Identification in Auction Models with Interdependent Costs

Paulo Somaini

Stanford Graduate School of Business and National Bureau of Economic Research

This paper provides a nonparametric identification result for procurement models with asymmetric bidders, dependent private information, and interdependent costs. For risk-neutral bidders, the model's payoff-relevant primitives are the joint distribution of private information and each bidder's full-information expected cost. The joint distribution of bids identifies the joint distribution of signals. First-order conditions identify the expected cost conditional on tying with at least one competitor for the lowest bid. I show identification of each bidder's full-information cost, using variation in competitors' cost shifters that are excludable from bidders' own full-information costs, and generate variation in the set of competitors' signals that induce a tie for the lowest bid. I estimate the relevant payoff primitives using data from Michigan highway procurements and evaluate policies that affect the winner's curse's severity.

I. Introduction

In procurement auctions with interdependent costs the information about each bidder's cost is scattered among all bidders. Bidders should realize that the result of the auction is informative about competitors' information and that they win the projects deemed too costly or undesirable by other participants. This adverse-selection phenomenon is usually referred to as the winner's curse in the auctions literature. The effect of procurement

I am grateful to Frank Wolak, Liran Einav, Han Hong, Nikhil Agarwal, Tim Armstrong, Lanier Benkard, Arthur van Benthem, Steven Berry, Timothy Bresnahan, George Bulman, Dominic Coey, Ignacio Esponda, Amit Gandhi, Matt Gentry, Phil Haile, Tong Li, Alejandro

Electronically published September 10, 2020
[*Journal of Political Economy*, 2020, vol. 128, no. 10]

policies on equilibrium behavior and outcomes depends crucially on the nature and extent of the winner's curse and on how bidders adjust their bids in response to it.

While the interdependent-costs model has been analyzed extensively in the theoretical literature, the empirical literature has mainly analyzed two polar cases. At one extreme is the private-costs model, where each bidder knows his own costs and there is no room for the winner's curse. At the other extreme is the pure common-cost model, where all bidders would incur exactly the same cost if they win but are uncertain about it. Intermediate cases have received little empirical attention, mainly because of the lack of identification results.[1] In fact, Laffont and Vuong (1996) provide a negative result. They show that any joint distribution of bids that is rationalizable by an interdependent-costs model is also rationalizable by some private-costs model.

One of the main insights of the theoretical analysis of models with interdependent costs is that changes in competitive conditions have different effects on equilibrium bid strategies, depending on the true information structure. In models with strong interdependence, the effect of the winner's curse is important, and bidders may bid less aggressively in more competitive environments. If interdependence is weak or nonexistent, so is the effect of the winner's curse, and bidders may bid more aggressively in more competitive environments. Haile, Hong, and Shum (2004) use these insights to develop a test for the null hypothesis of private costs against an alternative of interdependence exploiting variation in the number of participants. I provide a positive identification result, using richer variation in competitive conditions generated by continuous cost shifters, for example, distance from each bidder's plant to the project location.

The objects to be identified are the primitives of the Bayesian game that represents the procurement. When bidders are risk neutral and the rules of the auction are known, the game is fully defined by (1) the distribution of bidders' private information and (2) each bidder's full-information cost:

Molnar, Peter Reiss, Johannes Stroebel, Ali Yurukoglu, and Daniel Waldinger. I have also benefited from seminar participants at Stanford, Berkeley, the University of California Los Angeles, Northwestern, the University of Chicago, the University of Wisconsin, New York University, the University of Pennsylvania, Harvard, Massachusetts Institute of Technology, Yale, and Vanderbilt. The Martin Lee Johnson fellowship is gratefully acknowledged. Data are provided as supplementary material online.

[1] The theoretical literature includes Maskin and Riley (2000b), Athey (2001), McAdams (2003, 2007), and Reny and Zamir (2004). Laffont, Ossard, and Vuong (1995), Athey, Levin, and Seira (2011), Krasnokutskaya (2011), and Roberts and Sweeting (2013) are examples within the private-values/cost paradigm, while Hendricks and Porter (1988), Athey and Levin (2001), Bajari and Hortaçsu (2003), and Hendricks, Pinkse, and Porter (2003) are examples within the common-values/cost paradigm. Paarsch (1992) studied the two polar cases to decide between them. Hong and Shum (2002) is one of the few empirical papers that estimate a model with interdependent costs.

a function that returns the expected completion cost conditional on own and competitors' private information. While these payoff-relevant primitives are coarser than the joint distribution of costs and private information—regarded as the true primitive by Athey and Haile (2002, 2007)—they are sufficient to analyze the effects of many relevant policy changes (e.g., rules of the auction, reserve prices, subsidies) on outcomes such as bidding behavior, project allocation, and prices.

I restrict the technology and information of the general interdependent-cost model in three important ways. First, each bidder's private information can be summarized by a one-dimensional signal. Second, the joint distribution of bidders' signals is independent of the cost shifters. Third, each bidder's cost shifter affects his own full-information cost but not his competitors'.

I also assume that, for each configuration of cost shifters, the observed data are generated by the repeated play of the same equilibrium where bidders use monotone pure strategies. Reny and Zamir (2004) showed that this type of equilibrium exists in the general interdependent-cost model with affiliated signals. I extend the results in Maskin and Riley (2000b) and show that when the reserve price does not bind, bidders employ strictly monotone strategies in every monotone-pure-strategy equilibrium. I also extend the constructive proof of Reny and Zamir (2004) and show that cost shifters generate variation in equilibrium bid functions. Adequately choosing a vector of cost shifters allows the locus of signals where all bidders tie to be moved around so that it includes any desired configuration of signals.

The main results of the paper show identification of the primitives of the model. Because equilibrium strategies are strictly monotone, the copula of signals is identified from the copula of bids. The marginal distribution of signals can be normalized to be uniform without loss of generality. Identification of the full-information costs involves several steps. The first step exploits the bidder's first-order condition, which equates the marginal cost and the marginal revenue associated with winning with marginally higher probability. In first-price procurements, the marginal revenue is identified through standard arguments (Guerre, Perrigne, and Vuong 2000; Campo, Perrigne, and Vuong 2003). The marginal cost is the expected cost conditional on the event that the observed bid is pivotal, that is, the set of competitors' signals such that the bidder ties with at least one of them for the lowest bid. This set is identified from the marginal distributions of bids. Moreover, cost shifters are able to generate rich variation in this set. The main result of the paper shows that if cost shifters induce sufficient variation, the full-information cost is identified for all configurations of own and competitors' signals. While the main focus of this paper is on first-price auctions, the identification result can be extended to other sealed-bid formats, such as second-price and all-pay auctions.

Nonparametric estimators of the joint distribution of bids and full-information costs may suffer from the curse of dimensionality. In most empirical contexts, it may be useful to rely on some parametric and distributional assumptions. I employ a parametric model with a Gaussian information structure that is tractable and general enough to nest many important cases in the private- and interdependent-values paradigms. In an application to first-price highway procurements in Michigan, I use bidder's distance to the project as the bidder-specific cost shifter and estimate the model parameters using a multistep procedure that mirrors the identification argument.

The full-information cost estimates point to a structure that lies between the two polar paradigms of private and pure common costs. The effect of competitors' signals is statistically significant but of smaller magnitude than the effect of own signals. To illustrate the economic importance of the estimated degree of interdependence and the implied winner's curse, I perform a series of counterfactuals where the state Department of Transportation restricts participation and sets different reserve prices. Restricting participation reduces the severity of the winner's curse, and, under some conditions, it may even reduce procurement costs (Matthews 1984; Hong and Shum 2002). The estimates suggest that the effect of the winner's curse is strong enough that bidders bid more aggressively when participation is restricted, but not enough to generate cost savings.

This paper contributes to the literature on identification of auction models. The initial contributions focused on the private-values model with independent (Guerre, Perrigne, and Vuong 2000) or affiliated private information (Li, Perrigne, and Vuong 2002; Campo, Perrigne, and Vuong 2003). Laffont and Vuong (1996) showed that the affiliated- or interdependent-value model is not identified from the joint distribution of bids. Li, Perrigne, and Vuong (2000) studied the conditionally independent private-information model and showed identification of the two polar cases of private and pure common values. Hendricks and Porter (1988), Hendricks, Porter, and Wilson (1994), and Hendricks, Pinkse, and Porter (2003) used observable ex post values to test equilibrium bidding models in common-value environments. Hendricks, Pinkse, and Porter (2003) also proposed using observed ex post values to identify a symmetric pure common-value model under the assumption that bidders' posterior estimates of the common value are unbiased. They also noticed that in symmetric first-price auctions, the equilibrium distribution of bids around the reserve price can be used to distinguish between private and common values. Hill and Shneyerov (2013) developed a formal test based on this observation. Février (2008) considered identification of a particular class of pure common-values model. More recently, the literature has focused on cases where the researcher has access to the distribution of bids conditional on other covariates that introduce some exogenous variation in competitive or informational conditions. This additional source of variation can be used to test the null hypothesis of private values (Haile, Hong, and Shum 2004; Hortaçsu and Kastl 2012) and the null of pure common values (Hendricks, Pinkse, and Porter 2003, n. 2). It can also be used to identify attitudes toward risk (Guerre, Perrigne, and Vuong 2009; Campo et al. 2011), correlated private values in ascending auctions (Aradillas-López, Gandhi, and Quint 2013), and a selective entry process (Gentry and Li 2014). In this paper, I show that variation in cost shifters can be used to identify all the payoff-relevant characteristics of the interdependent-values model with risk-neutral bidders.

The literature on single-unit empirical auctions has made significant progress in analyzing various settings, including highway construction, timber sales, oil leases, cars, wines, stamps, and coins, among others. Auction models have also been useful in recent studies of defense contracts and consumer loans.[2] The results have yielded insights into the internal working of cartels, the competitive effects of entry, and the choice of optimal reserve prices. In modeling these markets, researchers often had to decide between the two extreme paradigms of bidder behavior, thus facing a trade-off between the economic phenomena that they could analyze. Under private values, there is no room for the winner's curse. In the common-value model, there is no room for improving the efficiency of the allocation, as all bidders have the same ex post value. The results of this paper expand the set of available tools and relax the constraints on modeling decisions so that the model can accommodate economic phenomena deemed relevant for the empirical setting and research question at hand.

The rest of the paper is organized as follows. Section II describes the general interdependent-cost framework. Section III shows the main identification results. Section IV presents an application using highway procurement data from Michigan, and section V presents the empirical results. Section VI concludes.

II. The Interdependent-Cost Model

A. *Primitives*

An auctioneer procures the completion of a project and runs a first-price sealed-bid auction between n risk-neutral bidders. There is no reserve

[2] Highway construction: Porter and Zona (1993), Krasnokutskaya (2011), Krasnokutskaya and Seim (2011), Balat (2017), and Bolotnyy and Vasserman (2019). Timber sales: Baldwin, Marshall, and Richard (1997), Athey and Levin (2001), Haile (2001), Athey, Levin, and Seira (2011), Aradillas-López, Gandhi, and Quint (2013), and Roberts and Sweeting (2013). Oil leases: Hendricks and Porter (1988, 1993), Hendricks, Porter, and Wilson (1994), Hendricks, Pinkse, and Porter (2003), and Compiani, Haile, and Sant'Anna (2020). Cars: Roberts (2013) and Larsen (2020). Wines: Ashenfelter (1989). Stamps: Asker (2010). Coins: Bajari and Hortaçsu (2003). Defense contracts: Bhattacharya (2018). Consumer loans: Cuesta and Sepúlveda (2019).

price. The cost to bidder i is a random variable C_i, and his information is summarized by a signal S_i. I adhere here to the convention of using capital and lowercase letters to denote random variables and their realized values, respectively. At the time of the auction, i knows s_i, the realization of his own signal, but is uncertain about the realization of the vector of competitors' signals $S_{-i} = [S_j]_{j \neq i}$ and own future project completion cost C_i. In other words, each bidder knows his own information but does not know his competitors'; moreover, his information allows him to make only an imperfect forecast of his own costs. Denote the full random vector of signals by $S = [S_i]_{i=1}^n$ and the vector of costs by $C = [C_i]_{i=1}^n$.

All bidders have access to the following public information: a set of bidder-specific cost shifters $x = [x_1, x_2, \dots, x_n]$ and a set of observable auction characteristics x_0. In principle, cost shifters and auction characteristics can be multidimensional. For example, the cost shifter of bidder i may include his distance to the project site and publicly observable measures of his backlog, utilization, and capacity. Auction characteristics x_0 may include publicly available estimates of the cost and duration of the project. To keep notation simple, however, I consider a case in which each bidder has a single-dimensional cost shifter, that is, $x_i \in \mathbb{R}$ and $x \in \mathbb{R}^n$. Observed heterogeneity x_0 will be conditioned upon and is hereafter omitted from notation.

The primitives of the model are the joint distribution of costs and signals conditional on public information. Its cumulative distribution function (CDF) is denoted by $F_{C,S|x}$, or F for short. The CDF of i's completion cost, given his information at the time of the auction, is $F_{C_i|s_i,x}$, and its expectation is $E(C_i|s_i, x)$. If i learns all competitors' signals, the CDF of his costs becomes $F_{C_i|s_{-i},s_i,x}$, with expectation $E(C_i|s_{-i}, s_i, x)$. This expectation is bidder i's full-information expected completion cost, or full-information cost, for short. All these distributions and expectations of costs are uniquely determined by the primitive F. Because signals have only ordinal meaning, their marginal distribution can be normalized to be uniform without loss of generality. The joint distribution of signals F_S is thus a copula.

The identification results in this paper require the following assumptions on bidders' technology and information.

Assumption A1. Cost shifters and signals are independent: $F_{S|x} = F_S$.

Assumption A2. Signals are one-dimensional random variables. The joint distribution F_S admits a continuously differentiable and bounded density function f.

Assumption A3. The full-information cost of bidder i is $c_i(s_{-i}, s_i, x_i)$, which does not depend on x_j for any $j \neq i$; that is,

$$c_i(s_{-i}, s_i, x_i) := E(C_i|s_{-i}, s_i, x_i) = E(C_i|s_{-i}, s_i, x). \tag{1}$$

For every x_i, $c_i(\cdot, \cdot, x_i)$ is bounded and continuously differentiable.

Assumptions A1 and A3 define the conditions that an observable auction characteristic has to satisfy to be considered a cost shifter instead of one of the characteristics in x_0. Assumption A1 states that the joint distribution of signals does not depend on cost shifters. Assumption A3 requires that bidders' cost shifters affect only their own full-information cost but not competitors'. For example, if x_i is distance to the project, these assumptions rule out that bidder i has systematically lower or higher costs when j is close to the project. They also rule out the case where i regards the signal of bidder j as more informative when j is close to the project.

Assumption A2 states that bidders summarize all private information in a single-dimensional variable. It rules out the possibility that bidder i receives two signals: one that is informative about his own costs (e.g., own equipment availability) and one that is informative about competitors' information and costs (e.g., conditions in the equipment rental market). The requirement that the density and full-information costs are bounded and continuously differentiable ensures that it is valid to differentiate under the integral sign when evaluating $(\partial/\partial s_i)E(C_i|S_{-i} \geq s_{-i}, s_i, x)$.

Most models considered in the single-unit auction literature are special cases of the interdependent-cost model. Write $F = F_{C|S,x}F_{S|x}$. Typical models impose conditions on either $F_{C|S,x}$ or $F_{S|x}$. Conditions on $F_{C|S,x}$ determine whether the model is in the private-cost paradigm. Conditions on $F_{S|x}$ determine how bidders' private information is distributed.

In private-costs models, each bidder knows his own expected completion cost, so i's full-information cost does not depend on competitors' signals; that is,

$$c_i(s_{-i}, s_i, x_i) = E(C_i|s_i, x_i) \tag{2}$$

for all s_{-i}, s_i, and x_i. The standard independent-private-values (costs) model assumes that $F_{S|x}$ is the product of its marginal distributions (Milgrom and Weber 1982, sec. 2.1), while the affiliated-private-values model allows signals to be affiliated (Li, Perrigne, and Vuong 2002). In pure common-costs models the cost of completing the project is common to all bidders. They all share exactly the same full-information cost (Rothkopf 1969; Wilson 1977; Hendricks and Porter 1988).

The interdependent-cost model is more general than the pure common-costs and the private-cost models (see Reny and Zamir 2004; McAdams 2007; Bergemann, Brooks, and Morris 2017). First, it allows for arbitrary interdependence. Each bidder may have a different full-information cost function, which may arbitrarily depend on competitors' signals. Second, private information need not be independent; that is, $F_{S|x}$ is not necessarily the product of the marginals. Third, bidders can be asymmetric. The function $F_{S|x}$ is not required to be exchangeable in its arguments, and $c_i(\cdot, \cdot, \cdot)$ may be different from $c_j(\cdot, \cdot, \cdot)$. Fourth, it allows for asymmetries introduced by observable cost shifters, which will be critical for the identification argument. Appendix A (apps. A–D are available online) presents examples of information structures within the interdependent-cost model.

B. *The First-Price Auction*

First-price sealed-bid auctions are modeled as Bayesian games. In the game theory jargon bidders are players, signals are types, and bids are actions. The payoff functions and joint distribution of signals are common knowledge. The payoff function of a risk-neutral bidder i is

$$u_i(b, s, x) = (b_i - c_i(s_{-i}, s_i, x_i))1(b_i \prec b_{-i}), \tag{3}$$

where b_i is the bid submitted by bidder i, $b = \{b_j\}_{i=j}^{n}$, and $b_i \prec b_{-i}$ denotes that b_i beats all other bids b_{-i}. I follow the literature (Reny and Zamir 2004; McAdams 2007) in assuming no entry costs and ruling out ties with a deterministic tie-breaking rule that awards the project to the bidder with the lower index. Therefore, $b_i \prec b_{-i}$ is equivalent to $b_i < b_j$ for $j < i$ and $b_i \leq b_j$ for $j > i$. The payoff of bidder i depends on competitors' bids (or actions) b_{-i} because they determine whether i is awarded the project. It also depends on own and competitors' signals through the full-information cost.

The Bayesian game is defined by the n payoff functions and the joint distribution of signals. Thus, $\langle\{u_j\}_{j=1}^{n}, F_S, x\rangle$ defines a Bayesian game. Two primitive models with different $F_{C|S,x}$ but identical full-information cost functions $\{c_j\}_{j=1}^{n}$ give rise to the same Bayesian game. In other words, all other features of $F_{C|S,x}$ not accounted for by the full-information cost are not payoff relevant.

A pure-strategy Bayes-Nash equilibrium of the game $\langle\{u_j\}_{j=1}^{n}, F_S, x\rangle$ is a profile of decision rules, or bidding functions, $\beta = \{\beta_j\}_{j=1}^{n}$ such that for every $i = 1, \dots, n$ and every s_i,

$$\beta_i(s_i) \in \operatorname*{argmax}_{b_i} E((b_i - c_i(S_{-i}, s_i, x_i))1(b_i \prec \beta_{-i}(S_{-i}))|s_i).$$

Following the standard definition of Bayes-Nash equilibrium (Mas-Colell, Whinston, and Green 1995), each bidder must be playing a best response to the conditional distribution of his competitors' strategies for each signal that he might end up receiving.

There is a large theoretical literature addressing the questions of existence of equilibrium in first-price auctions with interdependent costs (Maskin and Riley 2000b; Athey 2001; McAdams 2003, 2007; Reny and Zamir 2004). To use and expand these results, I introduce the following additional assumptions.

Assumption A4. Bidders' signals are affiliated, and $f(s)$ is strictly positive on $[0, 1]^n$.

Assumption A5. For every $x_i \in \mathbb{R}$, the full-information cost $c_i(\cdot, \cdot, x_i)$ is strictly increasing in own signal s_i. Moreover, either it satisfies the private-costs condition (eq. [2]) or $(\partial/\partial s_j)c_i(s_{-i}, s_i, x_i)$ is strictly positive for all $s \in [0, 1]^n$, $x \in \mathbb{R}^n$, $j \neq i$, and i.

Assumption A4 is assumption A.2 in Reny and Zamir (2004).[3] The requirement of a bounded and strictly positive density rules out distributions where the support of S_{-i} depends on S_i, for example, if S_1 and S_2 are perfectly correlated.

Assumption A5 was introduced by Maskin and Riley (2000b, see their eqq. [1] and [7]), and it is slightly stronger than the corresponding assumption in Reny and Zamir (2004, see assumption A.1(iii)). It rules out cases where some bidders' costs are private while others are interdependent. It also rules out cases where bidders' costs are private over some range of signals and interdependent over others.

Reny and Zamir (2004) established the existence of equilibria in monotone pure strategies for first-price auctions with asymmetric bidders, interdependent values, and affiliated one-dimensional signals. McAdams (2007) showed that every mixed-strategy equilibrium is outcome equivalent to a monotone-pure-strategy equilibrium, justifying the practice of restricting attention to monotone pure strategies.

Maskin and Riley (2000b) showed that in any monotone-pure-strategy equilibrium, the support of the winning bid W is an interval $[b_*, b^*]$ and that its CDF H_W is continuous on $[b_*, b^*)$.[4]

The equilibrium strategies considered by Reny and Zamir (2004) satisfy one additional property: bidders never bid below their expected cost conditional on winning, even if the probability of winning is zero. Proposition 1 below uses this property and the fact that there are no reserve prices to extend the results in Maskin and Riley (2000b).

Proposition 1. Assume A2–A5. For every $x \in \mathbb{R}^n$, the first-price auction game has an equilibrium in monotone pure strategies β where bidders never bid below their expected cost conditional on winning. For every such equilibrium of the game, there exists an interval $[b_*, b^*]$ such that (i) every bid $b \in [b_*, b^*]$ belongs to the support of bids of at least two bidders, (ii) $H_W(b_*) = 0$, (iii) H_W is continuous and strictly increasing on $[b_*, b^*]$, with

[3] They also assume that f is measurable on $[0, 1]^n$, which is implied by continuous differentiability in assumption A2.

[4] See eq. (7) in Maskin and Riley (2000b). They studied high-bid auctions and found that the distribution of winning bids may have an atom at the lower end point. Their result in procurement, or low-bid, auctions implies that there may be an atom at the higher end point b^*, which results in an H_W continuous on $[b_*, b^*)$ and $H_W(b_*) = 0$. Proposition 1 shows that these atoms cannot occur in the equilibrium strategies considered by Reny and Zamir (2004) when reserve prices do not bind.

$H_W(b^*) = 1$, and (iv) there are positive scalars κ_i for $i \in \{1, 2, \ldots, n\}$ such that $\beta_i(s_i') - \beta_i(s_i) \geq (s_i' - s_i)\kappa_i$ for all $0 \leq s_i < s_i' \leq 1$ and $\beta_i(s_i') \leq b^*$.

COROLLARY. Fix any $x \in \mathbb{R}^n$ and equilibrium monotone pure strategies β such that bidders never bid below their expected cost conditional on winning. Bidder i's inverse equilibrium bid function $\beta_i^{-1} : [b_*, b^*] \to [0, 1]$, defined as $\beta_i^{-1}(b) = \inf\{s_i \in [0, 1] : b < \beta_i(s_i)\}$, is a weakly increasing Lipschitz continuous function with Lipschitz constant κ_i^{-1}, $\beta_i^{-1}(b_*) = 0$, and $\beta_i^{-1}(b^*) \leq 1$. There is at least one bidder j with $\beta_j^{-1}(b^*) = 1$. If $n = 2$, then $\beta_i^{-1}(\cdot)$ is also strictly increasing on $[b_*, b^*]$.

The proof is in appendix B.1, along with an additional corollary summarizing the properties of equilibrium bid functions β. Proposition 2.5 in the online supplementary material shows that inverse bid functions are strictly increasing when $n > 2$ for a large family of primitives.

As remarked by Maskin and Riley (2000a), if $\phi_i := \beta_i^{-1}(b^*) < 1$, then bidder i has to be indifferent between participating and staying out at the signal ϕ_i. This means that for $b_i = \beta_i(\phi_i) \leq b^*$ and $s_{-i}(b_i) := \{\beta_j^{-1}(b_i)\}_{j \neq i}$,

$$E(C_i | s_i = \phi_i, S_{-i} \geq s_{-i}(b_i), x_i) = b_i.$$

Bids for each signal $s_i < \phi_i$ have to satisfy some optimality conditions. The expected profit of bidder i as a function of his bid b can be written as

$$\int_{\{\tau : \tau \geq s_{-i}(b)\}} (b - c_i(\tau, s_i, x_i)) f(\tau | s_i) d\tau. \tag{4}$$

It will be useful to consider the case where competitors' inverse bid functions are differentiable.[5] Differentiating with respect to b yields the following first-order condition:

$$\begin{aligned} &\Pr(S_{-i} \geq s_{-i}(b) | s_i) \\ &\quad - \sum_{j \neq i} \frac{\partial \beta_j^{-1}(b)}{\partial b} f(s_j | s_i) \int_{\{\tau : \tau \geq s_{-ij}(b)\}} \big(b - c_i([\tau, s_j], s_i, x_i)\big) f(\tau | s_j, s_i)\, d\tau \\ &\quad = 0, \end{aligned}$$

where $s_j = \beta_j^{-1}(b)$. Rearranging the expression yields

$$b - \frac{\Pr(S_{-i} \geq s_{-i}(b) | s_i)}{|(\partial/\partial m) \Pr(S_{-i} \geq s_{-i}(m) | s_i)|_{m=b}|} = \sum_{j \neq i} \pi_j^{(i)} E\big(C_i | s_i, s_j, S_{-ij} \geq s_{-ij}(b), x_i\big), \tag{5}$$

[5] Because inverse bid functions are Lipschitz continuous, they are differentiable almost everywhere. The formal results of this paper rely on the envelope theorem, which holds even if competitors' inverse bid functions are not differentiable and specializes to eq. (5) when they are.

where

$$\left|\frac{\partial}{\partial m}\Pr(S_{-i} \geq s_{-i}(m)|s_i)|_{m=b}\right| = \sum_{j\neq i}\frac{\partial\beta_j^{-1}(b)}{\partial b}f\left(s_j|s_i\right)\Pr\left(S_{-ij} \geq s_{-ij}(b)|s_i, s_j\right)$$

and

$$\pi_j^{(i)} = \frac{\left(\partial\beta_j^{-1}(b)/\partial b\right)f\left(s_j|s_i\right)\Pr\left(S_{-ij} \geq s_{-ij}(b)|s_i, s_j\right)}{\sum_{k\neq i}\left(\partial\beta_k^{-1}(b)/\partial b\right)f(s_k|s_i)\Pr(S_{-ik} \geq s_{-ik}(b)|s_i, s_k)}.$$

The optimality condition (5) states that the marginal expected revenue of increasing the probability of winning should equal the marginal cost. The left-hand side shows that bidder i's marginal revenue is b minus a markup term. The right-hand side is the expected cost conditional on the event where b is a pivotal bid—it ties with at least one competitor for the lowest bid. The event "b is a pivotal bid" or "ties with a competitor for the lowest bid" is the union of $n - 1$ events where i ties with bidder j and outbids all other bidders. As a result, the expected costs conditional on submitting a pivotal bid are a weighted average of the expected cost conditional on each of these events. The weights are given by the probability of tying with each competitor conditional on tying with at least one of them. These conditional probabilities depend on the competitors' bid functions and on the distribution of signals.

Cost shifters play a critical role. Under assumption A3, competitors' cost shifters are excluded from bidder i's full-information costs and from the right-hand side of equation (5). By inducing variation in equilibrium bid functions, they generate variation in the events "i ties with bidder j and outbids all the other bidders" and in the weights $\pi_j^{(i)}$ for $j \neq i$. Proposition 2 uses the proof techniques developed by Athey (2001) and Reny and Zamir (2004) to show that for each s_{-i}, it is possible to select a vector of cost shifters such that $s_{-i}(b) = s_{-i}$, keeping bidder i's own signal s_i and cost shifter x_i fixed. Let

$$\begin{aligned}\bar{c}_i(x_i) &:= c_i([1, \dots, 1], 1, x_i),\\ \underline{c}_i(x_i) &:= E(c_i(S_{-i}, S_i, x_i)|S_i = 0),\\ h_i &:= \sup_{x_i\in\mathbb{R}}\bar{c}_i(x_i) - \underline{c}_i(x_i).\end{aligned}$$

The first two objects are finite, by assumption A5. The following assumption ensures that h_i is also finite.

Assumption A6. (i) There exists a uniform bound $\sup_{x_i\in\mathbb{R}} c_i([1, \dots, 1], 1, x_i) - c_i([0, \dots, 0], 0, x_i) < \infty$. (ii) For every s, the function $c_i(s_{-i}, s_i, \cdot)$ is a strictly increasing bijection mapping $\mathbb{R}$ to itself.

The second part of assumption A6 implies that cost shifters introduce continuous variation in full-information costs and that full-information costs can be made arbitrarily low or high by appropriately choosing the level of cost shifter $x_i \in \mathbb{R}$. This condition is trivially satisfied when full-information costs are additively separable in the cost shifters.

PROPOSITION 2. Assume A1–A6. For all $x_i \in \mathbb{R}$ and $s \in (0,1]^n$, there are cost shifters $x_{-i} = [x_2, \ldots, x_n] \in \mathbb{R}^{n-1}$, a bid b, and equilibrium strategies $\{\beta_j\}_{j=1}^n$ of the game $\langle \{u_j\}_{j=1}^n, F_S, (x_i, x_{-i}) \rangle$ such that for all j, $\beta_j^{-1}(b) = s_j$. Moreover, $x_{-i} \in X_{-i}(x_i, s)$, where

$$X_{-i}(x_i, s) = \left\{ \begin{array}{c} x_{-i} \in \mathbb{R}^{n-1} : \forall\, j \neq i, \\ \underline{c}_j(x_j) \ \le\ \bar{c}_i(x_i) + (1 - \Pr(S_{-i} \ge s_{-i}|s_i))^{-1}\max_{k\neq i} h_k \\ \bar{c}_j(x_j) \ge \underline{c}_i(x_i) - \left(1 - \Pr\left(S_{-j} \ge s_{-j}|s_j\right)\right)^{-1}\max_{k\neq j} h_k \end{array} \right\},$$

and $b \in [\underline{c}_i(x_i), \bar{c}_i(x_i) + (1 - \Pr(S_{-i} \ge s_{-i}|s_i))^{-1}\max_{k\neq i} h_k]$.

Because $c_j(s_{-j}, s_j, \cdot)$ is a strictly increasing bijection, the set $X_{-i}(x_i, s)$ is a Cartesian product of bounded intervals $[\underline{x}_j, \bar{x}_j]$ for $j \neq i$. The proposition also bounds the equilibrium bid b, and as a result, it provides an upper bound to the markup that bidders set in equilibrium.

COROLLARY. If $b = \beta(s_i, x_i, x_{-i})$ and $b' > b$, then

$$(b' - b)\frac{\Pr(S_{-i} \ge s_{-i}(b')|s_i)}{\Pr(s_{-i}(b) \ \le\ S_{-i} \ngeq s_{-i}(b')|s_i)}$$

is bounded above by

$$\bar{c}_i(x_i) + (1 - \Pr(S_{-i} \ge s_{-i}|s_i))^{-1}\max_{k\neq i} h_k - \min_{k\neq i} E(C_i|s_i, s_k, S_{-ik} \ge s_{-ik}(b), x_i).$$

Taking the limit as $b' \to b$ yields an upper bound for the markup term in equation (5). The proofs of proposition 2 and its corollary are in appendix B.2.

C. *Observables*

I consider situations where the data consist of a sample of independent auctions. In each auction there is a draw of signals and costs from F. Each bidder submits a bid after observing his own private signal and all public information. The researcher observes bids and public information.

Let B_i denote the bid made by i, $H_{B_i|x}$ its CDF conditional on public information, and $H_{B|x}$ the joint distribution of bids conditional on public information. If bidder i decides not to participate in an auction or submits a bid with zero probability of winning, his bid is recorded as infinity; therefore, $\sup_{b_i\in\mathbb{R}} H_{B_i|x}(b_i)$ may be less than one.[6] The observed support of x is X^{o},

[6] Equilibrium strategies in Reny and Zamir (2004) prescribe $b^* < \beta_i(s_i) < \infty$ for $\phi_i < s_i$. Allowing these bids not to be recorded ensures that the identification argument does not rely on bids that win with zero probability.

and X_i^o is that of x_i. For simplicity, assume that $X^o = X_1^o \times \dots \times X_n^o$. It will be useful to define $Q_{B_i|x}$ as the quantile function of $B_i|x$ and $H_{M_i|B_i=b_i,x}$ as the CDF of $M_i = \min_{j\neq i} B_j$ conditional on $B_i = b_i$, $X = x$.

If bidder i plays a monotone pure strategy in every market configuration x, then his behavior is described by a bid function $\beta_i(s_i, x)$ that is increasing in its first argument. Bid functions $\beta = \{\beta_i\}_{i=1..n}$ describe how each bidder behaves under different private and public information. Let H be the collection $\{H_{B|x} : x \in X^o\}$. For a given F_S, it is said that β generates H if the repeated play of strategy profile β, given the joint distribution of signals F_S, generates the joint distribution of bids H, that is, if $H_{B|x}(b) = \Pr(\cap_{i=1}^n \{\beta_i(S_i, x) \le b_i\})$ for all (x, b).

Throughout the paper it is assumed that the observed data are generated by the repeated play of equilibrium strategies.

Assumption A7. The observables H are generated by β, given F_S, where β are bid functions that constitute a Bayes-Nash equilibrium in monotone pure strategies for every $x \in X^o$.

This assumption does not rule out existence of multiple equilibria, but it requires that bidders always use the same equilibrium strategies in all games indexed by x.[7] It rules out nonequilibrium behavior and equilibrium play of nonmonotone strategies if such an equilibrium exists.

III. Identification

This section discusses identification of the payoff-relevant characteristics of primitives: the joint distribution of signals and each bidder's full-information costs.

A. *Joint Distribution of Signals*

Assumption A7 provides the key links between observables and economic primitives. By monotonicity, bid functions and marginal bid distributions are inverses of each other: $H_{B_i|x}(\cdot) = \beta_i^{-1}(\cdot, x)$ and

$$H_{B|x}(b) = \Pr(\cap_{i=1}^n \{\beta_i(S_i, x) \le b_i\}|x) = F_{S|x}(s),$$

where $s_i = H_{B_i|x}(b_i)$ for all i and $b = [b_1, \dots, b_n]$. This expression implies the following lemma.

Lemma 1. Under assumptions A2 and A7,

i. if there is a vector b such that $s = [H_{B_i|x}(b_i)]_{i=1}^n$, then $F_{S|x}(s)$ is identified by $H_{B|x}(b)$; and

[7] Similar assumptions are required to identify and estimate dynamic models (Bajari, Benkard, and Levin 2007).

ii. if there are $x \in X^{o}$ and $b \in \mathbb{R}$ such that $H_{B_i|x}(\cdot)$ is continuous and $H_{B_i|x}(b) = 1$ for all i, then $F_{S|x}(s)$ is identified from $H_{B|x}$ for every $s \in [0,1]^n$ by

$$F_{S|x}(s) = H_{B|x}\left(\left\{Q_{B_j|x}(s_j)\right\}_{j=1}^{n}\right).$$

The first part implies that assumption A1 places testable restrictions on the distribution of bids across different vectors of cost shifters. Similarly, under assumptions A2–A5, the corollary of proposition 1 imposes testable restrictions on the marginal distribution of bids. The following lemma formalizes these intuitions.

LEMMA 2. Assume A2 and A7. The model has the following testable implications.

i. Assumption A1 implies that $H_{B|x}(b) = H_{B|x'}(b')$ whenever (x, x') and (b, b') are such that $H_{B_i|x}(b_i) = H_{B_i|x'}(b_i')$ for all i.
ii. Assumptions A2–A5 imply that for every $x \in X^{o}$, $H_{W|x}(\cdot)$ has support over an interval $[b_*, b^*]$, each $b \in [b_*, b^*]$ belongs to the support of bids of at least two bidders, and both $H_{W|x}(\cdot)$ and $\{H_{B_i|x}(\cdot)\}_{i=1}^{n}$ are continuous on $[b_*, b^*)$. Under the additional assumption that bidders never bid below the expected cost conditional on winning, $H_{W|x}(\cdot)$ and $\{H_{B_i|x}(\cdot)\}_{i=1}^{n}$ are continuous on $[b_*, b^*]$.

Proposition 2 and the second part of lemma 1 yield the following result on identification of the joint distribution of signals. The proof is in appendix B.3.

PROPOSITION 3. Under assumptions A1–A7, if $x_i \times X_{-i}(x_i, [1, \ldots, 1]) \subset X^{o}$ for some bidder i and cost shifter x_i, then there exist $x \in X^{o}$ and equilibrium bid functions $\{\beta_j\}_{j=1}^{n}$ such that for every $s \in [0,1]^n$, $F_S(s)$ is identified from the observables $H_{B|x}$ generated by $\{\beta_j\}_{j=1}^{n}$. Specifically,

$$F_S(s) = H_{B|x}\left(\left\{Q_{B_j|x}(s_j)\right\}_{j=1}^{n}\right).$$

B. *Full-Information Costs*

Another key implication of assumption A7 is that for every signal s_i, $b_i = Q_{B_i|x}(s_i)$ satisfies a best-response condition. For example, the first-order condition (5) holds when competitors' inverse bid functions are differentiable at b_i. The probability $\Pr(S_{-i} \geq s_{-i}(b_i)|s_i)$ that appears on the left-hand side of equation (5) is identified by $\Pr(M_i \geq b_i|b_i, x)$ because conditioning on s_i is equivalent to conditioning on $B_i = b_i$.[8] Therefore, the left-hand side is identified by

[8] See also Athey and Haile (2002) and Campo, Perrigne, and Vuong (2003).

$$\xi_i(b_i, x) \coloneqq b_i - \frac{\Pr(M_i \geq b_i | B_i = b_i, x)}{|(\partial/\partial m) \Pr(M_i \geq m | B_i = b_i, x)|_{m=b_i}|}. \quad (6)$$

The right-hand side is

$$\sum_{j \neq i} \pi_j^{(i)} E\big(C_i | s_i, s_j, S_{-ij} \geq s_{-ij}(b_i), x_i\big) = \lim_{\varepsilon \downarrow 0} E(C_i | S_i = s_i, S_{-i} \in L_\varepsilon(b_i, x), x_i), \quad (7)$$

where $L_\varepsilon(b_i, x) \coloneqq \{s_{-i}(b_i, x) \leq s_{-i} \ngeq s_{-i}(b_i + \varepsilon, x)\}$ is a subset of competitors' signals. It will be useful to work with this expression for small but positive ε. Each event $L_\varepsilon(b_i, x)$ can be described in terms of the observables. Let $R(b_i, x) = \{s_{-i} : s_j \geq H_{B_j|x}(b_i)\}$ denote the set of competitors' signals for which bidder i wins with bid b_i in configuration x. This set is an $(n - 1)$-dimensional rectangle in the space of signals. Then $L_\varepsilon(b_i, x) = R(b_i, x) \cap R(b_i + \varepsilon, x)^c$; that is, it is the set for which i wins with bid b_i but loses with bid $b_i + \varepsilon$. If $n = 3$, this set has an L shape in the unit square. To sum up, the left-hand side of equation (5) is identified from observables, while the right-hand side is the expectation of $c_i(S_{-i}, s_i, x_i)$ conditional on $S_{-i} \in L$ where L is a set that can be written in terms of observables.[9]

As noted by Haile, Hong, and Shum (2004), assumption A3 and the null hypothesis of private values imply that the marginal costs do not depend on x_{-i}, which is a testable prediction. The identification argument in this paper shows how to recover $c_i(s_{-i}, s_i, x_i)$ by exploiting rich variation in competitive conditions. Intuitively, x_{-i} can be thought as an $(n - 1)$-dimensional instrument that identifies the effect of S_{-i}, which also has dimension $n - 1$, on the full-information cost.

It is illustrative to consider a first-price auction between two bidders: 1 and 2. Because there is only one competitor to tie with, the event "tie with a competitor for the lowest bid" has a simple representation in the space of competitors' signals: $S_2 = H_{B_2|x}(b_1)$, that is, the event where bidder 2 receives signal $H_{B_2|x}(b_1)$, which is the signal that prompts 2 to bid b_1. Bidder 1's first-order condition can be written as

$$\begin{aligned} \xi_1(b_1, x) &= E(C_1 | S_1 = H_{B_1|x}(b_1), S_2 = H_{B_2|x}(b_1), x_1) \\ &= c_1(H_{B_2|x}(b_1), H_{B_1|x}(b_1), x_1). \end{aligned} \quad (8)$$

[9] Formally, eq. (5) is an integrodifferential equation that can be written more succinctly as $\xi_i = \Omega_i c_i$, where ξ_i is defined in eq. (6) and Ω_i is an integrodifferential operator that depends on observables such that for a function $c(s_{-i}, s_i, x_i)$, $\Omega_i c = (\partial/\partial t) E(c(S_{-i}, S_i, x_i) | S_{-i} \geq H_{B_{-i}|x}(t), S_i = s_i)$, where $s_i = H_{B_i|x}(b)$. The identification result below amounts to constructively showing the existence of the inverse linear operator Ω_i^{-1} and obtaining $c_i = \Omega_i^{-1} \xi_i$. See proposition 2.5 in the online supplementary material for the formal details about the construction of the inverse linear operator.

The identification argument is straightforward. The identified marginal revenue ξ_1 equals the full-information cost of bidder 1 evaluated at $s_1 = H_{B_1|x}(b_1)$, $s_2 = H_{B_2|x}(b_1)$, and x_1, where $x = [x_1, x_2]$. Evaluating this expression for different values of (x_2, b_1) results in the full-information cost evaluated at different pairs of signals while x_1 is held constant. The exclusion restriction (assumption A3) avoids confounding the effect of x_2 on the pairs of signals with a direct effect on costs. It follows that $c_1(\cdot, \cdot, x_1)$ is identified on $\{(s_1, s_2) : s_1 \in [0, \phi_1([x_1, x_2])], s_2 = H_{B_2|[x_1,x_2]}[Q_{B_1|[x_1,x_2]}(s_1)], x_2 \in X_2^o\}$.

If $n > 2$, the event "tie with a competitor for the lowest bid" has a more complex representation, because there is more than one competitor to tie with. It takes the form "S_{-i} such that $S_j = s_j$ for some competitor j and $S_k \geq s_k$ for all other competitors." Consider an auction with three bidders: 1, 2, and 3. Figure 1 shows the pair of competitors' signals that makes them both bid exactly b_1, along with $L_\varepsilon(b_1, x)$, an L-shaped set containing all competitors' signals such that their minimum bid is in $[b_1, b_1 + \varepsilon]$. The

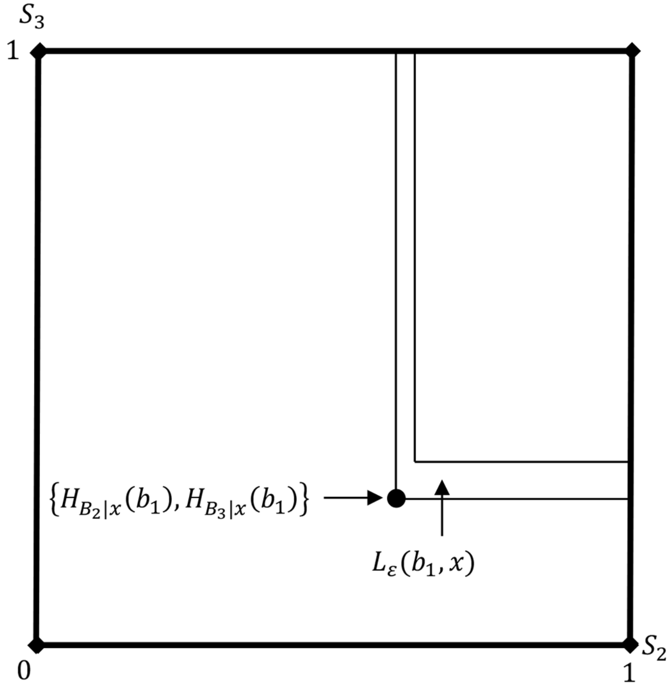

FIG. 1.—This figure shows the pair of competitors' signals that make them both bid exactly b_1, along with $L_\varepsilon(b_1, x)$, an L-shaped set containing all competitors' signals such that their minimum bid is in $[b_1, b_1 + \varepsilon]$. The bidder's first-order optimality condition identifies the expected cost conditional on this set.

right-hand side of equation (5) denotes the expected cost conditional on this set as $\varepsilon \to 0$. This limit can be written as

$$\sum_{j=2,3}\left[\frac{h_{B_j|x}(b_1)f(s_1, s_j)\Pr(S_{5-j} \geq s_{5-j}|s_1, s_j)}{\sum_{k=2,3}h_{B_k|x}(b_1)f(s_1, s_k)\Pr(S_{5-k} \geq s_{5-k}|s_1, s_k)}\right]E(C_1|s_1, s_j, S_{5-j} \geq s_{5-j}, x_1), \tag{9}$$

where $s_j = H_{B_j|x}(b_1)$ and $h_{B_j|x}$ is the density of $B_j|x$. The term in the square brackets is the probability that bidder i ties with bidder j conditional on tying with at least one competitor. These probabilities are identified from the observed distribution of bids. The objects of interest are the two terms $E(C_1|s_1, s_j, S_{5-j} \geq s_{5-j}, x_1)$ for $j = 2, 3$. Each of these terms is the expected cost conditional on tying with bidder j while underbidding the other bidder. The best-response condition implies that the identified marginal revenue $\xi_1(b_1, x)$ equals limit (9), which results in a single equation for two unknowns. The information provided by a single bid is insufficient to identify these terms. It is shown below that using bids under different values of competitors' cost shifters will generate additional information that identifies the expected costs conditional on winning: $E(C_1|s_1, S_2 \geq s_2, S_3 \geq s_3, x_1)$.

Consider again figure 1, and fix $\varepsilon > 0$. Keeping s_1 and x_1 constant, find a triplet $[x_2', x_3', t]$ so that $s_1 = H_{B_1|[x_1,x_2',x_3']}(t)$ and $L_\varepsilon(t, [x_1, x_2', x_3'])$ stacks on top of the previous L-shaped set. The expected cost conditional on the union of these two sets is equal to a weighted average of the expected cost conditional on each L-shaped set. The weights are given by the probability of each set, which is identified from the joint distribution of signals. The expected cost conditional on each L-shaped set can be approximated by $\xi_1(b_1, x)$ and $\xi_1(t, [x_1, x_2', x_3'])$, respectively. This process can be repeated to obtain a weighted average over the whole rectangle $\{S_j \geq s_j\}_{j=2,3}$, as shown by figure 2. Under assumption A1, signals are independent from cost shifters, and this average equals $E(C_1|s_1, S_2 \geq s_2, S_3 \geq s_3, x_1)$, which is the expected cost conditional on the event where 1 wins the auction. If $\varepsilon \to 0$, the approximation error for each L-shaped set vanishes and, under some technical conditions, the average becomes a Riemann integration.

If $E(C_1|s_1, S_2 \geq s_2, S_3 \geq s_3, x_1)$ is identified around a neighborhood of (s_2, s_3), then its derivative is also identified. Differentiating it with respect to s_2 and s_3 yields

$$\begin{aligned}&\frac{d^2E(C_1|s_1, S_2 \geq s_2, S_3 \geq s_3, x_1)\Pr(S_2 \geq s_2, S_3 \geq s_3|s_1)}{ds_2 \times ds_3}\\ &\quad = c_1([s_2, s_3], s_1, x_1)f_{S_2,S_3|s_1}(s_2, s_3),\end{aligned} \tag{10}$$

where $f_{S_2,S_3|s_1}$ is the density of competitors' signals conditional on $S_1 = s_1$. The full-information cost is obtained dividing the left-hand side of equation (10) by the density of signals, both identified objects.

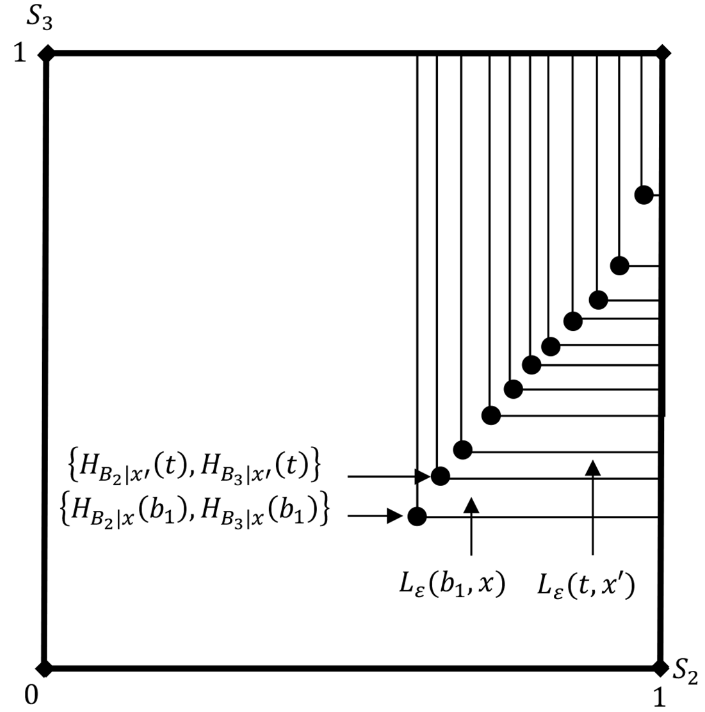

FIG. 2.—Variation in $[x_2, x_3]$ is used to find a different L-shaped set that stacks on top of $L_\varepsilon(b_1, x)$, holding s_1 and x_1 constant. This set is $L_\varepsilon(t, x')$ with $x' = [x_1, x_2, x_3]$. The expected cost over the union of $L_\varepsilon(b_1, x)$ and $L_\varepsilon(t, x')$ is equal to a probability-weighted average of the expected cost over each L-shaped set. These weights are identified from the joint distribution of signals. This process can be repeated to obtain a probability-weighted average over the whole rectangle $\{S_j \geq s_j\}_{j=2,3}$.

A formal identification proof based on these intuitions has to overcome a couple of technical challenges. First, it has to be possible to construct an ever-finer partition of the event $\{S_{-i} > s_{-i}\}$ composed of identified pivotal sets. Second, the process of averaging over these sets has to converge to the Riemann integral of the full-information cost over $\{S_{-i} > s_{-i}\}$.

The construction of an ever-finer partition of the event $\{S_{-i} > s_{-i}\}$ requires variation in competitors' cost shifters that induces variation in equilibrium bidding strategies. Keeping x_i and s_i constant, for each point $\hat{s}_{-i}$ along the integration path, it is necessary to find a vector of competitors' cost shifters x_{-i} and a bid b so that $H_{B_j|x_i,x_{-i}}(b) = \hat{s}_j$ for all $j = 1, \ldots, n$, and $Q_{B_i|x_i,x_{-i}}(s_i) = b$.

DEFINITION 1. The distribution H exhibits sufficient variation for bidder i at signal vector s^* and cost shifter $x_i \in X_i^{\text{o}}$ if for every $s \geq s^*$, there are cost shifters $x_{-i} \in X^{o}_{-i}$ and a finite bid b such that $H_{B_j|x_i,x_{-i}}(b) = s_j$ for all $j = 1, \ldots, n$, and $Q_{B_i|x_i,x_{-i}}(s_i) = b$.

There is one additional condition that the observables have to satisfy to ensure that the number of L-shaped sets that have to be stacked to cover the full rectangle is finite. This condition requires that if i bids b and b', respectively, when he receives signals s_i and s_i', then i's competitors bid in the interval $[b, b']$ with nonvanishing probability. In other words, it is required that bidder i face nonvanishing competition in that interval.

Definition 2. The distribution H exhibits nonvanishing competition for bidder i at signal vector s^* and cost shifter x_i if for every $\delta > 0$ there is a positive constant κ such that

$$\frac{H_{M_i|B_i=b,x}(b') - H_{M_i|B_i=b,x}(b)}{1 - H_{M_i|B_i=b,x}(b')} \geq \kappa\delta$$

for every (b, b', x_{-i}) such that $x_{-i} \in X^{o}_{-i}$, $b = Q_{B_i|x}(s_i)$ for some $s_i \geq s_i^*$, $H_{B_j|x}(b) \geq s_j^*$ for all $j \neq i$, and $b' = Q_{B_i|x}(H_{B_i|x}(b) + \delta)$.

The main theorem shows that with sufficient variation in observables and nonvanishing competition, the full-information costs are identified. The proof involves showing that the process described above converges to the expected value of the full-information cost conditional on sets of the form $\{S_{-i} > s_{-i}\}$.

Theorem 4. Under assumptions A1–A3 and A7, the full-information cost $c_i(s_{-i}, s_i, x_i)$ is identified for all $s \geq s^*$ if H exhibits sufficient variation and nonvanishing competition for bidder i at signal vector s^* and cost shifter x_i.

The detailed proof of theorem 4 is in appendix B.4. Appendix C contains two extensions to theorem 4. Appendix C.1 shows that it is possible to relax the assumption of independence between cost shifters and signals. Appendix C.2 shows that theorem 4 can be applied in a large class of auction formats that includes all-pay and second-price auctions. As long as the object is awarded to the bidder who submits the lowest bid, there will be an optimality condition analogous to equation (5) that ensures that the marginal expected revenue of increasing the probability of winning equals the marginal cost. This condition can be used to achieve identification.

C. *Extensions and Limitations*

1. Sufficient Variation and Nonvanishing Competition

Theorem 4 shows that the observable distribution of equilibrium bids contains all the necessary information to reconstruct the underlying information structure of a general interdependent-cost model without relying on any particular parametric or functional form restrictions. The result requires that observables exhibit sufficient variation and nonvanishing competition. It is natural to ask whether equilibrium bid functions typically generate observables that satisfy these conditions.

The following results show that in the first-price auction, observables typically exhibit nonvanishing competition and variation in equilibrium bid functions. The proofs are relegated to appendixes B.5 and B.6, respectively.

PROPOSITION 5. Under assumptions A1–A6, monotone equilibrium bid functions of the first-price auction generate observables H that exhibit nonvanishing competition for every bidder i, at every cost shifter $x_i \in X_i^{\mathrm{o}}$, and vector of signals $s^* > 0$.

PROPOSITION 6. Under assumptions A1–A6, for every bidder i, signal vector $s^* > 0$, and cost shifter $x_i \in X_i^{\mathrm{o}}$, the set of competitors' cost shifters

$$\tilde{X}_{-i}(x_i, s^*) = \left\{ \begin{array}{c} x_{-i} \in \mathbb{R}^{n-1} : \forall\, j \neq i, \\ \underline{c}_j(x_j) \leq \bar{c}_i(x_i) + \left(1 - \Pr\left(S_{-i} \geq s^*_{-i} | S_i = 1\right)\right)^{-1} \max_{k \neq i} h_k \\ \bar{c}_j(x_j) \geq \underline{c}_i(x_i) - \left(1 - \Pr\left(S_{-j} \geq s^*_{-j} | S_j = 1\right)\right)^{-1} \max_{k \neq j} h_k \end{array} \right\} \subset \mathbb{R}^{n-1}$$

is bounded. Provided that $\tilde{X}_{-i}(x_i, s^*) \subset X^{\mathrm{o}}_{-i}$, for every $s \geq s^*$, there are cost shifters $x_{-i} \in X^{\mathrm{o}}_{-i}$, a finite bid b, and equilibrium strategies that generate observables H satisfying $H_{B_j|x_i,x_{-i}}(b) = s_j$ for all $j = 1, \ldots, n$.

Proposition 6 stresses that variation of competitors' cost shifters over a bounded set may be sufficient to identify $c_i(\cdot, \cdot, x_i)$ for all $s \geq s^*$. It is also evident that $\tilde{X}_{-i}(x_i, s) \subset \tilde{X}_{-i}(x_i, s')$ if $s \geq s'$. Identification of $c_i(\cdot, \cdot, x_i)$ for lower values of signals may require a larger support of cost shifters. The set $\tilde{X}_{-i}(x_i, s)$ grows unboundedly when at least $n - 1$ components of the vector of signals approach zero.[10] Identification of $c_i(\cdot, \cdot, x_i)$ at these extreme values of signals may require cost shifters to vary over an unbounded support or some extrapolation from less extreme values.

Proposition 6 would imply that H exhibits sufficient variation if $H_{B_i|x}(\cdot)$ were also strictly increasing at b. It would be sufficient to show that all equilibrium strategies $\beta_i(\cdot, x)$ are continuous on the interval $[0, \phi_i(x)] \subset [0, 1]$ and that $\beta_i(\phi_i(x), x) = b^*(x)$, where $b^*(x)$ is the highest end point of the support of the winning bid $W|x$. The theoretical literature has shown that first-price auctions with nonbinding reserve prices have continuous equilibrium bid functions in asymmetric private-values models (Lebrun 2002), symmetric interdependent models (Milgrom and Weber 1982; Bergemann,

[10] It is not possible to generate an equilibrium where $\beta_i(s_i, x) = b_*$ for a strictly positive s_i because that would imply that bidder i bids b_* for all signals below s_i, which leads to an atom at b^*, which contradicts the results in Maskin and Riley (2000b). It is possible to generate an equilibrium such that $\beta_i(s_i, x) = b_* + \varepsilon$ for a small and positive ε and $s_i > 0$. As $\varepsilon \to 0$ while s_i is fixed, this equilibrium requires an increasingly large asymmetry between bidder i and its competitors.

Brooks, and Morris 2017), and two-bidder general asymmetric interdependent models that satisfy the assumptions in this paper (Lizzeri and Persico 2000; Maskin and Riley 2000a). It is an open question whether and under what conditions those results can be extended to the general interdependent-cost model with more than two bidders.

Another open question is whether every observable H that satisfies the testable implications of assumption A1 and exhibits nonvanishing competition and sufficient variation can be rationalized by a general interdependent-cost model satisfying assumptions A1–A3. Intuitively, the identification result is derived from a set of first-order necessary conditions for a local extremum. The uniquely identified primitives may fail to satisfy the conditions for local and global maximization. In the independent-private-value model, for example, theorem 1 in Guerre, Perrigne, and Vuong (2000) shows that some distributions of bids cannot be rationalized because they would imply that pseudovaluations are decreasing in bids and that the second-order conditions for local maximization do not hold.

The identification result can be used to shed light on the two open questions mentioned above. Any observable H that satisfies the testable conditions of assumption A1 can be represented by a copula F and a pair of functions $\gamma(t, x)$ and $b(t, x)$ such that $t = H_{W|x}(b(t, x))$ and $\gamma_j(t, x) = H_{B_j|x}(b(t, x))$ for every j. Intuitively, for a fixed vector of market conditions x, $b(t, x)$ returns the tth quantile of the distribution of winning bids, and $\gamma(\cdot, x)$ is a monotone parametric curve that maps $[0, 1] \to [0, 1]^n$ and returns the vector of marginal CDF of bids evaluated at the tth quantile of the distribution of winning bids. Sufficient variation and nonvanishing competition can be verified directly on γ. Proposition 2.5 in the online supplementary material shows that, given an affiliated copula F and a function $\gamma(\cdot, \cdot)$ satisfying some regularity conditions that imply nonvanishing competition and sufficient variation, there exists a nonempty open convex cone of functions $b(\cdot, \cdot)$, denoted by B', such that $b \in B'$ if and only if observables H, represented by (F, γ, b), are rationalizable by an interdependent-cost model with full-information costs that are strictly monotone in signals. This result has two additional implications. First, it is possible to construct primitives of the interdependent-cost model and equilibrium strategies that generate observables satisfying the continuity condition required to close the gap between H exhibiting sufficient variation and proposition 6. Second, if $b \in B'$, the first-order conditions used to identify the primitives of the model are also sufficient for optimality. In other words, the local and global optimality conditions do not impose any additional testable restrictions. Conversely, if $b \notin B'$, then the identified primitives violate the strict-monotonicity assumption in A5; thus, local and global optimality conditions may impose additional testable restrictions. In the general interdependent model, violations of the monotonicity assumption do not imply a violation of optimality as in the independent-private-values case.

2. Reserve Prices

Lemmas 1 and 2 and theorem 4 hold irrespective of whether reserve prices bind. On the one hand, binding reserve prices can limit the ability of cost shifters to generate sufficient variation in equilibrium bid functions. For example, if reserve prices are set so that the probability of awarding the contract is always below 0.9, it is not possible to find (x, b) such that $\{H_{B_j|x}(b) = s_j\}_{j=1}^{n}$ when $\Pr(S \geq s) < 0.1$. This prevents identification of $E(C_i|s_i, S_{-i} \geq s_{-i}, x_i)$ and $\Pr(S \geq s)$ for any such vector of signals. It may still be possible to identify $E(C_i|s_i, S_{-i} \geq s_{-i}, x_i)$ and the full-information costs for some subset of the set of signals such that $\Pr(S \geq s) \geq 0.1$ if every integration path eventually reaches a point such that the reserve price binds. On the other hand, if reserve prices vary for exogenous reasons, they can generate additional variation in observables that can help to identify full-information costs under limited variation in cost shifters. A complete analysis of these possibilities would require a better understanding of how reserve prices affect equilibrium bid functions in interdependent-cost auctions than is currently available.

3. Entry Costs

In interdependent-cost auctions, bidders may find it optimal not to participate if they receive a high signal. At the maximum signal at which they participate, they must be indifferent between staying out and submitting the bid prescribed by equilibrium strategies. In the absence of entry costs, this indifference condition is $b = E(C_i|\phi_i, S_{-i} \geq s_{-i}, x_i)$, where ϕ_i is the maximum signal at which i participates and $s_{-i} = \{H_{B_j|x}(b)\}_{j\neq i}$. This condition is used in the proof of theorem 4 to finalize the integration procedure. If there are positive entry costs, the indifference condition may change. The auctions literature identifies two polar cases for entry costs. Samuelson (1985) considered an entry-cost model in which bidders make their participation decision after learning their signals. Therefore, if the entry cost is K, the indifference condition at the maximum signal for which bidder i is observed to participate is $(b - E(C_i|\phi_i, S_{-i} \geq s_{-i}, x_i))\Pr(S_{-i} \geq s_{-i}|\phi_i) - K = 0$. Rewriting this condition as

$$\int_{\{\tau : \tau \geq s_{-i}(b)\}} (b - c_i(\tau, \phi_i, x_i) - K\delta_1(\tau)) f(\tau|\phi_i)\, d\tau = 0,$$

where $\delta_1(\tau)$ is a Dirac delta function at the vector $[1, 1, \ldots, 1]$, stresses that models with entry costs will be observationally similar to models where the full-information costs exhibits a sharp increase as competitors' signals approach the point $[1, 1, \ldots, 1]$. To distinguish models with sharply increasing full-information costs from models with entry costs, it will be necessary to identify $E(C_i|s_i, S_{-i} \geq s_{-i}, x_i)$ for $s_{-i} \to [1, 1, \ldots, 1]$. One feature that the two types of models share is that in equilibrium bidders may find it optimal not to participate when they receive signals close to 1 even if reserve prices do not bind. As a result, while proposition 6 still applies, equilibrium bid functions may exhibit discontinuities that generate observables with insufficient variation to identify $E(C_i|s_i, S_{-i} \geq s_{-i}, x_i)$ and discriminate between the two types of model.

Levin and Smith (1994) considered a model in which bidders make their participation decision before learning their signals. It is useful to think about bidders having to pay an information acquisition cost to learn their signal. Suppose first that information acquisition costs are deterministic and that a bidder's decision to pay this cost is a function of the vector of auction characteristics x. It is straightforward to identify the set of bidders who pay the cost at each x. Bidders may still decide not to participate after learning their signal, which results in the same indifference condition as in the case of no entry costs. Theorem 4 can be applied with no modifications, but the assumptions of proposition 6 fail. Cost shifters may not be able to generate sufficient variation. Intuitively, it may be impossible to find a tie for a configuration of signals with low s_i and very high s_{-i}. The competitive conditions have to be such that the probability of winning for bidder i will be close to zero, even if he received a good signal. Therefore, i may decide not to pay the information acquisition cost. Now consider the case where the information acquisition cost is random and unobserved to the econometrician. If a bidder is observed not to participate, it is impossible to know whether he was discouraged by the information acquisition cost or by the signal. This is a violation of the monotonicity assumption in A7, by which nonparticipation can be interpreted as an infinite bid that results from a high signal in the auction. Theorem 4 does not apply.

4. Identified Counterfactuals

The identified payoff-relevant primitives are sufficient to compute the equilibrium in counterfactual situations that can be represented by a Bayesian game with primitives that depend only on the identified information structure (e.g., changes in reserve prices, subsidies to some preferred bidders, participation fees, changes in auction format from first-price to second-price auction, random allocation among the two lowest bidders). As long as the counterfactual situation does not require that bidders make bidding or bargaining decisions after they learn additional information, their payoff functions still depend only on their full-information costs. For example, consider a counterfactual policy that allows resale or subcontracting after bidders learn their costs. Suppose that at the subcontracting stage there is no private information and that each bidder's publicly known cost is his full-information cost evaluated at the realized signals plus an idiosyncratic ex post shock. Resale market opportunities are more profitable when the variance of the ex post shock is larger, because the expected minimum competitors' cost is lower. Because bidders should take into account the possibility of subcontracting (or resale, as in Haile 2001), this policy results in a more competitive auction environment if the variance of the ex post shock is large. This variance is not an identified characteristic of the model. Therefore, it is not possible to compute this counterfactual using the identified information structure.

5. Practical Considerations

As Athey and Haile (2007) point out, the ideal of understanding the empirical content of an economic model without relying on parametric or distributional assumptions does not necessarily imply a preference for nonparametric estimation methods. In finite samples, nonparametric methods can be subject to the curse of dimensionality. In applied work, there is a clear trade-off. On the one hand, one would like to control for economically relevant covariates and use a rich model that captures all the economically relevant effects. On the other hand, the richer the model, the more severe the curse of dimensionality becomes. When discussing parametric assumptions in discrete choice models, Berry and Haile (2014, 1752) argue that "the functional form and distributional assumptions typically used in practice play their usual roles: approximation in finite samples and compensation for the gap between the exogenous variation available in practice and that required to discriminate between all nonparametric models."

In a general nonparametric model, a bidder's full-information cost depends on all other competitors' signals. Identification of the effect of $n-1$ competitors' signals on one bidder's full-information costs requires variation in, at least, $n-1$ excluded shifters or instruments. In my application to highway procurements in Michigan, I assume that the model has a Gaussian information structure. This assumption allows me to parameterize cost interdependence by a single parameter. I use distance from the bidder's plant to the project site as a cost shifter. For fixed plant locations, the full vector of bidders' distances exhibits two-dimensional variation because it is a function of the geographical coordinates of the project location. This variation suffices to keep a bidder's distance to the project fixed and takes advantage of the variation in competitive conditions induced by the second dimension as the project location varies along a circumference centered at the bidder's plant. The model is still fully asymmetric. Bidders may have signals that allow them to predict their own costs with different levels of precision. For example, the full-information cost of a bidder that has access to a perfectly informative signal will depend only on this signal and on the bidder's cost shifter.

This bidder will not be exposed to the winner's curse. As the informativeness of the signal decreases, the effect of competitors' signals on the full-information cost and the bidder's exposure to the winner's curse increase.

IV. Winner's Curse in Michigan Highway Procurement

The Michigan Department of Transportation (MDOT) uses first-price sealed-bid auctions to award around 1,050 highway construction and maintenance contracts per year, at a total cost of 1.2 billion dollars. On each monthly letting date, 150–200 firms submit sealed bids for one or more of 50–70 contracts. Firms may participate in as many auctions per letting date as allowed by their work-type and financial prequalification status. The work-type prequalification status of a firm is a list of all the classes of work that the firm can perform. There are 52 work classifications, and the typical firm is prequalified for 6–10 of these. The financial rating of a firm is the maximum dollar amount of contracts it can have pending with the MDOT. While the work-type classification is public information, the financial rating of a firm is confidential.

Contracts are advertised for at least 45 days before the letting date, allowing bidders to have detailed information about them. The information about future projects is less precise. The MDOT publishes a 3-month projection of future projects and a 5-year Transportation Improvement Plan, but these are subject to frequent changes and updates.

Before submitting a bid, firms download the technical plan and submit a form to become eligible to bid. The MDOT keeps an updated list of both eligible bidders and plan holders that is publicly available on its website. Firms may submit the eligibility form as late as 5:00 p.m. on the day preceding the letting date and may not appear in the list of eligible bidders before the bid submission deadline. Moreover, eligible bidders and plan holders often choose not to bid. As a result, firms are unable to predict with certainty the set of participants in an auction. Thus, it is likely that firms base their expectations of competition largely on the location and technical characteristics of the project.

Each contract describes a list of tasks that the contractor has to perform. A task specifies a description and a quantity, for example, earth excavation, 600 cubic yards. For each task, the MDOT engineer sets a unit price, hence the total estimated cost of the task is the unit price times the quantity. The engineer's estimate for the contract is the sum of all the tasks' estimated costs. Bidders submit a unit bid for each task, and the total bid is the inner product of unit prices and quantities. The bidder with the lowest total bid wins. If there are no modifications to the contract, the MDOT pays the winner its total bid upon project completion.

Although there is no formal reserve price, the procuring agency has the option to reject all bids if the lowest bid exceeds 110% of the MDOT engineer's estimate. In the case that the bids are rejected, the project can be revised and offered in a future letting date. If the agency accepts a bid over that limit, it must justify in writing why the estimate was not correct or why the bids were excessive. From 2001 to 2010, the lowest bid exceeded 110% of the engineer's estimate in 11% of cases, and 15% of these were rejected.

This paper focuses on contracts where contractors must be prequalified to perform hot-mixed asphalt (HMA) work. HMA is the pavement material used in 96% of all paved roads and streets in the United States. It consists of asphalt or bitumen and mineral aggregate heated and mixed in a plant. The mix must be trucked to the project site, laid on the road, and compacted while the mix is sufficiently hot (above 275°F/135°C). The temperature of the mix at the time of compaction is key to the quality of the pavement mat. Once the mix falls below 175°F/79°C, it cannot be further compacted, and a poorly compacted mat will age faster. HMA pavement projects are rarely performed during winter in Michigan for this reason.

Trucking time from the plant to the project location is an important determinant of costs not only because of transport costs but also because of the cooling process. During transport, the surface layer of the mix cools faster than the inner mass. Once the mix is dumped into the paver and laid on the road, these temperature differentials may persist and result in cool spots in the pavement mat that cannot be properly compacted. These problems can be mitigated by incurring additional labor and rental costs, for example, by using a material transfer vehicle to remix on site. Thus, firms that own plants located close to the project have lower transportation costs and lower costs associated with excess cooling of asphalt.[11]

There are several sources of interdependence that may arise even in standard, small, and straightforward projects involving the application of HMA. At the time of the auction, bidders may have imperfect information about payoff-relevant conditions at the time of execution. Obvious candidates include weather and traffic conditions. The likelihood and conditions of potential contract renegotiation may also be an important

[11] There may be other types of cost shifters. For example, Balat (2017) uses backlog as a shifter of costs in a dynamic auction model of highway procurements in California. I did not find a significant effect of backlog on bidding behavior. Paving projects are carried out throughout the year in California, which may lead to different dynamic considerations. Asphalt plants procure their aggregates from local quarries. Variation in local prices of aggregates over time could also be used as a cost shifter affecting a particular plant. Moreover, asphalt specifications can vary across projects. Asphalt specifications determine the mix of aggregates of different sizes. Bidders have access to local quarries that specialize in some type of aggregate but not in others. The description of the project would then contain not only the location of the project but also a vector of shares for different aggregate size. If each bidder has an ideal specification, this variation could also be used to augment the effective dimension of the distance vector if it is necessary to identify a model that allows for richer interactions among signals.

source of interdependence (Bajari, Houghton, and Tadelis 2014). The future prices of inputs and availability of subcontractors are also unknown and common determinants of completion costs.

A. *Data*

Data for all auctions and bids from 2001 to 2010 are available from the MDOT. For each auction, the data include the project's description, location, and prequalification requirements; the engineer's estimate of the total cost of the project; and the list of participating firms and their bids. To obtain the geographical coordinates for each project location, I match the road names in the description to the database of roads available at the Michigan Geographical Data Library.

The locations of each firm's plants and mineral aggregate quarries were obtained from several sources: the MDOT contractor directory; individual firm searches using OneSource North American Business Browser, Dun & Bradstreet's Hoovers, and yellowpages.com; firms' websites; and the data on firms collected by Einav and Esponda (2008). A firm location was included in the final data set if it appeared in at least two sources or if it was listed explicitly on a firm's website.

Of the 10,522 MDOT auctions that were run from January 2001 to December 2010, 3,851 auctions required the prime contractor or one of the subcontractors to be prequalified to perform work with HMA, and in 1,925 of the auctions this is the only prequalification requirement for the prime contractor. Table 1 shows the main descriptive statistics of this set of auctions. The typical auction has three bidders. The median engineer's estimate is around \$633,000, while the median winning bid is around \$586,000. It is convenient to normalize bids with respect to the engineer's estimate. Let the normalized bid be $b = \text{Bid}/\text{Engineer} - 1$. The median normalized winning bid is 6% below the engineer's estimate. The median participant is located 38 km (26 miles) from the project, while the median winner is only 22 km (14 miles) away. The average "money left

TABLE 1
Descriptive Statistics: Engineer's Estimate, Bids, and Distances

Variable	*N*	Mean	Standard Deviation	P5	Median	P95
Engineer's estimate (\$000)	1,925	1,145	1,783	119	633	3,899
Lowest bid (\$000)	1,925	1,090	1,771	111	586	3,737
Participants	1,925	3.04	1.3	2	3	6
(2nd-lowest/lowest bid − 1) × 100	1,852	8.2	8.7	.5	5.9	23.4
(Lowest/engineer − 1) × 100	1,925	−6	12.5	−25.9	−6.8	15.2
Distance of winner (km)	1,816	32	40	2	22	91
Distance of bidder (km)	18,778	51	50	4	38	138

Note.—P5, P95 = 5th and 95th percentiles, respectively; 2nd-lowest = second-lowest bid; engineer = engineer's estimate. There were 73 auctions with only one bid and 109 won by a firm for which I did not find any verifiable location.

on the table," or how much higher the second-lowest bid is relative to the lowest, is about 8%.

It is interesting to observe how normalized bids vary with the number of actual participants. Table 2 shows that the average winning bid is decreasing with the number of participants but that the average bid is not. In a standard symmetric independent-private-cost model with exogenous participation, both the expected lowest bid and the expected bid should be decreasing in the number of bidders. Of course, in this setting it is likely that the symmetry and exogenous-participation assumptions fail.

I construct the normalized minimum competitors' bid for each bidder and auction in the sample. If a bidder faces no competition in an auction, then this variable is set to the maximum observed finite normalized bid in the sample. The minimum competitors' bid is set equal to the winning bid in the few cases where the winning bid exceeded 110% of the engineer's estimate and was rejected by the procurement agency. For each bidder i and auction t, I observe the pair of normalized bid and minimum competitors' bid (b_{it}, m_{it}), which are assumed to be a random draw of (B_i, M_i) distributed according to $H_{M,B|\mathrm{lat}_t,\mathrm{lon}_t}$, where lat_t and lon_t denote the latitude and longitude of the location of the project in auction t.

B. Estimation

I estimate a model with a Gaussian information structure. Signals are normalized to have a standard normal marginal distribution and are denoted by $Z \sim N(0, \Sigma)$. The correlation matrix Σ has a factor structure $LL' + \Lambda$, where L is an $n \times l$ loading matrix and Λ is a positive diagonal matrix that

TABLE 2
NUMBER OF PARTICIPANTS AND NORMALIZED BIDS

		Winning Bid		Average Bid	
Participants	Auctions	Mean	SE	Mean	SE
1	73	1.2	1.0	1.2	1.0
2	658	−5.0	.5	.0	.6
3	727	−6.1	.4	2.1	.5
4	247	−6.5	.8	3.3	.9
5	116	−9.9	1.0	−.9	1.0
6	54	−10.6	1.9	.3	2.2
7	28	−9.4	2.3	1.5	2.3
8	14	−11.1	3.6	.0	4.1
9	6	−15.2	5.1	−5.0	4.9
10	2	−17.9	3.8	−3.3	8.4

NOTE.—Bids were normalized by the engineer's estimate: (Bid/engineer'estimate − 1) × 100. The average winning bid and average submitted bid are tabulated by the number of participants in the auction. For example, there were 658 auctions with only two participants. In these auctions, the normalized winning bid was on average 5% below the engineer's estimate, while the average submitted bid was equal to it. While the average winning bid decreases with the number of participants, the average submitted bid does not.

ensures that the elements on the main diagonal of Σ are all 1. Let x_0 be a vector of observable characteristics and x_i be bidder i's distance to the project. One implication of the Gaussian information structure (see app. A for details) is that the full-information cost is

$$E(C_i|Z = z, x_0, x_i) = \alpha_{ic} + \alpha_{i0}'x_0 + \alpha_{i1}x_i + \alpha_{i2}\sum_{j\neq i}\tilde{\sigma}_{ij}z_j + \nu_i z_i, \qquad (11)$$

where parameter $\tilde{\sigma}_{ij}$ is the ijth element of $(-\Sigma)^{-1}$. Affiliation implies that $\tilde{\sigma}_{ij} \geq 0$.

The parameterization is convenient for three reasons. First, it allows for a flexible parametric structure on the correlation of signals. Second, the full-information cost of each bidder is additively separable in competitors' signals, which will simplify the estimation procedure. Third, the degree of interdependence is summarized by one parameter per bidder: α_{i2}. This parameter is identified using the exclusion restriction on competitors' distance, which, after controlling for own distance, exhibits only one-dimensional variation. The Gaussian information structure nests the private-costs case when $\alpha_{i2} = 0$ and the interdependent-cost case when $\alpha_{i2} > 0$, and it can be falsified when $\alpha_{i2} < 0$.

Every observed bid b_i has to satisfy equation (5), which provides the key link between the model and the data. The markup on the left-hand side can be written in terms of observables H as

$$\iota_i(b_i, x|H) = \frac{\Pr(M_i \geq b_i|B = b_i, x)}{(\partial/\partial m)\Pr(M_i \geq m|B = b_i, x)|_{m=b_i}},$$

which depends on the joint distribution of bidder i's bid and the minimum competitors' bid conditional on covariates x. Under the Gaussian model, the right-hand side is

$$\sum_{k\neq i}\pi_j^{(i)}E(C_i|Z_{ik} = z_{ik}(b_i, x), Z_{-ik} \geq z_{-ik}(b_i, x), x_0, x_i),$$

where $z_{ik}(b_i, x) = \{z_j(b_i, x)\}_{j\in\{i,k\}}$, $z_{-ik}(b_i, x) = \{z_j(b_i, x)\}_{j\neq i,k}$, and $z_j(b_i, x)$ is the Gaussian signal that prompts bid b_i from bidder j. Using the law of iterated expectations, substituting $E(C_i|Z, x_0, x_i)$ from equation (11), and rearranging yields

$$\alpha_{ic} + \alpha_{i0}'x_0 + \alpha_{i1}x_i + \nu_i z_i(b_i, x) + \alpha_{i2}k_i(b_i, x|H),$$

where

$$k_i(b_i, x|H)$$
$$= \sum_{j\neq i}\tilde{\sigma}_{ij}\left[\pi_j^{(i)}z_j(b_i, x) + \sum_{k\neq i,j}\pi_k^{(i)}E\big(Z_j|Z_{ik} = z_{ik}(b_i, x), Z_{-ik} \geq z_{-ik}(b_i, x)\big)\right].$$

This term depends on the observables H. The weights $\pi_j^{(i)}$ and cutoffs $z_k(b_i, x)$ are those implied by the marginal distribution of bids conditional on x evaluated at b_i. The parameters $\tilde{\sigma}_{ij}$ are those implied by the joint distribution of bids. Therefore, the random variable B_i satisfies

$$B_i = \alpha_{ic} + \alpha_{i0}' X_0 + \alpha_{i1} X_i + \alpha_{i2} k_i(B_i, X|H) + \iota_i(B_i, X|H) + \nu_i Z_i, \tag{12}$$

where $Z_i = z_i(B_i, X)$. The term $\iota_i(B_i, X|H)$ is the markup due to market power that arises because bidder i can increase his bid and still enjoy some positive probability of winning. The term $\alpha_{i2} k_i(B_i, X|H)$ is the markup due to the winner's curse. It arises because bidder i knows that he wins the auction only when competitors draw high signals.

Let $\nu_{i\tau}$ be the τth quantile of $\alpha_{ic} + \nu_i Z_i$. By independence of signals and cost shifters,

$$\Pr(\alpha_{ic} + \nu_i Z_i \leq \nu_{i\tau}|X) = \Pr(\alpha_{ic} + \nu_i Z_i \leq \nu_{i\tau}) = \tau.$$

Using equation (12) to substitute $\alpha_{ic} + \nu_i Z_i$ yields

$$E[1(B_i - \alpha_{i0} X_0 - \alpha_{i1} X_i - \alpha_{i2} k_i(B_i, X|H) - \iota_i(B_i, X|H) - \nu_{i\tau} \leq 0) - \tau|X] = 0.$$

This expression implies a set of moment conditions that define the parameters $(\alpha_i, \nu_{i\tau})$, where $\alpha_i = [\alpha_{i0}, \alpha_{i1}, \alpha_{i2}]$. There are two issues that have to be addressed before proceeding to obtain an empirical analog for estimation. First, when bidder i decides not to participate, B_i is censored. The econometrician observes only that the bidder chose a bid higher than any other bid that yields a positive probability of winning. Second, the copula Σ, competitors' strategies β, and the distribution of competitors bids H are not observed but can be consistently estimated.

The censoring problem introduced by nonparticipation can be addressed following the arguments in Buchinsky and Hahn (1998). Monotonicity of the bid functions implies that there will be a cutoff signal such that only signals below it prompt finite bids. This cutoff is $\Phi^{-1}(\phi_i(X|H))$, where Φ is the CDF of a standard normal distribution and $\phi_i(X|H)$ is the probability of participation. The notation stresses that the cutoff is specific to the bidder and to the full vector of cost shifters. Conditional on participation, that is, $\Phi(Z_i) \leq \phi_i(X|H)$, the probability that $\alpha_{ic} + \nu_i Z_i \leq \nu_{i\tau}$ is

$$\begin{aligned}\pi_i(\tau, X|H) &:= \Pr(\alpha_{ic} + \nu_i Z_i \leq \nu_{i\tau}|X, \Phi(Z_i) \leq \phi_i(X|H)) \\ &= \min\left(\frac{\tau}{\phi_i(X|H)}, 1\right).\end{aligned} \tag{13}$$

Using equation (12) to substitute $\alpha_{ic} + \nu_i Z_i$ and noting that $B_i < \infty$ whenever $\Phi(Z_i) \leq \phi_i(X)$ yields the censoring-adjusted moment

$$E[1(B_i - \alpha_{i0}x_0 - \alpha_{i1}X_i - \alpha_{i2}k_i(B_i, X|H) - \iota_i(B_i, X|H) - \nu_{i\tau} \leq 0) \\ - \pi_i(\tau, X|H)|X, B_i < \infty] = 0. \tag{14}$$

A similar moment can be constructed for different values of τ. The moment conditions depend on H, which describes the distribution of competitors' bids and the distribution of competitors' signals conditional on Z_i. Let $\hat{H}$ be an initial estimator of H. I estimate the parameters $(\alpha_i, \nu_{i\tau})$ by $(\hat{\alpha}_i, \hat{\nu}_i)$, which solves the sample minimization problem:

$$\min_{(\alpha_i, \nu_{i\tau})} \| M_n(\alpha_i, \nu_{i\tau}, \hat{H}) \|,$$

where $M_n(\alpha_i, \nu_{i\tau}, \hat{H})$ equals

$$\frac{1}{T_i}\sum_{t=1}^{T_i}\left[1(b_{it} - \alpha_{i0}x_{0t} - \alpha_{i1}x_{it} - \alpha_{i2}k_{it} - \iota_{it} - \nu_{i\tau} \leq 0) - \pi_i(\tau, x_t|\hat{H})\right]x_t,$$

where $t = 1, \ldots, T_i$ indexes all the auctions where i participates, $k_{it} = k_i(b_{it}, x_t|\hat{H})$, and $\iota_{it} = \iota_i(b_{it}, x_t|\hat{H})$. Instead of estimating α_{ic} and ν_i, I estimate $\nu_{i\tau}$, which corresponds to the τth quantile of $\alpha_{ic} + \nu_i Z_i$. While M_n is not smooth in parameters $(\alpha_i, \nu_{i\tau})$, $E[M_n(\alpha_i, \nu_{i\tau}, \hat{H})]$ is. Chen, Linton, and Van Keilegom (2003) provided conditions that ensure consistency and asymptotic normality of this class of estimators and asymptotic validity of the bootstrap. These conditions are verified in section 3 of the online supplementary material.

The parameters can be estimated using the techniques developed for instrumental variable quantile regressions. I use the procedure proposed by Chernozhukov and Hansen (2006) in remark 5 and Chernozhukov and Hansen (2008) in comment 3. The ordinary least squares projection of k_i on $(1, x_0, x_i, x_{-i})$ is used as the instrument for k_i. The constant, x_0, and x_i are used as instruments for themselves. I construct a grid of values $\{\alpha_j, j = 1, \ldots, J\}$ and run for each j the censoring-adjusted τ-quantile regression of $b_i - \iota_i - \alpha_j k_i$ on the set of included instruments $(1, x_0, x_i)$ and the instrument for K_i. Let $\hat{\alpha}(\alpha_j, \tau)$ be the coefficients on $(1, x_0, x_i)$ and $\hat{\gamma}(\alpha_j, \tau)$ be the coefficient on the instrument. I choose $\hat{\alpha}_2(\tau)$ as the value among $\{\alpha_j, j = 1, \ldots, J\}$ that sets $\hat{\gamma}(\alpha_j, \tau) = 0$. The estimator for the remaining coefficients is given by $\hat{\alpha}(\hat{\alpha}_2(\tau), \tau)$. I calculate standard errors for the coefficients of interest by bootstrap. I run the same multistep procedure for 200 bootstrap samples of 1,925 auctions each and report the standard deviation of each parameter across samples.

This procedure relies on a first-step estimate of H. The rest of this section follows the ideas in Guerre, Perrigne, and Vuong (2000) and Campo, Perrigne, and Vuong (2003) of estimating the distribution of bids flexibly in a first step. In particular, I describe how to estimate the joint distribution of signals and the marginal distribution of bids to compute $\pi_i(\tau, x_t|\hat{H})$, $\iota_{it} = \iota_i(b_{it}, x_t|\hat{H})$, and $k_{it} = k_i(b_{it}, x_t|\hat{H})$ for each observed bid.

1. Entry Probability and Joint Distribution of Signals

Monotonicity of the bid function implies that the copula of bids conditional on the vector of cost shifters is equal to the joint distribution of signals. The copula is estimated semiparametrically (Hubbard, Li, and Paarsch 2012). In the first stage, the univariate conditional marginals are estimated flexibly. In the second stage, the copula parameters are estimated by likelihood methods.

I estimate the univariate conditional marginals using splines. Instead of conditioning on the full vector of competitors' distances, I condition on the geographical coordinates (lat, lon) of the project location. These coordinates fully determine the vector of distances, as firms' plants locations are fixed. The marginal distribution of $B_i|\text{lat, lon}$ is estimated as

$$\hat{F}_{B_i|\text{lat,lon}}(\text{bid}) = \sum_{j=1}^{K_{\text{lat}}}\sum_{k=1}^{K_{\text{lon}}}\sum_{b=1}^{K_{\text{bid}}} \hat{c}_{jkb}\theta_j(\text{lat})\theta_k(\text{lon})\theta_b(\text{bid}), \tag{15}$$

where θ_j, θ_k, and θ_b are cubic B-splines and the coefficients $\hat{c}_{jkb}$ solve the following problem:

$$\min_{c\in[0,1]^{K_{\text{lat}}\times K_{\text{lon}}\times K_{\text{bid}}}} \sum_{t=1}^{T}\left(1(b_{it} \le \text{bid}) - \textstyle\sum_{j=1}^{K_{\text{lat}}}\sum_{k=1}^{K_{\text{lon}}}\sum_{b=1}^{K_{\text{bid}}} c_{jkb}\theta_j(\text{lat}_t)\theta_k(\text{lon}_t)\theta_b(\text{bid})\right)^2.$$

The restriction that $c \in [0,1]^{K_{\text{lat}}\times K_{\text{lon}}\times K_{\text{bid}}}$ ensures that the predicted probabilities lie between zero and one. The set of splines in each dimension are chosen to cover the full range of locations observed in the data. I use $K_{\text{lat}} = K_{\text{lon}} = 15$ in the sample, which results in one knot every 87 km from east to west and one knot every 91 km from south to north. For K_{bid} I created 1,000 equally spaced knots between the minimum and maximum (finite) normalized bids observed in the data.

For each bidder i and auction t, I construct $\hat{s}_{it} = \hat{F}_{B_i|\text{lat,lon}}(b_{it})$. This is an estimate of the signal that prompted bid b_{it}. If bidder i does not participate in auction t, I obtain $\hat{F}_{B_i|\text{lat,lon}}(\bar{b})$, where $\bar{b}$ is the maximum finite normalized bid observed in the data. This is an estimate of $\phi_i(x|H)$, the threshold signal where bidder i is indifferent between participating and staying out. Nonparticipation implies that the realized signal was above the threshold. The sample analog to equation (13) is

$$\pi_i(\tau, x_t|\hat{H}) = \min\left(\frac{\tau}{\hat{F}_{B_i|\text{lat,lon}}(\bar{b})}, 1\right).$$

I construct a new data set consisting of T realizations of a vector of signals that may be truncated because of nonparticipation. The likelihood of the observed data can be written explicitly for a correlation matrix with a factor structure: $\Sigma = LL' + \Lambda$. Recall that L is an $n \times l$ loading matrix that can be normalized so that $L_{jk} = 0$ if $j < k$. Here, Λ is a positive diagonal matrix that ensures that the elements on the main diagonal of Σ are all one. This formulation reduces the number of free parameters in the correlation matrix Σ from $n(n-1)/2$ to $nl - l(l-1)/2$. I estimate the parameters L using a simulated maximum likelihood estimator proposed by Kamakura and Wedel (2001).

2. Markup Due to Market Power

Following Guerre, Perrigne, and Vuong (2000) and Campo, Perrigne, and Vuong (2003), one can estimate the marginal cost that rationalizes each observed bid as the left-hand side of equation (5). This expression depends on the probability of winning and its derivative. The probability of winning with the observed bid b is estimated as

$$\hat{H}(M_i \geq b | b, \text{lat}, \text{lon}) = \sum_{j=1}^{K_{\text{lat}}} \sum_{k=1}^{K_{\text{lon}}} \sum_{m=1}^{K_m} \hat{c}_{jkm} \theta_j(\text{lat}) \theta_k(\text{lon}) \theta_m(b),$$

where θ_j, θ_k, and θ_m are the cubic B-splines and the coefficients $\hat{c}_{jkm}$ solve the following problem:

$$\min_{c \in [0,1]^{K_{\text{lat}} \times K_{\text{lon}} \times K_m}} \sum_{t=1}^{T} \left(1(m_{it} \geq b_{it}) - \sum_{j=1}^{K_{\text{lat}}} \sum_{k=1}^{K_{\text{lon}}} \sum_{m=1}^{K_m} c_{jkm} \theta_j(\text{lat}_t) \theta_k(\text{lon}_t) \theta_m(b_{it}) \right)^2.$$

The restriction on c ensures that the fitted probabilities lie between zero and one. The coefficients $\hat{c}_{jkm}$ provide the best linear predictor of the event that bidder i wins the auction with bid b conditional on (b, lat, lon) subject to these constraints. The derivative $\hat{h}(M_i \geq b | b, \text{lat}, \text{lon})$ is estimated in a similar way, except that the $\hat{c}'$ coefficients solve

$$\min_{c \in \mathbb{R}_+^{K_{\text{lat}} \times K_{\text{lon}} \times K_m}} \sum_{t=1}^{T} \left(\frac{1}{h} K\left(\frac{m_{it} - b_{it}}{h} \right) - \sum_{j=1}^{K_{\text{lat}}} \sum_{k=1}^{K_{\text{lon}}} \sum_{m=1}^{K_m} c_{jkm} \theta_j(\text{lat}_t) \theta_k(\text{lon}_t) \theta_m(b_{it}) \right)^2,$$

where K is a kernel function and h is a bandwidth. Instead of fitting whether bidder i won the auction with bid b_i, the coefficients fit whether the minimum competitors' bid was sufficiently close to b_i. The constraints on c ensure that the fitted density is positive. The set of splines for each geographical dimension is identical to those in the previous step. I use five splines for bids that are chosen so that the knots partition the support of bids into five segments of the same length. The locations of the knots vary across bidders, depending on the support of their bids. I use a quartic kernel and a bandwidth equal to 0.1, that is, 10% of the engineer's estimate; I follow a leave-one-out cross-validation procedure with a log-likelihood criterion to choose the bandwidth.

The estimate of the marginal cost that rationalizes each observed bid is the bid minus the markup due to market power:

$$m\hat{c}_{it} = b_{it} - \iota_i\left(b_{it}, x_t|\hat{H}\right).$$

The markup due to market power can be computed as

$$\iota_i\left(b_{it}, x_t|\hat{H}\right) = \frac{\sum_{j=1}^{K_{\text{lat}}}\sum_{k=1}^{K_{\text{lon}}}\sum_{m=1}^{K_m}\hat{c}_{jkm}\theta_j(\text{lat})\theta_k(\text{lon})\theta_m(b_{it})}{\max\left(\sum_{j=1}^{K_{\text{lat}}}\sum_{k=1}^{K_{\text{lon}}}\sum_{m=1}^{K_m}\hat{c}'_{jkm}\theta_j(\text{lat})\theta_k(\text{lon})\theta_m(b_{it}), \zeta\right)},$$

where ζ is a vanishing positive constant that corrects for estimation error in the denominator that could result in very small densities.

3. Markup Due to Winner's Curse

I estimate the unconditional probability that j ties for the lowest bid with i:

$$\tilde{\pi}_{jt}^{(i)} = \sum_{j=1}^{K_{\text{lat}}}\sum_{k=1}^{K_{\text{lon}}}\sum_{m=1}^{K_m}\hat{c}_{jkm}\theta_j(\text{lat}_t)\theta_k(\text{lon}_t)\theta_m(b_{it}),$$

where the coefficients $\hat{c}_{jkm}$ solve

$$\min_{c\in\mathbb{R}_+^{K_{\text{lat}}\times K_{\text{lon}}\times K_m}}\sum_{t=1}^{T}\left(\frac{1}{h}K\left(\frac{b_{jt}-b_{it}}{h}\right)\prod_{k\neq i,j}1(b_{kt}\geq b_{it}) - \sum_{j=1}^{K_{\text{lat}}}\sum_{k=1}^{K_{\text{lon}}}\sum_{m=1}^{K_m}c_{jkm}\theta_j(\text{lat}_t)\theta_k(\text{lon}_t)\theta_m(b_{it})\right)^2.$$

The coefficients fit whether competitor j bids sufficiently close to b_i and all other competitors bid strictly above. The conditional probability that bidder i ties with bidder j conditional on tying with at least one competitor, $\pi_{jt}^{(i)}$, is estimated as

$$\hat{\pi}_{jt}^{(i)} = \frac{\tilde{\pi}_{jt}^{(i)}}{\sum_{k\neq i}\tilde{\pi}_{kt}^{(i)}}.$$

The markup due to the winner's curse can be computed by

$$k_{it}\left(b_{it}, x_t|\hat{H}\right) = \sum_{j\neq i}\tilde{\sigma}_{ij}\left(\left[\pi_{jt}^{(i)}z_{jt}^{(i)}\right] + \sum_{k\neq i,j}\pi_{kt}^{(i)}E\left(Z_j|z_{kt}^{(i)}, z_{it}^{(i)}, \left\{Z_m\geq z_{mt}^{(i)}\right\}_{m\neq i,k}\right)\right).$$

The previous step obtained estimates of $\pi_{jt}^{(i)}$. I estimate $z_{jt}^{(i)}$ by $\hat{z}_{jt}^{(i)} = \Phi^{-1}(\hat{F}_{B_j|\text{lat}_t,\text{lon}_t}(b_{it}))$, where $\hat{F}_{B_j|\text{lat}_t,\text{lon}_t}(\cdot)$ was estimated by equation (15). The expectation of Z_j conditional on $z_{kt}^{(i)}$, $z_{it}^{(i)}$, and $\{Z_m\geq z_{mt}^{(i)}\}_{m\neq i,k}$ can be calculated numerically, given the estimated joint distribution of signals. The procedure in section IV.B.1 yields estimates of $\tilde{\sigma}_{ij}$.

V. Results

I focus on the 19 firms that were observed participating with higher frequency in the final sample. The rest of the firms are grouped together as a fringe twentieth bidder, and I keep only the lowest of their bids. The

TABLE 3
CORRELATION OF SIGNALS

	Factor 1		Factor 2		Signal Decomposition		
	Coefficient	Standard Error	Coefficient	Standard Error	Factor 1	Factor 2	Individual
Bidder 1	.410***	.025	.173***	.054	.17	.03	.80
Bidder 2	.449***	.020	.125**	.061	.20	.02	.78
Bidder 3	.434***	.035	.218***	.055	.19	.05	.76
Bidder 4	−.356***	.057	.322***	.074	.13	.10	.77
Bidder 5	.346***	.036	.291***	.046	.12	.08	.80
Bidder 6	.221**	.103	.428***	.090	.05	.18	.77
Bidder 7	.140	.101	.458***	.074	.02	.21	.77
Bidder 8	.071	.121	−.328***	.126	.01	.11	.89
Bidder 9	.370***	.072	−.021	.138	.14	.00	.86
Bidder 10	.022	.082	.455***	.033	.00	.21	.79
Bidder 11	.382***	.073	.215*	.128	.15	.05	.81
Bidder 12	.037	.176	.099	.163	.00	.01	.99
Bidder 13	.351***	.060	.107	.116	.12	.01	.87
Bidder 14	.298***	.077	.173*	.103	.09	.03	.88
Bidder 15	.363***	.058	.204**	.087	.13	.04	.83
Bidder 16	.126	.210	.294	.241	.02	.09	.90
Bidder 17	.225***	.072	.026	.090	.05	.00	.95
Bidder 18	−.444***	.086	.279**	.136	.20	.08	.73
Bidder 19	−.075	.171	.475***	.103	.01	.23	.77
Bidder 20			.504***	.018		.25	.75

NOTE.—Loading matrix estimates obtained by simulated maximum likelihood, following Kamakura and Wedel (2001). Standard errors are obtained from 200 bootstrap samples. Each bidder's signal can be written as $Z_i = \hat{L}_{i1}\text{Factor}_1 + \hat{L}_{i2}\text{Factor}_2 + \text{Individual}_i$, where Factor_j, $j = 1, 2$, and Individual_i are jointly independent normal random variables with variances 1, 1, and $1 - \hat{L}_{i1}^2 - \hat{L}_{i2}^2$, respectively. The total variance of Z_i can be decomposed into three terms: $\hat{L}_{i1}^2$, $\hat{L}_{i2}^2$, and $1 - \hat{L}_{i1}^2 - \hat{L}_{i2}^2$.
* $p < .1$.
** $p < .05$.
*** $p < .01$.

results below show the firm-specific parameter estimates for the distribution of signals and full-information costs.

A. *Correlation of Signals*

I estimated the correlation structure for $l = \{0, 1, 2, 3\}$ factors and used the Akaike information criterion (AIC) to select $l = 2$.[12] Table 3 shows the estimates of the parameters of the loading matrix L and the elements on the diagonal of Λ. It also shows a decomposition of the total signal variance in the l factors and in the bidder-specific component. The first factor

[12] The AIC selects $l = 2$: $\text{AIC}_1 - \text{AIC}_2 = 120$, and $\text{AIC}_3 - \text{AIC}_2 = 6.5$. The Bayesian information criterion (BIC) is inconclusive between $l = 1$ and $l = 2$ but strongly rejects $l = 3$: $\text{BIC}_1 - \text{BIC}_2 = -0.28$, and $\text{BIC}_3 - \text{BIC}_2 = 121$. I added region-specific factors to test the null hypothesis of independence of signals and cost shifters. I can reject the null at the 10% level when $l = 1$ but not when $l = 2$.

introduces correlation in the signals of bidders 1–3, 5, 6, 9, 11, 13–15, and 17. The implied correlation among their signals ranges between 0.11 and 0.29 (see table 4). This factor explains up to 20% of the signal variance for bidders in this group. It also introduces negative correlation between the signals of this group and those of bidders 4 and 18. The second factor introduces correlation between the signals of bidders 3–7, 10, 11, 15, 18, 19, and the fringe bidder. This factor explains up to 25% of the signal variance for bidders in this group. Many of the firms in this group also participate in construction projects (which are excluded from the sample), and they may have different technology and capabilities. In particular, firms 4 and 10 do not own asphalt plants. In paving projects, they have to buy asphalt from other firms.

B. Full-Information Costs

Tables 5 and 6 show the estimates of the full-information cost for bidders 1–12 and 13–19, respectively. The included exogenous covariates are a set of dummies for projects farther than 10, 50, and 100 km, a measure of road density in the vicinity of the project that allows for potential differential costs in highly urbanized areas,[13] and three constants: one for the 25th percentile of the distribution of costs, one for the difference between the 50th and the 25th percentiles, and one for the difference between the 75th and the 25th percentiles.

The coefficient labeled "common costs" in tables 5 and 6 corresponds to the parameter α_{i2}. To estimate this parameter, I use the vector of competitors' distances as instruments for the common cost k_{it}. Chernozhukov and Hansen (2008) suggest collapsing the vector into a single instrument. This is achieved by projecting the endogenous variable on a space of piecewise linear functions of the vector of distances with kinks at 10, 50, and 100 km and using the fitted value as the instrument for the instrumental variable quantile regression.

The estimate of the full-information cost for bidder 1 evaluated at the 25th percentile own signal and the 50th percentile competitors' signal for projects within 10 km of its plants is 73% of the engineer's estimate. For projects in the ranges 10–50 and 50–100 km, this cost increases by 2.5 and 15.7 percentage points (pp), respectively. Bidder 1 does not seem

[13] I used geographical information on all roads available at the Michigan Center for Geographical Information (https://gis-michigan.opendata.arcgis.com/datasets/all-roads-v17a). I computed each road segment's length and assigned a weight based on its classification (interstates, 3.5; freeways, 3; principal arterials, 2.5; minor arterials, 2.2; major collectors, 2; minor collectors, 1.5; and local roads, 1). I computed the density of roads using a 10-km bandwidth kernel and evaluated it at each project location. This measure was subsequently standardized. Rural areas have a standardized measure of approximately −1, sparse areas such as the city of Holland have 0, denser areas such as Grand Rapids have 1, and Detroit reaches 3.5.

TABLE 4
Correlation of Signals

	Bidders																		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
Bidder 2	.259***																		
Bidder 3	.275***	.287***																	
Bidder 4	−.115**	−.154***	−.110**																
Bidder 5	.240***	.243***	.274***	−.038															
Bidder 6	.210***	.197***	.247***	.077	.257***														
Bidder 7	.174***	.155***	.209***	.126	.232***	.295***													
Bidder 8	−.033	−.011	−.050	−.158*	−.084	−.151*	−.170**												
Bidder 9	.178***	.199***	.192***	−.170**	.147**	.089	.052	.038											
Bidder 10	.110**	.085	.140***	.178***	.176***	.256***	.271***	−.177**	−.002										
Bidder 11	.241***	.250***	.271***	−.085	.243***	.224***	.193***	−.052	.164***	.133									
Bidder 12	.036	.033	.043	.021	.047	.058	.058	−.032	.013	.052	.040								
Bidder 13	.195***	.208***	.216***	−.111	.184***	.151**	.120*	−.012	.148***	.068	.188***	.026							
Bidder 14	.181***	.187***	.204***	−.061	.183***	.170***	.147**	−.041	.122**	.102	.179***	.030	.141***						
Bidder 15	.226***	.235***	.255***	−.080	.228***	.210***	.181***	−.048	.154***	.125**	.224***	.037	.177***	.169***					
Bidder 16	.121	.112	.144	.060	.153	.185	.183	−.098	.046	.162	.131	.036	.086	.100	.123				
Bidder 17	.111***	.121***	.122***	−.084*	.099**	.072	.051	.008	.092**	.020	.105**	.011	.091***	.079**	.099**	.039			
Bidder 18	−.175***	−.218***	−.177**	.332***	−.095	.029	.088	−.153	−.215***	.155	−.143	.013	−.159**	−.105	−.135*	.032	−.112*		
Bidder 19	.065	.033	.092	.233*	.143**	.243***	.269***	−.195**	−.047	.275***	.093	.051	.030	.073	.087	.157	−.005	.222*	
Bidder 20	.113***	.083*	.145***	.214***	.190***	.285***	.304***	−.203**	−.013	.298***	.140	.058	.067	.108	.131**	.181	.016	.191*	.316***

Note.—Correlation matrix derived from the factor structure shown in table 3. Diagonal elements are all equal to 1.

* $p < .1$.
** $p < .05$.
*** $p < .01$.

to be more efficient in denser areas. The 50th (75th) percentile cost is 8.2 (18.1) pp higher than the 25th. Table 7 presents the estimated effect of own and competitors' signals on the full-information cost. All estimates are positive, which is consistent with the testable predictions of the interdependent-cost model. It is interesting to compare the effect of own and competitors' signals. A 1 standard deviation increase in the Gaussian signal received by bidder 1 increases its full-information cost by 10 pp. The effect of the same increase in a competitor's signal depends on the identity of the competitor. The average effect over competitors is 4 pp, and the maximum is 7.

The results for other bidders are qualitatively similar, but there are a couple of noteworthy differences that justify estimating these cost functions separately. While the private-cost hypothesis is rejected for bidders 1 and 2, it cannot be rejected for bidder 3. The coefficient on the common cost for bidder 3 is precisely estimated around zero. The same can be said about bidder 9. Another important difference is that while bidders 3, 5, and 10 seem to be more efficient in road-dense areas, bidder 9 seems to be more efficient in less dense areas. Because bidders 13–19 are not observed participating as regularly, their parameters are estimated less precisely. Nonetheless, the point estimates are comparable to those of firms that are observed participating more often.

C. *The Effect of Competition*

I consider cases where participation is restricted to a set of invited bidders. Inviting more bidders will always result in more competitive bidding and lower procurement costs in independent-private-cost auctions (Bulow and Klemperer 1996), but the results may be different in affiliated-private-cost or interdependent-cost models. Pinkse and Tan (2005) showed that affiliation of signals can offset the procompetitive effect of an additional bidder. This affiliation effect arises because bidders realize that winning in the presence of one more bidder implies that it is more likely that rivals' costs are high (because of affiliation), and that they can profitably increase their markups. Interdependent- or common-costs models have the additional anticompetitive effect of adverse selection, or the winner's curse (Matthews 1984; Hong and Shum 2002). Each bidder realizes that it wins in an adversely selected set of states of the world: when competitors have bad signals. Their expected completion cost conditional on winning is thus higher than the unconditional one. Bidding against an additional bidder worsens the selection problem, so bidders may react by increasing their bids to account for higher expected costs.

To evaluate the effects of increased competition on bidding behavior and procurement costs, I simulate the effects of a policy that restricts bidder

TABLE 5
FULL-INFORMATION COST ESTIMATES

Variable	Coefficient (1)	Standard Error (2)	Coefficient (3)	Standard Error (4)	Coefficient (5)	Standard Error (6)	Coefficient (7)	Standard Error (8)
	Bidder 1		Bidder 2		Bidder 3		Bidder 4	
Const p25	.730***	.031	.782***	.020	.860***	.025	.537*	.287
Dist > 10	.025**	.011	.038*	.020	.068***	.023	.162	.126
Dist > 50	.157***	.024	.105***	.022	.399***	.111	.111	.167
Dist > 100	.363**	.166	.240***	.058	−.021	.132	−.110	.174
Road density	−.008	.016	.010	.033	−.032***	.007	.101	.321
Const p50 − p25	.082***	.020	.054***	.017	.101***	.031	.205	.136
Const p75 − p25	.181***	.053	.126***	.032	.238***	.067	.910*	.502
Common costs	.500***	.190	.350***	.118	−.025	.092	.000	.161
Observations	1,096		936		411		72	
	Bidder 5		Bidder 6		Bidder 7		Bidder 8	
Const p25	.957***	.021	.927***	.124	.836***	.134	.763***	.118
Dist > 10	.016	.017	.085**	.038	.020	.042	.051	.049
Dist > 50	−.045	.160	.127***	.029	.034	.035	.067	.086
Dist > 100			.422	.309	.056	.051		
Road density	−.021**	.009	.265	.172	−.042	.164	.022	.051
Const p50 − p25	.082***	.019	.101	.074	.051*	.029	.169**	.066
Const p75 − p25	.186***	.034	.243	.151	.155***	.047	.422***	.140
Common costs	.175***	.051	.800*	.441	.525***	.147	.950	.700
Observations	399		374		249		280	

	Bidder 9		Bidder 10		Bidder 11		Bidder 12	
Const p25	.919***	.047	.940***	.065	.799***	.080	.833***	.038
Dist > 10	.096**	.045	.114**	.057	−.013	.029	−.003	.035
Dist > 50	.207***	.044	−.001	.046	.061	.055	.052	.107
Dist > 100	.142	.106	.421*	.217	−.012	.185		
Road density	.288***	.079	−.034**	.016	−.229*	.131	−.008	.086
Const p50 − p25	.111***	.025	.047***	.018	.070***	.025	.054	.041
Const p75 − p25	.278***	.061	.145	.095	.134***	.049	.154	.101
Common costs	.050	.056	.325***	.097	.500***	.174	.525	.332
Observations	273		122		148		162	

Note.—Estimates of the full-information cost parameters for each bidder. The estimates were obtained by instrumental variable quantile regression using three different quantiles and restricting all coefficients except the constant (const) from varying across quantiles. The chosen quantiles were 0.25 (p25), 0.50 (p50), and 0.75 (p75). For bidders 4, 8, and 10, I used quantiles $0.25Q$, $0.5Q$, and $0.75Q$, where Q is the maximum estimated probability of participation in any given auction (0.26, 0.69, and 0.41, respectively). Standard errors were obtained by bootstrap. Dist = distance (in kilometers) to project.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

TABLE 6

Full-Information Cost Estimates

Variable	Coefficient (1)	Standard Error (2)	Coefficient (3)	Standard Error (4)	Coefficient (5)	Standard Error (6)	Coefficient (7)	Standard Error (8)
	Bidder 13		Bidder 14		Bidder 15		Bidder 16	
Const p25	1.146***	.161	.971***	.091	.903***	.045	.981***	.109
Dist > 10	.010	.046	.022	.063	.004	.055	−.055	.064
Dist > 50	−.033	.051	.013	.106	.393***	.086	.035	.069
Dist > 100	.068	.228	.176	.209			.192	.295
Road density	.617***	.237	−.022	.036	−.130***	.038	−.066	.077
Const p50 − p25	.063	.038	.126**	.050	.094***	.028	.036	.040
Const p75 − p25	.127	.092	.279*	.168	.237***	.059	.101	.202
Common costs	.250**	.122	.075	.107	.025	.112	.450*	.246
Q	1.00		.60		.64		.44	
Observations	151		85		129		60	
	Bidder 17		Bidder 18		Bidder 19			
Const p25	.759***	.074	.356	.438	.885***	.072		
Dist > 10	.055	.038	.025	.145	.078	.066		
Dist > 50	.052	.077	.049	.112	.218**	.096		
Dist > 100			−.080	.131	−.123	.107		
Road density	.028	.142	−.513	.542	.019	.102		
Const p50 − p25	.077***	.025	.091*	.055	.060	.042		
Const p75 − p25	.137**	.069	.197*	.117	.225	.145		
Common costs	.400**	.193	.150	.129	.300**	.142		
Q	.96		.37		.34			
Observations	106		61		62			

Note.—Estimates of the full information cost parameters for each bidder. The estimates were obtained by instrumental variable quantile regression using three different quantiles and restricting all coefficients except the constant (const) from varying across quantiles. The chosen quantiles were 0.25 (p25), 0.50 (p50), and 0.75 (p75). For bidders 14–19, I used quantiles $0.25Q$, $0.5Q$, and $0.75Q$, where Q is the maximum estimated probability of participation in any given auction. Standard errors were obtained by bootstrap. Dist = distance in kilometers to project.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

TABLE 7
PRIVATE- AND COMMON-COST COMPONENTS

	Own Signal		Competitor (mean)		Competitor (maximum)	
	Coefficient	Standard Error	Coefficient	Standard Error	Coefficient	Standard Error
Bidder 1	.10***	.02	.04***	.01	.07***	.02
Bidder 2	.09***	.02	.03***	.01	.06***	.02
Bidder 3	.13***	.05	.00	.00	.00	.01
Bidder 4	.59*	.32	.00	1.15	.00	3.95
Bidder 5	.13***	.02	.01***	.00	.02***	.01
Bidder 6	.15	.11	.04*	.02	.09**	.04
Bidder 7	.12***	.04	.03***	.01	.07***	.02
Bidder 8	.32***	.12	.05	.05	.14	.15
Bidder 9	.18***	.04	.01	.03	.02	.10
Bidder 10	.10***	.04	.02***	.00	.05***	.02
Bidder 11	.10***	.03	.03***	.01	.06***	.02
Bidder 12	.07	.06	.03	.04	.06	.10
Bidder 13	.09*	.05	.02**	.01	.04	.03
Bidder 14	.25***	.09	.00	.01	.01	.02
Bidder 15	.20***	.05	.00	.01	.00	.02
Bidder 16	.07	.08	.02	.02	.05	.06
Bidder 17	.12***	.04	.04	.12	.09	.33
Bidder 18	.22*	.11	.07	.15	.20	.41
Bidder 19	.15	.10	.02**	.01	.06**	.02

NOTE.—Effect of a 1 standard deviation increase of own and competitors' Gaussian costs on each bidder's full-information cost. The effect of an increase in competitor's signal depends on its identity. The table reports mean and maximum effects of competitors' signals.
* $p < .1$.
** $p < .05$.
*** $p < .01$.

participation in 250 randomly drawn auctions from the 1,925 in the sample. I order bidders according to their estimated median full-information cost in each auction. I assume that only the first N bidders with the lowest median cost will be allowed to participate. The set of invitees will vary across auctions because bidders have different estimated costs, depending on the location of the project. Three different set of primitives/models are considered: the common-cost (CC) model, which uses the primitives estimated in the previous section; the affiliated-private-cost (APC) model, which sets all CC coefficients α_2 to zero; and the independent-private-cost (IPC) model, which also sets the parameters of the loading matrix to zero.[14] For each model and each $N = \{2, 3, \ldots, 10\}$, I compute a Bayes-Nash equilibrium, using a numerical algorithm described in appendix D, and simulate auction outcomes, drawing 250,000 signal realizations from the estimated joint distribution (in the APC or the CC models) or from independent distributions (in the IPC model).

[14] A more precise name for the APC model is "correlated private costs." The estimated joint distribution of signals has some negative correlations, which implies that the signals are correlated but not affiliated.

Figure 3 illustrates the effect of competition under the three models. To construct this figure, I averaged the transaction price across the 250,000 signal realizations and averaged over the 250 auctions, weighting by the engineer's estimate. I normalized this average by the duopoly outcome for each model. The figure plots the resulting average transaction price (relative to the duopoly) for each model and number of invitees. The three models predict that inviting an additional bidder reduces procurement costs. The IPC model has the largest effect, followed by the APC model. The CC model has the smallest effect. For example, increasing the number of invitees from two to seven reduces costs by 5.48 pp in the IPC model, 5.00 pp in the APC model, and 4.40 pp in the CC model. In all three models, the effect of inviting a ninth or a tenth bidder is negligible, because this bidder is typically quite inefficient.

Figure 4 shows the effect on the bid of the bidder with the lowest median full-information cost. While both IPC and APC models predict that this bidder will bid more aggressively when faced with more competition, the CC model predicts the opposite. This bidder reacts to increased competition by bidding higher as it faces a more severe winner's curse.

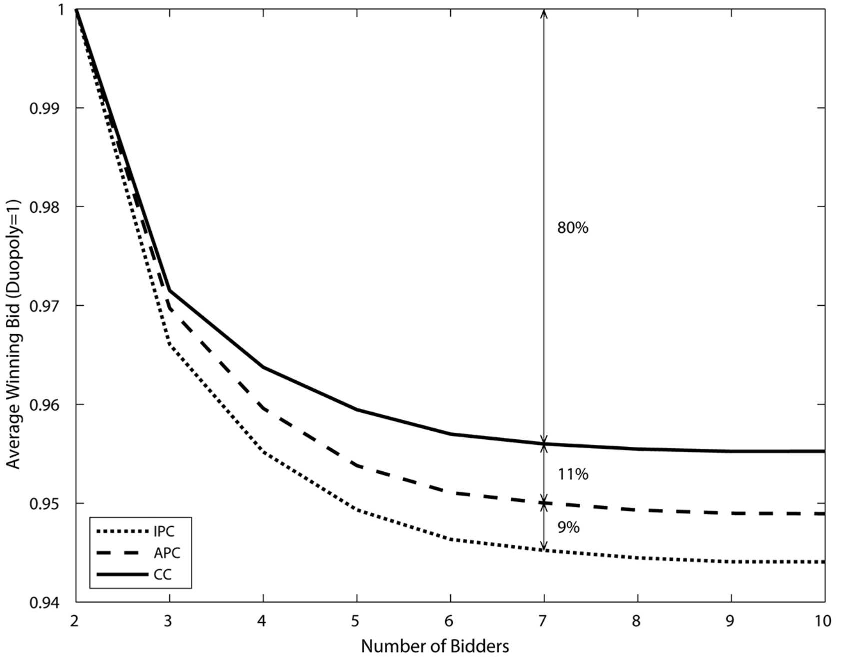

FIG. 3.—Effects of competition on procurement costs: procurement costs relative to the duopoly case in the independent-private-costs (IPC), affiliated-private-costs (APC), and common-costs (CC) models, average over 250 randomly selected auctions.

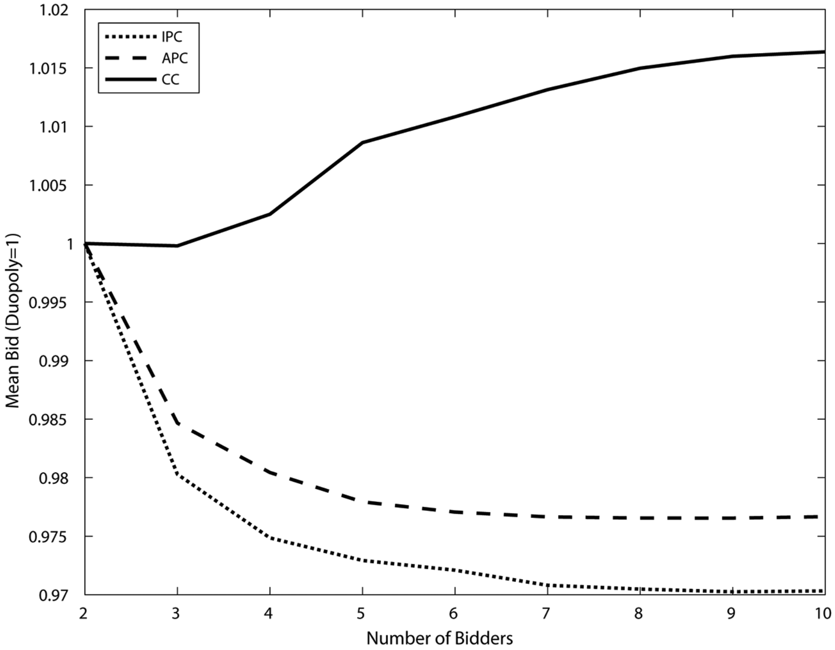

FIG. 4.—Effects of competition on bidding behavior: median bid relative to the duopoly case in the independent-private-costs (IPC), affiliated-private-costs (APC), and common-costs (CC) models, average over 250 randomly selected auctions.

To illustrate the magnitudes of the different effects at play, I decompose the total effect of competition on the CC model in four components. The *competitive effect* is the reduction in the average bid in the IPC model of the first invitee, that is, the bidder with the lowest median cost. The *affiliation effect* is the difference in this average bid reduction between the APC and the IPC models. The *winner's curse effect* is the difference in this average bid reduction between the CC and the APC models. The *sampling effect* is the difference between the reduction in the average bid of the first invitee and the reduction in procurement costs in the CC model.

Table 8 shows the results of this decomposition. The top panel compares the outcomes of inviting $N = 3, 4, \dots, 10$ bidders relative to inviting only two. The competitive effect has cost-saving effects that are increasing in the number of invitees. They range from 1.98 to 2.97 pp. Affiliation has an offsetting effect of 0.45–0.64 pp. The winner's curse has a stronger offsetting effect of 1.52–3.96 pp. These three effects combined make the overall response of the first invitee increasing in the number of bidders for all $N > 3$. The sampling effect has a strong cost-saving effect of 2.83–6.10 pp that ensures a total reduction in procurement costs

TABLE 8
Effects of Competition on Bidding Behavior

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Effects							
	3 vs. 2	4 vs. 2	5 vs. 2	6 vs. 2	7 vs. 2	8 vs. 2	9 vs. 2	10 vs. 2
Competitive	1.98	2.52	2.71	2.79	2.92	2.95	2.97	2.97
Affiliation	−.45	−.57	−.51	−.50	−.59	−.61	−.63	−.64
Winner's curse	−1.52	−2.21	−3.06	−3.37	−3.65	−3.84	−3.95	−3.96
Sampling	2.83	3.88	4.91	5.38	5.71	5.95	6.08	6.10
Total	2.84	3.63	4.05	4.30	4.40	4.45	4.47	4.47
	Incremental Effects							
	3 vs. 2	4 vs. 3	5 vs. 4	6 vs. 5	7 vs. 6	8 vs. 7	9 vs. 8	10 vs. 9
Competitive	1.98	.56	.19	.09	.14	.03	.02	.00
Affiliation	−.45	−.12	.06	.00	−.09	−.02	−.02	−.01
Winner's curse	−1.52	−.71	−.89	−.32	−.29	−.20	−.11	−.02
Sampling	2.83	1.08	1.08	.49	.35	.24	.13	.03
Total	2.84	.81	.44	.26	.10	.06	.03	.00

Note.—Decomposition of the cost-saving effect of inviting more bidders. The top panel compares inviting $N = 3, \ldots, 10$ bidders relative to inviting only two bidders. The bottom panel compares inviting $N = 3, \ldots, 10$ bidders relative to inviting $N - 1$ bidder. All units are normalized so that 100 equals the average procurement costs in the baseline case (2 and $N - 1$ bidders in the top and bottom panels, respectively).

between 2.84 and 4.47 pp. The bottom panel compares the outcomes of inviting N bidders relative to inviting $N - 1$. It shows that the magnitude of these effects decays rapidly as more and less efficient bidders are invited.

D. *Reserve Prices*

The existence of cost interdependencies also has an effect on the probability of successfully awarding a contract as a function of the reserve price. In private-values models, this is one minus the probability that all bidders' private costs are greater than the reserve. In interdependent-cost models, bidders also need to take into account the effect of the winner's curse. Figure 5 shows that the estimated cost interdependencies can reduce the success probability by as much as 0.1 when the reserve price is about 90% of the engineer's estimate. The figure also shows that reserve prices above 110% of the estimate rarely bind. They bind with probability 0.032 for the estimated affiliated-CC model. By federal rules, the MDOT has to justify in writing any project procured at a cost higher than that value. This analysis, together with the empirical observation that only 11% of the projects exceed this limit, suggests that this requirement is not very stringent.

Table 9 presents the MDOT surplus and the probability of awarding the contract under different reserve prices. The MDOT surplus is defined as its willingness to pay minus the expected procurement cost conditional

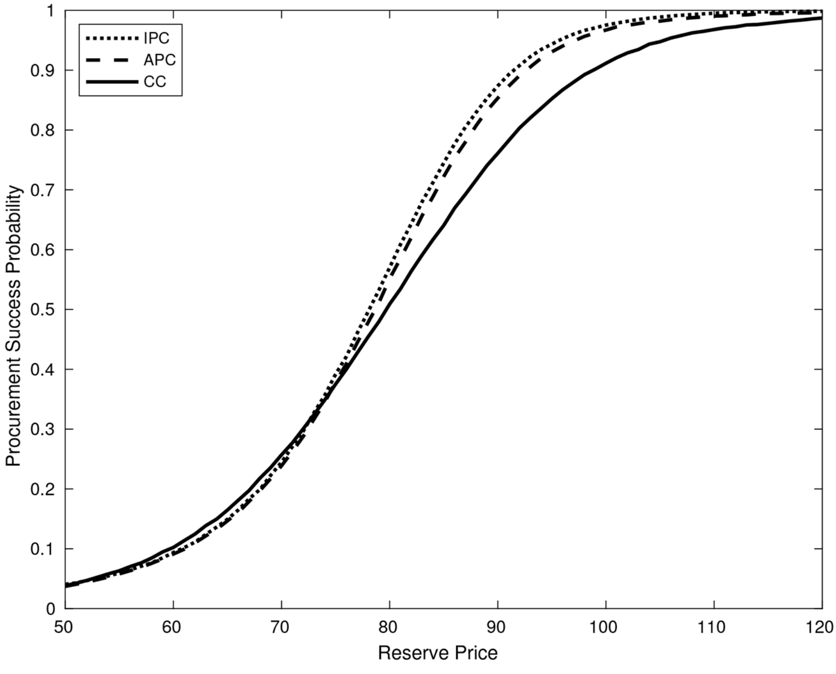

FIG. 5.—Reserve prices: probability of successfully procuring a project in the independent-private-costs (IPC), affiliated-private-costs (APC), and common-costs (CC) models. Reserve prices are expressed as percentage of the engineer's estimate; average over 250 randomly selected auctions.

on awarding the contract times the probability of awarding the contract. The second column shows the results under nonbinding reserve prices. By construction, the probability of awarding the contract is one. If the MDOT willingness to pay equals the engineer's estimate, setting a nonbinding reserve achieves a surplus of only 7% of the engineer's estimate. If the willingness to pay is 150% of the estimate, the surplus is 57%. In both cases, enforcing a 10% reserve would only marginally increase the MDOT surplus.

The MDOT can award only projects with positive interim expected net social benefits by setting the reserve price equal to its willingness to pay. In equilibrium, bidders will submit a bid only if their expected cost conditional on winning is below the reserve. Table 9 shows that setting the reserve price equal to the willingness to pay would result in a contract award with probability 0.912 and 0.997 when the willingness to pay is 100% and 150%, respectively. Because reserve prices affect the allocative efficiency of the mechanism, the socially optimal reserve price may be different from the MDOT's willingness to pay.

TABLE 9
Reserve Prices in the CC Model

	Nonbinding Reserve	Estimate + 10%	Efficient Reserve Price		Optimal Uniform Reserve			Optimal Auction-Specific Reserve			
WTP	Surplus	Surplus	Surplus	Pr	r (%)	Surplus (%)	Pr	Mean (r) (%)	SD(r) (%)	Surplus (%)	Pr
80%	−13	−11	1	.509	70	3	.254	70	2	3	.263
90%	−3	−1	4	.761	79	7	.478	79	3	7	.479
100%	7	8	10	.912	86	12	.675	86	3	12	.681
110%	17	18	18	.968	93	20	.814	93	4	20	.822
120%	27	28	27	.987	97	28	.884	98	5	29	.899
130%	37	37	37	.994	102	37	.930	103	6	38	.943
140%	47	47	47	.996	108	47	.963	107	7	47	.966
150%	57	57	57	.997	125	57	.992	110	8	57	.978

Note.—Willingness to pay for the project (WTP), procurement surplus (Surplus), and reserve prices (r) are expressed as a percentage of the engineer estimates. The contract award probability (Pr) is the probability that at least one bidder submits a bid below the reserve price. SD(r) = standard deviation of r.

The counterfactual simulations also show that if the MDOT had to set a uniform reserve price for every single procurement, it should set it equal to 86% (125%) of the engineer's estimate when its willingness to pay is equal to 100% (150%), achieving success with probability 0.67 (0.99). This policy results in a surplus of 12% (57%). In principle, the MDOT should be able to attain an even higher surplus if it sets an auction-specific reserve. The last set of columns shows that the magnitude of this improvement would be less than 1% of the engineer's estimate. The policy of always accepting bids below 110% of the engineer's estimate is consistent with a willingness to pay that exceeds 140% of the engineer's estimate. The results show that at these high valuations, the ability to set optimal reserve prices yields negligible gains relative to always setting a reserve price equal to the MDOT's willingness to pay.

VI. Conclusion

I provide a positive identification result for first-price procurements in the general interdependent-costs model. When bidders are risk neutral, the payoff-relevant primitives are the joint distribution of signals and each bidder's full-information costs. These primitives are sufficient to analyze the effects of most policy changes (e.g., rules of the auction, reserve prices, subsidies) on outcomes such as bidding behavior, project allocation, and prices. They may not be sufficient to analyze counterfactuals where the timing of the auction changes so that bidders are required to make decisions after some additional uncertainty in the model is resolved.

The identification result requires the following assumptions. First, each bidder's private information can be summarized by a real-valued signal. Second, the joint distribution of bidders' signals is independent of cost shifters. Third, each bidder's cost shifter affects his own full-information cost but not his competitors'. Fourth, the observed data are generated by the repeated play of the same equilibrium where bidders use monotone pure strategies.

The joint distribution of signals is identified from the joint distribution of bids because equilibrium bid functions are monotone in signals. The full-information cost is nonparametrically identified from variation in competitive conditions that generate enough variation in equilibrium bid functions. Intuitively, in models with strong interdependence, the effect of the winner's curse is important, and bidders may bid less aggressively in more competitive environments. The identification result exploits variation in bidder-specific cost shifters to vary the degree of competition that each bidder faces and identify cost interdependencies from their response to this variation.

I propose an estimation method that follows the indirect approaches proposed in the literature on estimation of private values and extend them to obtain estimates of the parameters that describe the primitives of the interdependent-cost model. I apply the estimator to bidding data from the MDOT and find that the estimated full-information cost is increasing with distance to the project and with competitors' signals. Despite being statistically significant, the effect of competitors' signals on expected costs is weaker than the effect of own signals, suggesting that the model is closer to pure private costs than to symmetric pure common costs. Policies that restrict the number of participants to ameliorate the winner's curse are successful at inducing more aggressive bidding among participants but fail at reducing overall procurement costs. In other words, the effect of the winner's curse is strong enough to make each bidder bid lower when there is less competition but not strong enough to compensate the procurer for the times when a bidder that would have won the auction was not allowed to participate. I also evaluate the effect of different reserve prices. The policy of always accepting bids below 110% of the engineer's estimate is consistent with a willingness to pay that exceeds 140% of the engineer's estimate. At these high valuations, the ability to set optimal reserve prices yields surplus gains of less than 1% of the engineer's estimate.

References

Aradillas-López, Andrés, Amit Gandhi, and Daniel Quint. 2013. "Identification and Inference in Ascending Auctions with Correlated Private Values." *Econometrica* 81 (2): 489–534.

Ashenfelter, Orley. 1989. "How Auctions Work for Wine and Art." *J. Econ. Perspectives* 3 (3): 23–36.

Asker, John. 2010. "A Study of the Internal Organization of a Bidding Cartel." *A.E.R.* 100 (3): 724–62.

Athey, Susan. 2001. "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information." *Econometrica* 69 (4): 861–89.

Athey, Susan, and Philip A. Haile. 2002. "Identification of Standard Auction Models." *Econometrica* 70 (6): 2107–40.

———. 2007. "Nonparametric Approaches to Auctions." In *Handbook of Econometrics*, vol. 6A, edited by James J. Heckman and Edward E. Leamer, 3847–965. Amsterdam: North-Holland.

Athey, Susan, and Jonathan Levin. 2001. "Information and Competition in U.S. Forest Service Timber Auctions." *J.P.E.* 109 (2): 375–417.

Athey, Susan, Jonathan Levin, and Enrique Seira. 2011. "Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions." *Q.J.E.* 126 (1): 207–57.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin. 2007. "Estimating Dynamic Models of Imperfect Competition." *Econometrica* 75 (5): 1331–70.

Bajari, Patrick, and Ali Hortaçsu. 2003. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND J. Econ.* 34 (2): 329–55.

Bajari, Patrick, Stephanie Houghton, and Steven Tadelis. 2014. "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs." *A.E.R.* 104 (4): 1288–319.

Balat, Jorge. 2017. "Highway Procurement and the Stimulus Package: Identification and Estimation of Dynamic Auctions with Unobserved Heterogeneity." Working paper, Yale Univ.

Baldwin, Laura H., Robert C. Marshall, and Jean-François Richard. 1997. "Bidder Collusion at Forest Service Timber Sales." *J.P.E.* 105 (4): 657–99.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2017. "First-Price Auctions with General Information Structures: Implications for Bidding and Revenue." *Econometrica* 85 (1): 107–43.

Berry, Steven T., and Philip A. Haile. 2014. "Identification in Differentiated Products Markets Using Market Level Data." *Econometrica* 82 (5): 1749–97.

Bhattacharya, Vivek. 2018. "An Empirical Model of R&D Procurement Contests: An Analysis of the DOD SBIR Program." Working paper, Dept. Econ., Northwestern Univ.

Bolotnyy, Valentin, and Shoshana Vasserman. 2019. "Scaling Auctions as Insurance: A Case Study in Infrastructure Procurement." Working paper, Harvard Univ.

Buchinsky, Moshe, and Jinyong Hahn. 1998. "An Alternative Estimator for the Censored Quantile Regression Model." *Econometrica* 66 (3): 653–71.

Bulow, Jeremy, and Paul Klemperer. 1996. "Auctions versus Negotiations." *A.E.R.* 86 (1): 180–94.

Campo, Sandra, Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. 2011. "Semiparametric Estimation of First-Price Auctions with Risk-Averse Bidders." *Rev. Econ. Studies* 78 (1): 112–47.

Campo, Sandra, Isabelle Perrigne, and Quang Vuong. 2003. "Asymmetry in First-Price Auctions with Affiliated Private Values." *J. Appl. Econometrics* 18 (2): 179–207.

Chen, Xiaohong, Oliver Linton, and Ingrid Van Keilegom. 2003. "Estimation of Semiparametric Models When the Criterion Function Is Not Smooth." *Econometrica* 71 (5): 1591–608.

Chernozhukov, Victor, and Christian Hansen. 2006. "Instrumental Quantile Regression Inference for Structural and Treatment Effect Models." *J. Econometrics* 132 (2): 491–525.

———. 2008. "Instrumental Variable Quantile Regression: A Robust Inference Approach." *J. Econometrics* 142 (1): 379–98.

Compiani, Giovanni, Philip A. Haile, and Marcelo Sant'Anna. 2020. "Common Values, Unobserved Heterogeneity, and Endogenous Entry in US Offshore Oil Lease Auctions." *J.P.E.* 128 (10): 3872–3912.

Cuesta, José Ignacio, and Alberto Sepúlveda. 2019. "Price Regulation in Credit Markets: A Trade-Off between Consumer Protection and Credit Access." Working paper. http://dx.doi.org/10.2139/ssrn.3282910.

Einav, Liran, and Ignacio Esponda. 2008. "Endogenous Participation and Local Market Power in Highway Procurement." Working paper, Stanford Univ. and New York Univ. Stern.

Février, Philippe. 2008. "Nonparametric Identification and Estimation of a Class of Common Value Auction Models." *J. Appl. Econometrics* 23 (7): 949–64.

Gentry, Matthew, and Tong Li. 2014. "Identification in Auctions with Selective Entry." *Econometrica* 82 (1): 315–44.

Guerre, Emmanuel, Isabelle M. Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica* 68 (3): 525–74.

———. 2009. “Nonparametric Identification of Risk Aversion in First-Price Auctions under Exclusion Restrictions.” *Econometrica* 77 (4): 1193–227.

Haile, Philip A. 2001. “Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales.” *A.E.R.* 91 (3): 399–427.

Haile, Philip A., Han Hong, and Matthew Shum. 2004. “Nonparametric Tests for Common Values in First-Price Sealed-Bid Auctions.” Working paper, Yale Univ., Duke Univ., and Johns Hopkins Univ.

Hendricks, Kenneth, Joris Pinkse, and Robert H. Porter. 2003. “Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions.” *Rev. Econ. Studies* 70 (1): 115–45.

Hendricks, Kenneth, and Robert H. Porter. 1988. “An Empirical Study of an Auction with Asymmetric Information.” *A.E.R.* 78 (5): 865–83.

———. 1993. “Bidding Behaviour in OCS Drainage Auctions: Theory and Evidence.” *European Econ. Rev.* 37 (2–3): 320–28.

Hendricks, Kenneth, Robert H. Porter, and Charles A. Wilson. 1994. “Auctions for Oil and Gas Leases with an Informed Bidder and a Random Reservation Price.” *Econometrica* 62 (6): 1415–44.

Hill, Jonathan B., and Artyom Shneyerov. 2013. “Are There Common Values in First-Price Auctions? A Tail-Index Nonparametric Test.” *J. Econometrics* 174 (2): 144–64.

Hong, Han, and Matthew Shum. 2002. “Increasing Competition and the Winner’s Curse: Evidence from Procurement.” *Rev. Econ. Studies* 69 (4): 871–98.

Hortaçsu, Ali, and Jakub Kastl. 2012. “Valuing Dealers’ Informational Advantage: A Study of Canadian Treasury Auctions.” *Econometrica* 80 (6): 2511–42.

Hubbard, Timothy P., Tong Li, and Harry J. Paarsch. 2012. “Semiparametric Estimation in Models of First-Price, Sealed-Bid Auctions with Affiliation.” *J. Econometrics* 168 (1): 4–16.

Kamakura, Wagner A., and Michel Wedel. 2001. “Exploratory Tobit Factor Analysis for Multivariate Censored Data.” *Multivariate Behavioral Res.* 36 (1): 53–82.

Krasnokutskaya, Elena. 2011. “Identification and Estimation of Auction Models with Unobserved Heterogeneity.” *Rev. Econ. Studies* 78 (1): 293–327.

Krasnokutskaya, Elena, and Katja Seim. 2011. “Bid Preference Programs and Participation in Highway Procurement Auctions.” *A.E.R.* 101 (6): 2653–86.

Laffont, Jean-Jacques, Hervé Ossard, and Quang Vuong. 1995. “Econometrics of First-Price Auctions.” *Econometrica* 63 (4): 953–80.

Laffont, Jean-Jacques, and Quang Vuong. 1996. “Structural Analysis of Auction Data.” *A.E.R.* 86 (2): 414–20.

Larsen, Bradley. 2020. “The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions.” *Rev. Econ. Studies*, forthcoming.

Lebrun, Bernard. 2002. “Continuity of the First Price Auction Nash Equilibrium Correspondence.” *Econ. Theory* 20 (3): 435–53.

Levin, Dan, and James L. Smith. 1994. “Equilibrium in Auctions with Entry.” *A.E.R.* 84 (3): 585–99.

Li, Tong, Isabelle Perrigne, and Quang Vuong. 2000. “Conditionally Independent Private Information in OCS Wildcat Auctions.” *J. Econometrics* 98 (1): 129–61.

———. 2002. “Structural Estimation of the Affiliated Private Value Auction Model.” *RAND J. Econ.* 33 (2): 171–93.

Lizzeri, Alessandro, and Nicola Persico. 2000. “Uniqueness and Existence of Equilibrium in Auctions with a Reserve Price.” *Games and Econ. Behavior* 30 (1): 83–114.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory*. New York: Oxford Univ. Press.
Maskin, Eric, and John Riley. 2000a. "Asymmetric Auctions." *Rev. Econ. Studies* 67 (3): 413–38.
———. 2000b. "Equilibrium in Sealed High Bid Auctions." *Rev. Econ. Studies* 67 (3): 439–54.
Matthews, Steven A. 1984. "'Information Acquisition in Discriminatory Auctions." In *Bayesian Models in Economic Theory*, edited by Marcel Boyer, and Richard Kihlstrom, 181–207. Amsterdam: North-Holland.
McAdams, David. 2003. "Isotone Equilibrium in Games of Incomplete Information." *Econometrica* 71 (4): 1191–214.
———. 2007. "Monotonicity in Asymmetric First-Price Auctions with Affiliation." *Internat. J. Game Theory* 35 (3): 427–53.
Milgrom, Paul R., and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089–122.
Paarsch, Harry J. 1992. "Deciding between the Common and Private Value Paradigms in Empirical Models of Auctions." *J. Econometrics* 51 (1–2): 191–215.
Pinkse, Joris, and Guofu Tan. 2005. "The Affiliation Effect in First-Price Auctions." *Econometrica* 73 (1): 263–77.
Porter, Robert H., and J. Douglas Zona. 1993. "Detection of Bid Rigging in Procurement Auctions." *J.P.E.* 101 (3): 518–38.
Reny, Philip J., and Shmuel Zamir. 2004. "On the Existence of Pure Strategy Monotone Equilibria in Asymmetric First-Price Auctions." *Econometrica* 72 (4): 1105–25.
Roberts, James W. 2013. "Unobserved Heterogeneity and Reserve Prices in Auctions." *RAND J. Econ.* 44 (4): 712–32.
Roberts, James W., and Andrew Sweeting. 2013. "When Should Sellers Use Auctions?" *A.E.R.* 103 (5): 1830–61.
Rothkopf, Michael H. 1969. "A Model of Rational Competitive Bidding." *Management Sci.* 15 (7): 362–73.
Samuelson, William F. 1985. "Competitive Bidding with Entry Costs." *Econ. Letters* 17 (1–2): 53–57.
Wilson, Robert. 1977. "A Bidding Model of Perfect Competition." *Rev. Econ. Studies* 44 (3): 511–18.