THE JOURNAL OF INDUSTRIAL ECONOMICS 0022-1821
Volume LXV June 2017 No. 2

COMMON VALUES AND LOW RESERVE PRICES*

Daniel Quint†

I show that the benefit of a high reserve price in a common-values ascending auction is lower than in the observationally equivalent private values setting. Put another way, when bidders have common values, empirical estimation based on a private-values model will overstate the value of a high reserve price. Via numerical examples, I show this same ranking typically applies to the level of the optimal reserve price as well, and often to the benefit of *any* reserve price, not just high ones. With common values, the optimal reserve can even be *below* the seller's valuation, which is impossible with private values.

I. INTRODUCTION

The workhorse model in auction theory – for both theoretical and empirical work – is the Symmetric Independent Private Values model. Among other things, this model suggests that a seller can always benefit from employing a reserve price strictly higher than her residual value for the unsold good, and often suggests that the profit-maximizing reserve price is fairly high – often much higher than reserve prices employed in practice. (Paarsch [1997] and Haile and Tamer [2003] are two examples.)

Subsequent theoretical work has shown, however, that a number of different deviations from the assumptions of the IPV model argue in favor of lower reserve prices. Correlation among bidder valuations (Quint [2008], Aradillas-López Gandhi and Quint [2013]), uncertainty about the exact value distribution (Kim [2013]), endogenous participation (Levin and Smith [1994], Samuelson [1985]), and competition between sellers (Peters and Severinov [1997]) have all been shown to reduce the reserve price that should be chosen by a profit-maximizing seller relative to the benchmark model – in some cases all the way down to the seller's residual value.

All of these, however, maintain the assumption of private values. In this paper, I compare the effect of a reserve price in ascending (English) auctions when bidders have private values, to its effect when bidder values are common, and bidders therefore face a winner's curse.

*Parts of this work were done while I was a visitor at eBay Research Labs in San Jose, California, and at the European University Institute in Florence, Italy. I am thankful to ERL and EUI for their support; to Dennis McWeeny for excellent research assistantship; and to Tom Blake, Dominic Coey, Peter Coles, Brad Larsen, Dimitriy Masterov, Alan Sorensen, Kane Sweeney, Steve Tadelis, Barr Yaron, the Editor, and two anonymous referees for helpful comments.
†Author's affiliation: Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, Wisconsin, U.S.A.
e-mail: dquint@ssc.wisc.edu

Part of the challenge of this exercise is finding the right *ceteris paribus* comparison. In a common values setting, bidders' valuations are inherently correlated, so comparing to a model with *independent* private values is potentially misleading. Instead, for a given common values environment, I define the private-values setting that would lead to identical bidding behavior in the absence of a reserve price.

I find that for an English auction in *any* interdependent or common values setting and a sufficiently high reserve price, both the likelihood of a sale and the expected revenue or profit are lower than under the analogous private values setting. (The same holds for *any* positive reserve price in a sealed-bid second-price auction.) I also explore a number of numerical examples, and find that the same revenue and profit rankings very often hold even outside of the range of reserves where the result is theoretically guaranteed, and that the optimal reserve price is similarly lower than it would be under private values.

My choice of comparison is meant to mirror the choice facing an empirical researcher. Given bid data from previous (reserve-free) auctions, the researcher could choose to rationalize the data via either a private or a common-values model. My results suggest that when the true environment has common values but the researcher mistakenly assumes it has private values, her counterfactuals will be biased in a particular direction: they will overstate the benefit of any high reserve price (and likely any reserve price), and likely overstate the level of the profit-maximizing reserve.[1]

One key feature of common value auctions, unlike private value auctions, is that bidders learn from each others' bidding. This overturns the usual private-values result that reserve prices below the seller's own valuation v_0 are always dominated. Under common values, this need not be true, as a reserve price of v_0 may truncate losing bids and thus reduce revenue from profitable sales. It's even possible that two bidders who would not bid at a reserve of v_0, would both bid at a lower reserve and then, seeing each other bidding, would both bid past v_0, creating a profitable sale that would have been prevented by a reserve of v_0. Thus, introducing a reserve price in an interdependent or common values setting has greater costs than it would in a similar private values setting, and these costs often seem to outweigh the benefits entirely. For many of the numerical examples I show in the paper, a reserve price of 0 maximizes revenue; when the seller's valuation v_0 is fixed and positive, the profit-maximizing reserve price is sometimes *below* v_0, sometimes even 0. When the seller's value is positively associated with the buyers', this introduces an additional winner's curse-type effect – the seller is more likely to retain the object when it is less valuable to her – which further decreases the value of a reserve price, and makes the optimal reserve price 0 in many cases. On the whole, relative to intuitions we have

[1] The private values environment corresponding to a given common values setting will have affiliated values, so Quint [2008] implies that if the researcher assumed a common values setting actually had *independent* private values, the bias would be in the same direction, and even larger.

from private values settings, when bidder values are common or interdependent, this suggests that reserve prices should be used much more cautiously, if at all.

II. CLOSELY RELATED LITERATURE

Laffont and Vuong [1996] observe that in sealed-bid auctions, common and private values cannot be distinguished from one another purely from bid data, as the two models are observationally equivalent. I use their logic in comparing a common values setting to the corresponding private values setting that would lead to identical bidding.

Perhaps as a result of Laffont and Vuong's finding, most of the empirical literature on auctions does not attempt to differentiate one model from the other empirically, and instead begins by assuming either one model or the other, but there are a few exceptions. Paarsch [1992] shows that the two models are distinguishable from each other under parametric distributional assumptions. For first-price auctions, Haile, Hong and Shum [2003] propose a test of common versus private values when there is exogenous variation in the number of bidders N, and Athey and Haile [2007], using ideas from Hendricks, Pinkse and Porter [2003], propose a test when there is a binding reserve price. For English auctions, Athey and Haile [2002] propose a test when there is variation in N, but also note that when N is fixed and values are known to be common, the model is not identified from observed bids.

Vincent [1995] shows that in a common values setting, if the seller's valuation is unknown but independent of the bidders' valuation, the seller can sometimes benefit *ex ante* from using a *secret* reserve price. (The seller would still be tempted to deviate and announce the reserve price when it is low, however, and therefore needs to be able to commit to keeping it secret.) Several papers noted in the introduction show deviations from the standard IPV model (while maintaining the assumption of private values) which favor lower reserve prices than the IPV benchmark.

III. SECOND PRICE AUCTIONS

While the focus of this paper is English (or ascending) auctions, much useful intuition can be gained from considering a simpler case: sealed-bid second price auctions. In this section, I show how second-price auctions under two different models of bidder valuations – one with interdependent values, and one with private values – which would generate the same bidding behavior in the absence of a reserve price, respond differently to the addition of a positive reserve price. In the next section, I will do the same for English auctions.

III(i). *Two Models of Valuations*

Model I – Interdependent Values. I use the standard interdependent values model with affiliated signals of Milgrom and Weber [1982]. Fix N the number of bidders. Let $X = \{X_1, \ldots, X_N\}$ be a set of signals drawn from a joint distribution which is symmetric and affiliated. Each bidder i learns the realization x_i of one signal X_i, but his valuation

$$V_i = u_i(X) = u\left(X_i, \{X_j\}_{j \neq i}\right)$$

depends (symmetrically) on the signals observed by the other bidders as well. As in Milgrom and Weber, $u : \Re^N \to \Re$ is assumed to be nonnegative, continuous, nondecreasing, strictly increasing in its first argument, and symmetric in its last $N-1$ arguments, with $\mathbb{E}u_i(X) < \infty$. Since the point of this paper is to contrast the interdependent case with the private values case, I assume u_i has nondegenerate dependence on $\{X_j\}_{j \neq i}$, and will often refer to this model as the common values case, although pure common values $(u_i(X) = u_j(X))$ is not assumed.

For a given bidder i, define $v(x, y)$ as the expectation of his value V_i, conditional on the realization of his own signal and the highest signal of his opponents,

$$v(x, y) = \mathbb{E}\left\{u\left(x, \{X_j\}_{j \neq i}\right) \middle| X_i = x,\ \max_{j \neq i} X_j = y\right\}$$

(Note that the expectation over $\{X_j\}_{j \neq i}$ is taken conditional on $X_i = x$, since the signals may be correlated.) In the absence of a reserve price, Milgrom and Weber establish that bidding in the symmetric equilibrium of the second price auction is

$$b_i(X_i) = v(X_i, X_i)$$

It's straightforward to show that if one's opponents play this strategy, a bidder could do no better even if he knew the highest bid submitted by his opponents.[2]

Model II – Private Values. Laffont and Vuong [1996] observe that in second-price auctions without a reserve price, bid data cannot be used to distinguish common values from private values, because however bidders chose to bid under common (or interdependent) values, they could alternatively have been bidding their valuations in a private-values setting. Using this logic, we define the private-values setting that is 'observationally equivalent' to the interdependent-values setting above: the signals $\{X_i\}$ have the same distribution as above, each bidder observes a single signal X_i, but his valuation is now

$$V_i = v(X_i, X_i)$$

[2] The second price auction also has asymmetric equilibria, including some 'collusive-looking' equilibria with low revenue; like the literature, I focus on the symmetric equilibrium, which is unique.

regardless of the realization of the other signals. Without a reserve price, bidders have a weakly dominant strategy of bidding their values, and therefore $b_i(X_i) = v(X_i, X_i)$, so the two models produce identical bidding behavior.

III(ii). *The Effect of a Reserve Price*

The effect of introducing a reserve price $r > 0$ in a private-values setting is straightforward. Bidders continue to bid their valuations, as long as those valuations are above r; if not, they do not bid. Thus, if no bidder has a valuation exceeding r, the reserve price prevents a sale, and the seller retains the object; if one bidder has a valuation exceeding r, then he pays r rather than the second-highest valuation.

In a common (or interdependent) values setting, the effect of a reserve price is more complicated. In particular, *the highest reserve price a bidder is willing to meet is lower than the bid he would make in the absence of a reserve*. This is because of the winner's curse: a bidder willing to meet a reserve price r must be willing to win the object even if no other bidder is willing to bid r. On the other hand, when contemplating raising ones bid, a bidder conditions on his bid's being pivotal, i.e., on another competitor having an equally high signal, partly mitigating this curse.

Thus, in a second-price auction with reserve price r and common values, bidder i still bids $v(X_i, X_i)$ if he bids, but only bids if $X_i \geq x^*$, where x^* is defined implicitly by

$$r = \mathbb{E}\left\{u\left(x^*, \{X_j\}_{j \neq i}\right) \middle| X_i = x^*, \ \max_{j \neq i} X_j < x^*\right\}$$

(A bidder with signal $X_i = x^*$ therefore loses whenever another bidder bids, and pays r for a prize whose expected value is r when no other bidder bids.) By iterated expectations, if we let $X^{(k)}$ denote the k^{th} order statistic of $\{X_1, \ldots, X_N\}$, the definition of x^* is equivalent to

$$r = \mathbb{E}\left\{v(x^*, X^{(2)}) | X^{(1)} = x^*\right\}$$

Since by definition $X^{(2)} \leq X^{(1)}$, the right-hand side is strictly less than $v(x^*, x^*)$ except in degenerate cases. Thus, a bidder willing to bid more than r in the absence of a reserve price, may still not be willing to bid at all with a reserve price of r.[3] On the other hand, in a private values setting, a bidder's willingness to bid (in the absence of a reserve price) is exactly the reserve price he would be willing to meet. This leads to the following results:

[3] This is the logic behind Vincent's [1995] result that a secret reserve price can sometimes be beneficial: when the seller's valuation is high and she therefore wants to set a high reserve price, two bidders might be deterred from bidding who might otherwise have bid above that price.

Theorem 1. Fix any common values setting and any reserve price $r > 0$. In a second price auction with reserve price r...

1. the likelihood of a sale is lower than in the corresponding private values setting
2. expected revenue is lower than in the corresponding private values setting
3. if r is greater than the seller's residual valuation v_0, expected profit is lower than in the corresponding private values setting

All theorems are proved in the supplemental materials available on the *Journal's* editorial website. Theorem 1 says that if you had bid data from second-price auctions in a setting with *common* values, but chose to estimate a structural model under the assumption of *private* values and used it to evaluate a reserve price counterfactual, you would underestimate the reduction in sales that would follow from introducing a reserve price, and therefore overestimate the benefit of introducing a reserve.

IV. ENGLISH AUCTIONS

Next, I turn to English auctions. It's well know that with private values, English and second-price auctions are strategically equivalent. With common values, however, this is not the case, since bidders update their beliefs about their own valuations based on how their opponents bid.

I'll employ the same strategy for English auctions as I did for second price auctions: begin with a general interdependent values setting, imagine I observed bids in English auctions without reserve prices, use those bids to define an observationally-equivalent private values setting, and then compare the impact of introducing a reserve price across the two settings. Three things make this exercise more complicated for English auctions:

1. Since an English auction ends before the winner reveals his willingness to pay, I'll need to make a decision about what his valuation should be in the private-values setting.
2. In a common values setting, bidders condition on each others' behavior, so bids (and therefore imputed private valuations) are 'more correlated' than information based on a single signal. This ends up implying that while *high* reserve prices are more likely to be met under the private values model, *low* but positive reserve prices may be more likely to be met under the common values model. (It will become clear why shortly.)
3. In a common values setting, a reserve price that does not set the price (because two or more bidders bid) still effects the price paid. If any losing bidders choose not to bid, their exact signals cannot be inferred by the bidder who eventually sets the price; via a 'linkage principle' effect, this

decreases expected price in the common values model, while it has no effect in a private values setting.

To address the first point, I imagine that after each auction, I was able to interview the winner and find out how high he had been planning to bid if his last opponent had not dropped out when he did, and use this as the winner's valuation in the corresponding private values setting. (This is data one might conceivably have in settings where bidders use automated proxy bids, such as on eBay.) To address the second, I will define a threshold reserve $\tilde{r}$ above which I can unambiguously sign the revenue and profit rankings between the two models. I will also use numerical examples to show that even when $\tilde{r}$ is high (and therefore the theoretical result is weak), the same revenue and profit rankings very often hold for reserves below $\tilde{r}$. The third point simply adds steps to the proof of the result.

IV(i). *Two Models of Valuations*

As in Milgrom and Weber [1982], I model English auctions as full information button auctions: the price starts low and rises continuously, bidders remain active until they choose to irreversibly drop out of the bidding, and bidders know who is currently active and at what price inactive bidders dropped out.

Model I – Interdependent Values. The model is the same as above. N bidders receive affiliated signals $\{X_i\}$ and have valuations $V_i = u(X_i, \{X_j\}_{j \neq i})$. Symmetric equilibrium bidding is described by Milgrom and Weber [1982, sections 5 and 7].[4] In the absence of a reserve price, equilibrium bidding can be briefly summarized as follows. At each point in the auction, the signals of the bidders who have already dropped out are correctly inferred by the remaining bidders. Given those signals, each bidder bids up to the price at which he would be exactly indifferent about buying the object if all his remaining opponents turned out to have signals matching his own.

Model II – Private Values. To create the observationally equivalent private values environment, I imagine we observe equilibrium English-auction bidding (with no reserve price) under the first model – including the price at which the winner planned to drop out – but interpret each bidder's bid as his private value.[5] Formally, let $i(k)$ denote the label of the k^{th} highest signal, so that

[4] Bikhchandani, Haile and Riley [2002] show that a continuum of symmetric, separating equilibria exist, but they all lead to the same outcome. They also lead to the same values of $v_{i(1)}$ and $v_{i(2)}$ in the private values model described below, and therefore the same outcome in that model. Thus, the multiplicity is not important for our purposes, and I focus on the equilibrium described by Milgrom and Weber.

[5] An alternative assumption about the winner's private value would be to imagine we had an independent measure of the winner's *ex post* surplus and use that, i.e., set $v_{i(1)}$ below equal to $u_{i(1)}(X)$. Under that assumption, Theorem 2 below would still hold under one change: modifying

$X_{i(k)} = X^{(k)}$. Let $X^{(1)} \geq X^{(2)} \geq \cdots \geq X^{(N)}$ denote the order statistics of $\{X_i\}$, and $x^{(1)} \geq \cdots \geq x^{(N)}$ their realization. Then for a given realization x of X, I define bidder valuations as

$$\begin{aligned}
v_{i(N)} &= u\left(x^{(N)}, x^{(N)}, x^{(N)}, x^{(N)}, \ldots, x^{(N)}, x^{(N)}, x^{(N)}\right) \\
v_{i(N-1)} &= u\left(x^{(N-1)}, x^{(N-1)}, x^{(N-1)}, x^{(N-1)}, \ldots, x^{(N-1)}, x^{(N-1)}, x^{(N)}\right) \\
v_{i(N-2)} &= u\left(x^{(N-2)}, x^{(N-2)}, x^{(N-2)}, x^{(N-2)}, \ldots, x^{(N-2)}, x^{(N-1)}, x^{(N)}\right) \\
&\vdots \\
v_{i(2)} &= u\left(x^{(2)}, x^{(2)}, x^{(3)}, x^{(4)}, \ldots, x^{(N-2)}, x^{(N-1)}, x^{(N)}\right) \\
v_{i(1)} &= u\left(x^{(1)}, x^{(1)}, x^{(3)}, x^{(4)}, \ldots, x^{(N-2)}, x^{(N-1)}, x^{(N)}\right)
\end{aligned}$$

Each bidder learns the realization v_i of his own valuation V_i.

Observational Equivalence of the Two Models when $r = 0$. By construction, a bidder's private value v_i in the second model is exactly the price at which he would drop out on the equilibrium path under the first model in the absence of a reserve price. In addition, the winning bidder's private value in the second model is the price at which he would have planned to drop out in the first model, had the second-highest bidder continued bidding.[6] Thus, the two models generate the same equilibrium bids – and would still appear identical if we knew at what price the winner had planned to drop out.

IV(ii). *The Effect of a Reserve Price*

In the private values case, bidding in an auction with reserve price r is again straightforward: each bidder i bids if $v_i \geq r$, and drops out at price v_i.

In the common values case, the threshold signal x^* above which bidders are willing to bid in an English auction with reserve price r is the same as in a second price auction, and is defined (as above) by

$$r = \mathbb{E}\left\{u\left(x^*, \{X_j\}_{j \neq i}\right) \middle| X_i = x^*, \max_{j \neq i} X_j < x^*\right\}$$

the definition of $\tilde{r}$ to be the highest crossing point of the CDFs $F_{U^{(1)}}$ and $F_{R^{(1)}}$, rather than $F_{V^{(1)}}$ and $F_{R^{(1)}}$.

[6] As noted above, we could alternatively let the winner's private value match his *ex post* surplus from winning in the first; with a change in the definition of $\tilde{r}$, the results would go through.

(Once again, a bidder with the threshold signal expects to be outbid if anyone else bids, and therefore must be indifferent when he is the only bidder willing to meet the reserve price.) Since bidders with signals below x^* do not reveal their signals through bids, bidders who do bid condition only on $X_j < x^*$ for bidders who don't bid; otherwise, equilibrium bidding among those who bid is the same, with each bidder dropping out when the price reaches the level at which he would in expectation be indifferent to winning, conditional on the information revealed by the bidders who have already dropped out, if all his remaining active opponents had received the same signal as he.

Define $Rev_{CV}(r)$ and $Rev_{PV}(r)$ as the expected revenue under the two models, as a function of the reserve price r; and fixing the seller's residual valuation v_0, define $\pi_{CV}(r)$ and $\pi_{PV}(r)$ as expected profits under the two models. My main theoretical result will be that $Rev_{CV}(r) < Rev_{PV}(r)$ and $\pi_{CV}(r) < \pi_{PV}(r)$ when r is sufficiently high. Since the observational equivalence described above implies $Rev_{CV}(0) = Rev_{PV}(0)$ and $\pi_{CV}(0) = \pi_{PV}(0)$, this means the benefit of a high reserve price is less under common values than under private values – or when the 'truth' is common values, a private-values analysis will overstate the benefit of a high reserve price. After, I will explore several numerical examples, under which these rankings often hold at *all* reserve price levels, and optimal reserve prices are lower under common values as well.

First, I need to formalize what it means for r to be 'sufficiently high' for the results to hold. Consider three random variables,

$$V^{(1)} = u\left(X^{(1)}, X^{(1)}, X^{(3)}, \ldots, X^{(N)}\right)$$

$$U^{(1)} = u\left(X^{(1)}, X^{(2)}, X^{(3)}, \ldots, X^{(N)}\right)$$

$$R^{(1)} = \mathbb{E}_{X^{(2)},\ldots,X^{(N)}|X^{(1)}}\left\{u(X^{(1)}, X^{(2)}, X^{(3)}, \ldots, X^{(N)})\right\}$$

and let $F_{V^{(1)}}$, $F_{U^{(1)}}$, and $F_{R^{(1)}}$ denote their probability distributions.

Under the common values model, for a given realization $x^{(1)}$ of $X^{(1)}$, the expectation $\mathbb{E}\left\{u(x^{(1)}, X^{(2)}, X^{(3)}, \ldots, X^{(N)})|X^{(1)} = x^{(1)}\right\}$ is the highest reserve price at which a bidder with signal $x^{(1)}$ would be willing to bid under the common values model. Thus, $F_{R^{(1)}}(r)$ is the probability that no bidder would be willing to bid given a reserve price of r, and therefore the probability that a reserve of r would fail to be met under common values.

Under the private values model, $u(x^{(1)}, x^{(1)}, x^{(3)}, \ldots, x^{(N)})$ is the imputed valuation of the winning bidder, and thus the highest reserve price at which he would bid under the private values model. $F_{V^{(1)}}(r)$ is therefore the probability that a reserve of r fails to be met under the private values model.

What will be crucial for us is that $F_{V^{(1)}}(r) \leq F_{R^{(1)}}(r)$ for r sufficiently high. To see why, first note that $U^{(1)}$ is a mean-preserving spread around $R^{(1)}$, and thus, we would expect $F_{U^{(1)}}$ to typically be above $F_{R^{(1)}}$ at low values and below

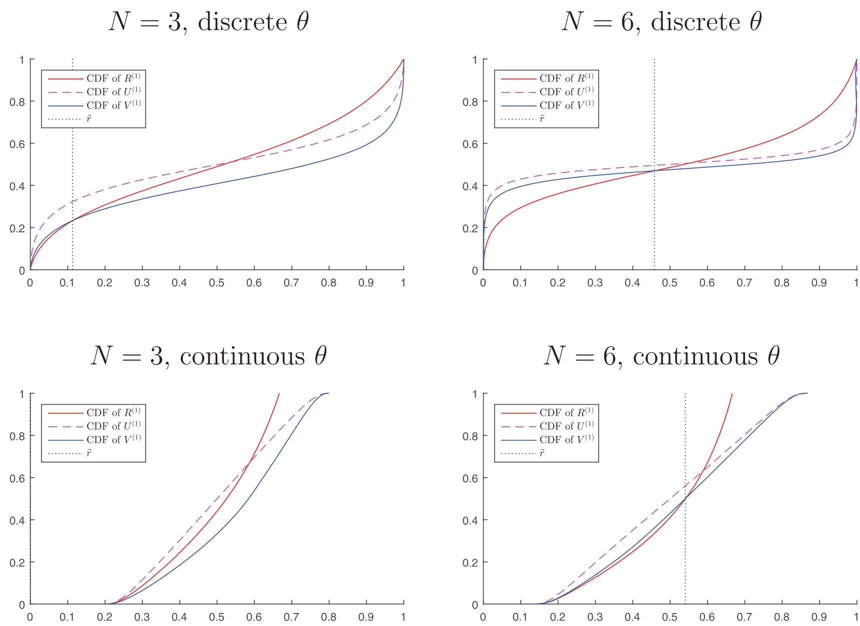

Figure 1
Determination of $\tilde{r}$ in Numerical Examples [Colour figure can be viewed at *wileyonlinelibrary.com*]

it at high values.[7] On the other hand, since $V^{(1)} \geq U^{(1)}$ for each realization of X, $F_{V^{(1)}} \leq F_{U^{(1)}}$ everywhere. Thus, we would expect $F_{V^{(1)}}(r) \leq F_{R^{(1)}}(r)$ for high values of r, while the rankings for low r are ambiguous. I define $\tilde{r}$ as the highest crossing point of these two CDF's; that is, $\tilde{r}$ is defined such that $F_{V^{(1)}}(r) \leq F_{R^{(1)}}(r)$ for every $r > \tilde{r}$.

Definition 1. $\tilde{r} \equiv \sup\left\{r : F_{V^{(1)}}(r) > F_{R^{(1)}}(r)\right\}$.

Figure 1 illustrates the CDF's of $V^{(1)}$, $U^{(1)}$, and $R^{(1)}$, and therefore how $\tilde{r}$ is defined, for the two numerical examples from the next section: the 'discrete θ' example presented in Section V(i), and the 'continuous θ' example presented in Section V(iv), each with $N = 3$ and $N = 6$. In each case, it's easy to see that $F_{U^{(1)}}$ (the dashed line) crosses $F_{R^{(1)}}$ (the lighter solid line) just once, near the middle of the support, and that $F_{V^{(1)}} < F_{U^{(1)}}$ (the darker solid line) is below the dashed line everywhere. In the two examples with $N = 3$, the difference between $F_{V^{(1)}}$ and $F_{U^{(1)}}$ (the impact of changing one bidder's signal) is large, making $F_{V^{(1)}}$ below $F_{R^{(1)}}$ on most of the range of valuations. (In the bottom-left pane, $F_{V^{(1)}}$ and $F_{R^{(1)}}$ never cross, so $\tilde{r} = 0$.) In the two examples with $N = 6$, one bidder's signal has less effect, so $F_{V^{(1)}}$ and $F_{U^{(1)}}$ are closer together, making $\tilde{r}$ larger.

[7] While $F_{R^{(1)}} \succ_{SOSD} F_{U^{(1)}}$ does not guarantee that the two CDF's cross only once, they do in all the numerical examples I've examined.

For reserve prices above $\tilde{r}$, revenue and profit rankings across the two models are clear:[8]

Theorem 2. Fix any common values setting and any reserve price $r \geq \tilde{r}$. In an English auction with reserve price r...

(i) the likelihood of a sale is lower than in the corresponding private values setting
(ii) expected revenue is lower than in the corresponding private values setting
(iii) if $r \geq v_0$, expected profit is lower than in the corresponding private values setting

Of course, the applicability of Theorem 2 depends on how restrictive the assumption $r \geq \tilde{r}$ is. In the next section, I examine a range of numerical examples, solved via simulation, to illustrate the level of $\tilde{r}$ as well as the profit curves $\pi_{CV}(r)$ and $\pi_{PV}(r)$. As Figure 1 suggests, when N is small, $\tilde{r}$ tends to be close to 0, and so Theorem 2 is quite informative. (When $N=2$, for example, $\tilde{r}$ is always 0.) When N is large, $\tilde{r}$ tends to be high; in those instances, Theorem 2 only applies to very high reserve prices, which may be unrealistic unless the seller's own valuation v_0 is quite high.

However, the numerical examples also show that the revenue and profit rankings of Theorem 2 very often hold even for r below $\tilde{r}$ – in most examples, for *all* r. We will also see that the profit-maximizing reserve price is typically lower under a common values model than under the corresponding private values model. In fact, with common values, the profit-maximizing reserve price is sometimes below the seller's own valuation – which is impossible with private values.

V. NUMERICAL EXAMPLES

In this section, I offer two numerical examples in which I explicitly calculate the level of $\tilde{r}$, compare $\pi_{CV}(r)$ to $\pi_{PV}(r)$, and compare the optimal reserve prices under the two models. In both examples, bidders have pure common values. (In Section VI(iii), I'll consider another example where valuations are a mix of common and private.) There is an underlying 'state of the world' θ which is every bidder's *ex post* valuation, and bidder signals are *i.i.d.* conditional on the value of θ.[9]

[8] Note that here, as in Theorem 1, I compare auctions across the two models of valuations at the same reserve price, rather than each at its own optimal reserve. This is because I imagine a seller who is unsure which is the correct model. Note also that if $r^*_{CV} \equiv \arg\max_r \pi_{CV}(r) \geq \tilde{r}$ (or in a second-price auction), optimality implies $\max_r \pi_{PV}(r) \geq \pi_{PV}(r^*_{CV}) \geq \pi_{CV}(r^*_{CV}) = \max_r \pi_{CV}(r)$.

[9] In both examples in this section, the conditional distribution of bidder signals is increasing in θ via the strict MLRP, so bidder signals are affiliated, and $E(\theta|X)$ is increasing in X. Thus, if we think of valuations as $V_i = E(\theta|X)$ instead of $V_i = \theta$, both examples fit within the Milgrom-Weber framework.

V(i). *Setup*

For the first example, θ takes the values 0 and 1, with equal probability, and conditional on θ, bidder signals $\{X_i\}$ are $i.i.d.$ draws from a distribution $F(\cdot|\theta)$ on [0, 1], where

$$f(s|\theta) = \begin{cases} \alpha s^{\alpha-1} & \text{if } \theta = 1 \\ \alpha(1-s)^{\alpha-1} & \text{if } \theta = 0 \end{cases} \quad \text{and}$$

$$F(s|\theta) = \begin{cases} s^{\alpha} & \text{if } \theta = 1 \\ 1-(1-s)^{\alpha} & \text{if } \theta = 0 \end{cases}$$

Note that for $\alpha > 1, f(\cdot|1)$ is increasing and $f(\cdot|0)$ is decreasing, so lower signals are more likely when $\theta = 0$ and higher signals when $\theta = 1$. A higher value of α means $f(\cdot|1)$ is more skewed toward high signals and $f(\cdot|0)$ more skewed toward low signals, so the signals are more informative about θ (and also more highly correlated with each other).

For the common values model, the bidders all have valuation θ. The corresponding private values model is defined as in the previous section.

V(ii). *Effect of Reserve Price on Bidding Behavior*

For an illustration of how a reserve price effects bidding under the two models, we first consider bidding given a fixed realization of signals. Suppose that $N = 5$ and $\alpha = 2$, and that the realized signals are $X_1 = 0.82$, $X_2 = 0.65$, $X_3 = 0.50$, $X_4 = 0.35$, and $X_5 = 0.18$. (These are approximately the unconditional medians of each order statistic.[10]) In an English auction with no reserve price, the first three bidders would drop out at prices

$$u(0.18, 0.18, 0.18, 0.18, 0.18) \approx 0.0005$$
$$u(0.35, 0.35, 0.35, 0.35, 0.18) \approx 0.0181$$
$$u(0.50, 0.50, 0.50, 0.35, 0.18) \approx 0.1057$$

with the remaining bidders correctly inferring the values of X_5, X_4, and X_3 as these bidders dropped out; bidder 2 would then drop out at $u(0.65, 0.65, 0.50, 0.35, 0.18) \approx 0.2896$, and bidder 1 would be pleased, having been prepared to bid up to $u(0.82, 0.82, 0.50, 0.35, 0.18) \approx 0.7104$ had bidder 2 kept going that long. Thus, for the analogous private values setting, I consider $V_1 = 0.71$ and $V_2 = 0.29$.

[10] Of course, this particular realization of signals is relatively unlikely – more typically, either all five signals will be higher (because $\theta = 1$), or all five will be lower (because $\theta = 0$). Still, this realization makes for a nice illustration of the effect of r.

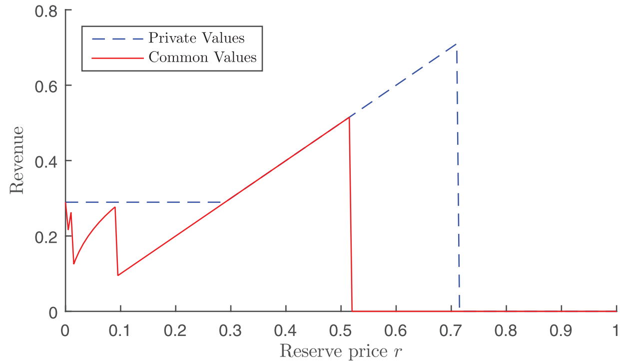

Figure 2
Revenue at 'Median' Signal Realizations as a Function of r (N = 5, α = 2) [Colour figure can be viewed at *wileyonlinelibrary.com*]

Figure 2 contrasts the effect of adding a reserve price under the two models, given this particular set of realized signals. The dashed curve shows revenue under the private values model. When $r < 0.29$, bidders 1 and 2 both bid, so the reserve price does not bind; revenue equals bidder 2's bid, which is his valuation, 0.29. For r between 0.29 and 0.71, bidder 1 bids, and pays the reserve price. When $r > 0.71$, nobody bids, and revenue is 0.

The solid curve shows revenue under the common values model, where the reserve price has more complicated effects. The highest reserve at which each bidder is willing to bid are $\mathbb{E}\left\{u(X)|X^{(1)} = x_i\right\} \approx 0.00002, 0.001, 0.01, 0.09$, and 0.52. So for $r \in [0, 0.00002)$, all five bidders bid, and the outcome is the same as with no reserve. For $r \in [0.00002, 0.001)$, however, bidder 5 does not bid, and rather than learning X_5, the others only learn that $X_5 < x^*$. For $r \in [0.001, 0.01)$, bidders 4 and 5 do not bid; for $r \in [0.01, 0.09)$, bidders 3, 4 and 5 do not bid. (Within each of these intervals, revenue is increasing in r, since x^* increases with r, and so the inference made about the 'truncated' signals becomes less negative.) For $r \in [0.09, 0.52)$, only bidder 1 bids, so he pays the reserve price; for $r \geq 0.52$, nobody bids and revenue is zero.

For this particular realization of signals, every reserve price r gives weakly less benefit under the common values model than under the analogous private values setting. This need not always be the case – for some signal realizations, certain reserve prices are more beneficial under common values – but Theorem 2 says that at least for reserve prices above $\tilde{r}$, those signal realizations which are less favorable under common values will dominate in expectation.

This example also demonstrates the magnitude of the winner's curse. Bidder 2 received a signal $X_2 = 0.65$, which is nearly twice as likely when $\theta = 1$ as when $\theta = 0$. Nevertheless, when considering whether to bid at a given reserve price, he worries about winning when all other bidders had signals low enough to not bid. Although $E(\theta|X_2 = 0.65) = 0.65$, $E(\theta|X_2 = 0.65, \{X_1, X_3, X_4, X_5\} < 0.65) \approx 0.09$; so while bidder 2's best guess on his own is that the prize is worth 0.65, he refuses to bid even at a reserve price of 0.1.

V(iii). *Effect of Reserve Price on Expected Revenue and Profit*

Of course, our main interest is not the outcome for a particular realization of signals, but *ex ante* expected outcomes. Figures 3, 5, 6, and 7 compare expected auction outcomes at different reserve prices across the two valuation models for various parameterizations of this example. Outcomes were calculated via numerical simulation; details are given in Appendix A(i). For each chart, the x-axis is reserve price, ranging from 0 to 1 (the support of valuations), and the y axis is expected profits (or expected revenue in the case of Figure 3). Outcomes under the common values model, $\pi_{CV}(r)$ (or $Rev_{CV}(r)$), and its maximizer r^*_{CV}, are shown as the lighter curve; results for the corresponding private values model ($\pi_{PV}(r)$ or $Rev_{PV}(r)$ and its maximizer r^*_{PV}) are shown as the darker curve. A dashed vertical line indicates the value of $\tilde{r}$, above which Theorem 2 applies; a shorter vertical line indicates the seller's valuation, v_0.

For Figures 3, 5 and 6, the level of signal precision α is fixed at 2. Figure 3 shows expected revenue, as a function of reserve price, under the two different models, for various values of N. Some things to note:

- $\tilde{r}$ (the dashed black line) is 0 at $N = 2$, and then increasing in N. Thus, Theorem 2 covers the widest range of reserve prices when N is small – exactly when reserve prices are most significant. Also note that the revenue-maximizing reserve price under the private values model, r^*_{PV}, is above $\tilde{r}$ for $N \leq 5$.
- While Theorem 2 only applies for $r \geq \tilde{r}$, the revenue ranking $Rev_{CV}(r) \leq Rev_{PV}(r)$ holds everywhere – at all r both above and below $\tilde{r}$, for each N considered in Figure 3.
- For $N \geq 3$, the revenue-maximizing reserve price under the common values model is $r = 0$, while the revenue-maximizing reserve price under the private values model is substantial.
- As N grows, the effect of reserve price vanishes under private values, but not under common values. When $N = 20$, revenue under private values is almost perfectly flat over the entire range of possible reserve prices; but a reserve price of 0.16 (optimal under private values) would cause a 6% loss in revenue under common values.

Note that fixing r, as N grows, expected revenue in the common values case is increasing toward 0.5, but gets there much more slowly than in the private values case. Figure 4 helps illustrate why, for the reserve price $r = 0.20$. In the common values case, as N increases, x^* increases as well, since the winner's curse a bidder must account for gets more severe. The left pane shows that as N increases, the likelihood of nobody's bidding increases, driven by an increase in

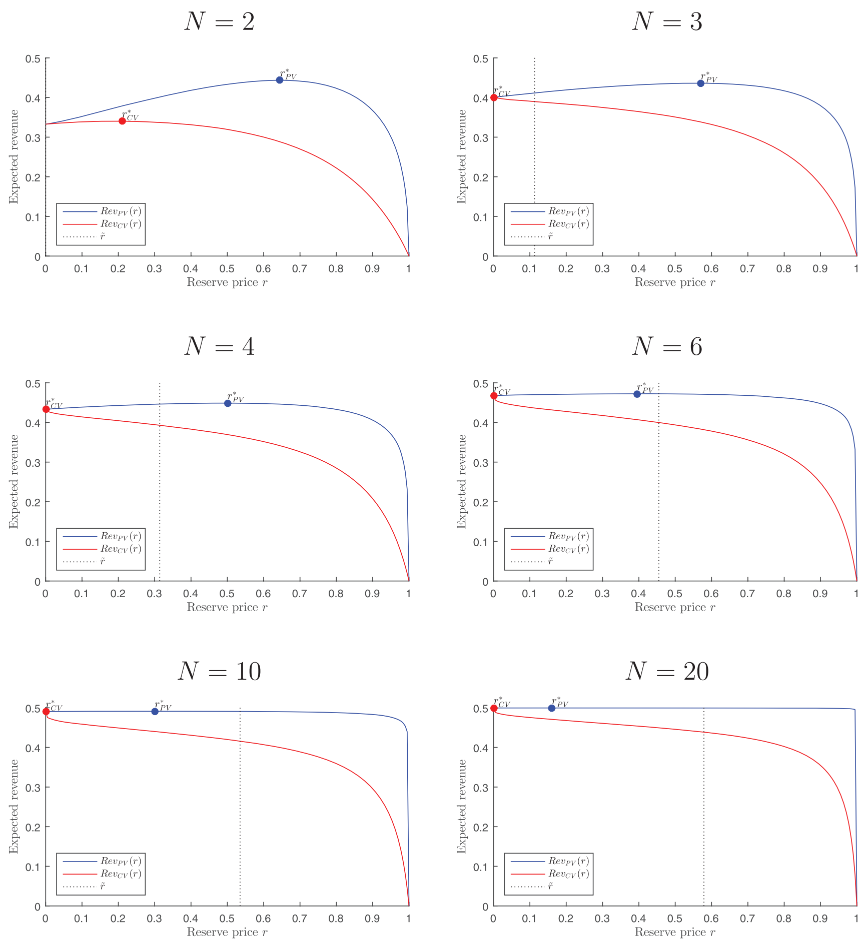

Figure 3
Expected Revenue as a Function of r for Various N (α = 2) [Colour figure can be viewed at *wileyonlinelibrary.com*]

the likelihood of no bids when $\theta = 0$. (The likelihood of no bids when $\theta = 1$ is decreasing, but is already so low that the decrease has less effect.)

The right pane of Figure 4 shows expected revenue when $r = 0.20$ in the private values (top curve) and common values (bottom curve) cases, as N changes.[11] The dashed line shows expected revenue (under either model) when no reserve price is used. The lighter line shows expected revenue under the common values

[11] In the private values case, the top two valuations $V^{(1)}$ and $V^{(2)}$ are, approximately, the expected value of θ conditional on all the realized signals; as N grows, these approach 1 with high probability when $\theta = 1$, and 0 when $\theta = 0$, so expected revenue quickly converges to 0.5 for any interior r.

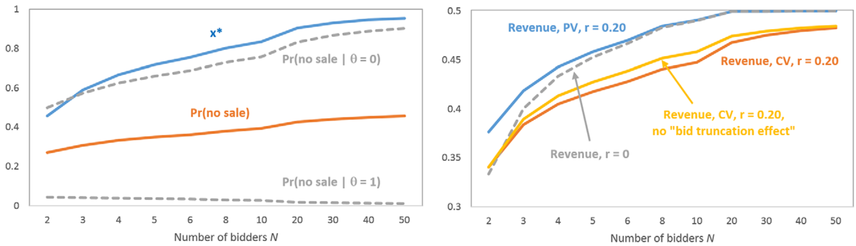

Figure 4
Effect of N on Outcomes with Fixed Reserve Price r = 0.20 [Colour figure can be viewed at *wileyonlinelibrary.com*]

model if bidders only bid when their signals exceed x^*, but those who do bid somehow learn the signals of the bidders who don't – that is, expected revenue if we shut down the effect the reserve price has on revenue through the truncation of losing bids. Thus, the difference between the lowest and next-lowest curves is the magnitude of this 'bid truncation effect,' while the difference between the lighter solid line and the dashed line is the revenue loss (or gain when N is small enough) due to the more obvious effects of a reserve price – the tradeoff between the possibility of no sale and the higher price when exactly one bidder bids.

Figure 5 compares expected profit across the two models, assuming the seller's valuation v_0 is 0.20. (Recalling that $E(V_i) = E(\theta) = \frac{1}{2}$ under the common values model, this means the seller's valuation is 40% of the expected buyer valuation.) Things to note in Figure 5:

- As in Figure 3 with revenue, the profit ranking – $\pi_{CV}(r) \leq \pi_{PV}(r)$ – holds nearly everywhere. The only exception is for $N = 2$ at some reserve prices below v_0, at which $\pi_{CV}(r) > \pi_{PV}(r)$. (This is because a sale is more likely under the private values model, but sales at prices below v_0 are unprofitable.)
- With $v_0 = 0.20$, a strictly positive reserve is optimal under common values as well as under private values, but $r^*_{CV} < r^*_{PV}$. Under both common and private values, the optimal reserve price is decreasing in N;[12] and as N grows, the increase in profits from setting r optimally (relative to setting $r = 0$) gets small quickly under both models.
- With $v_0 = 0.20$, when N is large ($N = 10$ and $N = 20$), the optimal reserve price under common values is *strictly below* v_0– the seller benefits from setting r low enough to risk selling at a loss. (Under private values, the optimal reserve is always above v_0.)

[12] Under an *independent* private values model, r^*_{PV} would be the same across N; but this need not hold when values are correlated, as they are here.

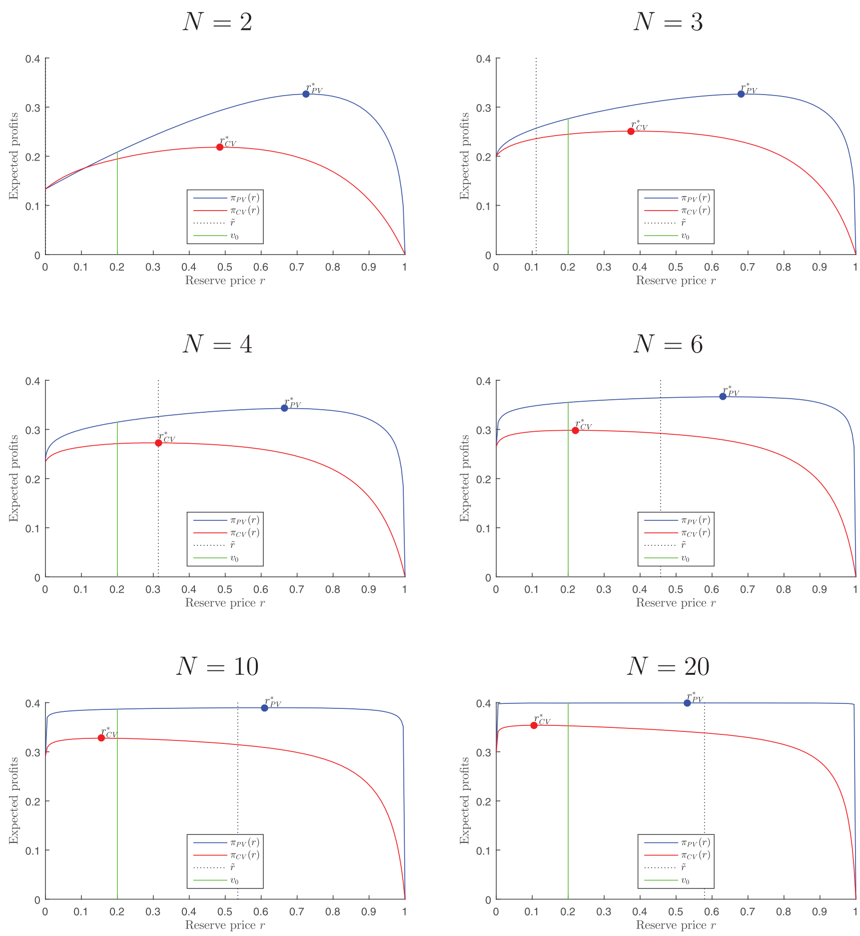

Figure 5
Expected Profit as a Function of r for Various N ($v_0 = 0.20$, $\alpha = 2$) [Colour figure can be viewed at *wileyonlinelibrary.com*]

Figure 6 considers the case $N = 5$ and $\alpha = 2$, and examines profit under different values of v_0. Some things to note in Figure 6:

- $\tilde{r} \approx 0.41$, so Theorem 2 applies to reserve prices above 0.41, a range which includes r^*_{PV} for all values of v_0, but only includes r^*_{CV} when v_0 is quite high.
- For each value of v_0, $r^*_{CV} < r^*_{PV}$; and for each value of v_0 and every value of r (both above and below $\tilde{r}$), $\pi_{CV}(r) \leq \pi_{PV}(r)$.
- When v_0 is small but positive, the profit-maximizing reserve price under the common values model is strictly less than v_0.

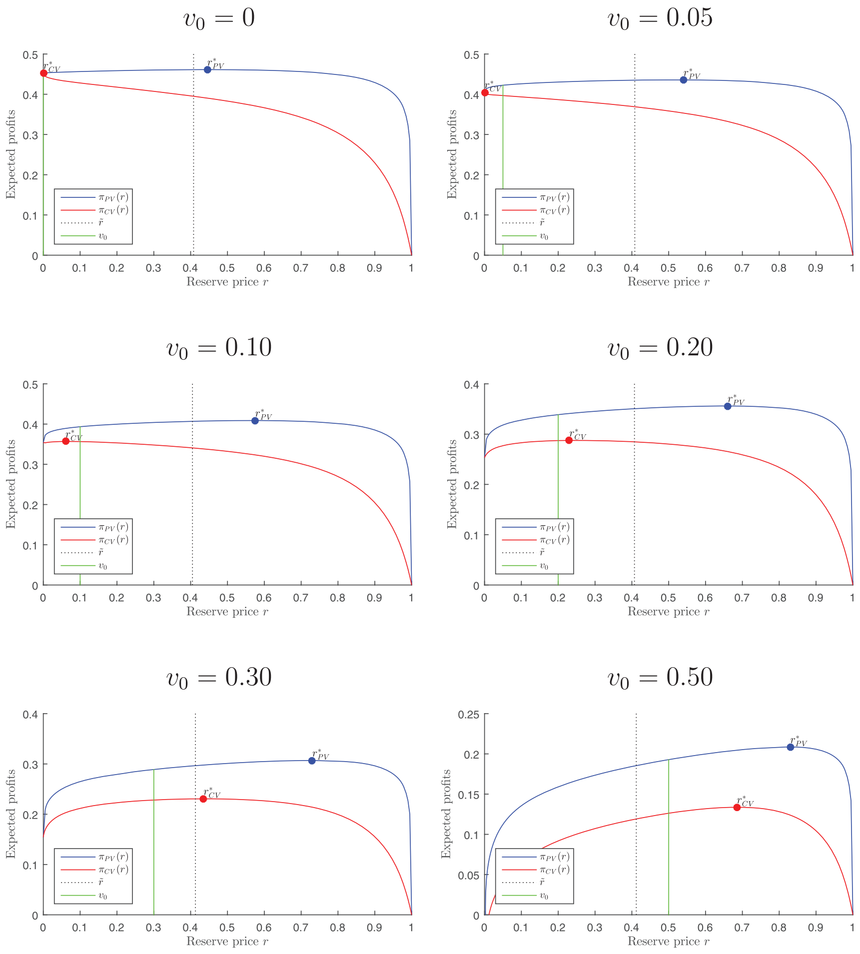

Figure 6
Expected Profit as a Function of r for Various v0 (N = 5, α = 2) [Colour figure can be viewed at *wileyonlinelibrary.com*]

Also note that when $v_0 = 0.05$, the finding that $r^*_{CV} = 0$ is not an artifact of a coarse grid: even a reserve price of 0.00001 would reduce expected profit.[13]

- Except when v_0 is quite high, the gain from setting the reserve price optimally under the common values model, relative to not using a reserve at all, is fairly small.

[13] When $N = 5$, a reserve of 0.00001 would still require a bidder to have a signal of at least 0.16 to bid, and would therefore cause nearly 30% of bidders not to bid when $\theta = 0$.

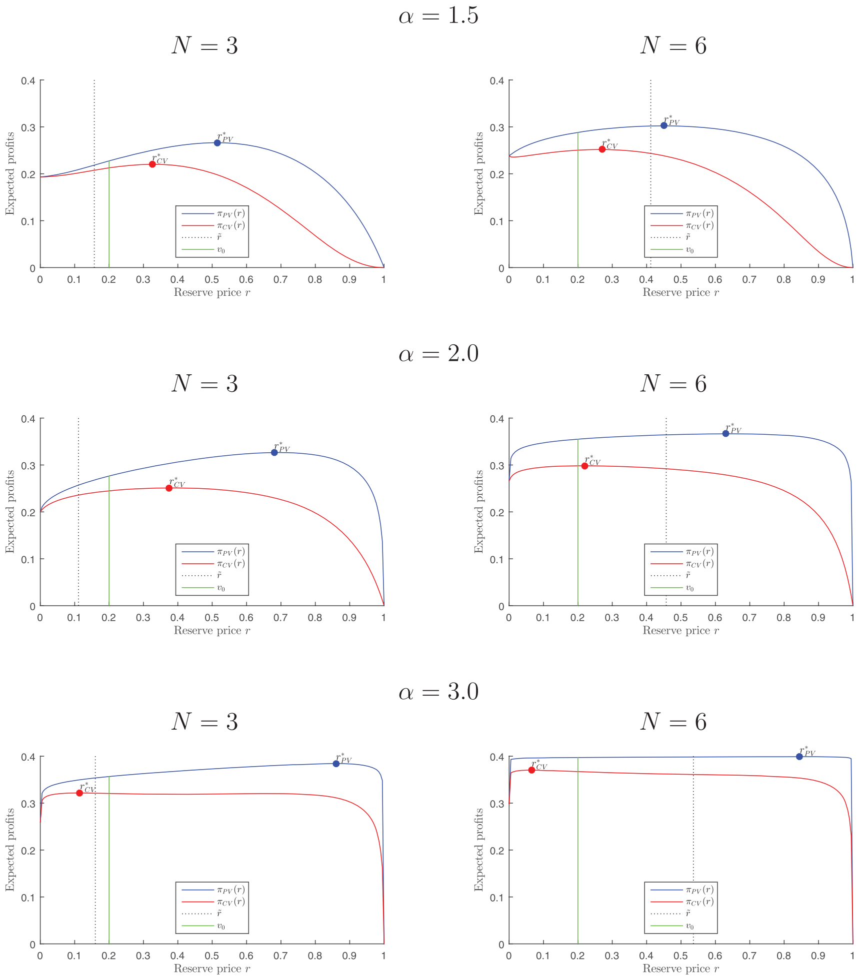

Figure 7
Expected Profit as a Function of r for Various α and N ($v_0 = 0.20$) [Colour figure can be viewed at *wileyonlinelibrary.com*]

Figure 7 shows how expected profit varies with α, the precision of the bidders' signals. Things to note:

- As individual signals get more precise, the effects of a reserve price get smaller.
- Across both values of N and all values of α tested, $\pi_{CV}(r) \leq \pi_{PV}(r)$ for all r, and $r^*_{CV} < r^*_{PV}$.
- r^*_{CV} is below v_0 when signals are very precise (high α).

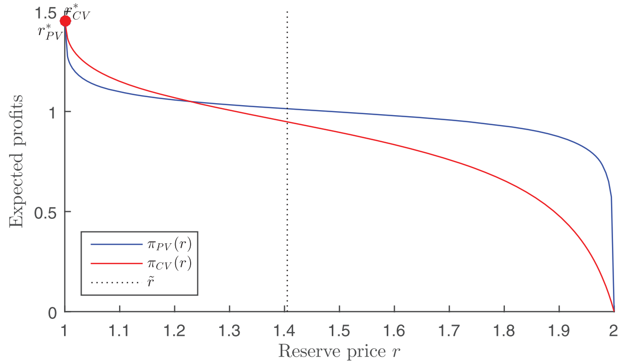

Figure 8
Expected Profit, N = 5, α = 2, $v_0 = 0$, Common Value $V_i = 1 + \theta$ [Colour figure can be viewed at *wileyonlinelibrary.com*]

Finally, Table I in Appendix A(ii) summarizes a few key measurements from each graph in Figures 3, 5, 6 and 7. The table shows the value of the optimal reserve price (in terms of the increase in expected profit relative to $r = 0$) for both the private and common value models, as well as the gain or loss under the common values model if the reserve price optimal under private values were used (again relative to no reserve). The patterns that emerge are these:

- First, the optimal reserve price nearly always gives less than half as much benefit under common values as under private values.
- And second, when v_0 is low, mistakenly using r^*_{PV} when there are actually common values is worse than using no reserve price; although when v_0 is high, mistakenly using r^*_{PV} is not so bad.

Aside from these, two patterns are consistent across Figures 3, 5, 6, and 7: the profit-maximizing reserve price is lower under common values in every case; and in every case, $\pi_{CV}(r) \leq \pi_{PV}(r)$ for every $r \geq v_0$, not just those above $\tilde{r}$. (The only case so far with $\pi_{CV}(r) > \pi_{PV}(r)$ *anywhere* is for some values of $r < v_0$ in the $N = 2$ case in Figure 5.)

However, it is worth noting that at least this last result does not hold universally – that is, it is not true that for every possible common values model and its corresponding private values model, $\pi_{CV}(r) \leq \pi_{PV}(r)$ for every $r \geq v_0$. For a counterexample to this possible conjecture, fix $N = 5$ and $\alpha = 2$, and modify the baseline example by letting the common value be $V_i = 1 + \theta$ rather than θ (and modify the private value analogue accordingly). For $v_0 = 0$ and r below about 1.2, the common values model predicts higher profit than the private values model. Figure 8 illustrates this example.

What's happening here is this. Any reserve price above v_0 trades off a cost (giving up some profitable sales) versus a benefit (improving the price when a sale occurs). For $r < \tilde{r}$, the private values model predicts a greater decrease in the likelihood of the sale; while for any r, it predicts a more favorable impact on price. By inflating valuations to $1 + \theta$ instead of θ, this example increases bids by 1 as well, while leaving the seller's valuation unchanged. Thus, the value of the sales given up by setting a reserve is magnified, relative to the changes in price when a sale occurs. Thus, for r just above 1, the reserve price hurts more under private values than under common values.

Of course, Figure 8 also shows that this reversal is not so practically relevant, since both models suggest that a nonbinding reserve $r \leq 1$ is optimal. I have not been able to find an example yet where π_{CV} is above π_{PV} (or Rev_{CV} is above Rev_{PV}) at reserve prices that are 'better than' $r = 0$, nor an example where the profit-maximizing reserve price is higher under common values than under private values. Still, I have not been able to show that such cases are impossible.

V(iv). *An Alternative Model with Continuous θ*

To make sure that the effects found so far do not hinge on the discreteness of θ and the resulting extremeness of valuations, I consider a second example with continuous-valued θ. This time, θ is distributed uniformly on [0, 1]. Once again, bidder signals take values in [0, 1], and are $i.i.d.$ conditional on the value of θ, this time with density

$$f(s|\theta) \quad = 1 + 4\left(\theta - \tfrac{1}{2}\right)\left(s - \tfrac{1}{2}\right)$$

Thus, when $\theta < \frac{1}{2}$, $f(\cdot|\theta)$ is decreasing, so lower signals are more likely; and when $\theta > \frac{1}{2}$, $f(\cdot|\theta)$ is increasing, and higher signals are more likely.

One difference between this example and the previous one is that signals are not 'unboundedly strong.' Even a signal of $X_i = 0$ or $X_i = 1$ leaves a chance of a wide range of possible θ, and as a result, $E(\theta|X)$ does not have full support. This means that for reserves below a certain level, no bidders will be deterred from bidding; and for reserves above a certain level, nobody will bid.[14] Another feature of this model is that for $N < 5$, the CDF's of $R^{(1)}$ and $V^{(1)}$ never cross, so $\tilde{r} = 0$, and Theorem 2 therefore applies to *all* reserve prices.

Figures 12 and 13 and Table II in Appendix A(iii) show expected revenue and profit for this example, for various values of N and v_0, and summarize key measurements. The main takeaways from this example are similar to those from the discrete-θ example:

- $\pi_{CV}(r) \leq \pi_{PV}(r)$ nearly everywhere[15]

[14] The maximum reserve at which everybody bids with probability 1 depends on N: it is $r \approx 0.20$ for $N = 3$, $r \approx 0.145$ when $N = 5$, $r \approx 0.115$ when $N = 7$, and $r \approx 0.085$ when $N = 10$. The minimum reserve at which nobody bids is $r = \frac{2}{3}$ regardless of N.

[15] For small N and v_0 sufficiently high (not shown), $\pi_{CV} > \pi_{PV}$ for some $r < v_0$; and for $N = 10$, $Rev_{CV} > Rev_{PV}$ for some positive reserve prices below $\tilde{r}$ (shown in figure 12, pane 7), but at reserves

- $r^*_{CV} < r^*_{PV}$, except when both are 0
- r^*_{CV} is sometimes below v_0, and sometimes 0 even when $v_0 > 0$

Thus, as with discrete θ, I find that most of the time, the optimal reserve price, and the expected profit at any given positive reserve price, is lower under common values than under the corresponding private values model. Table II shows that under this example, when v_0 is 0, 0.10, or even 0.20, even the optimal reserve price gives basically no benefit under common values, and when $v_0 = 0.35$, the benefit is about half what it would be under private values, once again similar to the findings above for discrete θ.

VI. EXTENSIONS

VI(i). *When v_0 Depends on θ (the 'Seller's Curse')*

Up to now, I've assumed v_0 is constant – the seller has a fixed valuation for the unsold object. In many settings – for example, if the seller's valuation is based on the same use as the buyers' valuations, or on resale prospects – it is natural to think the seller's valuation would vary along with the buyers'. With a positive reserve price, this introduces an adverse selection problem analogous to the winner's curse: the seller is most likely to retain the object exactly when it is least valuable to her. This curse further reduces the benefit of a reserve price.

Figure 9 illustrates this problem, comparing $\pi_{CV}(r)$ when $v_0 = \beta\theta$ to the case where v_0 is fixed at $\beta E(\theta) = \frac{1}{2}\beta$, using the discrete-$\theta$ example (with $\alpha = 2$) from the previous section. (Figure 14 in Appendix A(iii) does the same for the continuous-θ example.) At $r = 0$, of course, the two give the same outcome, since the object is always sold; but for any $r > 0$, expected profits (and therefore the benefit of a reserve price) are substantially lower when the seller's valuation depends on θ. In most of these cases, $r = 0$ was not optimal when v_0 was fixed, but becomes optimal when v_0 depends on θ.

It is worth noting one case – the third pane, corresponding to the discrete-θ example with $N = 3$ and a high seller valuation. While the value of a reserve price goes down significantly when v_0 is a function of θ rather than fixed, the optimal reserve price actually goes up. This is because when $v_0 = 0.7\theta$ and the reserve price is already being set reasonably high, the seller effectively 'gives up' on selling when $\theta = 0$, and optimizes primarily for the case where $\theta = 1$ and her residual valuation is 0.7; the optimal reserve price turns out to be right around 0.7, although the expected profit is only slightly higher than from setting $r = 0$. (This optimal reserve is still below the seller's optimal reserve under the analogous private values model with $v_0 = 0.35$, which is about 0.78.)

giving lower revenue than $r = 0$. Aside from these exceptions, $\pi_{CV} \leq \pi_{PV}$ and $Rev_{CV} \leq Rev_{PV}$ everywhere in every parameterization I've tried.

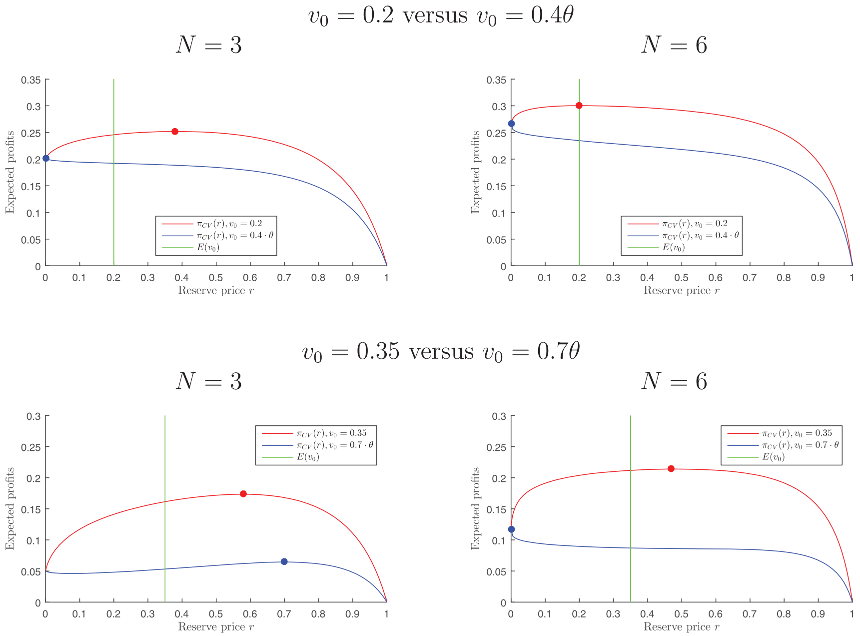

Figure 9

Expected Profit, Common Values, Discrete θ, Fixed v0 versus $v_0 = \beta\theta$ [Colour figure can be viewed at *wileyonlinelibrary.com*]

VI(ii). *Seller's Incentives to Disclose Information*

If the seller's valuation depends on θ, it might also be natural to suppose that the seller, like the buyers, may have some private information about θ. If the seller's information is verifiable, the result from Milgrom and Weber [1982, Theorem 18] extends to our setting, and the seller can always gain *ex ante* by committing to reveal her information (and adjusting the reserve price accordingly).[16]

Numerical examples suggest that the seller is likely to gain from revealing her information even if she does not alter the reserve price based on that information. Consider the discrete example from earlier, with $\alpha = 2$ and $v_0 = \beta\theta$. Suppose now that the seller also receives a signal S about θ; for simplicity, suppose the seller's signal is binary, and matches θ with probability $\frac{3}{4}$.[17]

[16] The result from Milgrom and Weber says that for any reserve price $\hat{r}$ with threshold signal $x^*(\hat{r})$, the seller can do weakly better by always revealing the value s of her own information S and then setting the reserve price $r(s)$ that, conditional on s, gives the same threshold signal $x^*(r(s)|s) = x^*(\hat{r})$ as before. While the result in Milgrom and Weber is for v_0 fixed, this new policy does not alter the set of signal realizations (S, X) under which the object is sold, so the dependence of v_0 on X has no effect on expected profit and the result therefore carries over to this setting.

[17] That is, $\Pr(S = 0|\theta = 0) = \Pr(S = 1|\theta = 1) = \frac{3}{4}$, and $\Pr(S = 1|\theta = 0) = \Pr(S = 0|\theta = 1) = \frac{1}{4}$.

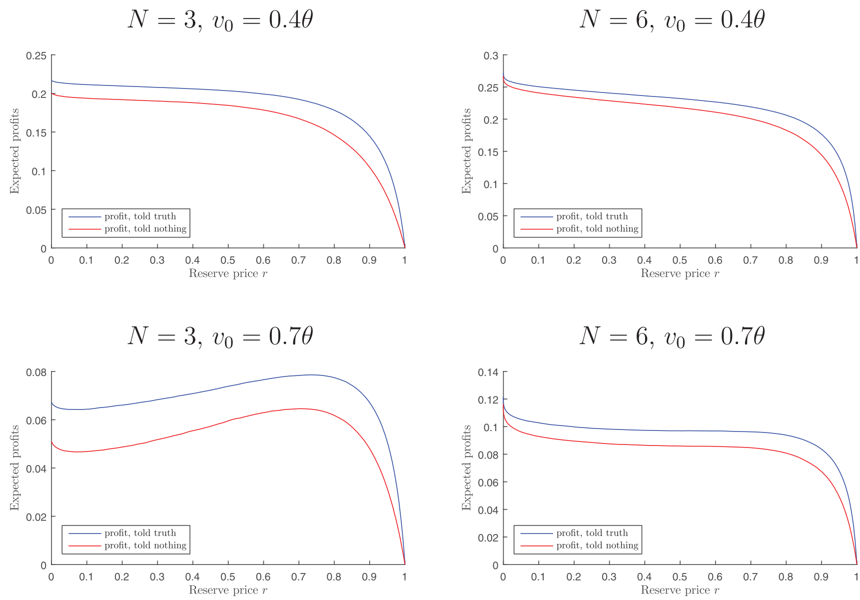

Figure 10
Seller's Incentive to Reveal Verifiable Information (Discrete θ, α = 2) [Colour figure can be viewed at *wileyonlinelibrary.com*]

Figure 10 compares the seller's expected profit at each reserve price when she reveals S (the top curve) to when she does not reveal S (the bottom curve). In all four examples – $N=3$ and 6, $v_0 = 0.4\theta$ and 0.7θ – the seller's expected profit is strictly higher at *every* reserve price when she reveals S. The increase in profits from revealing S and setting reserve optimally ranges from 2% (when $N=6$ and $v_0 = 0.4\theta$) to 22% (when $N=3$ and $v_0 = 0.7\theta$).

In three cases, the optimal reserve price is $r=0$ whether or not S is revealed (and regardless of its realization if it is), so no adjustment is necessary. In the remaining case ($N=3$ and $v_0 = 0.7\theta$), nearly all of the seller's benefit comes from the disclosure of the information, not from the subsequent re-optimization of the reserve price. That is, relative to a policy of not disclosing S and setting the reserve price optimally at 0.705, the seller can increase expected profit by 21.4% by committing to disclose S and still setting $r=0.705$, and only an additional 0.8% by switching to the new optimal reserve price for each realized value of S.[18] The extension of Milgrom and Weber's result guarantees the seller can gain by revealing her information and adjusting the reserve price accordingly; this example suggests the gain can be substantial, and that nearly all of it comes

[18] When S is not disclosed, the optimal reserve of $r=0.705$ gives expected profit of 0.0646. If the seller discloses S and still sets $r=0.705$, expected profit rises to 0.0784. If she discloses S and sets r optimally – at 0.753 when $S=1$ and 0.648 when $S=0$ – expected profit is 0.0790.

simply from revealing the information, not from the subsequent reserve price adjustment.[19]

VI(iii). *One Final Example – A Simple Linear Model*

Finally, I consider one other example, where bidder signals are independent and valuations have additive private and common value components. Let $\{X_i\}$ be distributed independently and uniformly on [0, 1], and let

$$V_i = (1-\lambda)\frac{N+2}{2N}X_i + \lambda\frac{\sum_{j=1}^{N} X_j}{N}$$

That is, a bidder's valuation is $1-\lambda$ times his own signal, plus λ times the average signal, so $\lambda = 0$ corresponds to pure private values and $\lambda = 1$ to pure common values. (The $\frac{N+2}{2N}$ term is a normalization to ensure that in the absence of a reserve price, expected revenue is constant as λ changes, so that in some sense we're comparing apples to apples. However, this example is *not* an instance of the general model presented in Section IV: the joint distribution of equilibrium bids changes with λ, so the $\lambda = 0$ case is not the 'observationally-equivalent private values case' of other values of λ. This is a completely separate model being offered to show that the results seem to hold more broadly.)

Figure 11 illustrates this example, giving the expected revenue and profit curves for five values of λ (0, 0.25, 0.5, 0.75, and 1) for several combinations of N and v_0.

In this example, the relevant rankings across models can all be established analytically. For a given reserve price r, it's straightforward to calculate that the bidding threshold x^*, when it is interior, satisfies $r = E(V_i | X_i = X^{(1)} = x^*) = \frac{N+2-\lambda}{2N}x^*$, and therefore that

$$x^* = \frac{2N}{N+2-\lambda}r$$

which is strictly increasing in λ. From this, I show the following:

Result 1. In the linear-independent example, expected revenue at $r = 0$ is constant in λ. For any $r > 0$ giving a nonzero probability of sale,

[19] If the seller had access to *unverifiable* information, she would always have an incentive to report the best possible news, so 'cheap talk' disclosure would not be credible, but reserve price could be used as a costly signal. Jullien and Mariotti [2006] and Cai, Riley and Ye [2007] study the problem of seller signaling via reserve price in a slightly different setting (see also Lamy [2010]). In the supplemental materials, I illustrate the equilibria (both separating and pooling) of the example in this section (with $N = 3$ and $v_0 = 0.7\theta$); I find that all separating equilibria give the seller *less than half* the profit she would get from the optimal pooling equilibrium, and therefore that sellers appear to have little incentive to acquire unverifiable information.

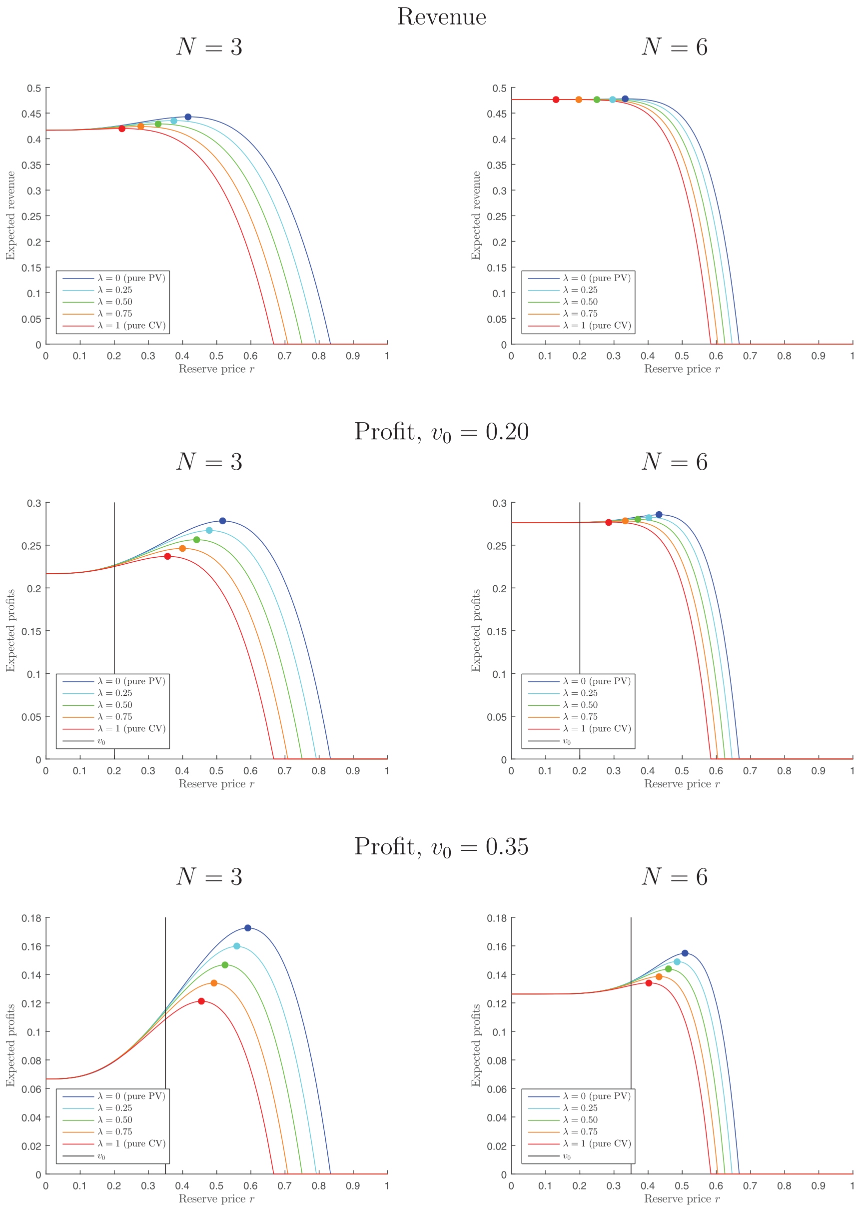

Figure 11

Expected Revenue and Profit for Linear-Independent Example [Colour figure can be viewed at *wileyonlinelibrary.com*]

(i) the probability of a sale is strictly decreasing in λ
(ii) expected revenue is strictly decreasing in λ
(iii) if $r \geq v_0$, expected profit is strictly decreasing in λ
Further, the revenue and profit-maximizing reserve prices are both strictly decreasing in λ.

The calculation is shown in the supplemental materials. Note that if signals were *i.i.d.* draws from any arbitrary distribution F rather than the uniform, the three numbered parts of this result would still hold, provided the private-value part of bidder valuations was normalized appropriately.[20]

VII. CONCLUSION

When bidders have private values, a reserve price offers a seller a way to avoid unprofitable sales, but also a way to further increase profits, by trading off a lower likelihood of sale against a higher average price. In this paper, I show this tradeoff is generally less favorable when bidders have interdependent or common values. I offer theoretical results that a high reserve price is less likely to be met, and less profitable, when bidders have common rather than private values; and I offer simulation results showing that the same profit ranking very often holds for *any* positive reserve price, not just high ones, and that the profit-maximizing reserve is typically lower under common values as well. Put another way, when bidders have interdependent values, analysis based on the assumption of private values is likely to overestimate both the optimal reserve price and the benefit of setting it. Thus, common values can be added to the list of departures from the standard workhorse model which would favor lower reserve prices – and which might help to explain the low reserve prices often observed empirically.

One reason for these findings is that beyond the usual tradeoff between likelihood of sale and minimum price, a reserve price also has an added cost under common values: it can reduce the expected price paid even when it does not bind, by concealing the bids of losing bidders. As a result, under common values, the profit-maximizing reserve price is sometimes below the seller's own valuation, and sometimes 0 even when the seller's valuation is positive. These effects are further magnified when the seller's valuation is interdependent with that of the buyers. Thus, while a reserve price is often an effective tool to increase seller profits in environments where buyers are confident of their own willingness to pay (private values), this paper shows they should be used much more cautiously – if at all – when bidder values have a significant common component.

[20] Like the $\frac{N+2}{2N}$ term in the uniform case, the normalization would be to make expected revenue when $r=0$ independent of λ for each realization of the second-highest signal. Specifically, this would mean defining valuations as $V_i = (1-\lambda)h(X_i) + \lambda \frac{1}{N}\sum_{j=1}^{N} X_j$, where $h(x) = \frac{1}{N}\left(2x + (N-2)E(X_j | X_j < x)\right)$.

APPENDIX

A(i). *Details of Simulations*

For the first example (discrete θ), note that by Bayes' Law,

$$\mathbb{E}\{\theta|X=x\} = \frac{\frac{1}{2}\prod_{i=1}^{N}\alpha x_i^{\alpha-1}}{\frac{1}{2}\prod_{i=1}^{N}\alpha x_i^{\alpha-1}+\frac{1}{2}\prod_{i=1}^{N}\alpha(1-x_i)^{\alpha-1}}$$

$$= \frac{1}{1+\prod_{i=1}^{N}\frac{(1-x_i)^{\alpha-1}}{x_i^{\alpha-1}}} \tag{1}$$

and

$$\mathbb{E}\left\{\theta|X^{(1)}=x_1\right\} = \frac{\frac{1}{2}\alpha x_1^{\alpha-1}(x_1^{\alpha})^{N-1}}{\frac{1}{2}\alpha x_1^{\alpha-1}(x_1^{\alpha})^{N-1}+\frac{1}{2}\alpha(1-x_1)^{\alpha-1}(1-(1-x_1)^{\alpha})^{N-1}}$$

$$= \frac{1}{1+\frac{(1-x_1)^{\alpha-1}}{x_1^{\alpha-1}}\left(\frac{1-(1-x_1)^{\alpha}}{x_1^{\alpha}}\right)^{N-1}} \tag{2}$$

and

$$\mathbb{E}\left\{\theta|(X_1,\ldots,X_k)=(x_1,\ldots,x_k),\{X_{k+1},\ldots,X_N\}<x^*\right\}$$

$$= \frac{\frac{1}{2}\prod_{i=1}^{k}\alpha x_i^{\alpha-1}\prod_{i=k+1}^{N}(x^*)^{\alpha}}{\frac{1}{2}\prod_{i=1}^{k}\alpha x_i^{\alpha-1}\prod_{i=k+1}^{N}(x^*)^{\alpha}+\frac{1}{2}\prod_{i=1}^{k}\alpha(1-x_i)^{\alpha-1}\prod_{i=k+1}^{N}(1-(1-x^*)^{\alpha})}$$

$$= \frac{1}{1+\left(\frac{1-(1-x^*)^{\alpha}}{(x^*)^{\alpha}}\right)^{N-k}\prod_{i=1}^{k}\frac{(1-x_i)^{\alpha-1}}{x_i^{\alpha-1}}} \tag{3}$$

For the simulations, the following was done in *Matlab*:

1. 500,000 sets of simulated signal realizations were randomly generated, by picking $\theta \in \{0,1\}$ at random for each simulation, generating N independent uniform random variables $\{\varepsilon_i\}$, and letting $X_i = F^{-1}(\varepsilon_i|\theta)$.
2. For each set of realized signals, the highest two private valuations $V^{(1)} = E(\theta|X = (x^{(1)}, x^{(1)}, x^{(3-N)}))$ and $V^{(2)} = E(\theta|X = (x^{(2)}, x^{(2)}, x^{(3-N)}))$ were calculated via (1).
3. For each $r \in \{0, 0.005, 0.010, 0, 015, \ldots, 0.995, 1.000\}$, x^* was calculated via *Matlab's* numerical solver as the solution to $E(\theta|X^{(1)} = x^*) - r = 0$ via (2).
4. For each simulation and each r, the outcome was calculated under each model by determining which bidders have signals above x^* (common values) or a valuation above r (private values) and calculating what price would be paid, if any (using (3) for the common values case), and subtracting the seller's valuation v_0 in case of a sale. Revenue and profit curves were then produced by averaging across simulations, and optimal reserve prices were determined as the grid point giving the highest revenue/profit.
5. The same simulated valuations were used to calculate $R^{(1)}$ and $V^{(1)}$ as defined in the text, and $\tilde{r}$ was calculated as the highest crossing point of their empirical CDF's.

For the second example (continuous θ), the simulations were done the same way, but calculation of $E(\theta|X)$ and $E(\theta|X^{(1)})$ were a bit more laborious. Letting P denote the prior on θ and $p(X|t)$ the conditional density of X given $\theta = t$, Bayes' Law gives us

$$E(\theta|X) = \int_0^1 \theta \frac{P(\theta)p(X|\theta)}{\int_0^1 P(t)p(X|t)dt} d\theta \quad = \quad \frac{\int_0^1 tp(X|t)dt}{\int_0^1 p(X|t)dt}$$

since P is uniform. So

$$\begin{aligned} E(\theta|(X_1,\ldots,X_k) &= (x_1,\ldots,x_k), \{X_{k+1},\ldots,X_N\} < x^*) \\ &= \frac{\int_0^1 t \prod_{i=1}^k f(x_i|t) \prod_{i=k+1}^N F(x^*|t)dt}{\int_0^1 \prod_{i=1}^k f(x_i|t) \prod_{i=k+1}^N F(x^*|t)dt} \end{aligned}$$

Plugging in $f(x_i|t) = 1 + 4(t - \frac{1}{2})(x_i - \frac{1}{2})$ and $F(x^*|t) = x^* \left(1 + 4(t - \frac{1}{2})(\frac{1}{2}x^* - \frac{1}{2})\right)$ (calculated via integration), some algebra and a change of variables allow this to be rewritten as

$$\frac{1}{2} + \frac{1}{4} \frac{\int_{-2}^2 T \prod_{i=1}^N (1 + Ta_i)dT}{\int_{-2}^2 \prod_{i=1}^N (1 + Ta_i)dT}$$

where $a_i = x_i - \frac{1}{2}$ for $i = 1, 2, \ldots, k$ and $a_i = \frac{1}{2}x^* - \frac{1}{2}$ for $i > k$.

Setting $a_1 = x - \frac{1}{2}$ and $a_i = \frac{1}{2}x - \frac{1}{2}$, then, we integrated both numerator and denominator (separately for each N – *Matlab* yielded closed-form integrals) to get an expression for $E(\theta|X^{(1)} = x)$, and found x^* by solving $E(\theta|X^{(1)} = x) - r = 0$ numerically within *Matlab*. To calculate simulated bids, we simplified the last expression for $E(\theta|X)$ to

$$E(\theta|X) = \frac{1}{2} + \frac{1}{4} \cdot \frac{\frac{2^4}{3}A_1 + \frac{2^6}{5}A_3 + \frac{2^8}{7}A_5 + \frac{2^{10}}{9}A_7 + \ldots}{4 + \frac{2^4}{3}A_2 + \frac{2^6}{5}A_4 + \frac{2^8}{7}A_6 + \frac{2^{10}}{9}A_8 + \ldots}$$

where

$$A_1 = \sum_i a_i$$

$$A_2 = \sum_{i<j}\sum a_i a_j$$

$$A_3 = \sum_{i<j<k}\sum\sum a_i a_j a_k$$

$$A_4 = \sum_{i<j<k<l}\sum\sum\sum a_i a_j a_k a_l$$

and so on, with A_n therefore equal to $\prod_i a_i$ and $A_{n+1} = A_{n+2} = \cdots = 0$, and wrote code to calculate this within *Matlab* for each simulation. The simulations were then done just as in the discrete-θ case.

A(ii). *Some Vital Statistics*

TABLE I
SUMMARIZING FIGURES 3, 5, 6 AND 7

Figure 3 – discrete θ, $\alpha = 2$, $v_0 = 0$ (expected revenue)						
N	2	3	4	6	10	20
$\pi_{PV}(r^*_{PV})/\pi_{PV}(0) - 1$	33.1%	8.8%	3.6%	0.9%	0.1%	0.0%
$\pi_{CV}(r^*_{CV})/\pi_{CV}(0) - 1$	2.2%	0.0%	0.0%	0.0%	0.0%	0.0%
$\pi_{CV}(r^*_{PV})/\pi_{CV}(0) - 1$	−13.0%	−14.9%	−14.8%	−12.3%	−10.9%	−8.3%
Figure 5 – discrete θ, $\alpha = 2$, $v_0 = 0.20$						
N	2	3	4	6	10	20
$\pi_{PV}(r^*_{PV})/\pi_{PV}(0) - 1$	146%	63%	47%	38%	34%	33%
$\pi_{CV}(r^*_{CV})/\pi_{CV}(0) - 1$	65%	25%	17%	12%	12%	18%
$\pi_{CV}(r^*_{PV})/\pi_{CV}(0) - 1$	44%	14%	6%	4%	7%	13%
Figure 6 – discrete θ, $N = 5$, $\alpha = 2$						
v_0	0.00	0.05	0.10	0.20	0.30	0.50
$\pi_{PV}(r^*_{PV})/\pi_{PV}(0) - 1$	2%	8%	15%	41%	99%	*
$\pi_{CV}(r^*_{CV})/\pi_{CV}(0) - 1$	0%	0%	1%	14%	49%	*
$\pi_{CV}(r^*_{PV})/\pi_{CV}(0) - 1$	−14%	−12%	−9%	5%	39%	*
Figure 7 – discrete θ, $v_0 = 0.20$						
N	3	3	3	6	6	6
α	1.5	2.0	3.0	1.5	2.0	3.0
$\pi_{PV}(r^*_{PV})/\pi_{PV}(0) - 1$	38%	63%	49%	27%	38%	34%
$\pi_{CV}(r^*_{CV})/\pi_{CV}(0) - 1$	14%	25%	24%	6%	12%	25%
$\pi_{CV}(r^*_{PV})/\pi_{CV}(0) - 1$	3%	14%	20%	1%	4%	19%

Notes:* When $N = 5$, $\alpha = 2$, and $v_0 = 0.50$ in the discrete-θ example, $\pi_{PV}(0) = \pi_{CV}(0) < 0$, so the ratios are excluded because they are meaningless.

A(iii). *Continuous-θ Simulation Results*

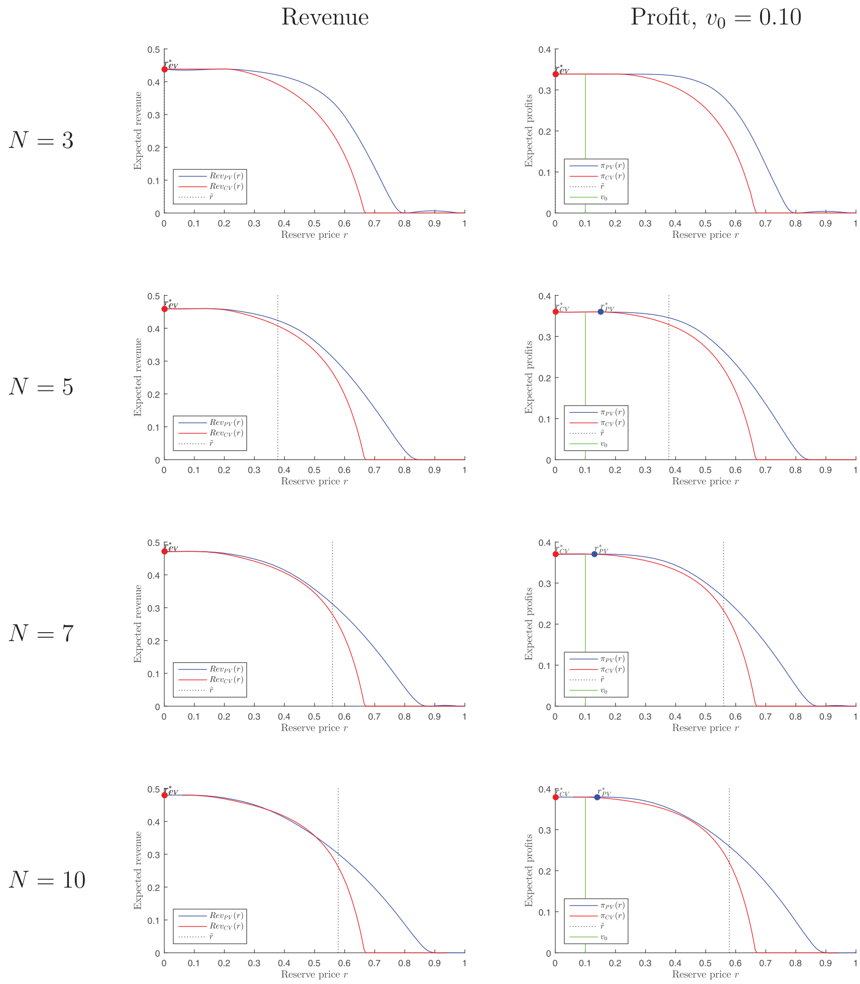

Figure 12
Expected Revenue and Profit, Continuous-θ Example [Colour figure can be viewed at *wileyonlinelibrary.com*]

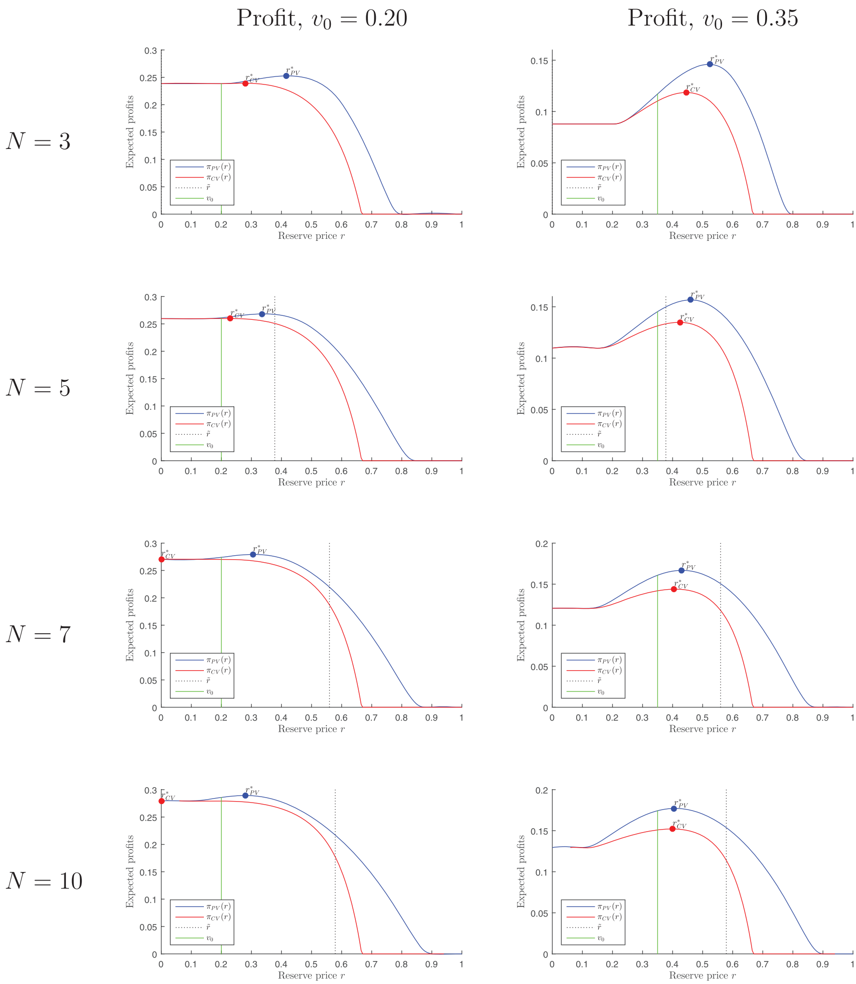

Figure 13

Expected Revenue and Profit, Continuous-θ Example (cont'd) [Colour figure can be viewed at *wileyonlinelibrary.com*]

TABLE II
VITAL STATISTICS FROM FIGURES 12 AND 13

Figure 12 – continuous θ, low v_0								
N	3	5	7	10	3	5	7	10
v_0	0	0	0	0	0.10	0.10	0.10	0.10
$\pi_{PV}(r^*_{PV})/\pi_{PV}(0) - 1$	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
$\pi_{CV}(r^*_{CV})/\pi_{CV}(0) - 1$	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
$\pi_{CV}(r^*_{PV})/\pi_{CV}(0) - 1$	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	−0.4%
Figure 13 – continuous θ, high v_0								
N	3	5	7	10	3	5	7	10
v_0	0.20	0.20	0.20	0.20	0.35	0.35	0.35	0.35
$\pi_{PV}(r^*_{PV})/\pi_{PV}(0) - 1$	6.0%	3.3%	3.2%	3.5%	66%	43%	38%	37%
$\pi_{CV}(r^*_{CV})/\pi_{CV}(0) - 1$	0.3%	0.1%	0.0%	0.0%	35%	23%	19%	18%
$\pi_{CV}(r^*_{PV})/\pi_{CV}(0) - 1$	−0.5%	−1.7%	−1.2%	−0.7%	26%	22%	19%	18%

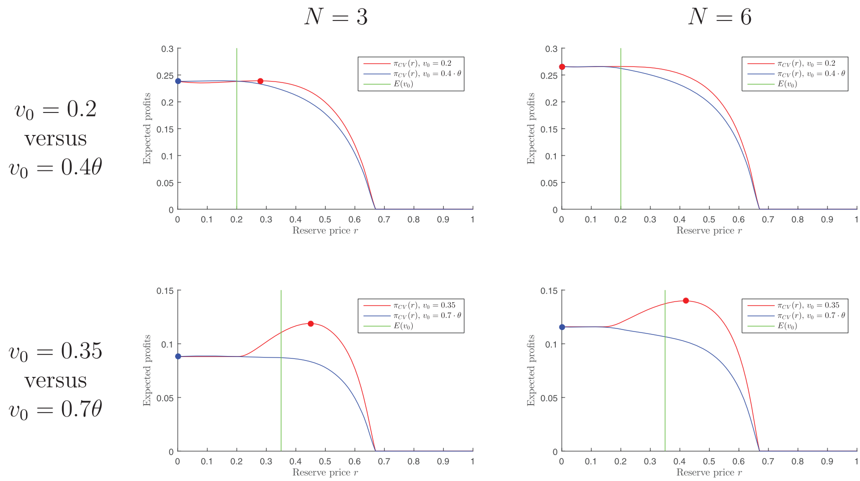

Figure 14
Expected Profit, Common Values, Continuous θ, Fixed v_0 versus $v_0 = \beta\theta$ [Colour figure can be viewed at *wileyonlinelibrary.com*]

REFERENCES

Aradillas-López, A.; Gandhi, A. and Quint, D., 2013, ‘Identification and Inference in Ascending Auctions with Correlated Private Values,’ *Econometrica*, 81(2), pp. 489–534.

Athey, S. and Haile, P., 2002, ‘Identification of Standard Auction Models,’ *Econometrica*, 70(6), pp. 2107–2140.

Athey, S. and Haile, P., 2007, ‘Nonparametric Approaches to Auctions,’ in Heckman, J., and Leamer, E. (eds.), *Handbook of Econometrics*, *Vol. 6A* (North Holland Press, Oxford, Oxon., England), pp. 3847–3965, Chapter 60.

Bikhchandani, S.; Haile, P. and Riley, J., 2002, ‘Symmetric Separating Equilibria in English Auctions,’ *Games and Economic Behavior*, 38(1), pp. 19–27.

Cai, H.; Riley, J. and Ye, L., 2007, ‘Reserve Price Signaling,’ *Journal of Economic Theory*, 135(1), pp. 253–268.

Haile, P.; Hong, H. and Shum, M., 2003, 'Nonparametric Tests for Common Values at First-Price Sealed-Bid Auctions,' NBER working paper, see *http://www.nber.org/papers/w10105*.

Haile, P. and Tamer, E., 2003, 'Inference with an Incomplete Model of English Auctions,' *Journal of Political Economy*, 111(1), pp. 1–51.

Hendricks, K.; Pinkse, J. and Porter, R., 2003, 'Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions,' *Review of Economic Studies*, 70(1), pp. 115–145.

Jullien, B. and Mariotti, T., 2006, 'Auction and the Informed Seller Problem,' *Games and Economic Behavior*, 56(2), pp. 225–258.

Kim, D.-H., 2013, 'Optimal Choice of a Reserve Price under Uncertainty,' *International Journal of Industrial Organization*, 31(5), pp. 587–602.

Laffont, J.-J. and Vuong, Q., 1996, 'Structural Analysis of Auction Data,' *American Economic Review*, 86(2), pp. 414–420.

Lamy, L., 2010, 'Corrigendum to "Reserve Price Signaling",' *Journal of Economic Theory*, 145(6), pp. 2498–2504.

Levin, D., and Smith, J., 1994, 'Equilibrium in Auctions with Entry,' *American Economic Review*, 84(3), pp. 585–599.

Milgrom, P., and Weber, R., 1982, 'A Theory of Auctions and Competitive Bidding,' *Econometrica*, 50(5), pp. 1089–1122.

Paarsch, H., 1992, 'Deciding between the Common and Private Value Paradigms in Empirical Models of Auctions,' *Journal of Econometrics*, 51(1-2), pp. 191–215.

Paarsch, H., 1997, 'Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales,' *Journal of Econometrics*, 78(2), pp. 333–357.

Peters, M., and Severinov, S., 1997, 'Competition among Sellers Who Offer Auctions Instead of Prices,' *Journal of Economic Theory*, 75(1), pp. 141–179.

Quint, D., 2008, 'Unobserved Correlation in Private-Value Ascending Auctions,' *Economics Letters*, 100(3), pp. 432–434.

Samuelson, W., 1985, 'Competitive Bidding with Entry Costs,' *Economics Letters*, 17(1-2), pp. 53–57.

Vincent, D., 1995, 'Bidding Off the Wall: Why Reserve Prices May Be Kept Secret,' *Journal of Economic Theory*, 65(2), pp. 575–584.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at *http://wileyonlinelibrary.com/journal/joie* or via the Journal's website, *http://www.jindec.org*