Indicative Bidding in Auctions with Costly Entry*

Daniel Quint, University of Wisconsin
Kenneth Hendricks, University of Wisconsin and NBER

September 3, 2015

Abstract

When selling a business by auction, investment banks frequently use indicative bids – non-binding preliminary bids – to select a limited number of bidders to participate. We show that if participation is costly, indicative bids can be informative: symmetric equilibrium exists in weakly-increasing strategies, but bidders "pool" over a finite number of bids, so the highest-value bidders are not always selected. We construct such equilibria for both first-price and English auctions. We also characterize equilibrium play when the number of potential bidders is large, and show that both revenue and bidder surplus are higher than when entry is unrestricted.

*The authors thank Serafin Grundl for excellent research assistantship, and Eric Budish, Peter Cramton, Navin Kartik, Paulo Somaini, Sheridan Titman, and attendees at numerous seminars and conferences for helpful comments.

1 Introduction

Over the last twenty years, the number of mergers and acquisitions (M&A) deals among U.S. companies has ranged from six and ten thousand per year, and the total value of these deals from \$400 billion to \$1.5 trillion per year.[1] Using SEC filings, Boone and Mulherin (2007, 2009) have constructed samples of hundreds of these deals. Their work indicates that in roughly half of the deals sellers negotiate exclusively with a single buyer, while in the other half sellers administer an auction.

However, an M&A auction is not like any that has been studied in the economics literature. The novel aspect is the process by which sellers screen and select the set of eligible buyers who can bid. It is best described using an example given in Boone and Mulherin (2007). In 1998, Instron, a worldwide leader in the material testing industry, hired the Beacon Group to explore options for a sale of the company. Beacon Group contacted 49 potential acquirers; 23 signed confidentiality agreements and in exchange received confidential information packets. Ten potential buyers submitted preliminatry indications of interest and acquired more information. Of these ten, the five who seemed willing to pay the most were invited to submit formal but nonbinding proposals specifying an estimated price and other details of the proposed deal. Two buyers submitted proposals: Kirtland Financial Partners, a private equity firm, submitted a bid of \$22 per share, and another buyer, identified in the SEC filings as Bidder A, submited a bid of \$22.50 per share. These two buyers were each invited to perform extensive due diligence, which included examining legal and accounting documents, touring facilities, and interviewing Instron's senior management. Kirtland and Bidder A then submitted bids of \$20.50 per share and \$20.00 per share, respectively. After further negotiations, Kirtland raised its bid to \$22.00 per share. Bidder A was informed that its bid was no longer competitive, but refused to raise its offer, and Kirtland won the auction. Since the bidders had multiple chances to raise their offers in response to those of their opponents, the auction somewhat resembled

[1] 2015 Mergerstat Review

an open-outcry English auction, but one in which bidders did not know the number or identities of their rivals.

The example illustrates several important features of the screening process. A large number of buyers are contacted. Based on their private valuations of the target, many buyers opt out, but several are sufficiently interested to proceed. However, before submitting bids, these buyers have to conduct due diligence, a process whose cost, depending on the size and complexity of the deal, can run into the millions of dollars. Thus, M&A auctions can be thought of as auctions with costly entry, with due diligence being the main entry cost. Often, even when it is buyers who directly incur entry costs, sellers bear them indirectly through their effect on equilibrium entry and bidding,[2] and therefore have an incentive to limit the number of buyers who can enter. But sellers also want to ensure that the "right" buyers are selected, namely, those likely to have the highest willingness to pay. Many sellers solve this problem by asking potential buyers how much they would be willing to pay, and permitting only the buyers who indicate the highest willingness to pay to enter. These "indicative bids" are costless, because they are never paid by the buyers; and they are non-binding, because they do not restrict in any way the real offers that a buyer may subsequently make in the auction. Despite this absence of commitment, the widespread use of indicative bids in M&A auctions suggests that sellers find them informative.

Our main goal in this paper is to analyze the effectiveness of indicative bids for screening and selecting bidders in auctions with costly entry. Clearly, if indicative bids are informative, the seller has every incentive to pay attention to them: revenue is highest when the buyers chosen are those with the highest willingness to pay. But the incentives for buyers to honestly report their willingness to pay are less clear. If entry costs were zero, then entry could be thought of as a free option, and we would expect indicative bidding to unravel to all buyers submitting the highest possible bid. (Morever, if entry were costless, the seller would have no reason to restrict it.) However, if entry costs are positive, then buyers' incentives are partly aligned with those of the

[2]See, for example, Levin and Smith (1994) and Cremer, Spiegel and Zheng (2009).

seller: the seller wants to restrict the number of buyers who incur the entry cost, and buyers want to avoid being selected and paying this cost if they are unlikely to win. Thus, low-value buyers will be willing to separate themselves from high-value buyers, by bidding less than the highest possible bid.

We use a version of the two-stage, value updating model developed by Ye (2007) to study the impact of indicative bidding on auction outcomes. In the first stage, N buyers draw their private values and are asked to simultaneously submit indicative bids or opt out. The seller commits to selecting the n potential buyers, typically two or three, who send the highest indicative bids, with ties broken randomly. If fewer than n buyers submit indicative bids, then all of them advance; if all buyers opt out, then no sale occurs. In the second stage, the entrants pay the costs of due diligence, possibly update their valuations, and simultaneously submit binding bids in the auction. Ye shows that a fully separating equilibrium fails to exist in this model. We are interested in characterizing the symmetric equilibria that do exist.

Our central result is that indicative bids can yield a partial sorting of buyers based on their private valuations of the asset (i.e., their types) in a broad class of environments. Fixing a countable bid space, we show that a symmetric equilibrium is a partition of the space of buyers' types. There is a finite upper bound on the number of indicative bids that are used, and only the lowest M bids are used with positive probability. Buyers with types in the same element of the partition submit the same indicative bid, and buyers in higher elements submit higher bids. Thus, the equilibrium leads to a partial sorting of the buyers, which helps the seller select high-value buyers with greater likelihood. We show how to construct the equilibrium for both first- and second-price auctions, provide general existence and uniqueness results for the latter, and conduct comparative static exercises.

How well does the indicative bidding mechanism perform? A natural benchmark is an auction in which entry is unrestricted and the seller sets optimal reserve prices. Buyers decide on the basis of their private information whether or not to enter the auction, pay the entry cost, possibly update their valuations, and submit binding bids. Our main result is that indicative bidding yields both greater revenue and greater buyer surplus than the unrestricted auction when the number of potential buyers is large. Through numerical examples, we find that this result appears to hold even when the number of buyers is small. Thus, the standard tradeoff in optimal auctions between efficiency and revenue is not present when entry is costly and sellers use indicative bids to screen and restrict entry.

This paper connects and contributes to two literatures: auctions with costly entry and cheap talk games. Much of the theoretical literature on auctions with costly entry deals with how the seller should restrict entry when buyers have no private information prior to entry. In this setting, the seller does not have to worry about selecting buyers, which is the central issue of our paper. Cremer, Spiegel, and Zheng (2009) characterize the optimal mechanism, and show that the seller can use entry fees and subsidies to extract all buyers' surplus. In the absence of such payments, Bulow and Klemperer (2009) demonstrate that auctions with unrestricted entry are less efficient than sequential mechanisms, but typically generate more revenue. Earlier papers include McAfee and McMillan (1987), Levin and Smith (1994), Burguet and Sakovics (1996), Menezes and Monteiro (2000), and Compte and Jehiel (2007).

Samuelson (1985) was the first to consider an environment in which buyers know their private values prior to the entry decision. In his model, bidders face bid preparation costs rather than information acquisition costs. Ye (2007) considers the more general case in which entry involves value updating, and shows that an entry rights auction can be used to induce efficient entry. Fullerton and McAfee (1999) had introduced this idea in the context of research tournaments (similar to all-pay auctions), and showed that a properly designed auction for entry rights can efficiently select the most qualified participants for the tournament. Lu and Ye (2014) characterize the optimal two-stage mechanism, which can be implemented using an entry rights auction. Our contribution to this literature is to use "cheap talk" rather than binding bids to resolve the coordination problem faced by buyers. In a later section, we discuss why entry rights auctions may be implausible to implement in M&A settings.

Our indicative bidding equilibria are similar to the "cheap talk" equilibria of Crawford and Sobel (1982): indicative bids are monotonic in buyers' initial information, but only a finite number of different bids are used in equilibrium, and different types of buyers "pool" on the same bid. In their seminal paper, Crawford and Sobel show that cheap talk can improve the ex ante payoffs of both parties when the sender has information relevant to the receiver's decision problem. Farrell and Gibbons (1989) and Matthews and Postlewaite (1989) similarly show that cheap talk can be informative prior to bilateral bargaining, and can therefore expand the set of equilibrium payoffs. Aside from considering cheap talk in a new setting, our contribution to this literature is to introduce a natural kind of commitment into a cheap talk setting, which sharpens the predictions of the model. As Farrell and Gibbons (1989) observed, in standard cheap-talk games, the receiver cannot commit to choosing a particular outcome as a function of the messages he receives. The messages therefore derive meaning only from the receiver's interpretation of them, and the receiver must act optimally given that interpretation. In our setting, we assume the seller commits both to the rules of the auction (which is standard) *and* to how he will select entrants based on the indicative bids received. In particular, we assume the seller commits to selecting those who send the highest indicative bids. This commitment to a monotone selection rule eliminates much of the multiplicity of equilibria that arises in cheap talk games. In particular, it rules out a "babbling" equilibrium, and any equilibrium where adverse off-equilibrium-path beliefs are used to deter unused messages.[3]

The paper proceeds as follows. Section 2 presents the basic model without value updating. In Section 3, we use a more restricted version of this model to show how to construct symmetric equilibria in both first- and second-price auctions. Section 4 establishes general existence results for the second-price auction and provides conditions for uniqueness. In Section 5, we evaluate the performance of the indicative bidding mechanism against the alternative of un-

[3]Navin Kartik and Joel Sobel (private communication) have similarly proved that if one imposes monotonicity on both the sender's and receiver's strategies in a standard cheap talk setting and then applies iterated weak dominance, this uniquely selects the "most informative" equilibrium.

restricted entry, and establish general results for the case where N is large. In Section 6 we extend our results to a more general information structure. Section 7 discusses alternative mechanisms for restricting and selecting entrants. Section 8 concludes.

2 Model

In this section, we describe the environment and present our model of an indicative bidding mechanism. For easier exposition, our base model assumes that bidders know their valuation for the asset perfectly prior to due diligence; thus, due diligence consists of verifying the information they already have, rather than acquiring new information. As we discuss in Section 6, this is also isomorphic to a model where bidders do receive new information during due diligence, but about a component of value common to all bidders. In Section 6, we extend the results to the case where bidders performing due diligence receive new idiosyncratic information about their private valuation.

There are $N \geq 3$ potential bidders, indexed by i. Each bidder i has a private value S_i for the asset, where S_i is a real-valued random variable. Bidders' values are drawn independently from a common distribution F, with support $[\underline{S}, \overline{S}]$ and density f. We assume that f is bounded below and has a finite derivative on its support. Each bidder knows her own value but not those of the other bidders.

The indicative bidding mechanism consists of a cheap talk stage and an auction stage. In the cheap talk stage, the seller asks potential buyers if they are interested in bidding for the asset and, if so, how much they are willing to pay. These bids are not binding. An indicative bid can be a real number, although in reality the seller often asks bidders to report a range in which they believe their willingness to pay is likely to fall. We formalize this stage of the game by assuming that buyers simultaneously send messages to the seller from a set of available messages $\{0, 1, ..., \overline{M}\}$, where 0 is "opt out", or decline to participate, and $\overline{M}$ is allowed to be either finite or infinite. Note that $\overline{M}$ is a parameter of choice for the seller. The substantive restriction here is that the set of messages is bounded below (i.e., there is a lowest "opt-in" message), fully ordered, and countable. Given this restriction, there is no loss of generality in assuming that the message space consists of non-negative integers.[4]

The seller selects bidders for the auction stage based on the messages received. We assume the seller commits to the maximal number of bidders n, and if more than n bidders opt in, commits to selecting the bidders who sent the highest messages, breaking ties randomly. If all bidders opt out, then the game ends with no sale. In addtion, the seller commits not to disclose the messages that the bidders send him, nor his response to those messages. Therefore, each bidder knows only whether or not she has advanced to the auction, and does not know which (if any) of the other bidders have advanced.

Participation in the auction stage means conducting due diligence to verify the value of the asset and then submitting a binding bid. Due diligence involves sending teams of lawyers and accountants to a data room that contains the contracts and financials of the business. Their main goal is to ensure there are no hidden liabilities or "skeletons in the closet." We assume the cost of due diligence is the same for all buyers, and we denote it by c and refer to it as the entry cost. To ensure that the game is nontrivial, we assume $c < \overline{S}$.

We assume that bidders who opt in are required to participate if selected. Within our model, if it were allowed, a bidder would sometimes want to opt in but then decline to participate if she faced the maximal level of competition. An easy way to prevent this is for the seller to commit not to disclose how many bidders opted in. In reality, sellers often adopt such a nondisclosure policy for a different reason – to avoid weakening their bargaining position when only one bidder is seriously interested in buying the asset. Thus, in practice, both during and after due diligence, bidders are often kept in the dark as to the number of competitors they face.[5]

Under the above assumptions, the primitives (N, F, c) fully describe the environment facing the seller. The indicative bidding mechanism faced by the

[4]As we discuss later, symmetric equilibrium fails to exist when the set of messages is continuous.

[5]Subramanian (2010) offers some funny stories on how this is accomplished.

bidders is defined by $(N, F, c, \overline{M}, n)$ and the rules of the auction, all of which are common knowledge. A pure strategy for bidder i consists of a message function

$$\tau_i : [\underline{S}, \overline{S}] \longrightarrow \{0, 1, .., \overline{M}\}$$

that maps the set of types to the set of messages, and a bid function

$$\beta_i : [\underline{S}, \overline{S}] \longrightarrow \mathbb{R}_+$$

that maps the set of types to the real line (possible auction bids). Mixed strategies are mappings from $[\underline{S}, \overline{S}]$ to probability distributions over these same spaces. The **support** of a first-stage strategy τ is defined to be the set of messages played with positive probability by a positive measure of types.

Note that when a bidder learns she has advanced to the auction, she will update her beliefs about how many other bidders advanced and the distribution of their types; in equilibrium, her auction bid $\beta_i(s_i)$ must be optimal given these new (correct) beliefs. However, since a bidder only gets to bid when selected for the auction, and since selected bidders do not learn any new information other than that they have been selected, we can alternatively think of the bidder as choosing both a first-round message m and a bid b simultaneously, rather than optimally choosing b after being selected. This, it turns out, will simplify some of the analysis.

Our objective is to characterize symmetric equilibria. Consequently, we need only define the expected payoff to a bidder when her opponents all play a common strategy. Let $v_{\tau,\beta}(m, b; s_i)$ denote the expected payoff to bidder i given her own type is s_i, she sends message m and bids b if selected, and her opponents all play the strategy (τ, β). Then a symmetric Bayes Nash equilibrium is a strategy (τ, β) such that

$$v_{\tau,\beta}(\tau(s_i), \beta(s_i); s_i) \quad \geq \quad v_{\tau,\beta}(m', b'; s_i)$$

for all $s_i \in [\underline{S}, \overline{S}]$, $m' \in \{0, 1, .., \overline{M}\}$, and $b' \in \Re_+$.

The indicative bidding game can be thought of as a cheap talk game with commitment. The messages of the bidders (i.e., senders) influence which action the seller (i.e., receiver) takes but, given that action, they do not affect the payoffs of the players directly. In the standard cheap talk game, the receiver chooses an action that, given his beliefs, is his best response to the messages sent. By contrast, in the indicative bidding game (as is standard in auctions), the seller commits to the mechanism. In particular, he commits to selecting those bidders who send the highest messages, and to ignoring messages outside a particular set.[6]

3 Equilibrium Construction (Restricted Model)

In this section, we consider a more restricted version of the model, and use it to illustrate the construction of symmetric equilibria for indicative bidding mechanisms with both first- and second-price auctions. For second-price auctions, we will show later that the construction generalizes, guaranteeing symmetric equilibrium existence (and, under an additional condition, uniqueness) for the general model above.

Throughout this section, we assume that the set of possible types is $[\underline{S}, \overline{S}] = [0, 1]$, and the distribution of types F is the uniform distribution on that interval; and we assume that $n = 2$, that is, only two bidders will advance to the auction. We restrict attention to first-round strategies τ which are pure strategies, non-decreasing, and have support $\{0, 1, \ldots, M\}$ for some finite $M \leq \overline{M}$ (even if $\overline{M}$ itself is infinite).[7] This means that higher types send higher messages, and that no messages are "skipped": if message k is used, then all lower messages must also be used. Any strategy with these properties induces

[6]Ex interim, once the messages are received, it might be in the seller's interest to deviate from this mechanism – to allow more bidders to advance, or to consider messages outside of the permitted set if they were used – but we assume he is able to credibly commit to the announced mechanism.

[7]These restrictions turn out to be nearly without loss of generality, as the next section will show: any symmetric equilibrium requires strategies to be non-decreasing and have finite support with no skipped messages. And while mixed-strategy equilibria exist, only a measure-zero set of types mix, so any symmetric equilibrium turns out to be payoff-equivalent to a pure-strategy one.

a partition of the type space: that is, for any such strategy, there is a series of thresholds $\alpha = (\alpha_0, \alpha_1, \ldots, \alpha_M)$, with $0 < \alpha_0 < \alpha_1 < \cdots < \alpha_M = 1$ and

$$\tau(s) \quad = \quad \begin{cases} 0 & \textit{if} \quad s \in [0, \alpha_0) \\ & \vdots \\ m & \textit{if} \quad s \in (\alpha_{m-1}, \alpha_m) \\ & \vdots \\ M & \textit{if} \quad s \in (\alpha_{M-1}, 1) \end{cases} \tag{1}$$

Given (1), we will say that the thresholds α **describe** τ. While τ uniquely determines α, α defines τ only up to its value at the threshold types: for $m < M$, $\tau(\alpha_m)$ could be either m or $m+1$, and if $M < \overline{M}$, $\tau(1)$ could be any of $\{M, M+1, \ldots, \overline{M}\}$. However, since these types are zero-probability events, α contains all the information about τ necessary to calculate expected payoffs.

3.1 Second-Price Auctions

When the seller uses a sealed-bid second-price auction or an English auction, then whatever message she sent in the first round, each bidder who advances has a dominant strategy to bid her type in the auction itself. That is, in second-price auctions,

$$\beta(s) \quad = \quad s.$$

Bidder i's payoff from advancing to the auction is therefore $s_i - c$ if she advances alone, and $\max\{0, s_i - s_j\} - c$ if she advances against an opponent with type s_j. These payoffs have two important implications. One is that other than the top interval $(\alpha_{M-1}, 1)$, the intervals of types sending each opt-in message cannot be too narrow.[8] This effect of entry costs on the bidder's incentive

[8] As we will show below, for $m < M$, a bidder with type α_m must be indifferent in equilibrium between sending messages m and $m+1$. Sending message $m+1$ increases her risk of advancing against an opponent stronger than herself, which leads to a payoff of $-c$. This added cost must be offset by a higher chance of advancing against another opponent who sent message m, and therefore has type $s_j \in (\alpha_{m-1}, \alpha_m)$. If $\alpha_m - \alpha_{m-1} < c$, this event would not give a positive payoff either, and a bidder with type α_m would strictly prefer to send the lower message.

to report her type truthfully is analogous to the bias between the sender and receiver preferences in Crawford and Sobel (1982).[9]

A second, closely related, consequence is that there is a finite upper bound on M, the number of opt-in messages used in equilibrium, even if $\overline{M}$ is large or infinite. It may seem counterintuitive that a bidder with the highest type would not want to separate herself by sending a higher message (if one were available) and advancing for certain. But the increase in her probability of advancing would arise solely from breaking ties against opponents who are sending the highest message M, and are therefore in the same interval as the highest type. As more messages get used, this interval gets sufficiently small that the highest type's payoff against a randomly selected opponent in this interval is negative. When this is the case, she will not want to separate herself by sending a higher message.

To characterize the expected payoff to a bidder from sending a particular message, we must first specify the expected payoff from advancing to the auction. Suppose that bidder i has type s_i, and that the highest first-round message sent by any of her opponents is m. Then if she advances to the auction, her expected payoff will be

$$V(s_i, [a, b]) \quad \equiv \quad \int_a^b [\max\{0, s_i - s\} - c] d\left(\frac{s-a}{b-a}\right) \tag{2}$$

where (a, b) is the interval of types who send message m. (Even if multiple opponents sent message m, the one chosen to advance is chosen randomly, so the interim distribution of her opponent's type is still uniform.) Note that bidder i's payoff, *conditional on advancing*, does not depend on the message she herself sent, only on the message sent by her opponent.

Earlier, we defined $v_{\tau,\beta}(m, b; s_i)$ as a function of bidder i's type, strategy (m, b), and the strategy (τ, β) played by her opponents. For second-price

[9]We thank Navin Kartik for pointing out this link between the two models. This also shows why (as shown in Ye (2007)) even if the message space were continuous, a fully revealing equilibrium could not exist: if a bidder's opponents were all signaling their types truthfully, she would want to underreport her own type by at least c, to avoid advancing in scenarios where she could not recoup her entry costs.

auctions, we assume that $\beta = I$ (the identity function) and $b = s_i$; noting also that bidder i's payoff depends on τ only through α, we define $v_\alpha(m, s_i) \equiv v_{\tau,I}(m, s_i; s_i)$. For any opt-in message $m > 0$, we can write v_α as

$$\begin{aligned} v_\alpha(m, s_i) \quad &= \quad \alpha_0^{N-1}(s_i - c) \\ &+ \sum_{m'=1}^{m-1} \sum_{j=1}^{N-1} \binom{N-1}{j} (\alpha_{m'} - \alpha_{m'-1})^j (\alpha_{m'-1})^{N-1-j} V(s_i, [\alpha_{m'-1}, \alpha_{m'}]) \\ &+ \sum_{j=1}^{N-1} \binom{N-1}{j} (\alpha_m - \alpha_{m-1})^j {\alpha_{m-1}}^{N-1-j} \left(\frac{2}{j+1}\right) V(s_i, [\alpha_{m-1}, \alpha_m]) \\ &+ \sum_{j=0}^{N-2} (N-1) \binom{N-2}{j} (1 - \alpha_m)(\alpha_m - \alpha_{m-1})^j (\alpha_{m-1})^{N-j-2} \left(\frac{1}{j+1}\right) V(s_i, [\alpha_m, 1]) \end{aligned} \tag{3}$$

Equation (3) groups the different profiles of opponent messages under which bidder i advances into four terms:

- The first term is the payoff to bidder i when all other bidders opt out. She advances alone and pays 0 for the asset.

- The second term covers events in which at least one opponent does not opt out, and m', the highest message of bidder i's opponents, is less than m. If j opponents send that message, then one of them is drawn at random to advance. Bidder i advances for sure and earns an expected payoff of $V(s_i, [\alpha_{m'-1}, \alpha_{m'}])$.

- The third term covers events in which j other bidders send message m and the remaining bidders send messages less than m. In this case, the probability that bidder i advances is $2/(j+1)$. If she advances, she bids against a type drawn randomly from the interval $[\alpha_{m-1}, \alpha_m]$, and obtains an expected payoff of $V(s_i, [\alpha_{m-1}, \alpha_m])$.

- The final term covers scenarios in which one of the bidders sends a message higher than m, j opponents send the message m, and the rest send messages less than m. In this case, bidder i advances with probability $1/(j+1)$ against a type drawn from $[\alpha_m, 1]$, and earns $V(s_i, [\alpha_m, 1])$.

When two or more opponents send messages higher than m, bidder i does not advance, so there is no term corresponding to payoffs in that case. When $m = 1$, the second term in equation (3) vanishes, and when $m = M$, the last term vanishes. Finally, the payoff to opting out, $v_\alpha(0, s_i)$, is 0.

In a symmetric equilibrium, bidder i's best reply to τ is τ. This requires each threshold type α_m to be indifferent between sending message m or message $m + 1$.[10] Thus, α must satisfy a system of M equations

$$v_\alpha(m+1, \alpha_m) - v_\alpha(m, \alpha_m) \quad = \quad 0 \quad \text{for} \quad m \in \{0, .., M-1\}. \qquad (4)$$

Next, we show how we find solutions to these M indifference conditions; after that, we will discuss one more condition that must be satisfied in a symmetric equilibrium.

Finding thresholds

For intuition, consider first the case $M = 2$, so that we are looking for just two thresholds, α_0 and α_1, which satisfy $v_\alpha(1, \alpha_0) = 0$ and $v_\alpha(2, \alpha_1) = v_\alpha(1, \alpha_1)$.

We begin with the observation that for a given value of α_1, the difference $v_\alpha(2, \alpha_1) - v_\alpha(1, \alpha_1)$ is strictly single-crossing from above in α_0. This means that given a value of α_1, there is a unique value of α_0 satisfying $v_\alpha(2, \alpha_1) - v_\alpha(1, \alpha_1) = 0$. Call this value $\alpha_0(\alpha_1)$.

Now, if we pick a value of α_1 close to 1, then $\alpha_0(\alpha_1)$ will be fairly high, and $v_\alpha(1, \alpha_0)$ will turn out to be strictly positive. By decreasing α_1, however, we force $\alpha_0(\alpha_1)$ downwards, eventually getting close enough to 0 that $v_\alpha(1, \alpha_0)$ must be negative. By continuity, then, there is a value of α_1 in between such that $\alpha = (\alpha_0(\alpha_1), \alpha_1)$ satisfies $v_\alpha(1, \alpha_0) = 0$. Since $v_\alpha(2, \alpha_1) - v_\alpha(1, \alpha_1) = 0$ by construction (this is how $\alpha_0(\alpha_1)$ was defined), these two thresholds satisfy both indifference conditions.

For general M, the construction is similar. For any $m \in \{1, 2, \ldots, M-1\}$, the difference $v_\alpha(m+1, \alpha_m) - v_\alpha(m, \alpha_m)$ depends only on the three thresholds

[10] For ϵ small, $\tau(\alpha_m - \epsilon) = m \longrightarrow v_\alpha(m+1, \alpha_m - \epsilon) \leq v_\alpha(m, \alpha_m - \epsilon)$, and $\tau(\alpha_m + \epsilon) = m + 1 \longrightarrow v_\alpha(m+1, \alpha_m + \epsilon) \geq v_\alpha(m, \alpha_m + \epsilon)$; (3) implies $v_\alpha(m+1, s_i)$ and $v_\alpha(m, s_i)$ are continuous in s_i, implying equality at $s_i = \alpha_m$.

$\{\alpha_{m-1}, \alpha_m, \alpha_{m+1}\}$, as the other thresholds all drop out;[11] and is strictly single-crossing in α_{m-1}. This means that a candidate value of α_{M-1} (along with the definition $\alpha_M = 1$) determines a unique candidate value of α_{M-2}; the values of α_{M-1} and α_{M-2} uniquely determine α_{M-3}; α_{M-2} and α_{M-3} determine α_{M-4}, and so on down to α_0. Once α_0 is reached, we can check whether $v_\alpha(1, \alpha_0)$ is positive or negative, and adjust our starting value of α_{M-1} accordingly. For a given set of primitives, as long as M is "not too big" (in a sense we define formally in the next section), we can always find a set of thresholds $\alpha_0 < \alpha_1 < \cdots < \alpha_{M-1}$ satisfying the M indifference conditions (4). We will define M^*, roughly, as the maximal value of M for which such thresholds exist.

One more condition to check

As we show in the next section, if $M = \overline{M}$, then any solution to the system of indiference equations (4) is an equilibrium. Thus, if $\overline{M} \leq M^*$, the construction above (with $M = \overline{M}$) gives a symmetric equilibrium for the indicative bidding game with $\overline{M}$ opt-in messages. However, if $M < \overline{M}$, there is one additional condition that must be satisfied in equilibrium: a bidder with the highest possible type must not benefit from deviating to an unused message, i.e., $v_\alpha(m', 1) \leq v_\alpha(M, 1)$ for any $m' > M$.[12] By sending such a message, say message $M+1$, the bidder would advance for sure; by sending message M, she would still be selected for sure unless two or more rivals also send message M, in which case she would be selected with probability less than 1. Thus, the only scenarios where message $M+1$ gives a higher probability of advancing than message M are scenarios where, upon advancing, the bidder faces an opponent who sent message M, i.e., an opponent with type above α_{M-1}. Thus, ruling

[11] Thresholds below α_{m-1} matter only in determining bidder i's opponent when all other bidders send messages below m; in that event, bidder i would advance for sure whether she sent message m or $m+1$, and the price she pays will be the same in either case.

[12] In a pure cheap-talk setting without commitment, such deviations could be ruled out via off-equilibrium-path beliefs that led to such deviators never being selected. Since we assume the seller commits to a monotone selection rule, such deviations need to be ruled out separately – which eliminates the multiplicity typically found in cheap talk models.

out such deviations is equivalent to requiring

$$0 \quad \geq \quad V(1, [\alpha_{M-1}, 1])$$

The right-hand side is strictly decreasing in α_{M-1}, so the latter condition is equivalent to $\alpha_{M-1} \geq \overline{\alpha}$, where $\overline{\alpha}$ is the solution to $V(1, [\overline{\alpha}, 1]) = 0$. This is analogous to the No Incentive To Separate (NITS) condition introduced by Chen, Kartin, and Sobel (2008).[13] In their model, NITS is an equilibrium refinement, and selects a unique equilibrium – the one with the maximal number of messages. In our setting, NITS is not a refinement, but a necessary condition that must hold in any equilibrium with $M < \overline{M}$.

Above, we defined M^* as the largest value of M for which we are able to satisfy all the indifference conditions (4). As it happens, the condition $\alpha_{M-1} \geq \overline{\alpha}$ holds at the thresholds satisfying the indifference conditions when $M = M^*$, but not when $M < M^*$. This means that the "candidate equilibrium" thresholds found above will indeed be an equilibrium if either $M = \overline{M}$ (because NITS is unnecessary) or $M = M^*$ (because NITS is satisfied), but not otherwise.

Thus, for F uniform, as well as for many other distributions, this construction leads to a symmetric equilibrium which is unique (up to the messages sent by the threshold types). For fully general F, as we will show below, existence is guaranteed but uniqueness is not. Even then, however, there is a sense in which the potential multiplicity is a very tractable problem. Our procedure reduces the search for symmetric equilibria to a finite number of one-dimenisonal searches; finding them all via exhaustive search is therefore computationally quite feasible.

We use a numerical example to illustrate the properties of the symmetric equilibrium in second-price auctions. We let F be the uniform distribution on $[0, 100]$ instead of $[0, 1]$, so that the numbers are easier to read; and let $c = 5$

[13]In their setting, NITS is the condition that the *lowest*-type sender would not choose to reveal his type truthfully if he could, while in our model, the condition is on the highest type. As noted above, the direction in which the sender/bidder would like to misrepresent his type goes in opposite directions in the two models.

and $N = 5$. For each value of $\overline{M}$, there is a unique symmetric equilibrium. Table 1 reports the equilibrium thresholds for various values of $\overline{M}$.

Table 1: $N = 5$, $F = U[0, 100]$, $c = 5$, $n = 2$, various $\overline{M}$

Opt-in messages available ($\overline{M}$):	1	2	3+
α_3	–	–	100.00
α_2	–	100.00	98.12
α_1	100.00	83.79	83.64
α_0	51.50	49.45	49.42
Revenue	53.67	57.21	57.26
Bidder Surplus	16.96	15.44	15.42
Total Surplus	70.63	72.65	72.68

When $\overline{M} \geq 3$, the equilibrium uses only messages $\{0, 1, 2, 3\}$; if we try to construct an equilibrium with more messages, the construction fails because the lowest threshold would have to be negative. Therefore, the maximum number of opt-in messages that will be used is 3. When $\overline{M} < 3$, all the available messages are used in equilibrium. In these cases, the highest type would like to separate but cannot do so because of the constraint that the seller imposes on the size of the message space.

Note also that when $\overline{M} \geq 3$, the intervals are narrower at higher messages. In this sense, there is finer sorting at the top of the type space than at the bottom. This property turns out to hold fairly generally when $\overline{M} \geq M^*$ – when the number of available messages is not a binding constraint. (As we note later, it follows from the same property that leads to uniqueness of the symmetric equilibrium.)

As $\overline{M}$ increases, the opt-in threshold α_0 decreases. As a result, each bidder's probability of opting in increases with $\overline{M}$, which we refer to as the *participation effect*. In addition, as $\overline{M}$ increases and more messages are used, the bidders sort more effectively, and the selected bidders are more likely to be those with the highest types. We refer to this effect as the *selection effect*. Both these effects favor higher revenue. Higher participation means the seller is more likely to sell the asset at a positive price (since this requires at least two bidders to opt in), and greater selection implies that the seller is likely to sell for a higher price. Bidder surplus, however, goes the opposite direction. Bidders benefit heavily from being the only one to opt in, which is more likely when $\overline{M}$ is lower; and they benefit from less effective sorting, since it increases the chance they do not face the toughest possible competition. Still, the increase in revenue appears to dominate the decrease in bidder surplus: in every example we've solved, total surplus is increasing in $\overline{M}$.

3.2 First-Price Auctions

As noted above, when a bidder advances to the auction stage, she updates her beliefs about the types of her opponent, conditioning on the fact that she advanced. This updating depends on the message she sent in the first stage: a bidder who advanced after sending a high message will have different beliefs than a bidder who advanced after sending a low message. In the case of a second-price auction, these beliefs were irrelevant, since bidding was in dominant strategies. In a first-price auction, however, these beliefs matter, which complicates the analysis. Nevertheless, similar equilibria still exist, and we will now show how to construct them.

We again restrict attention to first-stage strategies τ which are weakly increasing in a bidder's type and have finite support $\{0, 1, .., M\}$, and which are therefore described by a finite set of thresholds α; and we restrict attention to bid functions β which are strictly increasing and continuous in a bidder's type.[14] We will suppose a symmetric equilibrium exists in this type of strategy, and explore its implications. These conditions turn out to be sufficient as well as necessary for equilibrium, and we will then show how to construct strategies that satisfy them.

We begin by expressing a bidder's expected payoff, $v_{\tau,\beta}(m, b; s_i)$, in terms of its component parts. Letting $\sigma \equiv (\tau, \beta)$ be shorthand for the other bidders'

[14] Again, this is without loss: strategies violating these conditions cannot be played in symmetric equilibrium. Bidders who opt out do not bid, so we normalize $\beta(s) \equiv 0$ for $s < \alpha_0$. As we discuss below, $\beta(\alpha_0)$ will be 0 in any symmetric equilibrium, so this will not violate continuity.

strategy, we can write

$$v_\sigma(m,b;s_i) \quad = \quad (s_i - b)Q_\sigma(m,b) - L_\alpha(m)c$$

where $L_\alpha(m)$ is the probability a bidder advances to the auction if she sends message m (which depends on σ only through α) and $Q_\sigma(m,b)$ the probability she advances to the auction and then wins it given message m and bid b.

Next, we show how L_α and Q_σ are determined by the opponent bidders' strategy $\sigma = (\tau, \beta)$. We can write the probability of advancing, L_α, as

$$\begin{aligned} L_\alpha(m) \quad = \quad & (\alpha_{m-1})^{N-1} + \sum_{j=1}^{N-1} \tbinom{N-1}{j} (\alpha_m - \alpha_{m-1})^j (\alpha_{m-1})^{N-1-j} \left(\tfrac{2}{j+1}\right) \\ & + \sum_{j=0}^{N-2} (N-1) \tbinom{N-2}{j} (1-\alpha_m)(\alpha_m - \alpha_{m-1})^j (\alpha_{m-1})^{N-j-2} \left(\tfrac{1}{j+1}\right). \end{aligned}$$

The first term is the probability that all of a bidder's opponents submit messages below m, in which case she advances for sure. The second term covers the events in which j opponents send message m and the rest send lower messages, in which case the bidder advances with probability $2/(j+1)$. The last term is the sum of the probabilities of the events in which one opponent submits a higher message, j opponents send message m, and the rest send lower messages, in which case she advances with probability $1/(j+1)$. Note that the probability of advancing depends only on α_{m-1} and α_m, the thresholds associated with the message m, and not on the other thresholds. This property will be important when we construct the equilibrium.

Given her opponents' strategy σ, the probability $Q_\sigma(m,b)$ that a bidder advances and wins is a function of both her message and her bid. We will suppose that she plays an "on-the-equilibrium-path" combination of message and bid, so that $b \in [\beta(\alpha_{m-1}), \beta(\alpha_m)]$ (or $m = \tau(\beta^{-1}(b))$). In that case, when she advances against an opponent with type below α_{m-1}, she wins for sure; and when she advances against an opponent with type $s_j \in [\alpha_{m-1}, \alpha_m]$, she advances whenever $\beta(s_j) < b$, or $s_j < \beta^{-1}(b)$, which occurs with probability $\frac{\beta^{-1}(b)-\alpha_{m-1}}{\alpha_m - \alpha_{m-1}}$. (When she advances against an opponent with type above α_m, she loses for sure, so these events don't contribute to Q_σ.) Combining this with the probability of advancing against different types, this gives

$$Q_\sigma(m,b) \quad = \quad (\alpha_{m-1})^{N-1} + \sum_{j=1}^{N-1} \binom{N-1}{j} (\alpha_m - \alpha_{m-1})^j (\alpha_{m-1})^{N-1-j} \left(\tfrac{2}{j+1}\right) \left(\tfrac{\beta^{-1}(b)-\alpha_{m-1}}{\alpha_m-\alpha_{m-1}}\right).$$

(Note that this expression is not valid for "off-equilibrium-path" combinations of m and b.) Also note that, like L_α, $Q_\sigma(m,b)$ depends on τ only through the two nearest thresholds α_{m-1} and α_m.

Finally, we define $u_\sigma(s_i)$ as a bidder's expected payoff given her type and equilibrium play, $u_\sigma(s_i) = v_\sigma(\tau(s_i), \beta(s_i); s_i)$. Since (τ, β) must be a best-response to itself,

$$u_\sigma(s_i) \quad = \quad \max_{m,b} \{(s_i - b)Q_\sigma(m,b) - L_\alpha(m)c\}$$

We can apply the envelope theorem: the objective function depends on s_i directly only through its first term, so

$$u_\sigma(s_i) \quad = \quad u_\sigma(0) + \int_0^{s_i} Q_\sigma(\tau(s), \beta(s))ds$$

In equilibrium, bidders with type 0 must opt out, and so $u_\sigma(0) = 0$. Equating the two expressions for u_σ then gives

$$u_\sigma(s_i) \quad = \quad \int_0^{s_i} Q_\sigma(\tau(s), \beta(s))ds \quad = \quad (s_i - \beta(s_i))Q_\sigma(\tau(s_i), \beta(s_i)) - L_\alpha(\tau(s_i))c$$

As noted above, $\{\alpha_{m-1}, \alpha_m\}$ are enough to determine $L_\sigma(m)$ and $Q_\sigma(m, \beta(s_i))$ for $s_i \in [\alpha_{m-1}, \alpha_m]$; so if $\{\alpha_0, \alpha_1, \ldots, \alpha_m\}$ are known, then the last expression can be rearranged to give

$$s_i - \beta(s_i) \quad = \quad \frac{1}{Q_\sigma(\tau(s_i), \beta(s_i))} \left(\int_0^{s_i} Q_\sigma(\tau(s), \beta(s))ds + L_\alpha(\tau(s_i))c \right) \qquad (5)$$

and $\beta(s_i)$ is uniquely determined.

This, then, gives us the pieces to try to construct a symmetric equilibrium. With second-price auctions, we worked "down from the top"; here, we work "up from the bottom." We begin by choosing a candidate value for α_0, and set $\beta(\alpha_0) = 0$.[15] (Monotonicity of τ and β, the indifference conditions (4), $\beta(\alpha_0) = 0$, continuity of β, and (5) together constitute necessary and sufficient conditions for symmetric equilibrium, and are all satisfied by our construction.)

The indifference condition at α_0 can be written as

$$v_\sigma(1, 0; \alpha_0) \quad = \quad \alpha_0^{N-1} \cdot \alpha_0 - L_\alpha(1) \cdot c \quad = \quad 0$$

While it's not immediately obvious, $L_\alpha(1)$ is strictly increasing in α_1, and so given α_0, this gives a unique value of α_1 satisfying the indifference condition at α_0.

Once $\{\alpha_0, \alpha_1\}$ are known, this determines $Q_\sigma(\tau(s_i), \beta(s_i))$ for $s_i \in (\alpha_0, \alpha_1)$, and therefore, via (5), $\beta(s_i)$ as well; since β must be continuous, $\beta(\alpha_1)$ is therefore determined. Since β cannot jump discontinuously at α_1, the indifference condition at α_1 can then be written as

$$\left(Q_\sigma(2, \beta(\alpha_1)) - Q_\sigma(1, \beta(\alpha_1))\right)(\alpha_1 - \beta(\alpha_1)) \quad = \quad (L_\alpha(2) - L_\alpha(1))\, c$$

Crucially, since a bidder with type α_1 will never win against an opponent who sent message 2, α_2 does not enter into $Q_\sigma(2, \beta(\alpha_1))$; so the left-hand side is determined by $\{\alpha_0, \alpha_1\}$. Thus, the only term that depends on α_2 is $L_\alpha(2)$, which is strictly increasing in α_2; so this last expression determines α_2.

Once $\{\alpha_0, \alpha_1, \alpha_2\}$ are all known, this determines $\beta(s_i)$ up to α_2, and in particular, $\beta(\alpha_2)$. Once $\beta(\alpha_2)$ is known, the indifference condition at α_2 uniquely pins down α_3. Thus, a choice of a candidate value for α_0 uniquely determines each successive threshold α_m, along with the bid function, if these strategies are to be part of a symmetric equilibrium.

We continue to iterate in this way until α_M is reached. α_M must be equal

[15] Since β is monotonic, bidders with types α_0 only win when they enter unopposed; knowing this, if $\beta(\alpha_0)$ were positive, then bidding below $\beta(\alpha_0)$ would be a profitable deviation for a bidder with type α_0.

to $\overline{S} = 1$ in equilibrium; if it is too low or too high, we adjust our initial guess at α_0. Once we find a starting point α_0 at which $\alpha_M = 1$, this gives us an equilibrium for the case where $M = \overline{M}$.[16]

Since $L_\alpha(m) \geq Q_\sigma(m, b)$, equation (5) implies that bidders always "shade" their bids by more than c, so that when they win, they earn strictly positive surplus. This "no regret" property implies that the NITS condition required in equilibrium when $M < \overline{M}$ can never hold in the case of first-price auctions. Since β is monotonic, a bidder with the highest possible type $s_i = 1$ will always win for sure when she advances, and always earns a positive net surplus from winning; so given an opportunity, she would always increase her indicative bid to be sure to advance. (Contrast this with a second-price auction, where a bidder can win but still fail to recoup her participation costs if her opponent's type is sufficiently close to her own.) Thus, unlike with second-price auctions, first-price auctions do not seem to offer an upper bound on the number of messages that can be used in equilibrium, and symmetric equilibrium can only exist when $M = \overline{M} < \infty$.

The numerical example from earlier can be used to illustrate the properties of the symmetric equilibrium in first-price auctions. In that example, $c = 5$ and $N = 5$, and bidder valuations are independently drawn from a uniform distribution on the interval $[0, 100]$. Table 2 reports the equilibrium thresholds for various values of $\overline{M}$.

For each value of $\overline{M}$, there is a unique symmetric equilibrium. The intervals are narrower at higher messages, so there is finer sorting at the top of the type space, particularly at higher values of $\overline{M}$. As in the second-price mechanism, as $\overline{M}$ increases, expected revenue increases due to the participation and selection effects, bidder surplus decreases, and total surplus increases. However, even though bidders now use all of the available messages, the impact

[16]We don't have a formal proof that this procedure will always work. A key step in the construction is that $L_\alpha(m)$, which depends only on α_{m-1} and α_m, possesses a single-crossing property in α_m. This is the case in our restricted model, so we know that *at most* one value of α_m will satisfy the indifference equation at α_{m-1}. But we do not know for sure that *any* value necessarily will satisfy the indifference condition. Still, using this procedure numerically to find equilibria in our restricted model for various values of c, N, and M, we have not yet encountered a case where an equilibrium fails to exist.

Table 2: $N = 5$, $F = U[0, 100]$, $c = 5$, $n = 2$, first-price auction, various $\overline{M}$

	Opt-in messages available ($\overline{M}$)					
	1	2	3	4	5	10
α_{10}	–	–	–	–	–	100.00
α_9	–	–	–	–	–	99.94
α_8	–	–	–	–	–	99.87
α_7	–	–	–	–	–	99.73
α_6	–	–	–	–	–	99.44
α_5	–	–	–	–	100.00	98.83
α_4	–	–	–	100.00	97.71	97.56
α_3	–	–	100.00	95.21	94.91	94.89
α_2	–	100.00	89.92	89.27	89.21	89.20
α_1	100.00	78.47	76.96	76.82	76.81	76.81
α_0	51.50	48.54	48.26	48.24	48.23	48.23
Revenue	53.67	57.40	57.69	57.71	57.71	57.71
Bidder Surplus	16.96	15.45	15.33	15.32	15.32	15.32
Total Surplus	70.63	72.84	73.02	73.03	73.03	73.03

of the additional messages becomes minimal quite quickly: the lower thresholds asymptote, as do expected payoffs. The additional messages only serve to divide up the very top of the type space more finely. For example, when $\overline{M} = 10$, the top four messages are used only by bidders with types above 99.44, and 90% of bidders who opt in use message 1, 2 or 3. In fact, up to the precision shown in the table, revenue and bidder surplus do not change with $\overline{M}$ once $\overline{M}$ is above 4. Thus, even though there is no "natural" upper bound on the number of messages, this property turns out not to be payoff-relevant. In this example, revenue and total surplus are marginally higher under the first-price than under the second-price mechanism, but this has nothing to do with the extra messages; it appears to be driven by the equilibria shown in Table 2 having a slightly lower opt-in threshold α_0, and having thresholds which are more evenly distributed across the interval $[0, 100]$, leading to more participation and better sorting. (Consistent with our earlier discussion, this also leads to bidder surplus being lower in the first-price mechanism.)

4 Equilibrium Existence (General Model)

In this section, we return to the more general model introduced in section 2 and investigate existence and uniqueness of a symmetric equilibrium for general indicative bidding mechanisms with second-price auctions. In particular, we allow F to be any well-behaved distribution on $[\underline{S}, \overline{S}]$, and n, the number of bidders advancing to the auction, to be greater than 2. We also consider mixed strategy equilibria along with pure strategy ones. We first establish necessary and sufficient conditions for symmetric equilibria. We then prove that a symmetric equilibrium exists, and provide a sufficient condition on F for uniqueness. Proofs are given in the appendix. Since bidding is in dominant strategies, we will assume throughout that $\beta(s_i) = s_i$, and suppress the dependence of payoffs on β.

We begin by showing that the main conditions we imposed on τ when considering the simple model are necessary conditions for equilibrium.

Lemma 1. *If τ is a symmetric equilibrium strategy, then τ is nondecreasing*[17] *and has support $\{0, 1, 2, \ldots, M\}$ for some finite M.*

Given these two necessary conditions, τ can be described by a finite set of thresholds α, as in equation (1), but with $\overline{S}$ and $\underline{S}$ replacing 1 and 0 as the top and bottom of the support. While we allow for mixed strategies, Lemma 1 implies that only a measure zero of bidder types mix, so we stick to the simpler notation of pure strategies. Monotonicity of τ implies that $\tau(\alpha_m)$ must be either m, $m+1$, or some mix of the two, and $\tau(\overline{S})$ can be any mix over $\{M, M+1, \ldots, \overline{M}\}$.

As in the previous section, the description of a strategy by its thresholds makes it easier to express the expected payoff to a bidder from sending a given message. However, we will need to adjust those expressions to allow for the cases where more than two bidders advance to the auction. We start once again by defining the net payoff that a bidder can expect to earn in the auction conditional on advancing. Suppose that when bidder i advances to the

[17] For mixed strategies, this means increasing via the strong set order: if $\tau(s)$ puts positive probability on m and $\tau(s')$ puts positive probability on m', then $s' > s$ implies $m' \geq m$.

auction, she faces k opponents with types drawn from the interval $[a, b]$ (and possibly additional opponents with types lower than a). Then the CDF of her highest opponent's type is $\left(\frac{F(\,\cdot\,)-F(a)}{F(b)-F(a)}\right)^k$, and her expected payoff is therefore

$$V(s_i, k, [a, b]) \quad \equiv \quad \int_a^b [\max\{0, s_i - s\} - c] \, d\left(\frac{F(s) - F(a)}{F(b) - F(a)}\right)^k .$$

Note that the expected price that bidder i pays if she wins depends on the number of opponents who sent the highest message, but not on bidder i's message, the messages of opponents who advanced with lower messages, or the messages of opponents who did not advance.

Given her own type s_i and her opponents' strategy τ, we can write bidder i's expected payoff from sending message m (and bidding s_i if selected) as

$$\begin{aligned}
v_\alpha(m, s_i) \quad &= \quad F(\alpha_0)^{N-1}(s_i - c) \\
&+ \sum_{m'=1}^{m-1} \sum_{h=1}^{N-1} \binom{N-1}{h} (F(\alpha_{m'}) - F(\alpha_{m'-1}))^j F(\alpha_{m'-1})^{N-1-h} \\
&\qquad V(s_i, \min\{n-1, h\}, [\alpha_{m'-1}, \alpha_{m'}]) \\
&+ \sum_{j=1}^{N-1} \binom{N-1}{j} (F(\alpha_m) - F(\alpha_{m-1}))^j F(\alpha_{m-1})^{N-1-j} \\
&\qquad \min\left\{1, \tfrac{n}{j+1}\right\} V(s_i, \min\{n-1, j\}, [\alpha_{m-1}, \alpha_m]) \\
&+ \sum_{k=1}^{n-1} \sum_{j=0}^{N-2} \binom{N-1}{k} \binom{N-1-k}{j} (1 - F(\alpha_m))^k (F(\alpha_m) - F(\alpha_{m-1}))^j F(\alpha_{m-1})^{N-1-k-j} \\
&\qquad \min\left\{1, \tfrac{n-k}{1+j}\right\} V(s_i, k, [\alpha_m, \overline{S}])
\end{aligned} \tag{6}$$

Similar to equation (3) earlier, the first line covers opponent type profiles where bidder i is the only one to opt in; the second term covers profiles where the highest opt-in message sent by i's opponents is $m' < m$; the third line covers profiles where the highest opponent message is m; and the fourth line covers profiles where the at least one opponent sends a message above m. As in the restricted model, $v_\alpha(m, s_i)$ does not depend on the thresholds higher than α_m: in the event that bidder i sends message m and advances to the second round, every opponent with type above α_m, and who therefore sends a message higher than m, will have advanced as well, so the exact messages they sent are irrelevant. $v_\alpha(m, s_i)$ does depend on the thresholds below α_{m-1}: when bidder i's strongest opponent sends a message lower than m, there is some chance that he will tie with at least $n-1$ other bidders and not advance; the likelihood that this happens depends on how the interval $[\underline{S}, \alpha_{m-1})$ is partitioned. However, since bidder i is certain to advance in this event, the payoff associated with it is the same whether she sends message m or $m+1$. As a result, the thresholds below α_{m-1} drop out of the difference $v_\alpha(m+1, s_i) - v_\alpha(m, s_i)$, which then depends only on the three thresholds $\{\alpha_{m-1}, \alpha_m, \alpha_{m+1}\}$.

Next, we show that two additional conditions are both necessary and sufficient for τ to be a symmetric equilibrium.

Lemma 2. *Suppose a strategy τ is nondecreasing, and described by a finite series of thresholds $\underline{S} < \alpha_0 < \alpha_1 < \cdots < \alpha_{M-1} < \alpha_M = \overline{S}$ for some $M \leq \overline{M}$. Then τ is a symmetric equilibrium if and only if*

1. *$v_\alpha(m+1, \alpha_m) = v_\alpha(m, \alpha_m)$ for each $m \in \{0, 1, \ldots, M-1\}$, and*
2. *either $M = \overline{M}$ or $v_\alpha(M+1, \overline{S}) \leq v_\alpha(M, \overline{S})$, and $\tau(\overline{S}) = M$ if $v_\alpha(M+1, \overline{S}) < v_\alpha(M, \overline{S})$*

The first condition is that threshold types – types at the boundary between two messages – must be indifferent between the two. If $M < \overline{M}$, then bidders have a credible way to signal that they are high-type and guarantee selection; so the final condition requires that, unless $M = \overline{M}$, such deviations to unused messages must not be profitable. As noted earlier, this condition is the analogue of the No Incentive To Separate (NITS) condition introduced by Chen, Kartin, and Sobel (2008).

Lemma 2 simplifies the search for symmetric equilibria to a search for sets of thresholds α that satisfy the M indifference conditions and the terminal condition. Next, we establish that we can find such thresholds, proving existence of symmetric equilibrium.

As noted above, the difference $v_\alpha(m+1, s_i) - v_\alpha(m, s_i)$ depends only on s_i, α_{m-1}, α_m, and α_{m+1}. Further, when we evaluate this difference at $s_i = \alpha_m$, it satisfies a strict single-crossing property in α_{m-1}; so for a given choice of α_m and α_{m+1}, there is a unique value of α_{m-1} that satisfies $v_\alpha(m+1, \alpha_m) = v_\alpha(m, \alpha_m)$. We will call this value $\alpha^*(\alpha_m, \alpha_{m+1})$:

Definition. *Given an environment (N, F, c, n) and thresholds α_m and α_{m+1}, define*

$$\alpha^*(\alpha_m, \alpha_{m+1})$$

as the unique value of α_{m-1} on $[\underline{S}, \alpha_m]$ satisfying $v_\alpha(m+1, \alpha_m) = v_\alpha(m, \alpha_m)$, or as $-\infty$ if no such value exists.

This definition will allow us to proceed as in the previous section and construct candidate equilibria from the top down. We set $\alpha_M = \overline{S}$, guess at a value of α_{M-1}, and then find the unique value of α_{M-2} that makes α_{M-1} indifference between messages $M-1$ and M; then find the unique value of α_{M-3} that makes α_{M-2} indifferent between messages $M-2$ and $M-1$; and so on, until all thresholds down to α_0 are determined. However, instead of choosing a value of M and then checking to see if the solution satisfies the terminal condition and lies in the interval $[\underline{S}, \overline{S}]$, we can first characterize the values of M for which our procedure works.

Definition. *Given N, F, c, and n, define M^* as the largest value of M such that if we define $\tilde{\alpha}_{M-1} = \tilde{\alpha}_M = \overline{S}$, and iteratively define*

$$\tilde{\alpha}_m \quad = \quad \alpha^*(\tilde{\alpha}_{m+1}, \tilde{\alpha}_{m+2})$$

for each $m = M-2, M-3, \ldots, 0$, then all of these $\tilde{\alpha}_m > \underline{S}$ and $v_{\tilde{\alpha}}(1, \tilde{\alpha}_0) > 0$.

In the appendix, we define M^* in a slightly more formal way (and prove that it is indeed a well-defined value), but the two definitions are equivalent. Intuitively, M^* is the maximal number of thresholds $\alpha_0 < \alpha_1 < \ldots < \alpha_{M^*-1}$

we can fit into the type space $[\underline{S}, \overline{S}]$, letting $\alpha_{M^*-1} = \alpha_{M^*} = \overline{S}$ but satisfying the indifference condition at α_m for each $m = 1, 2, \ldots, M^* - 1$, and satisfying $v_\alpha(\alpha_0, 1) > 0$. With M^* defined, we get the following result:

Theorem 1. *Fix the environment (N, F, c) and the indicative bidding mechanism $(n, \overline{M})$. Define M^* as above given N, F, c, and n.*

1. *If $\overline{M} < M^*$, then a symmetric pure-strategy equilibrium exists using the messages $\{0, 1, 2, \ldots, \overline{M}\}$.*

2. *If $\overline{M} \geq M^*$, then a symmetric pure-strategy equilibrium exists using the messages $\{0, 1, 2, \ldots, M^*\}$.*

3. *If the density function f of F is nonincreasing on $[\underline{S}, \overline{S}]$, then the symmetric equilibrium is essentially unique.*[18]

Thus, a symmetric equilibrium always exists with $\min\{M^*, \overline{M}\}$ messages; and if f is decreasing, this is the unique symmetric equilibrium.[19] The uniqueness result depends on an additional monotonicity property of α^*, which can only be established when f is nonincreasing.[20] This monotonicity property can then be used to show that for a given M, there is a unique α satisfying all the indifference conditions, as well as the fact that the NITS condition $v_\alpha(M+1, \overline{S}) \leq v_\alpha(M, \overline{S})$ is always violated for $M < M^*$, which together imply uniqueness.

[18] That is, all symmetric equilibria have the same support and are described by the same series of thresholds, and are therefore payoff-equivalent from an ex ante perspective.

[19] In fact, for uniqueness to hold, f does not need to be decreasing over the entire type space $[\underline{S}, \overline{S}]$; it suffices for f to be nonincreasing on the upper part of its support $[\alpha_1, \overline{S}]$.

[20] Formally, the property ensures that when the intervals $[\alpha_m, \alpha_{m+1}]$ and $[\alpha_{m+1}, \overline{S}]$ get wider – in the probability weight sense – so does the interval $[\alpha^*(\alpha_m, \alpha_{m+1}), \alpha_m]$. This implies that during our construction of intervals, as the top interval $[\alpha_{M-1}, \overline{S}]$ gets wider, so does every lower interval $[\alpha_m, \alpha_{m+1}]$. This then leads to uniqueness of symmetric equilibrium. It also implies the fact noted earlier that when $M = M^*$, equilibrium intervals are narrower at the top of the type space – that is, $F(\alpha_{m+1}) - F(\alpha_m)$ is decreasing in m.

5 Efficiency and Revenue

In this section, we return to the numerical example used earlier to explore two points: the tradeoffs involved from allowing more bidders to advance to the auction, and a comparison between indicative bidding mechanisms and auctions with unrestricted entry. We then show that most of the results shown numerically for the example – specifically, the Pareto-optimality of $n = 2$, and the Pareto-dominance of indicative bidding mechanisms over unrestricted auctions – can be proven to hold generally when the number of potential bidders N is sufficiently large.

5.1 How Many Bidders to Advance

First, we examine the effect of the number of bidders advancing to the second round. We use the example from earlier – where F is the uniform distribution on $[0, 100]$, $c = 5$, and $N = 5$ – and focus on the case where $\overline{M}$ is sufficiently large that $M = M^*$. Table 3 reports the equilibrium partition and payoffs for each value of n ranging from 2 to 5.

Table 3: $N = 5$, $F = U[0, 100]$, $c = 5$, various n, $M = M^*$

Bidders advancing to second round (n):	2	3	4	5
α_3	100.00	–	–	–
α_2	98.12	100.00	100.00	–
α_1	83.64	91.04	94.26	100.00
α_0	49.42	53.93	54.82	54.93
Revenue	57.26	56.19	55.87	55.82
Bidder Surplus	15.42	14.27	13.99	13.96
Total Surplus	72.68	70.46	69.86	69.78

As n increases, a bidder's chance of being selected when there is strong competition goes up. As a result, bidders are less willing to participate and less willing to "separate" themselves from lower types, so the participation threshold α_0 goes up and the number of opt-in messages used in equilibrium goes down. (In the example, as n increases from 2 to 5, the expected number of bidders who opt in decreases from 2.53 to 2.25, and the number of opt-in messages used falls from 3 to 1.) Both of these reduce revenue, and the greater incidence of entry costs reduces bidder surplus. Thus, both sides of the market benefit from eliminating redundant entry costs by restricting entry more tightly.

The results of Tables 1 and 3 suggest that the seller faces no tradeoff between efficiency and revenues in choosing the details of the indicative bidding mechanism used: both revenue and total surplus are increasing in $\overline{M}$ and decreasing in n. In this particular setting, this implies that a seller should not constrain the number of messages that bidders can send, but should restrict the number of bidders who can advance to no more than two. We conjecture that this result holds more generally (we have not found any counterexamples), but have not been able to prove it.[21]

5.2 Comparison to an Unrestricted Auction

In this section, we evaluate the performance of the indicative bidding mechanism. The benchmark we use is the Samuelson (1985) model of costly entry, in which the seller does not ration entry, but allows bidders to choose independently whether or not to enter. The timing is the same as in our model: bidders learn their type, decide simultaneously whether to enter, and then those who chose to enter incur the cost of due diligence and submit binding bids. We focus on the case of a second-price auction, so bidding truthfully remains a dominant strategy. (By revenue equivalence, the outcome would be identical with a first-price auction.) The symmetric equilibrium involves a cut-off strategy, in which bidders enter when their type s_i is above some threshold γ, which we refer to as the entry threshold. A bidder with type $s_i = \gamma$ who chooses to enter will lose the auction unless she is the only entrant, in which

[21]The main difficulty lies in quantifying the selection effect. As either $\overline{M}$ or n changes, the resulting equilibrium partitions are not ordered – neither is a refinement of the other – so there are always some specific type profiles for which fewer messages or more bidders advancing would yield higher expected revenue.

case she will pay 0 for the asset; so γ satisfies the indifference condition

$$0 = F(\gamma)^{N-1}\gamma - c.$$

which defines the unique symmetric equilibrium.

Comparing an indicative bidding mechanism to an unrestricted auction from the seller's point of view, the main tradeoff is between the cost of not selling the asset at all and the cost of not selecting the bidders with the highest types. Bidders are more willing to opt in under indicative bidding, since they have a chance of avoiding entry costs when stronger bidders opt in as well, so the likelihood of the asset being sold rises. On the other hand, since ties can occur, the seller sometimes fails to select the bidders with the two highest valuations. To illustrate this tradeoff and how it varies with N, we return to our example from the previous sections, where F is the uniform distribution on $[0, 100]$ and $c = 5$. We compute the equilibrium entry threshold γ for the unrestricted auction, and based on that, calculate equilibrium expected payoffs. In computing the equilibrium payoffs of the indicative bidding mechanisms, we focus on the case where $n = 2$, and $\overline{M}$ is either large to not be a binding constraint (second price auction) or for payoffs to have asymptoted (first price auction). Table 4 compares outcomes under the three mechanisms.

The values of α_0 reported in Table 4 imply that for either indicative bidding mechanism, both the probability of a sale (at least one bidder opting in) and the probability of positive revenue (at least two bidders opting in) increase with N. The former is good for efficiency and latter for revenues. For the second price mechanism, the equilibrium number of messages used, M^*, is weakly decreasing in N, so the selection effect goes in the opposite direction. However, the participation effect is sufficiently strong that both revenue and total surplus increase with N, while bidder surplus decreases with N.

The opt-in thresholds for the indicative bidding mechanisms are always lower than the entry threshold for the auction with unrestricted entry. This implies that, for each N, the number of bidders opting in under indicative bidding first-order stochastically dominates the number of bidders entering

Table 4: Various N, $F = U[0, 100]$, $c = 5$, comparing different mechanisms

Standard auction with unrestricted entry

Potential buyers (N):	3	4	5	7	10	20	50
Entry threshold γ	36.84	47.29	54.93	65.18	74.11	86.09	94.18
Revenue	42.76	50.67	55.82	62.11	67.14	73.37	77.31
Bidder Surplus	21.38	16.89	13.96	10.35	7.46	3.86	1.58
Total Surplus	64.14	67.57	69.78	72.46	74.60	77.23	78.88

Indicative bidding, second price auction, $n = 2$, $\overline{M} \geq M^*$

Potential buyers (N):	3	4	5	7	10	20	50
M^*	3	3	3	3	2	2	1
α_0 with $M = M^*$	33.75	42.51	49.42	59.55	69.18	83.22	92.99
Revenue	42.99	51.51	57.26	64.49	70.46	77.96	83.34
Bidder Surplus	22.34	18.22	15.42	11.87	8.92	5.13	2.39
Total Surplus	65.33	69.73	72.68	76.37	79.38	83.09	85.73

Indicative bidding, first price auction, $n = 2$, $\overline{M}$ "sufficiently large"

Potential buyers (N):	3	4	5	7	10	20	50
$\lim \alpha_0$ as $\overline{M}$ grows	33.16	41.53	48.23	58.21	67.90	82.30	92.64
Revenue	43.13	51.82	57.71	65.17	71.31	79.15	84.29
Bidder Surplus	22.35	18.17	15.32	11.68	8.65	4.73	2.09
Total Surplus	65.48	69.99	73.03	76.85	79.97	83.88	86.38

under unrestricted entry. Hence, the probability of a sale and the probability of two or more bidders is higher under indicative bidding, which is good for efficiency and revenue. Of course, there is still a tradeoff: since entry into the second round of the indicative bidding mechanism is capped at two bidders, the expected revenue conditional on at least two entrants is lower than it would be with unrestricted entry, as the seller may select the "wrong" bidders, reducing both efficiency and revenues. Nonetheless, once again, the participation effect appears to consistently dominate the selection effect: for each N, revenue, bidder surplus and total surplus are all higher under indicative bidding.[22]

5.3 The "Large N" Case

Next, we consider what happens when the number of potential bidders N gets large.

In the Samuelson model, it is straightforward to show that as N goes to infinity, the cutoff $\gamma(N)$ converges to $\overline{S}$, and the probability that no bidders enter goes to a constant $\theta \equiv c/\overline{S} < 1$. This implies that as N goes to infinity, the distribution of the number of bidders in the auction converges to a Poisson distribution with parameter $-\ln\theta$. Equilibrium payoffs can be computed from this distribution. They depend on the probabilities of three events: the probability of no sale, the probability of a sale with no revenue, and the probability of a sale with revenue. In the latter case, revenue is $\overline{S}$, since in the limit only bidders with values arbitrarily close to $\overline{S}$ advance. Under the Poisson distribution, the probability of no bidders entering is θ, and the probability of only one bidder is $-\theta\ln\theta$, so in the limit, expected revenue is given by

$$R = (1-\theta+\theta\ln\theta)\overline{S}.$$

Expected total surplus is simply gross bidder surplus minus expected costs, or

$$W = (1-\theta)\overline{S} + (\ln\theta)c,$$

and bidders surplus is zero.

In deriving the limit results for the indicative bidding mechanism, we focus on the case of second-price auctions. The following Lemma establishes that, as

[22] This dominance result does depend upon the message space being unconstrained. A comparison with the payoffs reported in Tables 1 and 2 reveals that the revenue ranking reverses when the number of opt-in messages is constrained to be one: if $\overline{M} = 1$, either indicative bidding yields less revenue than an unrestricted auction. The indicative bidding mechanism still attracts more participation than the unrestricted auction, but this is overshadowed by the revenue loss from the "wrong" bidders being selected.

N gets large, only messages 0 and 1 are played in the symmetric equilibrium of this mechanism.

Lemma 3. *Fix (F, c, n). There exists an $\widehat{N} < \infty$ such that if $N > \widehat{N}$, there is a unique symmetric equilibrium, with $M = 1$.*

The proof is in the appendix. The rough intuition is that if an equilibrium existed with $m = 2$, a bidder with type α_1 would be indifferent between sending messages 1 and 2, which requires $\alpha_1 - \alpha_0 > c$. But as N gets large, $\alpha_0 < \overline{S} - c$ would mean a bidder with type close to α_0 would have no chance of advancing alone, and would therefore do better opting out than sending message 1, generating a contradiction.

Given Lemma 3, when N is large, the equilibrium is characterized by the lone indifference condition $v_\alpha(1, \alpha_0) = 0$. It is not difficult to show that, as N goes to infinity, the opt-in threshold $\alpha_0(N)$ goes to $\overline{S}$, but at a rate such that the probability of all bidders opting out goes to a constant ϕ which is bounded away from 0 and 1. Thus, the number of bidders who opt in converges in distribution to a Poisson distribution with parameter $-\ln \phi$. As in the Samuelson model, equilibrium payoffs are easily calculated from this distribution.

Since $\alpha_0 \to \overline{S}$ as N gets large, the selection effect becomes negligible, as all bidders who advance are arbitrarily close in valuations; but the participation effect does not. For this case, we are able to prove several results.

Theorem 2. *Fix F and c. When N is sufficiently large:*

1. *Revenue, bidder surplus, and total surplus are decreasing in n;*

2. *For any n, revenue, bidder surplus, and total surplus are above that of an unrestricted auction;*

3. *Entry is below the efficient level, but above the level of an unrestricted auction.*

Result 1 implies that, when N is large, $n = 2$ is Pareto optimal. Result 2 follows from the fact that lowering costs by capping the number of bidders in the auction increases the bidders' participation rate. A higher participation rate increases revenues, because it makes the event of two or more bidders entering the auction more likely; and it raises total surplus, because it makes the event of a sale more likely. Total surplus also increases because, when more than n bidders opt in, only n advance, so entry costs incurred are lower. Finally, while bidder surplus in both mechanisms is zero in the limit, we show that it is of order $1/N$, calculate the leading term, and prove that it is greater in the indicative bidding mechanism.

The inefficiency part of Result 3 may seem surprising. In the limit, since all bidders who consider entering have identical valuations, the model is similar to the entry model of Levin and Smith (1994), who show that entry in their setting (an unrestricted auction) is efficient. The difference is that with indicative bids, one bidder's decision to opt in has an additional effect relative to an unrestricted auction: she may "displace" another bidder, who would have been selected in her absence. This only occurs when at least $n-1$ other bidders opted in; when N is large, the displaced bidder would therefore have had a negative expected payoff, so displacing him imposes a positive externality. Since bidders who opt in impose a positive net externality, this implies participation is below the efficient level.

As noted above, when N gets large, there remains a nonvanishing chance that exactly one bidder opts in, and buys the asset at price 0. Thus, it might seem that, even in the limit, the seller could increase revenue by using a reserve price. However, Theorem 2 implies that this intuition is wrong. When N is large, bidder surplus is effectively zero, so ex ante, the seller is capturing all of the surplus. A reserve price would reduce participation, leading to lower total surplus and therefore lower revenue. On the other hand, subsidizing participation would increase participation, thereby increasing total surplus, and because N is large, the seller would capture this additional surplus.

Corollary 1. *Fix F, c, and n. When N is sufficiently large, a positive reserve price would strictly decrease revenue; a small bidder subsidy would strictly increase revenue.*

6 Value Updating During Due Diligence

An interesting extension of our model is one in which due diligence provides new information about the value of the asset. Ye (2007) introduced such a value updating model, and we adopt his approach. Each bidder's private value is equal to $S_i + T_i$; bidder i knows S_i initially, but learns T_i only after advancing to the auction and performing due diligence. The realizations of both signals are private information for bidder i. Let F_T denote the joint distribution of $\{T_i\}$. We assume that $\{T_i\}$ are symmetrically distributed and independent of $\{S_i\}$, but they need not be independent of each other.

A special case is when $\{T_i\}$ are perfectly correlated, and therefore the same for all bidders. In this case, $T_i = T$ represents the component of value that is common to all bidders, which bidders learn from due diligence. Bidders who advance compete away any rents associated with learning the realization of T, so assuming more than one bidder advances, each bidder's surplus is still $\max\{0, s_i - S^*\}$, where S^* is the highest of the initial types s_j of the other bidders who advanced. Thus, the extended model in which $\{T_i\}$ are perfectly correlated is equivalent to our baseline model.[23] We can therefore think of our baseline model as covering the case where new information learned during due diligence is about a component of value common to all bidders.

If, however, the new signals $\{T_i\}$ are not perfectly correlated, then bidders are getting new idiosyncratic information about their private values; in this case, the second private signal is a source of additional rents for the bidders who advance, and increases the expected payoffs from opting in. It also gives the seller an added incentive to allow more bidders to advance, since higher n means more draws from the distribution of T_i – any bidder could turn out to be the strongest once the $\{T_i\}$ are realized. The main technical difficulty in analyzing the model with private value updating is that the indifference conditions characterizing the equilibrium partition become more complicated as more bidders advance. In the baseline model (or with common value updat-

[23]The only difference is that the payoff to a bidder from advancing alone is now $s_i + E(T_i) - c$ instead of $s_i - c$, which can easily be accommodated in the baseline model by shifting the distribution of S_i to the right by $E(T_i)$.

ing), a bidder only cares about the highest rival type that advances, because she loses for sure when that type is higher than hers and wins for sure (at a price equal to that type) when it is lower than hers. By contrast, when updating is about private values, bidders' auction payoffs depend on the types of all rivals who advance, since a bidder can still win against a bidder with a higher initial type or lose against a bidder with lower initial type. As a result, the number of terms that need to be accounted for in the difference $v_\alpha(m+1,\alpha_m) - v_\alpha(m,\alpha_m)$ increases geometrically with the number of bidders who can advance. Therefore, for tractability purposes, we restrict the analysis of the private value updating model to the case of $n = 2$.

We also impose one additional restriction on the information learned after entry:

$$E\{\max\{0, T_i - T_j\}\} \quad < \quad c.$$

This restriction states that in expectation, the information rents accruing based on the information gathered during due diligence are not sufficient to cover its costs. It implies that a bidder prefers to only advance to the auction when her initial type is higher than that of her second-round opponent. Note that this is automatically satisfied in the case of perfect correlation, since $T_i = T_j$. We refer to this inequality as the "small rents" restriction.

Given these two restrictions, we can prove the following result:

Theorem 3. *Fix F, F_T, N, and c, let $n = 2$, and assume small rents.*

1. *A symmetric equilibrium exists in which* $\max\{M^*, \overline{M}\}$ *opt-in messages are used.*

2. *If F is uniform, then the symmetric equilibrium is unique.*

We can also obtain large-N results for the value updating model, and they are analogous to those obtained for the baseline model: as N grows, only messages 0 and 1 get used in equilibrium; entry is below the efficient level, but above the level of an unrestricted auction; revenue and bidder surplus are both strictly higher than in an unrestricted auction; and a reserve price would lower revenue, but a small bidder subsidy would increase it. In addition, for N large, symmetric equilibrium exists for an indicative bidding game with any n, but revenue and total surplus are both decreasing in n.

Theorem 3, and the corresponding large-N results, were proven in an earlier working-paper version of this paper. The main conclusion we draw is that the results obtained in the baseline model, or with pure common value updating, are robust to small changes in the information environment.

7 Alternative Mechanisms

Ye (2007) and Lu and Ye (2014) propose that sellers use an entry rights auction to screen potential bidders for the asset sale. Ye (2007) shows that in a private value updating model, an entry rights auction has a symmetric monotonic equilibrium, which can be used to induce efficient entry. Lu and Ye (2014) characterize the optimal two stage mechanism for this environment, and show it can be implemented using an entry-rights auction followed by a second-price "handicap auction," where a bidder's advantage depends on her first-round bid.

How big are the gains from using bids for entry rights rather than indicative bids to select bidders? To investigate this issue, we return to our baseline model. Without value updating, the optimal mechanism can be implemented using an auction for entry rights in which just one bidder advances and gets the right to claim the asset for free – in effect, the "real" auction is held prior to due diligence. In the numerical example we used above, expected revenue is maximized at a reserve price of \$47.50, and total surplus is maximized at a reserve price of 0. We find that the indicative bidding mechanism captures about 92% of the revenue generated by the optimal mechanism, and a similar percentage of the total surplus generated by the surplus-maximizing mechanism. Thus, the indicative bidding mechanism does well, but the entry rights auction can still do significantly better.

However, the example also highlights the main problem with an entry rights auction. Consider the value updating environment, but with $\{T_i\} = T$ perfectly correlated. In an entry rights auction, bidders must bid based on the expected value of T, and the winner pays for the asset *before* she conducts due diligence and learns the realization of T. But in that case, due diligence is no longer useful, and bidders would be at risk in two ways. First, they will sometimes overpay for assets that ex post turn out to be worth substantially less than their ex ante value, which would make them liable for the losses and vulnerable to a shareholder lawsuit.[24] Second, they face an adverse selection problem. Sellers with worthless assets would have an incentive to enter the market, and all sellers would have an incentive to overstate the value of T (or its ex-ante distribution from the point of view of the bidders). Sellers with legitmate, "high-T" assets, on the other hand, would prefer to be paid based on the realized value of T rather than its ex-ante distribution, and would favor mechanisms in which bidders bid *after* performing due diligence; so entry rights auctions would likely attract a worse mix of assets. A mechanism using indicative bids avoids these problems: since no money is committed until after due diligence, the winning bidder never pays more than the ex post value of the asset, and sellers with worthless assets have no incentive to enter.

The above argument explains why bidders want to conduct due diligence prior to bidding. However, it does not explain why each bidder conducts her own due diligence. If the issue is simply one of verifying the common value component of the asset, then the bidders could hire a third party to perform due diligence and avoid the duplication of costs.[25] The fact that they do not do so suggests that the due diligence process is to some extent bidder-specific, and the information generated is at least partly about idiosyncratic components of value. In that respect, due diligence in corporate takeovers appear more analogous to timber cruises than to seismic surveys.[26]

[24]Bhattacharya, Roberts and Sweeting (2014) make a similar point by noting that an entry rights auction may be less efficient than a standard auction with unrestricted entry if entry costs are high and bidders are risk-averse.

[25]We thank Jakub Kastl for raising this issue.

[26]In timber auctions, bidders typically perform their own cruises to estimate the value of the lot, even when a government cruise report is available. In contrast, in oil and gas auctions, bidders typically buy their seismic surveys from geophysical firms, who specialize in this kind of activity, and who may supply reports to multiple rival bidders.

8 Conclusion

We have shown that indicative bidding can be informative, and often leads to greater efficiency and higher revenue than an auction with unrestricted entry, particularly when the number of bidders is large. We have shown this within a fairly stylized model. But we expect the general insight to hold in a wider range of trading environments and communication protocols: when there are many buyers and entry costs are high, a seller can benefit from using indicative bids to "thin the field," and then either bargain with or hold an auction among a smaller number of buyers.

References

1. 2015 Mergerstat Review (Excerpt), from Business Valuation Resources, LLC, downloaded Sept 1 2015 from http://www.bvresources/pdfs/2015MergerstateReviewExcerpt.pdf

2. Bhattacharya, Vivek, James Roberts, and Andrew Sweeting (2014), Regulating Bidder Participation in Auctions, *RAND Journal of Economics* 45(4)

3. Boone, Audra, and Harold Mulherin (2007), How Are Firms Sold? *Journal of Finance* 62(2)

4. Boone, Audra, and Harold Mulherin (2009), Is There One Best Way to Sell a Company? Auctions Versus Negotiations and Controlled Sales, *Journal of Applied Corporate Finance* 21(3)

5. Bulow, Jeremy, and Paul Klemperer (2009), Why Do Sellers (Usually) Prefer Auctions? *American Economic Review* 99(4)

6. Burguet, Roberto, and Jozsef Sakovics (1996), Reserve Prices Without Commitment, *Games and Economic Behavior* 15(2)

7. Chen, Ying, Navin Kartik, and Joel Sobel (2008), Selecting Cheap-Talk Equilibria, *Econometrica* 76(1)

8. Compte, Olivier, and Philippe Jehiel (2007), Auctions and Information

Acquisition: Sealed Bid or Dynamic Formats? *RAND Journal of Economics* 38(2)

9. Crawford, Vincent, and Joel Sobel (1982), Strategic Information Transmission, *Econometrica* 50(6)
10. Cremer, Jacques, Yossi Spiegel, and Charles Zheng (2009), Auctions with Costly Information Acquisition, *Economic Theory* 38(1)
11. Farrell, Joseph, and Robert Gibbons (1989), Cheap Talk Can Matter in Bargaining, *Journal of Economic Theory* 48(1)
12. Fullerton, Richard, and R. Preston McAfee (1999), Auctioning Entry into Tournaments, *Journal of Political Economy* 107(3)
13. Levin, Dan, and James Smith (1994), Equilibrium in Auctions With Entry, *American Economic Review* 84(3)
14. Lu, Jingfeng, and Lixin Ye (2014), Optimal Two-stage Auctions with Costly Information Acquisition (working paper)
15. Matthews, Steven, and Andrew Postlewaite (1989), Pre-play Communication in Two-Person Sealed-Bid Double Auctions, *Journal of Economic Theory* 48(1)
16. McAfee, Preston, and John McMillan (1987), Auctions with Entry, *Economics Letters* 23(4)
17. Menezes, Flavio, and Paulo Monteiro (2000), Auctions with Endogenous Participation, *Review of Economic Design* 5(1)
18. Samuelson, William (1985), Competitive Bidding with Entry Costs, *Economics Letters* 17(1-2)
19. Subramanian, Guhan (2010), *Negotiauctions*, W. W. Norton & Co. (New York)
20. Ye, Lixin (2007), Indicative Bidding and a Theory of Two-Stage Auctions, *Games and Economic Behavior* 58(1)