BID DISTRIBUTION AND TRANSACTION COSTS

Mats A. Bergman and Johan Y. Stake

ABSTRACT: We show that with costly but unlimited entry and ex ante symmetric firms the spread between the lowest and the second lowest bid will be a lower-bounds estimate for seller-side transaction (entry) costs under reasonable distributional assumptions on production costs. The bid spread approximates the winner's expected profit and with costly entry the ex-post profit of winning balances the aggregated costs of entering into the procurement auction. Using empirical data from Swedish procurements we estimate total seller-side transaction costs to be, on average, at least 5-6 percent of tender values and possibly as large as 10-15 percent. Given four bidders in the average procurement, this means per-firm (per-bidder) transaction costs of 1.5 – 3 percent of the bid value. This is in the same order of magnitude as previous estimates, which were based on surveys and time studies.

INTRODUCTION

Following Coase's seminal work (1937) a huge but mainly theoretical or qualitatively empirical literature on transaction costs has evolved. Much of the literature has focused on the interplay between institutional forms and transaction costs, based on the premise that competitive pressure will favour institutional forms that successfully limit transaction costs. From Shelanski and Klein's (1995) survey of empirical research on transaction costs it is clear that the main thrust of the empirical work is to use proxies for the level of transaction costs to explain institutional form. Among the most important such proxies are the transaction specificity of investments, the level of uncertainty, transaction complexity and transaction frequency.[1]

Few studies specifically try to *measure* the transaction costs and even fewer try to measure transaction costs in public procurement. Based on surveys of time use, the EU Commission (2011) estimates the average transaction costs for a public procurement to € 28 000. The EEA authorities spend an average 22 days per procurement and the bidding firms require 16 days. With approximately six bids per procurement total costs can be estimated from reasonable assumptions on staff cost per unit of time. Costs are higher for works (e.g., building and construction) contracts than for service or supplies (goods) contracts.[2] Overall, the procurement transaction costs correspond to little more than one percent of the value of public procurement, but for relatively low-value procurements the transaction costs can be a sizeable share of all costs.[3]

---

[1]See also Holm, 2011.

[2] Holm reports similar values for the authorities' costs in Swedish public procurements.

[3]The EU study compares the estimated per-procurement cost with the *average* procurement value. It may, in fact, be more reasonable to relate the per-procurement cost with the *median* procurement value as discussed below.

We use bid data and bid spreads to estimate firms' procurement costs. The analysis is based on the assumptions of unlimited but costly entry and ex-ante symmetry between the participating firms. Firms will enter procurement auctions (submit bids) until the expected profit from doing so is zero – i.e., until economic profits are completely dissipated. A lower limit for the expected profit of the winner can be estimated from bid spreads.

PRELIMINARIES

We interpret the cost of preparing and submitting a bid, the bidding transaction cost, as a sunk entry cost that firms must pay to participate in a procurement auction. If firms must pay an entry cost to participate in a procurement market and if the firms are ex ante symmetric, we can expect that firms will enter the market and submit bids up to the point where there is no more profit from entering the market. At this point the expected profit, conditional on winning the contract, just compensates for the bidding cost. If the procurement attracts n bidders and if the winner's expected profit (not considering bidding costs) is π, then the bidding cost should be approximately π/n.

The number of bidders, n, is easy to observe. The profit cannot be observed directly, since we cannot observe the firms' costs, but we are able to infer something about the costs from the bidding behaviour. For example, if all firms add the same margin on top of their cost, then the difference between the winner's bid and the second-lowest bid will be equal to the difference between the two firms' costs.

Furthermore, from the revenue-equivalence theorem we know that the expected bid of the winning firm in a first-price auction will be equal to the expected cost of the second-lowest bidder.[4] Therefore, if

---

[4] Assuming risk neutral bidders, that individual values (or signals) are independently drawn from a common strictly increasing and atom-less distribution, that the auction mechanism always allocates the contract to the lowest-cost bidder and that a bidder with the highest feasible cost expects

bid-cost margins are equal for all firms the difference in bids would be a good estimate of the profit of the winner.

Unfortunately, firms' bid-cost margins will *not* be equal, since the optimal margin will depend on the firm's cost. However, for many distributional assumptions the bid spread will provide a lower-bounds estimate of the winner's expected profit. To say more we now turn to some specific cost distributions.

Recovering bidding costs from bid spreads and bid-distribution assumptions

We explore the relation between bid spreads and bidding costs under different distributional assumptions, starting with the uniform distribution. We assume that the bidding firms know the distribution from which their individual costs will be drawn and that they must pay the bidding cost before the cost is realized. The costs are private information but the distribution from which the costs are drawn is common knowledge. After the costs have been realized the bidders decide the bids.

Proposition 1: Uniform distribution with known support. With ex ante symmetric risk-neutral bidders that must pay a sunk entry cost before their production costs are independently drawn from the uniform distribution over $[\underline{c},\bar{c}]$,the expected profit of the winning bidder(gross of bidding costs), π_1, is related to the expected bids according to

$$E[\pi_1]=E[b_1]-E[c_1]=\frac{n}{n-1}(E[b_2]-E[b_1])$$

where b_i is the bid of the *i*:th lowest bid, c_i is the cost of that firm and n is the number of bidders. The per-firm bidding cost d is related to the expected cost according to

---

zero surplus; the two latter requirements are met by all standard auction formats. See Klemperer, 2004, p. 17 and 43.

$$E[\pi_1]/(n+1) \geq d \geq E[\pi_1]/n$$

Hence, $d \geq E[b_2-b_1]/n$

Proof. We use the revenue equivalence theorem and its consequence, that the expected profit of the winner will be the same irrespective of the auction format. In a second-price procurement auction each bidding firm i will submit a bid equal to its cost, such that $b_i = c_i$ so that the expected price is equal to the expected bid of the second-lowest bidder. The expected value of the k:th highest value among n independent draws from the uniform distribution on $[\underline{c},\bar{c}]$ is[5]

$$\underline{c} + \left(\frac{n+1-k}{n+1}\right)(\bar{c}-\underline{c}) \tag{1}$$

Hence, the expected lowest and second-lowest cost, respectively, will be

$$E[c_1] = \underline{c} + \left(\frac{1}{n+1}\right)(\bar{c}-\underline{c}) \tag{2}$$

$$E[c_2] = \underline{c} + \left(\frac{2}{n+1}\right)(\bar{c}-\underline{c}) \tag{3}$$

By the revenue equivalence theorem the expected profit of the winner in a first-price procurement auction will be

$$E[\pi_1] = E[c_2] - E[c_1] = \underline{c} + \left(\frac{2}{n+1}\right)(\bar{c}-\underline{c}) - \underline{c} - \left(\frac{1}{n+1}\right)(\bar{c}-\underline{c}) = \frac{1}{n+1}(\bar{c}-\underline{c}) \tag{4}$$

[5] See Klemperer, 2004, p. 52.

For a specific cost realization, c_1, the lowest-cost firm should, by the revenue equivalence theorem, bid so that the expected profit will be the same as in a second-price procurement. I.e., conditional on winning, bidder *i* will bid a value equal to its expected payment in a descending auction. Given that c_1 is the lowest value, the other *n-1*bidders' costs will be uniformly distributed on $[c_1, \bar{c}]$so that, by equation 1,

$$E[c_2 | c_1] = c_1 + \left(\frac{n-1+1-(n-1)}{n-1+1} \right)(\bar{c} - c_1) = c_1 + \frac{1}{n}(\bar{c} - c_1) \quad (5)$$

It follows from the revenue equivalence theorem that

$$b_1(c_1) = c_1 + \frac{1}{n}(\bar{c} - c_1) \quad (6)$$

The expected profit of the winning bidder is given by

$$E[\pi_1 | c_1] = \frac{1}{n}(\bar{c} - c_1) \quad (7)$$

The second lowest bidder will bid as if its cost were, in fact, the lowest cost. Again using the revenue-equivalence theorem we have:

$$b_2 = c_2 + \frac{1}{n}(\bar{c} - c_2) \quad (8)$$

Inserting equation (3) in equation (8) and taking expectations yields

$$E[b_2] = E[c_2] + \frac{1}{n}(\bar{c} - E[c_2]) = \underline{c} + \left(\frac{2}{n+1} \right)(\bar{c} - \underline{c}) + \frac{1}{n}\left(\bar{c} - \underline{c} - \left(\frac{2}{n+1} \right)(\bar{c} - \underline{c}) \right) = \\ = \frac{(n-1)^2 \underline{c} + (3n-1)\bar{c}}{n(n+1)} \quad (9)$$

It follows that

$$E[b_2 - b_1] = E[b_2 - c_2] = \frac{(n-1)^2\underline{c} + (3n-1)\bar{c}}{n(n+1)} - \left[\underline{c} + \left(\frac{2}{n+1}\right)(\bar{c} - \underline{c})\right] = \frac{(n-1)(\bar{c} - \underline{c})}{n(n+1)}$$

(10)
Comparing equations (4) and (10) we see that

$$E[\pi_1] = E[b_1] - E[c_1] = \frac{n}{n-1}\left(E[b_2] - E[b_1]\right) \qquad (11)$$

This proves the first part of the proposition. The second part of the proposition follows directly from profit maximization, from the assumption of ex ante symmetric firms and from the sunk cost of entry, d.◊

The expected profit of the winner will, for two reasons, be larger than the expected difference between the winning bid and the second lowest bid. First, the difference in bids will be smaller than the difference in costs, since firms with higher costs will optimally bid with a smaller mark up above the cost. Second, because of indivisibility the average ex-post profit may be strictly larger than the bidding cost. As *n* increases the expected difference in bids will approach the expected profit of the winning firm.

The intuition for the first effect is that the bids will depend on how the bidder's own cost is located in the range of possible costs and that this knowledge will impact on the optimal bidding strategy. If the own cost is low the bidder will optimally chose a relatively high margin; if the own cost is high the bidder will optimally chose a relatively low margin. The bidder trades the profit conditional on winning against the probability of winning. The higher the own cost the more densely packed are the other bidders *if* the firm indeed has the lowest cost.

It follows from Proposition 1 and equation (2) that, for a given bid spread, the equilibrium number of firms increases as the bid cost *d* decreases and that, for a given bid cost, the equilibrium number of firms grows as the range of the distribution increases. Furthermore, for the uniform distribution the number of bidders grows in proportion to the square root of the range of the distribution, $\underline{c} - \bar{c}$, divided by the per-bid transaction cost.

Proposition 2: Uniform distribution with moving frame. Assume that ex ante symmetric and risk-neutral bidders must pay a sunk entry cost *d* before their production costs *x*+*c* are realized, where the common cost *x* and individual costs *c* are independently drawn from the uniform distribution over $[\underline{x}, \bar{x}]$ and $[\underline{c}, \bar{c}]$, respectively, and where $[\bar{x} - \underline{x}]$ is large relative to $[\bar{c} - \underline{c}]$. Realized cost $x+c_i$ is observed by bidder *i* but remains unknown to the other bidders. Then the expected profit of the winning bidder gross of bidding costs, π_1, is related to the expected bids according to

$$E[\pi_1] = E[b_1] - E[c_1] = (E[b_2] - E[b_1])$$

with notation as above. The per-firm bidding cost *d* is related to the expected cost according to

$$E[\pi_1]/(n+1) \geq d \geq E[\pi_1]/n$$

Hence, $d \geq E[b_2 - b_1]/n$

Proof. Since $[\bar{x} - \underline{x}]$ is large relative to $[\bar{c} - \underline{c}]$ no inferences as to the size of the individual component *c* can be drawn from observing the realized cost *x*+*c*.[6] A bidder that observes its own costs $x+c_i$ must

[6]The value of *x* can be thought of as a common-value component and the uniform distribution over $[x + \underline{c}, x + \bar{c}]$ as a private-value component. The setting is still a private-value setting in the sense that the bidder will know its own cost for sure, so there is no winner’s curse. Instead of drawing inferences from the fact that the bidder wins to the size of its costs, the bidder can now draw inferences to the size of the *rivals’* costs. Being a

then assume that, conditional on winning the procurement, its individual cost component is equal to the expected individual cost of a winning bid. From equation (2) it follows that

$$E\left[x \middle| x+c_i, c_i < c_j\right] = x + c_i - \underline{c} - \left(\frac{1}{n+1}\right)\left(\bar{c} - \underline{c}\right)$$
$$E\left[c_i \middle| x+c_i, c_i < c_j\right] = \left(\frac{1}{n+1}\right)\left(\bar{c} - \underline{c}\right) \qquad (12)$$

for *1≤j≤n* and *j≠i*

Substituting $E\left[c_i \middle| x+c_i, c_i < c_j\right]$ and $E\left[x \middle| x+c_i, c_i < c_j\right]$ into equation (6) gives the optimal bid as

$$b_i(x+c_i) = x + c_i + \frac{1}{n+1}\left(\bar{c} - \underline{c}\right) \qquad (13)$$

That is, all bidders will add the same margin, given by the last term in equation (13) The proposition follows from subtracting the bid associated with cost c_j from the bid associated with cost c_i and from Proposition 1.◊

Since the bidder cannot know if its cost is relatively high or relatively low, it will bid as if it is equal to the expected lowest cost. Based on this assumption, no bidder can do better than to add the optimal bid-cost margin associated with that cost level. If all bidders reason the same way the difference in bids will equal the difference in costs.

We turn now to the normal distribution, since it is more reasonable to assume that the bids are normally distributed than to assume that they are uniformly distributed. However, explicit solutions for the order statistics of the normal distribution exist only for special cases.

---

winner is, in this setting, good news, since the conditional expected individual costs of the rivals will be higher than the unconditional expected individual costs.

For example, if n=2 with standard deviation σ and mean m the order statistics are[7]

$$E[c_1] = m - \frac{\sigma}{\sqrt{\pi}}$$
$$E[c_2] = m + \frac{\sigma}{\sqrt{\pi}} \tag{14}$$

The expected bid of the winning bidder will equal the expected cost of the bidder with the second lowest cost, so that $E[b_1$=$E[c_2]$. The expected bid of the losing bidder will be equal to the expected value of a draw from the normal distribution truncated from below at $E[c_2]$. That is[8]

$$E[b_2] = m + \sigma\lambda\left((c_2 - m)/\sigma\right)$$
$$\lambda(y) = \frac{\phi(y)}{1 - \Phi(y)} \tag{15}$$

where φ is the density function of the standard normal distribution and Φ is the corresponding cumulative distribution function. It follows immediately that

$$E[c_2] - E[c_1] = \pi_1 = \frac{2\sigma}{\sqrt{\pi}} \tag{16}$$

It is also easy to show that

7See Ahsanullah and Nevzorov , 2005, p 122.

8See entry for truncated normal distribution at Wikipedia.

$$E[b_2] - E[b_1] = \sigma \left[\frac{\phi\left(\frac{1}{\sqrt{\pi}}\right)}{1 - \Phi\left(\frac{1}{\sqrt{\pi}}\right)} - \frac{1}{\sqrt{\pi}} \right] \quad (17)$$

Although the probability density function can be expressed analytically, there is no simple way to represent the cumulative density function.[9]For σ=1, for example, the expected differences in bids will be approximately 0.62, while the expected profit will be approximately 1.13.[10]Harter (1961) has tabulated order statistics for the standard normal distribution. These values can be used to tabulate the expected bid spread if costs are normally distributed.

The low-cost bidder will bid equal to the second-lowest order statistics, $E[c_2]$. The second-lowest bidder will not, however, bid equal to the third-lowest order statistics, since that bidder will bid as if it were actually the low-cost bidder. Specifically, the second-lowest bidder will, in expectation, bid *less* than the third-lowest order statistics, since it will bid the lowest-order statistics for a normal distribution that is truncated from the left at its own cost. If we can establish that the difference between the third-lowest and the second-lowest order statistics is no larger than the difference between the lowest and the second-lowest, then we can conclude that the difference in bids will be smaller than the difference in costs

---

[9] Of course it can be represented as an integral of the probability density function; it can also be represented with an *erf* error function.

[10]The explicit function for the probability density function is $\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2}$ for the standard normal distribution and $\phi(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-m)^2}{2\sigma^2}}$ for the normal distribution with mean m and variance σ2. For the standard normal, the expected value of x, given that x is larger than $\frac{1}{\sqrt{\pi}}$, is

$\phi(1/\sqrt{\pi})/(1-\Phi(1/\sqrt{\pi})) \approx 1.1878$.

for the two lowest bidders. That means that the expected difference in bids between the lowest and second-lowest bidder will be lower than the expected profit of the winner, which, in turn, will be equal to the total transaction costs – for symmetric firms with costly entry.

Table 1. Expected bid spreads for standard normal distribution

Number of bidders	1st order statistics	2nd order statistics	3rd order statistics	Profit of winner*	Upper bound for bid spread**	Profit/bid spread (lower bound)
2	-0.564	0.564	n/a	1.128	n/a	n/a
3	-0.846	0	0.846	0.846	0.846	1
4	-1.029	-0.297	0.297	0.732	0.594	1.23
5	-1.163	-0.495	0	0.668	0.495	1.35
6	-1.267	-0.642	-0.202	0.625	0.440	1.42
7	-1.352	-0.757	-0.353	0.595	0.404	1.47
8	-1.424	-0.852	-0.473	0.572	0.379	1.51

* Equal to 2nd minus 1st order statistics
** Equal to 3rd minus 2nd order statistics
Order statistics from Harter (1961)

For n = 2 we know from above that the ratio between the expected profit and the expected bid spread is approximately 1.82 (=1.13/0.62). For n = 3 we know that the true expected bid spread will be *lower* than 0.846 and, hence, that the expected profit will be *larger* than the expected bid spread. According to the table this is true also for four to eight bidders and according to Harter's tabulation it is also true for up to 100 bidders. Harter only provides values for a few examples of more than 100 drawings from the normal distribution; also for these cases is it true that the difference between the second and first order statistics is larger than the difference between the third and the second.
Using differences in bids as an estimate of the winner's profit will, therefore, be a conservative estimate also for a normal distribution:

Proposition 3. Normal distribution. With ex ante symmetric risk-neutral bidders that must pay a sunk entry cost before their production costs are independently drawn from the same normal distribution the expected difference between the second lowest and the lowest bids is larger than n times the per-bidder transaction cost for all n up to at least 100.

Proof. The expected difference between the third lowest and the second lowest cost is smaller than the expected difference between the second lowest and the lowest cost for three to 99 bidders, according to the tabulation in Harder (1961).

By the revenue-equivalence theorem the expected bid of the second lowest bidder is equal to the first order statistics for *n-1* bidder for the normal distribution truncated from the left at the second-order statistics for *n* bidders for the normal distribution. This is less than the first order statistics for *n-2* bidders over the same distribution and this, in turn is equal to the third-order statistics for the original normal distribution. It follows that the expected bid of the second lowest bidder is lower than the expected third-lowest cost and hence, $d \geq E[b_2 - b_1]/n$.◊

GENERAL DISTRIBUTIONS

From the revenue-equivalence theorem, the expected bid of the lowest bidder in a first-price auction will be equal to the expected cost of the firm with the second-lowest cost, while the expected bid of the latter firm will be calculated as follows: Assume that the second-lowest bid were actually the lowest bid and then calculate the expected second-lowest bid in this hypothetical situation; this is what the second-lowest bidder will bid. That is

$$\begin{aligned} E[\pi_1] &= E[B_1 - C_1] = E[C_2 - C_1] \\ E[B_2 - B_1] &= E[C_2 | C_1 = E[C_2]] - E[C_2] \end{aligned} \tag{18}$$

The expected bid of the second-lowest bidder will tend to be closer to its expected costs than the expected bid of the lowest bidder will be to *its* cost for two reasons. First, if there is an upper limit to the distribution, since the second-lowest bidder is bidding under the assumption that all other bids are higher that his or her bid, the other firms' bids must be packed more densely over a smaller range. This is the effect seen in the uniform distribution as shown above.

Second, the difference can be smaller also because the probability density function typically assumes larger values closer to the median of the distribution and because the second-lowest cost will often be closer to the median. Then the expected difference between the own cost and the expected cost of the hypothetical second-lowest bidder will be smaller. This is roughly true to the extent that *f(x)/(1-F(x))* is rising with *x* – since the conditional average (average conditional on *x* being larger than *a*) is calculated as

$$\int_a^\infty \frac{xf(x)}{1-F(a)}dx \tag{19}$$

For example, to the left of the hump in a normal distribution *f(x)/(1-F(x))* is rising. Then the difference between the lowest and the second lowest bid will tend to *under*estimate the profit of the winning bid since the expected difference between the third lowest bid and the second lowest bid is smaller than the expected difference between the second lowest and the lowest bid.
Asymmetric bidders – a linear example

In practice bidders may not be symmetric. Assume that the dominant firm (or the incumbent) has a cost that is drawn from the uniform distribution over [0,2], while rival firms' costs are drawn from the uniform distribution over [1,2]. The optimal bid strategy of a dominant firm facing a single rival is

$$B_1 = 1.5 \quad \text{if} \quad C_1 \le 1$$
$$B_1 = \frac{C_1+2}{2} \quad \text{if} \quad C_1 > 1 \tag{20}$$

(We now let index 1 represent the dominant firm, rather than the low-cost bidder.) The rival's optimal strategy is

$$B_2 = \frac{C_2+2}{2} \tag{21}$$

The dominant firm will win with probability 0.75. The expected cost of the winner is

$$0.5 \cdot 0.5 + 0.5 \cdot \frac{4}{3} \approx 0.917$$

and the expected winning bid is

$$0.5 \cdot 1.5 + 0.5 \cdot \frac{5}{3} = \frac{19}{12} \approx 1.585$$

while the expected losing bid is

$$0.5 \cdot 0.5 \cdot \left(\frac{3}{2} + 2\right) + 0.5 \cdot 0.5 \cdot \left(\frac{5}{3} + 2\right) = \frac{43}{24} \approx 1.792$$

Hence the expected difference between the losing and the winning bid can be calculated as

$$E\left[B_{lose} - B_{win}\right] = \frac{43}{24} - \frac{19}{12} = \frac{5}{24} \approx 0.21$$

The expected profit of the rival firm can be calculated as

$$E\left[\pi_2\right] = \frac{1}{4}\left(\frac{5}{3} - \frac{4}{3}\right) = \frac{1}{12} \approx 0.083$$

That is, in an asymmetric situation it is no longer true that the difference between the lowest two bids is necessarily smaller than the expected profit of the *marginal* participant. The dominant firm will bid relatively low to increase its probability of winning, since winning is very profitable if costs are low. The rival firm will win infrequently but will have relatively high costs and will therefore optimally bid high. To be willing to participate, its transaction costs must be low (lower than its expected profit) and can, as the example shows, be lower than the expected differences in bids.

To conclude, for most distributions the difference between the two lowest bids is a conservative estimate of the total seller-side bidding transaction cost – if values are private and firms are ex-ante symmetric. With asymmetric bidders the conjecture underlying this study no longer holds.

Estimating bidding costs from the bid distribution

Assuming costly but unlimited entry into the procurement auction and ex-ante symmetric bidders, *n* times the fixed cost of participating in the auction should be equal to the expected profit of the winner. This, in turn, should equal the expected difference between the winner's cost and the cost of the second lowest bidder. From the above analysis, the difference between the lowest and the second lowest *bid* is a conservative estimate of the expected profit of the winning bidder.

Consequently, for procurement *i* the bid cost B_i is estimated as

$$B_i = \frac{b_{i2} - b_{i1}}{n_i} \tag{22}$$

where b_{ij} is the *j*:th lowest bid in procurement *i* and n_i is the number of bidders in the procurement.

We assume that the bidder must incur the bidding cost before it learns its own cost of production. To make empirical analysis possible the bids must be normalized. We divide each bid with the average bid for the corresponding procurement. The difference between the lowest and second-lowest bid will be used as a lower-bounds estimate for the total bidder transaction cost as a percentage of the average bid (including losing bids). This fraction can easily be recalculated so as to be a percentage of the winning bid. By dividing with *n* it is also possible to estimate (a low bound of) the per-bidder relative transaction cost.

Our method is not applicable to procurements with a single bid. We make separate estimates for procurements with different number of bidders.

As noted above, the equilibrium number of bids is larger when the bid cost is small and when the distribution from which costs are drawn is dispersed. The number of bids will grow less than proportionally to the bid spread.

THE DATA

Our dataset was designed to be a representative sample of all Swedish public procurements conducted by government authorities. We sampled 20 out of approximately 300 municipalities and counties and 20 out of close to 400 central authorities, using random sampling with sampling weights proportional to the sizes of the authorities. For each selected authority we randomly sampled 20 procurements that were held in 2007 or 2008 if the authority made more than 20 procurements during those years. If the number of procurements was 20 or less we included all procurements.

This gave us a sample of approximately 650 procurements with a response rate of approximately 97 percent. Since many procurements involve the tendering of multiple contracts, the sample includes approximately 4000 contracts and about 15000 bids. For each procurement and bid we have detailed information on the type of product, the procurement mechanism, the duration of the contract and so on.

For 570 of the contracts the selection criterion was lowest price. We identify the lowest and the second lowest bids, calculate the per-contract average bid and standard deviation. We also calculate the relative bid spread S_i as

$$S_i = \frac{b_{i2} - b_{i1}}{\bar{b}_i}$$

where $\overline{b}_i$ is the average bid for procurement *i*. To reduce the effect of coding errors we eliminated outliers, defined as contracts with a bid spread of more than 1.5. Procurements with a single bid could not be used and in addition there were a few procurements where only the winning bid was available. This gave us 459 useable procurements. In Figure 1 we plot the bid spread.

Figure 1. Relative bid spread for all price-only procurements

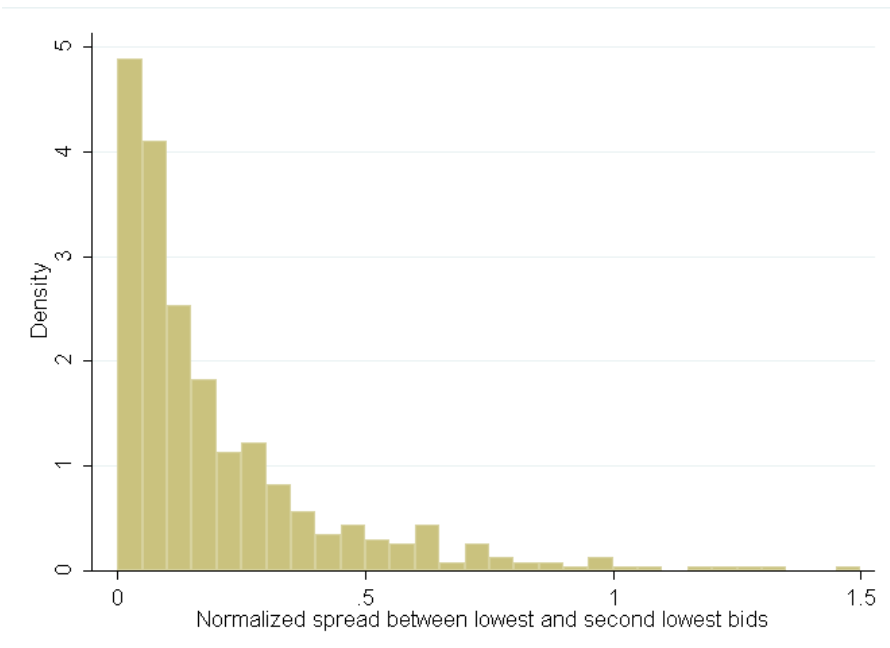

As can be seen, the histogram resembles that of an exponential distribution. Figure 2 shows the relative bid spreads broken down by the number of bids.

Figure 2 Relative bid spread and number of bidders

a) 2-5 bids (185 observations)

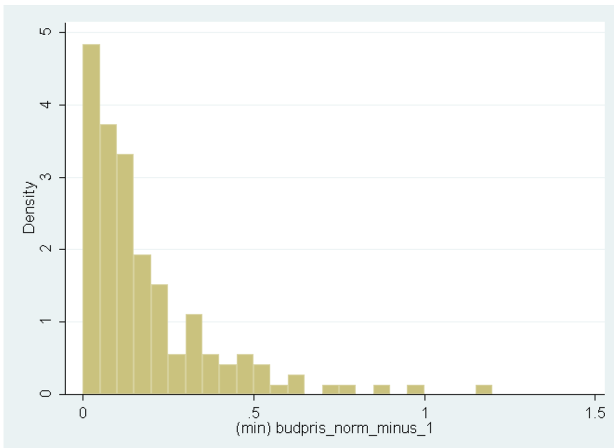

b) 6-10 bids (122 observations)

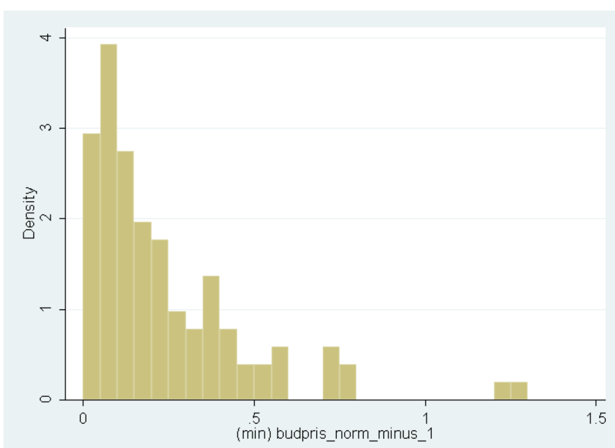

c)11-18 bids (112 observations)

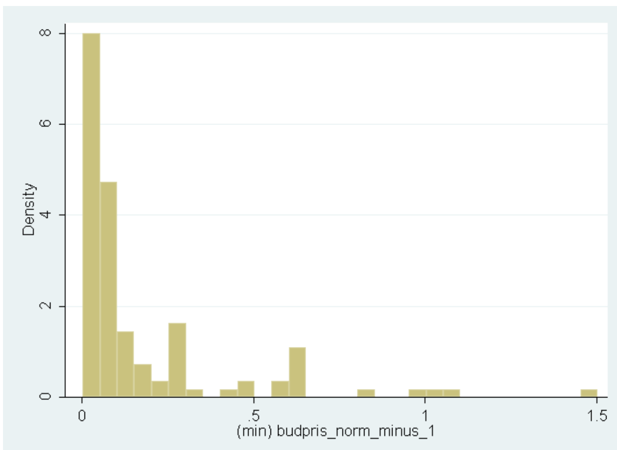

d) 19 or more bids (114 observations)

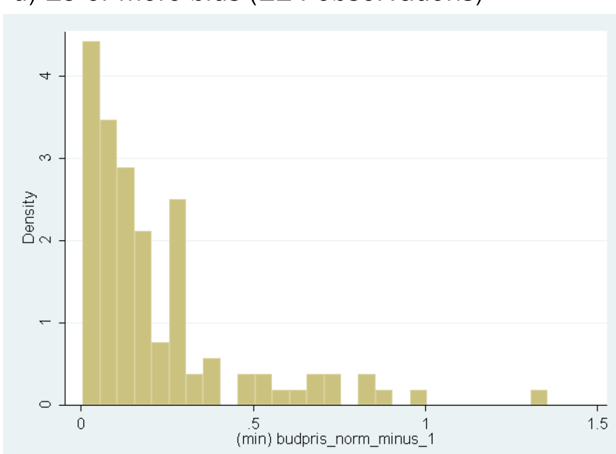

The similarity to an exponential distribution remains in all these histograms. Except for figure 2b, the histograms are close to monotonically decreasing. In all diagrams there appear to be a clustering of bids around 30 percent higher than the lowest bid.

Figure 2.e) Histogram log(bid spread)

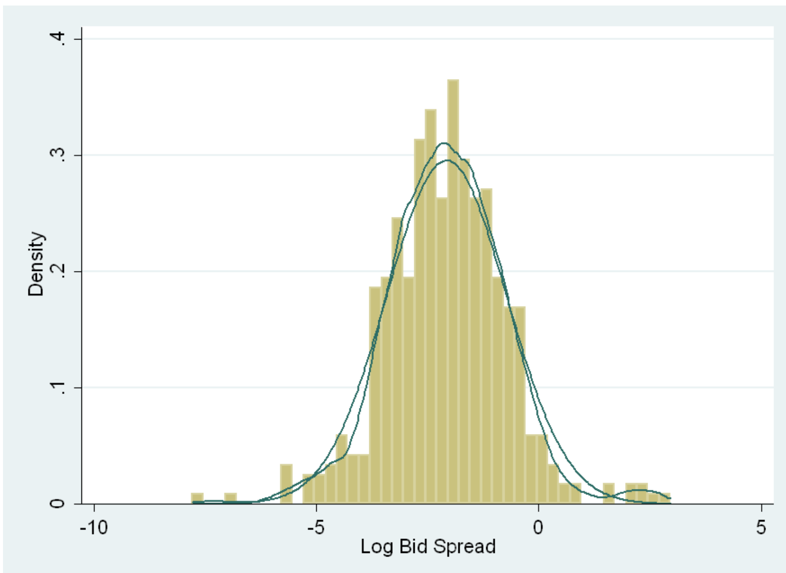

To investigate if the bid spread can be assumed to be lognormally distributed, we take the log of all bid spread. Figure 2.e shows the histogram of the log of bid spread for all contracts together with a fitted kernel density plot and a standard normal distribution curve (the smother one) fitted to the same mean.

The distribution of bid spread can be assumed to be lognormally distributed, since there are only marginal differences in the fitted and the standard normal distribution curve. The consistent difference between our fitted density plot and the normal curve is that our sample has a slightly smaller variance, σ^2.

To further explore the distribution of the bids, Figure 3 also plots the logged values of bids including the standard normal distribution. To make the distributions comparable across procurements we subtracted the per-contract average of the log of the bid. That is, the log-normalized bids are defined as

$$\beta_{ij} = \ln b_j - \bar{\beta}_i$$

$$\bar{\beta}_i = \frac{1}{n_i} \sum_j \ln b_{ij}$$

Implicitly we are assuming that the standard deviation increases in proportion to the scale, an assumption that is roughly true.

Figure 3. Normalized distribution of the log of all bids

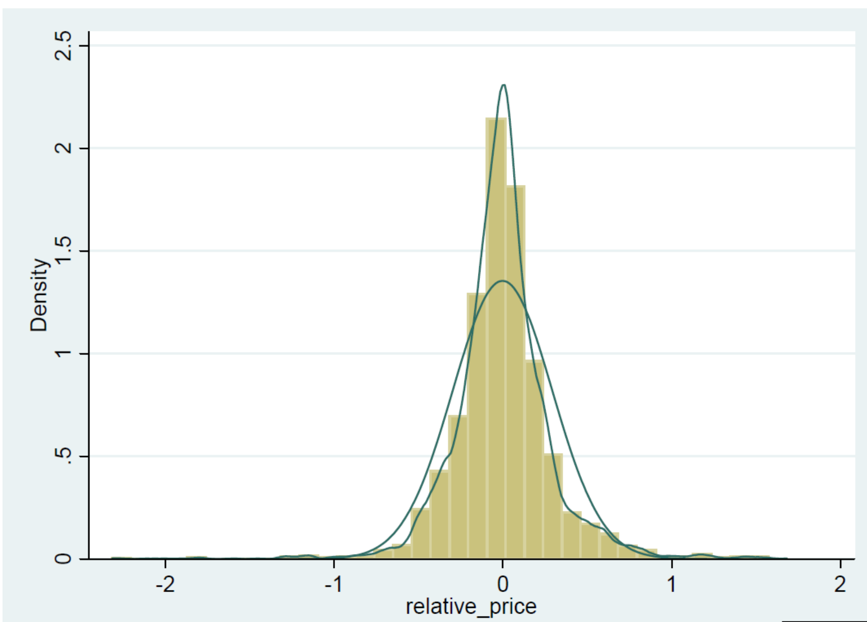

In the graph two lines show a kernel density function fitted to the empirical sample (the more peaked curve) and a normal distribution with the same mean and variance as the sample

RESULTS

For all procurements the mean and median relative bid spreads are 38 and 13 percent, respectively. Excluding relative bid spreads equal to 1 or larger gives 19 and 12 percent, respectively. We calculate the mean and median relative bid spread separately for tenders that receive two, three, four etcetera bids. The results are shown in Figure 4.

Figure 4. Mean and median relative bid spreads (share of bid, vertical axis) by number of bids (horizontal axis)

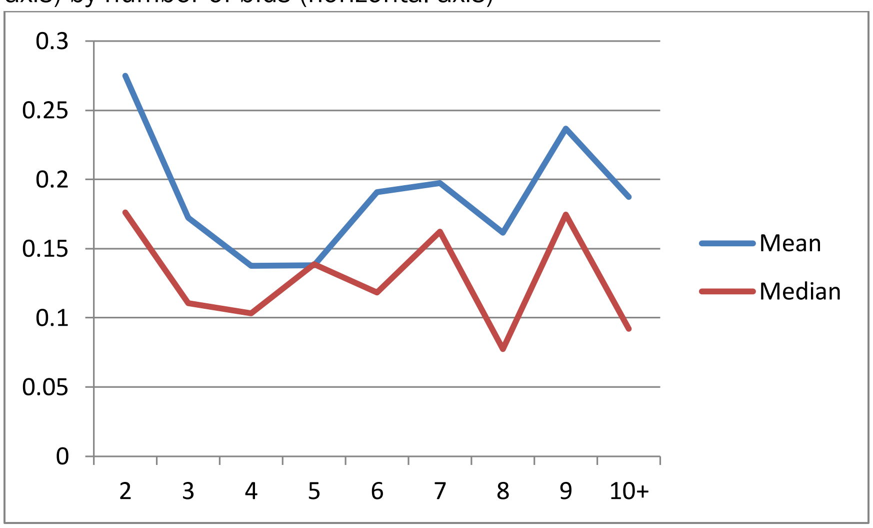

There is no obvious trend in the relation between the number of bids and the relative bid spread. A regression analysis with relative bid spread as the dependent variable and number of bids and the winning bid as explanatory variables confirms the visual impression; there is no significant relation between the variables. The median value for the relative bid spread fluctuates between just below 10 percent and a little more than 15 percent; the average of the mean over the different bid counts is 13 percent.

We have argued that a lower bound for the per-bid transaction cost can be estimated as the difference between the second-lowest and the lowest bid (the bid spread), divided by the number of bids. Since, according to Figure 4, the sum of the transaction costs for all bidders is roughly constant, irrespective of the number of bids, the estimated per-bidder transaction cost decreases with the number of bidders. (Conversely, a falling per-bidder transaction cost allows a larger number of bidders to recover their bidding costs in expectation.) Figure 5 shows the per-bidder (relative) transaction cost, estimated as mean and median, respectively, relative bid spread divided by the number of bidders.

Figure 5. Relative per-bid transaction cost (% of bid, vertical axis) and number of bidders (horizontal axis). Relative transaction cost estimated from mean and median bid spread.

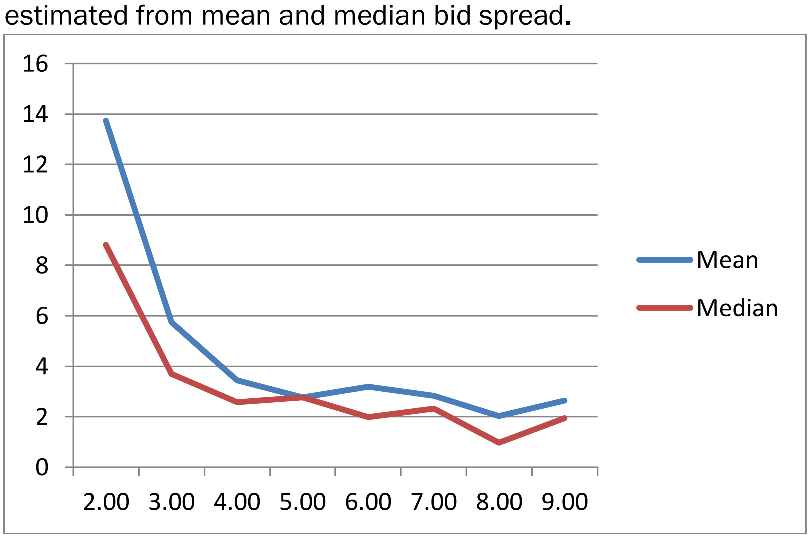

We consider median relative bid spread a conservative estimate of the sum of the bidders' transaction cost. Our estimate can be compared with the EU Commission's (2011) estimated € 28 000. As a fraction of threshold value, €193 000, this corresponds to 14.5

percent; as a fraction of the median value, about €400 000, it corresponds to 7 percent.[11] Holm (2011) estimates the transaction cost to approximately SEK 28 000 or €3 000 for procurements with a value below the threshold but above 15 % of the threshold, corresponding to between 1.5 and 9.7 percent. Assuming a median value of SEK 600 000 for procurements in this range, the relative transaction cost would correspond to 4.7 percent of the median value.

To summarize, our estimate of the transaction cost is about twice as high as those obtained in previous studies; studies that relied on a combination of surveys and valuation of time. This can be due to previous studies underestimating transaction costs. Alternatively, our method may overestimate the costs, for example because asymmetry is prevalent and because with asymmetric firms it need no longer be true that the bid difference is a lower-bounds estimate of transaction costs.

Further investigating the bid spread, we perform a regression of different explanatory variables on our dependent variable log(bid spread), see table 2 below.

Table 2. Estimations of log(bid spread)

Variable	Coefficient	Variable	Coefficient
Contract length	1.5716*** (.4614)	Multiple + frame	-3.6301** (1.6192)
Option1	-1.742** (.7829)	A-service	2.2574** (.90601)
Option2	-.1183 (.6719)	B-service	-.3063 (.9991)
Threshold value	-4.221*** (.8708)	Good+A-service	1.775** (.82869)
No. of bids	.2027*** (.0429)	Good+B-service	(omitted)
No. of winners	-.2692 (.2304)	Unclear good	2.727 (4.198)
Appealed	-1.988* (1.121)	Pro rata (variable) payment	-1.638* (.9795)
Mean bid	-1.26e-07*** (2.80e-08)	Fixed + pro rata payment	(omitted)

[11]See p. 116.

Multiple contracts	1.499* (.8799)	Constant	-4.438*** (.7572)
Frame agreement	-.4293 (.9516)		
N	322		
R^2	0.3144	Adj R^2	0.2784

* Indicates significance at the 90% level, ** Indicates significance at the 95% level,
*** Indicates significance at the 99% level.
Defaults are procurement of goods, no multiple contracts, no frame agreement and fixed payment.

As seen here there are a number of variables that influence the bid spread. The length of the contract influences bid spread positively while the length of the first contract-extension period (Option1) has the opposite effect. The number of bids significantly influences the bid spread, as does the multiple-contracts dummy and the multiple-contracts-and-framework-agreement dummy. Procurements above the threshold value have a smaller bid spread than procurements below it and the bid spread falls with the average bid, consistent with transactions costs partially being fixed and partially variable with the value of the procurement. Procurements of A services and A-services in combination with goods increases bid spread, while pro-rata (variable) payments reduce the spread.

We find some support for the notion that asymmetry increases bid spreads, since asymmetry is likely to be larger for procurements of long contracts and of services. With long contracts and with procurements of services it is more likely that there exists an incumbent with an informational advantage.

PROCUREMENT OF ROAD CONSTRUCTION

We cannot directly estimate the asymmetry between bidders in terms of winning probabilities in our dataset. The products are diverse and few bidders show up repeatedly in the sample. However, in a sample with more homogenous products we can expect the same bidders to appear repeatedly and we can empirically address the degree of asymmetry.

We have access to a dataset consisting of 1455 procurements made by the Swedish Road Administration during the 2000s. The average procurement attracted 4.3 bids, for a total of 6271 bids. Of these, we excluded a few consortia bids and a few bids with missing information, which gave us a net sample of 6236 bids made by 201 different firms.[12] Four firms – NCC, Peab, Skanska and Vägverket Produktion – accounted for two thirds of the bids and eleven other firms with at least 50 bids each accounted for a further 20 percent of the bids.

We grouped the firms into categories: firms that submitted a single bid in our sample, firms that submitted two bids, firms with three or four bids, firms with five to seven bids and so on. In each of these categories at least 60 bids were observed. Figure 6 shows the fraction of winning bids per category. As the figure shows, large firms have a slightly better chance of winning a procurement, with a winning probability of around 25 percent versus about 20 percent for small firms. That is, bidders' positions are only slightly asymmetric in the road construction industry.[13]

Figure 6. Number of submitted bids (horizontal axis) and probability of winning (%, vertical axis)

---

[12] Furthermore, about 60 procurements that received a single bid could not be used in the analysis.

[13] Removing the outlier firm Seiten shifts the trend line upwards by 1-2 percentage points, but has only a marginal impact on the slope coefficient.

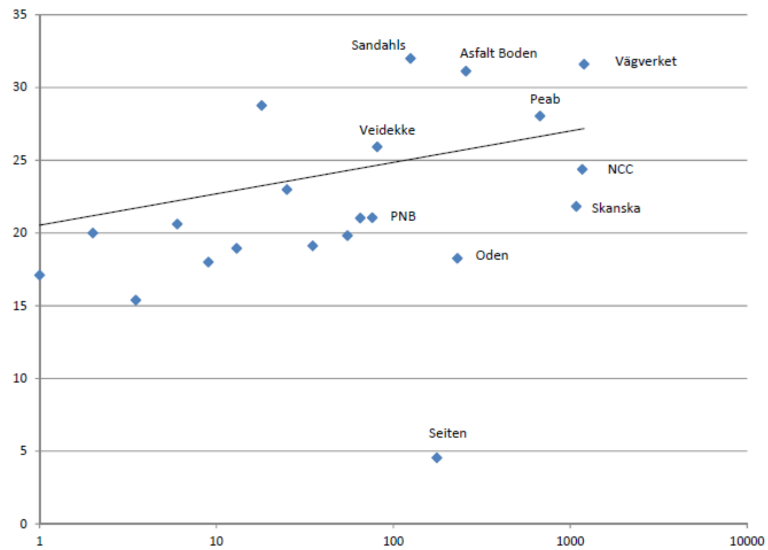

We analyzed the bids and the bid spread the same way as for the generic sample. The median spread between the lowest and second lowest bid was 5.4 percent – less than half of the bid spread in the more general sample.[14] If we ignore the modest asymmetry indicated by Figure 6 we can conclude that the total transaction cost is at least 5.4 percent of the tender value or 1.3 percent per firm. This value is close to the transaction costs estimated by the EU Commission (2011) and Holm (2011), if their estimates are related to *median* procurement values.

Since the average number of bidders is about the same as in the general sample – about four – we conclude that either the cost variation is smaller (and the transaction costs are smaller) in the construction industry or the bidding is more asymmetric in the general sample so that the large bid spread corresponds to a situation where the most successful firms face less competition. The

[14]The average bid spread was 8.7 percent. This ratio is likely to be somewhat inflated due to coding errors, although we checked and corrected the most obvious errors, i.e., those we initially coded to be larger than 1.6.

general sample includes numerous procurements of services where there is an incumbent, a fact that points in the direction of the latter explanation.

Figure 7 shows the bid spread as a function of the number of bidders. Compared to the general sample it has a much more pronounced U-shape, with minimum bid spreads in procurements that attracted 6-8 bids.

Figure 7. Mean and median relative bid spreads by number of bids, road constructions

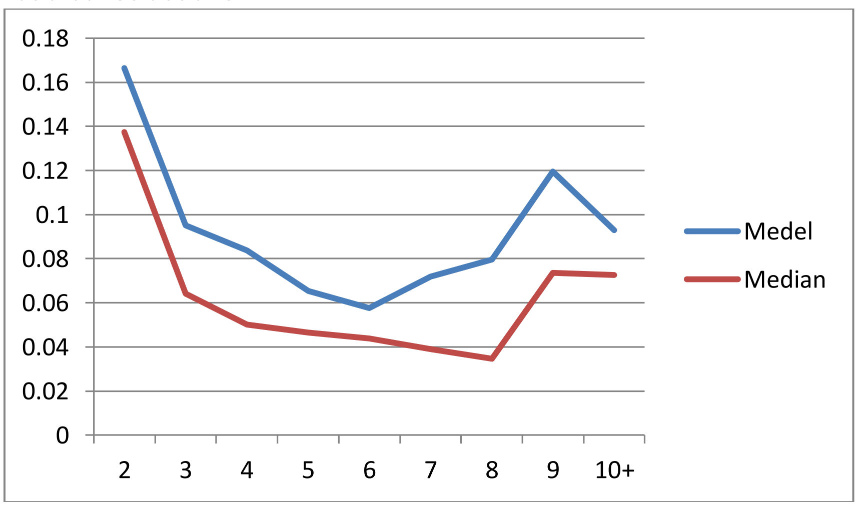

[Yet to do: Explore further how factors related to bidder asymmetry, such as Firm size and Incumbency (as proxied by duration of contract) impacts on bid spread.]

CONCLUSIONS

We show that as long as one is willing to entertain the assumptions of costly but unlimited entry and ex-ante symmetric firms, a lower bound for the bidder-side transaction costs can be estimated from the spread between the second lowest and the lowest bid. Using this method we find that average total bidder transaction costs are (more than) about 20 percent in a general sample of procurements. In the same sample we find a median value of about 12. That is, the (lower-bounds) estimate of bidder-side transaction costs is strongly skewed to the right:

The estimated transaction cost is likely to be biased upward when bidders are asymmetric. For this reason we have also studied what is likely to be a more symmetric (in terms of win probability) sample of road construction procurements. Here, the estimated average is about 12 percent and the median value is about 5 percent. The latter value is close to estimates obtained by a more traditional method, i.e., time studies.

REFERENCES

Ahsanullah, Mohammad and Valery B. Nevzorov, 2005, *Order statistics: examples and exercises*, Nova Science Publishers, New York.

Coase, Ronald, 1937, The Nature of the Firm, *Economica* 4 (16): 386–405

EU Commission, 2011, Evaluation Report, Impact and Effectiveness of EU Public Procurement Legislation, Commission staff working paper (draft version).

Harter, H. Leon, 1961, Expected Values of Normal Order Statistics, *Biometrica*, 48, 151-165.

Holm, Jerker, 2011, *En analys av transaktionskostnaderna vid offentlig upphandling*, Report to the Swedish Public Procurement Committee.[In Swedish.]

Klemperer, Paul, 2004, *Auctions: Theory and Practice*, Princeton University Press, Princeton.

Shelanski, Howard and Peter Klein, 1995, Empirical Research in Transaction Cost Economics: A Review and Assessment, *Journal of Law, Economics & Organization*, 11, 335-361.