Computationally Convenient Distributional Assumptions for Common Value Auctions

Michael B. Gordy*
Board of Governors of the Federal Reserve System

January 21, 1997

Abstract

Although the mathematical foundations of common value auctions have been well understood since Milgrom and Weber (1982), equilibrium bidding strategies are computationally complex. Very few calculated examples can be found in the literature, and only for highly specialized cases. This paper introduces two sets of distributional assumptions that are flexible enough for theoretical and empirical applications and yet permit straightforward calculation of equilibrium bidding strategies.

*The views expressed herein are my own and do not necessarily reflect those of the Board of Governors or its staff. I thank Preston McAfee and Daniel Vincent for helpful suggestions and Margaret Kyle for excellent research assistance.

Please address correspondence to the author at Division of Research and Statistics, Mail Stop 153, Federal Reserve Board, Washington, DC 20551, USA. Phone: (202)452-3705. Fax: (202)452-5295. Email: ⟨mgordy@frb.gov⟩.

In many of the most prominent auction settings, bidders share a common ex-post valuation for the good for sale, but ex-ante do not know what that value is. Examples of *common value* ("CV") auctions occur in primary markets for Treasury bills, timber tracts and oil exploration leases. The mathematical foundations of CV auctions have been well understood since Milgrom and Weber (1982) ("M-W"), and there is a large experimental and empirical literature as well.[1]

Nonetheless, examples of calculated equilibrium bidding strategies are quite rare. Although the M-W solution to the CV auction has an intuitively straightforward form, in practice it is difficult to specify distributional assumptions that yield closed form or nearly closed form solutions, even for the simplest case of a sealed-bid CV auction of a single unit with an exogenous number of bidders. Matthews (1984, §5) provides a model with closed form solution in which the value of the good is drawn from a Pareto distribution. Matthews derives a number of interesting theoretical results from his solution. For empirical work, however, this specification is likely to be too restrictive, because the Pareto density is strictly decreasing over its entire support. In most empirical settings, one would expect the density for the value of the good to have some sort of bell shape (not necessarily symmetric) centered near its mean. Paarsch (1992, §2.1.1) presents a similar specification with similar limitations.[2]

The few other known examples in the literature are still more restrictive. Fully closed form solutions exist for a class of models in which each bidder draws an independent random variable and the value of the good is the average of the draws (see, e.g., Vincent 1995). These models are peculiar in that public information cannot be distinguished from private signals, and are thus not well suited for empirical work. Closed form solutions also can be found when the prior is assumed to be diffuse; see Engelbrecht-Wiggans and Weber (1979) and Levin and Smith (1991) on this special

[1]For a general introduction to the auction literature, see McAfee and McMillan (1987) and Milgrom (1989).

[2]Preston McAfee has pointed out that a binomial common value, combined with a power function distribution for the signal, allows for computationally straightforward examples of bidding strategies. This too is quite restrictive.

case. For most CV auctions, however, the assumption of no public information is untenable.[3]

Perhaps surprisingly, numerical solution by direct methods is typically quite difficult. Computation of the M-W equilibrium strategy involves nested definite integrals. In general, some of the integrands will tend to infinity at the lower limit of integration, which complicates numerical solution. Furthermore, integration can magnify any inaccuracy in the integrand, which itself may be the result of a numerical approximation. Laffont and Vuoung (1993) suggest a method based on simulated moments, but it is not clear how the simulator could be constructed without resolving the same numerical problems as arise in direct methods.

This paper presents two sets of distributional assumptions that yield computationally convenient solutions for both first-price and second-price sealed-bid common value auctions. Each set of assumptions is consistent in spirit with the M-W framework, in that there is a non-diffuse public prior and private signals are conditionally independent and increasing in expected value with the true value of the good. Each also is parametrically flexible, in that the mean and precision of the prior and the precision of private signals can each be varied independently, and the equilibrium strategy can be computed for any number of bidders $n \geq 2$. This parametric flexibility is especially valuable for empirical applications.

Section 1 briefly summarizes the M-W framework and solution under risk neutrality. Sections 2 and 3 show how the M-W equilibrium strategy is easily computed under two sets of distributional assumptions. Comparative statics for the bidding strategies are explored in Section 4. Comparative statics for bidder profits and seller revenue are discussed in Section 5. In Section 6 I show that a reserve price can easily be accomodated under these specifications. Further extensions and applications are discussed in the Conclusion.

[3] In U.S. Treasury auctions, for example, the great majority of relevant information is incorporated in publicly observed pre-auction when-issued prices.

1 Milgrom-Weber common value auction model

In the simplest CV auction, n risk-neutral bidders compete for a single good of unknown value V. The bidders share a prior distribution on its ex-post value, and each bidder receives a private signal X such that $X_i|V$ is independent of $X_j|V$ for all bidders $i \neq j$. Each bidder submits a single sealed bid. The good is awarded to the highest bidder, who pays either her own bid (under first-price rules) or the bid of the next highest bidder (under second-price rules).

Let $f(v)$ and $F(v)$ denote the density and c.d.f. of the prior distribution, and let $g(x|v)$ and $G(x|v)$ denote the conditional density and c.d.f. of a bidder's signal. Without loss of generality, I take the perspective of bidder 1. Let Y_1 denote the highest among the other bidders' signals $X_2, \ldots, X_n$. Then M-W (Theorem 14) show that the equilibrium bidding strategy for the first-price auction is given by the first-order linear differential equation

$$B_1{}'(x) = (\nu(x,x) - B_1(x))\lambda(x,x) \tag{1}$$

where

$$\nu(x,y) \equiv E[V|X_1 = x, Y_1 = y]$$

$$\lambda(y,x) \equiv g_{Y_1}(Y_1 = y|X_1 = x)/G_{Y_1}(Y_1 = y|X_1 = x)$$

and where g_{Y_1} denotes the density of order statistic Y_1. The boundary condition for equation (1) is $B_1(\underline{x}) = \nu(\underline{x}, \underline{x})$, where $\underline{x}$ is the infimum of the support of g. The solution to the differential equation can then be written as

$$B_1(x) = \int_{\underline{x}}^{x} \nu(t,t)\lambda(t,t)L(t|x)dt \tag{2}$$

where

$$L(t|x) \equiv \exp\left(-\int_t^x \lambda(s,s)ds\right).$$

The equilibrium strategy for the second-price auction is simply

$$B_2(x) = \nu(x,x). \tag{3}$$

2 Gamma Prior, Inverse Gamma Signal (GIG)

Let the prior for V be the Gamma(α,β) distribution for $\alpha > 0, \beta > 0$. The p.d.f. is given by

$$f(v) = \frac{\beta^\alpha}{\Gamma(\alpha)} v^{\alpha-1} \exp(-\beta v).$$

Let the private signals be conditionally independent and distributed "inverse Gamma," i.e., such that $1/X$ given $V = v$ is distributed Gamma$(\tau, \tau v)$. The parameter τ measures the precision of the signal. It is straightforward to check that the expectation of X is increasing in the realization of V, and that the signals satisfy the affiliation property discussed in M-W.[4] Thus, our distributional choices satisfy the assumptions in M-W's analysis. It is also easily checked that the posterior distribution of V is also Gamma, i.e., $V|(X = x) \sim$ Gamma$(\alpha + \tau, \beta + \tau x^{-1})$.

For tractability, it is necessary to restrict τ to the set of positive integers.[5] In this case, the conditional c.d.f. of X is given by

$$G(x|v) \equiv \Pr(X \le x|V = v) = 1 - \Pr(X^{-1} \le x^{-1}|V = v) = e_{\tau-1}(\tau v x^{-1}) \exp(-\tau v x^{-1})$$

where $e_k(z) \equiv 1 + z + z^2/2! + \ldots + z^k/k!$ (see Abramowitz and Stegun, eds 1968, 6.5.1,6.5.13). In

[4]The conditional distribution of private signals satisfies the monotone likelihood ratio property, which M-W show is a sufficient condition (see pages 1098-99).

[5]The Gamma distribution with an integer first parameter is also known as the Erlang distribution.

calculating the functions ν and λ, the only difficulty is the distribution for order statistic Y_1, which is

$$G_{Y_1}(Y_1 = y|V = v) = G(X = y|V = v)^{n-1} = e_{\tau-1}(\tau v x^{-1})^{n-1} \exp(-(n-1)\tau v x^{-1}).$$

Let $c_j^{(k,n)}$ denote the coefficient on z^j in the polynomial expansion of $e_k(z)^n$. As e_k has only $k+1$ terms, the polynomial expansion is quickly computed. The integrals needed to obtain $g_{Y_1}(Y_1 = y|X = x)$ and other constituent parts of ν and λ can then be expressed as a finite sum of integrals of the form $\int_0^\infty v^{a-1} \exp(-bv) dv$ which has solution $\Gamma(a)/b^a$.

For notational convenience, let $(a)_k$ be the Pochhammer symbol, i.e., $(a)_0 = 1, (a)_1 = a, (a)_k = (a)_{k-1}(a + k - 1)$, and define the function

$$\psi(x, a, b, \tau, m, n) \equiv \sum_{j=0}^{m\cdot(\tau-1)} c_j^{(\tau-1,m)} (a + (n-m)\tau)_j \left(\frac{\tau}{bx + n\tau}\right)^j.$$

Calculations outlined in Appendix A.1 give

$$\begin{aligned} B_2(x) = \nu(x, x) &= \frac{(\alpha + 2\tau)x}{\beta x + n\tau} \frac{\psi(x, \alpha + 1, \beta, \tau, n-2, n)}{\psi(x, \alpha, \beta, \tau, n-2, n)} && (4) \\ \lambda(x, x) &= \frac{(n-1)}{\mathrm{B}(\alpha + \tau, \tau)} \frac{\tau^\tau}{x(\beta x + n\tau)^\tau} \frac{\psi(x, \alpha, \beta, \tau, n-2, n)}{\psi(x, \alpha, \beta, \tau, n-1, n)}. && (5) \end{aligned}$$

where B denotes the beta function, $\mathrm{B}(a, b) \equiv \Gamma(a)\Gamma(b)/\Gamma(a + b)$.

Computation of $B_2(x)$ is quite simple and fast. All summation terms in the ψ function are positive, so floating point round-off error is not an issue. In general, there does not appear to be a closed form solution for $B_1(x)$, but numerical solution presents only minor difficulty. In practice, solution of differential equation (1) by Runge-Kutta seems to be the easiest way to obtain $B_1(x)$. At the boundary point $\underline{x} = 0$, $\nu(0, 0) = 0$ so $B_1(0) = B_2(0) = 0$. Although $\lambda(0, 0) = \infty$, the gradient of B_1 at $x = 0$ is bounded by

$$0 \leq B_1'(0) \leq \nu(0,0)\lambda(0,0) = \frac{(n-1)}{\mathrm{B}(\alpha+\tau,\tau)} \frac{(\alpha+2\tau)}{\tau n^{\tau+1}} \frac{\psi(0,\alpha+1,\beta,\tau,n-2,n)}{\psi(0,\alpha,\beta,\tau,n-1,n)} < \infty.$$

It makes no discernable difference to the overall solution whether $B_1'(0)$ is set to its upper or lower bound. I find that solution using Matlab on a SparcStation5 typically takes less than two seconds.[6]

Special case: When $\tau = 1$, the conditional distribution of the signal simplifies to the exponential distribution. In this case, ν and λ both have simple forms, the function $L(t|x)$ has closed form solution, and even $B_1(x)$ has a simple closed form expression. These are given by

$$\begin{aligned}
B_2(x) = \nu(x,x) &= \frac{(\alpha+2)x}{\beta x + n} \\
\lambda(x,x) &= \frac{(n-1)(\alpha+1)}{x(\beta x + n)} \\
L(t|x) &= \left(\frac{t}{x}\frac{\beta x + n}{\beta t + n}\right)^{(\alpha+1)(n-1)/n} \\
B_1(x) &= \left(1 - \frac{1}{1+(\alpha+1)(n-1)/n}\right)\frac{(\alpha+2)x}{\beta x + n}.
\end{aligned}$$

Special case: When $n = 2$, $B_2(x)$ has the simple form

$$B_2(x) = \nu(x,x) = \frac{(\alpha+2\tau)x}{\beta x + 2\tau}$$

which holds for any $\tau \in \Re^+$. The remaining calculations to obtain $B_1(x)$ do not appear to simplify significantly.

[6]Although time for evaluation of ν and λ increases with τ and n, total time for Runge-Kutta solution need not. Rather, time to solution seems to be determined mainly by the degree of curvature near the origin. Matlab and C routines are available from the author.

3 Beta Prior, Negative Log Gamma Signal (BNLG)

It is often desirable to assume that the value of the good has bounded support. For example, the discount price of a Treasury bill must always be in $[0,1]$. An appealing and computationally convenient assumption is that the prior on V is the Beta(α,β) distribution for $\alpha>0,\beta>0$. Let the private signals be conditionally independent and distributed "negative log Gamma," i.e., such that $-\ln(X)$ given $V=v$ is distributed Gamma$(\tau,\tau v)$. It is easily checked that the signals are affiliated and increase in expectation with the realization of V, so the M-W framework applies.

Application of Bayes' rule gives the posterior of V as a confluent hypergeometric distribution (see Appendix A.2), which is conjugate for gamma-distributed signals. Remaining calculations are similar to those in Section 2. The conditional c.d.f. of order statistic Y_1 is expanded as a finite sum,

$$\begin{aligned} G_{Y_1}(Y_1=y|V=v) &= G(X=y|V=v)^{n-1} = e_{\tau-1}(-\tau v\ln(x))^{n-1}\exp((n-1)\tau v\ln(x)) \\ &= \exp((n-1)\tau v\ln(y))\sum_{j=0}^{(\tau-1)(n-1)} c_j^{(\tau-1,n-1)}(-\tau v\ln(y))^j \end{aligned}$$

where $c_j^{(k,n)}$ is (as above) the coefficient on z^j in the polynomial expansion of $e_k(z)^n$. The counterpart in this specification of the ψ function in the GIG specification is

$$\begin{aligned} &\theta(x,a,b,\tau,m,n) \equiv \mathrm{B}(a+(n-m)\tau,b) \\ &\quad\cdot\sum_{j=0}^{m\cdot(\tau-1)} c_j^{(\tau-1,m)}(-\tau\ln(x))^j\frac{(a+(n-m)\tau)_j}{(b+a+(n-m)\tau)_j}M(b,b+a+(n-m)\tau+j,-n\tau\ln(x)) \end{aligned} \tag{6}$$

where $M(a,b,z)$ is the confluent hypergeometric function (see Abramowitz and Stegun, eds 1968,

§13) and B is the beta function. Straightforward calculations, outlined in Appendix A.2, give

$$B_2(x) = \nu(x,x) = \frac{\theta(x,\alpha+1,\beta,\tau,n-2,n)}{\theta(x,\alpha,\beta,\tau,n-2,n)} \tag{7}$$

$$\lambda(x,x) = \frac{(n-1)\tau^{\tau}}{\Gamma(\tau)} \frac{(-\ln(x))^{\tau-1}}{x} \frac{\theta(x,\alpha,\beta,\tau,n-2,n)}{\theta(x,\alpha,\beta,\tau,n-1,n)}. \tag{8}$$

Although there does not appear to be a closed form solution for $B_1(x)$, numerical solution can be obtained with only minor difficulty. At the boundary point $\underline{x} = 0$, $\nu(0,0) = 0$ so $B_1(0) = B_2(0) = 0$.[7] The gradient for $B_1(x)$ is infinite at $x = 0$, so Runge-Kutta must be initialized at some small $\epsilon > 0$, say $\epsilon = 10^{-8}$.[8] On a SparcStation5, solution in Matlab requires three to 25 seconds for parameter values in the range explored in Section 5.

Special case: When $\tau = 1$, the conditional distribution of the signal simplifies to the power function distribution, $G(x|v) = x^v$, and $\nu(x,x)$, $\lambda(x,x)$ and $L(t|x)$ have simple closed form expressions. The integral in equation (2) does not appear to have a closed form solution even in this special case, so $B_1(x)$ still must be solved numerically.

4 Comparative Statics for Bidding Strategies

The complexity of CV auctions for both the bidder and the theorist is inherent in the bidders' statistical filter. If each bidder receives an unbiased private signal of the true value, then the winning bidder's signal is an order statistic of the set of private signals, and thus biased upwards. In order to avoid winner's curse, the bidder must interpret her private information as if she had the highest of the n bidders' signals. Some intuition for this complicated Bayesian problem can be

[7]As $x \to 0$, the argument $-n\tau\ln(x)$ goes to infinity in the argument of the hypergeometric function in equation (6). Apply the limiting form given in Abramowitz and Stegun, eds (1968, 13.1.4) to show that ν must go to zero.

[8]The starting value $B_1(\epsilon)$ is bounded by $0 = \nu(0,0) < B_1(\epsilon) < \nu(\epsilon,\epsilon)$, and it makes very little difference to the overall solution whether the lower or upper bound is used.

gained from exploring how bidding strategies change with changes in the parameters.

Figure 1 shows the effect on $B_1(x)$ in the GIG specification of changes in the parameters (α, β, τ, n). The upper right panel shows the effect of changing the precision of private signals while holding constant $\alpha = 8, \beta = 2, n = 5$. Increasing τ makes the signal more precise, and so increases the weight given to the signal. Therefore, $B_1(x)$ twists counter-clockwise so that the bid for low signal values decreases and for high signal values increases. The upper left panel shows the effect of varying n. Increasing n increases the potential for winner's curse, which should lower bids, but also increases competition among the bidders. At low signal values, expected bidder rents are low, so the former effect dominates. At higher signal values, expected rents to the winner increase, so the competition effect becomes stronger. Therefore, the bidding strategy twists counter-clockwise.

It is convenient to re-parameterize (α, β) as $(\beta\mu, \beta)$, so that the mean is μ and changing β changes only the variance μ/β.[9] The lower left panel shows that increasing μ increases $B_1(x)$. The effect of changing β, which acts as a precision parameter for the prior, is the opposite of the effect of increasing τ. As β increases, the bidder places greater weight on the public information relative to private signals. Therefore, the lower right panel shows that $B_1(x)$ twists clockwise so that the bid for low signal values increases and for high signal values decreases.

I find similar comparative statics for the GIG second-price auction and the BNLG first- and second-price auctions.

5 Comparative Statics for Bidder Profits and Seller Revenue

Common value auctions do not, in general, yield neat comparative statics for bidder profits and seller revenue. It might be expected, for example, that expected seller revenue strictly increases

[9] If (α, β) is re-parameterized as $(\mu^2/\sigma^2, \mu/\sigma^2)$, then mean μ can be varied independently from variance σ^2. However, the resulting comparative statics are less clean because changing either μ or σ^2 changes the *shape* of the prior distribution in unintuitive ways.

Figure 1: Bidding Strategies for GIG First-Price Model

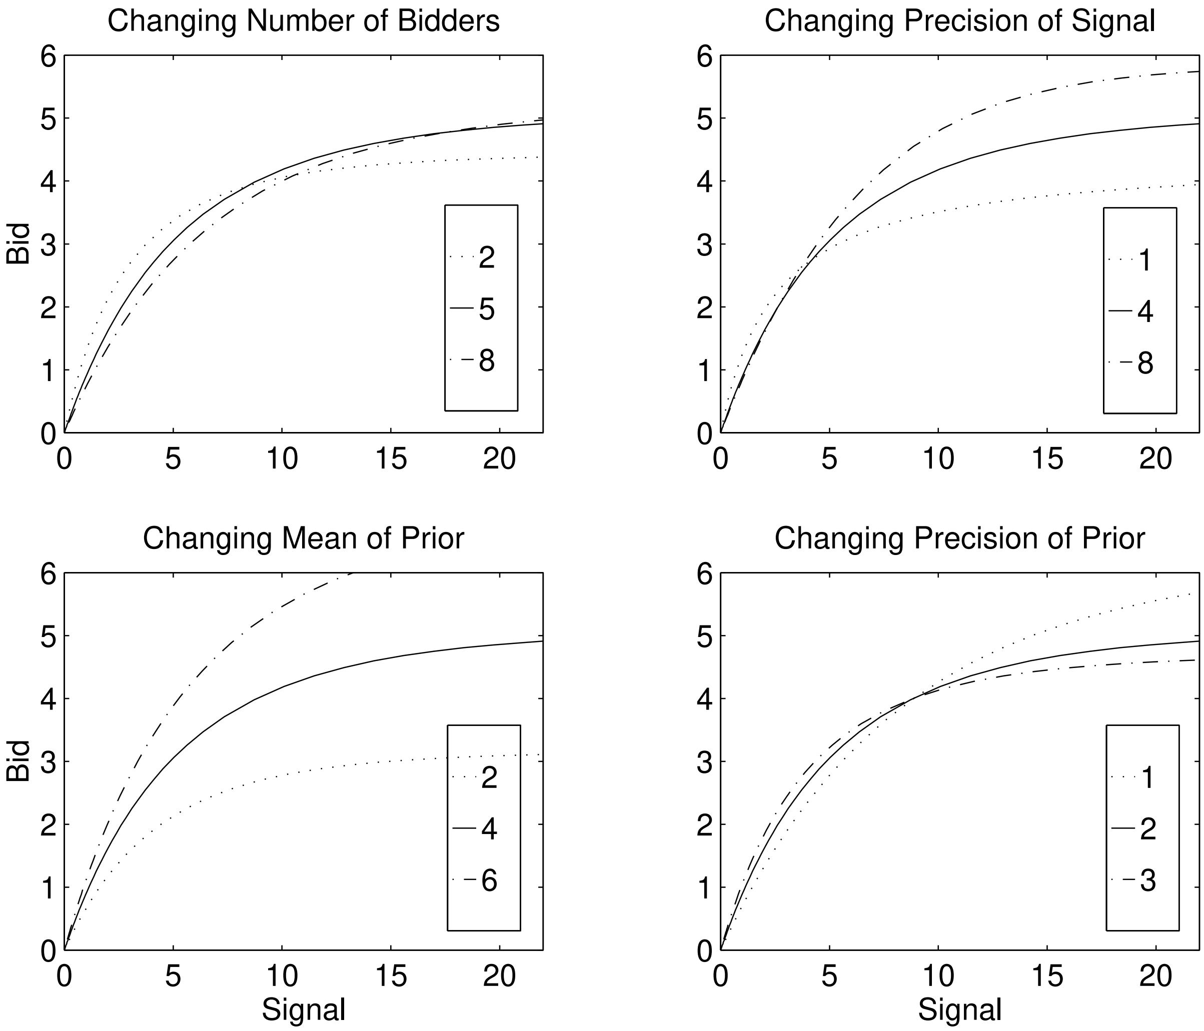

In the upper left panel, $n \in \{2, 5, 8\}$ with other parameters held constant at $\alpha = 8, \beta = 2, \tau = 4$. In the upper right panel, $\tau \in \{1, 4, 8\}$ with other parameters held constant at $\alpha = 8, \beta = 2, n = 5$. In the lower left panel, α is set to $\beta\mu$ for $\mu \in \{2, 4, 6\}$ with $\beta = 2, n = 5, \tau = 4$ are held constant. In the lower right panel, α is set to $\beta\mu$ for $\beta \in \{1, 2, 3\}$ with $\mu = 4, n = 5, \tau = 4$ held constant.

with the number of competing bidders. While it is true that competition must drive expected seller revenue to $E[V]$ as $n \to \infty$, Matthews (1984) shows that seller revenue falls with n at low values of n in his Pareto example. Similarly, intuition from the mechanism design literature might lead one to expect a bidder's expected rent to increase with the magnitude of his signal. Here too, counter-examples are available. In this section, I explore these comparative statics for the GIG and BNLG specifications, and find no evidence of pathological behavior.

In the first-price auction, the bidder's expected profit for a given signal x is

$$\begin{aligned}
\pi_1(x) &= \int_0^x (\nu(x,y) - B_1(x)) g_{Y_1}(y|x)dy \\
&= \int_0^x \int_V v(g_{Y_1}(y|v)f(v|x)/g_{Y_1}(y|x))dv g_{Y_1}(y|x)dy - B_1(x)G_{Y_1}(x|x) \\
&= \int_V vG(x|v)^{n-1}f(v|x)dv - B_1(x)G_{Y_1}(x|x).
\end{aligned}$$

Under the GIG specification, this expression reduces to

$$\pi_1(x) = \left(\frac{\beta x + \tau}{\beta x + n\tau}\right)^{\alpha+\tau} \left(\frac{(\alpha+\tau)x}{\beta x + n\tau}\psi(x, \alpha+1, \beta, \tau, n-1, n) - B_1(x)\psi(x, \alpha, \beta, \tau, n-1, n)\right).$$

It is difficult to sign the derivative $\pi_1{}'(x)$ analytically, but in numerical explorations it appears always to be positive. In Figure 2, I plot bidder profit as a function of x for parameter values used in Figure 1. In all cases, the bidder's expected profit increases with the level of the signal. Profit decreases with the number of competing bidders and increases with the precision of private information.[10]

[10] For the parameter values used in the right window of Figure 2, $\pi_1(x)$ increases at least through $\tau = 20$. Intuition suggests, however, that $\pi_1(x)$ must eventually decline with τ. At $\tau \to \infty$, signals become perfectly precise, so the true value becomes common knowledge. In the limit, the bidders face Bertrand competition, and so receive zero rent.

Figure 2: Comparative Statics for Bidder Profits

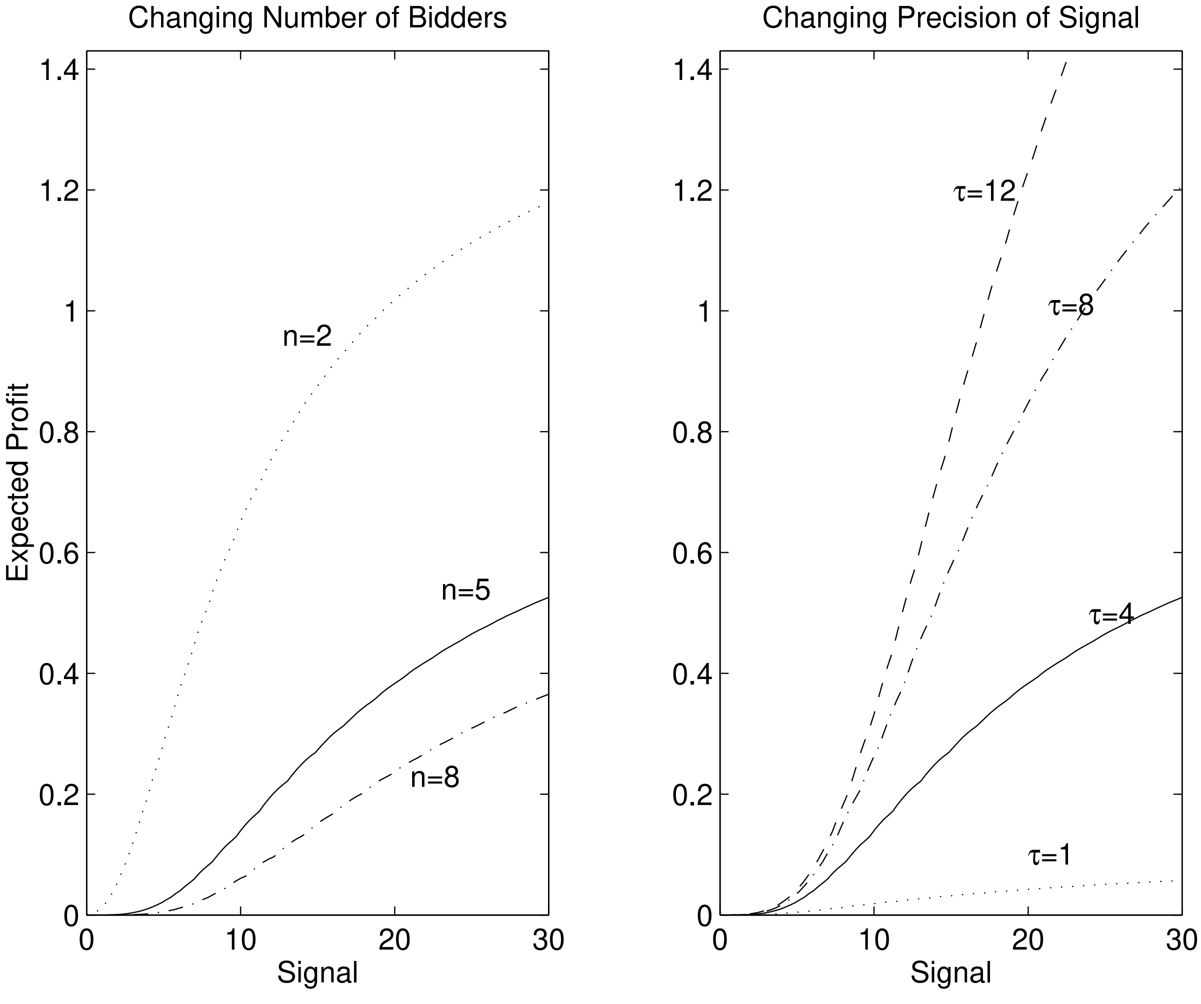

In the left panel, $n \in \{2, 5, 8\}$ with other parameters held constant at $\alpha = 8, \beta = 2, \tau = 4$.
In the right panel, $\tau \in \{1, 4, 8, 12\}$ with other parameters held constant at $\alpha = 8, \beta = 2, n = 5$.

Under the BNLG specification, $\pi_1(x)$ takes the form

$$\pi_1(x) = \frac{x^{\tau\cdot(n-1)}(\theta(x,\alpha+1,\beta,\tau,n-1,n) - B_1(x)\theta(x,\alpha,\beta,\tau,n-1,n))}{\mathrm{B}(\alpha+\tau,\beta)M(\beta,\alpha+\tau+\beta,-\tau\ln(x))}$$

and displays the same comparative statics as the GIG case. In the second-price auction, the bidder's expected profit for a given signal x is

$$\pi_2(x) = \int_0^x (\nu(x,y) - B_2(y))g_{Y_1}(y|x)dy.$$

This integral appears not to have a closed form solution under either specification, but can be solved by quadrature.

Expected seller revenue is in practice most easily estimated by Monte Carlo simulation. For each specification, I estimate $E[R_j]$ for $j \in \{1,2\}$, $n \in \{2,3,4,6\}$, $\tau \in \{1,2,3,5,8\}$, and a variety of α, β pairs.[11] Without exception, I find that $E[R_j]$ increases with n.

6 Reserve Prices

Reserve prices are commonly employed by sellers to raise expected revenue. Extension of the M-W model to include reserve prices is useful both for policy purposes, e.g., to approximate an optimal reserve price, and for structural estimation of the model when the auction data include an announced reserve price. Reserve prices can be accommodated under both the GIG and BNLG specifications with little additional computational complexity.

The announcement of a reserve price r induces potential bidders to drop out when private information indicates that the good is worth less than the reserve price. Given r, let $x^*(r)$ be the

[11] For GIG, I set $(\alpha,\beta) = (\mu\xi,\xi)$ for each combination of $\mu \in \{0.5,1,2,4,8\}$ and $\xi \in \{1,2,4,8\}$. For BNLG, I set $(\alpha,\beta) = (\mu\xi,(1-\mu)\xi)$ for each combination of $\mu \in \{0.05,0.20,0.50,0.8,0.95\}$ and $\xi \in \{1,4,8,25\}$. I draw $T = 30000$ random auctions for each set of parameters, and take average revenue.

screening level for x below which the bidder will drop out. This is given in M-W as

$$\begin{aligned} r &= E[V|X_1 = x^*, Y_1 < x^*] = \int_V v f(v|X_1 = x^*, Y_1 < x^*) dv \\ &= \frac{\int_V v G(x^*|v)^{n-1} g(x^*|v) f(v) dv}{\int_V G(x^*|v)^{n-1} g(x^*|v) f(v) dv}. \end{aligned}$$

For the GIG specification, x^* has implicit solution

$$r = \frac{(\alpha+\tau)x^*}{\beta x^* + n\tau} \frac{\psi(x^*, \alpha+1, \beta, \tau, n-1, n)}{\psi(x^*, \alpha, \beta, \tau, n-1, n)}.$$

For the BNLG specification, the solution is

$$r = \frac{\theta(x^*, \alpha+1, \beta, \tau, n-1, n)}{\theta(x^*, \alpha, \beta, \tau, n-1, n)}.$$

Note the close structural similarity of these screening equations both to one another and to the respective equations for $\nu(x, x)$.

Once the screening level is obtained, solution of the equilibrium strategies is straightforward. In the first-price auction, one solves differential equation (1) with initial value $B_1(x^*) = r$. The second-price solution is unchanged, except that bidders with $x < x^*$ drop out.

Conclusion

This paper introduces two sets of distributional assumptions that yield computationally convenient equilibrium strategies in the first- and second-price sealed bid common value auction. These results present new opportunities for structural estimation of auction models. Both specifications can be computed quickly and accurately, and can be parameterized flexibly. Moreover, both specifications are plausible representations of important classes of CV auctions. The gamma prior in the GIG specification is convenient for modeling an auction of a good with a natural lower bound on value (typically zero) but not a natural upper bound, e.g., an oil exploration tract. The beta prior in the BNLG specification is appropriate when there are natural lower and upper bounds, e.g., a Treasury bill. Both the gamma and beta allow for a density that is roughly bell-shaped and centered near the mean. Both specifications easily accomodate a reserve price, and both appear to be well behaved with respect to comparative statics for bidder profits and seller revenue.

The restriction of precision parameter τ to the set of positive integers can be overcome easily. The conditional density of the signals changes smoothly with τ, so the bidding strategy must as well. Therefore, bidding strategies for non-integer values of τ can be approximated by interpolation.

Some further extensions are possible. In some CV auctions, multiple units are offered for simultaneous sale. The equilibrium bidding strategy is similar in form to the single unit case, but with the order statistic Y_m (where m is the number of units) replacing Y_1 in functions ν and λ (see Bikhchandani and Huang 1989). The GIG and BNLG specifications can be extended to the multiple unit case, although possibly with numerical problems.[12] Both specifications can also be used to derive computationally feasible equilibrium solutions when bidders have constant absolute risk aversion.

A Intermediate calculations

This appendix lists the most important intermediate results not reported in the main text. As above, V represents the true value with c.d.f. F (p.d.f. f), X_1 the signal of our representative bidder (without loss of generality) with c.d.f. G (p.d.f. g), and Y_1 the highest of the other signals $X_2, \ldots, X_n$. To simplify notation, lower case x and y will stand for realizations of X_1 and Y_1.

[12] Specifically, the ν and λ functions will contain negative, as well as positive, terms. Round-off error occurs in taking the difference of two terms of roughly equal magnitude.

A.1 GIG specification

The conditional and unconditional distributions of the bidder's signal X_1 are given by

$$
\begin{aligned}
g(x|v) &= \frac{(\tau v)^{\tau}}{\Gamma(\tau)} x^{-(\tau+1)} \exp(-\tau v x^{-1}) \\
G(x|v) &= e_{\tau-1}(\tau v x^{-1}) \exp(-\tau v x^{-1}) \\
g(x) &= \int_0^{\infty} g(x|v) f(v) dv = \frac{\beta^{\alpha}\tau^{\tau}}{\mathrm{B}(\tau,\alpha)} \frac{x^{\alpha-1}}{(\beta x+\tau)^{\alpha+\tau}} \\
G(x) &= \int_0^{\infty} G(x|v) f(v) dv = \left(\frac{\beta x}{\beta x+\tau}\right)^{\alpha} \psi(x,\alpha,\beta,\tau,1,1)
\end{aligned}
$$

Generalizing the ψ function to

$$
\tilde{\psi}(y,x,a,b,\tau,m,n) \equiv \sum_{j=0}^{m\cdot(\tau-1)} c_j^{(\tau-1,m)} (a+(n-m)\tau)_j \left(\frac{\tau}{y(b+\tau x^{-1}+(n-1)\tau y^{-1})}\right)^{j}
$$

allows compact expression of the distribution of Y_1 given X_1 as

$$
\begin{aligned}
g_{Y_1}(y|x) &= \int_0^{\infty} g_{Y_1}(y|v) f(v|x) dv = (n-1)\int_0^{\infty} G_{Y_1}(y|v)^{n-2} g_{Y_1}(y|v) f(v|x) dv \\
&= \frac{(n-1)\tau^{\tau}(\beta+\tau x^{-1})^{\alpha+\tau}}{y^{\tau+1}\Gamma(\alpha+\tau)\Gamma(\tau)} \\
&\quad \cdot \sum_{j=0}^{(n-2)(\tau-1)} c_j^{(\tau-1,n-2)} (\tau y^{-1})^{j} \int_0^{\infty} v^{\alpha+2\tau+j-1} \exp(-(\beta+\tau x^{-1}+(n-1)\tau y^{-1})v) dv \\
&= \frac{(n-1)\tau^{\tau}}{y^{\tau+1}\mathrm{B}(\alpha+\tau,\tau)} \frac{(\beta+\tau x^{-1})^{\alpha+\tau}}{(\beta+\tau x^{-1}+(n-1)\tau y^{-1})^{\alpha+2\tau}} \tilde{\psi}(y,x,\alpha,\beta,\tau,n-2,n) \\
G_{Y_1}(y|x) &= \int_0^{\infty} G_{Y_1}(y|v) f(v|x) dv = \int_0^{\infty} G(y|v)^{n-1} f(v|x) dv \\
&= \left(\frac{\beta+\tau x^{-1}}{\beta+\tau x^{-1}+(n-1)\tau y^{-1}}\right)^{\alpha+\tau} \tilde{\psi}(y,x,\alpha,\beta,\tau,n-1,n)
\end{aligned}
$$

Noting that $\tilde{\psi}(x,x,a,b,\tau,m,n) = \psi(x,a,b,\tau,m,n)$ and re-arranging lead to the simple forms for $\nu(x,x)$ and $\lambda(x,x)$ in Section 2.

A.2 BNLG specification

The mechanics of the BNLG calculations are quite similar to those of the GIG case. The true value V is integrated out of distributions for the signals by employing the identity

$$\int_0^1 v^{a-1}(1-v)^{b-1}\exp(zv)dv = \exp(z)\mathrm{B}(a,b)M(b,a+b,-z) \tag{9}$$

where $M(a,b,z)$ is the confluent hypergeometric function (also written ${}_1F_1$) and B is the beta function (see Abramowitz and Stegun, eds 1968, 13.2.1,13.1.27). The conditional and unconditional distributions of the bidder's signal X_1 are given by

$$\begin{aligned}
g(x|v) &= \frac{(\tau v)^\tau}{\Gamma(\tau)}\frac{(-\ln(x))^{\tau-1}}{x}\exp(\tau v\ln(x)) \\
G(x|v) &= e_{\tau-1}(-\tau v\ln(x))\exp(\tau v\ln(x)) \\
g(x) &= \frac{\tau^\tau}{\Gamma(\tau)}\frac{(\alpha)_\tau}{(\alpha+\beta)_\tau}(-x\ln(x))^{\tau-1}M(\beta,\beta+\alpha+\tau,-\tau\ln(x)) \\
G(x) &= \frac{x^\tau}{\mathrm{B}(\alpha,\beta)}\theta(x,\alpha,\beta,\tau,1,1)
\end{aligned}$$

The form of the posterior $f(v|x)$ may be of independent interest. Define the *confluent hypergeometric distribution* by the p.d.f.

$$f_{\mathrm{CH}}(z|a,b,s) \equiv \frac{z^{a-1}(1-z)^{b-1}}{\mathrm{B}(a,b)}\frac{\exp(-sz)}{M(a,a+b,-s)}.$$

This generalizes the beta density using the confluent hypergeometric equation (9) in the way Armero and Bayarri (1994) use the ${}_2F_1$ function to generalize the beta to the Gaussian hypergeometric distribution. It is straightforward to show that if $Z \sim \mathrm{Beta}(\alpha,\beta)$ and signal $S|(Z=z) \sim \mathrm{Gamma}(\tau,z)$, then the posterior $Z|(S=s) \sim \mathrm{CH}(\alpha+\tau,\beta,s)$. Furthermore, the CH distribution itself is conjugate for gamma signals, so that if Z is distributed $\mathrm{CH}(\alpha,\beta,s_1)$ and signal $S_2|(Z=z) \sim \mathrm{Gamma}(\tau,z)$, then $Z|(S_2 = s_2) \sim \mathrm{CH}(\alpha+\tau, \beta, s_1+s_2)$. The CH distribution appears to be new to the literature.

Applying the definition of the CH distribution, $f(v|x)$ may be written as

$$f(v|x) = g(x|v)f(v)/g(x) = f_{\mathrm{CH}}(v|\alpha+\tau, \beta, -\tau\ln(x)).$$

The conjugate property of the CH distribution for gamma signals underpins the tractability of the BNLG specification.

Generalizing the θ function to

$$\tilde{\theta}(y, x, a, b, \tau, m, n) \equiv \mathrm{B}(a+(n-m)\tau, b)$$
$$\cdot \sum_{j=0}^{m\cdot(\tau-1)} c_j^{(\tau-1,m)}(-\tau\ln(y))^j \frac{(a+(n-m)\tau)_j}{(b+a+(n-m)\tau)_j} M(b, b+a+(n-m)\tau+j, -\tau(\ln(x)+(n-1)\ln(y))).$$

allows compact expression of the distribution of Y_1 given X_1 as

$$\begin{aligned} g_{Y_1}(y|x) &= \frac{(n-1)\tau^\tau}{\Gamma(\tau)} \frac{(-\ln(x))^{\tau-1}}{x} \frac{y^{(n-1)\tau}\tilde{\theta}(y, x, \alpha, \beta, \tau, n-2, n)}{\mathrm{B}(\alpha+\tau, \beta)M(\beta, \beta+\alpha+\tau, -\tau\ln(x))} \\ G_{Y_1}(y|x) &= \frac{y^{(n-1)\tau}\tilde{\theta}(y, x, \alpha, \beta, \tau, n-1, n)}{\mathrm{B}(\alpha+\tau, \beta)M(\beta, \beta+\alpha+\tau, -\tau\ln(x))} \end{aligned}$$

Noting that $\tilde{\theta}(x, x, a, b, \tau, m, n) = \theta(x, a, b, \tau, m, n)$ and re-arranging lead to the simple forms for $\nu(x, x)$ and $\lambda(x, x)$ in Section 3.

References

Abramowitz, Milton and Irene A. Stegun, eds, *Handbook of Mathematical Functions* number 55. In 'Applied Mathematics Series.', National Bureau of Standards, 1968.

Armero, C. and M.J. Bayarri, "Prior assessments for prediction in queues," *The Statistician*, 1994, *43* (1), 139–153.

Bikhchandani, Sushil and Chi-fu Huang, "Auctions with Resale Markets: An Exploratory Model of Treasury Bill Markets," *Review of Financial Studies*, 1989, (2), 311–340.

Engelbrecht-Wiggans, Richard and Robert J. Weber, "On the Non-Existence of Multiplicative Equilibrium Bidding Strategies," Discussion Paper 523, Cowles Foundation April 1979.

Laffont, Jean-Jacques and Quang Vuoung, "Structural Econometric Analysis of Descending Auctions," *European Economic Review*, April 1993, *37* (2/3), 329–341.

Levin, Dan and James L. Smith, "Some Evidence on the Winner's Curse: Comment," *American Economic Review*, 1991, *81* (1), 370–375.

Matthews, Steven, "Information Acquisition in Discriminatory Auctions," in Marcel Boyer and Richard E. Kihlstrom, eds., *Bayesian Models in Economic Theory*, Vol. 5 of *Studies in Bayesian Econometrics*, Amsterdam: Elsevier Science Publishers, 1984, pp. 181–207.

McAfee, R. Preston and John McMillan, "Auctions and Bidding," *Journal of Economic Literature*, Summer 1987, *XXV*, 699–738.

Milgrom, Paul R., "Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, 1989, *3* (3), 3–22.

___ **and Robert J. Weber**, "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1982, (50), 1089–1122.

Paarsch, Harry J., "Deciding Between the Common and Private Value Paradigms in Empirical Models of Auctions," *Journal of Econometrics*, 1992, *51*, 191–215.

Vincent, Daniel R., "Bidding Off the Wall: Why Reserve Prices May Be Kept Secret," *Journal of Economic Theory*, April 1995, *65*, 575–584.