

Entry and Competition in Takeover Auctions

Matthew Gentry* and Caleb Stroup†

December 23, 2014

Abstract

We show that in many cases target shareholders would obtain higher prices if their company were sold in a negotiated sale, rather than in an auction. Accounting for the endogenous determination of the size and composition of the bidder pool, we show that possible bidders in takeover auctions face substantial uncertainty prior to their entry into an auction, and that fewer than half of invited potential acquirers choose to participate in competitive bidding for a target. We show that higher pre-entry uncertainty encourages participation in competitive bidding, thus making auctions preferable to negotiations when uncertainty is high. Uncertainty reduces the effectiveness of upward bid-shading in negotiations to deter potential competitors, so negotiations are preferable to auctions when the selling company is relatively opaque to potential bidders. Our results call into question claims that target directors violate their fiduciary duty by selling a company via a negotiated transaction, even in the absence of a formal market check.

KEYWORDS. Takeover Auctions; Mergers and Acquisitions

JEL CLASSIFICATION NUMBERS: D44 G44.

*Department of Economics, London School of Economics. Email: m.l.gentry@lse.ac.uk.

†Department of Economics, Grinnell College. Email: stroupca@grinnell.edu.

A commonly held view is that auctions, in contrast to negotiated sales, yield higher average prices for shareholders of target companies: acquirers prefer negotiated transactions while sell side advisors regularly prescribe broad based auctions. For example, Wasserstein (2000) reports that “A wide-ranging auction generally maximizes value . . . sophisticated bidders will do their best to circumvent the auction format,” and the “acquisition criteria” section of Warren Buffet’s annual reports states “We don’t participate in auctions.” A recent survey showed that many buyers overwhelmingly prefer to participate in negotiated purchases but sell companies via auction (Auction Process Roundtable, Mergers and Acquisitions, December 2006, pp. 31-32).

The view that auctions revenue dominate negotiations has its origin in Bulow and Klemperer (1996), who show that auctions always in principle yield higher revenue than negotiated sales, a theoretical conclusion that Bulow and Klemperer (2009) demonstrate extends to the situation where negotiation bids are shaded upward to deter entry by potential competitors.

Recent research has found that financial markets react similarly to auctions and negotiations, which could be interpreted as evidence that the impact of potential competition on deterrence bids in negotiations is sufficient to generate prices similar to auctions (Boone and Mulherin (2007), Aktas, de Bodt and Roll (2010)). Yet there still exists little empirical evidence about how a particular firm *should* be sold. One approach to answering this question would be to compare observed deal premia resulting from transactions structured as auctions with deal premia resulting from negotiations. Such an approach would have the potential to yield insights about how firms *are* sold, but there are at least two reasons why it cannot be informative about how a firm should be sold.

First, the relative optimality of auctions and negotiations depends critically on the size and composition of the pool of participating bidders, yet in practice this pool is not exogenously given. As we show, less than half of invited potential bidders choose to participate in takeover auctions. The relative performance of auctions and negotiations depends on

the quantities that determine potential bidders' entry decisions, which include uncertainty about realizable synergies with the target and the costs of overcoming it, but these quantities are not directly observed in the data. A structural approach is thus required to characterize how a firm should be sold (e.g., Gorbenko and Malenko (2013), Roberts and Sweeting (2013)). We show that failure to account for endogeneity in the size and composition of the entering bidder pool leads to systematic overestimation of the return to auctions relative to negotiations.

Second, while takeover auctions have become relatively standardized in practice (Hansen (2001)), transactions involving a single bidder, which are typically classified as "negotiations," can take a variety of observationally indistinguishable forms, each of which produce different levels of expected revenue for target shareholders, since negotiations are not homogenous in their ability to induce high offer prices. An observed single bidder sale could, for example, reflect either a successful one-shot negotiation or it could reflect a successful first stage in a sequential negotiation.

We overcome these challenges in two ways. First, we develop and estimate a structural empirical framework that recovers estimates of takeover market unobservables that determine the mapping between observed bids and the distribution of all bidder valuations on the one hand and the mapping between the distribution of entering bidder valuations and the distribution of potential bidder valuations on the other. The estimates allow us to quantify how the answer to the question "How *should* a firm be sold?" systematically depends on potential bidders' entry decisions. Second, we use the estimates to characterize how deal premia would change if the targets sold via an auction were instead sold using one of several well defined negotiation procedures, the effectiveness of which are characterized by the takeover market primitives recovered by our estimation procedure.

In the framework, potential bidders differ in their valuations for the target and are invited to participate in a standard takeover auction. Each bidder's valuation comprises a target specific common component and an unobserved bidder and target specific asset

complementarity, with uncertainty mitigated upon entry through due diligence conducted on the target. Our estimation framework introduces several innovations to the empirical finance literature. First, it incorporates and allows us to estimate any level of average pre-entry uncertainty faced by potential bidders. This parameter is required to characterize endogenous entry patterns that determine the size and composition of the pool of entering bidders. Second, our structural empirical approach simultaneously accommodates endogeneity in the major decisions made by a seller and potential buyers including the target's choice of sale procedure, each potential bidder's decision to participate in the auction, and strategic bidding by entrants, all of which are conditioned on information about entry costs, the size of the potential bidder pool, the average degree of pre-entry uncertainty faced by potential bidders, and variation across potential bidders in realizable synergies (asset complementarities net integration costs) associated with purchase of the target. Third, we introduce a procedure widely used in the empirical structural auction literature that makes our estimates robust to possible endogeneity along a wide array of sale-level dimensions necessarily unobservable to a researcher but not to market participants.

We begin the analysis by estimating a generalization of our model along on hand collected data that we obtain from takeover filings submitted to the Securities and Exchange Commission. The estimated primitives permit a comparison of the relative performance of auctions and negotiations but also reveal new insights about takeover markets.

Our main findings are as follows. First, the estimates imply the existence of high uncertainty faced by potential bidders about their valuations for the target, with pre-entry beliefs embodying more noise than information. Invited potential entrants with unfavorable initial beliefs decline to participate, even though many would have discovered information upon entry that would have caused them to revise their valuations upward. High pre-entry uncertainty thus implies that high-value bidders are regularly absent from the participating bidder pool, and we show that the potential bidder with the highest *ex post* valuation of the target declines to participate in about 36% of takeover auctions.

Second, we use the estimated primitives to quantify how and to what extent pre-entry uncertainty impairs the ability of takeover auctions to elicit high prices for target shareholders. Pre-entry uncertainty differently affects the size and the composition of the entering bidder pool. We refer to the negative effect of uncertainty on the endogenous participation of relatively high-valuation potential bidders described above as the “composition effect.” At the same time, by degrading quality of the entering bidder pool through reduced participation by high value bidders, uncertainty induces additional entry by potential entrants who have initially unfavorable beliefs about their valuations for the target, some of whom discover high valuations for the target upon entry. We refer to this positive effect of uncertainty on expected takeover revenue as the “size effect.” The relative magnitude of these two competing effects is an empirical question, which our structural econometric approach allows us to quantify. We find that a reduction of pre-entry uncertainty from its average level to zero would raise expected takeover auction revenue by less than 3%. Auctions are thus surprisingly resilient to the high level of pre-entry uncertainty that exists in takeover markets, and this is because the negative effect of uncertainty on the average quality of the entering bidder pool is partially offset by its positive effect on the overall size of the entering bidder pool. Our estimates thus provide support for the conventional notion that the ability to generate a large pool of competing bidders is what makes takeover auctions a powerful tool for creating value for target shareholders.

Third, we formally demonstrate that failing to account for endogeneity in the size and composition of the entering bidder pool systematically leads to overestimation of expected takeover auction revenue relative to a negotiation auction. This is because while the relative optimality of auctions and negotiations depends on the ability of each to leverage potential competition, auctions and negotiations do so differently. As described above, auctions leverage potential competition through the endogenous entry patterns that generate a large and competitive pool of entering bidders. Negotiations leverage potential competition when a standing bidder shades up their offer price to deter entry by additional competitors. Pre-

entry uncertainty crucially impacts the relative performance of auctions and negotiations through its influence on endogenous entry in auctions and because it determines the effectiveness of deterrence bidding in negotiations. As described above, auctions are relatively robust to the presence of high average pre-entry uncertainty, but high uncertainty reduces the effectiveness and incidence of deterrence bidding in a negotiated sale. Overestimation of pre-entry uncertainty - or equivalently, failure to account for endogenous entry patterns - biases expected relative revenue in favor of auctions.¹ Our estimates imply that a researcher who failed to account for pre-entry uncertainty and endogenous entry would over-estimate the relative return to auctions by about 6%.

Fourth, we account for endogenous determination of the size and composition of the entering bidder pool and show that in many firms would obtain higher if they were sold via a negotiation rather than an auction. To do this we conduct two counterfactual comparisons. We first compare expected revenue from holding a takeover auction with expected revenue from conducting a one-shot negotiation followed by a market check (i.e., a “go shop”) where a standing bid is publicly posted and potential bidders are invited to submit a higher bid. This structure is a stylized version of negotiation structures widely used in practice (e.g., Subramanian (2008), Wasserstein (2000)). The one-shot negotiation with a market check thus presents a simple and realistic alternative to standard takeover auctions that incorporates pressure on current negotiating bidders to shade up their offers to deter entry by potential competitors (e.g., Bulow and Klemperer (2009)).

We also compare expected auction revenue with expected revenue arising from a canonical one-on-one negotiation that allows the target to terminate negotiations with a standing bidder and to successively negotiate with additional bidders. Similar sequential mechanisms have been widely examined in the theoretical and empirical literature on optimal sale design (e.g., Fishman (1988), Betton and Eckbo (2000), Horner and Sahuguet (2007),

¹We formally show that failing to account for limited participation and endogenous entry is tantamount to the unrealistic boundary assumption of infinite pre-entry uncertainty.

Dimopoulos and Sacchetto (2011)). The sequential negotiation also leads to a situation where a standing bidder shades up their offer price to deter potential competitors. We find that on average, targets would have obtained 2.5% higher deal premia by structuring the sale of their company as a sequential negotiation rather than as an auction. Sophisticated negotiation mechanisms thus have the potential to leverage potential competition more effectively than can traditional auction-based procedures. At the same time, we show that these cross-sectional averages mask dramatic variation across targets in the relative returns to auctions and negotiations: while the majority of targets would have obtained higher revenue via a negotiated transaction there exist a small fraction of targets that would obtain significantly higher deal premia through sale via auction, and we show that differences across targets in the relative optimality of auctions and negotiations can systematically be explained by differences across takeover markets in the average degree of uncertainty faced by potential bidders and the costs of overcoming it, with these two variables jointly explaining about 40% of the variation across targets in the relative optimality of auctions and negotiations.

Our work is related to several papers. Boone and Mulherin (2007) analyze 400 takeovers of large public targets and show that about half are structured as auctions. Though our sample is larger and involves more recent takeovers, we also find - using their method to classify auctions - that less than half of takeovers are structured as auctions and that less than half of invited bidders choose to participate in takeover auctions. They use regression analysis to study market reactions to takeovers but, unlike our study, does not examine deal premia, which are more important for our purpose of asking which sale procedures lead to higher sale prices.

This paper is also related to Aktas, de Bodt and Roll (2010), who construct proxies for potential competition and in a regression context show that these proxies are positively related to observed bid premia. At the same time, they do not answer the question of how this effect impacts the relative desirability of negotiations and auctions as we do.

Our work is most similar to Gorbenko and Malenko (2013), who build and estimate a

structural econometric model on hand collected data drawn from SEC statements to recover information about the distribution of bidder valuations, and to Roberts and Sweeting (2013), who estimate a model of government timber auctions with entry. Gorbenko and Malenko (2013) take as given the size and composition of the bidder pool and seek to understand the role played by different bidder types rather than to characterize how uncertainty and endogenous entry impact the takeover markets.

The paper is organized as follows. Section I provides a background on takeover auctions with endogenous entry. Section II develops the baseline framework. Section III develops the empirical generalization of the baseline framework. Section IV describes our hand collected data and reports summary statistics. Section V characterizes takeover environments and uses estimated primitives to understand takeover market efficiency. Section VI compares the performance of auctions and negotiations. Section VII concludes.

I. Institutional Background

This section describes a typical takeover auction. Takeover auctions follow a relatively standardized format, which has been discussed extensively elsewhere (see, for example, Hansen (2001), Gorbenko and Malenko (2013)). We restrict attention to aspects relevant to our study.

An auctioning board recruits a sell side advisor to identify and contact potential bidders, i.e., firms with a possible willingness and ability to acquire the target. Potential bidders are contacted individually and invited to participate in competitive bidding for the target, and participation becomes formal when a potential bidder signs a confidentiality agreement, which determines conditions under which non public information about the target is disclosed to the bidder. Access to non public information allows the bidder to conduct due diligence (a costly examination of the target's finances, operations, and business prospects) before submitting an indication of interest, or a formal bid, since in practice acquirer valuations depend both on a common component (the stand alone value of the target) but also on asset

complementarities specific to a particular merging firm pair. Due diligence may take up to several months and involves analysis of the bidder's management and in house deal team, and also by the buy side advisor's deal team. Due diligence typically focuses on aspects of the target's non public operations relevant for valuating pair specific asset complementarities and post merger integration costs and includes analysis of supply chains, software and machine technology, R&D overlap, intellectual property, marketing programs, potential technology transfer, retiree pension and medical benefits, debt covenants, complementarities in strategic operations, customer perceptions of both companies, the compatibility of corporate cultures and other human resources, strategic reactions of competitors, and customer perceptions of the two companies, among others. From the perspective of a potential acquirer, entry into a takeover auction is thus costly, both in terms of direct pecuniary costs and advisor fees, but also in terms of non-pecuniary costs associated with foregone acquisition opportunities while negotiations are ongoing, risk of reputational capital if negotiations fail, potential revelation of proprietary information if a competing bidder wins the takeover competition, and diversion of the management, board, and deal team's time.

Information about several aspects of the sale process are specifically restricted by confidentiality agreements. An entering bidder is typically precluded from revealing the fact that the target is up for sale, the value of its indications or bids, or the fact that it is participating in the takeover auction (Kirman (2008)). Potential bidders thus decide whether to enter the auction without knowing whether other firms have entered and, upon entry, must make bids without knowing how many other firms have entered or the value of their bids. This fact dramatically simplifies analysis of entry and bidding decisions, since it eliminates confounding signaling or timing effects that would arise if entry decisions were observed concurrently by potential entrants, and it also precludes jump bidding or other activities designed to signal or deter other bidders. This fact also rules out potentially complex forms of collusive behavior (Rosenbaum and Pearl (2009)).

Competing bids are not generally disclosed, even by the target. Bidders instead receive

feedback about their offers by receiving communication from the target’s board who indicate whether a bid is “adequate” (i.e., above the target’s reservation price or in striking distance of the highest standing bid). After feedback is delivered, bidders with low bids either raise their bids or exit the auction, and remaining bidders submit more competitive bids. As before, the target may respond to a bid by indicating that a binding offer is inadequate, and this process repeats until the bidder with the highest value is identified. If the highest bid is above the target board’s reservation price, the deal is announced publicly. This bidding structure most closely resembles an ascending auction with a reserve price in which bidders successively drop out until the bidder with the highest valuation remains (e.g., Subramanian, p. 59 (2011)).

II. Baseline Model Specification

A. Information and Entry

A takeover auction j is initiated when N_j potential bidders $j = \{1, \dots, J\}$ are contacted and invited to participate in competitive bidding for a target. It well known in auction theory that expected profits are strictly increasing in the number of potential bidders, so is not surprising that, when asked to conduct a broad based auction, sell side advisors generally seek to identify and contact all available potential bidders though, as we will see, many potential bidders will endogenously decline to participate. The set of potential bidders is thus viewed as determined by exogenously given characteristics of the target and its industry (e.g., size, market positioning, industry consolidation), though of course the set of participating bidders will be endogenously determined.

Each potential bidder next chooses whether to enter the auction. Potential bidders formally enter by executing confidentiality agreements with the target and conducting due diligence at cost c_j (e.g., Hansen (2001), Boone and Mulherin (2007)).

Each of the n_j entering bidders engage in competitive bidding for the target, with bidding based on valuations discovered during the entry process. Sale occurs if the final

purchase price is greater than the target’s reservation value V_{0j} .

We now develop a tractable approach to parameterizing a potential bidder’s information about potential asset complementarities and integration costs associated with the acquisition. Let V_{ij} denote potential bidder’s i ’s valuation for target j , which is observed after a potential bidder enters, receives access to nonpublic information, and conducts due diligence. Valuations depend both on a common stand alone component (M_j) and an idiosyncratic asset complementarity net of integration costs (ν_{ij}), specific to a particular acquirer and target pair. Following Gorbenko and Malenko (2013), we specify the unconditional distribution of valuations V_{ij} among potential acquirers for target j as $V_{ij} = M_j \exp\{\nu_{ij}\}$.

The distribution of V_{ij} reflects a heterogeneity across bidders along an array of dimensions (e.g., industrial or product market similarity to the target, strategic vs financial bidders, etc.), that determine asset complementarities and integration costs, which are in practice likely to be different across bidders. The ν_{it} are drawn independently from a Gaussian distribution with sale specific mean μ_{vj} and variance σ_{vj}^2 . By allowing primitives to be sale specific, both components of target valuations are allowed to be correlated through both the vector of observed target and market characteristics (X_j) and a vector of target and market level unobservables (e.g., μ_{vj}). Our empirical implementation accommodates both forms of correlation.

Each potential acquirer i observes M_j and a private signal S_{ij} of its uncertain valuation V_{ij} prior to entry. Conventional studies of auctions with entry have focused on one of two knife-edge cases: no pre-entry information ($S_{ij} \perp V_{ij}$), which has its origins in Samuelson (1985), and perfect pre-entry information ($V_{ij} = h(S_{ij})$) for some function $h(\cdot)$, originally developed by Levin and Smith (1994). These assumptions have important implications for how a company should be sold: the assumption of no pre-entry information, employed by Bulow and Klemperer (1996) implies that potential bidders randomly choose to enter into competitive bidding, which in turn generates the famous conclusion that auctions always generate higher returns to target shareholders than do negotiations.

These polar extreme assumptions simplify auction models, but in practice the degree of pre-entry uncertainty faced by potential bidders is never directly observed, yet entry behavior (and as a consequence the relative optimality of auctions and negotiations) depends crucially on it. The framework laid out above, in contrast, implies that $S_{ij} = V_{ij} \exp\{\varepsilon_{ij}\}$, where errors ε_{ij} are Gaussian white noise with sale specific standard deviation σ_{ε_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. A potential bidder's *ex ante* pre-entry uncertainty about their valuation of the target is parameterized using the following noise-to-signal ratio:

$$\alpha_j \equiv \sigma_{\varepsilon_j}^2 / (\sigma_{v_j}^2 + \sigma_{\varepsilon_j}^2) \in [0, 1]. \quad (1)$$

This definition of the noise-to-signal ratio implies an α_j closer to 0 indicates a more informative signal, in which case potential bidders place stronger credence on their pre-entry signals.

Our general entry framework thus overcomes the need to make an extreme assumption about the average level of potential bidder's pre-entry information: rather than making an extreme assumption about of unobserved pre-entry information and entry behavior by accommodating virtually any entry structure subject to the weak constraint that higher signals on average signal higher values. By imbedding the informativeness of signal precisions into the empirical model, we are able to directly estimate it to recover the first empirical measure of the degree to which potential bidders have confidence about their relative values for the target prior to entry, which in turn allows us to characterize entry patterns in corporate takeover auctions.

B. Expected Entry Profits

This section characterizes an entrant's expected profit, which depends on equilibrium in the bidding stage. This profit will be used in the next section to characterize equilibrium in the entry stage. All variables are sale specific, though we suppress the j subscripts in this section to ease exposition.

Entering bidders compete in a standard auction for the target. Recall that bids in takeover auctions are in practice sealed in the sense that standard confidentiality agreements prevent bidders from having access to information about competing bids or the number of entering bidders. The dominant strategy in such an environment, for an entrant with value realization v_i is to bid until the current posted price reaches v_i , and to exit when the target indicates that a bid $b_i > v_i$ is required to remain in competitive bidding. This structure mirrors an ascending button auction. If no entrant has a valuation above the target's reservation price, the auction ends when the final bidder drops out, and the auction results in no sale. Otherwise, bidding continues until the purchase price reaches the maximum of the second highest entrant valuation and the target's reservation price, at which point competitive bidding concludes.

Let $F^*(\cdot|N)$ be the CDF of the equilibrium distribution of valuations among n entering bidders and let V_{0t} be drawn a target specific reserve distribution $F_0(\cdot)$ with $V_{0t} = M_t \exp\{\nu_{0t}\}$ and where ν_{0t} is normally distributed with parameters $\{\mu_{\nu_{0t}}, \sigma_{\nu_{0t}}\}$. Appendix A shows that the expected profit of an entrant with valuation v_i is given as

$$\pi^*(v_i; n, N) = \int_0^{v_i} F_0(y) \cdot F^*(y|N)^{n-1} dy. \quad (2)$$

C. Entry Behavior

This section characterizes a potential bidder's entry decision, which depends in potentially complicated ways on a potential bidder's expectations of its own valuation and

those of other bidders, the expected number of competitors faced upon entry, competitors' potential synergies with the target, and the target's reservation value. We now show how this complexity can be appreciably reduced by recognizing that a potential bidder's optimal entry decision can be characterized by a signal threshold rule where the potential bidder enters if its expected profit from participating in the auction is greater than zero.

It can be shown that any equilibrium of the form considered here has a representation in threshold strategies. An important consideration is that in the knife-edge case $S_i \perp V_i$ there may exist equilibria in which bidders can do no better than randomizing their entry decisions. Gentry and Li (2014) show that when $S_i \perp V_i$ randomization on the basis of a threshold is equivalent to any other rule for randomization, so focusing on the threshold equilibrium involves no loss of generality.

We seek a symmetric pure strategy Bayesian Nash equilibrium in which entry decisions can be characterized by a signal threshold s_N^* such that bidder i chooses to enter if and only if $S_i \geq s_N^*$. Let $F^*(\cdot; s_N^*)$ denote the distribution of valuations conditional on the event

$$S_i \geq s_N^* : F^*(v; s_N^*) = F(v|S_i \geq s_N^*) = \frac{1}{1 - F_s(s_N^*)} \int_{s_N^*}^{\infty} F(v|t) f_s(t) dt. \quad (3)$$

The following identity will be useful: for any (v, s^*) ,

$$(1 - F_s(s^*)) F^*(v; s^*) = \int_{s^*}^{\infty} F(v|t) f_s(t) dt = F_v(v) - F(v, s^*). \quad (4)$$

Independence of signals implies the total number of entrants n follows a binomial distribution

based on entry probability $[1 - F_s(s_N^*)]$:

$$\Pr(n|N, s_N^*) = \binom{N}{n} F_s(s_N^*)^{N-n} \cdot [1 - F_s(s_N^*)]^n. \quad (5)$$

Further, by construction, the probability any given entrant draws a value below v is $F^*(v|N) = F^*(v; s_N^*)$. 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 , this acquirer forecasts profits $\Pi(s_i; s_N^*, N)$. Expanding this term yields,

$$\begin{aligned} &= E_V [E_n[\pi^*(v_i; n, N)|n \geq 1|S_i = s_i] \quad (6) \\ &= \int_0^\infty \int_0^v F_0(y) \left[\sum_{n=1}^N \binom{N-1}{n-1} F_s(s_N^*)^{N-n} ([1 - F_s(s_N^*)]F^*(y; s_N^*))^{n-1} \right] dy dF(v|s_i) \\ &= \int_0^\infty \int_0^v [F_0(y) [F_s(s_N^*) + (1 - F_s(s_N^*))F^*(y; s_N^*)]^{N-1}] dy dF(v|s_i) \\ &= \int_0^\infty \int_0^v [F_0(y) [F_s(s_N^*) + F_v(y) - F(y, s_N^*)]^{N-1}] dy dF(v|s_i), \end{aligned}$$

where the third equality follows by properties of binomial series.

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

$$\Pi(s_i; s_N^*, N) = \int_0^\infty [1 - F(v|s_i)] \cdot F_0(y) \cdot [F_s(s_N^*) + F_v(y) - F(y, s_N^*)]^{N-1} dy. \quad (7)$$

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

in s_N^* by the identity

$$F_s(s_N^*) + F_v(y) - F(y, s_N^*) = F_s(s_N^*) + \int_{s^*}^{\bar{s}} F(v|t) f_s(t) dt \quad (8)$$

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

We now pause to discuss several aspects of the intuition behind this expression. First, ex ante expected profit $\Pi(s_i; s_N^*, N)$ is increasing in s_i : a potential entrant who receives a high signal is more likely to be a relatively high valuation bidder and is thus more likely to win the auction upon entry and, conditional on winning, is likely to obtain higher surplus. Second, this effect is higher when pre-entry uncertainty is low, since in that case a potential bidder places stronger credence on their signal as an indicator of realizable synergies net of integration costs.

Third, expected profit is increasing in s_N^* , all else equal, since a higher equilibrium signal threshold implies less entry by all potential bidders, which results in a smaller set of post-entry competitors, again raising a potential entrant's probability of winning and the surplus that might obtain from winning.

Finally, expected entry profits are decreasing in the set of potential bidders, since more competition immediately decreases the probability a given signal reflects a winning underlying valuation. Thus, all else equal, potential bidder i prefers a higher own signal, prefers potential rivals to have a lower probability of entry, and prefers to purchase a target situated in an industry where there is less potential competition.

We now characterize equilibrium entry. 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:

$$\Pi(s_N^*; s_N^*, N) - c = 0; \quad (9)$$

that is, a marginal potential bidder with signal $S_i = s_N^*$ must be indifferent between entering

and not. $\Pi(s_i; s_N^*, N)$ is increasing in its first argument and strictly increasing in its second, so the break even condition (9) 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 form of the entry rule, the distribution of valuations among bidders choosing to enter in equilibrium is $F^*(v; s_N^*) = F(v|S_i \geq s^*)$. Thus the signal threshold s_N^* is sufficient to characterize equilibrium entry and bidding behavior.

III. Empirical Generalization

We now generalize the framework to develop a structural empirical model. First, define the observed deal premium obtaining from auction of target j as the sale price normalized by the target's market value four weeks prior to announcement, and denote this variable as p_j . Let $sale_j$ be an indicator variable taking a value of unity if an initiated auction results in sale. Our aim is to recover information about the sale specific vector of fundamental primitives $\theta_j = \{\mu_{vj}, \sigma_{vj}, \mu_{0j}, \sigma_{0j}, c_j, \alpha_j\}$, conditional on observing auctions resulting in sale.

Our empirical approach addresses two empirical challenges. The first challenge is that any study of observed takeover auctions is conducted on a nonrandom sample due to the existence of unobserved failed auctions not resulting in public announcement. While manually collecting data on the pre-announcement sale process, we encountered reports of previous failed sale attempts. The omission of failed auctions can lead to unobserved differences between the distribution of potential entrant valuations and the distribution of entering bidders' valuations even after explicitly conditioning on target characteristics since the distribution of target characteristics conditional on sale differs from the unconditional unobserved distribution of target characteristics. This form of sample selection operates through two channels. First, conditional on any realization θ_j , selection based on sale increases the likelihood of observing higher sale prices and the likelihood of observing a higher number of entrants. Second, since the vector θ_j is heterogeneous across targets in the population, the distribution of θ_j conditional on sale would, without correction, be biased toward realizations

that increase the conditional probability that a sale occurs.

The second empirical challenge is that while our empirical approach circumvents endogeneity in the choice of sale procedure by taking as given the decision to conduct an auction, the sale-level primitives are imperfectly predicted by observable characteristics, which could lead to bias in the estimated parameters. These concerns are mitigated by the fact that we explicitly control for sale level observables, but we cannot rule out the possibility that unobserved sale specific characteristics play a non trivial role in determining entry patterns. For example, better targets thus attract more entry and at the same time lead to higher prices. Failure to correct for this potential endogeneity in target characteristics would lead to bias in the estimated fundamental parameters.

To see how we address this concern, allow targets to differ both along observable dimensions (X_j) and along unobservable dimensions that influence sale-level primitives θ_j but which are never directly observed. To flexibly accommodate both forms of heterogeneity, we allow θ_j to be drawn unobservably from a joint distribution $g(\cdot)$, which depends on X_j through a vector of parameters Γ . This approach can provide an arbitrarily precise approximation of a traditional likelihood framework in which θ_j is deterministic given X_j , while simultaneously accommodating the far more important case where sale characteristics differ in ways fundamentally unobservable to the econometrician.

To gain intuition, first consider how we would implement estimation when θ_j is completely determined by X_j (e.g., $\theta_j = X_j\Gamma$). Let $\Pr(\text{sale}_j|N_j, \theta_j)$ denote the probability an auction with N_j bidders for a target with characteristics θ_j results in sale, and $\Pr(n_j, p_j, \text{sale}_j|N_j; \theta_j)$ denote the probability of the joint event “ n_j of N_j bidders enter, the sale price is p_j , and the auction results in sale” for a target with characteristics θ_j . Setting $\theta_j = X_j\Gamma$, the likelihood contribution of target j at parameters Γ would then be:

$$L(p_j, n_j|\text{sale}_j, N_j, X_j, \Gamma) = \frac{\Pr(p_j, \text{sale}_j, n_j|N_j; \theta_j = X_j\Gamma)}{\Pr(\text{sale}_j|N_j; \theta_j = X_j\Gamma)}. \quad (10)$$

Our bidding model yields analytic forms for $\Pr(\text{sale}_j|N_j, \theta_j)$ and $\Pr(n_j, p_j, \text{sale}_j|N_j; \theta_j)$ (see Appendix B). Thus estimation of Γ would simply reduce to maximizing (10) across targets in the sample.

Now consider estimation in the more realistic case when sale primitives θ_j are only imperfectly predicted by X_j . In this case, as is well known in the auction literature, failure to account for variation in θ_j over and above that predicted by X_j overestimates the variance of valuations among bidders and therefore lead to upward biased estimates of parameters such as entry costs (e.g., Krasnokutskaya (2012)). More generally, the econometrician is unlikely to observe all factors determining target values, while some of these may be observable to entrants. Explicitly accounting for such unobserved sale heterogeneity thus adds an additional dimension of robustness to the analysis. We proceed as follows. Let $g(\theta|X_j, \Gamma)$ denote the distribution of θ_j given X_j at parameters Γ . Integrating out unobserved θ_j and adjusting for selection of targets on the basis of sale, we thus ultimately obtain the observable sale likelihood function

$$L_j(p_j, n_j|\text{sale}_j, N_j, X_j, \Gamma) = \frac{\int \Pr(p_j, \text{sale}_j, n_j|N_j; \theta_j) g(\theta_j|X_j, \Gamma) d\theta_j}{\int \Pr(\text{sale}_j|N_j; \theta_j) g(\theta_j|X_j, \Gamma) d\theta_j}. \quad (11)$$

As above, our primitives of interest are the parameters Γ , which describe the relationship between parameters θ_j and observables X_j . Accounting for unobservable sale-level differences, however, these now describe the entire *distribution* of θ_j given X_j rather than just its mean. Maximization of the likelihood function (11) with respect to Γ yields an estimate $\hat{\Gamma}$ robust to both rich unobserved sale-level heterogeneity and sample selection based on sale. In turn, since these characterize the entire distribution $g(\theta|X_j, \Gamma)$ rather than just its mean, we ultimately obtain the ability to perform counterfactuals along both observable and unobservable dimensions of sale-level heterogeneity. This is a novel feature in our analysis and yields a rich set of predictions not available with a traditional maximum likelihood

analysis.

Direct evaluation of the likelihood function (11) using standard simulation procedures is computationally prohibitive. We circumvent this difficulty by applying the importance sampling procedure proposed by Akerberg (2009) and implemented by Roberts and Sweeting (2013). Conceptually, this procedure involves drawing a large random sample of primitives $\{\theta_r\}_{r=1}^R$ from any proposal density $\tilde{g}(\cdot)$; standard choices for $\tilde{g}(\cdot)$ would be normal or uniform distributions, though convergence implies the estimated parameters are insensitive to the initial proposal distribution. Taking logs of (11) and simulating integrals by the importance sample $\{\theta_r\}_{r=1}^R$ from $\tilde{g}(\cdot)$, we obtain the following tractable sale-level log likelihood, maximization of which yields our estimated parameters:

$$\begin{aligned} \ln L(p_j, n_j | \text{sale}_j, N_j, X_j, \Gamma) &= \ln \left(\sum_{r=1}^R \Pr(p_j, \text{sale}_j, n_j | N_j; \theta_r) \frac{g(\theta_r | X_j, \Gamma)}{\tilde{g}(\theta_r)} \right) \\ &\quad - \ln \left(\sum_{r=1}^R \Pr(\text{sale}_j | N_j; \theta_r) \frac{g(\theta_r | X_j, \Gamma)}{\tilde{g}(\theta_r)} \right). \end{aligned} \quad (12)$$

Note given $\{\theta_r\}_{r=1}^R$, (12) depends on Γ only through the density function $g(\cdot | X, \Gamma)$, which leads to significant computation savings relative to the initial formulation (11), since we need only compute $\Pr(p_j, \text{sale}_j, n_j | N_j; \theta_r)$ and $\Pr(\text{sale}_j | N_j; \theta_r)$ once for each θ_r .

We now pause to conceptually describe how maximization of this function recovers sale-level primitives (a more formal and detailed description can be found in Appendix B). The initial draws $\{\theta_r\}_{r=1}^R$ yield a sample of *hypothetical* targets, where (as usual in importance sampling) weighting by $1/\tilde{g}(\theta_r)$ corrects for the fact that these are drawn from \tilde{g} rather than g . Maximization of (12) with respect to Γ is then equivalent to choosing *weights* $g(\cdot | X_j, \Gamma)$ on these hypothetical targets to maximize the likelihood of the observed data. In other words, we first generate a universe of possible targets and next choose Γ to select the subset of empirically relevant targets. Akerberg (2009) shows $\hat{\Gamma}$ to be a consistent estimator of Γ .

IV. Data and Summary Statistics

We analyze a set of corporate takeovers announced between January 1, 2000 and December 31, 2009 drawn from the Securities Data Corporation (SDC) mergers and acquisitions database and satisfying the following set of conditions:

- The target is a publicly traded U.S. company listed in the S&P 1500
- The deal value greater than \$1 million
- The acquirer owns 100% as a consequence of the deal
- Financial data on the target is available from Standard and Poor's *Compustat* database

We used proxy statements submitted by the target or acquirer to the Securities and Exchange Commission (SEC) to manually collect information on the number of potential and participating bidders in each auction.² For a takeover to be included in the sample, we required these background sections to be available on the SEC Edgar online filing system.

We required data on the final sale price relative to the target's share price four weeks prior to announcement - the deal premium - to be available from Thomson's *SDC Platinum*. Information on winning bids was manually recorded from takeover press releases and proxy statements and used to cross check reported premia data from SDC. In our context, possible misvaluation of reported stock bids would show up as a form of measurement error that would be captured by our estimated heterogeneity distributions, described below. As a robustness check, we also estimated our model parameters on the sample of all cash bids and found similar results.

This sampling procedure yields 980 takeovers. Following Boone and Mulherin (2007) and Gorbenko and Malenko (2013), we define participating bidders as those signing a confidentiality agreement with the target, and we classify takeovers as auctions if multiple buyers

²This information is 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>

were contacted by the target. Table 1 reports the number of auction and negotiated transactions for each year in our sample.

[[[**Insert Table I About Here**]]]

Target characteristics come from the *Compustat* database. We obtain information on firm size defined as the book value of the target's total assets, market leverage, q ratio (market to book), cash and intangibles relative to target book assets, and the 4-digit SIC code. We follow standard assumptions used in the corporate finance literature to filter out implausible or unreasonable values. Specifically, we exclude from the sample observations with q ratio in excess of 10, instances in which the ratio of the winning bid to the target's value under the current management is less than unity and instances with bid premia above 200%. We also exclude nonclassifiable establishments from the sample. After applying these filters, our estimation sample contains 565 takeover auctions.

Table 2 reports the average number of potential bidders, entrants, and deal premia, for auctions, negotiations, and the full sample. Negotiations and auctions yield similar average deal premia. Almost 60% takeovers involve targets sold via auction, and among these only 43% of invited potential bidders elect to participate in competitive bidding for the target. Limited participation is thus an important stylized feature of the data.

[[[**Insert Table III About Here**]]]

Table 3 reports average characteristics of targets sold via auction and negotiation, and shows that these firms are very similar in their market to book ratios, cash to asset ratios, leverage, and intangibles to asset ratios. The top row shows that large targets tend to be sold via negotiation, with targets sold via auction having total assets averaging \$1.60 billion while targets sold via negotiation having total assets averaging \$2.95 billion. This size differential is consistent with the view that because there may exist only a single exogenously given potential bidder with the ability to finance the deal or integrate the target into its operations, the largest targets will be sold via negotiation.

V. Characterizing Takeover Environments

In this section, we estimate the structural model to recover distributions of the sale-level primitives that characterize the takeover environments in which both auctions and negotiations take place. Before proceeding, we pause and ask what new information these estimates might convey about takeover markets.

First, information about these primitives cannot be recovered via standard regression techniques, since the relevant quantities such as the average degree of pre-entry uncertainty are never directly observed, yet these primitives are required to characterize revenue that would accrue to targets under the counterfactual scenario. A potential bidder's entry decision, for example, depends on the entire distribution of potential bidders values, yet neither observed winning bids nor direct proxies for potential competition are sufficient to recover this distribution of non-entrant or losing bidder valuations, so a structural approach is required to characterize how entry and competition impact takeover markets takeover markets.

Second, our framework's generality nests as possibilities many different takeover environments, and the estimated parameters can provide a characterization of this environment in practice. As an example, under maximal pre-entry uncertainty where potential bidders beliefs are pure noise, entry decisions are random and the distribution of entering bidder willingness to pay is identical to the distribution of potential bidder willingness to pay, and endogenous entry is irrelevant for characterizing takeover auction environments. Conversely, if pre-entry uncertainty is extremely low and bidders approximately know their value for the target prior to entry, then only the bidders with the highest willingness to pay will enter, and in some cases only the highest-valuation bidder will enter, but in this case an auction would in the data appear as a single-bidder negotiated sale. A structural econometric framework is thus required both to recover unobservable primitives but also to characterize the competitive environment in which takeover auctions operate.

A. Recovering Fundamental Takeover Market Primitives

We begin by recovering estimates of the fundamental parameter vector θ via maximization of (12) over the data vector $(n_j, p_j, X_j, N_j, sale_j)$ using the method described in Section III. Table IV reports quantiles of the estimated sale-level parameter distributions evaluated at mean values of observables and median values of the heterogeneity distributions (Panel A) and at quantiles of the posterior likelihood evaluated at median observables (Panel B).

[[[**Insert Table IV About Here**]]]

Table IV reports the results. Column (2) of Panel A shows that potential bidders differ dramatically in their ex post valuations for the target: The standard deviation of potential acquirer valuations (σ_{vj}) is 16%. This dramatic variation across potential acquirers in realizable synergies implies that the identity of the ultimate buyer potentially plays an important role in determining the ability of takeover markets to create value. Limited participation and endogenous entry are thus likely to play an important role in determining the extent to which auctions extract value for target shareholders.

The mean of the potential entrant valuation distribution in Column (1) of Panel A is 0.19: the average maximal willingness to pay of potential bidders is a 19% deal premium. Notice that this figure is significantly below the 42% mean deal premium observed in the data: ultimate acquirers have a much higher willingness to pay than randomly-selected potential bidders. Though this finding provides some evidence that corporate takeover sale processes are relatively effective in matching targets with high-value bidders, it does not by itself indicate whether even higher deal premia might be obtained in for example a world of perfect pre-entry information.

Column (4) reports moments of the distribution of pre-entry uncertainty. The mean estimate of α in Panel A is 0.64, which along with equation (1) implies that potential bidders pre-entry beliefs contain more noise than information, and that the extreme assumptions of perfect information ($\alpha = 0$) and maximal uncertainty ($\alpha = 1$) are not born out by the data.

The existence of non-trivial pre-entry uncertainty implies that some potential entrants are unaware about the existence of high asset complementarities (or conversely, low integration costs) achievable through acquisition of the target and may not elect to participate in competitive bidding. Thus, even if takeover markets are efficient in all other aspects, the “right buyer” is regularly absent from the pool of competing bidders. Takeover markets thus generate less value than they would in a world of perfect pre-entry information. When considered together with our finding that potential bidders’ valuations are widely dispersed, this raises the possibility that high pre-entry uncertainty substantially inhibits the ability of takeovers to create value for target shareholders, an issue we formally explore in Section V.B.

A strength of the structural econometric approach is that it allows recovery of the entire distribution of takeover market primitives, not just their average effect across all targets. Rows (3) and (4) of Panel A show that the aggregate estimates just described conceal dramatic variation across targets in the parameters that characterize individual takeover markets. For example, α is less than 0.45 for 25% of targets while also being greater than 0.86 for 25% of targets. Heterogeneity in the cross-section of takeover markets is thus an important feature of the data.

The mean of the entry cost distribution, reported in Column (3), also varies in the cross-section of targets and is equal to about 1% of the deal value. As a comparison, advisory fees charged to acquirers average approximately 0.5% of deal value. The next section characterizes the extent to which the presence of pre-entry uncertainty inhibits the ability of takeover auctions to obtain high sale prices.

B. Characterizing Takeover Auctions

This section characterizes endogenous entry patterns to ascertain the potential negative effect of pre-entry uncertainty on the ability of takeover auctions to induce participation in takeover auctions to generate a large and attractively-composed pool of entering bidders.

Pre-entry uncertainty impairs potential bidders’ ability to assess their values for the

target. Some potential bidders will have initially unfavorable beliefs about potential value creation and thus about their ability to win a bidding war upon entry. These potential bidders will decline to participate in a takeover auction, but had they elected to enter a few of them would have discovered high valuations for the target. While the decision to decline participation was ex ante rational it would have been suboptimal ex post. Our structural estimates allow us to calculate the probability that the “right buyer” (i.e. the potential entrant with the highest valuation) declines to enter a takeover auction. We compute this probability for a typical takeover environment with $\alpha = 0.64$ and find that the “right buyer,” declines an invitation to participate in competitive bidding about 36% of the time.

Given a fixed size of the entering bidder pool, an auction generates higher expected revenue if the entering bidder pool contains relatively more high bidders with a high willingness to pay. We refer to the influence of pre-entry uncertainty on the average quality of the bidder pool as the “composition effect” and define it as follows. The composition effect is measured as the difference between the expected price that would maintain if the distribution of the entering and potential bidder pools were identical.

To understand the importance of our definition of the composition effect, notice that the distribution of potential and participating bidders would be identical only when potential bidders’ decisions to enter were random, and as we have shown this occurs only in the limiting case where pre-entry signals are pure noise, i.e. when $\alpha = 1.0$. Our structural estimates identified $\alpha = 0.64$ as the typical case, so in practice entry decisions are non-random and exhibit disproportionate entry by potential bidders with relatively favorable pre-entry beliefs. Holding the size of the entering bidder pool fixed, the composition effect thus quantifies this effect.

To formally examine the composition effect we hold the overall level of entry constant and compute the probability that potential entrants with different values will elect to participate in competitive bidding. Specifically, we compute the relative probability the p th percentile highest value bidder enters for $\alpha = 0.64$. We next re-compute these entry

probabilities at various levels of pre-entry uncertainty. For tractability in what follows in this section we focus on a typical takeover auction, which henceforth indicates one defined by median values of σ_{vj} , and c_j with $V_{0j} = 1$, and mean target-level characteristics, which includes observables and our estimates of unobserved (Γ_j) sale-level heterogeneity and μ_v set to match the mean of observed and model deal premia. Our main results are robust across a wide array of parameter and covariate estimates. Section IV.C. below explores sale-level heterogeneity in more detail.

[[[**Insert Figure I About Here**]]]

Figure I reports the relative probability that different types of bidders elect to enter. As uncertainty rises, relatively high valuation bidders become more likely to participate, while relatively low valuation bidders become less likely to participate. This is because as uncertainty increases, relatively high valuation bidders frequently receive lower signals and choose not to participate while relatively low valuation bidders more frequently receive higher signals and choose to participate. Thus, holding the size of the entering bidder pool fixed, an increase in pre-entry uncertainty relative to its average level degrades the overall quality of the entering bidder pool.

In Table V we report estimates of the composition effect for different levels of pre-entry uncertainty and entry costs. For comparability with observed deal premia the estimates are expressed as a percent of the target’s share price. Looking across the first row, we see that the effect of endogenous entry on the *quality* of the entering bidder pool accounts for anywhere between 10.1 and zero percentage points. The ability to generate a well-composed pool of bidders is thus a quantitatively important determinant of observed deal premia. The most striking feature of this table is the monotonic decline in the composition effect in pre-entry uncertainty: as the informativeness signals becomes more noisy, the “right buyers” become less frequent participants in takeover auctions, and this reduces the ability of auctions to generate shareholder value.

[[[**Insert Table VI About Here**]]]

The finding that uncertainty leads to less entry by relatively high valuation bidder could be interpreted as implying that uncertainty unambiguously decreases expected takeover auction revenue and that the auction format is a poor way to sell an opaque target. Accounting for endogenous entry, however, this need not be the case since higher uncertainty could encourage overall entry even as it degrades the quality of the entering bidder pool.

To understand this “size effect,” note that in a world with pre-entry uncertainty, the ultimate absence of high-valuation bidders from the entering bidder pool raises other potential bidders’ expectations about their own prospects for winning and for paying lower prices upon entry. Formally, we define the size effect for an auction with n entering bidders as the effect of uncertainty on the expected deal premia holding constant the composition of the entering bidder pool, which we quantify relative to the distribution of the potential bidder pool. This definition implies that the observed takeover premium is equal to the sum of the composition and size effects.

[[[**Insert Figure II About Here**]]]

To quantify the size effect we iteratively compute the probability that individual invited bidders choose to enter, determined by the break even threshold (9) and combine these to obtain the expected fraction of invited potential bidders that elect to participate in the takeover auction. We begin by computing this quantity for the observed mean of uncertainty, $\alpha = 0.64$, and then iteratively alter α and re-compute the entry probabilities.

Figure II reports how the degree of participation changes with pre-entry uncertainty. The fraction of entering bidders increases monotonically in pre-entry uncertainty, with about half of invited bidders choose to participate when $\alpha = 0.64$. Over 90% of invited bidders participate in auctions of highly opaque targets, i.e. those for which α is close to one.

These results show that pre-entry uncertainty has two opposing effects on expected takeover revenue. Though pre-entry uncertainty leads to degradation in the average quality

of the entering bidder pool, it also leads to more entry overall, though without further information it is not clear which effect dominates in practice, i.e. whether and to what extent pre-entry uncertainty lowers takeover auction revenue, overall.

Does the high level of uncertainty present in takeover markets impair the ability of auctions to generate value for target shareholders? It is tempting to conclude that the answer to this question is “yes,” since auctions generate high revenue when high-value bidders are aware of possible asset complementarities and elect to enter (i.e., the “composition effect”) Yet the logic developed in this paper shows that when entry is endogenous and costly, uncertainty encourages entry by presenting *all* potential bidders with a competitive environment where losing the takeover auction is not a foregone conclusion.

[[[**Insert Figure III About Here**]]]

In Figure III, we examine how the size and composition effects together affect average expected takeover revenue. This figure shows that target shareholder revenue is monotonically increasing in pre-entry uncertainty. Thus, at high levels of uncertainty the positive effect of uncertainty on overall entry is larger than the negative effect of uncertainty on the quality of the entering bidder pool. This somewhat counterintuitive finding that target opacity raises expected auction revenue is in fact consistent with the conventional view that the virtue of takeover auctions is their ability to generate a large pool of bidders: the size effect is relatively large in uncertain takeover environments. Accounting for endogenous entry shows that this positive size effect comes hand in hand with the composition effect, so failing to account for its negative effect on takeover revenue will lead a researcher to draw unrealistically favorable conclusions about the ability of takeover auctions to generate revenue for target shareholders. Figure III shows that this upward bias becomes more extreme for less opaque targets.

Another important insight to draw from Figure III relates to the widespread practice of targets and their advisors to disseminate only public or vague information to potential bidders in pitch books as part of an invitation to participate in a takeover auction (e.g.,

Hansen (2001)). One might suppose that disseminating more information about the target's operations might reveal to potential bidders the existence of asset complementarities that might otherwise not be noticed, thus encouraging entry. The dominance of the size effect explains why such a policy might not be a good idea, since dissemination of more information would discourage entry, overall.

VI. Comparing Auctions and Negotiations

The takeover market primitives that determine the effectiveness of auctions also characterize how well negotiations generate revenue for target shareholders, since for a particular target both auctions and negotiations occur in the same competitive environment. In this section we use the estimated primitives to assess whether takeover premiums would have been different if auctioned targets had instead been sold via a negotiation. Before reporting the results, it may be instructive to consider what could be inferred from an approach that regressed observed takeover premiums on a variable that indicates the type of sale method chosen. There are two reasons why such an approach cannot provide information that is sufficient to answer the question about how a firm should be sold.

First, for any particular target, the relative benefits of auctions and negotiations on quantities not directly observable to researchers. For example, observed winning bids by themselves are not by themselves informative about the degree of potential competition, which is captured by the distribution of potential and actual entrant values. As a second example, the impact of additional potential entrants on bidding aggressiveness in negotiations depends crucially on both entry costs and the degree of uncertainty about potential synergies, neither of which are directly observed in the data. Observed deal premia may thus appear similar across auctions and negotiations in cross-sectional averages while one sale mechanism unambiguously revenue dominates the other for a particular target or a for a subset of targets.

Second, unlike auction sales, which have a fairly standardized structure, sales classified under the “negotiations” umbrella can in practice take a variety of specific forms that

are observationally indistinguishable from *ex post* data provided in SEC takeover proxy filings, which complicates the task of interpreting any results obtained by lumping together a heterogeneous set of single bidder sale procedures.

The structural approach circumvents both challenges: it is able to recover information about the unobservable primitives that characterize expected target shareholder revenue in both auctions and negotiations, and it does so by recovering from data on takeover auctions those features of the takeover environment such as average pre-entry uncertainty.

A. Negotiation Formats

This section formalizes two negotiation frameworks that have found support in the finance and economics literatures and whose expected revenue are determined by the fundamental takeover market primitives estimated in Section V.A. The first, a one-shot negotiation followed by a market check is a simple realistic alternative to a formal takeover auction that is used widely in practice (e.g., Subramanian (2008), Wasserstein (2000)). The second, which has been extensively studied in the theoretical and empirical literature on optimal sale design allows the target to terminate negotiations and to approach an additional bidder if a standing bidder’s best offer is not deemed adequate (e.g., Fishman (1988), Betton and Eckbo (2000), Horner and Sahuguet (2007), Dimopoulos and Sacchetto (2011), Roberts and Sweeting (2013)). Importantly, each of these negotiation frameworks incorporate the insight that in practice a standing bidder in a negotiation is not fully shielded from competition, since the target has the outside option of rejecting the negotiating bidder’s best offer as inadequate and negotiating with a different bidder instead (e.g., Aktas, de Bodt and Roll (2010), Bulow and Klemperer (2009)).

We now define a standard one-shot negotiation followed by a market check (a “go shop”) as follows: the target approaches potential buyer j with an invitation to participate in negotiations. Based on own signal S_j and the entry cost c , potential buyer j determines whether to enter the negotiation. Conditional on choosing to enter, potential buyer j learns its valuation V_j 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, the bid is publicly announced and all potential buyers are invited to simultaneously enter and make a higher bid for the target. Based on own signal S_j and the entry and the posted price acquirer j determines whether to enter the negotiation at cost c and make a bid. The bidder with the highest bid acquires the target.

The generalized sequential negotiation procedure proceeds in N rounds, one for each potential buyer. The sequence of events within each round n is as follows:

1. The target approaches potential buyer j with an invitation to participate in a negotiated transaction. Based on own signal S_j and the entry and bidding history of the game to date, potential buyer j determines whether to enter the negotiation at cost c .
2. Conditional on choosing to enter, potential buyer j learns its valuation V_j . If another negotiating bidder has previously entered, potential buyer j 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 the level at which the loser drops out.
3. Conditional on outbidding the current incumbent, potential buyer j 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 $n + 1$.

The separating equilibrium is one for which the jump bid submitted by a new incumbent at time j communicates the current standing value V_j to all subsequent potential entrants. For such separating behavior to be an equilibrium, potential buyer j must prefer truthful bidding based on V_j to impersonating any other valuation Z , which in turn implies a set of local best response conditions which must be satisfied by equilibrium bidding behavior. These restrictions uniquely define a separating equilibrium of the form desired, and that this is the only equilibrium to survive the standard D1 refinement of Cho and Kreps (1987). Intuitively, the D1 refinement specifies that beliefs place positive weight only on the types

“most likely” to deviate. More formally, for any two types v_j and v'_j and any equilibrium bid function $\beta(\cdot)$, if type v'_j strictly prefers bid $\beta(z)$ whenever type v_j weakly prefers bid $\beta(z)$, then equilibrium beliefs upon observing bid $\beta(z)$ should place no weight on type v_j . The D1 refinement is standard in analysis of signaling games, and yields a unique equilibrium whenever (as here) the underlying equilibrium payoffs satisfy an appropriate single crossing condition. As in the auction setup for our estimated model, the target has reservation value V_0 drawn from $F_v(\cdot; \theta)$, with the realization v_0 representing the point of departure for Round 1 bidding.

B. Comparing the Average Performance of Auctions and Negotiations

We begin by examining the average performance of auctions and negotiations and in the next section examine whether the same findings hold for all targets.

We first construct the unconditional posterior revenue distribution in the cross-section of targets for each of the three sale mechanisms described above by obtaining information on the primitives associated with the sale of target j and drawing the vectors (X_j, θ_j) from the prior likelihood function. We next use the takeover environment defined in Section II and fundamental parameters from Section V.A. to construct average expected revenue under each sale procedure. We obtain the unconditional posteriors of the resulting revenue functions, $R(X_j, \theta_j, \text{Sale Procedure})$, by repeating this procedure 10,000 times.

[[[**Insert Table VI About Here**]]]

Table VI reports moments of the resulting posterior distributions. Expected deal premia presented in the first row are 39.9% for auctions and are 40.2% for one-shot negotiations followed by a market check. This difference is extremely small, indicating that neither sale method obviously dominates in terms of expected revenue. This finding is consistent with the current jurisdictional interpretation of Delaware’s *Revlon* ruling that views the target’s board of directors as “auctioneers charged with getting the best price for the stockholders,” and where this duty can be satisfied when a negotiated sale is followed by a formal market

check.

Sequential negotiation revenue on the other hand is 3.5% higher than revenue from auctions. This number implies that a target of average size in our sample would be expected to obtain an additional 30.5 million dollars by structuring the sale as a sequential negotiation, rather than as a standard auction.

We next ask how the optimality of auctions or negotiations varies in the cross-section. To do this, we first define the revenue difference function $D(X_j, \theta_j, \text{Sequential}) = R(X_j, \theta_j, \text{Auction}) - R(X_j, \theta_j, \text{Sequential})$ and $D(X_j, \theta_j, \text{Market Check}) = R(X_j, \theta_j, \text{Auction}) - R(X_j, \theta_j, \text{Market Check})$, which measure for a particular target the difference between expected revenue obtaining in an auction, relative to revenue generated from a negotiated sale. The unconditional posterior distributions $D(X_j, \theta_j, \text{Sequential})$ and $D(X_j, \theta_j, \text{Market Check})$ are formed by drawing a vector, (X_j, θ_j) and constructing revenue for each of the sale mechanisms, for each target, and iterating this procedure 10,000 times.

[[[**Insert Figure IV About Here**]]]

Figure IV displays the cumulative distribution of the revenue difference functions. The blue line is the CDF of revenue difference between auctions and sequential negotiations and the green line is the CDF of revenue difference between auctions and one-shot negotiations followed by a market check. The mass of both distributions lies below zero, indicating that expected revenue is higher under negotiations for most targets: relative to an auction 75% and 77% of targets would obtain higher revenue from sale via a sequential negotiation and a one-shot negotiation followed by a market check, respectively. The revenue differences are extremely small for many firms, particularly for the one-shot negotiation followed by a market check, so in practice for many firms the difference between an auction and a negotiation is not quantitatively important. The revenue difference function for sequential negotiations, on the other hand, implies that this sale mechanism would yield a 3% increase in deal premia for almost 40% of targets, relative to auctions. At the same time, auctions generate substantially higher deal premia for a very small number of targets. Heterogeneity across firms in relative

optimality of auctions and sequential negotiations is thus an important feature of takeover markets, and so a single prescription is not appropriate for all targets.

C. Auctions and Negotiations for Different Targets

We now explore how endogenous entry patterns influence the relative performance of auctions and negotiations. As we have seen, in auctions potential bidders' participation decisions - and thus expected revenue - are crucially influenced by pre-entry uncertainty and the costs of overcoming it. As we will see, in negotiations, the strength of deterrence bidding depends systematically on target opacity. We thus begin by regressing the sequential negotiation revenue difference function $D(X_j, \theta_j, \text{Sequential})$ on polynomials of sale-level estimates of uncertainty and entry costs, α_j and c_j . The R-squared coefficient in this regression is 43%, indicating that pre-entry uncertainty and the costs of overcoming it can jointly explain almost half of the variation across targets in the relative performance of auctions and negotiations.

To explore this result, we use the estimated takeover market primitives to compute expected revenue differences across auctions and sequential negotiations for different values of α and c . We calculate the break even threshold for auctions and for sequential negotiations, we solve for the symmetric separating Perfect Bayesian equilibrium by backward induction, following the algorithm of Roberts and Sweeting (2013). In short, at each step, we first solve a differential equation induced by the first order condition of the bidding problem to find the optimal bid function in the separating PBE, next compute equilibrium expected profit induced by this bidding function, and finally find the break even signal threshold induced by the continuation play already determined.

Figure V presents the results. Red circles represent outcomes in which expected auction revenue is higher and filled blue circles represent cases in which sequential negotiations dominate. Hollow circles indicate cases in which simulation error is larger than estimated revenue differences. There are two main takeaways from this figure. First, expected revenue is generally higher for negotiations when entry costs are high. This is because high entry

costs disproportionately reduce entry in auctions, thus reducing their ability to produce a competitive pool of entering bidders.

[[[**Insert Figure V About Here**]]]

Second, in general auctions revenue dominate negotiations when pre-entry uncertainty is high. We have already seen that pre-entry uncertainty induces entry into competitive bidding in auctions, overall (Figure 2), thus increasing expected revenue in auctions (Figure 3). At the same time, the strength of negotiations is primarily is through the incentive they place on a standing negotiating bidder to shade up their offer price upward to deter entry by potential competitors. Yet pre-entry uncertainty reduces the effectiveness of such deterrence bidding, and in equilibrium a standing bidder shades up their offer price less when pre-entry uncertainty is higher. To quantify this second effect we estimated the average magnitude of upward bid shading (i.e., the deterrence effect) in sequential negotiations for different levels of pre-entry uncertainty for a typical target. The part of the deal premium attributable to deterrence bidding is 18.8% (just under half of realized deal premia) when pre-entry uncertainty is near-zero ($\alpha = 0.01$), but it falls monotonically to 7.9% as pre-entry uncertainty increases to its mean ($\alpha = 0.64$) and continues to fall in even more uncertainty takeover environments.

[[[**Insert Figure VI About Here**]]]

In Figure 6, we fix fixing c and compute expected revenue in auctions (blue line), go-shop negotiations (red line), and sequential negotiations (blue line) for different levels of pre-entry uncertainty.³The mean revenue difference across sale procedures is represented by the vertical distance between the lines with the average level of pre-entry uncertainty in

³We follow the computational procedure described above and in addition, for the negotiation followed by go shop, we first solve the entry equilibrium in the auction stage assuming separating equilibrium play in the first stage (i.e., assuming that the incumbent's value is revealed by his jump bid). We then solve for the first stage jump bidding function that induces truthful revelation by the first stage incumbent as in the sequential negotiation mechanism.

our sample marked by the dashed vertical line. Auctions raise the most revenue for opaque targets while sequential negotiations raise the most revenue for less opaque targets while go-shop negotiations fall somewhere in between.

VII. Conclusion

The ability of takeover auctions to generate value for target companies requires a well-composed and large pool of participating bidders, yet fewer than half of invited bidders participate in takeover auctions. This paper studies how endogenous entry patterns determine takeover revenue. Our estimates reveal that potential entrants in takeover auction environments face high pre-entry uncertainty about their values for a target, with beliefs containing more noise than information. Yet even as it leads to a degradation in the average quality of the entering bidder pool, high uncertainty encourages entry overall. Takeover auctions are thus highly resilient mechanisms for generating value for the shareholders of opaque targets. Conversely, the negotiations generate high prices when they encourage aggressive deterrence bidding, yet such upward bid-shading is more effective when pre-entry uncertainty is low, so negotiations are most effective for less opaque targets. These effects are hidden by cross-sectional averages, which show that auctions and go-shop negotiations produce similar revenue and that auctions and sequential negotiations differ in their revenue-generating ability by only 3%. These findings suggest an important role for future research that explores which sale mechanisms work for specific classes of target companies, rather than searching for a “one-size-fits-all” prescription. Our findings also indicate that while auctions often provide an intuitive contrast to negotiations, it matters which negotiation procedure a target actually employs.

We have found that even though the “right buyer” often elects not to participate in competitive bidding, takeover auctions are significantly better at efficiently matching acquirers and targets than would random assignment. Further research might attempt to quantify this effect and explore the ability of different sale mechanisms to efficiently allocate capital in the

market for corporate control, in addition to exploring how entry affects target shareholder revenue as we do.

Our methodology has uncovered a counterintuitive trade-off between maximizing revenue to target shareholders and maximizing overall takeover value creation: for many targets a decrease in uncertainty leads to higher value creation while also lowering expected deal premia. This finding rationalizes the widespread practice of withholding non public information from potential bidders prior to their entry into competitive bidding (Hansen (2001)). Thus, our estimates provide a new rationale for why targets may withhold or even obfuscate nonpublic information before a potential acquirer elects to participate in competitive bidding by signing a confidentiality agreement. Further researchers might attempt to quantify the relative impact of various factors that determine the amount of information disclosed by targets to potential bidders prior to their entry into takeover auctions.

By estimating fundamental takeover market parameters only on sales that were clearly structured as auctions, our paper examined how a firm should be sold to generate the highest expected price, and we found that some companies should be sold via auction while others might be sold via a negotiation. Further research might examine how characteristics of key actors who work for the target company, such as the CEO, affect the willingness of specific targets to structure their sale in a way that maximizes expected revenue for target shareholders.

Though our analysis focused on expected returns, we also obtained estimates of the risk of takeover premia and we found in Table VI that the standard deviation of expected revenue differs substantially across different sale mechanisms. Future research might explore more deeply how possible risk-return properties of different sale mechanisms influence the decision to conduct an auction or a negotiation.

Appendix A. Obtaining an Expression for Expected Profits Conditional on Entry

Let $Y_{k:n}$ denote the k th 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

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

We assume V_{0t} is drawn from a sale-specific reserve distribution $F_0(\cdot)$ with $V_{0t} = M_t \exp\{\nu_{0t}\}$ and where ν_{0t} is normally distributed with parameters $\{\mu_{\nu_{0t}}, \sigma_{\nu_{0t}}\}$. This significantly reduces the dimensionality of the parameter space, and is a natural simplification. Let $F^*(\cdot|N)$ be the CDF of the equilibrium distribution of valuations among entering bidders, and $H_n^*(\cdot|N)$ be the equilibrium CDF of the random variable $\max\{Y_{1:n-1}, V_0\}$:

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

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

$$h_n^*(v|N) = f_0(v) \cdot F^*(v|N)^{n-1} + (n-1)F_0(v)F^*(v|N)^{n-2}f^*(v|N), \quad (\text{A.3})$$

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

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

where the last equality follows from integration by parts.

Appendix B. Obtaining an Expression for the Likelihood Function

We begin with equation (10) and derive expressions for $\Pr(p_j, \text{sale}_j | n_j, N_j; \theta_j)$, $\Pr(n_j | N_j, \theta_j)$, and $\Pr(\text{sale}_j | N_j; \theta_j)$. Since these are each conditional on the primitives implied by θ_j , the corresponding probabilities follow directly from our auction model.

We begin by characterizing $\Pr(n_j | N_j, \theta_j)$. To do this, we make recourse to the signal threshold characterizing equilibrium entry behavior when N_j potential acquirers compete for a target with characteristics θ_j as $s^*(N_j; \theta_j)$. Equation (9) then becomes

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

Specification of a joint distribution $F_{vs}(\cdot; \theta_j)$ determines the marginal and conditional distributions $F_s(\cdot; \theta_j)$, $F_v(\cdot; \theta_j)$, and $F_{v|s}(\cdot | \cdot; \theta_j)$, at which point computation of $s^*(N_j, \theta_j)$ becomes a straightforward numeric exercise.

By construction, potential acquirers drawing signals $S_{ij} \geq s^*(N_j, \theta_j)$ elect to enter in equilibrium. Signal draws are independent given target characteristics θ_j , so the number of

entrants n_t follows a binomial distribution based on entry probability

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

This in turn implies

$$\Pr(n_j | N_j, \theta_j) = B(n_j; N_j, q(N_j, \theta_j)), \quad (\text{B.3})$$

where $B(n; N, q)$ is the binomial PDF corresponding to probability parameter q .

We now derive an expression for $\Pr(\text{sale}_j \cap p_j | n_j, N_j; \theta_j)$. 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 probability $\Pr(\text{sale}_j \cap p_j | n_j, N_j; \theta_j)$

$$\begin{aligned} &= \Pr(\text{sale}_j \cap Y_{2:n_j} = p_j | n_j, N_j, \theta_j) + \Pr(\text{sale}_j \cap V_{0j} = p_j | n_j, N_j, \theta_j) \quad (\text{B.4}) \\ &= \Pr(Y_{1:n_j} \geq p_j \cap Y_{2:n_j} = p_j \cap V_{0j} \leq p_j | n_j, N_j, \theta_j) \\ &\quad + \Pr(Y_{1:n_j} \geq p_j \cap Y_{2:n_j} \leq p_j \cap V_{0j} = p_j | n_j, N_j, \theta_j) \\ &= \left[n_j(n_j-1)F^*(p_j; N_j, \theta_j)^{n_j-2} [1 - F^*(p_j; N_j, \theta_j)] f^*(p_j; N_j, \theta_j) \right] \cdot F_0(p; \theta_j) \\ &\quad + \left[n_j F^*(p_j; N_j, \theta_j)^{n_j-1} [1 - F^*(p_j; N_j, \theta_j)] \right] \cdot f_0(p_j; \theta_j); \end{aligned}$$

where to streamline notation we let

$$F^*(v; N_j, \theta_j) = F^*(v; s^*(N_j, \theta_j)) = F(v | S_i \geq s^*(N_j, \theta_j)) \quad (\text{B.5})$$

denote the equilibrium distribution of valuations among entrants at (N_j, θ_j) .

Finally, we derive an expression for $\Pr(\text{sale}_j|N_j; \theta_j)$. 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:

$$\begin{aligned} \Pr(\text{sale}_t|N_t; \theta_t) &= \Pr(V_{0t} \leq Y_{1:N_t}|N_t, \theta_t) = 1 - \Pr(Y_{1:N_t} \leq V_{0t}|N_t, \theta_t) & (\text{B.6}) \\ &= 1 - \int_0^\infty [F_s^*(N_t; \theta_t) + F_v(v_0; \theta_t) - F_{vs}(v_0, s^*(N_t, \theta_t); \theta_t)]^{N_t} f_0(v_0, \theta_t) dv_0, \end{aligned}$$

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 .

Thus given a specification for the auction-level fundamentals $\{c(\theta), F_{vs}(\cdot; \theta), F_0(\cdot; \theta)\}$, computing the likelihood components $\Pr(p_j, \text{sale}_j, n_j|N_j; \theta_j)$, $\Pr(n_j|N_j, \theta_j)$, and $\Pr(\text{sale}_j|N_j; \theta_j)$ conditional on θ_j becomes a simple numeric exercise. Computation of the overall likelihood function (11) then involves integration of these objects over realizations of θ_j , which we do via simulation, which we describe next.

Appendix C. Details of the Estimating Procedure

Estimation based on the likelihood function requires repeated evaluation of the integrals

$$\int \Pr(p_j, \text{sale}_j|n_j, N_j; \theta) \Pr(n_j|N_j, \theta) g(\theta|X_j) d\theta \quad (\text{C.1})$$

and

$$\int \Pr(\text{sale}_j|N_j; \theta) g(\theta|X_j) d\theta \quad (\text{C.2})$$

for each target j . In principle, the objects $\Pr(p_j, \text{sale}_j|n_j, N_j; \theta)$, $\Pr(n_j|N_j, \theta)$, and $\Pr(\text{sale}_j|N_j; \theta)$ are known up to θ , so such evaluation is straightforward in theory, given a form for the heterogeneity distribution $g(\cdot)$.

Our objective in estimation is then to estimate the deep fundamental parameters Γ governing the distribution of θ_j in the population, using the selection-corrected likelihood relationship derived above. Our baseline results employ the truncated Gaussian specification of Roberts and Sweeting (2013), under which the elements of θ_j are drawn independently from the following distributions:

$$\begin{aligned}\mu_j &\sim N(\gamma_\mu \cdot X_j, \sigma_\mu^2) \\ \sigma_{vj} &\sim TN(\gamma_\sigma \cdot X_j, \sigma_\sigma^2; \tau, \infty) \\ c_j &\sim TN(\gamma_c \cdot X_j, \sigma_c^2; 0, \infty) \\ \alpha_j &\sim TN(\gamma_\alpha \cdot X_j, \sigma_\alpha^2; 0, 1),\end{aligned}$$

where $TN(\bar{E}, \bar{V}; a, b)$ denotes the truncation of a Gaussian distribution with mean \bar{E} and variance \bar{V} on the interval $[a, b]$, and $\tau > 0$ is a regularization constant which ensures the variance σ_{vj}^2 is bounded away from zero. The vector of parameters to estimate is thus $\Gamma = \{\gamma_\mu, \gamma_\sigma, \gamma_c, \gamma_\alpha; \sigma_\mu^2, \sigma_\sigma^2, \sigma_c^2, \sigma_\alpha^2\}$.

We also explore estimation under several alternative specifications for $g(\cdot|X_j)$, such as using a Beta distribution for the information parameter α_j , and Gamma or log normal distributions for the always positive parameters σ_{vj} and c_j . Results obtained under these alternatives are qualitatively similar to our baseline specification.

Direct evaluation of the likelihood function (11) is computationally prohibitive in practice since (C.1) and (C.2) 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 Akerberg (2009), which uses the principle of importance sampling to transform the complicated problem of repeated evaluation of the full likelihood (11) into the much simpler problem of repeated evaluation of $g(\theta|X_j) = g(\theta|X_j, \Gamma)$. 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 (C.1).

By standard importance sampling arguments, we can rewrite this integral as follows:

$$\begin{aligned}
& \int \Pr(p_j, \text{sale}_j | n_j, N_j; \theta) \Pr(n_j | N_j, \theta) g(\theta | X_j, \Gamma) d\theta & (C.3) \\
= & \int \left[\Pr(p_j, \text{sale}_j | n_j, N_j; \theta) \Pr(n_j | N_j, \theta) \frac{g(\theta | X_j, \Gamma)}{\tilde{g}(\theta)} \right] \tilde{g}(\theta) d\theta \\
= & \tilde{E} \left[\Pr(p_j, \text{sale}_j | n_j, N_j; \theta) \Pr(n_j | N_j, \theta) \frac{g(\theta | X_j, \Gamma)}{\tilde{g}(\theta)} \right],
\end{aligned}$$

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 $\{\tilde{\theta}_r\}_{r=1}^R$ is a random sample drawn from $\tilde{g}(\cdot)$, it follows that for large enough R

$$\begin{aligned}
& \int \Pr(p_j, \text{sale}_j | n_j, N_j; \theta) \Pr(n_j | N_j, \theta) g(\theta | X_j, \Gamma) d\theta & (C.4) \\
\approx & \sum_{r=1}^R \Pr(p_j, \text{sale}_j | n_j, N_j; \theta_r) \Pr(n_j | N_j, \theta_r) \frac{g(\theta_r | X_j, \Gamma)}{\tilde{g}(\theta_r)}.
\end{aligned}$$

If a new sample $\{\tilde{\theta}_r\}_{r=1}^R$ is drawn each time the integral (C.1) 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 motivates Akerberg (2009)'s reinterpretation of importance sampling: rather than obtaining new draws each time (C.1) is evaluated, draw a single large sample $\{\tilde{\theta}_r\}_{r=1}^R$ from $\tilde{g}(\cdot)$ once at the beginning of the estimation algorithm, and

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 of these. Maximization of the overall likelihood function (11) is then (approximately) equivalent to maximization of the simulated likelihood function (12) with respect to Γ , where θ is calculated internal to the maximization problem using the expressions derived in Appendix A.A, where evaluation of the likelihood function at different values of Γ requires only recalculation of the 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, Akerberg (2009) simulation is ideal; it mitigates the computational infeasibility that otherwise would be entailed by accommodating sample selection unobserved heterogeneity.

References

- Akerberg, Dan, 2009, A new use of importance sampling to reduce computational burden in simulation estimation, *Quantitative Marketing and Economics* 7, 343-376.
- Aktas, Nihat, Eric de Bodt and Richard Roll, 2010, Negotiations under the threat of an auction, *Journal of Financial Economics* 98, 241-255.
- Betton, Sandra, and B. Espen Eckbo, 2000, Toeholds, Bid Jumps, and Expected Payoffs in Takeovers, *Review of Financial Studies*, 13, 841-882.
- Boone, Audra L. and J. Harold Mulherin, 2007, How are firms sold? *Journal of Finance*, 62, 847-875.
- Bruner, Robert F., 2004, *Applied Mergers and Acquisitions*. Hoboken, NJ: John Wiley & Sons, Inc.
- Bulow, Jeremy and Paul Klemperer, 2009, Why do sellers (usually) prefer auctions? *American Economic Review* 99, 1544-1575.
- Bulow, Jeremy and Paul Klemperer, 1996, Auctions versus negotiations, *American Economic Review* 86, 180-194.
- Cho, In-Koo and David M. Kreps, 1987, Signaling games and stable equilibria, *Quarterly Journal of Economics* 102, 179-221.
- Dimopoulos, Theodosios, and Stefano Sacchetto, 2011, Preemptive bidding, target resistance, and takeover premiums, Working Paper.
- Fishman, M., 1988, A Theory of Preemptive Takeover Bidding," *RAND Journal of Economics*, 19(1), 88-101.
- French, Kenneth R., and Robert E. McCormick, 1984, Sealed bids, sunk costs, and the process of competition, *Journal of Business* 57, 417-441.
- Matthew Gentry and Tong Li, 2014, Identification in auctions with selective entry, *Econometrica* 82, 315-344.
- Gorbenko, Alexander S., and Andrey Malenko, 2013, Strategic and financial bidders in takeover auctions, *Journal of Finance*.
- Hansen, Robert, 2001, Auctions of companies, *Economic Inquiry* 39, 30-43.
- Horner, J., and N. Sahuguet, 2007, Costly Signalling in Auctions, *Review of Economic Studies*, 74, 173-206.

- Kirman, Igor, 2008, *M&A and Private Equity Confidentiality Agreements Line By Line*, Thompson, Aspatore.
- Krasnokutskaya, Elena, 2012, Estimating auction models with multi-dimensional unobserved heterogeneity, *International Economic Review*, 53(3).
- Levin, Dan, and James L. Smith, 1994, Equilibrium in auctions with entry, *The American Economic Review* 84, 585–599.
- Milgrom, Paul R. and Robert J. Weber, 1982, A Theory of auctions and competitive bidding, *Econometrica* 50, 1089-1122.
- Roberts, James W., and Andrew Sweeting, 2013, When Should Sellers Use Auctions? *American Economic Review* 103,1830-61.
- Rosenbaum, Joshua and Joshua Pearl, 2009, *Investment Banking - Valuation, Leveraged Buyouts, and Mergers & Acquisitions*. Hoboken, New Jersey.
- Samuelson, W. F., 1985, Competitive Bidding with Entry Costs, *Economics Letters*, 17, 53–57.
- Shleifer, Andrei, and Robert W. Vishny, 2003, Stock market driven acquisitions, *Journal of Financial Economics* 70, 295-311.
- Subramanian, Guhan, 2008, Go-Shops vs. No-Shops in Private Equity Deals: Evidence and Implications. *Business Lawyer*
- Subramanian, Guhan. (2011). *Dealmaking: The New Strategy of Negotiauctions*. WW Norton & Company.
- Wasserstein, Bruce, 2000, *Big Deal: Mergers and Acquisitions in the Digital Age*. New York, NY: Warner Business.

Figure 1. Uncertainty and the Composition of the Entering Bidder Pool




Figure 2. Uncertainty and Overall Entry




Figure 3. Uncertainty and Expected Deal Premium




Figure 4. Comparing Auctions and Negotiations




Figure 5. Uncertainty, Entry Costs, and Expected Revenue

This figure reports mean expected revenue across takeover environments with different levels of pre-entry uncertainty and entry costs. The estimates are constructed using mean observable characteristics and median unobservable characteristics (Median Γ) and the resulting baseline fundamental parameter estimates. Red circles indicate situations where auctions revenue-dominate negotiations, blue circles indicate situations where negotiations revenue-dominate auctions, and hollow circles indicate situations where simulation error is larger than estimated revenue differences.




Figure 6. Pre Entry Uncertainty and Expected Revenue

This figure reports mean expected revenue for auctions, sequential negotiations, and one-shot negotiations with market check against pre-entry uncertainty. The graph is constructed by fixing (μ, σ, c) at their mean values and setting $V_0=1$ and $N=8$. The red line is mean expected revenue under negotiations, the blue line is mean expected revenue under auctions, and the green line is mean expected revenue under one-shot negotiations with a market check.




Table I. Sample by Year

This table reports the number of takeovers of publicly-traded targets with deal value greater than 1 million U.S. dollars, where the acquirer owns 100 percent of the target as a consequence of the deal, and financial data on the target is available from Standard and Poor's Compustat database. The sample covers deals that satisfy these criteria and are announced between January 1, 2000 and December 31, 2009. We also require that takeover proxy statements for the firms be available from the Securities and Exchange Commission. The number of takeovers is reported for the full sample, for auction sales, and for negotiated transactions.

Year	Full Sample	Auction	Negotiation
2000	163	72	91
2001	135	81	54
2002	75	47	28
2003	104	67	37
2004	110	62	48
2005	92	52	40
2006	97	58	39
2007	103	61	42
2008	50	34	16
2009	51	31	20

Table II. Summary of the Sale Process

This table summarizes the sale processes of 982 takeovers. The first row reports summary statistics for the full sample, while the second and third report data for auctions and negotiations, respectively. The variable Contact reports the average number of contacted potential bidders for each sale mechanism and the variable Confidential reports the average number of invited potential bidders that sign confidentiality agreements with the target. Premium reports the average price paid by the winning bidder, relative to the target's share price four weeks prior to announcement.

Sample	Number	Contact		Confidential		Prem
		Mean	Median	Mean	Median	
Full Sample	980	8.61	2.0	3.92	1.0	42.59
Auctions	565	14.2	5.0	6.1	3.0	41.68
Negotiations	415	1.0	1.0	1.0	1.0	43.84

Table III. Target Characteristics

This table reports mean target characteristics for firms sold via auction and negotiation. Data are drawn from Standard and Poor's Compustat database. Standard deviations are reported in parentheses beneath the estimates. Size is equal to total asset value in millions of US dollars. The Market to book ratio is the market value of assets divided by the book value of assets. Cash, leverage, and intangibles to assets are respectively total cash, long-term debt plus short-term debt, and intangible assets all scaled by total book assets.

Variable	Auction	Negotiation
Size	1,603 (7,305)	2,952 (16,669)
Market to book	3.01 (11.10)	3.39 (5.52)
Cash to assets	0.14 (0.17)	0.15 (0.18)
Leverage	0.27 (0.33)	0.29 (0.36)
Intangibles to assets	0.10 (0.16)	0.11 (0.16)
Number of takeovers	565	415

Table IV. Sale-Level Parameters

This table shows moments of the estimated fundamental parameter distributions. Each panel reports parameter estimates at the mean, median, 25th percentile, and 75th percentile of the estimated parameter distribution. Panel A reports moments of estimated parameters for a representative takeover with average observable characteristics (mean X) and median unobservable characteristics (Median Γ). Panel B reports quantiles of the unconditional distribution of auction-level parameters θ across all auctions in the sample, accounting for uncertainty in estimates of structural parameters Γ implied by the estimated posterior distribution.

Panel A: Quantiles at mean X_j median Γ_j

	μ_j	σ_{vj}	c_j	α_j
Mean	0.1852	0.1578	0.0133	0.6409
Median	0.1854	0.1444	0.0122	0.6856
25th	0.0384	0.0967	0.0033	0.4446
75th	0.3324	0.2954	0.0318	0.8695

Panel B: Posterior quantiles across median X_j

	μ_j	σ_{vj}	c_j	α_j
Mean	0.2097	0.1581	0.0243	0.6407
Median	0.2089	0.1467	0.0120	0.6891
25th	0.0981	0.0981	0.0030	0.4441
75th	0.3786	0.2057	0.0317	0.8720

Table V. Estimates of the Composition Effect

This table reports estimates of the composition effect for different values of pre-entry uncertainty and entry costs. For comparability with observed deal premia the estimates are expressed as a percent of the target's share price. The composition effect, introduced in Section V.B., is measured as the difference between the expected price that would maintain if the distribution of the entering and potential bidder pools were identical. The estimates are constructed using mean observable characteristics and median unobservable characteristics (Median Γ) and the resulting baseline fundamental parameter estimates.

	Pre-entry uncertainty (α)				
	0.15	0.35	0.64	0.85	0.99
$c = 0.005$	10.1%	7.1%	2.2%	0.1%	0.0%
$c = 0.015$	12.2%	9.8%	5.4%	1.8%	0.0%
$c = 0.030$	12.9%	11.0%	7.3%	3.9%	0.1%

Table VI. Unconditional Counterfactual Estimates

This table reports counterfactual estimates comparing auctions with one-shot negotiations after which is a market-check (described in Section 6.A) and a sequential negotiation procedure (also described in Section 6.A). The estimates are constructed using mean observable characteristics and median unobservable characteristics (Median Γ) and the resulting baseline fundamental parameter estimates. The table reports means, medians, and standard deviations of the distribution of expected revenue for a given target.

	Auction	Market check Negotiation	Sequential Negotiation
Mean Premium	39.9%	40.2%	41.3%
Median Premium	33.1%	33.3%	34.9%
Revenue Std. Dev.	10.1%	9.1%	6.9%
Revenue Skewness	26.7	25.6	20.9
