16 July 2003

Auctions with Options to Re-Auction*†

Abstract

We examine a dynamic model of English auctions with independent private values. There is a single object for sale and it is not possible for the seller, who has a value of zero for the object, to commit not to sell in the future if a sale is not accomplished today. The seller may be able to commit to a reserve price, or make a cheap-talk announcement of a reserve price and secretly bid for the object herself in order to re-auction it in a later round with a new set of bidders. Bidders are "short-lived" in the sense that at the end of each round all existing bidders vanish and new bidders start arriving. This framework allows us to obtain existing results for one shot-auctions as special cases. This framework also allows us to capture some of the features of thick internet auctions and to obtain some new insights on the role of commitment, on optimal length and on socially optimal reserve prices that are not apparent from a one-shot auction perspective.

JEL: D44, D82.

Key words: shill bidding, re-auction option, reserve price, internet auctions.

Simon Grant
Department of Economics
Rice University
and
School of Economics
Australian National University

Atsushi Kajii
Institute of Social and Economic Research
Osaka University

Flavio Menezes
School of Economics
Australian National University
and
EPGE/FGV

Matthew J. Ryan
Department of Economics
University of Auckland

*While accepting responsibility for all remaining shortcomings of the paper, we would like to thank Paulo Monteiro and seminar participants at the Australian National University, University of New South Wales, University of Sydney and ESAM 2002 for their comments. Particular thanks are owed to Guillaume Rocheteau, who carefully read earlier drafts and offered many valuable suggestions. The financial support of ARC (Grant no. A000000055) is gratefully acknowledged.

†Corresponding Author: Flavio Menezes. Address: Australian National University, School of Economics (Building no. 26), Canberra, ACT, 0200, Australia. E-mail: Flavio.menezes@anu.edu.au. Phone: 61-2-61252651. Fax: 61-2-61255124.

1 Introduction

We examine a dynamic model of English auctions with independent private values. There is a single object for sale and it is not possible for the seller, who has a value of zero for the object, to commit not to sell in the future if a sale is not accomplished today. The seller may be able to commit to a reserve price, or make a cheap-talk announcement of a reserve price and secretly bid for the object herself in order to re-auction it in a later round with a new set of bidders. Bidders are "short-lived" in the sense that at the end of each round all existing bidders vanish and new bidders start arriving. This repeated auction structure captures some of the features of thick internet auctions for items such as used computer parts, software, etc.

This framework allows us to obtain existing results for one shot-auctions (such as the optimal auction results established by Myerson (1981) and Riley and Samuelson (1981)) as special cases.[1] Our framework also allows us to obtain some new insights on the role of commitment, on optimal length and on socially optimal reserve prices that are not apparent from a one-shot auction perspective.

We find that binding reserves continue to be valuable to the seller. In fact, both the optimal "secret" reserve – the one employed in the absence of commitment where any announcement by the seller is interpreted by bidders as "cheap-talk" – and the optimal binding (or "public") reserve are *both* higher than in the one-shot case. This is because the option value of the object exceeds its direct consumption value. The optimal "secret" reserve is thus raised to this higher option value level; while the optimal binding reserve also increases, because the opportunity cost of not selling in the current round is reduced.

It is a stylized fact that reserve prices are sometimes kept secret in auctions. For example, Horstmann and LaCasse (1997) argue that a common feature of many auctions, namely the seller's refusal to sell to the high-bidder, might be interpreted as evidence of the existence of a

[1] These results assume an exogenous, and commonly known, number of bidders, n, but the optimal reserve price is independent of n when all bidder valuations are drawn from the same distribution with the property that the hazard rate is strictly increasing. Moreover, the second-price auction with reserve remains optimal even if *bidders* are uncertain of n (McAfee and McMillan, 1987). Indeed, provided bidders have symmetric posteriors over the number of rivals they face, these authors show that the first-price auction with reserve will also remain optimal.

secret reserve.[2] These authors explain the existence of a secret reserve price as a signalling device about the (common) value of the object being sold. Such information could not be credibly transmitted to buyers by a publicly announced reserve price because of adverse selection and so the seller uses costly delay as a signalling device.[3] In our setting, with independent-private-values, such signalling is not possible. Instead, we offer an alternative rationalization for the rejection of high bids. A rejected high bid may indicate that the high bidder was in fact the seller. In the absence of a commitment technology[4] , seller bidding allows the seller to replicate the effect of a binding public reserve. Indeed, we show that seller bidding (or *shilling*) results in the *same expected revenue* as a binding public reserve. This equivalence is explained using standard one-shot auction arguments. This contrasts with a view held by many practitioners that shilling helps the seller. Indeed, eBay does not allow the seller to bid in its internet auctions. Similarly, the New South Wales State Government in Australia has moved to end shilling by requiring bidders to register and by limiting to one the number of bids a Vendor is allowed to make.

In our repeated auction framework, the seller's optimal reserve price, and the value of commitment, are natural generalizations of the one-shot results. However, the welfare conclusions are quite different. Because the good must eventually be sold, the question of allocative efficiency asks whether it is sold too quickly or too slowly. We show that the sale is made too quickly in the absence of a reserve price commitment (unless the seller bids on her own account). That is, the "secret" reserve price is *too low*. This contrasts with the one-shot scenario, in which the "secret" reserve – which is just the seller's direct consumption valuation of the good – is efficient. We also show that the privately optimal public reserve is too high from the social viewpoint, as in the standard one-shot scenario.

In the sequel we also consider the more general problem in which the seller chooses the duration

[2] Ashenfelter (1989) reports significant percentages of high-bid rejections in wine and fine art auctions. Jones, Menezes and Vella (1996) report high-bid rejection rates for wool auctions in Australia.

[3] Vincent (1995) argues that a secret reserve price might be used to increase bidder participation at the auction in a common value environment. His model, however, cannot explain why high bids are often rejected once an auction takes place.

[4] Typically explicit reserves are legally binding in the sense that the seller must sell if the reserve is exceeded. Our theory requires a commitment in the other direction: a commitment not to sell below the reserve. Thus, even though sellers might have access to a technology to implement the former type of commitment, it is less clear that sellers have access to a technology that implements the latter type of commitment.

of each auction round, as well as the reserve price. In eBay auctions, for example, the seller may choose a 3, 5, 7 or 10 day format. Unsurprisingly, if reserve prices are "secret", then the optimal length of a round of the auction is finite and strictly positive. For the "public" reserve auctions, we identify two countervailing forces at work. Shorter auctions reduce the time costs of the selling mechanism. However, longer auctions raise the probability of two or more bidders arriving, and hence increase the competition amongst bidders. If bidders cannot arrive in the same instance then in the limit, as the length of each round goes to zero, the likelihood that the object will be sold at the reserve goes to one: the auction collapses to a posted price mechanism.

Related Literature

The existing literature on repeated auctions has focussed on the case of long-lived bidders. For example, McAfee and Vincent (1997) consider an infinitely repeated, second-price IPV model in which the same n bidders participate in each round. As discount rates converge to unity, a version of the Coase conjecture takes hold, and the reserve price converges rapidly to zero over time. With bidder valuations distributed Uniformly on $[0, 1]$, McAfee and Vincent show that the value of commitment to a reserve price is seriously undermined by the inability to commit not to re-auction.

Burguet and Sákovics (1996) obtain a contrary result. They show that optimal reserves may be high (relative to the utility of the object to the seller) and valuable in a two-stage, first-price IPV model when buyers face participation costs, even with no discounting. They assume, however, that the seller does not to impose any reserve in the second-stage auction, despite having had access to a suitable commitment technology in the first period to implement a first-period reserve.

Haile (2000) also has a two period model where re-auctioning is assumed not to be possible after the second stage. Instead, a re-sale market opens in the second period, in which the seller commits not to participate. Since buyers are uncertain about their true valuations when they bid, this re-sale market is always active with some positive probability. Haile's main results concern the effect of the re-sale market on optimal bidding strategies in the (second-price) auction, but he also shows that high reserves may be valuable.

Therefore, when bidders are (infinitely) long-lived, and the seller can never credibly commit not to re-auction, it seems that reserve prices lose much of their value to the seller. However, the case of short-lived bidders – where each round attracts a new cohort of participants – has received comparatively scant attention. The motivation to study short-lived bidders comes from thick internet auction markets, where bidders who do not buy from a seller are likely to buy elsewhere and exit the market before this seller's next auction. As we demonstrate below, with short-lived bidders, the Coasian logic vanishes and optimal reserves are typically even higher than in the one-shot case. This optimal reserve, in the absence of a suitable commitment technology, can be implemented by seller bidding, but seller bidding cannot improve on the optimal public reserve auction.

Thus, our model dispenses with the untenable assumption of a seller commitment not to re-auction the object, but (a) implies a high value of commitment to a reserve price; and (b) provides a possible explanation for the phenomena of objects being passed-in in the absence of an announced reserve, even when a costless technology exists to bind the seller to the reserve.

A related paper is that of Wang (1993). He considers bidders who arrive according to a **homogeneous** Poisson process. The object is worth zero to the seller, there is no discounting but re-auctioning is allowed. The seller can sell the object either by auction or by posting a price. In the former option, the seller incurs a display cost at rate θ_d until an arriving buyer agrees to pay the posted price. In the latter option, the seller incurs a storage cost at rate θ_s until he sells the object. There is also a cost θ_a of running an auction and the seller chooses a future time at which the object is to be auctioned. If the object is not sold in the auction, the seller will keep it in storage and plan a future date for another auction. In our paper, we implicitly assume that, with the internet, displaying and storing costs are the same. In Wang, potential bidders arrive randomly, get notice of the auction and then return. Again internet auctions operate in a way that eliminates the need for passive bidders to make an active decision to return for the auction. While Wang's motivation is to compare posted-price selling and auctions, we investigate the effects of reserve price commitments. We emphasize that Wang's analysis can be recast in a way that allows us to offer some insights on the socially optimal reserve price and into the choice of the optimal length of the auction in thick internet markets. The model is presented in the next section. Our results are presented in section 3.

2 The model

An auction may be held in each of an infinite number of discrete periods. Any given round lasts T units of time[5] , with elapsed time in the current round denoted by t, an element of the interval $[0, T]$. No time elapses between the end of one auction round and the start of the next. In each round an English auction, modelled as a "button auction" (Milgrom and Weber, 1982), is used. Although we abstract from any fixed monetary cost of conducting an auction, there is one cost to re-auctioning in our model: any bidder present at time T in some round is assumed no longer to be present at time $t = 0$ of the next. This assumption is meant to capture a stylized feature of thick internet auctions.[6]

There are infinitely many potential buyers. Bidders arrive according to a stochastic process to be described below. Once buyers arrive, their values are determined by independent draws from a common cummulative distribution function F with strictly positive density f on $[0, 1]$. A buyer's own value is private information; but F is common knowledge. We shall assume throughout this paper:

Assumption 1 *The draws determining individuals' valuations are independent of the arrival process.*

We assume an arrival process with the following properties:

Assumption 2 *The probability of n arrivals in any auction round lasting T units of time is given by $p(n, T)$ and the expected number of arrivals for any given T is finite; that is, $\sum_{n=0}^{\infty} np(n, T) < \infty$.*

The first part of Assumption 2 makes the seller's problem stationary. For example, we could fix some stochastic process on $[0, T]$ that is used to determine the number of arrivals. The same process

[5] We shall relax this assumption, by making the duration of auction rounds endogenous, in section 3.6.

[6] In internet auctions, for example, where automated proxy bidding is common, there may be a standing bid at $t = T$ without the bidder being on-line. Also, the subsequent auction may potentially be conducted at a different site, with no "forwarding address" posted for bidders in the previous round.

is used in each round. Players arrive according to the stochastic process described above and at the end of each round they vanish and a new auction starts with bidders arriving according to the same process. That is, the distribution of the number of arrivals during a round is independent and identical across rounds. Alternatively, one might imagine a **homogeneous** Poisson arrival process on $[0, \infty)$ that is chopped into segments of length T. The homogeneity **property** implies that this is equivalent to restarting the process from time zero each T units of time.

The second part of Assumption 2 is a mild regularity assumption. This part of the assumption is also satisfied by the **homogeneous** Poisson process with **constant arrival** rate $\lambda > 0$.

We also assume that bidders play "myopically" (e.g. that each auction round attracts a different pool of bidders). The seller, however, has regard for the future beyond the current round, and a (continuous) rate of time preference $\rho > 0$. The seller values the object at zero.

We consider cases where a publicly announced reserve may or may not be enforceable, and the seller may or may not be able to shill bid.[7] When announced reserves are non-binding ("cheap talk") we say that the seller imposes a "secret reserve". We therefore use "public reserve" to refer to an announced reserve that is binding.

3 Results

3.1 Bidding strategies

The optimal bidding strategies are straightforward to determine. There is a weakly dominant strategy to bid up to one's valuation. For example, bidders could submit maximum bids equal to their values to proxy bidders in an eBay auction.

Due to the fact that bidders arrive at different times, proxy bidding slightly modifies the mechanics of the "button auction". If the current standing bidder faces no competition – either from a rival bidder, or from the seller's reserve price – he temporarily releases the "button" to stop

[7] Chakraborty and Kosmopoulou (2002) use the term "shill bidding " to describe the practice of sellers sending in bids as though they were legitimate bidders to influence the beliefs of the other bidders about the object's true value (i.e., erroneous signals sent by the seller to mislead other bidders in a common value auction). What we (and Wang, Hidvegi and Whinston, 2001) describe as shill bidding, Chakraborty and Kosmopoulou would term "phantom bidding", where the price is artificially jacked up to extract more surplus from the bidder with the highest value.

the price from rising further, but does not exit the auction. If a new bidder arrives, the original bidder depresses the "button" again. Once the bidder's valuation is reached, he permanently exits the auction.[8]

In summary, an optimal bidding strategy is described as follows. Let r denote the reserve price anticipated by bidders: r is the public reserve for auctions in which public reserves are posted and binding, or the anticipated secret reserve in a cheap-talk scenario. A bidder with value $v \in [r, 1]$ will submit bids in $[r, v]$ in order to meet the reserve and/or to remain the standing bidder. If bidding exceeds v this bidder will drop out of the auction. Bidders with values in $[0, r)$ expect zero surplus from participation in the auction. They may choose not to bid, or else submit a bid in $[0, r)$.

Note that the optimal bidding strategy is not affected by the presence or otherwise of shilling. Seller bids have the same consequences for bidders as "genuine" rival bids: if a seller bid stands at T, the object is passed in for re-auctioning.

3.2 The seller's decision problem

Observe that the seller's problem is stationary: it looks the same at the start of each auction round. Moreover, as bidders have a weakly dominant bidding strategy, the seller's problem is essentially non-strategic: it is a dynamic programming problem.

Leaving aside seller bidding for the moment, the seller's problem is to choose a reserve price, or acceptance rule, for each auction period. Because of the stationarity of the problem, and since the revenue in each round is bounded, an optimal strategy exists, and it can be written as a stationary policy function. So we shall focus on the strategy in which the same reserve price is applied in each period.

In a secret-reserve scenario, the reserve is secret, so we may think of the seller deciding on her acceptance rule at T: she may accept or reject the standing bid at that time.[9] With a public

[8] Some authors, such as Roth and Ockenfels (2001), have observed a tendency for "sniping", or last-minute bidding, in on-line auctions. However, the use of proxy bidding in eBay auctions avoids this. Other on-line auctions employ "soft" endings to the same effect. Auctions may be extended for short periods (say, 5 minutes) until no further activity is recorded; a sort of virtual "going...going..gone!". This, too, restores the usual second-price logic, and eliminates any incentive to "snipe". See Lucking-Reiley (2000).

[9] There is also no advantage from separately announcing the reserve, as it has no signalling value in the IPV

reserve price, the acceptance rule is set at $t = 0$. Moreover, the reserve will affect the bidders' strategies, as described above. In each case, however, the institutional assumptions we have made imply that the acceptance rule will take the form of a cut-off (reserve) price, r, such that the standing bid at T is accepted if and only if it is at least r.

We can now compute the seller's value function. Let $\Pi(r,T)$ be the probability that an auction round with reserve r ends with the object being passed in; and $R(r,T)$ be the expected revenue generated in one auction round with reserve r. Assumption 1 allows us to write:

$$\Pi(r,T) = \sum_{n=0}^{\infty} p(n,T) F(r)^n = p(0,T) + \sum_{n=1}^{\infty} p(n,T) F(r)^n$$

$$R(r,T) = \sum_{n=1}^{\infty} p(n,T) R_n(r)$$

where $R_n(r)$ is the expected revenue from a one-shot auction with n bidders and reserve r. This quantity may be expressed as:[10]

$$R_n(r) = \int_r^1 J(z)\, g_1^n(z)\, dz$$

context, and is thus ignored by bidders: it is "cheap talk".

[10] If $n = 1$ we have $g_1^n(z) = f(z)$, so

$$\begin{aligned} R_1(r) &= \int_r^1 J(z) f(z)\, dz \\ &= \int_r^1 z f(z) - [1 - F(z)]\, dz \end{aligned}$$

Since $zf(z) - [1 - F(z)] = \frac{d}{dz}[F(z) - 1]\, z$, we see that

$$R_1(r) = r[1 - F(r)]$$

which is the expected revenue from an auction with only one bidder.

For $n \geq 2$, let $g_2^n(z)$ be the density of the second order statistic from n random draws. Then we observe that

$$\begin{aligned} \int_r^1 z g_2^n(z)\, dz &= \int_r^1 n(n-1)[1 - F(z)] F(z)^{n-2} f(z)\, z\, dz \\ &= \int_r^1 n[1 - F(z)]\, z \frac{d}{dz} F(z)^{n-1}\, dz \\ &= -nF(r)^{n-1}[1 - F(r)]\, r + \int_r^1 J(z)\, g_1^n(z)\, dz \end{aligned}$$

using integration by parts in the last step. Therefore:

$$\begin{aligned} R_n(r) &= nF(r)^{n-1}[1 - F(r)]\, r + \int_r^1 z g_2^n(z)\, dz \\ &= \int_r^1 J(z)\, g_1^n(z)\, dz \end{aligned}$$

where

$$J(z) = z - \frac{[1 - F(z)]}{f(z)}$$

and $g_1^n(z)$ is the density of the *first* order statistic for n random draws from distribution F.

Write $v_0(r,T)$ for discounted expected revenue from the repeated auction with reserve r imposed in every round. By definition, $v_0(r,T)$ is the sum of discounted expected revenue from one shot auction, plus the discounted continuation value. Thus

$$v_0(r,T) = e^{-\rho T}[R(r,T) + \Pi(r,T)v_0(r,T)].$$

So we may therefore write the value function as

$$v_0(r,T) = \frac{e^{-\rho T}R(r,T)}{1 - e^{-\rho T}\Pi(r,T)}. \tag{1}$$

Lemma 1 *$v_0(r,T)$ is continuous in r.*

Proof. We first show that $\Pi_1(r,T)$ is continuous in r. Fix any r and $\varepsilon > 0$. By the continuity of F, there is $\delta > 0$ such that $|r - r'| < \delta$ implies that $|[F(r)] - [F(r')]| < \varepsilon$. Now

$$[F(r)]^n - [F(r')]^n =$$

$$([F(r)] - [F(r')])\left(F(r)^{n-1} + F(r)^{n-2}F(r') + ... + F(r')^{n-1}\right),$$

and $0 \leq \left(F(r)^{n-1} + F(r)^{n-2}F(r') + ... + F(r')^{n-1}\right) \leq n$ since $F(r)$ is a cumulative distribution. So $|r - r'| < \delta$ implies $\left|[F(r)]^n - [F(r')]^n\right| < \varepsilon n$ for all n. Now

$$\begin{aligned}
\Pi_1(r,T) - \Pi_1(r',T) &= \sum_{n=1}^{\infty} p(n,T)F(r)^n - \sum_{n=1}^{\infty} p(n,T)F(r')^n \\
&= \lim_{k\to\infty}\sum_{n=1}^{k} p(n,T)F(r)^n - \lim_{k\to\infty}\sum_{n=1}^{k} p(n,T)F(r')^n \\
&= \lim_{k}\sum_{n=1}^{k} p(n,T)\left(F(r)^n - F(r')^n\right).
\end{aligned}$$

So $|r - r'| < \delta$ implies that

$$\begin{aligned}
\left|\lim_{k}\sum_{n=1}^{k} p(n,T)\left(F(r)^n - F(r')^n\right)\right| &\leq \lim_{k}\sum_{n=1}^{k} p(n,T)\left|\left(F(r)^n - F(r')^n\right)\right| \\
&< \lim_{k}\sum_{n=1}^{k} p(n,T)\,n\varepsilon \\
&= \varepsilon\lim_{k}\sum_{n=1}^{k} p(n,T)\,n,
\end{aligned}$$

and by assumption $\lim_k \sum_{n=1}^{k} p(n,T)\, n$ is a finite constant. This shows that $\Pi_1(r,T)$ is continuous in r.

A similar argument shows that $R(r,T) = \sum_{n=1}^{\infty} p(n,T)\, R_n(r)$ is continuous in r, since $R_n(r)$ is continuous. Thus $v_0(r,T) = \frac{e^{-\rho T} R(r,T)}{1-e^{-\rho T}\Pi(r,T)}$ is continuous in r, since the denominator is always positive. ■

3.3 The value of a public reserve

Let us first compare, from the seller's point of view, a secret reserve auction with a public reserve auction. Is there value to the seller from committing to a public reserve price?

For the secret-reserve auction, the seller's strategy involves choosing an acceptance price (or reserve), r, being the lowest price at which she is prepared to sell the object in any round. The buyers' bidding strategies are as explained before. We consider a Nash equilibrium of the game; bidding is optimal given r, and r is chosen optimally given the bidding strategy. If buyers are bidding optimally given r, seller's payoff is $v_0(r,T)$ if r is chosen. Equilibrium then requires that

$$r = v_0(r,T). \tag{2}$$

Equation (2) says that the seller's reserve is equal to the value of re-starting the auction game, given the bidders' optimal responses to this reserve. If bidding fails to reach v_0, it is better to pass in the object; and accepting a bid above v_0 is preferable to re-auctioning the object.

Proposition 2 *There exists an equilibrium secret reserve, and all solutions to (2) are interior to* $[0,1]$.

Proof. Note that $v_0(0,T) > 0$ and $v_0(1,T) = 0$. That is, the seller expects a non-zero surplus in a "no reserve" auction; and setting a reserve of $r = 1$ ensures that the probability of achieving a sale in any given round is zero. Since $v_0(r,T)$ is continuous in r, $r = v_0(r,T)$ must have a solution and any solution must be an interior point. ■

For the public-reserve scenario, the seller's strategy again consists of choosing an acceptance price (reserve) r. However, r must now be stated publicly and committed to at the start of the auction round. The optimal reserve price therefore solves

$$\max_{r\in[0,1]} v_0(r,T) \tag{3}$$

Again, continuity of v_0 ensures the existence of a solution to (3).

The seller can clearly do no worse in the public-reserve auction, since the value of the secret-reserve auction is $v_0(r,T)$ evaluated at an r that satisfies (2). The seller would rather choose the *best anticipated* reserve, than choose a reserve that is *anticipated* and best given this *fixed bidder expectation*. Hence we have:

Proposition 3 *From the seller's point of view, the public-reserve-price scenario is preferable to the secret-reserve-price scenario.*

Proposition 3 can be understood by using insights from standard auction theory. From the point of view of a given time interval (auction round), we can take the seller's reservation value (her opportunity cost of reselling) as given. Thus, the seller's problem can be recast as single-object one-shot standard auction where the seller's valuation is equal to this opportunity cost.[11]

We know that the optimal auction in the standard auction literature can be implemented by an English auction with an optimally chosen public reserve. Therefore, if Proposition 3 were not true, then this would contradict the fact that public reserve prices are valuable.

In general, the optimal public reserve will be higher than the equilibrium secret reserve price under the following regularity assumption.

Assumption 3 $J(r) = r - \frac{[1-F(r)]}{f(r)}$ *is non-decreasing in* r*.*

Raising a reserve is more beneficial if done publicly, since it induces higher bids from high-value arrivals. Conversely, reductions in reserve are best kept secret, to improve the chances of a sale without encouraging bid reductions from high-value bidders facing no bidder competition. This creates a tendency for reserves to be higher when public and binding, than when they are secret. More precisely:

[11] For a formal argument see the proof of Proposition 4.

Proposition 4 *Let r^* denote an optimal public reserve and r^{**} an equilibrium secret reserve price. Suppose Assumption 3 holds. Then*

$$v_0(r^*, T) \;=\; r^* - \frac{[1 - F(r^*)]}{f(r^*)} \tag{4}$$

and hence $r^* > r^{**}$. Furthermore $r^* \in (0, 1)$.

Proof. Let $\hat{v}(T)$ denote $v_0(r, T)$ evaluated at an optimal public reserve price. Then

$$\begin{aligned} e^{\rho T}\hat{v}(T) &= \max_r \{R(r,T) + \Pi(r,T)\,\hat{v}(T)\} \\ &= \max_r \left\{ \sum_{n=1}^{\infty} p(n,T)\left[R_n(r) + F(r)^n \hat{v}(T)\right] \right\} + p(0,T)\,\hat{v}(T) \end{aligned}$$

But

$$\max_r R_n(r) + F(r)^n \hat{v}(T)$$

is the problem of choosing an optimal public reserve in a one-shot auction with n bidders and a seller valuation $\hat{v}(T)$ for the object. For *any* n, the FOC for this problem is

$$J(r) \;=\; \hat{v}(T), \tag{5}$$

and Assumption 3 implies that the problem is concave. Hence

$$\max_r \left\{ \sum_{n=1}^{\infty} p(n,T)\left[R_n(r) + F(r)^n \hat{v}(T)\right] \right\} \;=\; \sum_{n=1}^{\infty} p(n,T)\left\{ \max_r R_n(r) + F(r)^n \hat{v}(T) \right\}$$

and the solution to (5) solves this maximization problem also. Then (4) follows from (5). From (2), (3) and (4): $r^{**} = v_0(r^{**}, T) \leq v_0(r^*, T) \leq r^*$. It follows that $r^* \geq r^{**}$, with equality if and only if $1 - F(r^*) = 0$. But the latter is equivalent to $r^* = 1$, and Proposition 2 rules out $r^{**} = 1$. Therefore, $r^{**} < r^* < 1$. Finally, $r^* > r^{**}$ and Proposition 2 imply $r^* > 0$. ■

Re-arranging (4) gives

$$r^* \;=\; v_0(r^*, T) + \frac{[1 - F(r^*)]}{f(r^*)}$$

This matches the familiar expression for the optimal reserve in a one-shot IPV auction – see, for example, Riley and Samuelson (1981, Proposition 3) – except that the seller's valuation of the object, $v_0(r^*, T)$, is now the "option value" of retaining the object for re-auctioning. Since

$v_0(r^*, T) > 0$, the optimal public reserve is *higher* than in the one-shot case, since the seller's opportunity cost of not trading in any given round is greater than the consumption value of the object. With short-lived bidders, the Coasian logic of McAfee and Vincent (1997) is reversed.

3.4 Public reserve versus seller bidding

A seller bid submitted at time t has the same effect on bidder behavior as imposing a public reserve at t equal to this bid. The ability to shill bid is therefore equivalent to the ability to set a public reserve at $t = 0$ and alter it continuously throughout the auction. It follows that the seller is happy to forego access to a reserve price commitment technology provided she can bid in the auction.

Of course, many auctions, including on-line eBay auctions, have rules against seller bidding. The extent to which these rules are enforceable, especially in the context of virtual auctions, is highly debatable. In our framework, however, these regulations are entirely redundant, as the seller is *completely indifferent* between the seller-bidding and public-reserve-price formats. The additional strategic freedom afforded by seller bidding – essentially, the freedom to use a "flexible reserve" – adds no additional value.

Graham, Marshall and Richard (1990) showed that in a one-shot, English, IPV auction, if the seller has the opportunity to bid once all genuine bidding activity has ceased, her optimal bid is independent of the standing bid. We extend this result to our framework. The incentive for the seller to bid at t is shown to be independent of both the current standing bid and t, even though additional bidding activity may occur in $(t, T]$ and there is the possibility of re-auctioning the object.

Proposition 5 *Under Assumption 3 the seller-bidding and public-reserve-price auction formats generate the same* ex ante *expected revenue for the seller.*

The above proposition can be seen as a direct application of Graham, Marshall and Richard (1990) and so the proof is omitted. To see this let $\tilde{v}(T)$ be the value of the auction game when the seller uses an optimal shill bidding strategy. The usual dynamic programming logic implies that we may treat this as exogenous. The seller-bidding format is equivalent to the seller making a take it or leave it offer to the last remaining buyer (a final seller bid). Graham, Marshall and Richard (1990) then says that the optimal shill bid $\tilde{b}$ at $t = T$ is independent of the level of the standing bid at T. Hence, the optimal shill bid $\tilde{b}$ is identical to the optimal public reserve $\hat{r}$ (neither uses any of the information revealed during the auction). It follows that $\hat{v}(T) = \tilde{v}(T)$.

3.5 The surplus maximizing reserve price

In a one-shot auction, any non-zero reserve price is surplus reducing, since there is some chance that no trade occurs. Hence, privately optimal reserve price commitments are necessarily too high from a social efficiency point of view. With a repeated auction, this is no longer the case. Unless the reserve price is set at unity – which is precluded by Propositions 2 and 4 – the object will sell with probability one. The efficiency issue concerns whether the sale occurs too quickly or too slowly.

To be more precise, let us consider a Social Planner who must respect all the exogenous constraints of the auction mechanism. The Social Planner discounts at the same rate as the seller. At the end of any period, the Planner may therefore choose to allocate the object to any one of the arrivals during that period, or else wait one more round. We assume that the Social Planner observes the values of all arrivals during the current round, but does not know the values of future arrivals.[12] Which allocation rule maximizes the discounted expected value of total surplus?

As usual, stationarity implies that the Planner's rule will consist of a cut-off, r, such that the good is allocated to the highest value arrival during the current period if and only if the highest value exceeds r. Note, therefore, that the Social Planner's cut-off is a *valuation*, while the seller's is a *bid level*. However, in each case, the good is allocated in the current round if and only if a bidder arrives whose valuation exceeds the cut-off. In this sense, they are directly comparable for the purposes of determining the *allocative efficiency* of the auction.

Let $v_S(r, T)$ denote the Planner's value function. The Social Planner problem is of the "optimal stopping" variety. In each period, a maximum surplus is drawn, and the Planner must determine

[12] This seems to us the natural benchmark against which to assess the allocative efficiency of the auction mechanism. Alternatively, one may justify the lack of Social Planner omniscience by supposing that bidders do not receive their values for the object until they arrive. This value may depend, for example, on actions taken prior to arrival.

a cut-off level of surplus at which to stop the process. Observe that

$$e^{\rho T} v_S(r,T) \;=\; \left[1 - \Pi\left(r,T\right)\right] \mathbb{E}\left[z \,|\, z \geq r\right] \,+\, \Pi\left(r,T\right) v_S(r,T) \tag{6}$$

where $\mathbb{E}\left[z \,|\, z \geq r\right]$ is the expected value of the highest bidder valuation, conditional on this being at least r.

Since

$$\frac{d}{dr}\mathbb{E}\left[z \,|\, z \geq r\right] \;=\; \frac{-\pi(r,T)}{1-\Pi(r,T)}\left\{r \,-\, \mathbb{E}\left[z \,|\, z \geq r\right]\right\},$$

where

$$\pi(r,T) \;=\; \frac{\partial}{\partial r}\Pi\left(r,T\right) \;=\; \sum_{n=1}^{\infty} p\left(n,T\right) g_1^n\left(r\right),$$

we have

$$\left[e^{\rho T} \,-\, \Pi(r,T)\right] \frac{\partial}{\partial r} v_S(r,T) \;=\; \pi(r,T)\left[v_S(r,T) \,-\, r\right].$$

Therefore, the necessary first-order condition for an optimal choice of r is

$$\frac{\partial}{\partial r} v_S(r,T) \;= 0 \Longleftrightarrow v_S(r,T) = r \tag{7}$$

The Planner should allocate the object whenever the highest value arrival exceeds the continuation value, and this cut-off should also maximize the value function. We may immediately conclude:

Proposition 6 *The secret-reserve auction sells the object too quickly from a social efficiency point of view. That is, if r^{**} and $\hat{r}$ solve (2) and (7) respectively, then $r^{**} < \hat{r}$.*

Proof. It is clear that $v_0\left(r^{**},T\right) < v_S\left(r^{**},T\right)$, since the seller must share the (strictly positive) expected surplus with the buyer, while the Social Planner does not. Therefore, if $\hat{r} \leq r^{**}$ we have

$$v_S\left(\hat{r},T\right) \;=\; \hat{r} \;\leq\; r^{**} \;=\; v_0\left(r^{**},T\right) \;<\; v_S\left(r^{**},T\right).$$

But this contradicts the fact that $\hat{r}$ maximizes $v_S\left(r,T\right)$. ■

The intuition behind Proposition 6 is easy to see. Consider an auction round that ends with a high bid equal to r, which in turn is equal to the high bidder's valuation. Accepting the current high bid earns r for both the seller and the Social Planner, while passing-in earns the former $v_0(r,T)$ and the latter $v_S(r,T)$. Since the Social Planner anticipates the whole expected surplus from future trade, while the seller only anticipates a fraction – that is $v_0(r,T) < v_S(r,T)$ – the Social Planner is more patient. The secret-reserve auction sells too quickly.

Is the privately optimal *public* reserve too high or too low from a social point of view? In principle, the overall effect might seem ambiguous: the tendency to set too low a reserve may be offset by the gain to the seller of setting a higher reserve to extract additional surplus. However, we show next that the privately optimal public reserve is too high from a social point of view. Before we establish this result we need to establish a lemma.

Note that the seller's value function can be written as:

$$v_0(r,T) = \frac{\sum_{n=1}^{\infty} e^{-\rho T} p(n,T) \int_r^1 g_1^n(z) J(z)\, dz}{1 - \sum_{n=0}^{\infty} e^{-\rho T} p(n,T) F(r)^n}$$

Let $\tilde{p}(n,T) = e^{-\rho T} p(n,T)$,

$$\begin{aligned} \tilde{\Pi}(r,T) &= e^{-\rho T}\Pi(r,T) \\ &= \sum_{n=0}^{\infty} \tilde{p}(n,T) F(r)^n \end{aligned}$$

and

$$\tilde{\pi}(r,T) = \frac{\partial}{\partial z}\tilde{\Pi}(r,T) = \sum_{n=1}^{\infty} \tilde{p}(n,T) g_1^n(r)$$

Then:

$$\begin{aligned} \int_r^1 \tilde{\pi}(z,T)\, dz &= \tilde{\Pi}(1,T) - \tilde{\Pi}(r,T) \\ &= e^{-\rho T} - \tilde{\Pi}(r,T) \end{aligned}$$

and hence

$$\begin{aligned} v_0(r,T) &= \frac{\int_r^1 \tilde{\pi}(z,T) J(z)\, dz}{\int_r^1 \tilde{\pi}(z,T)\, dz + [1 - e^{-\rho T}]} \\ &= \left[\frac{\int_r^1 \tilde{\pi}(z,T)\, dz}{\int_r^1 \tilde{\pi}(z,T)\, dz + [1 - e^{-\rho T}]}\right] \int_r^1 J(z)\, \mu(z,r,T)\, dz \qquad (8) \\ &= \left[\frac{e^{-\rho T} - \tilde{\Pi}(r,T)}{1 - \tilde{\Pi}(r,T)}\right] \int_r^1 J(z)\, \mu(z,r,T)\, dz \qquad (9) \end{aligned}$$

where

$$\mu(z,r,T) = \frac{\tilde{\pi}(z,T)}{\int_r^1 \tilde{\pi}(z,T)\, dz} = \frac{\pi(z,T)}{\int_r^1 \pi(z,T)\, dz}$$

and

$$\pi(z,T) \;=\; \frac{\partial}{\partial z}\Pi(z,T) \;=\; \sum_{n=1}^{\infty} p(n,T)\, g_1^n(z)$$

Note that μ is a density on $[r,1]$, and $\tilde{\Pi}(r,T)$ is the discounted probability that no bidder with value at least r arrives in T units of time. Recall that $\Pi(r,T) = \sum_{n=0}^{\infty} p(n,T)\, F(r)^n$ is the undiscounted quantity.

Similarly, for the Social Planner's value function, we have:

$$v_S(r,T) \;=\; \left[\frac{e^{-\rho T} - \tilde{\Pi}(r,T)}{1 - \tilde{\Pi}(r,T)}\right] \int_r^1 z\mu(z,r,T)\; dz \tag{10}$$

Therefore:

$$v_S(r,T) - v_0(r,T) \;=\; \left[\frac{e^{-\rho T} - \tilde{\Pi}(r,T)}{1 - \tilde{\Pi}(r,T)}\right] \int_r^1 \left[\frac{1-F(z)}{f(z)}\right] \mu(z,r,T)\; dz \tag{11}$$

Observe that

$$\tilde{\Pi}_r(r,T) \;=\; e^{-\rho T} \sum_{n=0}^{\infty} p(n,T)\, g_1^n(r) \;>\; 0.$$

We can now establish our Lemma:

Lemma 7 $\frac{\partial}{\partial r}\left[\frac{e^{-\rho T} - \tilde{\Pi}(r,T)}{1 - \tilde{\Pi}(r,T)}\right] < 0.$

Proof.

$$\begin{aligned}
\frac{\partial}{\partial r}\left[\frac{e^{-\rho T} - \tilde{\Pi}(r,T)}{1 - \tilde{\Pi}(r,T)}\right] \;&=\; \frac{-\tilde{\Pi}_r\left[1-\tilde{\Pi}\right] + \tilde{\Pi}_r\left[e^{-\rho T} - \tilde{\Pi}\right]}{\left[1-\tilde{\Pi}\right]^2} \\
&=\; \frac{\tilde{\Pi}_r\left[e^{-\rho T} - 1\right]}{\left[1-\tilde{\Pi}\right]^2} \\
&<\; 0
\end{aligned} \tag{12}$$

since $e^{-\rho T} < 1$. ■

Proposition 8 *The socially optimal r is less than the privately optimal binding reserve when $[1-F(z)]/f(z)$ is decreasing.*

Proof. By direct calculation:

$$\frac{\partial}{\partial r}\left[v_S(r,T) - v_0(r,T)\right] \;=$$

$$\left[\frac{e^{-\rho T} - \tilde{\Pi}(r,T)}{1 - \tilde{\Pi}(r,T)}\right] \frac{\pi(r,T)}{\left(\int_r^1 \pi(z,T)\, dz\right)^2} \int_r^1 \pi(z,T) \left\{\left[\frac{1-F(z)}{f(z)}\right] - \left[\frac{1-F(r)}{f(r)}\right]\right\} dz$$
$$+ \frac{\partial}{\partial r}\left[\frac{e^{-\rho T} - \tilde{\Pi}(r,T)}{1 - \tilde{\Pi}(r,T)}\right] \int_r^1 \left[\frac{1-F(z)}{f(z)}\right] \mu(z,r,T)\, dz$$

The second term is negative by Lemma 7. If $[1 - F(z)]/f(z)$ is decreasing, then the first term is negative at any r, so

$$\frac{\partial}{\partial r} v_S(r,T) < \frac{\partial}{\partial r} v_0(r,T) \tag{13}$$

at any r. Furthermore, if $J(z)$ is increasing – which is implied by $[1 - F(z)]/f(z)$ decreasing – there is a unique interior maximum of $v_0(r,T)$ in r, and this is equal to the global maximum. From this fact and (13) it follows that the socially optimal reserve is less than the privately optimal one. ■

Intuition for this result is as follows. The second term reflects the Social Planner's relatively greater patience, since her option value (v_S) exceeds that of the seller (v_0), giving the Social Planner a greater incentive to increase r. The first term reflects the relative marginal incentive for rent-seeking through a higher reserve price. For the seller, this rent is

$$\frac{1-F(z)}{f(z)} = z - J(z).$$

Since the Social Planner's rent-seeking incentive is constant (zero), if the rent $[1 - F]/f$ is decreasing, the *relative marginal incentive* (of the Planner over the seller) is negative.

The result above extends Wang's (1996) Theorem 3 to allow for discounting and a general arrival process. Wang's Theorem 3 also states that if $\frac{1-F(z)}{f(z)}$ is increasing then we obtain the reverse result. However, it turns out that this case is empty as there does not exist a bounded distribution that satisfies this property as we show next.

Proposition 9 *There does not exist a bounded distribution that satisfies the property that both* $\frac{1-F(z)}{f(z)}$ *and* $J(z)$ *are increasing.*

Proof. Suppose $F(x)$ is differentiable with support $(0,1)$ and such that

$$\left(\frac{1-F(v)}{f(v)}\right)' > 0.$$

We want to find the family of distributions with support $[0,1]$ satisfying the above condition. If $\Pi(v) > 0, v \in (0,1)$ and let F be such that

$$\left(\frac{1-F(v)}{f(v)}\right)' = \Pi(v). \tag{14}$$

Define $\Pi(v) = \int_0^v \Pi(y)\,dy$. Integrating (14) we have

$$\frac{1-F(v)}{f(v)} = \Pi(v) + C,\ \ C \geq 0.$$

Taking the inverse

$$\frac{f(v)}{1-F(v)} = \frac{1}{\Pi(v)+C}.$$

Integrating on $[0,v]$ yields:

$$\log\left(\frac{1}{1-F(v)}\right) = -\log(1-F(v)) = \int_0^v \frac{dz}{\Pi(z)+C} \leq \frac{v}{C}.$$

Let $v \to 1$. The LHS $\to \infty$. Thus $C = 0$ and

$$\log\left(\frac{1}{1-F(v)}\right) = \int_0^v \frac{dz}{\Pi(z)}.$$

Rewriting

$$F(v) = 1 - \exp\left[-\int_0^v \frac{dz}{\Pi(z)}\right], \int_0^1 \frac{dz}{\Pi(z)} = \infty.$$

If $\Pi(\cdot)$ is bounded then $\Pi(v) < Kv$ and for all $a > 0$, $\int_0^a \frac{dz}{\Pi(z)} > \int_0^a \frac{dz}{Kz} = \infty$. Thus Π cannot be bounded and in particular $v - \frac{1-F(v)}{f(v)}$ cannot be increasing. ■

Given that we have established above that the equilibrium secret reserve price is too low and that the privately optimal public reserve is too high, it is of interest to determine which auction format comes closes to achieving the socially optimal expected surplus. In particular, are reserve price commitments socially desirable? For the special case of the uniform distribution of individuals' valuations on [0,1], $T = 1$ and a **homogeneous** Poisson arrival process, Figure 1 describes the difference between the social planner's value function evaluated at the privately optimal secret reserve (v_s^s) and the privately optimal public reserve (v_s^p) (as a percentage of the latter). For sufficiently high discount factors, the public reserve is relatively more efficient.[13]

[13] For the algebraic details, readers can refer to the longer working paper version in Grant, Kajii, Menezes and Ryan (2002).

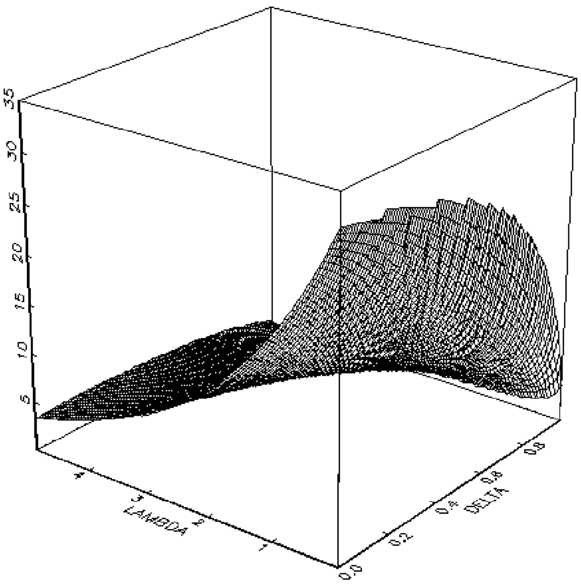

Figure 1: Social value of commitment: $100\left[v_s^s - v_s^p\right]/v_s^p$

3.6 Optimal auction length

Until now, we have fixed the length of auctions at T units of time. However, on-line auctions often allow sellers to choose their duration, at least within some bounds. It is therefore of interest to endogenize the length T of auction rounds. We shall allow $T = 0$ by defining $v_0(r, 0) := \lim_{T\downarrow 0} v_0(r, T)$.

Before endogenizing T, we must first revisit our interpretation of the bidder arrival process. Recall the two examples of processes satisfying Assumption 2 that were given in section 2:

(i) A process on $[0, T]$ with finite expected number of arrivals is given. Each round's arrivals represent the sample paths of independent draws from this process.

(ii) Auction rounds are described by dividing the interval $[0, \infty)$ into non-overlapping segments of length T. Arrivals within rounds are determined by a (homogeneous) Poisson arrival process on $[0, \infty)$.

It is straightforward to endogenize T within scenario (ii) which is a special case of (i), because of the homogeneity of this Poisson process. Importantly, in this case $v_0(r, 0)$ is equivalent to the value function when the seller uses a *posted price* mechanism, with price r posted at $t = 0$, and buyers arriving according to the underlying Poisson process. In particular:

$$v_0(r,\, 0) \;=\; \frac{r\lambda\left[1 - F(r)\right]}{\rho + \lambda\left[1 - F(r)\right]}$$

where λ is the Poisson arrival rate.

In the case of scenario (i), meaningfully endogenising T suggests the following natural revision:

(i′) A process on $[0, \infty)$ is given, having a finite expected number of arrivals in $[0, T]$ for any $T < \infty$. Given the length T for auctions, independent draws are made from this process for each round, with the sub-path on $[0, T]$ describing the arrivals for that round.

We now ask the question: Which value of $T \in [0, \infty)$ should the seller choose in order to maximize her expected discounted revenue? The following assumption suffices for $T = 0$ never to be optimal in a secret-reserve auction.

Assumption 4 *For each n, the function $p(n, T)$ is continuous in T. Moreover, the following two regularity conditions are satisfied:*

$$\lim_{T \downarrow 0} \frac{T}{1 - p(0, T)} \;>\; 0 \tag{15}$$

$$\lim_{k \to \infty} \frac{\sum_{n=2}^{\infty} p\left(n, T^k\right)}{p\left(1, T^k\right)} \;=\; 0, \textit{ for any decreasing sequence } \left\{T^k : k = 1, ...\right\} \textit{ with } \lim_{k \to \infty} T^k = 0 \tag{16}$$

The first condition in Assumption 4 ensures that, at any reserve price, the expected elapsed time until a sale is achieved remains bounded away from zero as T vanishes. The second condition would be implied by, for example, *uniformly orderliness* (Synder, 1975)[14] . It ensures that, as $T \downarrow 0$, the number of active bidders converges to unity conditional on a sale being achieved, and hence the expected sale price converges to r. Together, these two conditions deliver a version of the "Diamond paradox" (Diamond, 1971): seller surplus vanishes as $T \downarrow 0$.

Proposition 10 *Under Assumptions 1–4, $T = 0$ is not optimal in a secret-reserve-price auction.*

[14] The uniformly orderliness condition in our model can be stated as follows. Write $p(n, t, \zeta)$ is the probability of n arrivals in the interval $[t, t + \zeta]$. Then for any $\varepsilon > 0$ there exists a $\zeta > 0$ such that $\sum_{n=2}^{\infty} p(n, t, \zeta) \leq \varepsilon p(1, t, \zeta)$, for all $t \in [0, \infty)$.

Proof. First, note that the expected price conditional on making a sale at the end of round s is independent of s. Thus, given r and T, we can represent this expected price by $q(r,T)$. (If the probability of sale is zero, that is, $\Pi(r,T)=1$, we define $q(r,T)=0$.) We may therefore write:

$$v_0(r,T) = q(r,T)\left[\sum_{s=1}^{\infty}\Pi(r,T)^{s-1}\left[1-\Pi(r,T)\right]e^{-\rho sT}\right] \tag{17}$$

The square-bracketed term is the expected discount factor applied to the sales revenue.

We first show that

$$\lim_{T\downarrow 0} q(r,T) \le r \tag{18}$$

whenever $r>0$. See this as follows. Recall that $1-\Pi(r,T)$ is the probability of making a sale in any given round, with reserve r. Note also that $1-\Pi(r,T) \le 1-\Pi(0,T) = 1-p(0,T)$. Thus condition (15) implies that $1-\Pi(r,T)\to 0$ as $T\downarrow 0$. If there is some $\hat{T}>0$ such that $1-\Pi(r,T)=0$ for all $T<\hat{T}$, then (18) follows. Otherwise, there must exist an infinite sequence $\left\{T^k\right\}_{k=1}^{\infty}$ such that $T^k\downarrow 0$ as $k\to\infty$ and $1-\Pi\left(r,T^k\right)>0$ for all k. By continuity, it suffices to show:

$$\lim_{k\to\infty} q\left(r,T^k\right) \le r \tag{19}$$

But:

$$\begin{aligned} q\left(r,T^k\right) &\le \frac{p\left(1,T^k\right)\left[1-F(r)\right]}{1-\Pi\left(r,T^k\right)}r + \frac{\sum_{n=2}^{\infty} p\left(n,T^k\right)}{1-\Pi\left(r,T^k\right)} \\ &\le \frac{p\left(1,T^k\right)\left[1-F(r)\right]}{1-\Pi\left(r,T^k\right)}r + \frac{\sum_{n=2}^{\infty} p\left(n,T^k\right)}{p\left(1,T^k\right)\left[1-F(r)\right]} \end{aligned}$$

By (16), the second term vanishes, and so (19) follows.

We next show that

$$\lim_{T\downarrow 0}\sum_{s=1}^{\infty}\Pi(r,T)^{s-1}\left[1-\Pi(r,T)\right]e^{-\rho sT} < 1 \tag{20}$$

by virtue of (15). In fact:

$$\begin{aligned} \sum_{s=1}^{\infty}\Pi(r,T)^{s-1}\left[1-\Pi(r,T)\right]e^{-\rho sT} &= e^{-\rho T}\left[1-\Pi(r,T)\right]\sum_{s=0}^{\infty}\Pi(r,T)^{s}e^{-\rho sT} \\ &= e^{-\rho T}\left[1-\Pi(r,T)\right]\sum_{s=0}^{\infty}\left[\Pi(r,T)e^{-\rho T}\right]^{s} \\ &= \frac{1-\Pi(r,T)}{e^{\rho T}-\Pi(r,T)} \end{aligned}$$

Observe that

$$\frac{1-\Pi(r,T)}{e^{\rho T}-\Pi(r,T)} = \frac{\frac{1-\Pi(r,T)}{T}}{\frac{e^{\rho T}-1}{T}+\frac{1-\Pi(r,T)}{T}}$$

Since

$$\lim_{T\downarrow 0}\frac{e^{\rho T}-1}{T} = \rho > 0$$

and (15) implies

$$\lim_{T\downarrow 0}\frac{1-\Pi(r,T)}{T} \leq \lim_{T\downarrow 0}\frac{1-\Pi(0,T)}{T} = \lim_{T\downarrow 0}\frac{1-p(0,T)}{T} < \infty$$

it follows that (20) holds.

Combining (18) and (20) gives

$$\lim_{T\downarrow 0} v_0(r,T) < r$$

whenever $r>0$. Therefore, $v_0(r,0)=0$ iff $r=0$, which implies $T>0$ is optimal. ∎

When the seller cannot commit to her posted price, and she faces a "search cost" of finding another buyer (the expected delay until the next arrival), the only credible price announcement is zero. Absence of commitment deprives the seller of any effective market power. To acquire surplus from trade, she must hold an auction long enough that there is some non-zero probability of buyer competition.

However, for a public-reserve auction, matters are not so clear. The seller does not need to use non-zero auction length to overcome her price commitment problem. Of course, longer rounds do increase buyer competition and hence seller surplus as before, but they also delay trade. It is not clear *a priori* which T will provide an optimal balance between these two effects. However, for the case in which F is Uniform on [0,1], we can show that zero-length auctions (price posting under scenario (ii)) may be optimal with a public reserve. (For these details, in addition to several specific results for this distribution, readers can refer to the longer working paper version in Grant, Kajii, Menezes and Ryan (2002).). The optimality of price-posting (in the presence of commitment) does not undermine the relevance of the model for the analysis of internet auctions. On eBay, for example, sellers have the option of posting a "Buy It Now" price. In the absence of a standing bid (above the reserve), if an arriving bidder offers the "Buy It Now" price, it is automatically accepted and the auction is cancelled.

4 Concluding Remarks

The model presented here offers a natural generalization of the one-shot auction that allows for re-auctioning. As $\rho \to \infty$ the familiar one-shot results are recovered. However, the generalized framework is useful in a number of respects.

First and foremost, it will be rare that sellers will be able to credibly commit not to re-auction the object. In particular, internet auction sites allow a seller to re-auction quickly and at negligible cost.

Second, the potential to re-auction allows for a more realistic analysis of reserve prices. This potential generates a non-zero option value of retaining the object, and hence raises both "secret" and "public" reserves. However, with our short-lived bidders, a substantial value from a reserve price commitment remains. We are also able to confirm that seller bidding offers no advantage over a reserve price commitment in our setting.

Third, re-auctioning significantly alters the welfare properties of different auction mechanisms. "Secret" reserves are too low for allocative efficiency; and reserve price commitments may be socially preferable.

Finally, by modelling the bidder arrival process and endogenizing auction duration, our framework also nests price posting as a limiting case (when $T \to 0$). One may then observe the relative value of price versus time commitments for revenue generation. In the absence of price commitments, the logic of the *Diamond paradox* implies value in a commitment to $T > 0$, as this will increase buyer competition.

References

Burguet, R. and J. Sákovics, 1996, "Reserve Prices without Commitment," **Games and Economic Behavior 15**, 149-164.

Chakraborty, I. and G. Kosmopoulou, 2002, "Auctions with Shill Bidding, " mimeo, University of Oklahoma.

Durrett, R., 1991, **Probability: Theory and Examples**. Brooks/Cole Publishing, California.

Diamond, P., 1971, "A Model of Price Adjustment," **Journal of Economic Theory 3**, 156-168.

Graham, D. A., R. Marshall and J.F. Richard, "Phantom Bidding Against Heterogeneous Bidders," **Economics Letters 32**, 1990, 13-17.

Grant, S., A. Kajii, F. Menezes and M. J. Ryan, 2002, "Auctions with Options to Re-Auction," Discussion Paper 55, Tilburg University, Center for Economic Research.

Haile, P. A., 2000, "Auctions with Private Uncertainty and Resale Opportunities," forthcoming in **Journal of Economic Theory**.

Harstad, R. M., J. H. Kagel and D. Levin, 1990, "Equilibrium Bid Functions for Auctions with an Uncertain Number of Bidders," **Economics Letters 33**, 35-40.

Holt, C. A., Jr., 1979, "Uncertainty and the Bidding for Incentive Contracts," **American Economic Review 69(4)**, 697-705.

Horstman, I. and C. LaCasse, 1997, "Secret Reserve Prices in a Bidding Model with a Resale Option," **American Economic Review 87**, 663-684.

Jones, C., F. M. Menezes and F. Vella, 1996, "Auction Price Anomalies: Evidence from Wool Auctions in Australia." Australian National University Working Paper No. 303.

Lucking-Reiley, D., 2000, "Auctions on the Internet: What's Being Auctioned, and How?" **Journal of Industrial Economics 48(3)**, 227-252.

McAfee, R. P. and J. McMillan, 1987, "Auction with a Stochastic Number of Bidders," **Journal of Economic Theory 43,** 1-19.

McAfee, R. P. and D. Vincent, 1997, "Sequentially Optimal Auctions," **Games and Economic Behavior 18**, 246-76.

Milgrom, P. R. and R. J. Weber, 1982, "A Theory of Auctions and Competitive Bidding," **Econometrica 50**, 1089-1122.

Myerson, R. B., 1981, "Optimal Auction Design," **Mathematics of Operations Research 6(1)**, 58-73.

Porter, R. H., 1995, "The Role of Information in Offshore Oil and Gas Lease Auctions." **Econometrica 63(1)**, 1-28.

Riley, J.G. and W. F. Samuelson, 1981, "Optimal Auctions," **American Economic Review 71(3)**, 381-392.

Roth, A., E. and A. Ockenfels, 2001, "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet." Unpublished working paper, Harvard University.

Synder, D., 1975, **Random Point Processes**. John Wiley & Sons, New York.

Vincent, D., 1995, "Why Reserve Prices May Be Kept Secret," **Journal of Economic Theory 65(2)**, 575-584.

Wang, R., 1993, "Auctions versus Posted-Price Selling, " **American Economic Review 83(4)**, 838-851.

Wang, W., Z. Hidvegi and A. B. Whiston, 2001, "Shill Bidding in English Auctions," mimeo, University of Texas-Austin.

5 Supplementary Material for the Referees: The case of Uniformly distributed values

Our objective in this appendix is to identify the optimal secret and public reserve prices when F is the Uniform distribution and bidders arrive according to a Poisson process with arrival rate λ. We may then determine the value (to the seller) from a commitment to a public reserve. We also compute the social costs of the alternative auction formats, and the optimal auction length.

We first normalize $T = 1$ for convenience and write $v_0(r)$ for $v_0(r, 1)$. For the case of Uniformly distributed bidder valuations, $f(z) = 1$, $F(z) = z$. Hence, by using (1) we can show that $v_0(r)$ may be expressed as:

$$v_0(r) \;=\; \frac{e^{\lambda(1-r)}\left(1-\frac{2}{\lambda}\right) \;+\; \left(1-2r+\frac{2}{\lambda}\right)}{e^{\rho+\lambda(1-r)} \;-\; 1} \tag{21}$$

Straightforward calculations reveal that

$$v_0'(r) \;\gtrless\; 0$$

$$\Leftrightarrow \quad 2\delta e^{-\lambda(1-r)} \;\gtrless\; [2\delta - \lambda(1+\delta)] + 2\lambda r \tag{22}$$

(where $\delta \equiv \exp(-\rho)$).

5.1 Optimal public reserve

Consider the public-reserve-price scenario.

Proposition 11 *There is a unique optimal reserve in the public-reserve-price auction for any* $(\lambda, \delta) \in \mathbb{R}_{++} \times (0, 1)$*. This optimal reserve lies in* $(0.5, 1)$*, converging to* 0.5 *as* $\lambda \to 0$ *or* $\delta \to 0$*.*

Proof. Existence of an optimal public reserve, and the fact that any such reserve price is strictly less than 1, are guaranteed by Proposition 4. Uniqueness of the optimal public reserve follows from the facts that the left-hand side of (22) is convex in r, while the right-hand side is linear. The lower bound on the optimal reserve comes from equation (4) in Proposition 4, which implies

$$r^* \;>\; 1 - r^*$$

and hence $r^* > 0.5$.

Letting $r = 0.5$ in (22) gives

$$\frac{2\delta}{e^{\lambda/2}} \gtrless (2-\lambda)\delta \tag{23}$$

Both sides of this expression converge to 2δ as $\lambda \to 0$, or to 0 as $\delta \to 0$. This proves that $r = 0.5$ is optimal in each of these limiting cases, since we have already established the existence of a unique interior maximizer of $v_0(r)$. □

In a one-period version of this auction game, it is well-known that $r = 0.5$ is the optimal (enforceable) reserve for any $\lambda > 0$ (i.e. for any number of bidders). Hence, if $\delta \to 0$ the optimal reserve must converge to $r = 0.5$. A similar result obtains in a model in which the same bidders vie for the object in each round – see McAfee and Vincent (1997, p.250). However, in their model, the optimal reserve in period 1 converges to 0.5 *from below* as $\delta \to 0$ (*ibid.*, Figure 1(b)). When facing the same set of bidders each period, the inability to commit *not* to re-auction the object places downward Coasian pressure on the reserve.

If $\delta > 0$ and re-auction is possible, the seller has an incentive to raise her reserve above $r = 0.5$, since she receives a positive value even if the object is passed in. However, as the expected number of bidders in any given round goes to zero, the value of this re-auction option becomes negligible, so the optimal reserve again converges on $r = 0.5$.

Conversely:

Proposition 12 *For any $\delta \in (0,1)$, the optimal public reserve converges to $r = (1+\delta)/2$ as $\lambda \to \infty$; while for any $\lambda > 0$, the optimal public reserve converges to $r = 1$ as $\delta \to 1$.*

Proof. The left-hand side of (22) goes to zero as $\lambda \to \infty$. The right-hand side will only go to zero if $r \to (1+\delta)/2$.[15]

As $\delta \to 1$, (22) converges to the condition

$$e^{\lambda r} \gtrless e^{\lambda}\left[1 - \lambda\left(1 - r\right)\right].$$

The left-hand side of this expression is strictly increasing and strictly convex, rising to e^{λ} as $r \to 1$. The right-hand side is linear and strictly increasing in r, also rising to e^{λ} as $r \to 1$. When $r = 1$

[15] This puts an upper bound on the optimal reserve price (for given δ), since it is intuitive that r is increasing in λ. Indeed, if $r > (1+\delta)/2$, then the right-hand side of (22) strictly exceeds 2δ (for any $\lambda > 0$), while the left-hand side is no greater than 2δ for any $r \leq 1$ and $\lambda > 0$.

both sides have slope λe^{λ}. Therefore, the two sides are equal if and only if $r = 1$. □

Contrast this result with McAfee and Vincent (1997, p.251 and Figure 1(d)). In their model, with values drawn from the Uniform distribution on $[0, 1]$, the optimal first period reserve also rises with the number of bidders, n, but converges to the static solution, $r = 0.5$. In our model, a higher arrival rate λ augments the incentive to pass-in the object and search for higher value bidders. This reduces the pressure on the seller to sell in any given round, and pushes the reserve further and further above its optimal one-shot value. In the McAfee and Vincent model, with the same bidders each period, increasing the number of bidders raises the reserve for a quite different reason. Higher n raises the pressure on bidders to bid early, allowing the seller to raise her reserve. In the limit, the competitive pressure on bidders completely overcomes the Coasian dynamic, and the optimal one-shot reserve is achieved.

In contrast to our Proposition 12, McAfee and Vincent (1997, p.251 and Figure 1(b)) obtain that the optimal reserve declines with increases in δ in their model.

5.2 Optimal secret reserve

Let us now turn to properties of the secret-reserve auction. The first observation is the following:

Proposition 13 *There exists a unique equilibrium reserve in the secret-reserve auction for any* $(\lambda, \delta) \in \mathbb{R}_{++} \times (0, 1)$.

Proof. Use $v_0 = r$ in (21) and re-arrange the resulting expression to get

$$e^{\lambda}\left[z + 2\delta - \delta\lambda\right] \; = \; \delta e^{z}\left[2 + \lambda - z\right] \tag{24}$$

where $z = \lambda r$. The left-hand side of (24) is linear and strictly increasing in z, while the right-hand side is strictly increasing and strictly convex. Hence there will be at most two solutions to (24) in $z \in [0, \lambda]$. A necessary and sufficient condition for two solutions is that the LHS≥RHS at $z = 0$ and $z = \lambda$. The latter is always the case, but the former is so if and only if

$$e^{\lambda} \; \geq \; \frac{2 + \lambda}{2 - \lambda} \tag{25}$$

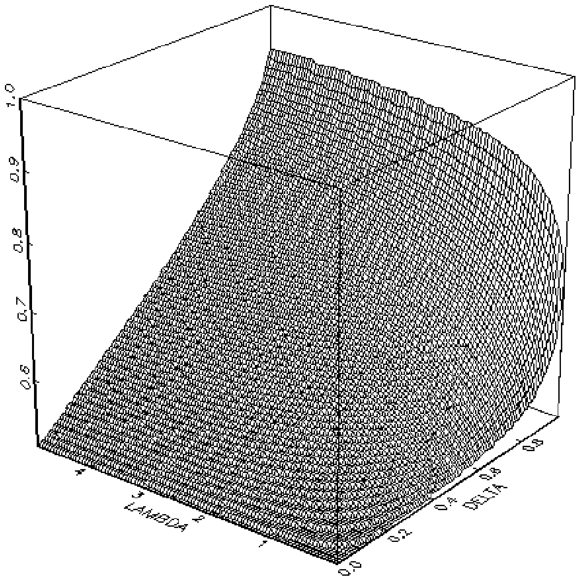

Figure 2: Optimal public reserve price

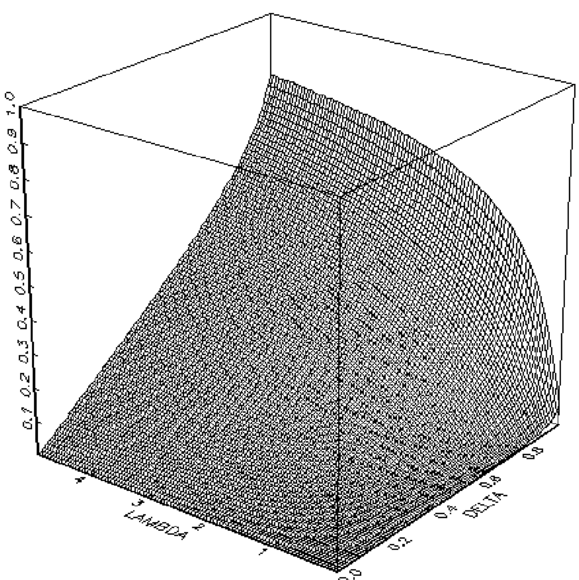

Figure 3: Equilibrium secret reserve price

In fact, one may verify that (25) fails for any $\lambda > 0$. Hence, the equilibrium secret reserve is unique. $\square$

Proposition 4 implies that the optimal secret reserve will be strictly smaller than the optimal public reserve. Furthermore, it is immediate from (24) that:

Proposition 14 *For any $\lambda > 0$, the equilibrium secret reserve $r \to 0$ as $\delta \to 0$.*

This is just convergence to the optimal secret reserve in the one-shot auction.

5.3 Numerical results

5.3.1 The case $T = 1$

Although explicit solutions for the optimal secret and public reserves are elusive, we may easily obtain them numerically. Figures 1 and 2 give the optimal reserves for a wide range of (δ, λ) values.[16]

The figures illustrate the behavior of the reserves for limiting values of the parameters. In relation to Proposition 11, we observe that convergence of the optimal public reserve to 0.5 as $\lambda \to 0$ is comparatively slow. Figure 1 uses λ values that are bounded below by $\lambda = 0.1$. One verifies Proposition 11 by taking λ values down to $\lambda = 0.005$, as in Figure 3.

Similarly, Proposition 12 is also obscured by the scale Figure 1. Indeed, it is clear that Propositions 11 and 12 imply "extreme" behavior of the surface near $(\delta = 1, \lambda = 0)$, with convergence to a non-smooth contour.

Figure 4 plots $r^P - r^S$, where r^P denotes the optimal public reserve, and r^S is the equilibrium secret reserve price. This is the increment in the reserve as a result of public announcements binding the seller.

Figure 5 illustrates the value of this commitment to a public reserve. It plots the gain from commitment, as a percentage of the secret reserve value. That is:

$$100 \left[\frac{v_0 \left(r^P \right) - v_0 \left(r^S \right)}{v_0 \left(r^S \right)} \right].$$

[16] Recall that $\delta = e^{-\rho}$.

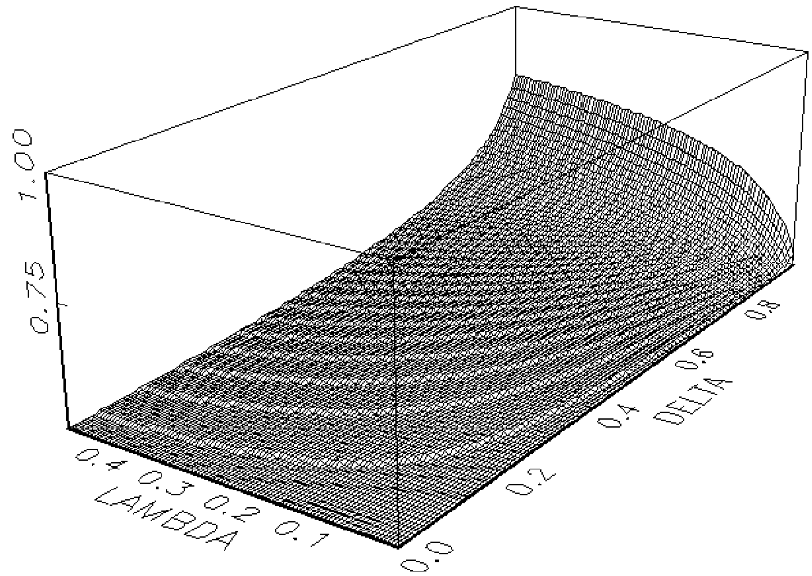

Figure 4: Optimal public reserve for low λ values

We observe that this value falls sharply as the average number of arrivals per period, λ, rises. The percentage gain is below 10% (for any δ) before λ reaches 5. Gains also tend to diminish with rises in δ, though much more slowly. In this sense, the inability to commit not to re-auction the object does devalue the commitment to a reserve price, but we do not see the dramatic decline in value observed by McAfee and Vincent (1997) in the case of long-lived bidders.

Since public reserves are higher than secret reserves, on average it will take longer to sell an object using the public-reserve-price format. The probability of concluding a sale in any given period is $1 - \Pi_1(r)$, which is equal to $1 - e^{-\lambda(1-r)}$ in the Uniform case. Let $p(r, \lambda)$ denote this quantity. It is decreasing in r for any λ, as one would expect. Figure 6 indicates the increment in this probability from using a secret-reserve-price auction: $p\left(r^S, \lambda\right) - p\left(r^P, \lambda\right)$.[17]

The expected number of periods until a sale is achieved is $p(r, \lambda)^{-1}$. Figure 7 indicates the expected extra delay (in numbers of periods) from using the public-reserve-price format rather than a secret-reserve-price auction: $p\left(r^P, \lambda\right)^{-1} - p\left(r^S, \lambda\right)^{-1}$.

[17] This difference depends on δ, as well as λ, since r^S and r^P depend on both parameters.

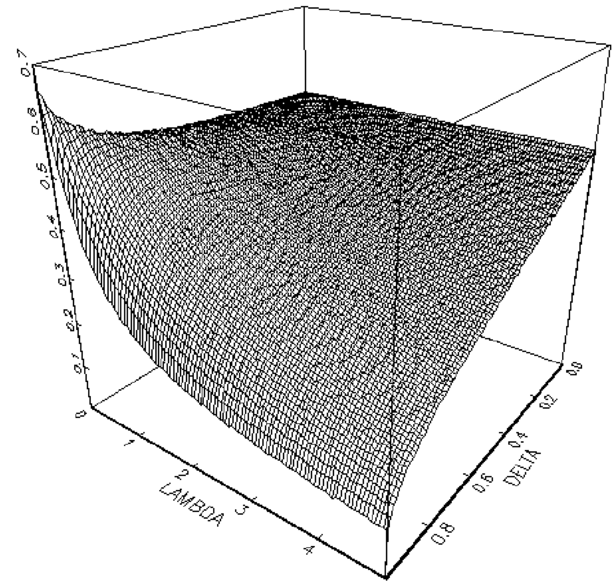

Figure 5: Public reserve less secret reserve

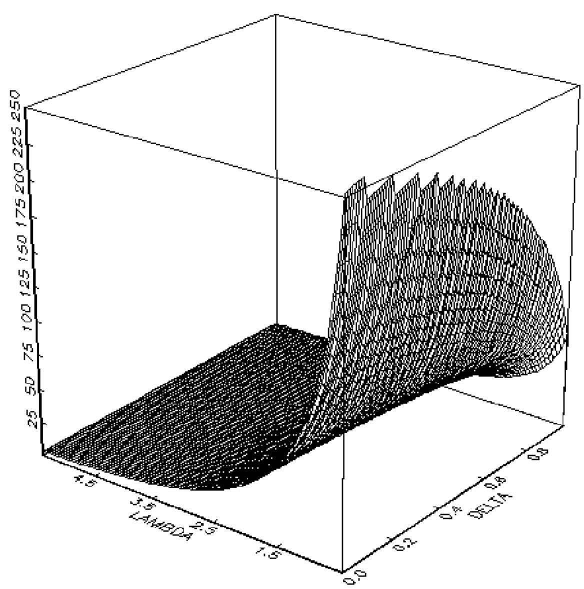

Figure 6: Value of commitment

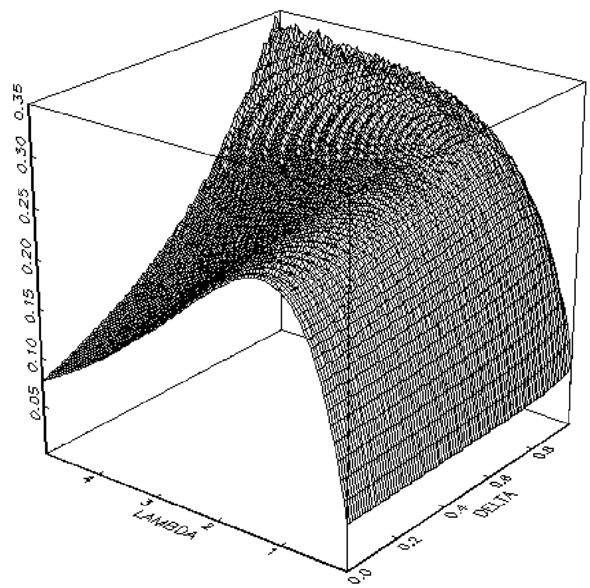

Figure 7: Extra per period sale probability from a secret reserve

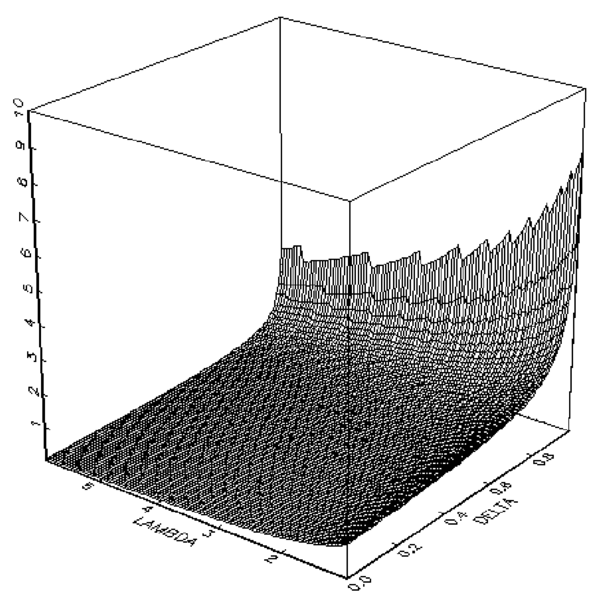

Figure 8: Expected extra delay from a public reserve

Finally, let us compare the optimal public reserve to the socially optimal reserve price. When $T = 1$ and F is Uniform on $[0, 1]$, the Social Planner's value function becomes

$$v_S(r) \;=\; \frac{\delta\left\{1 \;-\; re^{-\lambda(1-r)} \;-\; e^{-\lambda(1-r)}\left(r-\lambda^{-1}\right)\right\}}{1-\delta e^{-\lambda(1-r)}} \tag{26}$$

Figure 8 reveals that the optimal public reserve is too high. It graphs the difference between the optimal public reserve and the maximizer of (26).

Since we know (Proposition 6) that the equilibrium secret reserve price is too low, it is of interest to determine which auction format comes closest to achieving the socially optimal expected surplus. Are reserve price commitments socially desirable? Figure 9 provides the answer: it plots the difference between $v_S(r^S)$ and $v_S\left(r^P\right)$ (as a percentage of the latter). One observes that, for sufficiently high discount factors, the public reserve is relatively more efficient.

5.3.2 Endogenous T

For the special case in which F is Uniform, Figure 10 plots the value function $v_0(r, T)$ when $\rho = 0.1$ and $\lambda = 3$.

It is easy to observe that the optimal choice of (r, T) will have $T = 0$.[18] In fact, if we plot

$$\max_{r\in[0,1]} v_0(r, T)$$

[18] We obtain the same conclusion for all other parameter values that we have tried. In fact, one may show that

$$\left.\frac{dv_0\left(r^*(T), T\right)}{dT}\right|_{T\downarrow 0} < 0 \tag{27}$$

where $r^*(T)$ denotes the optimal public reserve given T. To do so, we first calculate

$$\begin{aligned} \frac{\partial v_0(r,T)}{\partial r} \;&=\; 0 \\ \Leftrightarrow \quad e^{-[\rho+\lambda(1-r)]T} \;&=\; e^{-\rho T} \;+\; \frac{\lambda T}{2}\left[2r \;-\; 1 \;-\; e^{-\rho T}\right] \end{aligned}$$

Now substutute this into the expression resulting from the calculation

$$\frac{\partial v_0(r,T)}{\partial T} \;\gtrless\; 0,$$

divide through by T^2 and simplify to obtain

$$\begin{aligned} &-\left[\rho+\lambda(1-r)\right]\left[\lambda \;+\; \lambda(1-2r)e^{-\lambda T(1-r)} \;+\; \frac{2\left(e^{-\lambda T(1-r)}-1\right)}{T}\right] \;\gtrless \\ &\left[\frac{e^{\rho T}-1}{T} \;-\; \frac{\lambda}{2}e^{-\rho T} \;+\; \lambda r\right]\left[\lambda(1-r)(\lambda T-2) \;-\; \frac{2\left(e^{-\lambda T(1-r)}-1\right)}{T}\right] \end{aligned}$$

Taking limits as $T \to 0$ we get

$$-\left[\rho+\lambda(1-r)\right]2\lambda \;\gtrless\; 0$$

Since $\rho + \lambda(1 - r) > 0$, (27) follows.

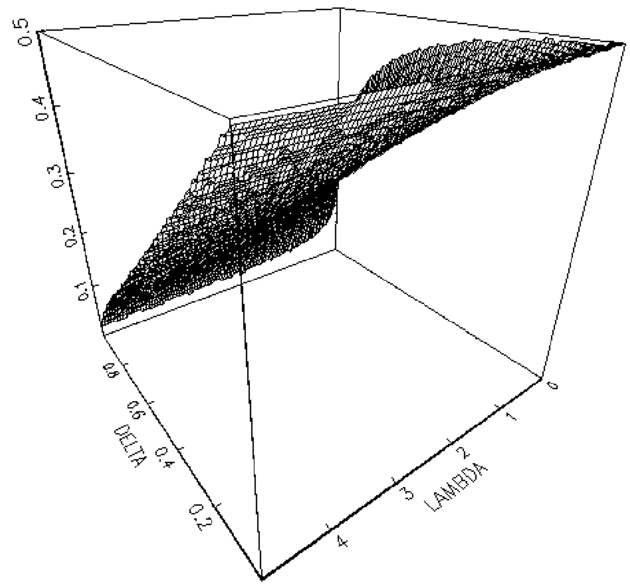

Figure 9: Optimal public reserve less socially optimal reserve.

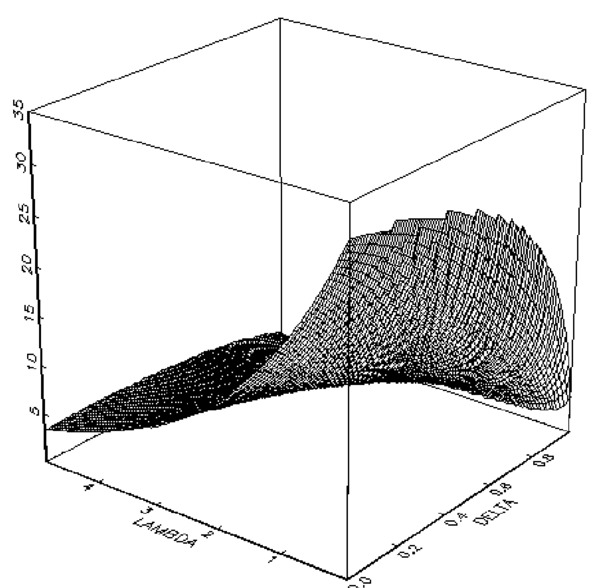

Figure 10: Social cost of commitment: $100\left[v_S\left(r^S\right)-v_S\left(r^P\right)\right]/v_S\left(r^P\right)$

as a function of T (again assuming that values are Uniformly distributed, $\rho = 0.1$ and $\lambda = 3$) we obtain Figure 11. Furthermore, straightforward calculations reveal that $v_0(r, 0)$ is maximized when

$$r = (1 + \theta) - \sqrt{\theta(1 + \theta)} \tag{28}$$

(where $\theta = \rho/\lambda$). Hence, with Uniformly distributed valuations, the optimal public-reserve-price auction has $T = 0$ and reserve price (28). We therefore have a complete characterization of the optimal auction (within the scenarios considered here) for the Uniform case.

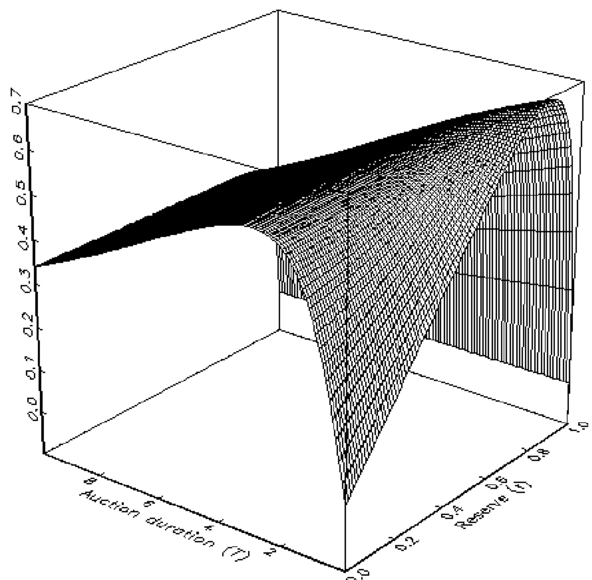

Figure 11: $v_0(r, T)$ when F Uniform, $\rho = 0.1$ and $\lambda = 3$

The optimality of $T = 0$ in a public-reserve-price-auction should be interpreted with care. It relies on our assumption that there is no delay *between* auction rounds, and no other fixed (monetary) cost to running another auction. In reality, at least one, and probably both, of these assumptions will be violated. For example, eBay charges sellers a small fee for each auction they run. Time or other costs of running many auctions will provide pressure to increase T.

On the other hand, the optimality of price-posting (in the presence of commitment) does not undermine the relevance of the model for the analysis of internet auctions. On eBay, for example,

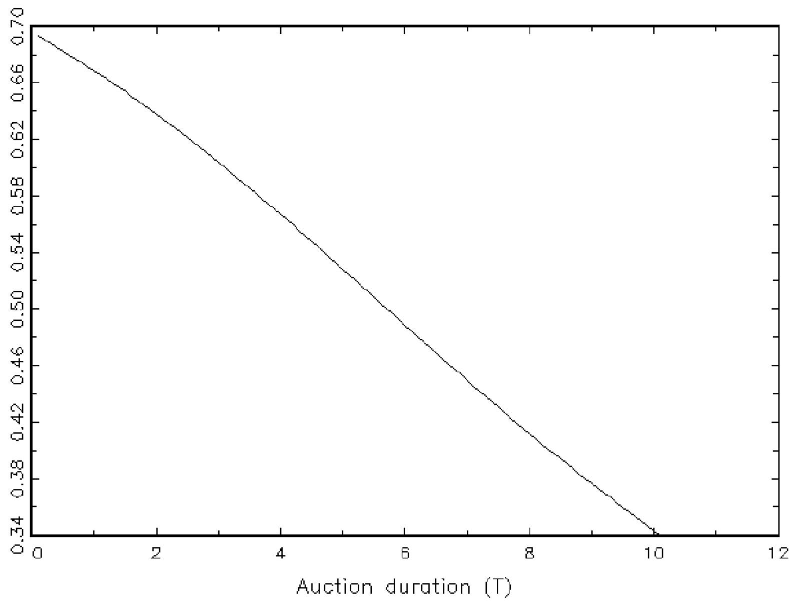

Figure 12: Auction value with a public reserve

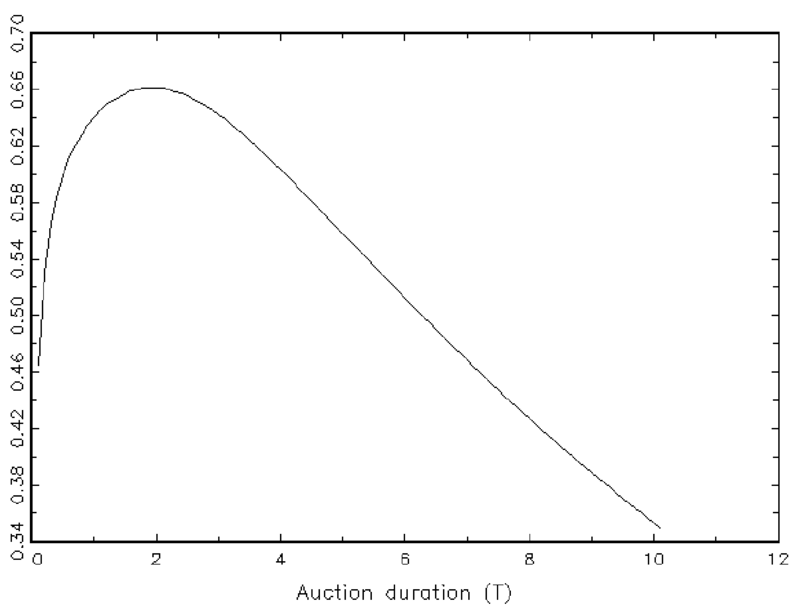

Figure 13: Auction value with a secret reserve

sellers have the option of posting a "Buy It Now" price. In the absence of a standing bid (above the reserve), if an arriving bidder offers the "Buy It Now" price, it is automatically accepted and the auction is cancelled.

Figure 12 illustrates $v_0(r(T), T)$, the value of a secret-reserve-price auction as a function of T. This has a maximum at an auction length around 2.