Matthew Gentry and Caleb Stroup

Entry and competition in takeover auctions

Article (Accepted version)
(Refereed)

Original citation:
Gentry, Matthew and Stroup, Caleb (2018) Entry and competition in takeover auctions. Journal of Financial Economics. ISSN 0304-405X (In Press)

DOI: https://doi.org/10.1016/j.jfineco.2018.10.007

Accepted Manuscript

Entry and Competition in Takeover Auctions

Matthew Gentry, Caleb Stroup

PII: S0304-405X(18)30296-4
DOI: https://doi.org/10.1016/j.jfineco.2018.10.007
Reference: FINEC 2981

To appear in: *Journal of Financial Economics*

Received date: 12 January 2017
Revised date: 16 January 2018
Accepted date: 29 January 2018

Entry and Competition in Takeover Auctions ☆

Matthew Gentry[1,*], Caleb Stroup[2]

Abstract

We estimate the degree of uncertainty faced by potential bidders in takeover auctions and quantify how it affects prices in auctions and negotiations. The high degree of uncertainty revealed by our structural estimation encourages entry in auctions but reduces a target's bargaining power in negotiations. In the aggregate, auctions and negotiations produce similar prices, even though auctions are preferred in takeover markets with high uncertainty, while the reverse is true for negotiations. Firm characteristics predict pre-entry uncertainty and thus are informative about the relative performance of auctions and negotiations for individual targets.

Keywords:
Mergers and acquisitions, Auctions, Structural estimation, Negotiations, Takeovers, Information frictions
JEL: G34, D44

1. Introduction

The corporate control market is one of the world's largest, with more than ten thousand U.S. companies sold in 2015, totaling $2.47 trillion in deal value. Yet scholars remain divided about how companies should be sold

☆We are grateful to JFE Co-editor Toni Whited and an anonymous JFE referee for suggestions that have substantially improved the paper.

*Corresponding author
Email addresses: m.l.gentry@lse.ac.uk (Matthew Gentry), castroup@davidson.edu (Caleb Stroup)
URL: https://sites.google.com/site/calebstroupeconomics (Caleb Stroup)
[1]London School of Economics and Political Science, Houghton Street, London WC2A 2AE. Tel. +44 (0) 20 - 7955 6213
[2]Davidson College, Box 7123, Davidson, NC 28035 (castroup@davidson.edu)

to maximize sale prices for target shareholders. This disagreement exists because the channels that determine competition in auctions and negotiations are not well-understood.

A commonly held view, originally formalized by Bulow and Klemperer (1996), is that direct competition among bidders generates high auction prices. For example, Wasserstein (2000) reports that "A wide-ranging auction generally maximizes value . . . sophisticated bidders will do their best to circumvent the auction format." Many companies prefer to sell via auction and acquire via negotiation (Auction Process Roundtable, Mergers and Acquisitions, December 2006, pp. 31-32). Given this perspective, it may seem puzzling that about half of companies are sold via single-bidder negotiations.

A different view, based on information frictions, suggests that negotiations outperform auctions. Some scholars have suggested that invited potential bidders face uncertainty about their values for a target and decline when invited to participate in a takeover competition. In this view, absence of potential competitors from the pool of entering bidders weakens direct competition in auctions (e.g., French and McCormick, 1984; Boone and Mulherin, 2007). Recent research has also suggested that negotiated transactions might even benefit from information frictions if a negotiating bidder with an informational advantage shades up its offer price to deter potential competitors (e.g., Fishman, 1988; Aktas, de Bodt, and Roll, 2010; Dimopoulos and Sacchetto, 2014).

Yet information frictions in takeover markets are not directly observable, a fact that has prevented researchers from quantifying the channels that determine the relative performance of auctions and negotiations. Faced with these difficulties, one might be tempted to simply compare observed revenue resulting from auctions with those arising from negotiated transactions. Such an approach might have the potential to yield insights about how firms *are* sold, but not about how firms *should* be sold, in part because the relative optimality of auctions critically depends on the endogenous size and composition of the pool of entering bidders.

We build and estimate a structural empirical framework that enables recovery of the deep takeover market characteristics that determine the performance of auctions and negotiations. Our point of departure is a takeover auction framework in which entering bidders compete for a target by offering successively higher prices until the winning bid is discovered. The key modification we introduce is that invited potential bidders endogenously choose whether to participate based on imperfect information about their heterogeneous values for the target. As we will show, in this framework the degree of information frictions plays an empirically important role in determining whether auctions or negotiations produce higher prices for target shareholders.

Our structural identification strategy relies on a fact, well-known in the econometrics literature, that inferences about information frictions and other takeover market primitives can be recovered from data on bids and entry patterns, which we hand-collect from takeover filings submitted to the Securities and Exchange Commission (SEC). The main insight behind this strategy is that pre-entry uncertainty affects the endogenous entry patterns that in turn determine the competitiveness of the entering bidder pool in takeover auctions (e.g., Roberts and Sweeting, 2013; Gentry and Li, 2014). Uncertainty weakens the link between invited bidders' heterogeneous ex post values and their pre-entry beliefs, thus affecting entry decisions that determine the equilibrium size and composition of the entering bidder pool. The relative magnitudes of the channels through which pre-entry uncertainty affects the competitiveness of the entering bidder pool depend on the magnitudes of the takeover market primitives we study, hence the need to discipline the model via structural estimation.

As a comparison with auctions, we focus on two negotiation procedures, each of whose performance is determined by the takeover market characteristics recovered by our structural estimation. These negotiation procedures are realistic, easy to implement, and have been studied widely in the finance and economics literatures (e.g., Fishman, 1988; Betton and Eckbo, 2000; Povel and Singh, 2006; Horner and Sahuguet, 2007; Bulow and Klemperer, 2009).[3] The first is the sequential negotiation, in which a target can successively negotiate with individual bidders until it reaches an adequate price. The second is a one-shot negotiation followed by an auction-style market-check (i.e., a "go-shop").

We use the structural estimates to quantify three channels through which uncertainty affects the relative performance of auctions and negotiations. First, the *composition effect* in auctions refers to degradation in the quality of the entering bidder pool that arises when pre-entry uncertainty lowers par-

[3]While one could certainly imagine hypothetical complex negotiation procedures, our goal is to analyze negotiation procedures that are realistic and can be easily implemented without the target or its investment bank having specific knowledge about potential bidders' values or beliefs, the degree of pre-entry uncertainty, or the average costs of entry.

ticipation rates among relatively high-value potential entrants. This channel captures the intuition that information frictions discourage entry, potentially impairing takeover auction performance. Second, the *size effect* in auctions refers to the fact that pre-entry uncertainty can encourage entry overall. This effect exists because absence of high-value bidders from the pool of entrants, just described, raises prospects for all potential entrants, thus encouraging entry overall. Ignoring the size effect, or presuming its magnitude to be small relative to the composition effect, leads to the belief that information frictions necessarily impair auction performance. Yet whether the size or composition effect dominates is an empirical question. The third channel, at work in negotiated transactions, is the *deterrence effect*, which refers to the possibility that a standing negotiating bidder shades up its offer price to deter potential competitors. The sign and magnitude of the deterrence effect is theoretically ambiguous, since an increase in uncertainty could in principle either strengthen or weaken a negotiating bidder's incentive to shade up its offer price.[4] The relative performance of auctions and negotiations thus depends on the signs and relative magnitudes of the size, composition, and deterrence effects, which we jointly quantify using the structural estimates.

Our main findings are as follows. First, we show that takeover competitions are characterized by endogenous entry, with fewer than half of invited bidders choosing to participate in takeover auctions. Second, information frictions in takeover markets are large, with potential bidders' pre-entry beliefs about their values containing more noise than information.

Third, we show that pre-entry uncertainty varies across takeover markets. Some takeover markets are characterized by nearly complete pre-entry uncertainty, while in others potential bidders have relatively precise information. We show that variation in pre-entry uncertainty is associated with the target observables. For example, pre-entry uncertainty rises with target size and leverage, falls with *q-ratio* (market-to-book) and cash holdings, while being U-shaped in the sales-to-asset ratio.

Fourth, we show that auctions and negotiations perform similarly overall, with sequential negotiations generating revenue about 1.13 percentage points higher than auctions, a figure corresponding to about \$10 million for a typical

[4] Conceptually, though pre-entry uncertainty unambiguously reduces the marginal effectiveness of upward bid shading, it is not obvious that this impairs negotiation performance in equilibrium since reduced marginal effectiveness could cause a bidder to shade its offer up even higher if the incentive to deter entry is strong.

target in our sample.

Fifth, we show that the aggregate similarity in the relative performance of auctions and negotiations masks heterogeneity arising from differences across targets in the degree of potential bidders' pre-entry uncertainty. We show empirically how differences across targets in the relative performance of auctions and negotiations depends on the size, composition, and deterrence effects in our sample. For example, our estimates reveal that the composition effect is small relative to the size effect, a fact implying that pre-entry uncertainty raises expected revenue in auctions by encouraging entry overall, even while degrading the composition of the entering bidder pool. This finding implies that takeover auctions counterintuitively benefit when potential bidders know less about their values, a result that stands in stark contrast to studies overlooking the size effect. We also show how pre-entry uncertainty affects target bargaining power in negotiations. Our estimates reveal that deterrence bidding accounts for about 13 percentage points of deal premia for a typical target in our sample. Pre-entry uncertainty weakens deterrence bidding, thus reducing takeover revenue in negotiations.

Taken together, these findings imply that auctions tend to produce higher prices in takeover markets with high pre-entry uncertainty, while the reverse is true for negotiations. Quantitatively, the difference in expected revenue between auctions and negotiations is about a percentage point for takeover markets with pre-entry uncertainty at the 25th percentile of the distribution. We show that the relative performance of auctions and negotiations varies systematically with target-level observables and find, for example, that large, highly leveraged targets that hold less cash are associated with more pre-entry uncertainty, thus tending to prefer auctions over negotiations.

This paper contributes to several literatures. First, we contribute to the literature attempting to quantify information frictions faced by potential bidders in takeover markets. Several studies have theoretically proposed the types of frictions we study (e.g., French and McCormick, 1984; Povel and Singh, 2006), or have inferred their possible existence from regressions involving proxies of information frictions (e.g., Boone and Mulherin, 2007; Rousseau and Stroup, 2015). Ours is the first to directly estimate pre-entry uncertainty faced by potential bidders in takeover markets and to show that takeover auctions can counterintuitively benefit from large information frictions, while the opposite is true for negotiations.

We also contribute to the growing literature that studies the channels that affect the relative performance of auctions and negotiations (e.g., Betton and Eckbo, 2000; Povel and Singh, 2006; Horner and Sahuguet, 2007; Aktas, de Bodt, and Roll, 2010). Some studies have theoretically examined the channels studied here or have provided indirect evidence for their possible existence, but without providing a direct comparison of their relative magnitudes (e.g., Aktas, de Bodt, and Roll, 2010). Our study is the first to empirically quantify the size, composition, and deterrence effects in takeover markets. By doing so, ours is thus the first study to empirically quantify the channels through which information frictions affect takeover markets and to use these to identify situations in which auctions or negotiations are more likely to produce higher prices for target shareholders.

Finally, we contribute to the broader debate about whether targets should be sold via auctions or negotiations. Specifically, assuming that potential bidders have no pre-entry information, Bulow and Klemperer (2009) show theoretically that formal auctions yield higher sale prices than do sequential negotiations. Viewed in light of their result, the widespread use of negotiations in practice could be viewed as a corporate governance failure. On the other hand, Roberts and Sweeting (2013) theoretically show that this auction dominance result hinges on the knife-edge assumption that potential bidders are completely uninformed about their values prior to entry. Whether uncertainty is sufficiently low to reverse auction dominance for some targets is an empirical question. Our estimates provide the first such evidence, showing that in many takeover environments sequential negotiations can in fact yield higher revenue than auctions. Importantly, our comparisons feature a level playing field between auctions and negotiations in the sense that for each target we use a single set of estimated takeover market primitives, thus permitting a comparison that mirrors corporate directors' decision to sell their company via either an auction or a negotiation.

Our paper is methodologically similar to two prior studies. In a structural auction framework, Gorbenko and Malenko (2014) recover takeover market unobservables using hand-collected data drawn from proxy statements submitted to the SEC. They take as given the size and composition of the entering bidder pool and explore possible valuation differences between strategic and financial bidders. Our focus is instead on the overall entry patterns that affect the relative ability of auctions to produce high prices for target shareholders. Roberts and Sweeting (2013) develop a framework permitting recovery of takeover market primitives when entry is endogenous and use it to recover primitives that characterize government timber auctions.

The paper is organized as follows: Section 2 provides background on takeover auctions. Section 3 develops the empirical framework. Section 4 describes the data and reports summary statistics. Section 5 recovers takeover market features. Section 6 analyzes the effects of information frictions on auctions and negotiations. Section 7 reports robustness checks. Section 8 concludes.

2. Institutional background

This section describes a typical takeover auction, focusing on aspects specific to the present study (see also Hansen, 2001).

A target board desiring an auction recruits a sell-side advisor to identify and individually contact potential acquirers, inviting them to participate in competitive bidding. Entry into the sale process occurs when a potential bidder signs a confidentiality agreement restricting dissemination of nonpublic information about the target's finances and operations. The provision of nonpublic information allows a bidder to discover its own value of the target, since nonpublic information about the target's finances, operations, and business prospects allows a bidder to assess possible asset complementarities and integration costs specific to possible combination of the two firms' operations. A potential bidder's value is thus pair-specific and possibly uncertain prior to entry.

These costly activities conducted by entrants are undertaken by the bidder's management, in-house deal team, financial advisor, and legal advisor, and focus on the discovery and assessment of nonpublic information permitting an assessment of the bidder's value for the target, i.e., asset complementarities and post-merger integration costs specific to a particular business combination. Discovery activities typically include analyses of supply chains, software and machine technology, research and development overlap, intellectual property, marketing programs, potential technology transfer, retiree pension and medical benefits, debt covenants, customer perceptions of the target and potential acquirer, compatibility of corporate cultures and other human resources, and the strategic reactions of competitors, among others. Entry into takeover auctions is thus costly, not only because of direct pecuniary costs and advisor fees, but also because of non-pecuniary costs associated with forgone acquisition opportunities during the time when negotiations are ongoing, because of reputation risk if negotiations fail, because of potential revelation of proprietary information if a competing bidder wins the takeover competition, and because the bidder's management, board, and deal team expend time that could otherwise be used for other productive projects.

Confidentiality agreements prevent bidders from revealing to other parties the fact that the target is for sale, their presence in the takeover auction, or their bids (e.g., Kirman, 2008). Potential bidders thus decide whether to participate, and what bids to offer, without access to specific knowledge about entry decisions or bids offered by other potential acquirers. In practice, the target provides entering bidders with feedback about the "adequacy" of their bids, and in some cases provides guidance about ranges of offers that will be considered seriously.[5]

After entry, takeover auctions proceed by multiple bidding rounds, with bidders raising their offers or dropping out after receiving unfavorable feedback from the target. In doing so, the target responds to bids by indicating bid adequacy, with this process repeating until the bidder with the highest value remains. If the highest offer price is above the target board's reservation value, the deal is announced publicly. This bidding structure most closely resembles an ascending auction in which bidders successively drop out until the highest bidder remains (e.g., Subramanian, 2011 p. 59).

3. Baseline Model Specification

3.1. Information and entry

A takeover auction of a target j is initiated when N_j potential bidders $i = \{1, ..., J\}$ are invited to participate in competitive bidding. Each of these N_j invited potential bidders decides whether to participate in the takeover auction. Potential bidders formally enter by signing confidentiality agreements with the target, at which point each entering bidder conducts due diligence on the target at cost C_j (e.g., Fishman, 1988; Hansen, 2001). The n_j entering bidders next engage in competitive bidding for the target, with

[5]The fact that neither potential bidders' participation decisions nor offered bids are generally disclosed simplifies our analysis of entry and bidding, since in takeover auctions it eliminates confounding signaling or timing effects that would otherwise arise if entry decisions were concurrently observed by other potential entrants (e.g., Rosenbaum and Pearl, 2009). Such issues would be of much greater concern for studies attempting to recover takeover market primitives from data on takeover negotiations, which are less structured, with many economically relevant features being unobservable to an econometrician.

bidding based on values discovered during due diligence. Entering bidders compete in a standard ascending auction. In such an environment, the dominant strategy of an entrant with value realization V_{ij} is to continue bidding until the current posted price reaches V_{ij}, and to exit when the target indicates that a bid $b_j > V_{ij}$ is adequate (e.g., Subramanian, 2011 p. 59). As a robustness check, we estimated an alternate bidding model and obtained similar results (Section 7, below). If the target's reservation price exceeds even the final bidder's value, the auction concludes with no sale. Otherwise, bidding continues until the purchase price reaches the maximum of the second-highest entrant value and the target's reservation price, at which point competitive bidding concludes.

Let V_{ij} denote potential bidder's i's ex ante unknown value of target j, discovered after entry. Values depend both on a common observable standalone component (M_j) and an idiosyncratic asset complementarity net of integration costs (ν_{ij}) specific to a particular bidder-target pair. This formulation follows the private-values literature on takeover auctions (e.g., Fishman, 1988; Gorbenko and Malenko, 2014).

Following Gorbenko and Malenko (2014), we specify the unconditional distribution of values V_{ij} among potential bidders as:

$$V_{ij} = M_j \exp\{\nu_{ij}\}, \tag{1}$$

where idiosyncratic bidder-target complementarities ν_{it} are drawn independently from a normal distribution with sale-specific mean μ_j and variance σ^2_{vj}. The sale-specific mean μ_j allows for correlation among values while heterogeneity in ν_{ij} reflects the fact that bidders differ in many dimensions (e.g., industrial or product market similarity to the target) that determine their pair-specific asset complementarities and integration costs.

Prior to entry, each bidder i observes a private signal S_{ij} of its (unknown ex ante) value V_{ij}, which informs its entry decision. This signal S_{ij} is related to the value V_{ij} as follows:

$$S_{ij} = V_{ij} \exp\{\varepsilon_{ij}\}, \tag{2}$$

where the error ε_{ij} is Gaussian white noise with variance $\sigma^2_{\varepsilon j}$. Since monotone transformations of a signal preserve information, the marginal distribution of S_{ij} is irrelevant; all that matters is the dependence between V_{ij} and S_{ij}. This result is important because it implies any normalization for S_{ij} generates identical empirical results if the copula between V_{ij} and S_{ij} is preserved. Let

α_j denote the average degree of potential bidders' ex ante uncertainty about their values, parameterized by the following noise-to-signal ratio:

$$\alpha_j \equiv \sigma^2_{\varepsilon j}/(\sigma^2_{vj} + \sigma^2_{\varepsilon j}) \in [0, 1]. \tag{3}$$

Conditional on observing signal $S_{ij} = s$ prior to entry, invited bidder i therefore expects to draw a final (normalized) value V_{ij} from the log-normal distribution

$$V_{ij}/M_j \sim \text{Log-Normal}(\alpha_j \mu_j + (1 - \alpha_j) \ln s,\ \alpha_j \sigma^2_{vj}). \tag{4}$$

Note that α_j close to zero indicates that pre-entry signals are relatively informative, while α_j close to one indicates that pre-entry signals are uninformative. As we will show, this noise-to-signal ratio α_j turns out to be a critical determinant both of the absolute revenue raised through takeover auctions and of the relative revenue ranking of auctions versus negotiations.[6]

The target is endowed with a private reservation value V_{0j}, drawn from the ex ante distribution of potential acquirer values: i.e., $V_{0j} = M_j \exp(\nu_{0j})$ with $\nu_{0j} \sim N(\mu_j, \sigma_j)$. This specification reflects both management's own assessment of the target's stand alone prospects and the fact that the target could always negotiate with a randomly selected potential bidder as an alternative to the current takeover competition. As a robustness check, we directly estimated the target's reservation value distribution and obtained similar results (Section 7, below).

Let $c_j = C_j/M_j$ denote the entry cost C_j as a fraction of target j's standalone value M_j. Since both C_j and c_j are indexed by a specific target, this normalization is without loss of generality. Taking observables M_j and N_j as given, the takeover environment for target j is summarized by the vector $\theta_j \equiv (\mu_j, \sigma_{\nu j}, c_j, \alpha_j)$.

[6]Conventional studies of auctions with endogenous entry have focused on two knife-edge cases. In the first, originally developed by Samuelson (1985), invited potential bidders have perfect knowledge about their own valuations. In the second, originally developed by Levin and Smith (1994), invited potential bidders have no knowledge about their own valuations. These polar assumptions have been traditionally used because they simplify auction theory. These knife-edge assumptions are problematic for the empirical study of takeover market performance because as we will show the average degree of pre-entry uncertainty influences endogenous entry patterns in auctions and deterrence bidding in negotiations.

3.2. Equilibrium

As described above, bidding strategies in the takeover auction itself are straightforward: conditional on entry, bids are linked with values through the bidding equilibrium. In what follows, we therefore focus on characterizing the entry equilibrium and the distribution of entering-bidder values, which determine the distribution of bids. This section describes key qualitative features of equilibrium entry behavior, aiming to build intuition for the subsequent analysis. A formal treatment of equilibrium is provided in Appendix A.

In any symmetric monotone Bayesian Nash equilibrium, the entry stage involves a threshold such that each invited bidder i enters if its signal s_{ij} is above s_j^*. This threshold s_j^* is determined endogenously by the condition that each potential bidder enters if expected post-entry profit is positive.

To see how s_j^* is determined, consider a potential entrant with signal realization s_{ij} facing $N_j - 1$ potential rivals $k \neq i$ who enter when $S_{kj} \geq s_j^*$. If potential entrant i chooses to enter, it incurs the entry cost c_j, learns its value v_{ij}, and bids, winning the auction if v_{ij} exceeds both the seller's reserve and the maximum value among entering rivals, paying a price p_j equal to the maximum among these in this event.

Conditional on entry, bidder i with signal s_{ij} expects to draw its value from the conditional distribution $F(v|s_{ij}, \theta_j)$. Meanwhile, for each rival $k \neq i$, i's expectations encompass two possibilities: either $S_{kj} < s_j^*$ and bidder k remains out, or $S_{kj} \geq s_j^*$ and bidder k enters. In the latter case, i expects k to draw values from the distribution $F(v|S_{kj} \geq s_j^*)$: i.e., the *selected* distribution of values among equilibrium entrants.

Now we return to the entry decision of potential entrant i. As we show formally in Appendix A, the expected post-entry profit of a potential bidder with signal realization s_{ij} is increasing in s_{ij}, increasing in $s_{k\neq j}^*$, and decreasing in N_j. In other words, all else equal, potential entrants prefer higher expected values (higher s_{ij}), less aggressive rival entry (higher $s_{k\neq j}^*$), and less potential competition (lower N_j). For any competition level N_j and target characteristics θ_j, there is therefore a unique signal threshold s_j^* such that a potential entrant drawing signal realization $S_{ij} = s_j^*$ earns expected post-entry profit equal to its entry costs c_j. This zero-profit condition yields the threshold s_j^* characterizing entry in symmetric Bayesian Nash equilibrium.

Holding all else constant, the entry threshold s_j^* will be increasing in both N_j and c_j: facing either more potential competition or larger entry costs, potential bidders will require more positive pre-entry information in order to enter. Since greater dispersion of post-entry values implies greater expected profit, and in view of our log-normal specification the variance of V_{ij} is increasing in both μ_j and $\sigma_{\nu j}$, s_j^* will generally be increasing in both μ_j and $\sigma_{\nu j}$. Finally, reducing α_j will lead all potential entrants to place more credence on their pre-entry signals, which in particular will lead the *marginal* entrant (i.e., the potential bidder with $S_{ij} = s_j^*$) to believe more strongly that it will draw a relatively low value conditional upon entry. Hence, all else equal, higher α_j (more uncertainty) will also tend to increase entry, though the magnitude of this effect depends on the full vector of fundamental parameters.

From the target's perspective, however, the comparative statics of the auction process are far less clear: both expected takeover revenue and the attractiveness of auctions relative to negotiations will depend critically, and not necessarily uniformly, on the specifics of an individual takeover auction. For instance, by more closely aligning potential bidders' pre-entry signals with their post-entry values, more precise pre-entry information (lower α_j) will increase the degree to which bidders with disproportionately high values self-select into entry. But it will also (as seen above) increase the entry threshold s_j^* and therefore lead to less entry overall. It is an empirical question which of these effects dominates, depending on both N_j and the entire vector of model primitives θ_j through the equilibrium threshold s_j^*. Similarly, as we show below, whether formal auctions or negotiations dominate for a given target will depend on the entire vector of primitives θ_j. Theory alone does not specify which combination of model parameters best describes takeover markets, hence the need to discipline the model with the empirical estimates.

3.3. Identification

This section provides a conceptual overview of how primitives of the takeover environment are identified. For simplicity's sake, we first consider identification within a sample of targets with similar characteristics θ_j. We then proceed to discuss complications induced by target-level heterogeneity. A more formal treatment of both issues can be found in Gentry and Li

(2014).[7]

The first insight, which is common to structural auction models such as Gorbenko and Malenko (2014), is that the distribution of bidder values can be recovered from observed bids. In the context of the endogenous entry model we consider here, this implies identification of the distribution of values among bidders electing to enter given θ_j and N_j; i.e., $F(v|S_{kj} \geq s_j^*)$ in the notation above. Within the specific parameterization we consider here, this distribution $F(v|S_{kj} \geq s_j^*)$ is determined by the parameters (μ_j, σ_{vj}) of the ex ante value distribution $F_0(\cdot)$, the parameter α_j governing precision of pre-entry information, and the equilibrium signal threshold s_j^* above which invited bidders elect to enter.

The second insight is that as expected post-entry profits fall, fewer invited potential bidders elect to enter. For a given level of potential competition N_j and fundamental parameters θ_j, the fraction of invited bidders who choose to enter will identify the equilibrium probability of entry $1 - F_s(s_j^*)$. This disciplines model parameters in two important respects. First, given μ_j, σ_{vj}, and α_j, the fraction of invited bidders choosing to participate is directly informative about entry costs c_j, since potential bidders choose to participate only when the expected payoff net of participation costs is positive. Second, as shown above, the equilibrium entry threshold s_j^* will be increasing in both c_j and the number of invited bidders N_j since potential bidders expecting higher costs or more competitors will require more favorable private signals in order to participate. Hence, all else equal, more invited bidders (higher N_j) will induce more selected distributions of values among entrants.

The third insight is that the degree to which the selected entrant value distribution $F(v|S_{kj} \geq s_j^*)$ changes with s_j^* is directly informative about the noise-to-signal ratio α_j. To see why, recall that as expected post-entry profits fall and s_j^* rises, fewer potential bidders choose to enter, but those that do enter have higher values on average ($F(v|S_{kj} \geq s_j^*)$ is increasing in s_j^*). Intuitively, higher pre-entry beliefs s_{ij} are required to rationalize entry. When α is small, the s_{ij} primarily reflect information about values, and relatively

[7]Gentry and Li (2014) study nonparametric identification in a general class of selective entry auctions nesting the ascending auction considered here. They formally show how the parameters characterizing endogenous entry patterns, for example, pre-entry uncertainty, can be recovered from a structural model based on observed variation in entry patterns and winning bids. They then verify that identification extends to the case of auction-level unobserved cross-sectional differences across targets.

high-value potential bidders disproportionately elect to participate, implying a steep gradient of $F(v|S_{kj} \geq s_j^*)$ in s_j^*. In a world of high pre-entry uncertainty, the s_{ij} are noisy and entry patterns become less tightly associated with values, implying a relatively flat gradient of $F(v|S_{kj} \geq s_j^*)$ in s_j^*. Since the degree to which the selected entrant value distribution $F(v|S_{kj} \geq s_j^*)$ changes with s_j^* is directly informative about the noise-to-signal ratio α_j, comparing changes in $F(v|S_{kj} \geq s_j^*)$ across targets with different numbers of invited bidders N_j pins down α_j. In turn, having determined α_j, knowledge of $F(v|S_{kj} \geq s_j^*)$ for any entry probability $1 - F_s(s_j^*)$ uniquely identifies $F_{0j}(\cdot)$ and hence its parameters (μ_j, σ_{vj}).[8] Thus, recovery of the distributions that characterize the composition of the invited and entering bidder pools, along with direct observation of their sizes in our data set, permits sharp identification of the average degree of pre-entry uncertainty.

As we describe in detail below, our econometric implementation accommodates heterogeneity across targets in ways that are observable to market participants but unobservable to the econometrician. In this case, formal identification is more involved, as observed outcomes potentially arise from unobserved mixtures across θ_j. Nevertheless, the intuition of the argument follows closely from the discussion above: All else equal, the distribution of transaction prices will be informative about the distributions of μ_j and σ_j, average entry frequencies will be informative about the distribution of entry costs c_j, and changes in prices across otherwise similar targets with different numbers of invited bidders N_j will be informative about α_j. In our structural implementation, we analyze all of these effects jointly within

[8] Formally, this is implied by the following identity, which follows from Equation A.6 in Appendix A:

$$F(v|S_{kj} \geq s_j^*) = \frac{F_{0j}(v) - F_{vs,j}(v, s_j^*)}{1 - F_{sj}(s_j^*)} = \frac{F_{0j}(v) - C_{\alpha_j}(F_{0j}(v), F_{sj}(s_j^*))}{1 - F_{sj}(s_j^*)},$$

where C_{α_j} is the *copula* describing dependence between pre-entry signals S_{ij} and post-entry valuations V_{ij}. Given our parametric assumptions above, one can show that C_{α_j} is a bivariate Gaussian copula with covariance parameter $1/\alpha_j$. Knowledge of α_j therefore determines C_{α_j}, from which identification of both $F(v|S_{kj} \geq s_j^*)$ and $F_{sj}(s_j^*)$ implies identification of $F_{0j}(v)$. Note that given a sufficiently large sample of targets with similar θ_j, we may apply this argument without parametric assumptions on the marginal distribution F_{0j} of V_{ij}, from which it follows that parameterization of the copula alone is sufficient for identification.

a conditional maximum likelihood framework, allowing all underlying model parameters to affect all observed outcomes including, for example, the fraction of invited bidders who participate, the cross-sectional distribution of the number of entrants, and the distribution of bids. Combined with a parameterization described below, this in turn permits recovery of the distribution of target-level primitives θ_j across targets in our sample.

3.4. Estimation Strategy

Let p_j denote the realized revenue from sale of target j, defined as the sale price normalized by the target's market value four weeks prior to sale announcement. As above, let N_j and n_j denote the number of invited and participating bidders, respectively, in target j's sale process. Conditional on the number of invited bidders N_j and target-level primitives $\theta_j = (\mu_j, \sigma_{\nu j}, \alpha_j, c_j)$, target-level outcomes (p_j, n_j) are determined endogenously through the formal entry and bidding process described in Sections 3.1 and 3.2.

Let X_j denote a vector of target-level characteristics potentially influencing target sale-level primitives θ_j. Elements of X_j include the target's industry, book value of total assets (current assets plus net property, plant, and equipment plus other noncurrent assets), market leverage, *q-ratio*, along with cash, intangibles, and sales, with these latter three variables scaled relative to book assets.

While targets clearly differ along many elements of X_j, takeover environments may also differ in other dimensions not observed by the researcher. To accommodate this, we allow target-level primitives $\theta_j = (\mu_j, \sigma_{\nu j}, \alpha_j, c_j)$ to vary across takeover markets based on both target-level observables and factors unobserved by the econometrician. Specifically, we assume that the distribution of primitives θ_j conditional on observables X_j is governed by a distribution $\theta_j \sim g(\theta_j|X_j, \Gamma)$, where Γ is a vector of structural parameters to be estimated. We specify this distribution $g(\theta_j|X_j, \Gamma)$ as follows:

$$\begin{aligned} \mu_j &\sim \text{Normal}(\text{mean} = \gamma_\mu X_j, \text{var} = \sigma_\mu^2) \qquad (5)\\ \sigma_{vj} - \tau &\sim \text{Gamma}(\text{mean} = \exp(\gamma_\sigma X_j), \text{shape} = \sigma_\sigma)\\ c_j &\sim \text{Gamma}(\text{mean} = \exp(\gamma_c X_j), \text{shape} = \sigma_c)\\ \alpha_j &\sim \text{Beta}(\text{mean} = \exp(\gamma_\alpha X_j)/\left(1 + \exp(\gamma_c X_j)\right), \text{var} = \sigma_\alpha^2) \end{aligned}$$

where Gamma (mean $= \mu$, shape $= \sigma$) denotes a Gamma distribution parameterized to have mean μ and shape parameter σ, Beta (mean $= \mu$, var $= \sigma^2$)

denotes a Beta distribution parameterized to have mean μ and variance σ^2, and $\tau > 0$ is a regularization constant which ensures that the variance parameter σ_{vj}^2 is bounded away from zero.[9] The elements of θ_j are independent conditional on X_j (i.e., the errors are independent).

Our goal is to use target-level outcomes (p_j, n_j, N_j) to conduct inference on the deep structural primitives $\Gamma_0 = \{\gamma_\mu, \gamma_\sigma, \gamma_c, \gamma_\alpha; \sigma_\mu^2, \sigma_\sigma^2, \sigma_c^2, \sigma_\alpha^2\}$ characterizing the distribution of $\theta_j|X_j$. Toward this end, first note that the sample of observed takeovers is a selected sample by construction, since auctions not resulting in sale are not reported to the SEC and are thus not observable. To account for this fact, we therefore consider estimation of Γ_0 based on conditional maximum likelihood with respect to the following event: n_j bidders enter and the final sale price is p_j, conditional on the target contacting N_j potential bidders and the auction resulting in sale. We construct the conditional likelihood function corresponding to this event as follows.

Let $sale_j$ be an indicator taking a value of one if an auction results in sale. Taking θ_j and N_j as given, let $\Pr(n_j|N_j;\theta_j)$ be the probability that n_j of N_j invited potential bidders ultimately enter, let $\Pr(sale_j|n_j, N_j, \theta_j)$ be the probability that the auction of target j leads to an announced sale, given entry by n_j of N_j invited potential bidders, and let $f_p(p_j|sale_j, n_j, N_j; \theta_j)$ be the density of the final price p_j, given entry by n_j of N_j potential bidders *and* that the auction results in an announced sale. Integrating over the target-level unobservables θ_j, we may then express the predicted likelihood of observing p_j, n_j conditional on N_j, X_j, and $sale_j = 1$ in terms of the unknown parameter vector Γ as follows:

$$
\begin{aligned}
&L_j(p_j, n_j|sale_j = 1, N_j, X_j; \Gamma) = \\
&\int f_p(p_j|sale_j, n_j, N_j; \theta_j) \cdot \Pr(sale_j|n_j, N_j; \theta_j) \cdot \Pr(n_j|N_j; \theta_j)\ g(\theta_j|X_j, \Gamma)\ d\theta_j \\
&\qquad\qquad \Big/ \int \Pr(sale_j|n_j, N_j; \theta_j) \cdot \Pr(n_j|N_j; \theta_j)\ g(\theta_j|X_j, \Gamma)\ d\theta_j. \qquad (6)
\end{aligned}
$$

We evaluate this expression via simulation by first deriving predicted equilibrium values of $f_p(p_j|sale_j, n_j, N_j; \theta_j)$, $\Pr(sale_j|n_j, N_j; \theta_j)$, and $\Pr(n_j|N_j; \theta_j)$ for many possible values of θ_j from the theoretical entry and bidding model above and then taking appropriate weighted averages over these to evaluate

[9] We also explore estimation under several alternative specifications, such as using truncated log-normal distributions for the parameters σ_{vj}, c_j, and α_j, and find similar results.

the integrals on the right-hand side of (6). We next maximize the (log of) the resulting likelihood function over Γ to recover conditional maximum likelihood estimates of the deep structural parameters Γ_0. Appendix B describes the computational details of this procedure.

3.5. Negotiation Formats

In this section, we formalize two commonly studied negotiation frameworks, which we compare to the auction benchmark using the takeover market primitives estimated using the structural procedure above. By conducting counterfactual comparisons using a single set of takeover market primitives, we circumvent the empirical challenge that transactions involving a single bidder can take a variety of strategically differing forms that produce different levels of expected revenue, but which are also observationally indistinguishable to an econometrician. Negotiation structures differ, for example, in the extent of competitive pressure exerted by potential competition. Even more troubling, a single-bidder sale observed in the data could reflect different negotiation structures, such as a successful one-shot negotiation or a successful first stage in a sequential negotiation, among other possibilities.

The first negotiation format we consider is a *sequential negotiation*, variants of which have been studied extensively in the theoretical literatures on takeovers and mechanism design (e.g., Fishman, 1988; Daniel and Hirshleifer, 2018; Betton and Eckbo, 2000; Povel and Singh, 2006; Horner and Sahuguet, 2007; Bulow and Klemperer, 2009; Aktas, de Bodt, and Roll, 2010; Dimopoulos and Sacchetto, 2014; Roberts and Sweeting, 2013). Relative to the standard auction outlined above, this sequential negotiation differs in two respects. First, rather than contacting all potential buyers at once, the target negotiates with one standing bidder at a time and successively approaches additional potential buyers in an attempt to obtain a higher price. Second, a standing negotiating bidder has market power in the sense that it has an informational advantage and can shade up its offer price to deter subsequent rival entry.

This sequential negotiation procedure proceeds in N rounds, one round for each potential buyer. Round $n = 1, ..., N$ begins when the target approaches potential buyer $i = 1, ..., N$ (ordered at random) with an invitation to participate in a negotiation. The sequence of events in round n then proceeds as follows:

1. Potential buyer i observes its private signal S_{ij} for the target. Based on this signal S_{ij} and the entry and bidding history up to round $n-1$, potential bidder i determines whether to enter the negotiation at cost c.
2. Conditional on choosing to enter, potential buyer i learns its value V_{ij}. If another negotiating bidder has previously entered, potential buyer i and the current incumbent compete in an ascending button auction for the right to remain in the auction. The loser of this bidding round exits and the winner becomes the incumbent, with the current standing price being the level at which the loser drops out.
3. Conditional on outbidding the current incumbent, potential buyer i may submit a bid above the current standing price. If submitted, this deterrence bid is observed by all subsequent potential buyers, and becomes the standing price in round $n+1$.

This sequential negotiation game corresponds to the game considered by Bulow and Klemperer (2009) in all but one respect: Whereas Bulow and Klemperer (2009) assume potential bidders have no private information about their values prior to entry, our framework allows potential bidders to observe partially informative pre-entry signals about their post-entry values. As we will see, accounting for endogenous entry patterns in the presence of such imperfect pre-entry information, most targets would in fact earn greater revenue from sequential sales than they would from auctions. This finding contrasts with Bulow and Klemperer (2009), who theoretically show that, in their environment without pre-entry information, auctions yield higher expected revenue than sequential negotiations.

The price in this sequential negotiation evolves as follows. Let b_{n-1} be the standing bid at the beginning of round n, and y_{n-1} be the value of the incumbent submitting bid b_{n-1}. Suppose that i elects to enter in round n, drawing value v_{ij} upon entry. Then, three outcomes are possible in round n. If v_{ij} is less than the current standing bid b_{n-1}, then i exits and the negotiation proceeds to round $n+1$ with standing bid $b_n = b_{n-1}$. If v_{ij} is greater than b_{n-1} but less than y_{n-1}, then i bids up the price to v_{ij} before exiting, and the negotiation proceeds to round $n+1$ with standing bid $b_n = v_{ij}$. Finally, if v_{ij} is greater than y_{n-1}, then the current incumbent bids up the price to y_{n-1} before exiting, and i becomes the new incumbent. To signal strength and thereby deter future entry, incumbent i may then submit a *deterrence bid* $b_n \geq y_{n-1}$. This deterrence bid b_n then becomes the standing price in round $n + 1$. The standing bid at the start of Round 1 (i.e., b_0) is the target's reservation valuation v_{0j}.

We focus on the unique *symmetric separating perfect Bayesian equilibrium* within this sequential negotiation game. This equilibrium has two key components. First, entry decisions by bidder i in round $n = 1, ..., N$ are described by a signal threshold $s^*_n(y_{n-1})$ such that potential buyer i enters if and only if $S_{ij} \geq s^*_n(y_{n-1})$, where y_{n-1} denotes the value of the standing bidder and $s^*_n(y_{n-1})$ is strictly increasing in y_{n-1}. Second, conditional on outbidding an incumbent with value y_{n-1} in round $n < N$, new incumbent i with value $v_{ij} \geq y_{n-1}$ submits a deterrence bid $b_n \geq y_{n-1}$ described by a symmetric monotone deterrence bidding strategy $\beta_n(v_{ij}, y_{n-1})$. This deterrence bidding strategy $\beta_n(\cdot, y_{n-1})$ embodies by the following trade-off: a new incumbent can credibly signal a higher value (and thereby deter rival entry) by submitting a higher deterrence bid, but must then pay a higher price conditional on winning with this higher bid. As described in detail in Appendix C, the resulting separating equilibrium exists, is unique, and implies deterrence bidding strategies $\beta_n(\cdot, y_{n-1})$ which, for round $n < N$ and any standing value y_{n-1}, are strictly monotone in v for all $v \geq y_{n-1}$. Furthermore, as shown by Roberts and Sweeting (2013), this separating equilibrium is the only perfect Bayesian equilibrium to survive standard refinements (no weakly dominated strategies, sequential equilibrium, and the D1 refinements of Cho and Sobel (1990) and Ramey (1996)) on equilibria of the underlying sequential negotiation game.

The second negotiation procedure, formally described in Appendix C, is a one-shot negotiation followed by a market check. This mechanism provides a simple alternative to a formal takeover auction and captures the possibility that a target may choose to conduct an auction-style process after concluding negotiations with a potential buyer (e.g., Wasserstein, 2000; Subramanian, 2008; Wang, 2016). In this mechanism, the target approaches a potential buyer with an invitation to participate in a negotiated sale. Based on its pre-entry beliefs and the entry cost, the potential buyer determines whether to participate in the negotiation. Conditional on choosing to enter, the potential buyer learns its value and submits a bid for the target. If, after negotiations between this bidder and the target conclude, the agreed-upon price is higher than the target's reservation value, then the bid is announced publicly and other potential bidders are invited to enter and to make more competitive bids for the target. Based on their own signals, the entry cost, and the posted price, additional bidders choose whether to enter and furnish a bid.

The bidder with the highest offer acquires the target.[10]

The equilibria in both negotiation structures described above capture the well-known potential advantage of negotiations in which a standing bidder has a first-mover advantage that allows it to deter potential competition by elevating its offer price (e.g., Aktas, de Bodt, and Roll, 2010).[11] We will show that the strength of such deterrence bidding—and thus the target's realized bargaining power in negotiations—crucially depends on information frictions in takeover markets, which determine the extent of an initial bidder's informational advantage over potential competitors (e.g., Povel and Singh, 2006).

4. Data and Summary Statistics

Our sample of takeovers comes from the Securities Data Corporation (SDC) mergers and acquisitions (M&A) database, filtered by the following criteria:

- The takeover was announced between January 1, 2000 and January 1, 2010;
- The winning bidder is publicly listed and obtains 100% of the target's shares as a result of the deal;
- Deal backgrounds are available on the SEC Electronic Data Gathering, Analysis, and Retrieval (EDGAR) online filing system;[12]
- The target is a publicly traded Standard & Poor's (S&P) 1500 company;

[10]For example, the proxy statement submitted by Disc Graphics to the SEC in conjunction with its sale to Main Street Resources in December 2002 indicates that the sale resulted from a nine-month negotiation between Disc Graphics and Main Street Resources, followed by an auction in which the target's investment bank invited 15 potential bidders to participate in a post-signing market check.

[11]The separating equilibria have the feature that a standing bidder's type can be inferred from its jump bid.

[12]Deal backgrounds are contained in SC-TOT, 14D-9, PREM 14C, DEFM 14C, DEFS 14A, and S-4 filings, available at http://www.sec.gov/edgar/searchedgar/companysearch.html.

- The deal is an auction,which we define as one having at least two invited bidders;

- Financial data on the target are available from the S&P Compustat database.

We also required data on the winning bid to be available in the Thomson M&A database. As a robustness check, we estimated the model on a set of bids manually recorded from press releases and found similar results. In our context, possible misvaluation of reported stock bids would reflect measurement error, captured by the heterogeneity distributions described in Section 3. For robustness we also estimated model parameters on the subsample of winning cash bids and found similar results.

Target characteristics come from the Compustat database and include the target's industry, book value of total assets, market leverage, *q-ratio*, along with cash, intangibles, and sales, with these latter three variables scaled relative to book assets. We exclude observations for which the winning bid is below the target's share price or greater than 200% above the target's share price four weeks prior to announcement, or for which the target's market value is greater than $1 trillion.

Table 1 reports sale characteristics for the 529 auctions in our sample, both for the entire sample (reported in the first row) and for each of the 12 Fama and French (1997) industries. The second two columns show that limited participation is a pervasive feature of takeover markets. Only 44% of invited potential bidders (on average, five out of 13) participate in takeover auctions. There is some variation in limited participation across industries: More than half of invited bidders participate in auctions of retail firms, while fewer than a third participate in auctions of energy companies. The fourth column shows that winning bids are on average higher in industries where a greater fraction of invited bidders participate, consistent with our model in which higher average potential bidder values lead to higher signals and more entry. Standard deviations of these variables are reported in parentheses. Wide variation in entry patterns and winning bids, even within specific industries, underscores our estimating approach is robust to the presence of unobserved heterogeneity not captured by target-level observables.

Table 2 reports descriptive statistics (means and standard deviations) of target characteristics. Targets in our sample are similar to average Compustat-listed and S&P 1500 firms.[13]

5. Characterizing takeover environments

Panel A of Table 3 reports the estimates obtained using the method outlined in Section 3. We begin by comparing specifications. The first two columns report estimates obtained from different versions of the numerical estimation procedure. The estimates in Column 1 are obtained using uninformed simulation draws, while those in Column 2 use simulation draws informed by the estimates in Column 1.[14] The estimates are similar, with some improvement gained from the additional iterations in Column 2.

The estimates in Columns 3 - 5 condition on the vector of target observables X_j, which includes target size, leverage, cash ratio, sales-to-asset ratio, intangibles-to-asset ratio, market-to-book ratio, and indicator variables for the 12 Fama-French industries. The vector of target observables enters the distribution of μ in Column 3, the distributions of μ and α in Column 4, and the distributions of μ, α, c, and σ in Column 5. The estimates of pre-entry uncertainty and other fundamental takeover market characteristics are similar across specifications, an expected finding given the fact that our identification strategy is based on information about endogenous entry patterns and bids, rather than target characteristics (see Section 3.3). At the same time, there is a possibility of overfitting the estimation with free parameters: The estimates in Column 5, for example, add an additional 68 parameters to the baseline specification. The appropriate specification test for covariate inclusion in our context is the likelihood ratio (LR) comparison of restricted versus unrestricted models. LR tests indicate that Column 4 is the appropriate specification, i.e., that the distributions of μ and α should be conditioned on target covariates.[15]

The main estimates of pre-entry uncertainty in Column 4 reveal that potential bidders have imperfect information about their values for the target,

[13] See, for example, Panel A of Table 3 in Moeller, Schlingemann, and Stulz (2004) for Compustat firms and Table 2 in Rousseau and Stroup (2015) for S&P 1500 firms.

[14] See Appendix B for a detailed description of the estimating procedure.

[15] Specifically, we conduct LR tests of the model excluding observables against the alterative where they enter in the distribution of μ (LR statistic 58.6 with p-value 1.7e-06). The null of observables entering μ is rejected against the alternative of observables entering both μ and α (LR statistic 35.2 with p-value 0.006). The LR statistic fails to reject the null of observables in μ and α against the alternative in which they enter all distributions (LR statistic 21.5 with p-value 0.95).

with pre-entry signals containing more noise than information ($\alpha = 0.58$). Likelihood ratio tests strongly reject the views that potential bidders are perfectly informed or completely uninformed about their values for a target (LR statistic 319.0 and p-value of 0.00 for $\alpha = 1$ and LR statistic 1750.1 and p-value of 0.00 for $\alpha = 0$). Entry costs, reported in the second row of Column 4, average about one percent of deal value.

These findings have important implications for the understanding of takeover markets: Together, imperfect pre-entry information and costly participation cause potential bidders with unfavorable pre-entry beliefs to decline participation, some of which would have discovered higher values for the target had they chosen to participate. Thus, pre-entry uncertainty causes relatively high-value potential buyers to be absent from the pool of participating bidders, an effect we empirically quantify below in Section 6.2. At the same time, the fact that pre-entry beliefs contain some information implies that high-value potential entrants are still disproportionately represented in the pool of entering bidders relative to the pool of potential bidders.

Panel B disaggregates the primitives and reveals significant variation across takeover markets. The standard deviation of pre-entry uncertainty is 0.13, and over a fourth of takeovers occur in environments where potential bidders' pre-entry beliefs primarily reflect information (i.e., $\alpha < 0.4$ at the 25th percentile of the distribution). As we will show in Section 6 below, pre-entry uncertainty systematically affects the relative performance of auctions and negotiations, implying that the cross-sectional variation in pre-entry uncertainty shown here results in differences across takeover markets in the relative performance of auctions and negotiations.

Table 4 reports the partial effects associated with individual covariates in the fundamental parameter distributions. In interpreting these, it is important to recall how our parameters are identified. The structural parameters in our framework are identified based on information about entry patterns and bids, observable by the econometrician ex post, i.e., after a sale process is completed. Yet our analysis also accounts for differences in takeover markets based on target-level characteristics (e.g., size) that can be observed ex ante, i.e., before the beginning of a sale process. While it is true that these observable target characteristics cannot by themselves identify our model, they could turn out to be systematically associated with the takeover market primitives recovered above. The partial effects in Table 4 quantify the nature of this association. Note, however, that insofar as one is ultimately interested in the overall associations between target observables and takeover prices, these marginal effects may be somewhat difficult to interpret, since covariates may be correlated with each other, while also affecting multiple fundamental parameters simultaneously.

To address these issues, we sort targets by 12 Fama-French industries and compute average pre-entry uncertainty within each industry. The findings are reported in Panel A of Table 5. Average pre-entry uncertainty ranges from 0.87 for nondurables to 0.35 in oil, gas, and coal. Panel A also reports 2.5 and 97.5 percentiles of the fundamental parameter distributions and reveals substantial variation across targets, even within each Fama-French industry. For example, potential bidders for a target at the 2.5th percentile in business equipment have near-perfect pre-entry information ($\alpha = 0.08$), in contrast to potential bidders for a target at the 97.5th percentile who are on average relatively uninformed ($\alpha = 0.75$).

In Panel B of Table 5 we sort targets in our sample based on target size, cash, leverage, q-ratio, and sales, then calculate average pre-entry uncertainty implied by the structural model across targets within each quantile. We would expect to observe less pre-entry uncertainty in sales of companies with high cash ratios, since cash reflects a component of total assets that is relatively less difficult to value. Panel B of Table 5 shows that pre-entry uncertainty does indeed fall with target cash holdings at every size quantile, averaging 0.72 among small, low-cash targets and 0.33 among large, high-cash targets. Pre-entry uncertainty also rises with target leverage, for example, from 0.39 to 0.62 for large targets. Pre-entry uncertainty is negatively associated with the q-ratio, for example, falling from 0.70 to 0.44 between the first and fifth quantiles for medium-sized targets. Having shown that pre-entry uncertainty varies in the cross-section of takeover markets, we now examine how this variation affects the overall performance of takeover auctions and negotiations for individual targets.

6. Comparing auctions and negotiations

6.1. Aggregate results

We now use the structural estimates to show how pre-entry uncertainty and endogenous entry affect takeover revenue in auctions and negotiations. This comparison is possible because the takeover market primitives recovered by our structural estimation characterize expected revenue for each target in both auctions and negotiations. The counterfactual comparisons are obtained using the structural model described in Section 3 and the fundamental parameter estimates from Section 5.[16]

The first row of Table 6 reports expected revenue for each sale procedure, aggregated across all takeovers. Average revenue is similar across auctions (41.17%), sequential negotiations (42.61%), and negotiations followed by a go-shop (41.48%). This finding has important implications for the understanding of how firms should be sold. Theoretical comparisons of auctions and negotiations have suggested that direct competition among entering bidders in auctions generates higher revenue for target shareholders than can be obtained from negotiated transactions, a conclusion that has been used to interpret the widespread prevalence of negotiations as reflecting managerial capture of corporate takeover processes (e.g., Bulow and Klemperer, 1996, 2009).[17] The belief that auctions always revenue-dominate negotiations would imply support for more stringent interpretations of Delaware's *Revlon* ruling, originally construed as directing target boards to act as "auctioneers charged with getting the best price for the stockholders."

The estimates also reveal two important sources of variation masked by aggregate comparisons of auctions and negotiations. The first of these is variation in realized revenue arising from unpredictable aspects of the sale process. The second row of Table 6 reports wide variability in revenue for a particular target, with the standard deviation of takeover auction revenue equal to 11%. This high variability in outcomes points to the fact that the pool of entering bidders is itself variable due to endogenous entry patterns in the face of pre-entry uncertainty. Indeed, the bottom row of Table 6 shows that the invited bidder with the highest value declines to participate in about a quarter of takeover auctions.

Negotiations have greater return predictability than auctions, with revenue standard deviation of 3% for sequential negotiations and 9% for negotiations followed by a go-shop (row two of Table 6). At the same time, the

[16]Formally, estimates are obtained from the posterior revenue distribution for each of the three sale mechanisms described above by drawing vectors (X_j, θ_j) from the prior likelihood and by using the structural environment defined in Section 3 and the fundamental parameters obtained in Section 5 at median observable and unobservable characteristics, setting μ_v to match observed revenue, $\nu_{0j} = 1$, and constructing expected revenue for each sale procedure, then repeating this procedure 10,000 times.

[17]Another approach in this debate has been to compare observed revenue in transactions structured as auctions with those structured as negotiations, but without tackling econometric concerns addressed by our structural procedure.

fourth row of Table 6 quantifies this effect and reveals that targets extract more surplus in sequential negotiations than in auctions. These effects occur because equilibrium bid shading in negotiations deters potential competitors, leading to greater surplus extracted for targets, but less variability in observed revenue. Consistent with this idea, the fifth row of Table 6 shows that the potential bidder with the highest value is less likely to acquire a target sold by a negotiation, relative to an auction. Taken along with the fact that target boards may be reluctant to initiate relatively risky sale processes, our finding that auction revenue is relatively unpredictable provides a new explanation for the widespread prevalence of negotiated transactions in corporate takeover markets.

The second source of variation is in the relative performance of auctions and negotiations across takeover markets. We construct the distribution of auction-negotiation revenue differences using the method described in Section 3 and find that auction revenue is greater than negotiation revenue for about a third of takeovers. The standard deviation of target-level revenue differences is 1.7%, or about \$39 million for a target at the 75th percentile of the size distribution. Our work is thus the first to disaggregate the relative performance of auctions and negotiations in this way, showing that either auctions or negotiations may be unambiguously preferable for a particular target, while at the same time both sale procedures perform similarly on aggregate. We next quantify the specific channels that generate these differences across individual takeover markets in our sample.

6.2. Information and the relative performance of auctions and negotiations

In this section, we quantify how pre-entry uncertainty affects the relative performance of auctions and negotiations in the cross-section of takeover markets. As we will see, pre-entry uncertainty operates through three competing channels whose signs and relative magnitudes determine whether a particular target in our sample would expect greater revenue by being sold in an auction or in a negotiated transaction.

The first channel is the *composition effect* in takeover auctions, which reflects the fact that some potential bidders decline to participate based on initially unfavorable pre-entry beliefs, even though some would have discovered high values had they chosen to participate in a takeover competition. Such invited bidders rationally decline to participate based on information available to them at the time, but their absence from the entering bidder pool is suboptimal ex post and degrades competition in auctions, even if takeover markets are efficient in other respects.

Does the composition effect play a quantitatively important role in determining takeover auction revenue? To answer this question, we define the composition effect as the difference between the winning bid and the winning bid that would have resulted had the distributions of potential and entering bidder values been identical. Notice that the distribution of potential and entering bidder values is identical only when potential bidders' decisions to enter are random. As we have shown, this occurs only in the limiting case wherein bidders have no pre-entry information about their values (i.e., when $\alpha = 1$). Our structural estimates identify $\alpha = 0.58$ as the typical case, so in practice some high-value potential bidders decline to participate. Holding constant the size of the entering bidder pool, the composition effect quantifies this effect.

Fig. 1 uses the structural estimates from our sample of takeovers to construct the probability of entry for different potential bidder types. The vertical axis measures the probability that the potential bidder with the Pth highest value enters, relative to an average potential bidder. When $\alpha = 0.3$, the top 10% of potential bidders are about 65% more likely to participate than an average potential bidder. But for takeover markets with greater pre-entry uncertainty, high-value potential bidders are less likely to participate. In light of the composition effect, the high degree of pre-entry uncertainty and limited participation, shown here, could be interpreted as direct evidence that information frictions impair the performance of takeover auctions, as suggested by some scholars.

The second channel through which endogenous entry affects takeover auctions is the *size effect*. It refers to the fact that pre-entry uncertainty encourages entry overall: Possible absence of high-value bidders, due to the composition effect, raises other potential bidders' prospects conditional upon entry, thus incentivizing their participation. But when takeover markets are characterized by imperfect information, some entrants discover higher values for the target upon entry, and their participation raises the overall competitiveness of the entering bidder pool. A larger pool of entering bidders can thus contain more high-value bidders overall, even if a lower fraction of entrants have relatively high values. Formally, we define the size effect for an auction with n entering bidders as the effect of endogenous entry on expected revenue, holding constant the composition of the entering bidder pool. This by definition isolates the effect of pre-entry uncertainty on overall entry, since it holds constant the relative probabilities with which high- and low-value potential bidders enter. By definition takeover revenue is equal to the sum of the size and composition effects.

Fig. 2 empirically quantifies the size effect for targets in our sample using the structural estimates obtained in Section 5. The vertical axis measures the proportion of invited bidders that choose to participate. These estimates reveal that the size effect is quantitatively important, with the average probability of entry for each potential bidder rising from 35% to 60% within the 25th and 75th percentiles of the distribution in pre-entry uncertainty (i.e., between $\alpha = 0.4$ and $\alpha = 0.77$).

The existence of the size effect, overlooked in previous research on takeover auctions, raises the possibility that information frictions could lead to higher takeover auction revenue if the size effect is quantitatively large relative to the composition effect. Though the size effect is unambiguous in sign, theory alone is uninformative about its magnitude, which depends on the full vector of estimated takeover market primitives through the entry equilibrium described in Section 3.

In Table 7, we use the structural estimates to compute the relative contribution of the size and composition effects to takeover auction revenue, for various levels of pre-entry uncertainty and entry costs. Table 7 shows that the composition effect—the negative effect of information frictions on the composition of the entering bidder pool—accounts for less than 10% of takeover auction revenue. This quantitative dominance of the size effect implies that takeover auctions are surprisingly resilient to the presence of information frictions, and even benefit from them, a finding that stands in stark contrast to studies that ignore the size effect, and which presume information frictions always impair takeover auction performance.

The third channel through which pre-entry uncertainty affects the relative performance of auctions and negotiations is the *deterrence effect*, which formalizes the idea that a negotiating bidder might shade up their offer price to deter potential competition (e.g., Fishman, 1988; Aktas, de Bodt, and Roll, 2010; Dimopoulos and Sacchetto, 2014). Formally, we quantify the deterrence effect in a sequential negotiation as the additional amount by which the buyer's bid exceeds the second-highest value among the set of bidders that participated at any point in the negotiation sequence.

Does pre-entry uncertainty raise or lower a target's bargaining power? In other words, does pre-entry uncertainty lead to more or less aggressive deterrence bidding? The answer to this question is theoretically ambiguous. To see why, consider a standing bidder's incentive to shade up its offer price.

The marginal effectiveness of a slightly higher bid depends on the extent to which it deters potential competition, which in turn depends on a potential competitor's decision to enter into negotiations if invited by the target. Pre-entry uncertainty lowers a potential competitor's sensitivity to an observed increase in the posted price offered by a standing bidder, thus lowering the marginal effectiveness of the deterrence bid. Yet, this could incentivize a standing bidder to shade up their offer even higher to mitigate this reduced effectiveness if the incentive to deter entry is strong. Whether this is the case depends on the full vector of estimated takeover market primitives through equilibrium in the sequential negotiation. It is thus an empirical question of whether pre-entry uncertainty leads to lesser or greater deterrence bids.

In Fig. 3, we use the estimated takeover market primitives to quantify the direction and strength of deterrence bidding as a function of variation in pre-entry uncertainty. In a typical takeover negotiation with $\alpha = 0.58$, deterrence bidding accounts for about 13 percentage points of target revenue (i.e., of deal premia). At the same time, takeovers with greater pre-entry uncertainty are associated with weaker deterrence bidding, a finding implying that negotiations tend to perform better in takeover environments with relatively low pre-entry uncertainty.

In Fig. 4, we use the structural estimates to jointly quantify these three channels to show how variation in pre-entry uncertainty across takeover markets affects the relative performance of auctions and negotiations. To do this, we compute the average expected revenue difference between sequential negotiations and auctions (solid line) and between go-shop negotiations and auctions (dashed line) for different levels of pre-entry uncertainty, using the takeover market primitives recovered from the structural estimation. The figure shows that auctions revenue-dominate negotiations when pre-entry uncertainty is high, near the 75th percentile of pre-entry uncertainty across targets in our sample (Table 3). Indeed, the majority of cross-sectional variation in the relative performance of auctions and negotiations in our sample is accounted for by information frictions operating through the size, composition, and deterrence channels just described.[18]

The finding that auctions perform relatively well when pre-entry uncer-

[18]This result was obtained by regressing the target-specific sequential negotiation revenue difference function on polynomials and interactions of the fundamental parameters. We then compute the ratio of R-squares from regressions that omit and include the informational parameters α_j and c_j, which is 0.86.

tainty is high has important implications for the understanding of how firms should be sold. First, having no knowledge about the size effect, some previous researchers studying corporate takeovers presumed that information frictions unambiguously impair auction performance. Our structural estimates reveal that the size effect is quantitatively more important than the composition effect, implying that takeover auctions counterintuitively benefit from pre-entry uncertainty. Second, though other studies have provided evidence for the possible existence of the deterrence effect, it was not previously known how potential competitors' uncertainty affects a target's bargaining power in negotiations. Our structural estimates have shown that the positive effect of deterrence bidding on negotiation performance is likely to be most pronounced in takeover environments with low information frictions.

We also check to see how variation in entry costs interacts with the effect of pre-entry uncertainty. In Fig. 5, we use the structural estimates to quantify the effect of entry costs and pre-entry uncertainty on the relative performance of auctions and negotiations. In the figure, dark circles represent cases in which auctions produce higher prices, while hexagrams represent cases where negotiations produce higher prices. Hollow circles indicate cases in which simulation error exceeds estimated revenue differences. Consistent with the results in Fig. 4, auctions perform relatively well when pre-entry uncertainty is high, while Fig. 5 reveals that higher entry costs raise the threshold level of pre-entry uncertainty at which auctions revenue-dominate negotiations.

6.3 Target observables and the relative optimality of auctions and negotiations

Up to this point we have shown how, through the size, composition, and deterrence effects, pre-entry uncertainty affects the relative performance of auctions and negotiations. The findings reported in Table 5 reveal that variation in pre-entry uncertainty across takeover markets is predictable by observable target firm characteristics. In this section, we examine how the relative performance of auctions and negotiations varies with target observables. To do this, we sort targets in our sample based on each characteristic and use the estimated fundamental parameters to determine the expected frequency with which each target expects to obtain relatively higher revenue from holding an auction.

Panel A of Table 8 reports box plots showing how the relative performance of auctions and negotiations varies across Fama-French industries. Variation in the relative performance of auctions and negotiations across industries is consistent with variation in pre-entry uncertainty across industries, shown in Section 5. For example, Panel A of Table 8 shows that the business equipment sector has the lowest fraction of targets for which auctions outperform negotiations, consistent with the low degree of pre-entry uncertainty in this industry, shown in Table 5. Conversely, the estimates in Table 5 uncovered high average pre-entry uncertainty among bidders for targets headquartered in consumer durables and nondurables, and Panel A of Table 8 shows that these industries have the highest proportion of targets for which auctions outperform negotiations.

Turning to target financials, Panel B of Table 8 reports the fraction of targets for which auctions are preferable, for the bottom and top quantiles of targets in our sample. Auctions are more likely to be preferred when leverage is high, cash holdings are low, and the q-ratio is high, consistent with the associations between these characteristics and pre-entry uncertainty, shown above.

Our findings about the relationship between target observables and the relative performance of auctions and negotiations differ from earlier research that did not account for endogenous entry and pre-entry uncertainty when asking how a firm should be sold. For example, Fidrmuc et al. (2012) do not find a relationship between q and targets' propensities to sell via auction, while our results show that pre-entry uncertainty falls with q, incentivizing sale via a negotiated transaction. Another difference is our finding that pre-entry uncertainty, and thus the relative performance of auctions, is positively associated with target leverage, as should be expected given that higher leverage can complicate a potential bidder's task of valuing assets. With respect to target size, Boone and Mulherin (2007) do not find a relationship between target size and the relative market reaction to announced auctions and negotiations. By focusing on winning bids and accounting for endogenous entry, we show that larger targets more frequently prefer auctions.

Our analysis also differs from work attempting to recover insights about optimal sale procedures based on ex post deal outcomes. For example, Boone and Mulherin (2007) attempt to explain a target's choice of sale mechanism by whether the winning bidder paid in cash. But studies in the corporate finance literature have shown that method of payment depends on the winning bidder's characteristics (e.g., Travlos, 1987; Eckbo, Giammarino, Heinkel, 1990; Gorbenko and Malenko, 2017). Our analysis shows that the identity of the winning bidder—and thus its characteristics—are an outcome of the entry process itself. In our structural framework, these ex post variables, such as cash payment, the initiator of a deal, or the relative size of the target and acquirer, lose their predictive ability once we condition on the other observables that enter our analysis.[19]

More generally, by using a structural model to directly quantify pre-entry uncertainty and its effects on the relative performance of auctions and negotiations, we build on earlier work that found similar shareholder returns across auctions and negotiations in the aggregate, but that did not empirically characterize cross-sectional variation in information frictions or show how they influence the relative optimality of auctions and negotiations for individual targets, as we do here. At the same time, our finding of large variation in pre-entry uncertainty across targets, only some of which is predicted by ex ante target characteristics, invites more research to verify whether certain shareholder empowerment movements should press for more stringent interpretation of statutes that could be construed as requiring that target boards conduct auction-style processes (Bebchuck, 1982; Revlon v. McAndrews and Forbes Holdings, 1986; Cramton and Schwartz, 1991).

Indeed, our structural estimates, identified based on information about entry patterns and bids, reveals significant cross-sectional variation across targets in takeover market primitives and in the relative performance of auctions and negotiations. At the same time, results reported in this section reveal that some portion of this total variation is explained by ex ante observable target characteristics, thus opening the door to further research that might explore alternate pathways to systematically predicting optimal sale procedures based on information that might be available to researchers before a sale process is initiated.

7. Robustness

[19] To show this, we regressed the model-predicted target-specific relative returns to auctions on three ex post deal outcomes: a dummy variable for whether the winning bidder paid in cash, a dummy variable for whether the deal was a tender offer, a dummy variable for whether the deal was unsolicited, and the relative size of the winning bidder and target. Three of the four variables are highly statistically significant in this regression, but the explanatory power of all four ex post variables vanishes when we control for the ex ante observable target characteristics from Table 8. This finding echoes other structural studies of takeover auctions that caution against the use of simple associations between observed deal outcomes and such variables without characterizing the endogenous behavior of actors whose behavior leads a researcher to observe them in the first place.

In this section, we provide several checks on our findings. First, we check to see whether the model performs well at predicting takeover auction patterns. The model predicts that the average number of invited bidders that choose to enter is 5.33, similar to the corresponding figure of 5.84 in Table 2. The model predicts that takeover revenue is 41.18% on average, similar to the 42.8% average revenue in our sample, shown in Table 2. The model thus performs reasonably well at predicting entry patterns and winning bids in takeover auctions.

Second, we check the sensitivity of our results to the format of the bidding game. Here, we assume that a subset of entering bidders with values above the target's reservation value compete in a first-price auction for the target.[20] Specifically, bidding among entering bidders starts from a seller-announced initial adequate price v_0 and proceeds via an ascending auction until the auction reaches a price $\overline{p} \geq v_0$ at which at most k bidders remain. Note that our baseline ascending auction obtains when $k = 1$. From this point, the remaining bidders compete in a first-price sealed bid auction with the highest bidder in this final round winning the target and paying its bid. However, the form of the bidding equilibrium will change, and this could in principle influence our interpretation of the data. Details on the bidding equilibrium are provided in the online internet appendix. The estimates are reported in Column 1 (for $k = 2$) and Column 2 (for $k = 3$) of online internet appendix Table A.1. The estimates are similar to the main estimates in Table 3.

Third, we examine the model's sensitivity to the assumption that potential bidders know N_j when deciding whether to participate in an auction. This assumption was approximated by the view that potential bidders or their advisors use knowledge about the target and its industry to assess the size of the potential bidder pool. We implement one such approximation in which each potential bidder's belief is based on a forecast of N_j based on target observables, which we construct using the predicted value from OLS regressions of N_j on the vector X_j.[21] Estimates obtained using this predicted

[20] Gorbenko and Malenko (2014) implement a technique developed by Haile and Tamer (2003) that allows them to estimate parameters in an ascending auction model without committing to a specific bidding format. Our focus on endogenous entry necessitates the use of a general model of entry behavior in which parameters cannot be recovered using the Haile and Tamer (2003) technique.

[21] This procedure is intentionally based on the crude assumption that the only informa-

value, reported in Column 3 of the online internet appendix Table A.1, are similar to the baseline estimates in Table 3.

Fourth, we re-estimate all parameters under alternative assumptions on the target's reservation value v_0. Specifically, we assume that v_0 is drawn from an exponential distribution with target-specific mean μ_{0j}, with μ_{0j} drawn independently from a Gamma distribution with mean and variance estimated by empirical model. Estimates obtained using this model are reported in Column 4 of the online internet appendix Table A.1. While these estimates are similar to the baseline estimates in Table 3, they imply a price distribution with a mean sale premium of 47.03%, which departs substantially more from both the data (42.80%) and the estimates implied by our baseline reserve specification (41.18%). This finding both bolsters confidence that our results are not driven by modeling assumptions about the unobserved target reservation value and underscores our empirical preference for the baseline model.

Fifth, we check the sensitivity of our results to the sample restriction that excludes deals with winning bids below the target's share price four weeks prior to announcement. The results, reported in Column 5 of the online internet appendix Table A.1, are similar to the baseline estimates reported in Table 3.

8. Concluding remarks

A large and well-composed pool of entering bidders creates competition and high takeover prices, yet in practice fewer than half of invited bidders choose to participate in takeover auctions. This paper estimates an empirical structural model that characterizes this endogenous entry, thus permitting

tion available to a potential bidder are the X_j and information about N_j from previous deals. To see how such possible uncertainty maps into model parameters, consider a hypothetical increase in a potential bidder's uncertainty about the number of invited potential bidders. Conditional on a given signal, uncertainty in N_j maps into uncertainty about the number of competing entrants, and thus about expected post-entry profits. Hence, all else equal, an increase in uncertainty about N_j will tend to appear in the model as an increase in α. Since in practice potential bidders and their advisers have access to information about the target not contained in X_j, this procedure provides a plausible lower bound on potential bidder's degree of knowledge about N_j. Since they are constructed based on the assumption that potential bidders have access to very little information about the target, the estimate α reported in online internet appendix Table A.1 can be interpreted as an upper bound on average pre-entry uncertainty.

recovery of takeover market characteristics that determine whether auctions or negotiations yield higher prices for individual target firms. The estimates reveal that potential bidders' pre-entry uncertainty is a quantitatively important feature of takeover markets, with pre-entry beliefs containing more noise than information on average. At the same time, our analysis reveals that pre-entry uncertainty varies dramatically across takeover markets, while also differently affecting auctions and negotiations. Contrary to what some scholars have proposed, information frictions can enhance the performance of auctions, even while reducing a target's bargaining power in negotiated transactions. Together, our findings imply that negotiations perform relatively well in takeover markets with low information frictions.

Our analysis also uncovers associations between pre-entry uncertainty and firm-level variables such as size and cash holdings. Future research might extend this line of exploration by investigating how and to what extent information frictions, endogenous entry patterns, and the relative performance of auctions and negotiations can be explained by firm- and market-level characteristics observable to researchers before a sale process begins.

Our study has compared takeover revenue resulting from a target's sale by auction with revenue from the same target's decision to be sold by a negotiated transaction. In doing so, our analysis places auctions and negotiations on an even playing field. An alternate line of argument in the debate about how firms should be sold hypothesizes the possible existence of non-price procedure-specific costs or benefits, for example, a target board's desire to quickly sell the company. Such hypothesized costs and benefits might cause a sale procedure to be preferred even if it produces lower expected revenue. To date, these lines of argument lack systematic evidence about the quantitative magnitudes of such possible intangible costs and benefits, a fact implying that any sale procedure might in principle be justified ex post by making recourse to the possible existence of such sale-specific features. Future research might thus add such non-price frictions posited in corporate finance literature to our model, for example, a cost of delay (e.g., Subramanian, 2008), or a cost of approaching additional bidders (e.g., Hansen, 2001). In the same sense that we have integrated the size, composition, and deterrence effects into a single structural model of endogenous participation, quantifying the relative magnitudes of these channels on an even playing field, future research might incorporate additional frictions, thus permitting their relative magnitudes to be quantified and directly compared with the effects studied here.

Appendix

A. Entry Equilibrium

First, we formally obtain an expression for a particular bidder's expected profits conditional on entry. All variables are sale specific, so in this section we suppress both j subscripts and the target-level primitives θ_j to ease exposition. All derivations below should be interpreted as conditional on a particular realization of target-level primitives θ_j.

Let s_N^* denote the equilibrium threshold characterizing entry behavior among the N potential bidders for target j, $F^*(\cdot|N) \equiv F(\cdot|S_i \geq s_N^*)$ be the c.d.f. characterizing the selected distribution of valuations for each of the n bidders electing to enter, and $F_0(\cdot)$ be the distribution of the target's reservation value V_0. Let $Y_{k:n}$ denote the kth highest valuation among n entering bidders, let $y_{k:n}$ denote the realization of this random variable, and let v_0 denote the realization of the target's reservation value V_0. If $y_{1:n} \geq v_0$, the target is sold at $p = \max\{y_{2:n}, v_0\}$ so conditional on realizations of all random variables, the surplus of bidder with valuation v_i is thus

$$\begin{aligned} \mathbf{1}[v_i &\geq \max\{y_{1:n-1}, v_0\}](v_i - p) \\ &= \mathbf{1}[v_i \geq \max\{y_{1:n-1}, v_0\}](v_i - \max\{y_{1:n-1}, v_0\}). \end{aligned} \tag{A.1}$$

Let $H_n^*(\cdot|N)$ be the equilibrium CDF of the random variable $\max\{Y_{1:n-1}, V_0\}$:

$$H_n^*(y|N) = F_0(y) \cdot F^*(y|N)^{n-1}. \tag{A.2}$$

By definition, $H_n^*(v|N)$ is the probability that a bidder with valuation v is the final standing bidder, with the associated density

$$h_n^*(v|N) = f_0(v) \cdot F^*(v|N)^{n-1} + (n-1)F_0(v)F^*(v|N)^{n-1}f^*(v|N), \tag{A.3}$$

describing the distribution of the bidder's outside option in this case, so the expected profit of an entrant with valuation v_i is thus

$$\begin{aligned} \pi^*(v_i; n, N) &= H_n^*(v_i|N) \int_0^{v_i} (v_i - y) \cdot \frac{h_n^*(y|N)}{H_n^*(y|N)} dy \\ &= \left[v_i H_n^*(v_i|N) - \int_0^{v_i} y \; h_n^*(y|N) \; dy \right] \\ &= \int_0^{v_i} F_0(y) \cdot F^*(y|N)^{n-1} \; dy, \end{aligned} \tag{A.4}$$

where the last equality follows from integration by parts.

Having obtained an expression for an entering bidder's expected profits, we now characterize the symmetric monotone pure strategy Bayesian Nash equilibrium. In any such equilibrium, entry decisions can be characterized by a signal threshold s_N^* such that bidder i chooses to enter if and only if $S_i \geq s_N^*$:

$$F^*(v; s_N^*) = F(v|S_i \geq s_N^*) = \frac{1}{1 - F_s(s_N^*)} \int_{s_N^*}^{\infty} F(v|t)\ f_s(t)\ dt, \tag{A.5}$$

The CDF of the distributions of valuations among entrants is then $F^*(v|N) = F^*(v; {s^*}_N)$. The following identity will be useful: for any (v, s^*),

$$(1 - F_s(s^*))\, F^*(v; s^*) = \int_{s^*}^{\infty} F(v|t)\ f_s(t)\ dt = F_v(v) - F_{vs}(v, s^*). \tag{A.6}$$

Independence of signals implies that the total number of entrants n follows a binomial distribution based on the entry probability $[1 - F_s(s_N^*)]$. Now consider the entry decision of potential acquirer i drawing signal realization $S_i = s_i$. Conditional on own signal s_i, the equilibrium threshold s_N^*, and total competition N, a potential bidder forecasts profits $\Pi(s_i; s_N^*, N)$. Expanding this term yields,

$$\begin{aligned}
&= E_V\left[E_n[\pi^*(v_i; n, N)|n \geq 1]|S_i = s_i\right] && \text{(A.7)}\\
&= \int_0^\infty \int_0^v F_0(y) \left[\sum_{n=1}^{N} \begin{pmatrix} N-1 \\ n-1 \end{pmatrix} F_s(s_N^*)^{N-n} \left([1 - F_s(s_N^*)] F^*(y; s_N^*)\right)^{n-1}\right] dy\ dF(v|s_i)\\
&= \int_0^\infty \int_0^v \left[F_0(y)\left[F_s(s_N^*) + (1 - F_s(s_N^*)) F^*(y; s_N^*)\right]^{N-1}\right] dy\ dF(v|s_i)\\
&= \int_0^\infty \int_0^v \left[F_0(y)\left[F_s(s_N^*) + F_v(y) - F(y, s_N^*)\right]^{N-1}\right] dy\ dF(v|s_i),
\end{aligned}$$

where the third equality follows by properties of binomial series.

Reversing the order of integration yields our main expression for *ex ante* expected profit for potential acquirer with Stage 1 signal $S_i = s_i$:

$$\Pi(s_i; s_N^*, N) = \int_0^\infty [1 - F(v|s_i)] \cdot F_0(y) \cdot [F_s(s_N^*) + F_v(y) - F(y, s_N^*)]^{N-1}\, dy. \tag{A.8}$$

$F(v|s_i)$ is decreasing in s_i by stochastic dominance, and $F_s(s_N^*) + F_v(y) - F(y, s_N^*)$ is increasing in s_N^* by the identity

$$F_s(s_N^*) + F_v(y) - F(y, s_N^*) = F_s(s_N^*) + \int_{s^*}^{\bar{s}} F(v|t)\ f_s(t)\ dt \tag{A.9}$$

and it is easy to show that $F_s(s_N^*) + F_v(y) - F(y, s_N^*) \in [0, 1]$.

We now characterize equilibrium entry decisions. Bidder i enters into competitive bidding if expected profit from doing so is positive, so the equilibrium threshold s_N^* must thus satisfy the break even condition:

$$\Pi(s_N^*; s_N^*, N) - c = 0. \tag{A.10}$$

In other words, a marginal potential bidder with signal $S_i = s_N^*$ must be indifferent between entering and not entering. $\Pi(s_i; s_N^*, N)$ is increasing in its first argument and is strictly increasing in its second argument, so the break even condition (A.10) has a unique solution s_N^*. Further, since $\Pi(s_i; s_N^*, N)$ is decreasing in N, this solution s_N^* is increasing in N. Finally, by the form of the entry decision rule, the distribution of valuations among entering bidders is $F^*(v; s_N^*) = F(v|S_i \geq s^*)$. The signal threshold s_N^* is thus sufficient to characterize equilibrium entry behavior.

B. Estimation Algorithm

Recall the objective of our structural estimation procedure: to recover the deep structural parameters Γ_0 governing the distribution $g(\theta|X_j, \Gamma_0)$ of the target-level characteristics θ_j, accounting for the facts that (1) individual realizations of θ_j are unobserved to the econometrician, and (2) we observe only auctions resulting in sale. Toward this end, we consider maximum likelihood estimation based on events of the following form: for a given target j, n_j of bidders enter and the final sale price is $P_j = p_j$, conditional on target j inviting N_j potential bidders and holding an auction which results in sale. Integrating over unobserved target-level characteristics θ_j, we thereby obtain a target-level structural conditional likelihood function (in Γ) of the form:

$$L_j(p_j, n_j|sale_j = 1, N_j, X_j; \Gamma) =$$
$$\int f_p(p_j|sale_j, n_j, N_j; \theta_j) \cdot \Pr(sale_j|n_j, N_j; \theta_j) \cdot \Pr(n_j|N_j; \theta_j)\ g(\theta_j|X_j, \Gamma)\ d\theta_j$$
$$\Big/ \int \Pr(sale_j|n_j, N_j; \theta_j) \cdot \Pr(n_j|N_j; \theta_j)\ g(\theta_j|X_j, \Gamma)\ d\theta_j, \tag{6}$$

In describing our procedure for estimating Γ based on equation (6), we proceed in four steps. First, we describe how we solve for the equilibrium entry threshold $s^*(N_j, \theta_j)$, which is the key prediction of the structural entry and bidding model above. Second, we discuss computation of the equilibrium objects $\Pr(n_t|N_t;\theta_j)$, $\Pr(sale_t|N_t;\theta_j)$, and $f_p(p_t|sale_t, n_t, N_t;\theta_j)$ appearing in (6). Third, we describe the importance sampling procedure by which we evaluate the integrals over θ_j appearing in the numerator and denominator of (6). Finally, we discuss the Markov Chain Monte Carlo algorithm by which we maximize (6) to recover point estimates and confidence intervals for $\hat{\Gamma}$.

B.1. Solving for the equilibrium entry threshold

Consider a takeover auction among N_j potential bidders competing for a target j with characteristics θ_j (observed to bidders, but unobserved to us). Let $s^*(N_j;\theta_j)$ denote the signal threshold characterizing equilibrium entry behavior in this takeover environment. From Equation (A.10), we know that we may compute this threshold $s^*(N_j;\theta_j)$ as the unique solution s^* to the breakeven condition:

$$c(\theta_j) = \int_0^\infty [1 - F_{v|s}(y|s^*;\theta_j)] \cdot F_0(y;\theta_j) \cdot H_{N_j}(y;\theta_j)dy, \tag{B.1}$$

where $H_{N_j}(\cdot; s^*, \theta_j)$ denotes the expected c.d.f. of the maximum valuation realized (through entry) among the $N_j - 1$ potential rivals of each bidder i, accounting for the fact that some of these rivals may not enter in equilibrium:

$$H_{N_j}(y; s^*, \theta_j) \equiv [F_s(s^*;\theta_j) + F_v(y;\theta_j) - F_{vs}(y, s^*;\theta_j)]^{N_j-1}.$$

Taking $\theta_j \equiv (\mu_{vj}, \sigma^2_{vj}, c_j, \alpha_j)$ as given, (V_j, S_j) are jointly log-normal with mean vector $[\mu_{vj}, \mu_{vj}]$ and variance-covariance matrix

$$\text{Var}\left(\begin{bmatrix} V_j \\ S_j \end{bmatrix}\right) = \begin{bmatrix} \sigma^2_{vj} & \sigma^2_{vj} \\ \sigma^2_{vj} & \frac{\sigma^2_{vj}}{1-\alpha_j} \end{bmatrix}.$$

Meanwhile, the conditional distribution of V_j given $S_j = s^*$ is normal with mean $\alpha_j\mu_{vj} + (1-\alpha_j)s^*$ and variance $\alpha_j\sigma^2_{vj}$. Given θ_j, computation of both $H_{N_j}(y; s^*;\theta_j)$ and $F_{v|s}(y|s^*)$ is therefore straightforward. For given θ_j and N_j, we may therefore solve (B.1) numerically to obtain the equilibrium entry threshold $s^*(N_j;\theta_j)$. In practice, we approximate this solution by interpolation over a fine grid in $\log s^*$, computing the right-hand side integral by the trapezoidal rule over a fine grid in $\log v$.[22]

B.2. Computing equilibrium objects in the likelihood function

With the equilibrium entry threshold $s^*(N_j;\theta_j)$ in hand, we turn to computation of the equilibrium objects $\Pr(n_t|N_t;\theta_j)$, $\Pr(sale_t|N_t;\theta_j)$, and $f_p(p_t|sale_t, n_t, N_t;\theta_j)$ appearing in (6).

Toward this end, first consider $\Pr(n_t|N_t;\theta_j)$. By construction, potential acquirers drawing signals $S_{ij} \geq s^*(N_j,\theta_j)$ elect to enter in equilibrium. The probability that any given potential acquirer elects to enter is therefore

$$q(N_j,\theta_j) = 1 - F_s(s^*(N_j,\theta_j);\theta_j). \tag{B.2}$$

Furthermore, conditional on θ_j, signal draws are independent across potential aquirers. Taking N_j and θ_j as given, the distribution of n_j therefore follows a binomial distribution with success probability $q(N_j,\theta_j)$:

$$\Pr(n_j|N_j,\theta_j) = \begin{pmatrix} N_j \\ n_j \end{pmatrix} q(N_j,\theta_j)^{n_j}(1-q(N_j,\theta_j))^{N_j-n_j}. \tag{B.3}$$

Next consider $\Pr(sale_j|N_j;\theta_j)$. By construction, the auction for target j ends in sale whenever at least one entering bidder draws a valuation above the seller's reservation value V_{0j}. It follows that:

$$\begin{aligned}\Pr(\text{sale}_j|N_j;\theta_j) &= \Pr(V_{0j} \leq Y_{1:N_j}|N_j,\theta_j) = 1 - \Pr(Y_{1:N_j} \leq V_{0j}|N_j,\theta_j) \\ &= 1 - \int_0^\infty \left[F_s(N_j;\theta_j) + F_v(v_0;\theta_j) - F_{vs}(v_0, s^*(N_j,\theta_j);\theta_j)\right]^{N_j} \; f_0(v_0,\theta_j)\, dv_0,\end{aligned} \tag{B.4}$$

where (as above) the term in brackets represents the probability that potential acquirer i either does not enter or enters but draws a valuation less than v_0. As above, taking N_j and θ_j as given, the right-hand side integral is straightforward to compute, yielding a numeric solution for $\Pr(\text{sale}_j|N_j;\theta_j)$.[23]

[22] In practice, we consider grids of 200 points in both $\log s$ and $\log v$, with grid support between the 10^{-6}th and $(1-10^{-6})$th quantiles of $\log S$ and $\log V$. Numerical simulations confirm that the resulting solution is quite accurate in practice.

[23] In practice, as in computing $s^*(N_j;\theta_j)$ above, we approximate this integral via the trapezoidal rule on a grid of 200 points in $\log v_0$, with grid points spaced evenly (in $\log v_0$) between the 10^{-6}th and $(1-10^{-6})$th quantiles of $\log V_0$.

Finally, consider the density $f_p(p_j|sale_j, n_j, N_j; \theta_j)$; i.e. the distribution of the sale premium p_j when target characteristics are θ_j, n_j of N_j potential bidders enter, and the auction results in sale. In characterizing this distribution, we adopt the following convention: through expressions such as

$$\Pr(sale_j \cap P_j = p|n_j, N_j; \theta_j) \tag{B.5}$$

we intend to indicate to the mixed joint density of the discrete random variable $sale_j$ and the continuous random variable P_j; i.e. more precisely,

$$\Pr(sale_j \cap P_j = p|n_j, N_j; \theta_j) := \lim_{h \downarrow 0} \frac{Pr(sale_j \cap P_j \in [p, p+h]|n_j, N_j; \theta_j)}{h}. \tag{B.6}$$

Applying this convention, we have by construction:

$$f_p(p_j|sale_j, n_j, N_j; \theta_j) = \frac{\Pr(sale_j \cap P_j = p_j|n_j, N_j; \theta_j)}{\Pr(sale_j|n_j, N_j; \theta_j)}. \tag{B.7}$$

$\Pr(sale_j|n_j, N_j; \theta_j)$ was characterized above, so having obtained an expression for $\Pr(sale_j \cap P_j = p_j|n_j, N_j; \theta_j)$ the argument will be complete.

By construction, a sale occurs when at least one entrant draws a valuation above the seller's reservation value v_{0j}. If only one entrant draws a valuation above v_{0j}, the transaction price p_j is the seller's reservation valuation v_{0j}. If at least two entrants draw valuations above v_{0j}, the transaction price p_j is the second highest entrant valuation $y_{2:n_j}$. Decomposing likelihoods of these events using properties of order statistics yields the overall mixed density $\Pr(sale_j \cap P_j = p_j|n_j, N_j; \theta_j)$

$$\begin{aligned}
&= \Pr(sale_j \cap Y_{2:n_j} = p_j|n_j, N_j, \theta_j) + \Pr(sale_j \cap V_{0j} = p_j|n_j, N_j, \theta_j) \\
&= \Pr(Y_{1:n_j} \geq p_j \cap Y_{2:n_j} = p_j \cap V_{0j} \leq p_j|n_j, N_j, \theta_j) \\
&\quad + \Pr(Y_{1:n_j} \geq p_j \cap Y_{2:n_j} \leq p_j \cap V_{0j} = p_j|n_j, N_j, \theta_j) \\
&= \left[n_j(n_j-1)F^*(p_j; N_j, \theta_j)^{n_j-2}[1 - F^*(p_j; N_j, \theta_j)]f^*(p_j; N_j, \theta_j)\right] \cdot F_0(p; \theta_j) \\
&\quad + \left[n_j F^*(p_j; N_j, \theta_j)^{n_j-1}[1 - F^*(p_j; N_j, \theta_j)]\right] \cdot f_0(p_j; \theta_j).
\end{aligned}$$

where as above $F^*(v|N_j, \theta_j)$ denotes the equilibrium distribution of valuations among entrants at (N_j, θ_j):

$$F^*(v|N_j, \theta_j) = F(v|S_i \geq s^*(N_j, \theta_j)) = \frac{F_v(v; \theta_j) - F_{vs}(v, s^*(N_j, \theta_j); \theta_j)}{1 - F_s(s^*(N_j; \theta_j); \theta_j)} \tag{B.8}$$

Again, taking θ_j and $s^*(N_j;\theta_j)$ as given, the distribution (B.8) can easily be computed as above. Having solved for the equilibrium entry threshold $s^*(N_j;\theta_j)$, computation of the equilibrium objects $\Pr(n_t|N_t;\theta_j)$, $\Pr(sale_t|N_t;\theta_j)$, and $f_p(p_t|sale_t,n_t,N_t;\theta_j)$ for given N_j, θ_j thus becomes a reasonably straightforward numerical exercise.

B.3. Importance sampling approach to integration over θ_j

Give numeric expressions for the equilibrium objects $\Pr(p_j, sale_j|n_j, N_j;\theta)$, $\Pr(n_j|N_j,\theta)$, and $\Pr(sale_j|N_j;\theta)$ obtained as above, we can in principle evaluate the likelihood (6) directly by computing the numerator and denominator integrals

$$\int \Pr(p_j, sale_j|n_j, N_j;\theta)\Pr(n_j|N_j,\theta)\ g(\theta|X_j,\Gamma)\ d\theta \tag{B.9}$$

and

$$\int \Pr(sale_j|N_j;\theta)\ g(\theta|X_j,\Gamma)\ d\theta. \tag{B.10}$$

Direct evaluation of the likelihood function is computationally prohibitive in practice since (B.9) and (B.10) depend on θ through the equilibrium condition (B.1), which itself requires solution of an equation involving integrals. We circumvent this challenge by implementing estimation via the simulated likelihood method of Ackerberg (2009)[24], which uses the principle of importance sampling to transform the complicated problem of repeated evaluation of the full likelihood into the much simpler problem of repeated evaluation of $g(\theta|X_j,\Gamma)$.

To illustrate the main idea of this method, let $\tilde{g}(\cdot)$ be any fixed *proposal distribution* over θ, and consider evaluation of the sale-level likelihood integral (B.9). By standard importance sampling arguments, we can rewrite this integral as follows:

$$\begin{aligned} & \int \Pr(p_j, sale_j|n_j, N_j;\theta)\Pr(n_j|N_j,\theta)\ g(\theta|X_j,\Gamma)\ d\theta && \text{(B.11)} \\ = & \int \left[\Pr(p_j, \mathrm{sale}_j|n_j, N_j;\theta)\Pr(n_j|N_j,\theta)\ \frac{g(\theta|X_j,\Gamma)}{\tilde{g}(\theta)}\right]\tilde{g}(\theta)\ d\theta \\ = & \ \tilde{E}\left[\Pr(p_j, \mathrm{sale}_j|n_j, N_j;\theta)\Pr(n_j|N_j,\theta)\ \frac{g(\theta|X_j,\Gamma)}{\tilde{g}(\theta)}\right], \end{aligned}$$

[24]Ackerberg, D., 2009. A new use of importance sampling to reduce computational burden in simulation estimation, *Quantitative Marketing and Economics* 7, 343-376.

where the expectation in the last line is taken with respect to the proposal distribution $\tilde{g}(\cdot)$ rather than the true distribution $g(\cdot|X_j, \Gamma)$. If $\left\{\tilde{\theta}_r\right\}_{r=1}^{R}$ is a random sample drawn from $\tilde{g}(\cdot)$, it follows that for large enough R

$$\begin{aligned} & \int \Pr\left(p_j, sale_j | n_j, N_j; \theta\right) \Pr\left(n_j | N_j, \theta\right) \, g(\theta | X_j, \Gamma) \, d\theta \qquad \text{(B.12)} \\ \approx & \sum_{r=1}^{R} \Pr\left(p_j, \text{sale}_j | n_j, N_j; \theta_r\right) \Pr\left(n_j | N_j, \theta_r\right) \, \frac{g(\theta_r | X_j, \Gamma)}{\tilde{g}(\theta_r)}. \end{aligned}$$

If a new sample $\left\{\tilde{\theta}_r\right\}_{r=1}^{R}$ is drawn each time the integral (B.9) is evaluated, this importance sampling procedure will of course do nothing to simplify computation. Note, however, that the parameters Γ now appear *only* in the distribution $g(\theta_r|X_j, \Gamma)$, which itself only affects weights on elements in a sum. This in turn motivates Ackerberg (2009)'s reinterpretation of importance sampling.

Specifically, rather than drawing $\{\tilde{\theta}_r\}_{r=1}^{R}$ anew each time (B.9) is evaluated, Ackerberg (2009) propose to draw a single large sample $\{\tilde{\theta}_r\}_{r=1}^{R}$ from $\tilde{g}(\cdot)$ at the beginning of the algorithm. Holding this sample $\{\tilde{\theta}_r\}_{r=1}^{R}$ fixed, we may then calculate the integrand elements $\Pr(p_j, \text{sale}_j|n_j, N_j; \theta_r)$, $\Pr(n_j|N_j, \theta_r)$, and $\Pr(\text{sale}_j|N_j; \theta_r)$ for each θ_r *once for all* prior to estimation. Holding these pre-computed objects fixed, computation of the importance-sampling approximation (B.12) to the integral (C.1) at different values of Γ requires only recalculation of the importance sampling weights $g(\theta_j|X_j, \Gamma)$. As costs of computing $g(\theta_r|X_j, \Gamma)$ are trivial relative to costs of recomputing equilibrium, this allows for vastly accelerated estimation even net of higher setup costs, with the added advantage that the simulated likelihood function is automatically smooth in Γ. For our purposes, therefore, Ackerberg (2009) simulation is ideal; it mitigates the computational infeasibility that otherwise would be entailed by accommodating sample selection unobserved heterogeneity.

In practice, we implement this importance sampling procedure in two steps as follows. As a first pass, we draw a candidate importance sample $\{\theta_r\}_{r=1}^{R}$ of size $R = 10000$ for each target j from a multivariate uniform distribution over the following intervals: for $\mu_{vj}^r \sim U[-0.5192, 1.1825]$ (corresponding to 4 standard deviations of price above and below the mean), $\sigma_{vj}^r \sim 10^{-6} + U[0, 0.5]$, $c_j^r \sim U[0, 0.1]$, and $\alpha_j^r \sim U[0, 1]$. We then maximize the log-likelihood to obtain a first-step estimate $\hat{\Gamma}_0$ for Γ_0, and re-draw a new importance sample $\{\theta_r\}_{r=1}^R$ (also of size 10,000) for each auction j from the resulting predicted distribution $g(\cdot|X_j, \hat{\Gamma}_0)$. Finally, we maximize the simulated likelihood implied by this more accurate importance sample to obtain our final estimator $\hat{\Gamma}$ for Γ_0.

B.4. Inference: Pseudeo-Bayesian Markov Chain Monte Carlo Algorithm

In view of the importance sampling algorithm above, a variety of algorithms are feasible to maximize the (simulated analogue to) the log-likelihood (6). In practice, however, we focus on a Markov Chain Monte Carlo procedure in the spirit of Chernozhukov and Hong (2003).[25] Specifically, starting from a given initial point Γ^0, we use a Markov Chain Monte Carlo (MCMC) algorithm to obtain a sample $\{\Gamma^k\}_{k=1}^K$ of parameters from the conditional likelihood $\prod_j L(p_j, n_j|sale_j, N_j, X_j; \cdot)$, interpreted as (proportional to) the Bayesian posterior over Γ induced by the observed sample $\{(p_j, n_j|N_j)\}_{j=1}^J$ in conjunction with a flat (uninformative irregular) prior over Γ. We then consider the resulting posterior sample $\{\Gamma^k\}_{k=1}^K$ as a basis for inference on Γ_0.

While not strictly necessary, this pseudo-Bayesian MCMC approach has several practical advantages. First, MCMC is a global maximization algorithm, which will eventually trace out the entire posterior of Γ (i.e. the entire likelihood function) from any initial point Γ^0. Second, under classical maximum likelihood regularity conditions, the Bernstein-Von Mises Theorem implies that (for any prior) any of the mean, median, and mode of $\{\Gamma^k\}_{k=1}^K$ will be asymptotically equivalent to the classical maximum likelihood estimator $\hat{\Gamma}$, with consistent classical standard errors provided by the standard deviation of $\{\Gamma^k\}_{k=1}^K$.[26] Third, even when classical maximum likelihood regularity conditions fail — for instance, in models which are only set, not point, identified — the posterior distribution traced out by the MCMC sample $\{\Gamma^k\}_{k=1}^K$ will still permit exact finite-sample Bayesian inference on Γ_0, leveraging only the information on Γ_0 revealed by the data. Finally, given the posterior sample $\{\Gamma^k\}_{k=1}^K$, conducting inference on any function $f(\Gamma)$ of Γ_0 is also straightforward — one need simply compute relevant quantiles of $f(\Gamma^k)$ over a (subsample of) $k = 1, ..., K$.

[25] Chernozhukov, V. and H. Hong, 2003. "An MCMC approach to classical estimation," *Journal of Econometrics* 115:2, 293-346.

[26] See, e.g., Chernozhukov and Hong (2003) for one discussion of this property.

In practice, we implement MCMC at each iteration using a block Metropolis-Hastings algorithm, with four parameter blocks corresponding to the parameters governing the distributions of μ_{vj}, σ_{vj}, c_j, and α_j, and a multivariate normal proposal distribution for each block. We begin with a burnin phase of $K = 20,000$ iterations, using acceptance-rejection probabilities over this burnin phase to tune the proposal variances for each of our four parameter blocks. We then run another $K = 60,000$ MCMC iterations, taking the resulting posterior sample $\{\Gamma^k\}_{k=1}^{60,000}$ as a basis for inference on Γ_0 as described above. The resulting MCMC appears to have good mixing properties, converging quickly (within 5000 burnin iterations) to the neighborhood of the maximum likelihood even for relatively high-dimensional Γ. The high numbers of burnin and regular iterations above therefore primarily reflect an abundance of caution – in practice, much shorter chain lengths give virtually identical results.

C. The Negotiation Mechanisms

As described above, our counterfactual experiment of primary interest concerns an N-round sequential negotiation mechanism. This mechanism proceeds as follows. Round $n = 1, ..., N$ begins when the target approaches potential buyer $i = 1, ..., N$ (ordered at random) with an invitation to participate. The following events then take place in round n:

1. Potential buyer i observes its private signal S_{ij} for the target. Based on this signal S_{ij} and the entry and bidding history up to round $n-1$, potential bidder i determines whether to enter the negotiation at cost c.
2. Conditional on choosing to enter, potential buyer i learns its valuation V_{ij}. If another negotiating bidder has previously entered, potential buyer i and the current incumbent compete in an ascending button auction for the right to remain in the auction. The loser of this bidding round exits and the winner becomes the incumbent, with the current standing price being the level at which the loser drops out.
3. Conditional on outbidding the current incumbent, potential buyer i may submit a bid above the current standing price. If submitted, this jump bid is observed by all subsequent potential buyers, and becomes the standing price in round $n+1$.

Let b_{n-1} be the standing bid at the beginning of round n, and y_{n-1} be the valuation of the incumbent submitting bid b_{n-1}. In view of the sequence of events above, these objects evolve as follows. Upon being contacted by the acquirer in round n, potential acquirer i observes its signal s_{ij} and (based on this plus the history of the game to date) decides whether to enter. If i remains out in round n, the game proceeds to round $n+1$ with standing bid $b_n = b_{n-1}$. Alternatively, if i elects to enter and draws valuation v_{ij} upon entry, then three outcomes are possible. First, if v_{ij} is less than the current standing bid b_{n-1}, then i exits and the negotiation proceeds to round $n+1$ with standing bid $b_n = b_{n-1}$. Second, if v_{ij} is greater than b_{n-1} but less than y_{n-1}, then i bids up the price to v_{ij} before exiting, and the negotiation proceeds to round $n+1$ with standing bid $b_n = v_{ij}$. Finally, if v_{ij} is greater than y_{n-1}, then the current incumbent bids up the price to y_{n-1} before exiting, and i becomes the new incumbent. To signal strength and thereby deter future entry, incumbent i may then submit a *jump bid* $b_n \geq y_{n-1}$. This jump bid b_n then becomes the standing bid in round $n+1$.

In practice, we focus on the unique *separating perfect Bayesian equilibrium* within this sequential negotiation game. This equilibrium has two key components. First, entry decisions by bidder i in round $n = 1, ..., N$ are described by a signal threshold $s_n^*(y_{n-1})$ such that potential buyer i enters if and only if $S_{ij} \geq s_n^*(y_{n-1})$, where y_{n-1} denotes the valuation of the standing bidder and $s_n^*(y_{n-1})$ is strictly increasing in y_{n-1}. Second, conditional on outbidding an incumbent with valuation y_{n-1} in round $n < N$, new incumbent i with valuation $v_{ij} \geq y_{n-1}$ submits a jump bid $b_n \geq y_{n-1}$ described by a symmetric monotone jump bidding strategy $\beta_n(v_{ij}, y_{n-1})$. A separating perfect Bayesian equilibrium is therefore a collection of round-specific entry threshold functions $(s_1(y_0), ..., s_N(y_{N-1})$ and round-specific bidding strategies $(\beta_1(v, y_0), ..., \beta_{N-1}(v, y_{N-1}))$ such that all players are best-responding at each information node. We characterize these strategies by backward induction as follows.

First consider the deterrence bidding decision of a new incumbent in round N. As the game concludes at the end of round N, a new incumbent in round N has no incentive to submit a deterrence bid. Conditional on knocking out an incumbent with standing valuation y_{N-1}, a new entrant drawing valuation $v \geq y_{N-1}$ therefore trivially submits bid $b_N = y_{N-1}$ for all $v \geq y_{N-1}$. In this event the new incumbent earns ex post payoff $v - y_{N-1}$.

Next consider the entry decision of the potential bidder contacted in round N. By hypothesis, we are considering a separating equilibrium in which prior new incumbents have played bidding strategies strictly monotone in their valuations. Observing the history of the game to date, the potential entrant in round N will therefore infer the standing valuation y_{N-1} of the incumbent at the end of round $N-1$. Conditional on drawing signal realization $S_i = s_i$ against an incumbent with standing valuation y_{N-1}, potential entrant N therefore expects post-entry profit

$$\Pi_N(s_i, y_{N-1}) = \int (V_i - y_{N-1})\, dF_{v|s_N}(V_i|s_i). \tag{C.1}$$

Potential entrant N will enter when expected profits exceed costs, i.e. when $\Pi_N^*(s_9, y_{N-1}) \geq c_j$. This breakeven condition in turn determines the breakeven threshold $s_N(y_{N-1})$, which by arguments similar to those in Appendix A can be shown to be a strictly monotone function of the current standing valuation y_{N-1}.

Now consider the deterrence bidding decision of a new incumbent in round $N-1$. Conditional on knocking out a prior incumbent with valuation y_{N-2}, a new incumbent i with valuation $v_i \geq y_{N-2}$ faces the following tradeoff: by submitting a higher bid, i may pretend to be a higher type and thereby deter entry by a potential competitor in round N, but doing so will require i to pay a higher cost conditional on winning in this event. Specifically, if rivals expect i to bid according to the strategy $\beta_{N-1}(\cdot, y_{N-2})$, and entry decisions by i's potential round-N rival are taken according to the threshold $s_N(\cdot)$ above, then we may write i's deterrence bidding problem as

$$\max_{z \geq y_{N-2}} \pi_N(v_i, z; \beta_{N-1}(z, y_{N-2})) \tag{C.2}$$

where $\pi_N(v_i, z; b_{N-1})$ denotes the expected profit, at the start of round N with standing bid b_{N-1}, of an incumbent with true valuation v_i but whom potential rivals *believe* to have valuation z:

$$\begin{aligned} \pi_N(v_i, z; b_{N-1}) &= s_N(z) \cdot (v_i - b_{N-1}) \\ &\quad + (1 - s_N(z)) \int_0^{v_i} (v_i - \max\{Y_N, b_{N-1}\}) dF(Y_N|S_N \geq s_N(z)). \end{aligned} \tag{C.3}$$

Note that the first term of $\pi_N(v_i, b_{N-1}; z)$ reflects i's profit from events in which round N entry is successfully deterred, whereas the second term represents i's expected profit in events where i's round-N rival enters but draws a valuation below v_i.

Taking a first-order condition of (C.2) with respect to i's type report z, we obtain:

$$\frac{\partial}{\partial b_{N-1}} \pi_N(v_i, z; \beta_{N-1}(z; y_{N-2})) \cdot \beta'(z, y_{N-2}) + \frac{\partial}{\partial z} \pi_N(v_i, z; \beta_{N-1}(z; y_{N-2})) = 0.$$

Enforcing the restriction that in equilibrium the strategy $\beta_{N-1}(\cdot, y_{N-2})$ must be such that it is optimal for i to report truthfully, this in turn implies the following differential equation characterizing the unknown deterrence bidding strategy $\beta_{N-1}(\cdot; y_{n-2})$:

$$\beta'(v_i; y_{N-2}) = -\frac{\frac{\partial}{\partial z}\pi_N(v_i, v_i; \beta_{N-1}(v_i; y_{N-2}))}{\frac{\partial}{\partial b_{N-1}}\pi_N(v_i, v_i; \beta_{N-1}(v_i; y_{N-2}))}. \tag{C.4}$$

Combined with the boundary condition $\beta(y_{N-2}; y_{N-2}) = y_{N-2}$, this in turn determines the function $\beta(\cdot; y_{N-2})$ describing equilibrium deterrence bidding in round $N-1$.

It is straightforward to show that, for given v_i, round-N profit $\pi_N(v_i, b_{N-1}; z)$ is strictly decreasing in b_{N-1} given z and strictly increasing in z given b_{N-1}. In other words, for given rival beliefs, i always prefers a strictly lower standing bid, and for a given standing bid, i always prefers rivals to believe she is a stronger type. Hence $\beta'(v_i; y_{N-2}) > 0$ above, which implies that the equilibrium deterrence bidding function $\beta(\cdot; y_{N-2})$ is strictly increasing as expected. This confirms that bidding in round $N-1$ is consistent with a strictly separating equilibrium, as desired.

Finally, consider the entry decision of potential bidder i with signal $S_i = s_i$ in round $N-1$, facing an incumbent with standing valuation y_{N-2}. Conditional on entry, the expected profit of this potential entrant will be equal to the optimal round N profit $\pi_N(v_i, v_i; \beta_{N-1}(v_i; y_{N-2}))$ described above, integrated over potential realizations V_i of v_i such that $V_i \geq y_{N-2}$:

$$\Pi_{N-1}(s_i; y_{N-2}) = \int_{y_{N-2}}^{\infty} \pi_N(V_i, V_i; \beta_{N-1}(V_i; y_{N-2}))\, dF_{v|s}(V_i|s_i). \tag{C.5}$$

It is again straightforward to show that the right-hand integrand must be increasing in V_i and decreasing in y_{N-2}. In view of the fact that V_i is stochastically increasing in s_i, this in turn implies that there will exist a unique threshold function $s_{N-2}(y_{N_2})$ such that i enters against an incumbent with standing valuation y_{N-2} only if $S_i \geq s_{N-2}(y_{N_2})$. We thereby obtain a complete characterization of entry and bidding behavior in round $N-2$.

Proceeding recursively in this fashion for rounds $N-3, N-4, ..., 1$, one ultimately obtains the desired series of strictly increasing entry functions $s_1(y_0), ..., s_N(y_{N-1})$ and strictly increasing bidding functions $\beta_1(v_i; y_0), ...,$ $\beta_{N-1}(v_i; y_{N-2})$ characterizing the unique symmetric separating perfect Bayesian equilibrium of the sequential negotiation game. Roberts and Sweeting (2013) furthermore show that the resulting separating equilibrium is the *only* perfect Bayesian equilibrium to survive standard refinements (no weakly dominated strategies, sequential equilibrium, and the D1 refinements of Cho and Sobel (1990) and Ramey (1996)) on equilibria of the sequential negotiation game.

The go-shop mechanism is similar to the sequential negotiation but with only a single round of deterrence bidding. Bidding starts at a standing bid b_0 equal to the target's reservation valuation. The target approaches one potential bidder i at random, who observes their signal realization s_i and based on this and the standing bid b_0 decides whether to enter. If bidder i enters, i observes its valuation v_i, and may submit a jump bid above b_0, which becomes the new standing bid b_1. Otherwise, the standing bid remains $b_1 = b_0$. The game then proceeds to a go-shop round, in which the target contacts the other $N-1$ bidders, who based on the history of the game and their private signals decide whether or not to enter. If at least two bidders ultimately enter with values above the standing bid b_1, the game concludes with an ascending auction to determine the winning bidder. Otherwise the game concludes at the standing bid b_1.

Equilibrium in the go-shop mechanism is similar to that in the sequential mechanism, but simpler, since there is only one round of deterrence bidding. As above, we look for a symmetric separating perfect Bayesian equilibrium. The $N-1$ entrants in the final round observe the history of the game to date, including the initial bid b_0, whether a bidder entered in Round 1, and the entrant's jump bid b_1 if one was submitted. Based on this, all potential entrants infer the standing valuation upon entry, which we denote y_1: either $y_1 = v_0 = b_0$ if the standing bid is the seller's reservation price, or y_1 equal to the incumbent's valuation if the incumbent submitted a jump bid. Entry by all $N-1$ bidders contacted in the go-shop stage therefore proceeds according to a threshold $s^*(y_1)$, determined by the condition that a go-shop entrant drawing signal realization $S_i = s^*(y_1)$ must just break even from entry when the standing value is y_1 and the other go-shop entrants enter according to the threshold $s^*(y_1)$:

$$\int_{y_1}^{\infty} [F_s(s^*(y_1)) + F_v(v) - F_{vs}(v, s^*(y_1))]^{N-2}\, dF_{v|s}(v|S_i = s^*(y_1)). \qquad \text{(C.6)}$$

Now consider Round 1 deterrence bidding by the incumbent. We seek a deterrence bidding strategy $B_I(v; b_0)$ which (for $v > b_0$) is strictly monotone in the incumbent's valuation v. As above, this strategy will be characterized by the condition that the incumbent's gains from pretending to be a higher type are just offset by the additional costs of a higher standing bid. Letting z denote the incumbent's pretended type, this yields the maximization problem

$$\max_z \Big\{ (v - \beta_I(z; b_0)) \cdot [F_s(s^*(z)) + F_v(\beta_I(z; b_0)) - F_{vs}(\beta_I(z; b_0), s^*(z))]^{N-1} + \int_{\beta_I(z;b_0)}^{v} (v - y) \frac{d}{dy} [F_s(s^*(z)) + F_v(y) - F_{vs}(y, s^*(z))]^{N-1} \, dy \Big\}, \quad (C.7)$$

where the first term reflects incumbent profents in the event that the incumbent faces no rival entrant with a valuation above the go-shop standing bid $b_1 = \beta_I(z; b_0)$, and the second reflects expected profits in the event that at least one go-shop rival enters with a valuation above the chosen standing bid $\beta_I(z; b_0)$. Taking a first-order condition with respect to z and enforcing the equilibrium condition $v = z$, we ultimately obtain a differential equation characterizing the derivative $B_I'(v; b_0)$ of $B_I(v; b_0)$, which together with the boundary condition $B_I(b_0; b_0) = b_0$ uniquely determines the equilibrium deterrence bidding strategy $B_I(\cdot; b_0)$. As above, it is straightforward to show that this strategy $B_I(\cdot; b_0)$ must be strictly increasing, which confirms that $B_I(\cdot; b_0)$ describes a separating perfect Bayesian equilibrium as desired.

References

Aktas, N., de Bodt, E., Roll, R., 2010. Negotiations under the threat of an auction. Journal of Financial Economics 98, 241–255.

Bebchuk, L., 1982. The case for facilitating competing tender offers: A reply and extension. Stanford Law Review, 23–50.

Betton, S., Eckbo E., 2000. Toeholds, bid jumps, and expected payoffs in takeovers. Review of Financial Studies, 13, 841–882.

Boone, A., Mulherin, H., 2007. How are firms sold? Journal of Finance, 62, 847–875.

Bulow, J., Klemperer, P., 2009. Why do sellers (usually) prefer auctions?

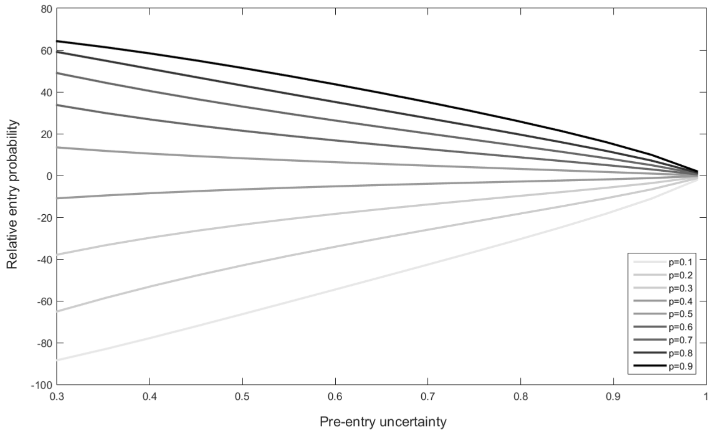

Figure 1: Pre-entry uncertainty and the composition of the entering bidder pool. This figure reports the probability that potential bidders with different valuations enter, relative to an average bidder. Specifically, each line plots the probability of entry (relative to the mean bidder) for a bidder whose value lies at the pth percentile of the potential bidder valuation distribution. The probabilities are computed using the estimates obtained in Section 5 and the method described in Section 3, for the sample of takeover auctions of public targets announced between January 1, 2000 and January 1, 2010.

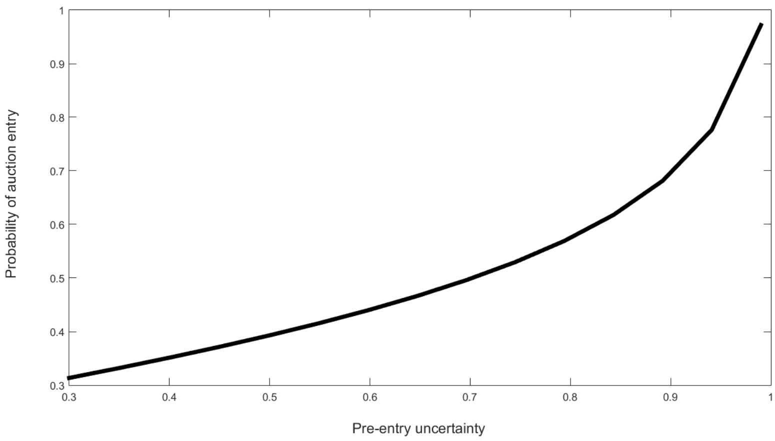

Figure 2: Pre-entry uncertainty and the size of the entering bidder pool. This figure reports the proportion of invited bidders that participate, on average, for takeover markets with different levels of pre-entry uncertainty. The estimates are obtained using median fundamental takeover market primitives obtained in Section 5 using the method described in Section 3, on the sample of takeover auctions of targets announced between January 1, 2000 and January 1, 2010.

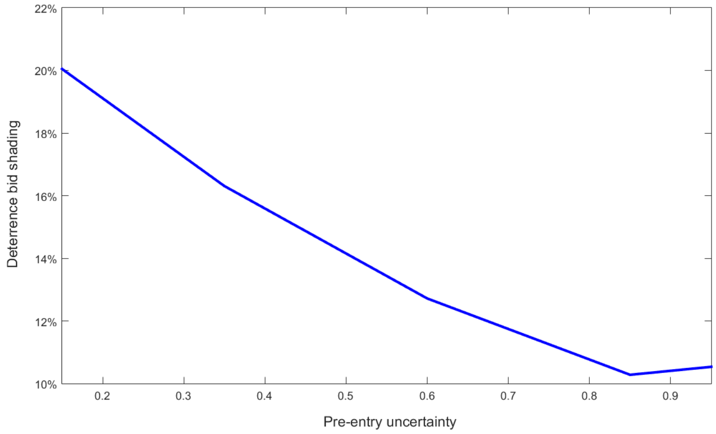

Figure 3: Deterrence bidding in takeover negotiations. This figure shows additional contribution to deal premia from deterrence bidding, for different levels of pre-entry uncertainty. The vertical axis measures deterrence bidding in takeover premia units, i.e., as a percent of the target's current share price. Formally, the deterrence effect is defined as the difference between actual sale revenue and counterfactual revenue that would have resulted from sale by an auction among the set of bidders participating at any point in the negotiation sequence, i.e., the maximum of the second-highest negotiating bidder's valuation or the target's reservation price. The estimates are obtained using the structural model and median fundamental takeover market primitives estimated from the sample of takeover auctions of public targets announced between January 1, 2000 and January 1, 2010.

Table 1
Sample descriptive statistics: summary of the sale process

This table summarizes descriptive statistics (number, means, and standard deviations) of sale characteristics in 529 takeover auctions from January 1, 2000 to January 1, 2010 for all auctions and across Fama and French (1997) industries. The variable Contact reports the average number of contacted potential bidders. The variable Confidential reports the average number of potential bidders that participate in a takeover auction by signing a confidentiality agreement with the target. The variable Winning Bid reports the average price paid by the winning bidder, normalized by the target's share price four weeks prior to the deal's announcement. Standard errors appear in parentheses.

	Number	Contact	Confidential	Winning bid
All auctions	529	13.291	5.844	1.428
		(19.374)	(7.479)	(0.339)
Consumer Nondurables	19	13.631	6.526	1.419
		(14.330)	(6.327)	(0.278)
Consumer Durables	5	23.400	9.000	1.325
		(29.065)	(14.560)	(0.320)
Manufacturing	26	24.384	11.115	1.483
		(42.002)	(18.313)	(0.424)
Oil, Gas, and Coal	16	19.375	5.687	1.236
		(20.806)	(6.508)	(0.123)
Chemicals, etc.	4	4.000	2.500	1.552
		(2.160)	(1.290)	(0.366)
Business Equipment	159	12.773	4.339	1.501
		(19.839)	(5.348)	(0.366)
Communications	13	10.230	7.076	1.443
		(12.234)	(8.097)	(0.457)
Utilities	7	14.571	6.428	1.394
		(9.519)	(5.826)	(0.276)
Wholesale, Retail	12	18.166	9.916	1.645
		(17.486)	(12.265)	(0.529)
Healthcare, Medical, etc.	57	12.473	5.473	1.462
		(16.546)	(5.590)	(0.338)
Finance	178	11.797	6.016	1.331
		(15.421)	(6.157)	(0.242)
Other	33	12.121	6.181	1.517
		(17.318)	(7.699)	(0.415)

Table 2
Target characteristics

This table reports descriptive statistics (means and standard deviations) for targets sold in 529 takeover auctions announced between January 1, 2000 and January 1, 2010 for the full sample and for each of the 12 Fama and French (1997) industries. The table reports data on target assets measured as the book value of total assets (current assets plus net property, plant, and equipment plus other noncurrent assets), market leverage, q ratio (market-to-book), along with cash, intangibles, and sales, with these latter three variables scaled relative to book assets. Standard deviations appear in parentheses.

	Assets	q	Leverage	Cash	Intangibles	Sales
All auctions	848.77	1.292	0.214	0.141	0.103	0.708
	(1452.14)	(1.647)	(0.248)	(0.176)	(0.166)	(0.797)
Nondurables	972.06	1.273	0.199	0.086	0.142	1.474
	(1311.90)	(0.882)	(0.201)	(0.147)	(0.176)	(0.818)
Durables	783.96	0.935	0.359	0.099	0.110	1.189
	(1272.22)	(0.430)	(0.336)	(0.100)	(0.120)	(0.680)
Manufacturing	708.54	1.357	0.150	0.105	0.122	1.261
	(1669.53)	(1.051)	(0.163)	(0.140)	(0.162)	(0.551)
Oil, Gas, and Coal	1490.39	1.146	0.338	0.050	0.015	0.704
	(1689.33)	(0.423)	(0.177)	(0.072)	(0.027)	(0.869)
Chemicals, etc.	287.71	1.001	0.165	0.053	0.040	2.195
	(190.66)	(0.103)	(0.191)	(0.061)	(0.055)	(1.587)
Business Equipment	350.44	1.713	0.062	0.246	0.162	0.846
	(721.03)	(2.263)	(0.126)	(0.183)	(0.187)	(0.441)
Communications	962.66	1.437	0.371	0.094	0.307	0.705
	(973.00)	(0.729)	(0.287)	(0.109)	(0.227)	(00.893)
Utilities	2504.67	0.934	0.377	0.013	0.087	1.439
	(2533.90)	(0.285)	(0.126)	(0.014)	(0.142)	(2.672)
Wholesale, Retail	638.49	1.903	0.152	0.109	0.099	2.596
	(862.93)	(2.767)	(0.269)	(0.143)	(0.121)	(1.493)
Healthcare	563.00	2.529	0.104	0.241	0.119	0.629
	(1529.87)	(1.756)	(0.150)	(0.210)	(0.179)	(0.454)
Finance	1350.21	0.465	0.365	0.039	0.020	0.195
	(1734.26)	(0.449)	(0.264)	(0.066)	(0.062)	(0.438)
Other	516.07	1.311	0.257	0.193	0.166	0.982
	(1044.48)	(0.779)	(0.298)	(0.237)	(0.232)	(0.634)

Table 3
Takeover market primitives

This table reports estimates of the takeover market primitives recovered by the structural model described in Section 3, on the data described in Section 4. Panel A reports parameter means across all targets in our sample, and medians in brackets. The parameters of interest are the mean (μ) and variance (σ) of the potential entrant value distribution, average entry costs (c), and the average degree of pre-entry uncertainty (α). Column 1 reports estimates from a specification that employs a uniform proposal distribution without conditioning on target observables. The estimates in Columns 2-5 iterate once on the proposal distribution as described in Appendix B. The estimates in Column 2 do not condition on target observables. The estimates in Column 3 condition (μ) on the vector of target observables. Estimates in Column 4 condition both (μ) and (α) on the vector of target observables. Estimates in Column 5 condition all fundamental parameters on target observables. Panel B reports characteristics of the distribution of model parameters across takeover markets: the standard deviation and the 25th and 75th percentiles of the estimated parameter distributions.

Panel A: Point estimates

	(1)	(2)	(3)	(4)	(5)
Pre-entry uncertainty (α)	0.6200	0.6260	0.6367	0.5830	0.6082
	[0.6595]	[0.6479]	[0.6592]	[0.6035]	[0.6203]
Average entry cost (c)	0.0151	0.0109	0.0115	0.0091	0.0528
	[0.0089]	[0.0044]	[0.0046]	[0.0037]	[0.0187]
Average potential bidder valuation (μ)	0.2024	0.2818	0.2834	0.2841	0.2660
	[0.2026]	[0.2819]	[0.2832]	[0.2840]	[0.2662]
Spread of valuations across potential bidders (σ)	0.1539	0.0598	0.0593	0.0581	0.0753
	[0.1261]	[0.0543]	[0.0544]	[0.0524]	[0.0558]

Panel B: Cross-sectional variation in takeover market parameters

	α	c	μ	σ
Standard deviation	0.1280	0.0121	0.2106	0.0330
25th Percentile	0.4010	0.0007	0.1308	0.0342
75th Percentile	0.7785	0.0113	0.4376	0.0753

Table 4
Estimates of partial effects

This table reports average marginal effects for individual γ's. The marginal effects are constructed by holding X constant at the median and drawing a sample of γ's from the posterior (n=100). The marginal effect of X at γ is defined as the average effect of either a unit change in X (for dummy variables) or a one standard deviation change in X (for continuous variables) on the mean of θ. Standard errors of the marginal effects are reported in parentheses beneath coefficient estimates. a, b, and c denote statistical significance of individual coefficients at the 5

	Median X			Median X	
	Posterior γ			Posterior γ	
Variable	μ_v	α	Variable	μ_v	α
Size	-0.125^a	-0.032	Oil, Gas, Coal	-0.094^a	-0.069^a
	(0.040)	(0.038)		(0.026)	(0.023)
q	-0.011	-0.039^c	Chemicals, etc.	0.018	-0.027^a
	(0.034)	(0.028)		(0.029)	(0.013)
Leverage	0.054^c	0.011	Business Equipment	-0.004	-0.192^a
	(0.034)	(0.026)		(0.068)	(0.041)
Cash	0.028	-0.013	Telephone & Television	-0.013	0.010
	(0.043)	(0.024)		(0.028)	(0.018)
Intangibles	-0.070^a	-0.004	Utilities	-0.021	-0.032
	(0.028)	(0.025)		(0.025)	(0.032)
Sales	-0.046^c	-0.006	Wholesale/Retail	0.0318	-0.004
	(0.037)	(0.029)		(0.037)	(0.020)
Nondurables	-0.013	0.0210	Healthcare	-0.047	-0.063^a
	(0.037)	(0.030)		(0.048)	(0.031)
Durables	-0.034^c	0.015	Finance	-0.176^a	-0.078^a
	(0.026)	(0.0364)		(0.067)	(0.036)
			Other	-0.011	-0.025
				(0.044)	(0.030)

Table 5
Target characteristics and average pre-entry uncertainty

This table shows average pre-entry uncertainty in different subsamples, using the main estimates from Column 4 of Table 3 and the model-implied weights. Panel A reports target-level averages of pre-entry uncertainty for each industry, and pre-entry uncertainty at the 2.5 and 97.5 percentiles of the parameter distributions. Panel B sorts target-level fundamental parameters into quantiles based on observable balance sheet characteristics and reports the average level of pre-entry uncertainty across targets within each size-observable quantile.

Panel A. Variation in pre-entry uncertainty (α) within and across industries

Industry	Mean	2.5%	97.5%	Industry	Mean	2.5%	97.5%
Nondurables	0.87	0.46	0.99	Communication	0.79	0.33	0.99
Durables	0.84	0.24	1.00	Utilities	0.54	0.04	0.99
Manufacturing	0.80	0.43	0.99	Wholesale, Ret.	0.73	0.30	0.98
Oil, Gas, and Coal	0.35	0.03	0.78	Healthcare	0.58	0.19	0.93
Chemicals, etc.	0.50	0.07	0.94	Finance	0.67	0.29	0.94
Bus. Equipment	0.38	0.08	0.75	Other	0.70	0.29	0.98

Panel B: Quantile sorts of α on balance sheet characteristics

			Cash					Leverage		
	Low	2	3	4	High	Low				High
Small	0.72	0.77	0.59	0.57	0.52	0.54	0.53	0.58	0.64	0.74
2	0.74	0.64	0.58	0.50	0.48	0.51	0.53	0.63	0.66	0.66
3	0.65	0.69	0.67	0.47	0.41	0.50	0.53	0.61	0.64	0.68
4	0.61	0.66	0.60	0.51	0.41	0.52	0.51	0.55	0.61	0.65
Big	0.60	0.63	0.58	0.56	0.33	0.39	0.40	0.59	0.61	0.62
			q					Sales		
Small	0.67	0.59	0.60	0.51	0.53	0.60	0.56	0.59	0.51	0.61
2	0.68	0.57	0.59	0.51	0.48	0.64	0.53	0.51	0.51	0.63
3	0.70	0.66	0.57	0.52	0.44	0.68	0.65	0.45	0.55	0.64
4	0.68	0.64	0.60	0.46	0.51	0.67	0.59	0.51	0.54	0.57
Big	0.64	0.63	0.53	0.63	0.49	0.64	0.54	0.52	0.50	0.72

Table 6
Unconditional counterfactual estimates

This table reports means, medians, and standard deviations of expected revenue accruing to target shareholders from auctions with endogenous entry, from one-shot negotiations followed by a go-shop auction, and from sequential negotiations, on a sample of takeover auctions announced between January 1, 2000 and January 1, 2010. The estimates are constructed using the structural model described in Section 3 at median observable and unobservable characteristics (Median Γ) and the resulting baseline fundamental parameter estimates.

	Auction	Negotiation with go-shop	Sequential negotiation
Expected revenue	41.17	41.48	42.61
Revenue S.D.	11.14	8.78	3.32
Revenue skewness	-2.63	-2.84	-0.52
Surplus extracted if sale (pct.)	84.55	85.65	88.40
Highest value potential bidder wins (pct.)	76.20	70.77	63.95

Table 7
The size effect

This table reports the percent of expected deal revenue attributable to the size effect, for different values of pre-entry uncertainty and entry costs. This size effect, defined in Section 6.2, is the difference between actual expected revenue and counterfactual expected revenue where entering bidders were randomly selected from the pool of potential bidders instead of self-selecting based on their pre-entry beliefs. The counterfactual revenue is computed by taking as given the number of actual entrants in an auction, and using the structural estimates to compute revenue that would have resulted if the same number of entrants had been randomly selected from the pool of potential bidders. The values reported in the table are obtained by dividing the size effect by expected auction revenue. The estimates are constructed using median observable and unobservable characteristics (Median Γ) and the resulting baseline fundamental parameter estimates recovered using the method described in Section 3.

	Pre-entry uncertainty (α)				
	0.15	0.35	0.60	0.85	0.99
$c = 0.005$	93.47%	93.88%	94.6%	95.88%	99.99%
$c = 0.010$	92.97%	93.38%	94.15%	95.40%	97.22%
$c = 0.030$	92.62%	93.08%	93.89%	95.18%	97.03%

Table 8
Auctions vs negotiations and target observables

This table shows differences in the relative performance of auctions and negotiations among different subsets of targets. The counterfactual estimates are obtained using the structural primitives recovered in Section 5 using the method described in Section 3, on the data described in Section 4. We form a decile grid in each fundamental parameter, and for each possible combination of fundamental parameters draw a heterogeneity vector from the posterior obtained in Section 5, applying model weights to obtain estimates of fundamental takeover market parameters that vary with target observables. The resulting target-level parameters are used to calculate expected revenue differences across takeovers. We sort all targets in our sample based on each characteristic and use the estimated model primitives to compute the fraction of targets within each category for which expected auction revenue is greater than expected revenue from a sequential negotiation.

Panel A: Industry patterns

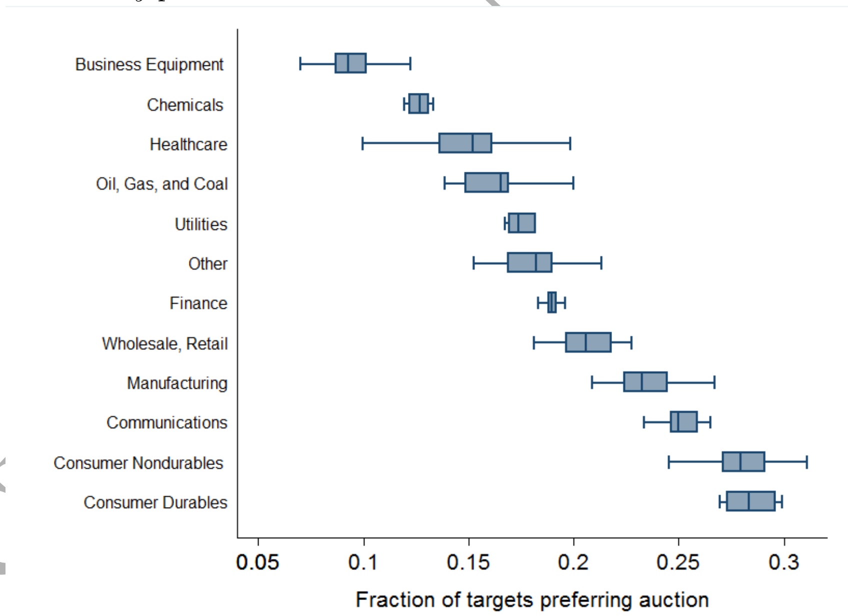

Panel B: Balance sheet characteristics

Characteristic	Bottom quantile of the Sample	Top quantile of the sample	Revenue difference correlation
Assets	14.75%	19.18%	-0.615
q	18.92%	12.19%	0.373
Leverage	13.34%	19.71%	-0.378
Cash	19.12%	11.51%	0.667
Intangibles	16.32%	16.03%	-0.019
Sales	18.16%	19.54%	-0.055

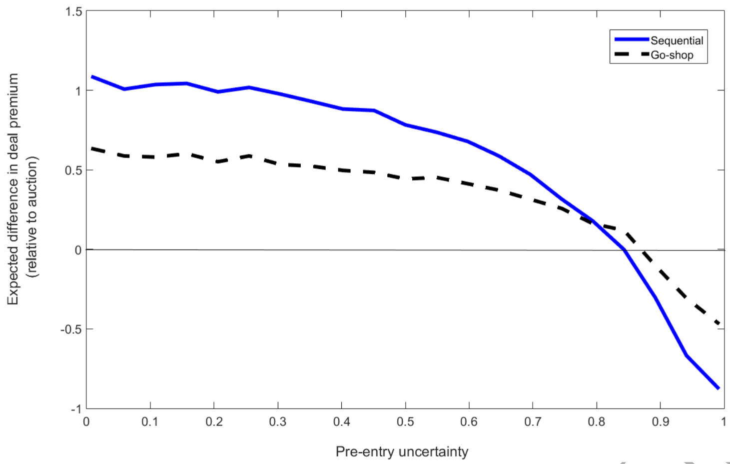

Figure 4: Negotiation revenue relative to auction. This figure reports the expected difference in takeover auction revenue between negotiations and auctions, as a function of pre-entry uncertainty. The estimates are obtained using the structural estimates and median fundamental takeover market primitives estimated from the sample of takeover auctions of public targets announced between January 1, 2000 and January 1, 2010.

American Economic Review 99, 1544–1575.

Bulow, J., Klemperer, P., 1996. Auctions versus negotiations. American Economic Review 86, 180–194.

Cho, I., Sobel, J., 1990. Strategic stability and uniqueness in signaling games. Journal of Economic Theory 50, 381–413.

Cramton, P., Schwartz, A.,1991. Using auction theory to inform takeover regulation. Journal of Law, Economics, & Organization, 7, 27–53.

Daniel, K., Hirshleifer, D., 1998. A theory of costly sequential bidding. Review of Finance, Forthcoming.

Dimopoulos, T., Sacchetto, S., 2014. Preemptive bidding, target resistance, and takeover premiums. Journal of Financial Economics, 114, 444–470.

Eckbo, B., Giammarino, R., Heinkel, R., 1990. Asymmetric information and the medium of exchange in takeovers: Theory and tests. Review of Financial Studies 3, 651–67.

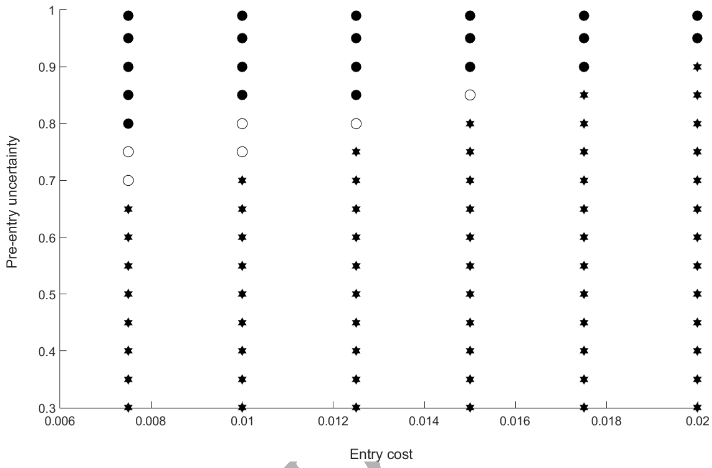

Figure 5: Uncertainty, entry costs, and expected revenue. This figure reports a comparison of revenue-maximizing sale mechanisms (auction versus sequential negotiation) for different levels of uncertainty and entry costs. The estimates are constructed using median observable and unobservable characteristics and the resulting baseline fundamental parameter estimates. Dark circles indicate takeover market parameters implying auctions revenue-dominate sequential negotiations, and hexagrams indicate the opposite. Hollow circles indicate situations where simulation error is larger than estimated revenue differences.

Fama, E., and French, K. 1997. Industry costs of equity. Journal of Financial Economics 43, 153–193.

Fidrmuc, J., Roosenboom, P., Paap, R., Teunissen, T., 2012. One size does not fit all: selling firms to private equity versus strategic acquirers. Journal of Corporate Finance, 18, 828–848.

Fishman, M., 1988. A theory of preemptive takeover bidding. RAND Journal of Economics, 19, 88–101.

French, K., McCormick, R., 1984. Sealed bids, sunk costs, and the process of competition. Journal of Business 57, 417–441.

Gentry, M., Li, T., 2014. Identification in auctions with selective entry. Econometrica, 82, 315–344.

Gorbenko, A., Malenko, A., 2014. Strategic and financial bidders in takeover auctions. Journal of Finance, 69, 2513–2555.

Gorbenko, A., Malenko, A., 2017. The timing and method of payment in mergers when acquirers are financially constrained. Review of Financial Studies, forthcoming.

Haile, P., Tamer, E., 2003. Inference with an incomplete model of english auctions. Journal of Political Economy, 111, 1–51.

Hansen, R., 2001. Auctions of companies. Economic Inquiry 39, 30–43.

Horner, J., Sahuguet, N., 2007. Costly signalling in auctions. Review of Economic Studies, 74, 173–206.

Kirman, I., 2008. M&A and private equity confidentiality agreements line by line, Thompson, Aspatore Books, Boston, MA.

Levin, D., Smith, J., 1994. Equilibrium in auctions with entry. The American Economic Review, 84, 585–599.

Moeller, S., Schlingemann, F., Stulz, R., 2004. Firm size and the gains from acquisitions. Journal of Financial Economics, 73, 201–228.

Povel, P., Singh, R., 2006. Takeover contests with asymmetric bidders.

Review of Financial Studies, 19, 1399–1431.

Ramey, G., 1996. D1 Signaling equilibria with multiple signals and a continuum of types. Journal of Economic Theory 69, 508–531.

Roberts, J., Sweeting, A., 2013. When should sellers use auctions? American Economic Review, 103, 1830–1861.

Rosenbaum, J. Pearl, J., 2009. Investment Banking - Valuation, Leveraged Buyouts, and Mergers & Acquisitions. John Wiley & Sons, Inc., Hoboken, NJ.

Rousseau, P., Stroup, C., 2015. Director histories and the pattern of acquisitions. Journal of Financial and Quantitative Analysis, 50, 671–698.

Samuelson, W., 1985. Competitive bidding with entry costs. Economics Letters, 17, 53–57.

Subramanian, G., 2008. Go-shops vs. no-shops in private equity deals: evidence and implications. Business Lawyer, 63, 729-760.

Subramanian, G., 2011. Dealmaking: the new strategy of negotiauctions. WW Norton & Company, New York.

Travlos, N., 1987. Corporate takeover bids, methods of payment, and bidding firm's stock returns. Journal of Finance 42, 943–963.

Wang, Z., 2016. Structuring M&A offers: auctions, negotiations and go-shop provisions. Unpublished Working Paper, Stanford University.

Wasserstein, B., 2000. Big Deal: Mergers and Acquisitions in the Digital Age. Warner Business, New York.