randomly from the same bids, except that the first priority winner is excluded. This process iterates until the supply is exhausted.[2]

The auction environment provides more technical challenge than the lottery; thus, an important contribution of this paper is to characterize its equilibria. An analytic solution is not always possible, but we identify a number of features that are highly useful to narrow the search when solving numerically.

Beyond this technical contribution, we demonstrate that the auction implements a form of 2nd degree price discrimination, where the bids act as an entry fee whose level is endogenously determined by competition among the bidders. The bidding can extract much of the consumer surplus, particularly when buyers have similar preferences, which also occurs in the standard analysis of two-part tariffs (as in Oi, 1971). In the extreme case in which preferences are identical, the seller can fully extract rent as long as at least one bidder will be unable to purchase any of the good.[3]

Prior work on all-pay auctions with a single indivisible prize (Baye, *et al*, 1993; Menicucci, 2006) has identified an *exclusion principle*, meaning that the seller can increase expected bid revenue by excluding some of the bidders from the auction. Indeed, if anyone, it is the bidders with the highest valuations who should be excluded. This occurs when the presence of one dominant bidder will discourage other participants from placing moderately high bids.

In our divisible good model, the exclusion principle must be altered. For instance, in our three-bidder example, excluding a bidder only generates greater revenue when the value of placing second is relatively small and thus less important to bidding strategies. Moreover, the seller may prefer to exclude a bidder with the lowest valuation rather than the highest.

[2] This lottery process is equivalent to a *raffle*. By purchasing more tickets (*i.e.* a larger access payment), the participant improves his odds of being drawn as the first winner; but this is always relative to the total number of tickets sold.

[3] This is similar to the outcome of *pay-to-bid auctions*, where bidders pay a fee to increase the current auction price. As modeled by Platt, Price and Tappen (2010), the pay-to-bid auction fully dissipates expected consumer surplus even though the typical closing price is well below retail. Several nascent internet sites have implemented such auctions; data from one such site provides strong evidence confirming the model.

Finally, we compare the performance of the two types of contests on the dimensions of revenue and efficiency. While it is natural to think that first priority *always* goes to the highest bidder (as in the auction), it is also reasonable to consider a less deterministic outcome. Other random factors may enter into the seller's decision, and with under-the-table access payments, bidders would not have legal recourse to enforce their claim to a particular priority. The lottery environment might capture this uncertainty, where the bidder can improve his chances with larger payments, but has no guarantee. Indeed, the lottery and the all-pay auction are each special cases of the *contest success function* introduced by Tullock (1975, 1980) in the context of a contest for a single indivisible prize. Fang (2002) provides an extensive comparison of these two cases.[4]

Our results are characterized in terms of each bidder's *net gain* in utility from placing first versus second, or second versus third. The auction is more lucrative than the lottery when the net gain (from placing first or from placing second) is somewhat close across bidders. As in Fang (2002), auctions inspire greatest competition among bidders when valuations are similar; but here, it may be enough to have similar valuations of second place, even if the valuation for first place is quite uneven.

In terms of efficiency, the results are less definitive. The lottery can offer small efficiency gains over the auction when the net gains are quite similar, but for most parameter values, the auction generates a larger total welfare. A significant exception arises whenever the seller optimally excludes the highest valuation bidder from the auction. While this generates more revenue than either the lottery or the auction without exclusion, it guarantees that the good will be allocated to bidders who value it much less. If so, the lottery performs much better since the high valuation bidder wins first priority with high probability.

The study of all-pay auctions initially examined competition for a single prize among various participants who may have the same (Hillman and Samet, 1987) or differing (Hillman and Riley, 1989; Baye, *et al*, 1996) valuations of the prize. Others (Clark and Riis, 1998; Barut and Kovenock, 1998; Siegel, 2009) have extended the all-pay auction to contests with multiple prizes; still, each prize is indivisible and each

[4]Taylor, *et al* (2003) undertake a related study with m identical, indivisible units offered to bidders with unit demand. However, their lottery consists of randomly selecting m participants with equal probability, without any effort on their part. Also, their auction has bidders valuations as private information. Quite the opposite of our result, they find that the lottery is more efficient than the auction when consumers are more homogeneous.

agent can win no more than one prize. Our paper is the first analysis of an all-pay auction in which agents may choose varying amounts of a perfectly divisible good.

Our paper also contributes to a recent strand of the all-pay literature on rank-order spillovers. Baye, *et al* (2009) consider two-person contests in which both the winner and the loser can be directly affected by the bid of the other person. Most closely related is Klose and Kovenock (2011), which examines an n bidder all-pay auction where each bidder's payoff is affected by the identity of the winner. Our model is more general in that the payoff to a bidder in third place depends on the identities (*e.g.* purchases) of first *and* second place; in addition, the value of second and third place can differ in our model, even when the first-place winner is held constant. On the other hand, our model is set in a service-order environment, which restricts the type of payoffs that may occur, while Klose and Kovenock allow for a wider variety of spill-overs between first place and the remaining competitors.

This paper is also related to *priority auctions* (surveyed in Hassin and Haviv, 2002, Sec. 4.5). These are queuing models in which a server must determine an order in which to serve randomly-arriving customers, who can acquire an early position in the queue through their bids. Typically, all of the customers can eventually be served; the purpose of a higher bid is to reduce delay by advancing ahead of other bidders.

While these queuing processes differ from our deterministic environment with a fixed number of customers and a short supply, there are some parallels in the results. For instance, if customers are identical, bidding reduces their expected welfare to zero, as long as some of the customers choose not join the queue in equilibrium (because they are indifferent between the expected cost of waiting and the eventual benefit from service), and this outcome is efficient. However, unlike our model, the priority auction is also efficient even if the benefit from service or cost of waiting varies across customers, because the lowest benefit or highest cost customers are the ones who abstain. In our environment, these customers still participate in the bidding and occasionally displace higher benefit customers.

We assume that bidders are fully informed of each others' preferences. Bidders also know beforehand the available supply and the per-unit price charged once a person gains access. The assumption of full information, while not common in first- and second-price auctions, is frequently used in all-pay auctions (including all auction articles mentioned above), as well as in other models of rent seeking. As commonly occurs in a complete-information auction environment, agents bid using mixed strategies in equilibrium. We also assume that resale is not possible.[5]

The paper proceeds as follows: Section 2 defines the environment and mechanism for the auction for priority access. Section 3 then proves existence of equilibrium and characterizes its basic features. In Section 4, we present the full solution in several special cases, including n homogeneous agents and 2 or 3 heterogeneous agents. Section 5 defines the lottery for priority access and compares its equilibrium outcomes to the auction outcomes from the preceding section, and section 6 concludes and offers alternative applications. All proofs appear in the Appendix.

2 Auction Model

A seller has S units of a perfectly divisible commodity, to be sold among n buyers via an auction for priority access. Each buyer submits a sealed bid b_i, and these are simultaneously opened by the seller. The buyer with the largest b_i is then allowed to purchase any amount q_i (up to S units of the good) at an exogenously set price $p \geq 0$ per unit. If any supply is available after this purchase, the buyer with the second largest bid is allowed to purchase from the remaining units at the same price p. This continues until supply is exhausted. Ties are broken randomly.

All players forfeit their bid, regardless of whether they obtain any supply (an *all-pay* auction). The model is only slightly more complicated if only those who obtain some supply are required to pay (a *pay-as-bid* auction); we do not pursue that analysis here. The bid is constrained to be positive but less than an exogenous maximum M.[6]

Buyers are assumed to have preferences that are quasi-linear with respect to money: $u_i(q_i, b_i) = v_i(q_i) - pq_i - b_i$. We assume $v_i(0) = 0$, and that $q_i \geq 0$. Moreover, marginal utility is diminishing in the auctioned commodity: $v'' < 0$. Thus, we can define $d_i(p) \equiv \arg\max_{q_i} u_i(q_i, b_i)$ as the quantity buyer i would purchase at price p if supply is not binding.[7] With the assumed utility, if available supply is less

[5]If resale were possible, the highest bidder would always purchase the entire supply then sell any portion that he does not wish to consume. The constraints on method of resale are of critical importance, though. If this first winner can only resell under the same rules as the original seller (that is, a constant per unit price of p, while auctioning the order of service), the outcome would be virtually the same. If the first winner has greater latitude in methods of resale, the outcome could vary greatly; but it also begs the question as to why the original seller was denied this latitude.

[6]In making use of this model, one would typically set M sufficiently large to never be binding, such as larger than the maximum value of winning. For existence purposes, it is convenient to have a bounded space of actions.

[7]Note that d_i is independent of b_i, which is precisely the purpose of assuming quasi-linear utility.

than $d_i(p)$, the consumer would optimally choose to purchase all the available supply. These preferences are commonly known among all buyers, allowing each to correctly predict how much supply will remain after a given order of buyers are served.

Thus, for a given profile of bids $b = \{b_i\}_{i=1}^n$, if agent i is not tied with another agent ($b_i \neq b_j$ for all $j \neq i$), he will obtain:

$$q_i(b) \equiv \min\left\{d_i(p),\ \max\left\{0,\ S - \sum_{j\colon b_j > b_i} d_j(p)\right\}\right\}. \tag{1}$$

If he is involved in a tie, the tied bidders are randomly ordered (with equal probability on each permutation) and their demand is filled in order. Thus, if k agents (including agent i) have a bid of b_i, then with probability $\frac{1}{k!}$, the outcome is:

$$q_i(b) \equiv \min\left\{d_i(p),\ \max\left\{0,\ S - \sum_{j\colon b_j > b_i} d_j(p) - \sum_{\substack{j\colon b_j = b_i,\\ \pi(j) < \pi(i)}} d_j(p)\right\}\right\}, \tag{2}$$

where $\pi(\cdot)$ is a permutation that creates a strict order of the k tied bidders. Note that the tie could be *inconsequential*, meaning that $q_i(b)$ is the same for any of the permutations. For instance, $q_i(b) = d_i(b)$ for all permutations when there is sufficient remaining supply for all the tied agents. Similarly, if all supply has been exhausted by agents with bids higher than b_i, $q_i(b) = 0$ for all permutations.

Typically, there are no pure strategy equilibria for this game, as will be shown later. Thus, we define each bidder's mixed strategy as a probability measure μ_i on $[0, M]$. By way of notation, let B_i be the support[8] of μ_i. Also let B_{-i} be the cross product of the strategy support of each agent besides i, and $\mu_{-i}(b_{-i})$ be the product of $\mu_j(b_j)$ for all $j \neq i$.

The Nash equilibrium of this game is a strategy profile μ^* such that for all i, any bid in the support B_i^* maximizes i's expected utility, given other players' strategies μ_{-i}^*. That is, for all $b_i^* \in B_i^*$ and all $\hat{b}_i \in [0, M]$,

Otherwise, the sunk cost of the bid could reduce wealth and affect demand; eliminating these wealth effects greatly simplifies bidding strategies.

[8]The *support* of a probability measure μ_i is the set B_i such that $\mu_i(B_i) = 1$ and $\mu_i(B') < 1$ for any proper closed subset $B' \subset B_i$.

$$EU_i(b_i^*, \mu_{-i}^*) \equiv \int_{B_{-i}^*} \Big(v_i(q_i(b_{-i}, b_i^*)) \;-\; pq_i(b_{-i}, b_i^*) \Big)\, d\mu_{-i}^*(b_{-i}) \;-\; b_i^*$$

$$\geq \int_{B_{-i}^*} \Big(v_i(q_i(b_{-i}, \hat{b}_i)) \;-\; pq_i(b_{-i}, \hat{b}_i) \Big)\, d\mu_{-i}^*(b_{-i}) \;-\; \hat{b}_i.$$

3 Existence and Characterization of Equilibrium

As is often the case for games with continuous action spaces, this auction typically does not have a Nash equilibrium in pure strategies. This is because payoffs can be discontinuous if multiple agents choose the same bid. By bidding an infinitesimal amount more, an agent can avoid the tie and ensure that he gets first access to the goods — and if supply is constraining, this creates a discrete jump in utility. This is proven formally in Proposition 2, Claim 2. However, we can prove existence in mixed strategies in full generality.

Proposition 1. *For any $p \geq 0$, a mixed strategy equilibrium exists.*

Beyond existence, we can also significantly narrow the set of potential equilibria by establishing (in the following proposition) key features of equilibrium strategies. Indeed, these observations are quite useful as one seeks to calculate such an equilibrium, as is done in Section 4.

Proposition 2. *If μ^* is an equilibrium strategy in an auction for priority access,*

1. *For all i, $B_i^* \subset \cup_{j \neq i} B_j^* \cup \{0\}$.*
2. *For all i, there is no $b_i > 0$ such that $\mu_i^*(\{b_i\}) > 0$.*
3. *$\cup_i B_i^*$ is a connected set, and $0 \in \cup_i B_i^*$.*
4. *For all i, $EU_i(\mu^*) \geq v_i(0)$.*
5. *If $\mu_i^*(\{0\}) > 0$ for some i and $S \leq \sum_{j \neq i} d_j(p)$, then $EU_i(\mu^*) = v_i(0)$.*
6. *If $\mu_i^*(\{0\}) > 0$ for some i and $\sum_{j \neq i} d_j(p) < S \leq \sum_j d_j(p)$, then*

$$EU_i(\mu^*) = u_i\left(S - \sum_{j \neq i} d_j(p), 0\right).$$

7. If $\max B_i^* = \max \cup_j B_j^*$ *for some* i*, then*

$$EU_i(\mu^*) = u_i\left(\min\{d_i(p), S\}, \max B_i^*\right).$$

The first claim establishes that an agent will not include a particular bid in his support unless it also appears in the support of some other player. There is no reason to do so, since in such a case, one could reduce his bid without changing the expected amount of supply available for purchase. The exception, of course, is 0, since it is not possible to reduce the bid any further.

The next observation is that atoms[9] never occur except possibly at zero. If a consequential tie were to occur with positive probability (*i.e.* two agents have an atom on the same bid), at least one of them would strictly prefer to break the tie. If consequential ties almost never occur, two possibilities exist. On the one hand, it could be that if the agent with an atom reduced his bid, he would obtain the same purchases with the same probabilities but lower bid cost. If not, there must be some other agent bidding just below the atom, whose outcome depends on his order relative to the first. But this second agent can do strictly better by raising his bid just above the atom, for then he reduces his chances of purchasing behind the first while incurring an infinitesimal increase in bid cost.

An immediate corollary of the second claim is that no pure strategy equilibrium exists. The only exception is when all agents set $\mu_i(\{0\}) = 1$, and that could only be sustained if supply is weakly greater than aggregate demand.

The third claim states that there can be no gaps in the aggregate support; otherwise, agents bidding just above the gap would do well to reduce their bid inside of the gap. By the same reasoning, the aggregate support always includes 0.

The fourth claim identifies that equilibrium expected utility is bounded below by autarky. This is because agents are always free to bid nothing, $b_i = 0$. The worst that can happen in such a case is that the agent comes in last and obtains nothing. Note that this claim only applies to expected utility; once the randomization on mixed strategies occurs, agents may find themselves having a high bid yet still being beaten by some other agent and left without any opportunity to make a purchase. Yet other ex-post realizations may place the agent first in line in spite of a small bid.

The last three claims are particularly useful in computing equilibria, since they

[9]An atom of a measure μ_i occurs at strategy b_i if $\mu_i(\{b_i\}) > 0$.

pin down an agent's utility relative to certain aspects of his strategy. Whenever an atom occurs in his mixed strategy, his expected utility is the same as if he bid nothing and won exactly the leftovers after all others are served (or zero if there is insufficient supply to have leftovers). On the other hand, anyone whose support includes the highest bid receives an expected utility as if he always wins the first priority (claiming either his full demand or, if not possible, the entire supply) and bidding his highest bid for certain.

4 Examples

To illustrate the equilibrium behavior and provide more detailed analysis, we now examine several analytically tractable special cases. Our goal is to demonstrate key elements that are common to any solution (including numerical solutions for environments with more bidders or greater heterogeneity), as well as to demonstrate the process of computing the equilibrium strategies.

The real innovation of our model occurs with 3 or more heterogeneous bidders, since then bidders must consider not only whether they are outbid but by whom. Even so, we begin by considering the cases of n homogenous bidders and of 2 heterogeneous bidders, which can be analyzed by adapting prior results in the literature. These highly tractable cases serve as a benchmark, providing useful intuition for threads shared in common with the case of three heterogeneous bidders, while also providing contrast for the distinct behavior in that more complex environment.

4.1 n Homogeneous Bidders

Suppose all agents have the same underlying utility function, u_i. In this environment, agents have equal demand d for the good; thus, the highest k bidders, where $k \cdot d \leq S < (k+1)d$, will receive their full demand. The $k+1^{\text{th}}$ bidder will receive the leftovers, and all others will receive nothing.

With the assumption of homogenous bidders, our model can be translated into the environment of Barut and Kovenock (1998). They present an all-pay auction for multiple indivisible prizes, and all agents agree on the value of the various prizes. Here, the top k bidders receive a prize which everyone agree is worth $w \equiv v(d) - pd$. The $k+1^{\text{th}}$ prize is valued by all bidders at $x \equiv v(r) - pr$ where $r \equiv S - kd$, and all others receive a prize of $v(0)$, which we normalize to 0. An application of their results provides the following characterization. From this point on, we express the equilibrium mixed strategies as cumulative distribution functions, $F_i(b)$, rather than probability measures μ_i.

A symmetric equilibrium always exists,[10] in which all agents play the same mixed strategy. If $S < d$, then only the first place bidder receives any amount of the good. The unique equilibrium strategy, expressed as a cumulative distribution function, is $F_i(b) = \left(\frac{b}{x}\right)^{\frac{1}{n-1}}$ for all i, which has support $B_i^* = [0, x]$.

If a larger supply is available ($1 \leq k < n-1$), the symmetric equilibrium strategy $F_i(b)$ is implicitly expressed as:

$$b = x \binom{n-1}{k} F_i(b)^{n-k-1}(1 - F_i(b))^k + w \sum_{m=0}^{k-1} \binom{n-1}{m} (1 - F_i(b))^m (F_i(b))^{n-m-1},$$

where the parenthetical element represents the binomial coefficient of $n-1$ choose k. This will have a support of $B_i^* = [0, w]$ and has no atoms. With a large enough supply that $k = n - 1$, the equilibrium strategy is represented by $F_i(b) = 1 - \left(\frac{w-x-b}{w-x}\right)^{\frac{1}{n-1}}$ for all i, which has a support of $B_i^* = [0, w - x]$. Figure 1 illustrates this equilibrium strategy for three homogeneous bidders, for various levels of S.

Applying Claim 7 of Proposition 2, note that $EU_i(\mu^*) = 0$ (or $EU_i(\mu^*) = x$ if $k = n-1$). The expected expenditures on bids is $R = kw + x$ (or $R = (n-1)(w-x)$ if $k = n - 1$).

While this solution in the homogenous case is not novel, its application to our environment is. We now consider the efficiency and revenue implications of an auction for priority access.

One might address allocative efficiency on two levels. If we take as given that the quantities are split into k packages of size d and one of size r, then since everyone has the same utility from these packages, any distribution of them is equally efficient (*i.e.* on the extensive margin). However, if we allow alternative divisions of the S units, this allocation is certainly inefficient (*i.e.* on the intensive margin) because $v(\cdot)$ exhibits diminishing marginal utility; agents who end up with nothing after the

[10] When $1 \leq k < n-1$, there are also a continuum of equilibria with asymmetric mixed strategies, where our identical agents nonetheless employ differing mixed strategies. We refer the reader to Barut and Kovenock (1998) for a full characterization of these strategies. Indeed, despite their complexity, they result in the same expected utility and expenditure as the symmetric case, and thus add little to the performance of the auction.

Figure 1: Equilibrium Bidding Strategies, 3 Bidder Example

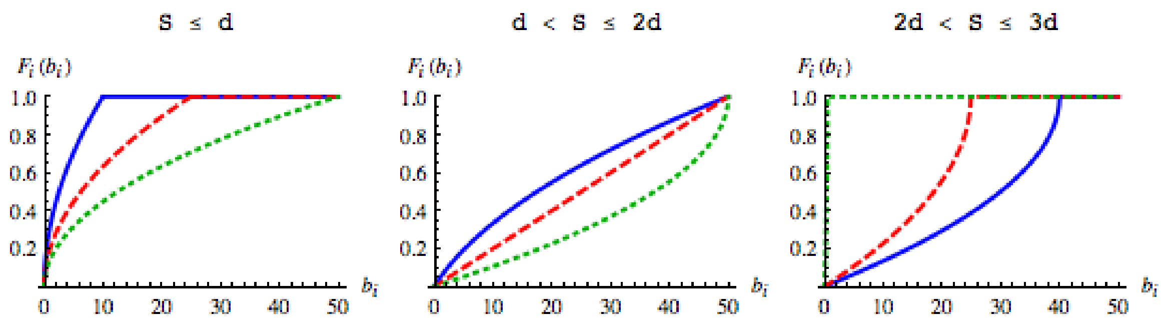

Notes: These plots depict CDF of the mixed strategy employed by all three bidders in the symmetric equilibrium, depending on the amount of good available. The three panes show when there is only enough supply for 1st place ($S \leq d$), for 1st and 2nd place ($d < S \leq 2d$), or some for all three ($2d < S \leq 3d$). We fix $w = 50$, and vary the utility from leftovers: $x = 10$ (solid), $x = 25$ (dashed), and $x = 50$ (dotted).

auction will value a marginal unit of the good higher than those who obtained d units.

Also note that since $EU_i = 0$, rent is fully dissipated when $S \leq (n-1)d$. In other words, if at least one agent will be unable to make purchases, they will all bid such that ex-ante, they are indifferent between participating in the auction or not.

Thus the auction for priority access implements a two-part tariff. The traditional two-part tariff requires the seller to know the consumer surplus of his customers, lest he charge too high an entry fee and exclude them from the market. Here, the seller can be ignorant of consumer surplus, as competition among bidders will endogenously determine the right entry fee.

Indeed, the seller could potentially earn more profit under the regulated price, p, than the market clearing price, p_{mkt}. This occurs when $(p_{mkt} - p) \cdot S < R$. In the case of linear demand, we can establish a more concrete result: the auction will be more profitable as long as at least half the market is still served.

Proposition 3. *Assume $x = 0$ and each bidder has linear demand for the good. Then the auction for priority access generates more revenue than the market clearing price if and only if $k > \frac{n}{2}$.*

The assumption that $x = 0$ (there are no leftovers for the $k+1^{\text{th}}$ bidder) is merely for expositional convenience and has little effect on the result.

One should note that this establishes the best case scenario. As we introduce heterogeneity, the auction will not achieve full rent extraction (which is also true of two-part tariffs in general). Even so, it is interesting that the artificially low price could benefit the seller, subverting the intention of the policy maker who presumably imposed the price control to reduce costs to consumers. In fact, the seller needs the price control as a commitment device. Suppose the seller were to voluntarily post a price below market clearing (in hopes of generating an auction for priority access). He would face the constant temptation to raise it when buyers with unmet demand offer higher prices; anticipating this, buyers would prefer to abstain from the auction and offer a higher price.

Throughout, we have treated S and p as exogenous, but we can ask what would be the most beneficial values (to the seller) for these parameters. Note that auction revenue is maximized when $S = (n-1)d$; that is, one should include all bidders in the auction, but supply enough to serve all but one of them. Moreover, if p is equal to marginal cost, total welfare (per customer served) is maximized; and since all of this surplus is captured by the seller, this is also profit maximizing.

At these values, the auction gives nearly the same result as the Oi two-part tariff (Oi, 1971) with two exceptions: first, one client is left unfulfilled here, where none would be under the Oi Tariff (which seeks full efficiency by excluding no one). Second, the entry tariff in our model is not a fee explicitly set by the monopolist, but rather endogenously through the competition among bidders.

4.2 2 Heterogeneous Bidders

The auction for priority access becomes even more interesting when bidders differ in their utility from and/or demand for the good. We begin by providing a complete analysis of two-person auctions. Our solution can be considered a minor generalization of a two-bidder all-pay auction of a single indivisible good, with the alteration that losing need not have zero value. Again, the solution is not novel (it can be seen in Baye, *et al*, 2009), but its application to our environment provides further insight in the interpretation of efficiency and revenue results which are present (though less obvious) in the more general case.

The unique mixed strategy equilibrium can be characterized in terms of each agent's value of being first or second. We use the following notation to represent these:

$$w_i \equiv v_i(q_i) - pq_i \text{ where } q_i = \min\{d_i(p), S\}$$
$$x_{ij} \equiv v_i(q_{ij}) - pq_{ij} \text{ where } q_{ij} = \min\{d_i(p), \max\{0, S - d_j(p)\}\}$$

Thus, w_i represents i's utility from realized consumption when he has first priority, and x_{ij} when he has second priority, having been beaten by bidder j. While the identity of j is obvious for the two-person case, we introduce this notation in preparation for the three-person example. Note that $w_i \geq x_{ij}$. On the other hand, comparisons between w_i and either w_j or x_{ji} could go either direction. Without loss of generality, we will assume that $w_1 - x_{12} \geq w_2 - x_{21}$; if this were not the case, we could just relabel the two agents. The difference $w_i - x_{ij}$ indicates the net gain to i from consuming first rather than second.

Proposition 4. *The unique equilibrium is* $F_1(b) = \frac{b}{w_2 - x_{21}}$ *and* $F_2(b) = \frac{b + w_1 - x_{12} - w_2 + x_{21}}{w_1 - x_{12}}$, *resulting in* $EU_1 = w_1 - w_2 + x_{21}$ *and* $EU_2 = x_{21}$.

Note that both distributions share the same support $[0, w_2 - x_{21}]$, and agent 2 has an atom at 0. In effect, the agent who has more to lose by being second (bidder 1) bids more aggressively: his mixed strategy first-order stochastically dominates the other bidder's. Indeed, he is more likely to win the first priority: his bid is higher than agent 2's with probability $1 - \frac{(w_2 - x_{21})}{2(w_1 - x_{12})}$, and this becomes more likely the larger the difference in net gains.

This has the flavor of an efficiency result: the total utility generated from consumption is larger when agent 1 consumes first than when agent 2 does (*i.e.* $w_1 + x_{21} \geq w_2 + x_{12}$), and that outcome is more likely the bigger this difference. However, the less efficient outcome always occurs with positive probability. Moreover, as with homogenous bidders, the outcome is always inefficient on the intensive margin, since marginal utilities will not be equal even when 1 does consume first.

It is interesting to note that agent 1 may not fare better than agent two in equilibrium; when $w_1 < w_2$, agent 2 receives higher expected utility. For instance, this could occur if agent 2 gets a lot of utility from the good whether buying first or second, but agent 1 gets very little utility from buying second, perhaps because agent 2 exhausts the available supply.

As with homogeneous bidders, the auction implements a form of second degree price discrimination, with two twists. First, note that rents are not fully extracted (even if the losing bidder is unable to purchase anything, $x_{ij} = 0$). This is because the bidders only compete away the *smaller* of the two net gains. This can be seen in the expected utility of the two bidders: $w_i - (w_2 - x_{21})$. This similarly occurs in standard two-part tariffs with heterogeneous consumers.

Second, as noted above, the ex-post total surplus depends on who is served first, creating inefficiency whenever bidder 2 wins. Thus, rather than merely transferring surplus, the auction introduces some deadweight loss. The expected bid expenditures from the two agents will be $R = \frac{(w_2 - x_{21})}{2} + \frac{(w_2 - x_{21})^2}{2(w_1 - x_{12})}$. This revenue always increases with $w_2 - x_{21}$, and decreases with $w_1 - x_{12}$ up until the two net gains are equal. The latter fact is somewhat surprising, but as bidder 1 has more to gain from winning and thus has a larger expected bid, his opponent is discouraged from bidding as much, which reduces total revenue.

4.3 3 Heterogeneous Bidders

With more than two bidders, heterogeneity among bidders introduces another complication: bidders are concerned not only about how many people place higher bids, but also which bidders place those bids. After all, the residual supply for second place may depend on who is in first place. While enriching behavior, this added feature also hampers any attempt to characterize all equilibria. Even with just three bidders, the precise solution depends on a comparison of the various bidder payoffs, resulting in many cases.[11] Moreover, for some of these, the solution can only be numerically solved.

For brevity, we select a three-bidder example that provides a closed-form solution and yet is representative of the rich behavior possible with three bidders. We employ the same notation as the preceding two-person example. Assume that there are two types of bidders, with one of type 1 and two of type 2. The type 1 bidder will exhaust the entire supply, even if he places first. A type 2 bidder does not exhaust the supply if she places first, but will if she places second (beaten by the other type 2 agent). Thus, $x_{21} = 0$, as is the utility to any agent from placing third.

[11] This also occurred in Klose and Kovenock (2011), limiting them to present several representative examples. Their three-bidder environment does not directly map into ours, as they assume payoffs to the second- and third-place bidder are only affected by the identity of the first-place bidder. This cannot address the example we present here, in which third place is strictly less desirable than second place because the second-place bidder will exhaust the supply.

Equilibrium bid strategies will depend on parameters, and need not be unique.[12] For instance, if $w_2 \geq w_1 + x_{22}$, an equilibrium exists such that bidder 1 completely abstains, while both type 2 bidders behave as they would in the homogenous two-person auction. Similarly, if $w_1 \geq w_2 + x_{12}$, an equilibrium exists where one of the type 2 bidder abstains (*i.e.* $F_2(b) = 1$ for all $b \geq 0$), while the remaining bidders behave exactly as in the heterogneous two-person example.

Note that this latter solution is an asymmetric equilibrium, in that the identical type 2 bidders employ different strategies. These asymmetries are often part of an equilibrium with more than 2 heterogeneous bidders; indeed, a symmetric equilibrium may not exist for a particular set of parameters.

Rather than present all equilibria, we focus on a particularly tractable example in which all three participants share a common support of $[0, w_2]$, and both type 2 agents employ the same strategy. In Appendix A.8, we sketch the process by which one can solve for any of the equilibria (whether analytically or numerically).[13]

Proposition 5. *Assuming* $x_{21} = 0$*,* $w_1 \geq w_2$ *and* $2(w_1 - x_{12}) \geq w_2 - x_{22}$*, then*

$$F_1(b) = \frac{b(w_1 - 2x_{12})}{(w_2 - x_{22})\left(-x_{12} + \sqrt{x_{12}^2 + (w_1 - 2x_{12})(b + w_1 - w_2)}\right) + x_{22}(w_1 - 2x_{12})}$$

$$F_2(b) = \frac{-x_{12} + \sqrt{x_{12}^2 + (w_1 - 2x_{12})(b + w_1 - w_2)}}{w_1 - 2x_{12}},$$

for $b \in [0, w_2]$ *constitutes an equilibrium strategy profile, producing* $EU_1 = w_1 - w_2$ *and* $EU_2 = 0$*.*

One can compute a similar equilibrium without the imposed assumptions, but the expression becomes more cumbersome. The last assumption in particular is needed to make a symmetric equilibrium possible; without it, the type 2 bidders will have to employ distinct strategies (with one placing no weight on an interval $(0, z)$, where $z \leq w_2$).

[12] The multiple equilibria may not be revenue equivalent, either. Baye, *et al* (1996) demonstrate this for an auction of a single indivisible prize among heterogeneous bidders.

[13] Other equilibria are more complicated to express and do not add much insight. They tend to be an intermediate case between the equilibria presented, with some agent abstaining for some portion of the aggregate support. In numerous numerical computations, the revenue and efficiency properties of these other equilibria also lay between the full abstention equilibrium and the full participation equilibrium presented.

Note that the type 2 agents have an atom at $b = 0$ as long as $w_1 > w_2$. Also, one might expect bidder 1 to bid more aggressively ($F_1(b) < F_2(b)$ for all b), but this is not always the case. If $2(w_1 - x_{12}) < 2w_2 - x_{22}$, bidder 1 places less weight on high bids than bidder 2 does. In essence, x_{12} is large enough that bidder 1 is willing to accept second place and compete less aggressively for first.

In light of this, one might reasonably ask: could the seller increase revenues by excluding one of the bidders? In an all-pay auction of a single, indivisible prize, Baye, *et al* (1993) answered yes. For instance, in a three person auction, if two bidders have equal valuations while the third's is strictly higher, the latter should always be excluded. They label this effect the *exclusion principle.*

In our divisible-good environment, however, the exclusion principle must be altered. When x_{22} is near zero, the seller increases revenue by excluding bidder 1, who values first place and possibly even second place more than the remaining two bidders. However, for moderate values of x_{22} and x_{12}, excluding one of the type 2 bidders raises more revenue. This can occur even though $w_2 < w_1$ and $x_{22} < x_{12}$; that is, bidder 1 clearly values the prizes more, yet another bidder is excluded. Moreover, for larger values of either x_{22} or x_{12}, more revenue is raised by excluding none of three bidders.

Intuitively, the seller's choice to exclude is a question of whether he is deriving more revenue from competition for first place or competition for second place. When the value of second place is low, the original exclusion principle applies because bidders are almost exclusively motivated by their desire for first place. When the value of second place is high, it is as important to competition as first place; indeed, by excluding any bidder, the remaining two would bid less aggressively since the loser still receives a significant prize. This echos the result for homogenous bidders, where the seller maximizes revenue by ensuring that exactly one bidder will obtain nothing.

The intermediate case is also interesting. Here, the type 2 bidder is significantly more aggressive against a lone opponent of type 1 than she would be against a lone opponent of type 2 or against both type 1 and type 2. This is because type 2 must place first to win anything, since bidder 1 absorbs the entire supply. If the other type 2 bidder is included, the additional competition for first place is mitigated by the fact that second place will sometimes yield positive utility.

Expected revenue in this three-person example is a lengthy expression and is not useful for analytical comparisons; thus we illustrate these results in some representa-

Figure 2: Exclusion in Auctions for Priority Access

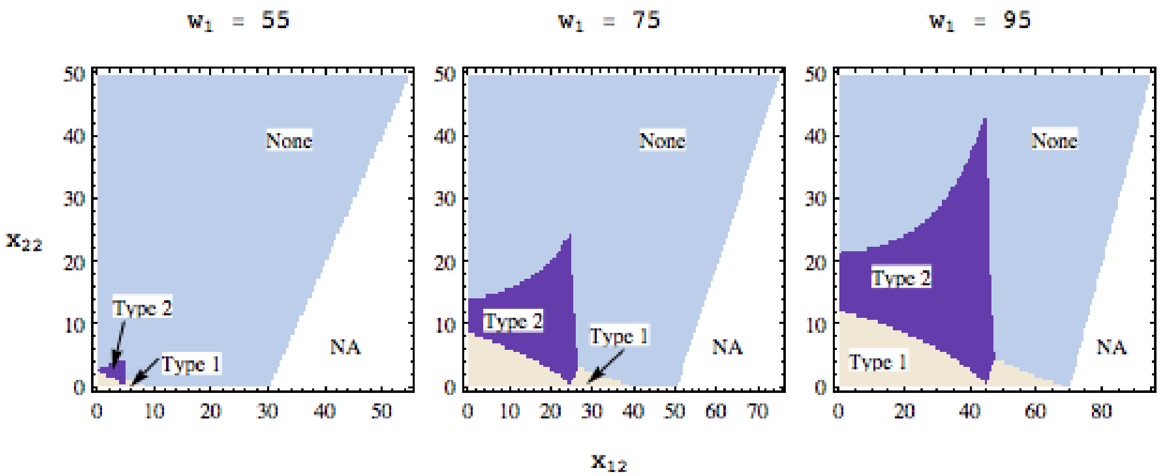

Notes: The shaded regions indicate parameter values for which auction revenue is highest by excluding the *Type 1* bidder, one of the *Type 2* bidders, or excluding *None*. The white *NA* region indicates parameters for which the example equilibrium does not exist. In all cases, we set $w_2 = 50$.

tive numerical examples. Figure 2 illustrates several cases. Note that as w_1 increases relative to w_2, the region of parameter space in which exclusion increases revenue is larger.

This three-bidder equilibrium also provides a nice setting to illustrate a key concern in auctions for priority access: a type 2 bidder in second place is greatly affected by whether he was beaten by the type 1 bidder (therefore getting 0) or the other type 2 (thus getting x_{22}). Compare this equilibrium to the symmetric equilibrium in the environment where all three bidders are of type 2. There, they would employ a mixed strategy (depicted in Figure 1, Panel 2) with the same support, but no atom:

$$\tilde{F}_2(b) = \frac{-x_{22} + \sqrt{x_{22}^2 + b(w_2 - 2x_{22})}}{w_2 - 2x_{22}}.$$

Comparing this equilibrium to the preceding, type 2 bidders place are more reluctant to bid when faced with the prospect of being shut out by the type 1 bidder. Indeed, when $w_1 - x_{12} \geq w_2 - x_{22}$, the strategy versus a type 1 and a type 2 is first-order stochastically dominated by the strategy versus two type 2 bidders: $F_2(b) > \tilde{F}_2(b)$ for all b. Even so, the solution with heterogeneous bidders still implements a form of second degree price discrimination, though the heterogeneity somewhat dampens its effectiveness.

5 Comparisons to Lotteries for Priority Access

We now contrast the auction for priority access to the proportional contest, which is frequently used in models of rent seeking. In this environment, each bidder chooses a bid b_i, and wins the contest with probability $\frac{b_i}{\sum_j b_j}$. In this sense, the bid is like a purchase of lottery or raffle tickets, where the participant can improve his odds of winning by buying more tickets.

Of course, in our environment, there is not a single prize to win. However, this same mechanism can be used to establish an order of priority.[14] After each agent chooses a bid, the first access is awarded to i with probability $\frac{b_i}{\sum_j b_j}$. If any supply remains, the second priority is awarded to k with probability $\frac{b_k}{\sum_{j \neq i} b_j}$, assuming i won the first round. This continues until either the supply or the bidders are exhausted. As before, bids are sunk; an unlucky participant pays his bid even if he is unable to purchase any of the good.

The key difference is that priority order is no longer a deterministic function of bids, as it is in the auction. In the lottery, a larger bid produces a proportionately larger probability of getting first access, but does not guarantee it. Another difference is that existence of equilibrium is not in question, as payoffs are continuous. In fact, lotteries often result in a unique pure strategy equilibrium.

We present solutions for the same cases analyzed in Section 4, and compare the results to their auction counterparts. Fang (2002) performed a similar comparison in the context of a single indivisible prize auctioned among heterogeneous bidders; this is most closely related to our two-person example, except that the value of second place can be positive in our model.

5.1 Lottery with n Homogeneous Bidders

We focus on symmetric strategies; others can sometimes exist but produce dramatically less revenue. Let b_i denote i's bid, while b_j denote the bid used by all other

[14] Clark and Riis (1996) introduced this process of using one bid in a Tullock success function to allocate k identical, indivisible prizes among n bidders with unit demand.

agents. The expected utility of agent i is:

$$\begin{aligned} EU_i(b_i, b_j) &= w\left(\sum_{s=0}^{k-1} \frac{b_i}{b_i + (n-s-1)b_j} \prod_{t=1}^{s}\left(1 - \frac{b_i}{b_i + (n-t)b_j}\right)\right) \\ &\quad + x\frac{b_i}{b_i + (n-k-1)b_j} \prod_{t=1}^{k}\left(1 - \frac{b_i}{b_i + (n-t)b_j}\right) - b_i. \end{aligned}$$

The first summation represents the k opportunities that agent i has to gain access while there is still enough supply to satisfy his demand. Note that, conditional on having lost on the first $s \leq k$ draws, the probability of winning on the $s + 1^{\text{th}}$ draw, $\frac{b_i}{b_i+(n-s-1)b_j}$, is increasing in s. This is because those who have already been given access are removed from the subsequent draws. The second term indicates the probability of being selected for exactly the $k + 1^{\text{th}}$ priority.

Proposition 6. *The symmetric equilibrium of a lottery for priority access with n homogeneous bidders is:*

$$b_i^* = \frac{x - w + ((n-k)w - x)\sum_{t=0}^{k} \frac{1}{n-t}}{n}$$

for all i, with $EU_i^* = \frac{kw+x}{n} - b_i^*$ and bid revenue nb_i^*. This always produces more expected utility and less expected revenue than the auction for priority access.

The lottery is not as effective as the auction at extracting rent when bidders are identical. The same comparison of auction versus lottery holds in Fang (2002) for a single indivisible prize. Intuitively, in the lottery, the awarding of priority order is not deterministic based on the bids; this element of randomness decreases competition among participants.

On the other hand, note that since pure strategies are employed, lottery bid revenue is deterministic. Thus, a seller might favor the lottery over the auction if he is sufficiently risk averse.

5.2 Lottery with 2 Heterogeneous Bidders

Again, we consider a two-person environment using the same notation as in Section 4.2. As before, let $w_1 - x_{12} > w_2 - x_{21}$. Bidder i has expected utility:

$$EU_i(b_i, b_j) = w_i \frac{b_i}{b_i + b_j} + x_{ij} \frac{b_j}{b_i + b_j} - b_i.$$

This game has the following unique Nash equilibrium; equilibrium expected utility is reported in the proof.

Proposition 7. *The unique equilibrium of a lottery for priority access with 2 bidders is:*

$$b_1^* = \frac{(w_1 - x_{12})^2 (w_2 - x_{21})}{(w_1 - x_{12} + w_2 - x_{21})^2} \quad \text{and} \quad b_2^* = \frac{(w_1 - x_{12})(w_2 - x_{21})^2}{(w_1 - x_{12} + w_2 - x_{21})^2},$$

with bid revenue $R = \frac{(w_1 - x_{12})(w_2 - x_{21})}{w_1 - x_{12} + w_2 - x_{21}}$. *This is always less efficient than the auction, and produces more revenue if and only if* $(\sqrt{2} - 1)(w_1 - x_{12}) > w_2 - x_{21}$.

Regarding efficiency, the probability that agent 1 receives first priority after the lottery is $\frac{w_1 - x_{12}}{w_1 - x_{12} + w_2 - x_{21}}$; thus, a larger difference in the net gain of the two bidders will create a higher probability of the more efficient ordering. However, an auction for priority access will achieve that outcome with even higher probability:

$$1 - \frac{(w_2 - x_{21})}{2\,(w_1 - x_{12})} > \frac{w_1 - x_{12}}{w_1 - x_{12} + w_2 - x_{21}} \qquad \Longleftrightarrow \qquad w_1 - x_{12} > w_2 - x_{21},$$

which holds by assumption. This demonstrates that even though mixed strategies are employed in the auction, the atom in agent 2's bid skews the outcome so that agent 1 wins more frequently than in the lottery.

A comparison of bid revenues shows that the lottery generates more than the auction if and only if $(\sqrt{2} - 1)(w_1 - x_{12}) > w_2 - x_{21}$. In words, if bidder 2's net gain is significantly smaller (*i.e.* at least 58% smaller) than bidder 1's net gain, the lottery will incite greater competition than the auction. Indeed, while the auction's revenue strictly decreases as $w_1 - x_{12}$ increases, the lottery's revenue strictly increases. In both mechanisms, bidder 2 backs off as $w_1 - x_{12}$ becomes larger. In the auction, bidder 1 does not change his strategy as his net gain increases, though he does increases his bid in the lottery. Intuitively, the random draw in the lottery draw encourages greater effort from bidder 1 to secure first priority. In a single prize contest (*i.e.* when $x_{12} = x_{21} = 0$), Fang (2002) reaches the same conclusion.

Finally, a comparison of expected utility reveals that bidder 2 fares strictly better under the lottery, rather than the auction. The comparison is ambiguous for bidder 1, who prefers the lottery whenever $(w_1 - x_{12})^3 > (w_1 - w_2)(w_1 - x_{12} + w_2 - x_{21})^2$.

In particular, if $w_1 < w_2$, this will always hold.

5.3 Lottery with 3 Heterogeneous Bidders

We now turn to the same three-person example used in Section 4.3. Solving for the lottery proceeds as before; however, the equilibrium solution must obtained numerically. We focus on the symmetric equilibrium, in which the type 2 bidders employ the same strategy.[15] Summarizing many numerical computations, our goal is to comment on three issues: exclusion, revenue, and efficiency.

First, excluding bidders from the lottery almost never increases revenue. In the auction mechanism, exclusion can sometimes help when aggressive bidding by one agent would cause the others to back off; thus, excluding that aggressive bidder can increase competition from the other two. In a lottery, however, bidders are less sensitive to each others' bids, since a low bid can still win over a high bid with some probability. The only exception occurs when w_2 is roughly two orders of magnitude smaller than w_1. Fang (2002) finds, in a single-prize setting with heterogeneous bidders, that exclusion can never increase revenue.

Second, the comparison of lottery and auction revenue bears a strong resemblance to the two-bidder example. The auction (using exclusion where beneficial) earns strictly more revenue than the lottery when w_1 is close to w_2. As these grow farther apart, a region of (x_{12}, x_{22}) parameters in which the lottery will dominate emerges and grows, which is illustrated in Figure 3. Note that this region involves high x_{22} and low x_{12} values. As before, the auction performs best when the agents are more similar in their valuations; remarkably, this is true even when the value of second priority is close, even if the value of first priority is not.

Finally, efficiency in this environment clearly depends on the four parameters, but we compute the expected total welfare (expected utility plus expected revenue) for both the lottery and the auction (employing the revenue-maximizing exclusion where useful). The results are depicted in Figure 4. The auction produces higher surplus in the intermediate range of x_{22}, while the lottery can be more efficient for very low or very high values of x_{22}. For the low values, this essentially coincides with where the auction would exclude bidder 1; this ensures an inefficient outcome since $w_1 > w_2 + x_{22}$ and/or $w_2 + x_{12} > w_2 + x_{22}$. The lottery performs better because it

[15] There are typically two other solutions to the system of first-order conditions, but the asymmetric solutions produced lower revenue for all parameters we attempted.

Figure 3: Revenue Comparison: Three Bidder Example

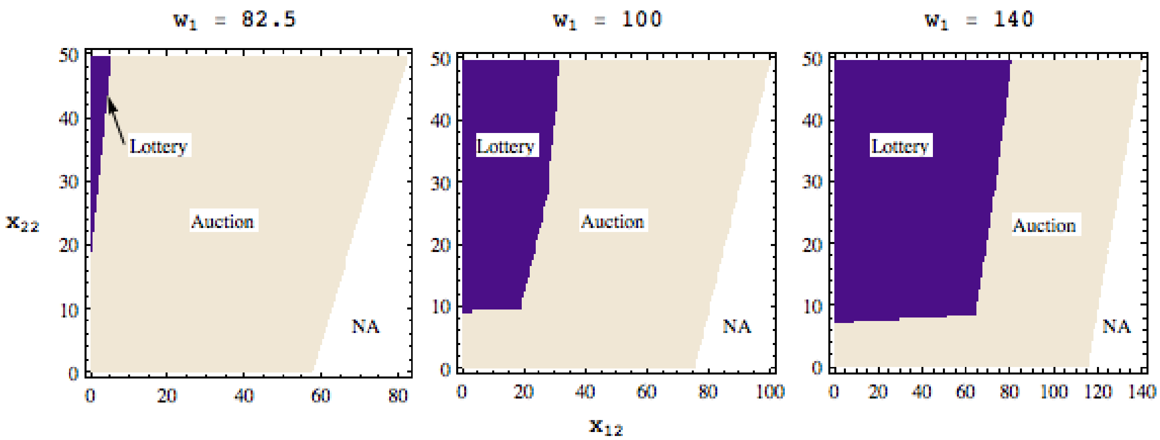

Notes: The shaded regions indicate parameter values for which the *Auction* or *Lottery* earn greater expected revenue. For $w_1 < 80$, the auction always dominates. The white *NA* region indicates parameters for which the example equilibrium does not exist. In all cases, we set $w_2 = 50$.

includes bidder 1. If no exclusion occurred, the auction would be more efficient than the lottery in this low range.

For the higher region, the efficiency gain of the lottery over the auction is minor — less than one percent, compared to 5 to 10% gains in the moderate region or even 50% gains in the low region. This is not surprising since all three outcomes (w_1, $w_2 + x_{12}$, or $w_2 + x_{22}$) are nearly the same in this high region.

6 Conclusion

Contests for priority access provide a method of allocating a good in short supply among potential buyers, using the bids to establish an order of service. Since there is still a marginal cost of procurement even after access is awarded, the first customer may not exhaust supply. We have shown that these contests can be effective methods for the seller to extract consumer surplus from his patrons, implementing a two-part tariff without the seller needing to know a consumer's surplus. For instance, if the landlord of a rent-controlled property can establish competition among prospective tenants via gifts or key money, he could appropriate much of the surplus in filling a vacancy, perhaps with more profit than in an unregulated housing market.

Figure 4: Efficiency Comparison: Three Bidder Example

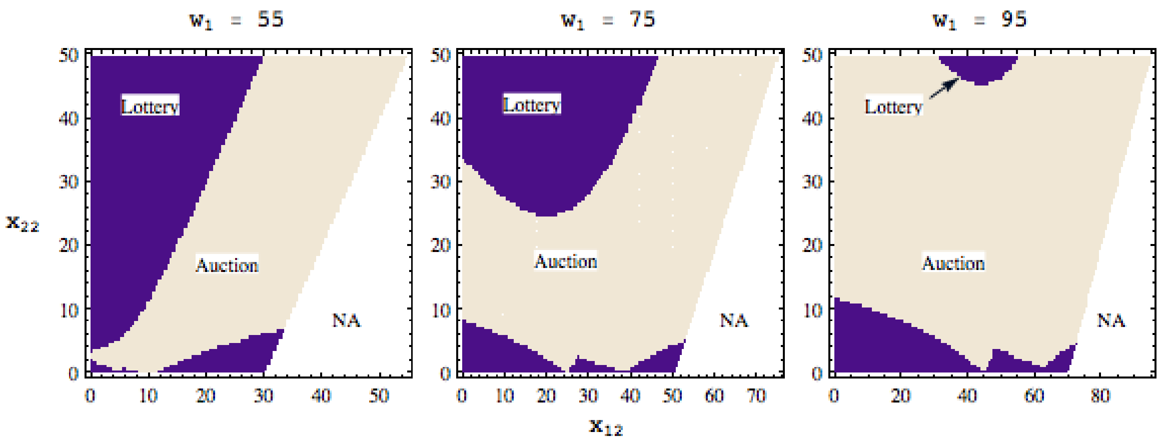

Notes: The shaded regions indicate parameter values for which *Auction* or *Lottery* generate higher total welfare. The white *NA* region indicates parameters for which the example equilibrium does not exist. In all cases, we set $w_2 = 50$.

In all-pay auctions for a single indivisible good, the seller can increase revenue by excluding a bidder who has the strictly highest valuation for the good. In an auction for priority access (with a perfectly divisible good), this exclusion principle must be altered. In our three-person example, the seller only excludes a bidder if placing second is relatively unimportant to bidders' strategies. Moreover, there are situations in which excluding one of the bidders with *lower* valuations generates greater revenue. This occurs because these bidders will compete more aggressively when their only chance of obtaining any good is to place first (when only facing the high valuation bidder), compared to having a chance for leftovers in second place (when facing both a high and a low valuation bidder).

From an efficiency standpoint, contests for priority access are doomed from the start in one sense. Some agents will be left without any of the good, even though their marginal utility from consumption is well above those who did procure a portion. Of course, this is where competitive markets excel, since prices align marginal incentives. These incentives do not arise in our auction or lottery because bids are a fixed cost, reflecting total rather than marginal surplus.

For this very reason, the auction for priority access still falls short of perfect price discrimination. For instance, in the best-case scenario, where bidders are homogeneous, the seller fully extracts the consumer surplus of the bidders who are served.

However, the customers who obtained nothing have a greater marginal value for the good than those who were served; thus total surplus would be higher if all customers received an equal portion of the good. Thus, the auction may not be the optimal mechanism for maximizing revenue in this environment, if some other mechanism can preserve incentives on the intensive margin *and* still capture the full consumer surplus.

Even on the extensive margin, the auction may occasionally award first priority to a customer with lower total surplus, due to the use of mixed strategies. However, as the valuations become more distinct (and hence, the potential loss becomes greater), the auction makes these errors less frequently. Moreover, the auction typically performs better on this dimension than the lottery for priority access; the biggest exception occurs when the seller excludes a bidder with the highest valuation.

Our model could easily be reinterpreted for other environments. For example, consider rent seeking. A politician has some supply of political capital (*i.e.* the ability to influence policy) that she can distribute. A number of special interest groups may want this capital used on their behalf, but the politician must somehow prioritize the issues she will address. Lobbyists can compete for access through their campaign contributions, with the largest contributor receiving first priority. The winning lobbyist, however, rarely depletes the entire supply.[16] Thus, we frequently see a politician catering to several special interest groups that were prominent in her election. Note also that the competing agents pay their bids regardless of the auction outcome (the unsuccessful lobbyist is not refunded his campaign donations). The access payment seems to be a particularly fitting allocation mechanism here, since directly selling the political capital on a per-unit basis would undoubtedly be seen as quid-pro-quo bribery, while an ex-ante donation is less suspect.

Another example is the grant application process. The organization offering the grant has some pool of money available. The applicant's bid for this grant money can be seen as the effort in drafting a proposal, with the best proposals (those with the highest bid) given first priority on the pool of monies. Even so, the best grant seldom

[16]For instance, the lobbyist may have to expend effort justifying the political capital to be used. Cotton (2010) presents a model in which lobbyists compete in an all-pay auction for an audience with the politician, but on winning, must present verifiable evidence supporting the policy they seek. In that sense, the politician could have identical preferences to the special interest group, with both agreeing when the marginal benefit of expending more political capital on that issue is outweighed by its opportunity cost.

absorbs the entire pool. This is because there is some marginal cost in justifying a larger budget for the proposal.

A similar competition describes consumer's queuing decisions for a good in short supply (such as with gasoline shortages or holiday door-buster sales). If the retailer opens at a particular time, each consumer bids by choosing how early to arrive and stand in line.[17] The earliest arrival will be the first to make purchases, but does not typically purchase the whole supply; perhaps the buyer has capacity or credit constraints preventing this.

In the context in which bids are collected by the seller, such as political rent seeking or key deposits to circumvent rent control, the auction allows the seller to capture rents. On the other hand, when applied to grant writing or queuing, the bids represent wasted resources (from a social standpoint). Indeed, applying our same results regarding expected bid revenue, the maximum amount of waste occurs when all but one customer will be able to buy what he wants. Indeed, society can benefit by having *fewer* units available so as to reduce the incentive to queue. Alternatively, one can eliminate the wasted time by allowing the price to adjust to its market-clearing level; indeed, consumers could find themselves spending less (when the saved time is included) at the market price than under the lower price ceiling, as demonstrated in Proposition 3.

If large amounts of resources are wasted via queuing, sellers would have strong incentives to introduce alternatives that capture the otherwise wasted rents. This can be seen in the debate over net neutrality. Bandwidth is a finite but divisible good, and in times of peak demand, it is typically allocated proportionally among users; for customers accessing large media files, the resulting wait could be sufficiently frustrating to make them abandon their downloads.

As an alternative, some service providers have sought the ability to prioritize certain content providers or customers (for a fee) as a means of allocating the scarce bandwidth. This tiered internet system could be seen as an auction for priority access, which we predict would capture much of consumer surplus. It is unclear whether this

[17] This was considered in Holt and Sherman (1982) and Taylor, *et al* (2003), with the assumption of unit demand and independent private valuations, which yields a pure strategy Bayesian Nash equilibrium. When interpreted as a model of strategic queuing, our paper yields results analogous to the competitive outcome in Platt (2009). In that environment, consumers take queue times as given and decide whether to wait in the queue to gain access to the market. The equilibrium queue time adjusts to make identical agents indifferent about participation. The aggregate queue time incurred in that setting would equal the expected bid revenue here.

would encourage expansion of bandwidth; a monopolist would want some regularly-occurring congestion to encourage competition for access, but as seen in our model, it only requires a few unsatisfied customers to maintain maximum expected revenue.

Of course, uncertainty plays a key role in internet service (and other) markets, as customers must anticipate their and others' data needs. Indeed, the assumption that bidders know each others' valuations, though common in all-pay auctions, is perhaps the most restrictive of our model. We anticipate that a parallel research agenda in a private valuation framework would share many features of our model, but we leave this study for future work.

A Proofs

A.1 Proof of Proposition 1

Proof. Existence is proven by applying Corollary 5.2 of Reny (1999). Three conditions are required to apply this result; the first two are clearly satisfied.

- The space of pure strategies $[0, M]^n$ must be compact and Hausdorff.
- The payoff function in pure strategies $v_i(q_i(b_{-i}, b_i)) - pq_i(b_{-i}, b_i) - b_i$ must be bounded and measurable on $[0, M]^n$.
- The mixed strategy game is *better-reply secure*. Formally, let $\bar{\mu}$ be any non-equilibrium mixed strategy profile. Let $\bar{u}$ be a profile of expected payoffs such that for some sequence $\mu^k \to \bar{\mu}$, $u(\mu_k) \to \bar{u}$. The game is better reply secure if there exists some player i and strategy $\hat{\mu}_i$ such that $u_i(\mu_{-i}, \hat{\mu}_i) > \bar{u}_i$ for all μ_{-i} within some open neighborhood of $\bar{\mu}_{-i}$.

The third condition is easily satisfied if $\bar{\mu}$ has no consequential ties occurring with strictly positive probability. If so, there is no discontinuity at $u(\bar{\mu})$. Since $\bar{\mu}$ is not an equilibrium profile, some agent i has a best response to $\bar{\mu}_{-i}$ that strictly increases his utility above $u_i(\bar{\mu})$. By keeping the neighborhood of $\bar{\mu}_{-i}$ sufficiently small, i's utility remains bounded above $u_i(\bar{\mu})$: even if nearby μ_{-i} introduce a consequential tie (and hence a discontinuity), the probability of that tie occurring is limited as small as needed. Hence, any drop in utility at such discontinuities can be kept arbitrarily small, and changes elsewhere are continuous.

Suppose instead that a consequential tie occurs with strictly positive probability under $\bar{\mu}$. In this case, it is not necessarily the case that $\bar{u} = u(\bar{\mu})$. Since $\bar{\mu}$ is not an equilibrium profile, there still exists some i with a best response to $\bar{\mu}_{-i}$ that strictly increases his utility above $u_i(\bar{\mu})$. If he is not involved in any of the consequential ties, the analysis from before still applies.

If the only agents not playing a best reply are involved in a consequential tie, the danger is that $\bar{u}$ may treat some of them as if they always win the tie — and there may not be a $\hat{\mu}_i$ for such a person that strictly provides more utility. However, since this is a consequential tie, at least one of the agents involved has positive probability of not receiving his full demand. That agent can strictly improve on $\bar{\mu}_i$ by shifting the strictly positive atom he placed on the consequential tie(s) to ϵ higher. There can only be a countable number of consequential ties, so he can do this without entering another tie. By doing so, he strictly wins each tie and obtains his full demand, but with an insignificant increase in his expected bid, thus obtaining a strictly higher payoff that $\bar{u}_i$.

With i playing the $\hat{\mu}_i$ so constructed, we can contain μ_{-i} to a small enough neighborhood so that i still receives more than $\bar{u}_i$. Even if they move some positive probability to the atoms of $\hat{\mu}_i$, we can restrict it to be small enough that it only slightly decreases the expected utility of i.

□

A.2 Proof of Proposition 2

Proof of claim 2.1. Assume there is some interval $(a_1, a_2) \subset [0, M]$ such that for all $b \in (a_1, a_2)$, $b \notin B_j^*$ for all $j \neq i$. Suppose agent i considers choosing some $b_i \in (a_1, a_2)$. If instead he chooses $b_i' = b_i - \epsilon \in (a_1, a_2)$, then for any $b_{-i} \in B_{-i}^*$, $q_i(b_i, b_{-i}) = q_i(b_i', b_{-i})$. In words, whichever bid profile is selected by the other players, both b_i and b_i' will have the same rank compared to the other bids, and thus obtain the same quantity for agent i.

But then $EU_i(b_i', \mu_{-i}^*) = EU_i(b_i, \mu_{-i}^*) + \epsilon$, since both integrate over the same outcomes, but b_i' does so with a smaller bid. Thus, no strategy $b_i \in (a_1, a_2)$ can be utility maximizing; hence $b_i \notin B_i^*$. □

Proof of claim 2.2. Suppose there exists a bid $a \in (0, M]$ such that $\mu_i^*(\{a\}) > 0$ for some agent i. Note that there can only be a finite number of atoms among the n agents; thus, there is some range $(a - \epsilon, a + \epsilon)$ in which there is no other atom, though other agents might have an atom at a.

Let $\tilde{q}_i(b)$ and $\underset{\sim}{q}_i(b)$ denote the most and least agent i could receive under bidding profile b. These only differ when a consequential tie occurs — if so, there is enough remaining supply to satisfy some but not all of those who bid the same as i, and the randomly selected permutation π determines who is served first.

Denote the set of bids where a is consequential to i as:

$$C_i(a) \equiv \left\{ b_{-i} \in B_{-i}^* \ : \ b_i = a \text{ and } \tilde{q}_i(b) > \underset{\sim}{q}_i(b) \right\} \tag{3}$$

Suppose $\mu_{-i}^*(C_i(a)) > 0$. Then agent i can increase his expected utility by shifting his atom $\mu_i^*(a)$ from a to $a+\epsilon$. By doing so, he ensures that he will always obtain $\tilde{q}_i(b)$ for all $b \in C_i(a)$, rather than sometimes obtaining $\underset{\sim}{q}_i(b)$. Moreover, for all $b \notin C_i(a)$, his quantity obtained will weakly increase. Thus, for ϵ sufficiently small, the strict increase in utility from added consumption will outweigh the small increase in bid.

If instead $\mu^*(C_i(a)) = 0$, then when i bids a, two possibilities exist. First, bidding $a-\epsilon$ could give the same outcomes with equal probability. This is to say, for any quantity q,

$$\mu^*_{-i}(\{b_{-i} \in B^*_{-i} \ : \ q_i(a, b_{-i}) = q\}) = \mu^*_{-i}(\{b_{-i} \in B^*_{-i} \ : \ q_i(a-\epsilon, b_{-i}) = q\}).$$

This would happen if no other agents have $(a-\epsilon, a) \subset B_j^*$, or if those who do have no effect on the remaining supply available for i (because the supply is already exhausted, or is so plentiful that both i and j can be satisfied). In such a case, $EU_i(a, \mu^*_{-i}) = EU_i(a-\epsilon, \mu^*_{-i}) - \epsilon$; i can reduce his bid cost while maintaining his average benefit from his opportunities to purchase.

Alternatively, if bidding $a-\delta$ for any $\delta \in (0, \epsilon)$ would have some impact on agent i's outcomes, then there must exist some agent j with $(a-\epsilon, a) \subset B_j^*$. If so, agent k's outcomes are *also* affected by whether agent i is allowed to purchase before him or not. Also, because they are part of an equilibrium best response, any $b_k \in (a-\epsilon, a)$ produces the same expected utility $EU_k(b_k, \mu^*_{-k})$. However, by bidding $a+\delta$ for some arbitrarily small $\delta > 0$, agent k can strictly increase his utility, because he has reduced by $\mu_i(\{a\})$ the probability of being outbid by agent i, which allows him the opportunity to purchase strictly more of the good, while incurring a slightly higher bid.

Thus, in all cases, we obtain a contradiction if agent i has an atom at a. Note that if an atom were at $a = 0$, the first case regarding consequential bids still apply ; however, if bids are almost always inconsequential bids, the latter arguments cannot be replicated since there is no bid below 0. □

Proof of claim 2.3. Suppose that there is some interval (a_1, a_2) such that for all $b \in (a_1, a_2)$, $b \notin B_j^*$ for all j. Pick an agent i such that $a_2 \in B_i^*$. By claim 2 of Proposition 2, there are no atoms at a_2. Thus, if agent i were to bid $b \in (a_1, a_2)$, he would achieve the same outcomes with the same probability as when bidding $b = a_2$, but with a lower bid. Hence this cannot be an equilibrium.

Thus, the aggregate support $\cup_i B_i^*$ must be connected, having ruled out any gaps.

Moreover, the same logic applies to an interval $(0, a_2)$, hence $0 \in \cup_i B_i^*$. □

Proof of claim 2.4. For any strategy μ_{-i}, $EU_i(0, \mu_{-i}) \geq v_i(0)$; the worst that can happen when $b_i = 0$ is that agent i never wins the opportunity to purchase any amount, but made no expenditure to get there. Thus, since $b_i = 0$ is always a feasible choice, then $EU_i(\mu_i^*, \mu_{-i}^*) \geq EU_i(0, \mu_{-i}^*) \geq v_i(0)$ in equilibrium. □

Proof of claim 2.5. Suppose $\mu_i^*(0) > 0$ for some i, $S \leq \sum_{j \neq i} d_j(p)$, and $EU_i(\mu^*) > v_i(0)$. This also means that $EU_i(0) > v_i(0)$; in other words, with positive probability, agent i is able to buy some positive amount of the good even when bidding $b_i = 0$.

If there were almost never consequential ties at $b = 0$, all agents $j \neq i$ would bid strictly more than 0 with probability 1. If so, $S \leq \sum_{j \neq i} d_j(p)$ implies that agent i would almost always receive 0 of the good, which contradicts.

So consequential ties must occur at $b = 0$ with positive probability; thus, some subset of the bidders also have an atom at 0. But then the same logic applies as in the proof of Claim 1.2: Any one of the agents who ties at 0 would strictly benefit by shifting their atom from 0 to ϵ. This ensures access to a strictly larger amount of the good and only requires an arbitrarily small increase in the bid. Thus μ^* would not be an equilibrium. □

Proof of claim 2.6. Suppose $\sum_{j \neq i} d_j(p) < S \leq \sum_j d_j(p)$. Again, we can rule out consequential ties at 0, because if they occurred with positive probability, any tied agent would have an incentive to raise his bid slightly.

Instead, suppose that there are almost never consequential ties at $b = 0$. Again, all other agents must bid strictly more than 0 with probability 1. That means when i bids 0, he receives $S - \sum_{j \neq i} d_j(p)$ almost surely. Hence, $EU_i(\mu^*) = u_i \left(S - \sum_{j \neq i} d_j(p), 0 \right)$. □

Proof of claim 2.7. Let $\hat{b}_i \equiv \max B_i^* = \max \cup_j B_j^*$. Bidding $\hat{b}_i$ will almost always result in i receiving first priority. Recall that there are no atoms if $\hat{b}_i > 0$, and if $\hat{b}_i = 0$ then there are no consequential ties and $S > \sum_j d_j(p)$. Thus $q_i(\hat{b}_i, b_{-i}) = \min\{d_i(p), S\}$ for almost all $b_{-i} \in B_{-i}^*$. Hence, i's expected utility must equal $EU_i(\hat{b}_i, \mu_{-i}^*) = u_i \left(\min\{d_i(p), S\}, \hat{b}_i \right)$. □

A.3 Proof of Proposition 3

Proof. Let each buyer have linear demand specified as: $p = c - a \cdot q_i$. At the controlled price p, individual demand is d; if the market clearing price p_{mkt} is charged, individual demand must be $\frac{S}{n}$. Hence,

$$p_{mkt} - p = c - a \cdot \frac{S}{n} - (c - a \cdot d) = a \cdot \left(d - \frac{S}{n}\right).$$

On the other hand, the total benefit w from purchasing d units at price p is defined from the individual's consumer surplus:

$$w = \frac{d \cdot (c - (c - a \cdot d))}{2} = \frac{a \cdot d^2}{2}.$$

Thus, the seller earns more profit under the price control if and only if: $a \cdot \left(d - \frac{S}{n}\right) \cdot S < k \cdot \frac{a \cdot d^2}{2}$. Note that by definition, $S = k \cdot d$. Thus, the inequality simplifies to $k > \frac{n}{2}$. □

A.4 Proof of Proposition 4

Proof. From Proposition 2, we know that F_1 and F_2 must share the same support (Claim 1), which includes 0 (Claim 3). Since expected utility must be constant throughout that support, we can solve for these c.d.f.s from the following two equations:

$$EU_1 = w_1 F_2(b) + x_{12}(1 - F_2(b)) - b \quad \text{and} \quad EU_2 = w_2 F_1(b) + x_{21}(1 - F_1(b)) - b,$$

obtaining:

$$F_1(b) = \frac{b - x_{21} + EU_2}{w_2 - x_{21}} \quad \text{and} \quad F_2(b) = \frac{b - x_{12} + EU_1}{w_1 - x_{12}}.$$

Suppose F_2 has no atom at 0. To get $F_2(b) = 0$ requires $EU_1 = x_{12}$. Then $F_2(b) = 1$ at $b = w_1 - x_{12}$. Since these two distributions share the same support (and there are no atoms above 0, due to Claim 2), $F_1(w_1 - x_{12})$ must equal 1 as well. This allows us to solve for $EU_2 = w_2 - w_1 + x_{12}$. But then $F_1(0) = \frac{w_2 - w_1 + x_{12} - x_{21}}{w_2 - x_{21}}$, which is admissible for a c.d.f. only if $w_1 - x_{12} = w_2 - x_{21}$, in which case this is equivalent to the proposed equilibrium (with substitution). If however $w_1 - x_{12} > w_2 - x_{21}$, then $F_1(0) < 0$ which contradicts the properties of a c.d.f.

Thus, consider when F_2 has an atom at 0. Claim 6 requires $EU_2 = x_{21}$, so $F_1(b) = \frac{b}{w_2 - x_{21}}$. The maximum bid in the support is $b = w_2 - x_{21}$; this in turn pins down $EU_1 = w_1 - w_2 + x_{21}$ in order to ensure $F_2(w_2 - x_{21}) = 1$. Note that $F_2(0) \geq 0$ while $F_1(0) = 0$. Moreover, bidder 2 would never want to bid more than $b = w_2 - x_{21}$; doing so would ensure he places first but give him less than x_{21} utils after the bid is paid. Similarly, bidder 1 would get strictly less than EU_1 by bidding more than the equilibrium support.

□

A.5 Proof of Proposition 5

Proof. In this three bidder environment, the expected utility from bidding b is expressed as follows:

$$\begin{aligned} EU_1(b) &= \left(F_2(b)\right)^2 w_1 + 2F_2(b)\left(1 - F_2(b)\right) x_{12} - b \\ EU_2(b) &= F_1(b)F_2(b)w_2 + F_1(b)\left(1 - F_2(b)\right) x_{22} - b \end{aligned}$$

It is a minor algebraic exercise to demonstrate that, after substituting for the proposed strategies $F_1(b)$ and $F_2(b)$, these simplify to $EU_1 = w_1 - w_2$ and $EU_2 = 0$ for all $b \in [0, w_2]$. Moreover, for either type, bidding $b > w_2$ yields strictly lower expected utility.

For this to be an equilibrium, we must assume that $w_1 \geq w_2$; otherwise $EU_1 < 0$, violating Claim 4 of Proposition 2. We must also assume $2(w_1 - x_{12}) \geq w_2 - x_{22}$ to ensure that $F_1(b)$ is weakly increasing and thus a well-formed CDF. □

A.6 Proof of Proposition 6

Proof. The equilibrium bid is found by taking the first order condition, setting the derivative with respect to b_i equal to zero. Then symmetry is imposed, replacing b_j with b_i. The second derivative confirms that this is a necessary and sufficient condition.

Expected bid revenue is always higher under the auction format. The auction produces more revenue when

$$\frac{(k+1)w}{(n-k)w-x} > \sum_{t=0}^{k} \frac{1}{n-t}.$$

But note that:

$$\begin{aligned}\frac{(k+1)w}{(n-k)w-x} &> \frac{k+1}{n-k} = \frac{k}{n-k} + \frac{1}{n-k} > \frac{k}{n-k+1} + \frac{1}{n-k} \\ &> \frac{1}{n} + \frac{1}{n-1} + \cdots + \frac{1}{n-k+1} + \frac{1}{n-k} = \sum_{t=0}^{k} \frac{1}{n-t},\end{aligned}$$

so this always holds.

By the same token, expected utility must be higher in the lottery if revenue is lower, since total surplus is equal. Indeed, when $k \leq n-1$, the auction gives expected utility of 0. The lottery, on the other hand, produces strictly positive utility, as evidenced by the fact that $b_i = 0$ was available but not selected. □

A.7 Proof of Proposition 7

Proof. Again, the first order condition is necessary and sufficient for utility maximization. The equilibrium bids are found by jointly solving the two first order conditions. Substituting these into the objective functions gives the equilibrium expected utility of:

$$EU_1^* = x_{12} + \frac{(w_1 - x_{12})^3}{(w_1 - x_{12} + w_2 - x_{21})^2} \quad \text{and} \quad EU_2^* = x_{21} + \frac{(w_2 - x_{21})^3}{(w_1 - x_{12} + w_2 - x_{21})^2}.$$

The analysis of efficiency and revenue are included in the text. □

A.8 Algorithm for Finding Equilibrium Auction Strategies

Here we briefly sketch an algorithm for finding the equilibrium mixed strategies in the auction for priority access. The process initiates with a guess (albeit an informed one) of each bidder's support, which will take the form $B_i = 0 \cup [a_i, A_i]$, where a_i and A_i are both real numbers.

One useful observation is that if two bidders are of the same type, their mixed strategies must be identical over the intersection of their supports (besides at $b = 0$)

in order for both to have constant expected utility across that shared support. Thus, asymmetries can only occur with one bidder having a bigger atom at 0 and abstaining for some range until the other bidder's cumulative distribution $F_i(b)$ equals the size of that atom, beyond which both bidders' c.d.f. are identical.[18] Thus, individuals with identical utility must have the same A_i.

If demand will be exhausted before the last bidder consumes, then typically $A_i = \min\{w_i, \max_{j\neq i} w_j\}$. This is because bidders are never willing to bid more than their value of first place, but if their value happens to be higher than all others, they never bid more than the second highest value (from Claim 1 of Proposition 2). At the same time, bidders typically include this A_i in their support because of the potential of placing last. If $w_i > \max \cup_{j\neq i} B_j^*$, then bidder i can guarantee himself first priority and a strictly positive payoff by bidding just above $\cup_{j\neq i} B_j^*$, which would be better than occasionally getting last place and no consumption.

Once this guess is made, expected utility for each bidder is pinned down by Claims 5, 6, and/or 7. We then set that equal to $EU_i(b)$ as a function of the unknown cumulative density functions, $F_i(b)$, as in the proof of Proposition 5. This formulates a system of n equations (for each b) with n unknowns ($F_i(b)$, for each b), which can sometimes be solved analytically, or if necessary, numerically.

Once the solution is found, one must verify that it is a well-formed cumulative distribution function. That is, it is possible to get a solution that violates one of the following rules. If so, this is usually fixed by adjusting a_i (perhaps creating asymmetries among identical bidders). In particular, we must check:

1. F_i is non-decreasing: $F_i'(b) \geq 0$ for all $b \in B_i$.

2. $F_i(A_i) = 1$.

3. $F_i(0) = F_i(a_i)$.

The last two ensure that there are no atoms except at $b_i = 0$, as required in Claim 2.

[18] Barut and Kovenock (1998) prove this for homogenous bidders with multiple prizes; the same proof (with minor adaptation) applies here to two bidders of the same type.

References

Barut, Yasar, and D. Kovenock. (1998) "The Symmetric Multiple Prize All-pay Auction with Complete Information," *European Journal of Political Economy* **14**, 627-644.

Baye, Michael, D. Kovenock, and C. de Vries. (1993) "Rigging the Lobbying Process: An Application of the All-pay Auction," *American Economic Review* **83**, 289-294.

Baye, Michael, D. Kovenock, and C. de Vries. (1996) "The All-Pay Auction with Complete Information," *Economic Theory* **8**, 291-305.

Baye, Michael, D. Kovenock, and C. de Vries. (2009) "Contests with Rank-order Spillovers," *Economic Theory*, forthcoming.

Clark, Derek, and C. Riis. (1996) "A Multi-Winner Nested Rent-Seeking Contest," *Public Choice* **87**, 177-184.

Clark, Derek, and C. Riis. (1998) "Competition over More than One Prize," *American Economic Review* **88**, 276-289.

Cotton, Chris. (2010) "Pay-to-Play Politics: Informational lobbying and Contribution Limits when Money Buys Access." Mimeo, University of Miami.

Fang, Hanming. (2002) "Lottery Versus All-pay Auction Models of Lobbying," *Public Choice* **112**, 351-371.

Hassin, Refael, and M. Haviv. (2002) *To Queue or Not to Queue: Equilibrium Behavior in Queuing Systems*, Kluwer Academic Publishers: Boston, MA.

Hillman, Arye, and J. Riley. (1989) "Politically Contestable Rents and Transfers," *Economics and Politics* **1**, 17-39.

Hillman, Arye, and D. Samet. (1987) "Dissipation of Contestable Rents by Small Numbers of Contenders," *Public Choice* **54**, 63-82.

Holt, Charles, and R. Sherman. (1982) "Waiting-Line Auctions," *Journal of Political Economy* **90**, 280-294.

Klose, Bettina, and D. Kovenock. (2011) "The All-Pay Auction with Complete Information and Identity-Dependent Externalities," Mimeo, Purdue University.

Menicucci, Domenico. (2006) "Banning Bidders from All-pay Auctions," *Economic Theory* **29**, 89-94.

Oi, Walter. (1971) "A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly," *Quarterly Journal of Economics* **85**, 77-90.

Platt, Brennan. (2009) "Queue-rationed Equilibria with Fixed Costs of Waiting," *Economic Theory* **40**, 247-274.

Platt, Brennan, J. Price, and H. Tappen. (2010) "Pay-to-Bid Auctions," *NBER Working Paper*, 15695.

Reny, Philip. (1999) "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games," *Econometrica* **67**, 1029-1056.

Siegel, Ron. (2009) "All-pay Contests," *Econometrica* **77**, 71-92.

Taylor, Grant, K. Tsui, and L. Zhu. (2003) "Lottery or Waiting-line Auction?" *Journal of Political Economy* **87**, 1313-1334.

Tullock, Gordon. (1975) "On the Efficient Organization of Trials," *Kyklos* **28**, 745-762.

Tullock, Gordon. (1980) "Efficient Rent Seeking," Ch. 6 in Buchanan, Tollison, and Tullock, *Toward a Theory of the Rent-Seeking Society*, Texas A&M University Press, College Station, Texas.