THE BIDDER EXCLUSION EFFECT

DOMINIC COEY, BRADLEY LARSEN, AND KANE SWEENEY

Abstract. We introduce a new, simple-to-compute test of independence of valuations and the number of bidders for ascending button auctions with symmetric, conditionally independent private values. The test involves estimating the expected revenue drop from excluding a bidder at random, which can be computed as a scaled sample average of a difference of order statistics. This object also provides a bound on counterfactual revenue changes from optimal reserve pricing or bidder mergers. We illustrate the approach using data from timber auctions, where we find some evidence that bidder valuations and the number of participants are not independent.

Date: July 2018.
We thank Gaurab Aryal, Matt Backus, Tom Blake, Jeremy Bulow, Alex Frankel, Amit Gandhi, Phil Haile, Bob Hammond, Jason Hartline, Jakub Kastl, Tatiana Komarova, Elena Krasnokutskaya, Nicola Lacetera, Jon Levin, Vadim Marmer, Dimitriy Masterov, Rob Porter, Dan Quint, Jimmy Roberts, Art Shneyerov, Paulo Somaini, Steve Tadelis, Caio Waisman, Glen Weyl, and seminar and conference participants at the 2013 Stanford-Berkeley IO Fest, Vanderbilt (2014), USC (2014), the Federal Trade Commission (2014), the 2014 International Industrial Organization Conference, the 2014 North American Econometric Society Meetings, the 2015 European Econometric Society Winter Meetings, the 2015 AEA Meetings, the 2015 Stanford Institute for Theoretical Economics, the 2015 Becker Friedman Advances in Price Theory Conference, the 2016 NBER IO Winter Meetings, and EC 2016 for helpful comments. This work was done while Coey and Sweeney were researchers at eBay Research Labs.
Coey: Facebook, Core Data Science; coey@fb.com.
Larsen: Stanford University, Department of Economics and NBER; bjlarsen@stanford.edu.
Sweeney: Uber, Marketplace Optimization; kane@uber.com.

1. Introduction

A number of recent innovations in empirical methodologies for auctions rely on the assumption that bidders' valuations are independent of the number of bidders participating in the auction. In these articles, it is assumed that when one additional bidder arrives at an auction that originally had n bidders, this additional bidder's valuation represents a random draw from the same data generating process that led to the original n bidders' valuations. In this article we demonstrate that this assumption is easily testable in no-reserve ascending (button) auctions with symmetric, conditionally independent, private values where bidders play the weakly dominant strategy of truthful bidding.[1] The data requirements are that the researcher observe two order statistics of bids and the number of participants. We also demonstrate a number of extensions to this test, including bounding counterfactual revenue under optimal reserve pricing or bidder mergers. We demonstrate that our environmental assumptions can be relaxed in a number of ways.

Throughout the article, we refer to the decrease in expected auction revenue when a random bidder is excluded from the auction as the *bidder exclusion effect*. In an ascending button auction with private values, this effect can easily be computed without the need to estimate a complex model, unlike many objects of interest in auction settings. In such an auction, with n bidders participating, if a bidder is excluded at random from the auction, with probability $\frac{n-2}{n}$ he will be one of the $n-2$ lowest bidders, and so his exclusion will not affect revenue. With probability $\frac{2}{n}$, he will be one of the two highest bidders, and revenue will drop from the second-highest to the third-highest value of the n bidders. The bidder exclusion effect is therefore $\frac{2}{n}$ times the expected difference between the second and third-highest values.

The bidder exclusion effect can yield several diagnostics for ascending auction settings. The first and foremost is that of testing the independence of bidder valuations and the number of bidders. By comparing the bidder exclusion effect in an n bidder auction to the

[1] Throughout the article, we will adopt the phrase *conditionally independent private values*, as used in Li et al. (2003), refering to a setting where bidders have private valuations that are correlated and where there exists a random variable U, unknown to bidders and to the econometrician, such that, conditional on U, bidders' valuations are independent. In the setting we focus on in the main body of the article—that of ascending button auctions—all of results also apply if this random variable U is *known* to the bidders, but still unobserved to the econometrician; such a setting is referred to in the literature as a setting of independent private values with unobserved heterogeneity. The distinction is inconsequential for our main results, but it is important for first price auctions, which we discuss in Appendix B.

actual decrease in revenue between n bidder and $n-1$ bidder auctions observed in the data, the researcher can test whether bidder valuations are indeed independent of the number of bidders. We demonstrate how this test can be performed in practice. The order statistic relationship we exploit here has been used elsewhere by maintaining the assumption that bidder valuations are independent of the number of bidders and instead testing for private vs. common values (Athey and Haile 2002).

Second, the bidder exclusion effect serves as a bound on the revenue gain to a seller from choosing the optimal reserve price, thus aiding the practitioner in deciding whether or not to adopt a reserve price at all. To do so, we rely on the result of Bulow and Klemperer (1996), that adding an additional random bidder does more to improve seller revenue than does an optimal reserve price. Third, the bidder exclusion effect can be used to bound the revenue losses to a seller from counterfactual mergers between bidders.

We evaluate the bidder exclusion effect in US timber auction data. In this setting, we first ask the following question: should the seller—in this case, the government—bother to compute an optimal reserve price? Computing an optimal reserve price can be computationally costly in practice, and mistakenly implementing too high a reserve price can be very detrimental to revenue. The bidder exclusion effect can provide the seller a simple tool for evaluating the size of the potential gains from optimal reserve pricing. We find that an upper bound on this gain is 13% of revenue on average in our data.

We then ask the question, if the seller does wish to compute an optimal reserve price, can she safely rely on exogenous variation in the number of bidders in doing so? Existing methods for computing bounds on the optimal reserve price itself rely on the assumption that bidder valuations are independent of the number of bidders in order to obtain tight, meaningful bounds (e.g., Haile and Tamer 2003, Aradillas-López et al. 2013, and Coey et al. 2017). We apply our test to our data and find evidence against the assumption that bidder valuations are independent of the number of bidders. However, after controlling for bidder asymmetries this evidence is less strong.

We highlight a number of extensions of the auction environments in which the bidder exclusion effect can be used. For example, we demonstrate that it can be computed in ascending button auctions with symmetric common values (i.e. when bidders have symmetric values and symmetric bidding strategies) or in ascending non-button auctions when bidders have private values but may potentially drop out below their values (such as in the setting of Haile and Tamer 2003). We also discuss extensions of these diagnostics to data from ascending auctions with potentially binding reserve prices and to first price auctions. For clarity of exposition—and because it is the most general case in which all of our results immediately apply—we focus the majority of the article on the case of no-reserve ascending button auctions with symmetric, conditionally independent, private values.

2. Related Literature

The use of the bidder exclusion effect to test the assumption that bidder values are independent of the number of bidders is related to several studies that rely on this type of independence for identification or for testing other aspects of auction models. Athey and Haile (2002, 2007) refer to this type of independence as "exogenous participation," and provide a one-sided test, related to ours, for testing private vs. common values, maintaining the assumption of exogenous participation. Aradillas-López et al. (2013) and Coey et al. (2017) refer to this type of independence as "valuations are independent of N," and they, along with earlier work by Haile and Tamer (2003), exploit this assumption to obtain tight bounds on counterfactual seller revenue and optimal reserve prices. Aradillas-López et al. (2016) provide an alternative test of the dependence between valuations and the number of bidders that has the advantage of only requiring that the transaction price be observed (as opposed to two order statistics of bids, as we require) but has the disadvantage of being more complex to compute than our test. Liu and Luo (2017) provide a test for such independence in first price auctions. Aradillas-López et al. (2013) and Aradillas-López et al. (2016) also demonstrate general conditions under which popular entry models, including those of Samuelson (1985), Levin and Smith (1994), and Marmer et al. (2013), will generate the type of dependence between valuations and the number of bidders that our test can detect.

The question of how much a seller would benefit by adopting an optimal reserve price—a quantity which the bidder exclusion effect can be used to bound—has been a counterfactual of interest for a number of empirical auction studies, such as Paarsch (1997), Li and Perrigne (2003), Li et al. (2003), Haile and Tamer (2003), Krasnokutskaya (2011), Tang (2011), Li and Zheng (2012), Aradillas-López et al. (2013), Roberts and Sweeting (2013, 2014), Bhattacharya et al. (2014), and Coey et al. (2017, 2018). A typical empirical approach to answering this question would rely on assumptions about the distribution of values and the information environment to estimate a detailed model, determine optimal reserve prices using the first-order condition for seller profit, and finally measure the revenue difference between the optimally designed auction and a no-reserve auction. An advantage of such an involved procedure, relative to ours, is that it could yield an estimate of the optimal reserve price itself, whereas our tool cannot. Our tool, however, circumvents the need for these steps and yields information about the revenue gain from choosing the optimal reserve price. The bidder exclusion effect, in bounding revenue, is therefore best thought of as being useful as a simple-to-compute initial diagnostic. Tang (2011) derives results for first and second price auctions and is particularly related to our revenue-bounding approach in spirit in that it provides a bound on counterfactual revenue without directly estimating valuations (as do Haile and Tamer 2003, Aradillas-López et al. 2013, Coey et al. 2017, and Chesher and Rosen 2017 in ascending auctions).

Many identification and testing results for first and second price auctions cannot be applied to ascending auctions because of the complicating factor that in ascending auctions the would-be bid of the highest-value bidder will not be observed. For this reason, empirical tools for ascending auction environments beyond independent private values settings have only recently become available. These results include Aradillas-López et al. (2013) and Coey et al. (2017), which apply to ascending auctions with correlated private values, and Hernandez et al. (2018), Freyberger and Larsen (2017), and Chesher and Rosen (2017), which apply to ascending auctions with separable unobserved heterogeneity. Our approach contributes to this literature by providing identification arguments for ascending auctions in a particular type of correlated private values environments, that of conditionally independent private values.

Bidder mergers have also been a focus of a number of studies in the empirical auctions literature, including Froeb et al. (1998), Waehrer (1999), Dalkir et al. (2000), Brannman and Froeb (2000), Tschantz et al. (2000), Waehrer and Perry (2003), Froeb et al. (2008), and Li and Zhang (2015). Our approach provides an initial diagnostic tool for such settings, allowing the researcher a quick-and-easy way to compute bounds on the effect of a counterfactual bidder merger on seller revenue.

In the spirit of Haile and Tamer (2003) and other bounds approaches that have followed (Tang 2011; Armstrong 2013; Komarova 2013; Aradillas-López et al. 2013; Gentry and Li 2013; Coey et al. 2017; Chesher and Rosen 2017; and others), our empirical approach does not seek to point identify and estimate the distribution of bidder values. Instead we draw inferences from functions of the value distribution that are point, or partially, identified. More broadly, our approach ties in closely to the recent literature on "sufficient statistics" for welfare analysis (Chetty 2009; Einav et al. 2010; Jaffe and Weyl 2013), which focuses on obtaining robust welfare or optimality implications from simple empirical objects without estimating detailed models.

3. Model Framework

We consider single-unit ascending button auctions with risk-neutral bidders and a risk-neutral seller, and we assume the auctions analyzed take place without a reserve price. We assume bidders have symmetric, conditionally independent, private values (CIPV). The equilibrium concept we consider throughout the main body of the article is equilibrium in weakly dominant strategies. Specifically, given that in a private values ascending button auction it is weakly dominant for bidders to bid their values, we assume that bidders do so.

Let N be a random variable denoting the number of auction participants and let n represent realizations of N. Let V_i denote bidder i's value and B_i his bid. For the subset of auctions that have exactly n bidders enter, let F^n denote the joint distribution of $V \equiv (V_i)_{i=1,\dots,n}$. Let f^n denote the joint density. By *bidder symmetry*, we refer to the case where F^n is exchangeable with respect to bidder indices.

The term *conditionally independent* is used to mean that bidders' values may be correlated in any given auction but that there exists some random variable U (unknown to bidders) such that bidders' values are independent conditional on U (Li et al. 2000). Our main results also apply to settings of unobserved auction-level heterogeneity as well as some other settings of correlated private values. Unobserved auction-level heterogeneity refers to settings where the random variable U is observed by the bidders but not by the econometrician and the realization of U in a given auction affects all bidders' valuations symmetrically. In the auctions literature, unobserved heterogeneity is frequently modeled as shifting valuations in an additively or multiplicatively separable fashion. All the results we derive in this article allow for *nonseparable* unobserved auction-level heterogeneity, assuming that the remaining, non-common component of bidder valuations is independent conditional on the realization of the unobserved auction-level heterogeneity.

Let $V^{1:n}, \ldots, V^{n:n}$ represent the bidders' valuations ordered from smallest to largest. Similarly, let the random variables $B^{1:n}, \ldots, B^{n:n}$ represent their bids ordered from smallest to largest. We will refer to $B^{n-1:n}$ and $B^{n-2:n}$ as the *second* and *third*-highest bids (and thus the phrase *highest bid* will refer to the drop out price of the highest-value bidder, which will not be observed in ascending auctions). We assume the researcher observes realizations of $B^{n-1:n}$, $B^{n-2:n}$, and N from a sample of i.i.d. auctions.[2] This data requirement may not be satisfied in many ascending auction settings, such as cases in which some bidding activity is not recorded by the auctioneer.

For $k \leq m \leq n$, let $B^{k:m,n}$ represent the k^{th} smallest bid in m bidder auctions, where the m bidders are selected uniformly at random from the n bidders in auctions that had exactly n bidders enter. Some remarks on this quantity are in order. We stress that this is a counterfactual if $m < n$: we assume that it is common knowledge amongst the remaining m bidders that $n-m$ of the original n bidders have been dropped, and that they are competing in an m-bidder auction, not an n bidder auction. The distribution of $B^{k:m,m}$ and $B^{k:m,n}$ for $m < n$ may be different, as different kinds of goods may attract different numbers of entrants, and bidders may value goods sold in auctions with m entrants differently from those sold in auctions with n entrants. Finally, $B^{k:m,m}$ and $B^{k:m}$ are the same random variable.

Our empirical strategy centers around three key variables. The first is the bidder exclusion effect. We define the bidder exclusion effect in n bidder auctions with no reserve price, $\Delta(n)$, as the expected fall in revenue produced by randomly excluding a bidder from those auctions. In ascending auctions, the bidder exclusion effect is:

$$\Delta(n) \equiv E(B^{n-1:n}) - E(B^{n-2:n-1,n}),$$

[2] This i.i.d. assumption refers to the relationship between auction observations in the data, not the relationship among valuations within a given auction. This assumption means that our framework does not apply to settings such as Backus and Lewis (2012), Hendricks and Sorensen (2015), and Coey et al. (2016) where there exist dynamic linkages across auctions affecting bidders' willingness to pay in a given auction.

that is, the expected second-highest bid in n bidder auctions, minus the expected second-highest bid in $n-1$ bidder auctions, where those $n-1$ bidder auctions are obtained by publicly dropping a bidder at random from n bidder auctions.

The second variable, $\Delta^{bid}(n)$, is the expected fall in revenue from dropping a *bid* (rather than a bidder) at random, assuming all other bids remain unchanged:

$$\Delta^{bid}(n) \equiv \frac{2}{n} E(B^{n-1:n} - B^{n-2:n}).$$

With probability $\frac{2}{n}$ one of the highest two bids will be dropped, and revenue will drop to the third-highest bid of the original sample, and with probability $\frac{n-2}{n}$, one of the lowest $n-2$ bids will be dropped, and revenue will not change. An advantage of focusing on a private values setting in which bidders bid their values is that $\Delta(n) = \Delta^{bid}(n)$. This is not necessarily the case in other environments discussed in the extensions in Section 7, such as a common values auction.

The third variable of interest, $\Delta^{obs}(n)$, is the *observed* difference in expected revenue between those auctions in which n bidders choose to enter, and those in which $n-1$ choose to enter:[3]

$$\Delta^{obs}(n) \equiv E(B^{n-1:n}) - E(B^{n-2:n-1}).$$

Unlike $\Delta(n)$, the quantities $\Delta^{bid}(n)$ and $\Delta^{obs}(n)$ are not counterfactual and can always be estimated as sample means using data on the two highest bids and the number of auction entrants. If the researcher observes a vector of auction-level characteristics, X, the researcher can estimate these objects conditional on X, estimating the sample analog of

$$\Delta^{bid}(n|X) \equiv \frac{2}{n} E(B^{n-1:n} - B^{n-2:n}|X) \tag{1}$$

$$\Delta^{obs}(n|X) \equiv E(B^{n-1:n}|X) - E(B^{n-2:n-1}|X). \tag{2}$$

Each of these objects can be computed using standard parametric or nonparametric approaches for estimating conditional means. Below, we will describe how these objects can be used for certain testing and counterfactual exercises.

[3]We use the term "observed" simply to help distinguish $\Delta^{bid}(n)$ from $\Delta^{obs}(n)$; the term is not meant to imply that $\Delta^{obs}(n)$ does not need to be estimated.

4. Testing If Valuations Are Independent of N

In this section we describe our test for the independence of valuations and the number of bidders. The intuition behind this test is as follows. If bidders' valuations do not vary systematically with the number of auction participants, then $n-1$ bidder auctions are just like n bidder auctions with one bidder removed at random. If the estimated bidder exclusion effect (the effect of randomly removing a bidder) is significantly different from the *observed* change in revenue between n and $n-1$ bidder auctions, this is evidence against independence of valuations and the number of entrants.

Let F_m^n denote the distribution of values of a random subset of m bidders, in auctions that actually had n participating bidders, where the $m \leq n$ bidders are drawn uniformly at random from the n bidders. Following Aradillas-López et al. (2013), we say that *valuations are independent of* N if $F_m^n = F_m^{n'}$ for any $m \leq n, n'$. Thus, if valuations are independent of N, we have $F_{n-1}^n = F_{n-1}^{n-1}$, and it follows that $E(B^{n-2:n-1,n}) = E(B^{n-2:n-1,n-1}) = E(B^{n-2:n-1})$, and

$$\begin{aligned}\Delta(n) &\equiv E(B^{n-1:n}) - E(B^{n-2:n-1,n}) \\ &= E(B^{n-1:n}) - E(B^{n-2:n-1}) \\ &\equiv \Delta^{obs}(n).\end{aligned}$$

As described in Section 3, in our environment, $\Delta^{bid}(n) = \Delta(n)$. Thus, testing the assumption that valuations are independent of the number of bidders involves comparing $\Delta^{bid}(n)$ and $\Delta^{obs}(n)$. We define $T(n)$ as

$$\begin{aligned}T(n) &\equiv \Delta^{obs}(n) - \Delta^{bid}(n) \\ &= \big(E(B^{n-1:n}) - E(B^{n-2:n-1})\big) - \frac{2}{n}E(B^{n-1:n} - B^{n-2:n}) \\ &= E\left(\frac{n-2}{n}B^{n-1:n} + \frac{2}{n}B^{n-2:n}\right) - E(B^{n-2:n-1}). \end{aligned} \tag{3}$$

The first term in the final expression is the expected revenue in n bidder auctions when one bidder is dropped at random, and the second term is the expected revenue in $n-1$ bidder auctions. The relationship in (3) is related to other recurrence relationships for order statistics (David and Nagaraja 1970). Athey and Haile (2002) propose using this same relationship between order statistics across samples of varying n as a test of private vs. common values. There, the authors maintain the assumption that valuations (and signals, in the common value case) are independent of N. The authors demonstrate that $T(n) < 0$ in a common values setting and $T(n) = 0$ in a private values setting; their test does not provide any interpretation for $T(n) > 0$ cases. In contrast, we maintain the assumption of private values and exploit this order statistic relationship to test the null hypothesis that valuations are independent of N. In our case, a statistically significant finding of $T(n) < 0$ can be interpreted as evidence that valuations are *higher*, in a stochastic dominance sense, in n bidder auctions than in $n-1$ bidder auctions. Similarly, a finding of $T(n) < 0$ can be interpreted as evidence that valuations are *lower*, in a stochastic dominance sense, in n bidder auctions than in $n-1$ bidder auctions. Throughout, we think of our test as a two-sided test, but in some cases the researcher may wish to implement the test as a one-sided test when the researcher is concerned about a particular sign of the dependence between valuations and the number of bidders.

As our approach requires observation of two order statistics of bids, it applies only to auctions with $n > 2$ bidders. Thus, the approach should not be used if the researcher is particularly concerned about a dependence between valuations and the number of bidders in auctions with only one or two bidders and has reason to believe that this dependence would not be detectable in data from auctions with $n > 2$. Auction-level unobserved heterogeneity may differ in one- or two-bidder auctions from $n > 2$ bidder auctions, for example. This does not pose a problem, however, as long as the test is interpreted correctly as providing information only directly for auctions with $n > 2$ bidders.

We test the null hypothesis $T(n) = 0$ using a two-sample t-test.[4] Let A_n represent the set of auctions with n entrants and let $b_j^{k:n}$ represent the $(k:n)$ order statistic of bids in auction j. The test statistic, $\widehat{T}(n)$, for this null is the sample analog of equation (3),

$$\widehat{T}(n) = \frac{1}{|A_n|}\sum_{j\in A_n}\left(\frac{n-2}{n}b_j^{n-1:n} + \frac{2}{n}b_j^{n-2:n}\right) - \frac{1}{|A_{n-1}|}\sum_{j\in A_{n-1}}(b_j^{n-2:n-1}). \tag{4}$$

[4]Standard techniques, like a Wald test or a Bonferroni correction, can be used to test $T(n) = 0$ for all n in some finite set. Note also that this test only uses information on the second and third-highest bids. If more losing bids are available and interpretable as the willingness-to-pay of lower-value bidders, this test could be made more powerful by including information from these losing bids. Intuitively, one could compare the revenue drop that would occur if k out of $n+k$ bidders were dropped at random to the actual revenue difference between n and $n+k$ bidder auctions. We address this idea in Section 7.

To perform the two-sample t-test, the variance of the first object in parentheses can be computed separately from the variance of the second object in parentheses, as the first comes from n bidder auctions and the second from $n-1$ bidder auctions. A simple regression-based form of this test is as follows. Let $y_j = \frac{n-2}{n} b_j^{n-1:n} + \frac{2}{n} b_j^{n-2:n}$ if $j \in A_n$ and $y_j = b_j^{n-2:n-1}$ if $j \in A_{n-1}$. Regress y_j on a constant and an indicator $\mathbb{1}(j \in A_n)$. The coefficient on the indicator is $\widehat{T}(n)$. Heteroskedasticity-robust standard errors for this regression would allow for the variance to differ in the n bidder and $n-1$ bidder auctions, as in the two-sample t-test.

If $\widehat{T}(n)$ is significantly different from 0, the test indicates the presence of dependence between valuations and N. This test is consistent against all forms of dependence that affect expected revenue (that is, if $\Delta^{bid}(n) \neq \Delta^{obs}(n)$ then the test rejects with probability approaching 1 as the number of auctions goes to infinity). Clearly, given that this test focuses only on expectations, it would not detect all types of dependence between valuations and the number of bidders: it is a test of a necessary condition of such independence, not a sufficient condition. This proposed test could be extended, and made more powerful, by comparing the entire distribution of $\frac{n-2}{n} B^{n-1:n} + \frac{2}{n} B^{n-2:n}$ to that of $B^{n-2:n-1}$, rather than only the means, but at a sacrifice of computational simplicity.

Appendix B shows Monte Carlo evidence on the power of this test relative to simply comparing mean values in $n-1$ and n bidder auctions in a model that allows for dependence between valuations and N (this model nests the entry model of Levin and Smith 1994). We find that the bidder exclusion test is a reasonably powerful alternative to this mean comparison test given that it uses considerably less data. Moreover, the bidder exclusion test is implementable with ascending auction data whereas the mean comparison test is not.[5]

Testing is also possible if valuations are assumed independent of N conditional on a set of observable auction characteristics X rather than unconditionally. The null hypothesis is $T(n|X) = 0$, where $T(n|X)$ is defined as

$$T(n|X) \equiv E\left(\frac{n-2}{n} B^{n-1:n} + \frac{2}{n} B^{n-2:n} \middle| X \right) - E(B^{n-2:n-1}|X). \tag{5}$$

[5]The bidder exclusion test uses the second and third-highest values in n bidder auctions and the second-highest value in $n-1$ bidder auctions, whereas the mean comparison test uses *all* n values in n bidder auctions and *all* $n-1$ values in $n-1$ bidder auctions. The mean comparison test cannot be implemented in ascending auctions because the highest valuation is never observed.

This hypothesis can be tested nonparametrically, without assuming any particular form for the conditional means. Chetverikov (2018), Andrews and Shi (2013), and Chernozhukov et al. (2013) develop inference procedures that apply to this setting.

A simple parametric version of this test is as follows. For a fixed n, specify the bidding equation for bidder i in auction j as

$$b_{nji} = \alpha_n + \beta X_j + \varepsilon_{nji}, \tag{6}$$

where X_j is a vector of observable characteristics of auction j and $\varepsilon_{nji} \perp\!\!\!\perp X_j$.[6] Then

$$\frac{n-2}{n} b_j^{n-1:n} + \frac{2}{n} b_j^{n-2:n} = a_1 + \beta X_j + e_{1j}, \quad j \in A_n \tag{7}$$

$$b_j^{n-2:n-1} = a_2 + \beta X_j + e_{2j}, \quad j \in A_{n-1} \tag{8}$$

where $a_1 = \alpha_n + E(\frac{n-2}{n}\varepsilon_{nj}^{n-1:n} + \frac{2}{n}\varepsilon_{nj}^{n-2:n})$, $a_2 = \alpha_{n-1} + E(\varepsilon_{(n-1)j}^{n-2:n-1})$, $E(e_{1j}) = E(e_{2j}) = 0$ and $e_{1j}, e_{2j} \perp\!\!\!\perp X_j$. After controlling for observables, a_1 determines the expected second order statistic (i.e. seller's revenue) when a bidder is removed at random from n bidder auctions and a_2 determines the expected second order statistic in $n-1$ bidder auctions when the actual number of bidders is indeed $n-1$. Testing the null hypothesis of equation (5) amounts to testing the null of $a_1 = a_2$.

We combine (7) and (8) as follows:

$$y_j = a_2 + (a_1 - a_2)\mathbb{1}(j \in A_n) + \beta X_j + e_{3j}, \quad j \in A_n \cup A_{n-1}, \tag{9}$$

where if $j \in A_n$, then $y_j = \frac{n-2}{n} b_j^{n-1:n} + \frac{2}{n} b_j^{n-2:n}$ and $e_{3j} = e_{1j}$, and if $j \in A_{n-1}$, then $y_j = b_j^{n-2:n-1}$ and $e_{3j} = e_{2j}$. This allows for a convenient regression-based test of the null hypothesis that valuations are independent of N. When $\beta = 0$, this test nests the no-covariates regression-based test described above.[7] We emphasize again that this parametric version is not the only way to implement our test; if desired, the test can also be implemented by allowing covariates X_j to enter in a more flexible parametric or nonparametric fashion.

[6] Athey et al. (2011) and Athey et al. (2013) also take the approach of specifying a parametric model directly for bids, rather than the underlying values. Given our button auction setting, specifying a model for bidding is equivalent to specifying a model for values.

[7] As described above, heteroskedasticity-robust standard errors in this regression test would allow for the variance of the unobserved term to differ in the n and $n-1$ bidder auction samples.

5. Using the Bidder Exclusion Effect to Inform Counterfactuals

In this section we demonstrate that the bidder exclusion effect can be used to compute bounds on several counterfactual objects of interest for auction studies. First, we provide a bound on the counterfactual improvement in revenue that a seller would receive by using an optimal reserve price. Second, we provide an approach for bounding the drop in revenue to a seller when bidders merge.

Bounding the Impact of Optimal Reserve Prices. The celebrated theorem of Bulow and Klemperer (1996) (Theorem 1) relates bidder entry to optimal auction design.[8] The authors demonstrate that an English auction with no reserve price and $n+1$ bidders is more profitable in expectation than *any* mechanism with n bidders.[9] On these grounds, they suggest that sellers may be better off trying to induce more entry than trying to implement a better mechanism. As they acknowledge, this interpretation may be problematic if the new bidders are weaker than the bidders who would have entered anyway (for example, if increased marketing efforts induce lower-value bidders to enter the auction). We propose an alternative interpretation of their theorem, namely that it can be used in empirical work to easily obtain upper bounds on the effect of improving auction design.

To apply the Bulow-Klemperer result, we are required to make one additional assumption imposed by Bulow and Klemperer (1996): that bidder valuations satisfy the *monotonicity of marginal revenue* property.[10] To define increasing marginal revenue, let $(V_1, \ldots, V_n)$ denote the private values of the n bidders in n bidder auctions, as above. The $n+1^{th}$ bidder, were he to enter, has a private value denoted V_{n+1}. We denote the marginal distribution of V_j by F_j^n, and the corresponding density by f_j^n. Define $\mathbf{V} \equiv (V_1, \ldots, V_{n+1})$, and $\overline{\mathbf{V}} = \mathbf{V}_{-(n+1)}$, that

[8]Following Bulow and Klemperer (1996) and most of the auction theory literature, we use "optimal" to mean optimal given a fixed set of participants. If entry is endogenous, then the mechanism's design may affect the number of participants. Optimal reserve prices for fixed and for endogenous entry may be different (McAfee and McMillan 1987; Levin and Smith 1994).

[9]The Bulow-Klemperer Theorem is frequently misunderstood as only applying to independent private values settings, when in fact it is stated for correlated private values settings as well as certain common values settings. We discuss this further in Section 7.

[10]Note that Bulow and Klemperer (1996) also assume risk neutrality, which we assume throughout, and that the seller's valuation for the good is less than that of all buyers. If this latter assumption does not hold, the Bulow-Klemperer result still applies, but the result would be modified to state that one additional bidder is better than an optimal reserve price at increasing the *seller's expected payment from bidders* rather than the seller's revenue (where the latter includes the seller's valuation of keeping the good and the former does not).

is, the values of bidders other than bidder $n+1$. Let $F_j^n(V_j|\mathbf{V}_{-j})$ and $f_j^n(V_j|\mathbf{V}_{-j})$ represent the distribution and density of bidder j's value conditional on competitors' values. Define $MR_j(\mathbf{V})$ and $\overline{MR}_j(\overline{\mathbf{V}})$ as

$$MR_j(\mathbf{V}) \equiv \frac{-1}{f_j^n(V_j|\mathbf{V}_{-j})} \frac{d}{dV_j} \left(\mathbf{V}(1 - F_j^n(V_j|\mathbf{V}_{-j}))\right) \tag{10}$$

$$\overline{MR}_j(\overline{\mathbf{V}}) \equiv \frac{-1}{f_j^n(V_j|\overline{\mathbf{V}}_{-j})} \frac{d}{dV_j} \left(\overline{\mathbf{V}}(1 - F_j^n(V_j|\overline{\mathbf{V}}_{-j}))\right) \tag{11}$$

We say bidders have "increasing marginal revenue" (as a function of their private values) if $V_j > V_i \Rightarrow MR_j(\mathbf{V}) > MR_i(\mathbf{V})$ and $\overline{MR}_j(\overline{\mathbf{V}}) > \overline{MR}_i(\overline{\mathbf{V}})$.[11] In the independent private values case this assumption simplifies to the function $MR(V_j) \equiv MR_j(\mathbf{V}) = \overline{MR}_j(\overline{\mathbf{V}}) = V_j - \frac{1-F_j^n(V_j)}{f_j^n(V_j)}$ being increasing in V_j.

We also note that the Bulow-Klemperer result is about *adding* a bidder, whereas the effect we can measure in the data is that of *removing* a bidder. We therefore prove the following result, which is a special case of results established by Dughmi et al. (2012) (Theorem 3.2). The specialization to our current single-item auction setting allows us to use only elementary mathematics, in contrast to Dughmi et al. (2012)'s proof, which relies on matroid theory. The proof is found in the Appendix.

Proposition 1. *In ascending button auctions with no reserve price where bidders have symmetric, conditionally independent, private values, if bidders' marginal revenue is increasing in their values then the absolute value of the change in expected revenue is smaller when adding a random bidder than when removing a random bidder.*

Our main optimal-revenue-bounding result then follows immediately:

Corollary 1. *In ascending button auctions with no reserve price where bidders have symmetric, conditionally independent, private values and increasing marginal revenue curves, then for all $n > 2$ the increase in expected revenue from using the optimal reserve price is less than $\Delta^{bid}(n)$.*

[11]Equivalently, bidders have decreasing marginal revenue, when marginal revenue is considered to be a function of bidder "quantity" (i.e. $(1 - F_j^n(V_j|\mathbf{V}_{-j}))$ and $(1 - F_j^n(V_j|\overline{\mathbf{V}}_{-j}))$) rather than of their values. Note that Bulow and Klemperer (1996) parameterize marginal revenue in terms of bidder "quantity" rather than their private values, so that their marginal revenue function is decreasing. For more on the interpretation of bidders' marginal revenue, see Bulow and Roberts (1989).

Proof. The result follows immediately from combining the Bulow-Klemperer Theorem with Proposition 1. □

We wish to emphasize that this result does not require independence between valuations and the number of bidders (the condition that can be tested using the results in Section 4), because it only relies on a given realization n of the number of bidders. We also note that the upper bound in Proposition 1 is not necessarily sharp. Although the Bulow and Klemperer (1996) bound is indeed sharp (i.e., there exists a limiting distribution of valuations for which adding a random bidder increases revenue by the same amount as an optimal reserve price would),[12] our bound will not necessarily be sharp due to potential slack in the bound given in Proposition 1, that is, due to the fact that we are considering the effect of removing a random bidder rather than that of adding a random bidder.

Bounding the revenue gains from optimal reserve prices can serve two key purposes. The first is to allow the practitioner to gauge whether or not to invest resources in determining/implementing an optimal reserve price. This can be useful in particular given the large losses that can result from charging too high a reserve price (see Kim 2013, Ostrovsky and Schwarz 2016, and Coey et al. 2018). The second is to allow researchers to compare quantitatively the effects of other interventions in auction environments to the benchmark of optimal reserve pricing. This concept is illustrated in Lacetera et al. (2016), where it is shown that the effect of a one-standard-deviation improvement in the ability of the human auctioneer at auto auctions raises revenue by $348 per auction, whereas the upper bound on the benefit of optimal auction design as measured through the bidder exclusion effect is $333. A similar comparison can be made to Tadelis and Zettelmeyer (2015), where the authors measure that information disclosure at similar auto auctions increases revenue by $643, Together, these results suggest that non-traditional instruments of auction design, such as information disclosure or high-performing auctioneers, can matter more than reserve prices for improving auction revenue—in spite of the primary focus in the existing literature on reserve pricing and similar instruments as the means of improving auction revenue.[13]

[12]Consider, for example, a symmetric IPV button auction and consider a sequence of distributions converging to $F(v) = v/(1+v)$. In the limit, the optimal revenue approaches the revenue with an additional bidder arbitrarily closely. We thank Jason Hartline for pointing out this fact.

[13]See Coey et al. (2014), the earlier working paper version of this study, for further discussion and results.

Bidder Mergers. The concept behind the bidder exclusion effect can also be used to bound above the expected fall in revenue resulting from a counterfactual bidder merger. We state our results in terms of a merger between two of the n bidders, but the results easily extend to more than two bidders merging. When bidders i and j merge, let M_k denote the willingness-to-pay of bidder $k \neq i, j$, and $M_{i,j}$ denote the willingness-to-pay of the joint entity. Our results in this subsection rely on the following assumption:

Assumption 1. *When any two bidders i and j merge, $M_k = V_k$ for all $k \neq i, j$, and $M_{i,j} \geq \max\{V_i, V_j\}$.*

Assumption 1 implies that mergers may result in increased production efficiencies and hence an *increased* willingness to pay of the merged entity, but the merger will not *decrease* this entity's willingness-to-pay.[14] Further intuition for this assumption comes from considering a procurement (reverse auction) setting rather than an increasing-price auction, where the analogous assumption would be that costs of merging bidders do not *rise* after the merger, but may *decrease* (due to economies of scale or other cost efficiencies). This assumption—or the even stronger assumption that Assumption 1 holds with equality—is satisfied in all previous auction merger studies of which we are aware (see Froeb et al. 1998, Waehrer 1999, Dalkir et al. 2000, Brannman and Froeb 2000, Tschantz et al. 2000, Waehrer and Perry 2003, Froeb et al. 2008, and Li and Zhang 2015, for example), but it certainly may not hold in all merger settings (for example, if the merger changes the willingness to participate of the non-merging entities), and its appropriateness should be evaluated on a case-by-case basis.[15]

It is important to note that Assumption 1 does not imply that the joint entity will necessarily end up *paying* more than $\max\{V_i, V_j\}$; this will only occur in cases where the joint

[14]The intuition derived in this section to analyze counterfactual bidder mergers can also be extended to analyzing the effects of counterfactual *collusion* among bidders. In the case of collusion, Assumption 1 is about two bidders colluding, rather than merger. Assumption 1 is satisfied in all models of efficient collusion (e.g. Mailath and Zemsky 1991, Krishna 2009, and Marmer et al. 2016). Waehrer and Perry (2003) explain that the efficient-collusion setting is equivalent to a bidder merger with no cost synergies. In some models of *inefficient* collusion the assumption is also satisfied (e.g. Graham et al. 1990 and Asker 2010), although in some it is not (e.g. von Ungern-Sternberg 1988 and Pesendorfer 2000, in which sidepayments are not allowed).

[15]There are certainly other plausible alternatives to Assumption 1. For example, rather than adopting the highest valuation of the merging bidders, the merged entity might instead adopt the higher of the two values with some probability α and the lower of the two values with probability $1-\alpha$. Here we maintain Assumption 1 throughout.

entity wins post-merger but would not have won (i.e. neither of two bidders merging would have won) in the absence of the merger. The actual price paid by the joint entity, when it wins, will be determined by the second-highest willingness to pay.

Under Assumption 1, we obtain the following result:

Proposition 2. *In ascending button auctions with no reserve price where bidders have symmetric, conditionally independent, private values and where Assumption 1 holds, then, for all $n > 2$, the decrease in expected revenue from two bidders merging is bounded above (i) by $\frac{1}{n-1}\Delta^{bid}(n)$ when the two bidders are randomly selected and (ii) by $E(B^{n-1:n} - B^{n-2:n})$ when the two bidders are not randomly selected.*

Proposition 2 distinguishes between the case where the merging bidders are randomly selected vs. non-randomly selected. The term *randomly selected* here means that each bidder in the merger does not know the other bidder's realized draw from the valuation distribution prior to deciding to merge. Randomly selecting the merging pair is equivalent to randomly selecting a pair of bidders, dropping the bidder with the lower value in the pair (and hence the highest-value bidder would never be dropped in this process), and then weakly raising the willingness-to-pay of the remaining bidders. As shown in the proof of Proposition 1, this process will only lead to a decrease in revenue when the lower-value bidder in the selected pair is the bidder whose value corresponds to $V^{n-1:n}$ (or, equivalently, the merger only leads to a revenue decrease when the pair contains the two highest-value bidders; a similar point is discussed in Froeb et al. 1998), which occurs with probability $\frac{1}{\binom{n}{2}} = \frac{2}{n(n-1)}$.[16] Thus, the upper bound in part (i) of Proposition 2 is smaller than $\Delta^{bid}(n)$, by a factor of $\frac{1}{n-1}$. The upper bounds in Proposition 2 are sharp; the bounds in (i) and (ii) can hold with equality when the inequalities in Assumption 1 bind.

In practice, it may be the case that bidders are not randomly matched to merge, and in these cases the wider bound in part (ii) can be useful, bounding the seller's loss using the unscaled expected gap between the second and third order statistics of bids. For example, in many settings it may be two high-value competitors who choose to merge. The tighter bound in (i) will only be an upper bound on the seller's loss if the merged entity is more likely to

[16] The proof also demonstrates that an even stronger result than that in Proposition 2 holds: the bounds hold in any given auction as well, not only in expectation. However, we state the proposition in terms of expectations for consistency of the exposition.

include low-value bidders than would a randomly formed merged entity. This might be the case, for example, if the merger is motivated by a bidder acquiring a smaller competitor.

These approaches to analyzing mergers in auctions relate to recent discussion in competition policy. Prior to the 2010 revision of the US Merger Guidelines, the Federal Trade Commission (FTC) and Department of Justice (DOJ) published a variety of questions for comment, one of which addressed how exactly unilateral effects should be evaluated in markets with auctions or negotiations (FTC 2009). In response to these questions, Moresi (2009) hints at a similar idea to what we propose: one can examine the second and third-lowest bids in a procurement auction (analogous to the second and third-highest in our ascending auction setting) to understand the pricing pressure created by bidders choosing to merge. These ideas also relate to FTC and DOJ analysis of mergers using auctions to model a market even when that market is not explicitly centrally run as an auction. Baker (1997) and Froeb et al. (1998) highlight examples of such uses by antitrust and competition authorities to quantify price effects in mergers of pharmacies, hospitals, mining equipment companies, and defense contractors; Dalkir et al. (2000) provides an example in health insurance markets.

6. Application: US Timber Auctions

Our empirical application uses US government timber auction data to illustrate how the bidder exclusion effect can be used to test for dependence between valuations and the number of bidders. The Forest Service's timber auction data has been used extensively in the empirical auctions literature, and is a natural context to demonstrate the applications of the bidder exclusion effect. After a brief description of the data, we address the question of whether the seller (in this case, the government), should bother investing the effort required to compute an optimal reserve price. We do so by using the bidder exclusion effect to bound the gains to the seller from optimal reserve pricing. We then ask the question, if the seller does wish to compute an optimal reserve price—or bounds on the optimal reserve price, using approaches such as Haile and Tamer (2003), Aradillas-López et al. (2013), or Coey et al. (2017)—can she, in doing so, safely assume that variation in the number of bidders is exogenous to bidder valuations (in order to obtain tighter, more meaningful bounds on the reserve price).

Data Description. Our data is the same as that used by Athey et al. (2013) and comes from ascending auctions held in California between 1982 and 1989 in which there were at least three entrants. There are 1,086 such auctions. These auctions had reserve prices, but they were low, only binding in 1.1 percent of cases, and we therefore treat the auctions as though they did not have reserve prices. For each auction, the data contains all bids, the number of bidders, and information on the bidders' identities, as well as auction-level information. These auction-level characteristics include appraisal variables (quintiles of the reserve price, selling value, manufacturing costs, logging costs, road construction costs, and dummies for missing road costs and missing appraisals), sale characteristics (species Herfindahl index, density of timber, salvage sale or scale sale dummies, deciles of timber volume, and dummies for forest, year, and primary species), and local industry activity (number of logging companies in the county, sawmills in the county, small firms active in the forest-district in the last year, and big firms active in the forest-district in the last year).

Bounding Counterfactual Revenue Changes. We first address the following question: should the seller—in this case, the government—at these timber auctions bother to compute an optimal reserve price? Computing an optimal reserve price requires effort in practice, and recent work has highlighted the extreme asymmetric payoff to sellers from mistakes in choosing reserve prices: setting too high a reserve price can lead to losses much larger in magnitude than the losses from setting too low a reserve price (see Kim 2013, Ostrovsky and Schwarz 2016, and Coey et al. 2018). Thus, the seller may find it useful to first gauge whether the investment in optimal auction design would be worthwhile, or whether she should instead simply run a no-reserve auction (which requires no pricing decision on the part of the seller). The bidder exclusion effect serves as a useful diagnostic for such an initial assessment.

We begin by estimating the bidder exclusion effect ($\frac{2}{n}E(B^{n-1:n} - B^{n-2:n})$) at various values of n. Figure 1 shows this quantity, both as a percentage of revenue, and in absolute terms.[17] Note that Figure 1 is simply an illustration of summary statistics of the data; it is only

[17]The larger confidence interval for the $n = 7$ auctions in the right panel of Figure 1 is driven by an outlier (an outlier in terms of its realization of $B^{n-1:n} - B^{n-2:n}$ but not as a percentage of revenue, $(B^{n-1:n} - B^{n-2:n})/B^{n-1:n}$, and thus in the left panel the effect is still precisely measured).

through the lens of our framework that these summary statistics have an interpretation as being informative about counterfactual revenue changes.[18]

The range of the bidder exclusion effect in the left panel of Figure 1 ranges from about 20% to less than 5%, decreasing sharply as n increases. Averaging over all values of n, the average bidder exclusion effect is 12% of auction revenue, with a standard error of 0.4%. Under the conditions of Proposition 1, the average increase in revenue from setting an optimal reserve price is therefore less than about 13% of revenue. We see this bound as large enough that a seller at these auctions would likely find it worthwhile to invest in determining an optimal reserve price. These estimated bounds on the revenue impact of optimal reserve prices are consistent with those found in previous structural timber auctions studies, but the bounds we compute require only a small fraction of the computational cost of these previous approaches and as such can serve as a useful initial diagnostic in practice.[19]

Testing Independence of Valuations and N. We now turn to the following question: If the seller does wish to compute the optimal reserve price itself—or bounds on the optimal reserve price—can she safely rely on exogenous variation in the number of bidders in doing so? Two prominent methods for obtaining bounds on the optimal reserve price in ascending auction settings—Haile and Tamer (2003) and Aradillas-López et al. (2013), as well as the method for asymmetric environments in Coey et al. (2017), all of which study timber auctions—rely on the assumption that valuations are independent of N in order to obtain tight, meaningful bounds on revenue and on the optimal reserve price itself, as we discuss in more detail below. Similarly, Brannman and Froeb (2000) study merger effects at timber auctions and build a model that relies on the assumption that valuations are independent of N. To address whether this assumption is reliable, we apply our test developed in Section 4.

[18]As discussed in Section 5, a quantity related to the bidder exclusion effect can also provide a bound on the fall in revenue from two bidders merging, given by $\frac{2}{n(n-1)}E(B^{n-1:n} - B^{n-2:n})$. In our data, we find this quantity to be 4% of revenue. A bound on the loss in seller revenue when instead two non-random bidders merge can be recovered simply by scaling the values in Figure 1 by a factor of $n/2$, yielding the expected gap between the second and third order statistics. Bidder mergers at timber auctions are studied in Brannman and Froeb 2000; Athey et al. 2011; Li and Zhang 2015.

[19]Optimal reserve prices and the gain from implementing the optimal reserve are examined in other timber auction settings in Paarsch 1997; Li and Perrigne 2003; Haile and Tamer 2003; Roberts and Sweeting 2014, 2013; Aradillas-López et al. 2013; Coey et al. 2017.

We begin with the simplest version of the test, without controlling for covariates. Table 1 displays the results. In the table, a_1 represents the expected second order statistic when a bidder is removed at random from n bidder auctions, a_2 represents the expected second order statistic in $n-1$ bidder auctions when the actual number of bidders is indeed $n-1$, and the test statistic is given by $\widehat{T}(n) = a_1 - a_2$. For most $n \in \{3, ..., 8\}$, $\widehat{T}(n)$ is insignificant, although at $n = 3$ and $n = 5$, the test statistic is significant and positive, indicating that dependence between valuations and the number of bidders may be a concern. Intuitively, a positive $T(n)$ indicates that bidders' values are higher in n than $n-1$ bidder auctions, as might be the case when goods that are more attractive (in a way that is unobservable to the econometrician) tend to draw many bidders.

Table 2 shows the results of the test conditional on auction characteristics. The objects a_1, a_2, and $T(n)$ are as in Table 1, but after controlling for the auction-level covariates listed above (the same controls as in Athey et al. 2013), following the parametric procedure described in Section 4).[20] Table 2 shows that there is stronger evidence for dependence of valuations and the number of bidders when controlling for auction characteristics than in the unconditional case. Conditional on auction characteristics, average revenue when a bidder is removed at random from n bidder auctions is higher than average revenue in $n-1$ bidder auctions when $n \in \{3, 4, 5, 7\}$, and this difference is significant at the 95% level. Again, one explanation for this would be positive selection: bidders' valuations appear to be higher in auctions with more participants. With $n \in \{6, 8\}$ the difference is negative and insignificant, consistent with a setting where valuations are independent of N. The joint null hypothesis of independence of valuations and the number of bidders across all $n \in \{3, \ldots, 8\}$ can be rejected at the 99.9% level.

Some bidders in timber auctions may be stronger than others. One common distinction in the literature is between mills, who have the capacity to process the timber, and loggers, who do not. Mills typically have higher valuations than loggers (e.g. Athey et al. 2011, 2013; Roberts and Sweeting 2014; Coey et al. 2017). The evidence of dependence in valuations and N above may be driven by differences in logger and mill entry patterns. We next turn

[20]Note that, due to the small sample size, this analysis only controls linearly for a number of different types of sales (as in Athey et al. 2011 and Athey et al. 2013). If the sample size were larger, it would be possible to evaluate the independence of bidder valuations and the number of bidders separately in subsamples of these different types of sales rather than grouping them all together.

to the question of whether evidence of dependence exists, even after restricting attention to a more homogeneous subset of bidders—in this case, loggers.[21]

Table 3 shows the results. The sample size is significantly smaller when restricting to auctions in which all entrants are loggers, which reduces the power of the test.[22] At $n = 4$, the test still rejects the null hypothesis that valuations are independent of N. However, the evidence on the whole is much weaker in the loggers-only sample: at $n \in \{3, 5, 6\}$ the difference is much smaller and insignificant, although these results should be interpretted with caution, as the smaller sample size may play a role. The joint null hypothesis of independence of valuations and the number of bidders across all n can no longer be rejected as it was before accounting for asymmetries. Together the results from Tables 2 and 3 suggest that dependence between valuations and the number of bidders may be a feature of timber auctions but is less pronounced after accounting for bidder types.

Implications of Test. Rejecting or failing to reject the independence of valuations and the number of bidders has a number of implications for analyzing auction data. Haile and Tamer (2003) and Aradillas-López et al. (2013), two of the most influential methodological innovations for ascending auction settings, both rely on the assumption that valuations are independent of the number of bidders in order to obtain meaningful bounds on seller profit and on optimal reserve prices.[23] In the setting of Haile and Tamer (2003), if the assumption of independence of valuations and the number of bidders is violated (as appears to be so in several of the cases examined in our data), the Haile and Tamer (2003) approach can still yield bounds on seller profit and on the optimal reserve price in symmetric IPV auctions, but these bounds will be looser than if the researcher were able to confidently exploit exogenous variation in the number of bidders.

[21]There are too few auctions without logger entrants (only 21) to present the same analysis for mills.

[22]One might be concerned that the restriction to logger-only auctions may be selecting on a potentially endogenous variable. The framework and results of Athey et al. (2013) suggest that this may not be a concern; the authors find that the estimated (asymmetric) value distributions are such that if loggers enter with positive probability, then all potential mills would enter. Thus, restricting attention to auctions with no mills entering is also restricting attention to auctions in which there were no potential mills available to enter. Within the Athey et al. (2013) framework, and in many other empirical auction settings, this number of potential entrants is considered to be exogenous (unlike the number of actual entrants).

[23]In addition to those methodologies discussed herein, other approaches found in Sections 5.3 and 5.4 of Athey and Haile (2007) also rely on the assumption that valuations are independent of the number of bidders.

In the symmetric correlated private values setting of Aradillas-López et al. (2013) (or in the related asymmetric correlated private values setting of Coey et al. 2017), the authors demonstrate that their two-sided bound on seller profits is only available when valuations are independent of the number of bidders. If valuations and the bidder count are positively dependent, a one-sided bound on seller revenue still holds, but no meaningful bound on the optimal reserve price is available. If valuations and the number of bidders are negatively dependent, bounds on seller revenue and optimal reserve prices will be unavailable or will be uninformative. The diagnostic test we propose herein can help the researcher determine whether or not she can exploit exogenous variation in the number of bidders to obtain meaningful bounds in these settings.

If the test rejects the independence of bidder valuations and the number of bidders, the researcher can incorporate this knowledge by placing more structure on the model, either by explicitly modeling a dependence between unobserved heterogeneity and the number of bidders, or by incorporating a model of bidders' entry decisions. Modeling entry into auctions has been a major innovation in recent work, including a number of studies of timber auctions (such as Athey et al. 2013, Roberts and Sweeting 2013, Aradillas-López et al. 2013, and Roberts and Sweeting 2014). The bidder exclusion effect test is one tool, among others, that can help the researcher in determining whether the additional modeling complexity required to explicitly account for entry is warranted.

7. Discussion of Extensions

In this section we discuss a number of different extensions of the uses of the bidder exclusion effect to more general environments.

Common Values. In the private values setting in the main body of the article, the change in auction revenue when one bidder is excluded can be computed by removing one bidder's bid, and calculating the fall in revenue assuming the other bids remain unchanged. This is not true with common values, as removing a bidder changes the remaining bidders' equilibrium bidding strategies. Theorem 9 of Athey and Haile (2002) can be used to show that in ascending button auctions with symmetric common values and symmetric bidding strategies, in any separating equilibrium, $\Delta(n) < \Delta^{bid}(n)$.[24] Removing one bidder's bid and assuming other bids remain unchanged ($\Delta^{bid}(n)$) thus overstates the decline in revenue from actually excluding a random bidder, because it does not account for the increase in bids due to the reduced winner's curse.

The test proposed in Section 4 can still be applied in the common values setting, but will only indicate dependence between valuations and N if $\widehat{T}(n)$ is significantly greater than zero, and not if it is significantly less than zero.[25] This test is consistent against forms of dependence in which bidders' values are *sufficiently increasing* with N. As highlighted in Aradillas-López et al. (2013, 2016), this type of positive dependence is particularly the kind that can occur in many popular entry models (see Aradillas-López et al. 2016 Appendix B.5). This one-sidedness feature of our test in this case is shared by the test in private values settings proposed in Aradillas-López et al. (2016), which the authors explain "has power against a fairly wide class of 'typical' violations of [valuations being independent of N]."

The application of the bidder exclusion effect to bounding counterfactual revenue under an optimal reserve price applies in common values settings as well. In particular, when bidders have affiliated signals, Bulow and Klemperer (1996) show that an auction with $n+1$ bidders and no reserve price still outperforms any "standard" mechanism with n bidders.[26] To apply the bidder exclusion effect to bound counterfactual revenue in common value auctions, however, we are required to make one additional assumption: that the magnitude of the

[24]Note that in discussing common values in this subsection, we depart from the main body of the article in our equilibrium concept and focus here on Bayes Nash Equilibria. Common value ascending auctions may have many equilibria; however, the results of Bikhchandani et al. (2002) demonstrate that this inequality is true in *any* symmetric, separating equilibrium of an ascending auction. Importantly, it will not necessarily hold in equilbria in asymmetric strategies.

[25]More precisely, in the case of common values, the exercise would be to test whether valuations *and signals* are independent of N, as valuations and signals are not synonymous in a common values setting.

[26]A "standard" mechanism in this context is one in which 1) losers pay nothing, 2) the bidder with the highest signal wins (if anyone) and pays an amount that increases in his own signal given any realization of other bidders' signals. Bulow and Klemperer (1996) highlight a result of Lopomo's (1995), which shows that an optimal mechanism in this class is an English auction followed by a final, take-it-or-leave-it offer to the high bidder (a reserve price). When bidders have correlated values, Crémer and McLean (1988), McAfee et al. (1989), and McAfee and Reny (1992) have provided examples of non-standard mechanisms that extract all bidder surplus and outperform an auction with a reserve price. Also, in later work, Bulow and Klemperer (2002) highlighted that the assumption of marginal revenues increasing in signals may be more stringent in common values settings than in private values settings, and the authors provided examples of common values settings in which the original Bulow and Klemperer (1996) result will not hold when the condition of increasing marginal revenues is violated.

change in revenue from adding a random bidder is smaller than the magnitude of the change from removing a bidder. With conditionally independent private values we prove this as a result (Proposition 1), but with common values it must be imposed directly as an assumption, making the counterfactual revenue bound in common values setting less appealing. However, what is appealing about this revenue bound in common value ascending auctions is that it yields a bound on revenue in a setting where no known identification (or even partial identification) results exist in the auction methodology literature for the optimal reserve price itself. We do not extend our bidder merger analysis to the common values case.

Non-button Ascending Auctions. In the button auction model of ascending auctions with private values (Milgrom and Weber 1982), bidders drop out at their values. As highlighted in Haile and Tamer (2003), in practice bidders' (highest) bids may not equal their values. For example, in English auctions, multiple bidders may attempt to bid at a certain price but only the first bidder the auctioneer sees may be recorded. Jump bidding or minimum bid increments may also lead to cases where bids do not equal values. Additionally, in some cases, a bidder may drop out at a low bid level, planning to participate again but never doing so as the bidding rises past her value.

A simple extension of our model that allows for some non-button behavior is to assume that the final auction price still represents the second-highest bidder's valuation, but that any lower bidder's bids are weakly lower than their valuations.[27] Under this low-bidding assumption, it is straightforward to demonstrate that $\Delta(n) \leq \Delta^{bid}(n)$. As with the common values case, where this inequality was instead strict, the test of exogeneity of N still applies but can only reject certain forms of exogeneity. The use of the bidder exclusion effect to bound counterfactual revenue under an optimal reserve price or under bidder mergers immediately extends to this low-bidding case. In Appendix B we demonstrate how the bidder exclusion effect can be computed in an environment that more closely (although still not perfectly) resembles that of Haile and Tamer (2003).[28]

[27]Athey and Haile (2002) argue, "...for many ascending auctions, a plausible alternative hypothesis is that bids $B^{n-2:n}$ and below do not always reflect the full willingness to pay of losing bidders, although $B^{n-1:n}$ does (since only two bidders are active when that bid is placed)."

[28]There, our environment still rules out some kinds of jump bidding that would be allowed within the more general Haile and Tamer (2003) framework. We note, however, that our environment throughout the article is more general than that of Haile and Tamer (2003) in another dimension, in that it allows for correlation between bidders' values.

Unobserved Number of Bidders. In some settings the number of bidders may not be known to the researcher. In ascending auctions, for example, not all potential bidders may place bids. A lower bound on the number of potential bidders may be known, however, such as in cases where the researcher only observes bids of bidders whose valuations exceed a reserve price. Let $\underline{n}$ represent this lower bound, such that for all realizations n of the random variable N, $n > \underline{n}$. In this case an upper bound on the average bidder exclusion effect, $E(\Delta(N))$, is given by averaging over realized order statistics from samples of unknown N, yielding $E(\Delta(N)) \leq \frac{2}{\underline{n}} E(B^{N-1:N} - B^{N-2:N})$. This quantity can then be used to compute an expected upper bound on the seller revenue changes from using an optimal reserve price or from mergers. As with the main presentation of the bidder exclusion effect in Section 3, this upper bound on the average bidder exclusion effect can also be estimated conditional on auction-level unobservables. In addition, the researcher may wish to estimate the lower bound $\underline{n}$ conditional on auction-level unobservables to obtain a better overall upper bound on the bidder exclusion effect. Note that the testing procedure described in Section 4 cannot be used if the number of bidders is unobserved because it explicitly requires observing realizations of N.

Asymmetric Bidders. If bidders have private values but are asymmetric—that is, their indices in the joint distribution of valuations are non-exchangeable—then the test for independence of valuations and N can stil be applied. However, the notion of independence between valuations and N is less straightforward in this case, and it less clear what it might mean if a test rejects the assumption. Intuitively, values may fail to be independent of N either because different bidders are more likely to enter depending on N, or because the same bidders enter but the value of the goods sold varies by N. We formalize and prove this statement in Appendix B, following the setup of Coey et al. (2017).

The bound on counterfactual revenue under an optimal reserve price no longer applies with bidder asymmetries, as the Bulow and Klemperer (1996) result requires bidder symmetry. In particular, it is difficult to conceptualize what type of bidder would be implied by the "additional random bidder" in the Bulow-Klemperer setting when bidders are asymmetric. The merger case, however, immediately applies even if bidders are asymmetric. This is useful given that Tschantz et al. (2000), Dalkir et al. (2000), Li and Zhang (2015), and others argue that allowing for bidder asymmetries is particularly important to capture realistic aspects of mergers in auction settings.

A Test Using All Bids From Button Auctions. Here we return to the framework addressed in the body of the article but consider the question of how, taking the button auction assumption seriously, the econometrician could use all bids—rather than just those of the second- and third-highest value bidders—to test for the independence of valuations and N.

Suppose the econometrician observes auctions in which there are n bidders, $n+1$ bidders, all the way to $\bar{n}$ bidders for some $\bar{n} > n$. For now, let n be fixed. For $k = 1, ..., \bar{n} - n$, the test described in Section 4 can be extended to compare revenue in auctions with n bidders to revenue in auctions with $n + k$ bidders where k bids have been randomly removed. This latter object can be computed as follows for an auction j in which $n+k$ bidders were present:

$$\frac{1}{\binom{n+k}{k}} \sum_{\iota=1}^{k+1} \binom{\iota}{\iota-1} \binom{n+k-(\iota+1)}{k-(\iota-1)} b_j^{n+k-\iota:n+k} \tag{12}$$

The denominator at the beginning of equation (12), $\binom{n+k}{k}$, indicates the number of different ways k bids could be dropped from $n + k$ bids. On the interior of the sum in equation (12), the order statistic $b_j^{n+k-\iota:n+k}$ can only be the revenue-setting bid if the set of k randomly removed bids *does not* include the $(\iota + 1)^{th}$-highest bid (that is, the bid $b_j^{n+k-\iota:n+k}$) and *does* include $\iota - 1$ out of the ι highest bids (this is captured by the term $\binom{\iota}{\iota-1}$ in equation (12)). The remaining $k - (\iota - 1)$ bids in the set of k bids dropped can come from any of the other $n + k - (\iota + 1)$ lower bids (this is captured by the term $\binom{n+k-(\iota+1)}{k-(\iota-1)}$ in equation (12)).

Still treating n as fixed, the parametric test presented in (9) can then be modified as follows.

$$y_j = \sum_{k=0}^{\bar{n}-n} a_k + \beta X_j + e_{kj} \tag{13}$$

where if $j \in A_{n+k}$, then y_j is replaced by the value from (12). The objects a_k are indicators for which value of k the auction revenue measure y_j comes from (where $k = 0$ means the actual revenue in n bidder auctions). The test for independence of valuations and the number of bidders can then be performed through a Wald test, testing simultaneously for whether all the coefficients a_k are equal to zero.

This same procedure can be performed for multiple values of n simultaneously, comparing, for example, three-bidder auctions to five-bidder auctions with two bidders removed *and* simultaneously comparing four-bidder auctions to six-bidder auctions with two bidders removed). To do so, the regression equations in (13) can be formed for each value of n and then stacked in a Seemingly Unrelated Regression (SUR, Zellner 1962). The Wald test would then use the estimated coefficients on all of the indicator variables a_k for each value of n and the SUR variance-covariance matrix (allowing the regression errors to be correlated within each value of the number of bidders actually present at the auction, $n+k$) to test the null hypothesis that all of these coefficients are equal to zero.

Other Extensions: Binding Reserve Prices and First Price Auctions. The testing procedure and optimal revenue bounding procedure can be applied in ascending auctions with binding reserve prices or in first price auctions. The presentation of these results requires a number of additional proofs, however, and we therefore present these results in Appendix B.

8. Conclusion

We developed a computationally simple test of independence of bidders' valuations and the number of bidders, a commonly invoked assumption in structural empirical auctions work. The test relies on computing the decrease in seller revenue from removing at random one of n bidders from the auction and comparing this quantity to the actual revenue difference between n and $n-1$ bidder auctions. We demonstrated that this quantity—the bidder exclusion effect—can also be used to bound counterfactual changes in revenue from the seller adopting an optimal reserve price or from bidders merging. We applied our proposed test to data from timber auctions, a setting in which the assumption of bidder valuations being independent of the number of bidders has been exploited in a number of studies. We found evidence to reject this independence assumption in our data. We derived our main results within a symmetric, conditionally independent private values settings at ascending button auctions, and then discussed a number of extensions to other environments.

We believe that this tool is also likely to be useful for other questions as well. For example, the bidder exclusion effect may be useful in multi-unit auction settings or internet search position auctions. As another example, the bidder exclusion effect can provide a simple specification check of standard assumptions in empirical auctions analysis: Under the assumption of independent private values in button auctions, one can invert the second-order statistic distribution to obtain an estimate of the underlying distribution of buyer valuations (Athey and Haile 2007) and simulate the revenue increase under an optimal reserve price; if the simulated revenue increase exceeds the bidder exclusion effect, the validity of either the assumption of independence or the assumption of private values—or both—is in question.

References

Andrews, D. and X. Shi (2013). Inference based on conditional moment inequalities. *Econometrica 81*(2), 609–666.

Aradillas-López, A., A. Gandhi, and D. Quint (2013). Identification and inference in ascending auctions with correlated private values. *Econometrica 81*(2), 489–534.

Aradillas-López, A., A. Gandhi, and D. Quint (2016). A simple test for moment inequality models with an application to english auctions. *Journal of Econometrics 194*(1), 96–115.

Armstrong, T. (2013). Bounds in auctions with unobserved heterogeneity. *Quantitative Economics 4*, 377–415.

Asker, J. (2010). A study of the internal organization of a bidding cartel. *American Economic Review*, 724–762.

Athey, S., D. Coey, and J. Levin (2013). Set-asides and subsidies in auctions. *American Economic Journal: Microeconomics 5*(1), 1–27.

Athey, S. and P. A. Haile (2002). Identification of standard auction models. *Econometrica 70*(6), 2107–2140.

Athey, S. and P. A. Haile (2007). Nonparametric approaches to auctions. *Handbook of Econometrics 6*, 3847–3965.

Athey, S., J. Levin, and E. Seira (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *Quarterly Journal of Economics 126*(1), 207–257.

Backus, M. and G. Lewis (2012). A demand system for a dynamic auction market with directed search. Working paper, Harvard University.

Baker, J. B. (1997). Unilateral competitive effects theories in merger analysis. *Antitrust 11*, 21.

Bhattacharya, V., J. W. Roberts, and A. Sweeting (2014). Regulating bidder participation in auctions. *RAND Journal of Economics 45*(4), 675–704.

Bikhchandani, S., P. A. Haile, and J. G. Riley (2002). Symmetric separating equilibria in english auctions. *Games and Economic Behavior 38*(1), 19–27.

Brannman, L. and L. M. Froeb (2000). Mergers, cartels, set-asides, and bidding preferences in asymmetric oral auctions. *Review of Economics and Statistics 82*(2), 283–290.

Bulow, J. and P. Klemperer (1996). Auctions versus negotiations. *American Economic Review 86*(1), 180–194.

Bulow, J. and P. Klemperer (2002). Prices and the winner's curse. *RAND Journal of Economics*, 1–21.

Bulow, J. and J. Roberts (1989). The simple economics of optimal auctions. *Journal of Political Economy*, 1060–1090.

Chernozhukov, V., S. Lee, and A. M. Rosen (2013). Intersection bounds: Estimation and inference. *Econometrica 81*(2), 667–737.

Chesher, A. and A. Rosen (2017). Incomplete english auction models with heterogeneity. Working paper, Centre for Microdata Methods and Practice, Institute for Fiscal Studies.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics 1*, 451–488.

Chetverikov, D. (2018). Adaptive tests of conditional moment inequalities. *Econometric Theory 34*(1), 186–227.

Coey, D., B. Larsen, and B. Platt (2016). A theory of bidding dynamics and deadlines in online retail. NBER Working Paper 22038.

Coey, D., B. Larsen, and K. Sweeney (2014). The bidder exclusion effect. NBER Working Paper 20523.

Coey, D., B. Larsen, K. Sweeney, and C. Waisman (2017). Ascending auctions with bidder asymmetries. *Quantitative Economics 8*(1), 181–200.

Coey, D., B. Larsen, K. Sweeney, and C. Waisman (2018). The simple empirics of optimal online auctions. NBER Working Paper 24698.

Crémer, J. and R. P. McLean (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 1247–1257.

Dalkir, S., J. W. Logan, and R. T. Masson (2000). Mergers in symmetric and asymmetric noncooperative auction markets: The effects on prices and efficiency. *International Journal of Industrial Organization 18*(3), 383–413.

David, H. A. and H. N. Nagaraja (1970). *Order Statistics.* Wiley Online Library.

Dughmi, S., T. Roughgarden, and M. Sundararajan (2012). Revenue submodularity. *Theory of Computing 8*(1), 95–119.

Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics 125*(3), 877–921.

Freyberger, J. and B. Larsen (2017). Identification in ascending auctions, with an application to digital rights management. NBER Working paper 23569.

Froeb, L., M. Shor, and S. Tschantz (2008). Mergers in auctions with an incumbent advantage. *Vanderbilt Law and Economics Research Paper* (08-24).

Froeb, L., S. Tschantz, and P. Crooke (1998). Mergers among asymmetric bidders: A logit second-price auction model. *Available at SSRN 69415*.

FTC (2009). Horizontal merger guidelines: Questions for public comment. Federal Trade Commission and U.S. Department of Justice.

Gentry, M. and T. Li (2013). Identification in auctions with selective entry. *Econometrica 82*, 315–344.

Graham, D. A., R. C. Marshall, and J.-F. Richard (1990). Differential payments within a bidder coalition and the Shapley value. *American Economic Review*, 493–510.

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

Hendricks, K. and A. Sorensen (2015). The role of intermediaries in dynamic auction markets. Working paper, University of Wisconsin.

Hernandez, C., D. Quint, and C. Turansick (2018). Estimation in english auctions with unobserved heterogeneity. Working paper, University of Wisconsin.

Jaffe, S. and E. G. Weyl (2013). The first-order approach to merger analysis. *American Economic Journal: Microeconomics 5*(4), 188–218.

Kim, D.-H. (2013). Optimal choice of a reserve price under uncertainty. *International Journal of Industrial Organization 31*(5), 587–602.

Komarova, T. (2013). Partial identification in asymmetric auctions in the absence of independence. *The Econometrics Journal 16*(1), S60–S92.

Krasnokutskaya, E. (2011). Identification and estimation of auction models with unobserved heterogeneity. *Review of Economic Studies 78*(1), 293–327.

Krishna, V. (2009). *Auction theory*. Academic press.

Lacetera, N., B. J. Larsen, D. G. Pope, and J. R. Sydnor (2016). Bid takers or market makers? the effect of auctioneers on auction outcomes. *American Economic Journal: Microeconomics 8*(4), 195–229.

Levin, D. and J. L. Smith (1994). Equilibrium in auctions with entry. *American Economic Review 84*(3), 585–599.

Li, T. and I. Perrigne (2003). Timber sale auctions with random reserve prices. *Review of Economics and Statistics 85*(1), 189–200.

Li, T., I. Perrigne, and Q. Vuong (2000). Conditionally independent private information in ocs wildcat auctions. *Journal of Econometrics 98*(1), 129–161.

Li, T., I. Perrigne, and Q. Vuong (2003). Semiparametric estimation of the optimal reserve price in first-price auctions. *Journal of Business & Economic Statistics 21*(1), 53–64.

Li, T. and B. Zhang (2015). Affiliation and entry in first-price auctions with heterogeneous bidders: An analysis of merger effects. *American Economic Journal: Microeconomics 7*(2), 188–214.

Li, T. and X. Zheng (2012). Information acquisition and/or bid preparation: A structural analysis of entry and bidding in timber sale auctions. *Journal of Econometrics 168*(1), 29–46.

Liu, N. and Y. Luo (2017). A nonparametric test for comparing valuation distributions in first-price auctions. *International Economic Review 58*(3), 857–888.

Lopomo, G. (1995). Optimality and robustness of the English auction. New York University, Stern School of Business Working Paper No. EC-95-03.

Mailath, G. J. and P. Zemsky (1991). Collusion in second price auctions with heterogeneous bidders. *Games and Economic Behavior 3*(4), 467–486.

Marmer, V., A. Shneyerov, and U. Kaplan (2016). Identifying collusion in english auctions. Working paper, University of British Columbia.

Marmer, V., A. Shneyerov, and P. Xu (2013). What model for entry in first-price auctions? a nonparametric approach. *Journal of Econometrics 176*(1), 46–58.

McAfee, R. P. and J. McMillan (1987). Auctions with entry. *Economics Letters 23*(4), 343–347.

McAfee, R. P., J. McMillan, and P. J. Reny (1989). Extracting the surplus in the common-value auction. *Econometrica*, 1451–1459.

McAfee, R. P. and P. J. Reny (1992). Correlated information and mechanism design. *Econometrica*, 395–421.

Milgrom, P. R. and R. J. Weber (1982). A theory of auctions and competitive bidding. *Econometrica*, 1089–1122.

Moresi, S. (2009). Bidding competition and the upp test. Horizontal Merger Guidelines Review Project-Comment, Project No. P092900.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research 6*(1), 58–73.

Ostrovsky, M. and M. Schwarz (2016). Reserve prices in internet advertising auctions: A field experiment. Working paper, Stanford University.

Paarsch, H. J. (1997). Deriving an estimate of the optimal reserve price: An application to british columbian timber sales. *Journal of Econometrics 78*(1), 333–357.

Pesendorfer, M. (2000). A study of collusion in first-price auctions. *Review of Economic Studies 67*(3), 381–411.

Pinkse, J. and G. Tan (2005). The affiliation effect in first-price auctions. *Econometrica 73*(1), 263–277.

Riley, J. G. and W. F. Samuelson (1981). Optimal auctions. *American Economic Review*, 381–392.

Roberts, J. W. and A. Sweeting (2013). When should sellers use auctions? *American Economic Review 103*(5), 1830–1861.

Roberts, J. W. and A. Sweeting (2014). Bailouts and the preservation of competition. *American Economic Journal: Microeconomics*, forthcoming.

Samuelson, W. F. (1985). Competitive bidding with entry costs. *Economics Letters 17*(1-2), 53–57.

Tadelis, S. and F. Zettelmeyer (2015). Information disclosure as a matching mechanism: Theory and evidence from a field experiment. *American Economic Review 105*(2), 886–905.

Tang, X. (2011). Bounds on revenue distributions in counterfactual auctions with reserve prices. *RAND Journal of Economics 42*(1), 175–203.

Tschantz, S., P. Crooke, and L. Froeb (2000). Mergers in sealed versus oral auctions. *International Journal of the Economics of Business 7*(2), 201–212.

von Ungern-Sternberg, T. (1988). Cartel stability in sealed bid second price auctions. *Journal of Industrial Economics*, 351–358.

Waehrer, K. (1999). Asymmetric private values auctions with application to joint bidding and mergers. *International Journal of Industrial Organization 17*(3), 437–452.

Waehrer, K. and M. K. Perry (2003). The effects of mergers in open-auction markets. *RAND Journal of Economics*, 287–304.

Zellner, A. (1962). An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association 57*(298), 348–368.

Figure 1. Bounding the Bidder Exclusion Effect

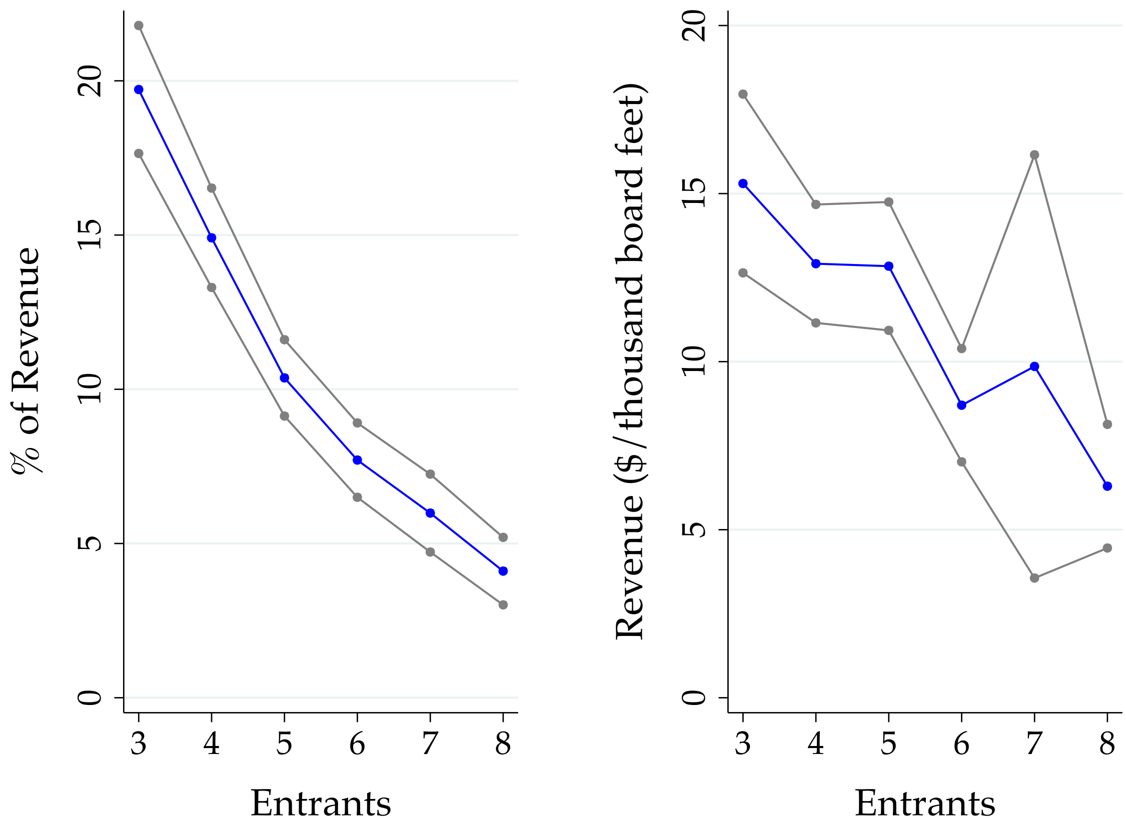

Notes: Graphs show point estimates and 95% pointwise confidence intervals for $\frac{2}{n}E(B^{n-1:n} - B^{n-2:n})$, for various values of n, the total number of entrants in the auction. Estimates in the left graph are expressed as a percentage of auction revenue. Estimates in the right graph are expressed in dollars per thousand board feet.

Table 1. Unconditional Tests for Dependence of Valuations and N, All Auctions

Entrants	3	4	5	6	7	8
a_1	78.15***	92.34***	119.75***	125.64***	123.16***	147.24***
	(6.36)	(4.62)	(4.81)	(11.14)	(6.69)	(9.65)
a_2	49.60***	78.37***	89.09***	119.66***	125.02***	130.46***
	(2.62)	(6.43)	(3.50)	(4.85)	(11.15)	(12.04)
$T(n) = a_1 - a_2$	28.55***	13.97	30.66***	5.98	-1.86	16.78
	(6.87)	(7.92)	(5.95)	(12.15)	(13.00)	(15.43)
Sample Size	497	496	456	350	243	164

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Table presents results of test for dependence of valuations and N unconditional on covariates, for various levels of the number of entrants.

Table 2. Conditional Tests for Dependence of Valuations and N, All Auctions

Entrants	3	4	5	6	7	8
a_1	-44.52	-120.64	3.34	-95.81**	-111.39	-172.47
	(59.38)	(74.72)	(24.72)	(32.93)	(71.81)	(159.06)
a_2	-64.84	-137.66	-10.47	-89.90**	-124.05	-167.05
	(59.39)	(75.24)	(25.01)	(31.83)	(72.75)	(145.70)
$T(n) = a_1 - a_2$	20.32***	17.02***	13.81***	-5.91	12.66*	-5.42
	(5.33)	(4.21)	(3.63)	(4.75)	(5.78)	(17.94)
Sample Size	497	496	456	350	243	164

Heteroskedasticity-robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Table presents results of test for dependence of valuations and N conditional on covariates as described in Section 4, for various levels of the number of entrants.

Table 3. Conditional Tests for Dependence of Valuations and N, Auctions with Only Loggers

Entrants	3	4	5	6
a_1	-158.28	-183.18	26.03	-285.34**
	(108.96)	(115.53)	(42.50)	(115.06)
a_2	-172.28	-222.73	8.53	-266.30
	(105.59)	(128.23)	(43.13)	(111.34)
$T(n) = a_1 - a_2$	14.00	39.54*	17.51	-19.04
	(16.65)	(19.29)	(9.92)	(25.66)
Sample Size	149	138	109	76

Heteroskedasticity-robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Table presents results of test for dependence of valuations and N conditional on covariates as described in Section 4, for various levels of the number of entrants, and only for auctions in which all entrants are loggers.

Appendix A. Proofs

Proof of Proposition 1.

Proof. We first prove that if $Z_1, \ldots, Z_{n+1}$ are iid random variables,

$$E(\max\{Z_1, \ldots, Z_n\}) - E(\max\{Z_1, \ldots, Z_{n-1}\}) \geq E(\max\{Z_1, \ldots, Z_{n+1}\}) - E(\max\{Z_1, \ldots, Z_n\}). \quad (14)$$

For any $z_1, \ldots, z_{n+1} \in \mathbb{R}^{n+1}$,

$$\max\{0, z_{n+1} - \max\{z_1, \ldots, z_{n-1}\}\} \geq \max\{0, z_{n+1} - \max\{z_1, \ldots, z_n\}\}, \quad (15)$$

implying

$$\max\{z_1, \ldots, z_{n-1}, z_{n+1}\} - \max\{z_1, \ldots, z_{n-1}\} \geq \max\{z_1, \ldots, z_{n+1}\} - \max\{z_1, \ldots, z_n\}. \quad (16)$$

Consequently for any random variables $Z_1, \ldots, Z_{n+1}$,

$$E(\max\{Z_1, \ldots, Z_{n-1}, Z_{n+1}\}) - E(\max\{Z_1, \ldots, Z_{n-1}\}) \geq E(\max\{Z_1, \ldots, Z_{n+1}\}) - E(\max\{Z_1, \ldots, Z_n\}), \quad (17)$$

because (16) holds for every realization $z_1, \ldots, z_{n+1}$ of $Z_1, \ldots, Z_{n+1}$. If the Z_i are iid, then $E(\max\{Z_1, \ldots, Z_{n-1}, Z_{n+1}\}) = E(\max\{Z_1, \ldots, Z_n\})$, yielding (14).

The expected revenue from any mechanism is the expected marginal revenue of the winning bidder (Myerson 1981). Ascending auctions assign the good to the bidder with the highest valuation, and therefore highest marginal revenue, because marginal revenue is increasing in valuations. It follows that expected revenue with n bidders is $E(\max\{MR(V_1), \ldots, MR\,(V_n)\})$. As $MR(V_1), \ldots, MR(V_{n+1})$ are iid random variables, we have

$$E(\max\{MR(V_1), \ldots, MR(V_n)\}) - E(\max\{MR(V_1), \ldots, MR(V_{n-1})\}) \geq E(\max\{MR(V_1), \ldots, MR(V_{n+1})\}) - E(\max\{MR(V_1), \ldots, MR(V_n)\}), \quad (18)$$

implying that Proposition 1 holds in independent private values settings. When there exists a random variable U such that bidder values $V_1, \ldots, V_n$ are iid conditional on U, if marginal revenue is increasing in values conditional on each realization of U, then the above proof applies conditional on each realization of U. Taking expectations over U, it follows that Proposition 1 holds in CIPV environments if bidders' marginal revenue curves are increasing in values conditional on each value of U. □

Proof of Proposition 2.

Proof. Let i and j represent the bidders who merge, and without loss of generality let $V_i \geq V_j$. Let $M^{1:n-1}, ..., M^{n-1:n-1}$ represent order statistics of $\{M_k\}_{k \neq i,j} \cup \{M_{i,j}\}$. Revenue in the presence of the merger will be given by $M^{n-2:n-1}$, and thus the revenue loss due to the merger is given by $E[B^{n-1:n} - M^{n-2:n-1}]$. If j, the lower-valued bidder in the joint entity, is such that $V_j = V^{k:n}$ for some $k \in \{1, ..., n-2\}$, then revenue will not drop due to the merger, because in this case

$$M^{n-2:n-1} \geq V^{n-1:n} \tag{19}$$

$$= B^{n-1:n}, \tag{20}$$

where the first line holds by Assumption 1. The only case where revenue may potentially drop is instead when j is such that $V_j = V^{n-1:n}$, in which case

$$M^{n-2:n-1} = V^{n-2:n} \tag{21}$$

$$\geq B^{n-2:n}, \tag{22}$$

where the first line holds by Assumption 1. Therefore, an upper bound the loss in revenue due to two bidders merging is given by

$$E[B^{n-1:n} - M^{n-2:n-1}] \leq E[B^{n-1:n} - B^{n-2:n}] \tag{23}$$

proving (ii).

The proof of (i) follows by noting that if two bidders are randomly selected, only *one* such pair of bidders i,j is such that j, the lower-valued bidder, has $V_j = V^{n-1:n}$ and the probability of selecting this pair is given by $\frac{1}{\binom{n}{2}} = \frac{2}{n(n-1)}$. Therefore,

$$E[B^{n-1:n} - M^{n-2:n-1}] \leq \frac{2}{n(n-1)} E[B^{n-1:n} - V^{n-2:n}] = \frac{1}{n-1} \Delta(n)$$

proving the result. □

Appendix B. Extensions of Bidder Exclusion Effect and Monte Carlo Simulations

Auctions in Which Reserve Prices Are Present

We consider ascending auctions with private values where bidders bid their values, and where there is a reserve price below which bids are not observed (i.e. a binding reserve price). We modify our notation accordingly: $\Delta(n, r)$ denotes the fall in expected revenue produced by randomly excluding a bidder from n bidder auctions, when the reserve price is r.

Proposition 3. *In ascending auctions with private values and a reserve price of r where bidders bid their value, for all $n > 2$ the bidder exclusion effect* $\Delta(n, r) = \frac{2}{n}E(B^{n-1:n} - \max(B^{n-2:n}, r)|r \leq B^{n-1:n}) \Pr(r \leq B^{n-1:n}) + \frac{1}{n} r \Pr(B^{n-1:n} < r \leq B^{n:n})$.

Proof. If $r \leq B^{n-1:n}$, then with probability $\frac{2}{n}$ dropping a bidder at random will cause revenue to fall from $B^{n-1:n}$ to $\max(B^{n-2:n}, r)$, so that in expectation revenue falls by $\frac{2}{n}E(B^{n-1:n} - \max(B^{n-2:n}, r)|r \leq B^{n-1:n})$. If $B^{n-1:n} < r \leq B^{n:n}$, then with probability $\frac{1}{n}$ dropping a bidder at random will cause revenue to fall from r to 0. If $B^{n:n} < r$, then dropping a bidder at random will not change revenue. These observations imply the result. □

This expression for $\Delta(n, r)$ can be estimated given observed data, as it does not depend on knowing the value of bids lower than the reserve price.

When the reserve price equals r in both n and $n-1$ bidder auctions, the expected revenue difference between those auctions is

$$E(\max(B^{n-1:n}, r)|r \leq B^{n:n}) \Pr(r \leq B^{n:n}) - E(\max(B^{n-2:n-1}, r)|r \leq B^{n-1:n-1}) \Pr(r \leq B^{n-1:n-1}). \tag{24}$$

If valuations are independent of N, then $F^n_{n-1} = F^{n-1}_{n-1} = F^{n-1}$ and hence expression (24) equals the expression for $\Delta(n, r)$ of Proposition 3. As in Section 4, we can test this hypothesis with a t-test, where the test statistic is formed by replacing expectations by sample averages. This test is consistent against forms of dependence between valuations and N that affect expected revenue, i.e. such that $E(\max(B^{n-2:n-1,n}, r)|r \leq B^{n-1:n-1,n}) \Pr(r \leq B^{n-1:n-1,n}) \neq E(\max(B^{n-2:n-1}, r)|r \leq B^{n-1:n-1}) \Pr(r \leq B^{n-1:n-1})$.

This test can be adapted to incorporate covariates. The null hypothesis is:

$$E\left(\mathbb{1}(r \leq B^{n-1:n})\left(\frac{n-2}{n}B^{n-1:n} + \frac{2}{n}\max\{B^{n-2:n}, r\}\right) + \mathbb{1}(B^{n-1:n} < r \leq B^{n:n})\frac{n-1}{n}r\middle|X\right)$$
$$= E(\mathbb{1}(r \leq B^{n-1:n-1})\max\{B^{n-2:n-1}, r\}|X). \tag{25}$$

This states that, conditional on covariates, revenue in n bidder auctions when one bidder is dropped at random equals revenue in $n-1$ bidder auctions. The regression-based test of Section 4 can be modified to test this restriction.

For the application to optimal mechanism design, we require an upper bound on $\Delta(n, 0)$. Using the fact that bids are non-negative, we can write this upper bound as the sum of three separate conditional expectations, one for each of the possible orderings of r, $B^{n-2:n}$, and $B^{n-1:n}$:

$$\begin{aligned}
\Delta(n,0) = & \frac{2}{n}E(B^{n-1:n} - B^{n-2:n}|r \leq B^{n-2:n})\Pr(r \leq B^{n-2:n}) \\
& + \frac{2}{n}E(B^{n-1:n} - B^{n-2:n}|B^{n-2:n} < r \leq B^{n-1:n})\Pr(B^{n-2:n} < r \leq B^{n-1:n}) \\
& + \frac{2}{n}E(B^{n-1:n} - B^{n-2:n}|B^{n-1:n} < r)\Pr(B^{n-1:n} < r) \qquad (26)\\
\leq & \frac{2}{n}E(B^{n-1:n} - B^{n-2:n}|r \leq B^{n-2:n})\Pr(r \leq B^{n-2:n}) \\
& + \frac{2}{n}E(B^{n-1:n}|B^{n-2:n} < r \leq B^{n-1:n})\Pr(B^{n-2:n} < r \leq B^{n-1:n}) \\
& + \frac{2}{n}r\Pr(B^{n-1:n} < r) \qquad (27)
\end{aligned}$$

The terms in (27) do not depend on knowing the value of bids lower than the reserve price, and can be estimated given observed data. The application to mergers can be extended analogously.

First Price Auctions

We now give upper and lower bounds on the bidder exclusion effect in first price auctions with symmetric IPV, and symmetric conditionally independent private values (CIPV). Unlike the ascending button auction case, in first price auctions there is a distinction between the CIPV environment and an IPV environment with unobserved auction-level heterogeneity; our results here only apply to the CIPV case. Let $b(V_i, F^n)$ denote bidder i's equilibrium bid, as a function of his value, V_i, and the distribution of bidders' valuations, F^n. We assume V_i is continuously distributed on some interval $[0, u]$. In this section we use subscripts to make explicit the distribution with respect to which expectations are taken, e.g. expected revenue with no reserve price is $E_{F^n}(b(V^{n:n}, F^n))$ in n bidder auctions and is $E_{F^n_{n-1}}(b(V^{n-1:n-1}, F^n_{n-1}))$ when one of the n bidders is randomly excluded. The bidder exclusion effect is $\Delta(n) \equiv E_{F^n}(b(V^{n:n}, F^n)) - E_{F^n_{n-1}}(b(V^{n-1:n-1}, F^n_{n-1}))$.

Proposition 4. *In first price auctions if i) bidders have symmetric independent private values, or ii) there is a random variable U common knowledge to bidders such that bidders have symmetric independent private values conditional on U, then* $E_{F^n}(b(V^{n:n}, F^n)) - E_{F^n_{n-1}}(b(V^{n-1:n-1}, F^n)) < \Delta(n) < E_{F^n}(b(V^{n:n}, F^n)) - E_{F^n_{n-1}}(b(V^{n-2:n-1}, F^n))$.

Proof. We first consider the case of symmetric independent private values. For the lower bound, note that in symmetric independent private values settings, equilibrium bids are strictly increasing in n: $b(v_i, F^n) > b(v_i, F^n_{n-1})$ (see, for example, Krishna 2009). This implies $E_{F^n_{n-1}} b(V^{n-1:n-1}, F^n) > E_{F^n_{n-1}} b(V^{n-1:n-1}, F^n_{n-1})$, and therefore $E_{F^n}(b(V^{n:n}, F^n)) - E_{F^n_{n-1}} b(V^{n-1:n-1}, F^n) < \Delta(n)$.

For the upper bound, we have

$$E_{F^n_{n-1}}(b(V^{n-2:n-1}, F^n)) < E_{F^n_{n-1}}(V^{n-2:n-1}) \tag{28}$$

$$= E_{F^n_{n-1}}(b(V^{n-1:n-1}, F^n_{n-1})). \tag{29}$$

The inequality holds because equilibrium bids are strictly less than values. The equality holds by revenue equivalence of first and second price auctions with symmetric independent private values. It follows that $\Delta(n) < E_{F^n}(b(V^{n:n}, F^n)) - E_{F^n_{n-1}}(b(V^{n-2:n-1}, F^n))$.

If values are symmetric and CIPV, then because U is common knowledge to bidders these lower and upper bounds hold conditional on every realization of U, and therefore hold unconditionally, taking expectations with respect to U. The bounds thus extend to the conditionally independent private values case. □

The lower bound above is the expected fall in revenue in n bidder auctions when one bid is removed at random, assuming the good will be sold at a price equal to the highest of the remaining bids. The upper bound is the expected fall in revenue in n bidder auctions when one bid is removed at random, assuming the good will be sold at a price equal to the *second* highest of the remaining bids. The following corollary characterizes these bounds more explicitly in terms of the bids from n bidder auctions.

Corollary 2. *In first price auctions if i) bidders have symmetric independent private values, or ii) there is a random variable U common knowledge to bidders such that bidders have symmetric independent private values conditional on U, then*

$$\frac{1}{n}\left(E_{F^n}(b(V^{n:n},F^n)) - E_{F^n}(b(V^{n-1:n},F^n))\right) < \Delta(n) \tag{30}$$

and

$$\Delta(n) < \frac{n-2}{n}\left(E_{F^n}(b(V^{n:n},F^n)) - E_{F^n}(b(V^{n-1:n},F^n))\right) + \frac{2}{n}\left(E_{F^n}(b(V^{n:n},F^n)) - E_{F^n}(b(V^{n-2:n},F^n))\right). \tag{31}$$

Proof. For the lower bound, note that with probability $\frac{n-1}{n}$ dropping a bid at random will not change the highest bid, and with probability $\frac{1}{n}$ the highest bid will drop from $b(V^{n:n},F^n)$ to $b(V^{n-1:n},F^n)$. For the upper bound, note that with probability $\frac{n-2}{n}$ the difference between the highest bid in the original sample and the second-highest bid after one bid has been dropped at random is $b(V^{n:n},F^n) - b(V^{n-1:n},F^n)$, and with probability $\frac{2}{n}$ it is $b(V^{n:n},F^n) - b(V^{n-2:n},F^n)$. □

Several remarks on these bounds are in order.

Remark 1. *The lower bound in Proposition 4 also holds under the more general setting of symmetric correlated private values, as long as equilibrium bids are strictly increasing in n.*[29]

Remark 2. *The upper bound in Proposition 4 also holds if bidders are risk-averse instead of risk-neutral, as first price auctions raise more revenue than ascending auctions with symmetric risk-averse bidders in IPV environments (Riley and Samuelson 1981).*

Remark 3. *In the CIPV case, if U is not common knowledge amongst bidders, then bidders' private information is correlated conditional on what they know at the time of bidding. This affects equilibrium bidding behavior and the argument of Proposition 4 does not hold.*

[29]Pinkse and Tan (2005) give conditions for this to hold.

Remark 4. *The upper bound of Proposition 4 can be replaced by* $E_{F^n}(b(V^{n:n}, F^n)) - E_{F^{n'}_{n-1}}(b(V^{n-2:n-1}, F^{n'}))$ *for any* $n' > n-1$*, as bids are below values in* n' *bidder auctions too. Consequently,* $\Delta(n) \leq E_{F^n}(b(V^{n:n}, F^n)) - \sup_{n'} E_{F^{n'}_{n-1}}(b(V^{n-2:n-1}, F^{n'}))$*.*

As with ascending auctions, the bidder exclusion effect can be used to test for dependence between valuations and N in first price auctions. Under the null hypothesis that valuations and N are independent, for all $n \geq 2$, $F^n_{n-1} = F^{n-1}_{n-1} = F^{n-1}$. This implies that the bidder exclusion effect $\Delta(n) \equiv E_{F^n}(b(V^{n:n}, F^n)) - E_{F^n_{n-1}}(b(V^{n-1:n-1}, F^n_{n-1}))$ equals $E_{F^n}(b(V^{n:n}, F^n)) - E_{F^{n-1}}(b(V^{n-1:n-1}, F^{n-1}))$. If the sample analog of $E_{F^n}(b(V^{n:n}, F^n)) - E_{F^{n-1}}(b(V^{n-1:n-1}, F^{n-1}))$—which is simply average revenue in n bidder auctions minus average revenue in $n-1$ bidder auctions—lies outside the sample analogs of the lower or upper bounds of Corollary 2, this is evidence against the null hypothesis. This test is consistent against violations of the null when values are "sufficiently" decreasing or increasing with n. Precisely, this is the case if $E_{F^{n-1}}(b(V^{n-1:n-1}, F^{n-1})) > E_{F^n_{n-1}}(b(V^{n-1:n-1}, F^n))$ or $E_{F^{n-1}}(b(V^{n-1:n-1}, F^{n-1})) < E_{F^n_{n-1}}(b(V^{n-2:n-1}, F^n))$. Again, the regression-based test of Section 6 can be modified to test that the null hypothesis holds conditional on observable covariates, rather than unconditionally.

The application to optimal mechanism design also works for first price auctions. The Bulow-Klemperer theorem is stated for ascending auctions, but by revenue equivalence also applies to first price auctions when bidders have symmetric IPV (or symmetric CIPV). Thus Proposition 1 extends to first price auctions, where the upper bound on the effect on expected revenue of improving mechanism design is given by Corollary 2. We do not extend the merger analysis to first price auctions.

Haile and Tamer (2003) Setting

Haile and Tamer (2003) make the following assumption:

Assumption 2. *Bidders (i) do not bid more than they are willing to pay and (ii) do not allow an opponent to win at a price they are willing to beat.*

As in the body of the article, $B^{k:n}$ represents the $(k:n)$ order statisic of bids. In a non-button auction setting, these bids are the last indication of willingness to pay by a bidder. As in Haile and Tamer (2003) we use $B^{n:n}$ to denote the final price of the auction, which can exceed the willingness-to-pay of the second-highest value bidder due to jump bidding or bid increments. Let $B^{n-1:n-1,n}$ represent the final auction price when one bidder is randomly removed from an n bidder auction.

We again use the notation $\Delta(n)$ to denote the bidder exclusion effect, i.e. the expected drop in revenue when a random bidder is excluded, which in this setting will be given by

$$\Delta(n) \equiv E(B^{n:n}) - E(B^{n-1:n-1,n})$$

In addition to Assumption 2, we make the following assumption:

Assumption 3. *Removing one of the $n-2$ lower bidders from the auction does not affect the auction price.*

Assumption 3 is unnecessary in the results stated in the body of the article, but in the Haile and Tamer (2003) case, the conditions in Assumption 2 are weak enough that they do not rule out some cases that could lead to a change in the final price at the auction. For example, it might be the case that the top bidder's likelihood of jump bidding is lower when one of the $n-2$ bidders is removed. However, Assumption 3 is also quite weak in practice, as the final price is set by back and forth activity between the top two bidders and hence is unlikely to be affected by dropped one of the $n-2$ lowest bidders.

We now state our result for this setting. Let τ represent the minimum bid increment.

Proposition 5. *In private values ascending auctions, if Assumptions 2 and 3 hold, then for all $n > 2$ the bidder exclusion effect $\Delta(n) \leq \frac{2}{n}E\left[B^{n:n} - (B^{n-2:n} - \tau)\right]$.*

Proof. With probability $\frac{n-2}{n}$, dropping a random bidder will have no effect on revenue, by Assumption 3, which will remain at $B^{n:n}$. With probability $\frac{2}{n}$ one of the highest two bidders will be dropped. In this case, $B^{n-1:n-1,n} \geq V^{n-2:n} - \tau$, because, if not, the $(n-2:n)$ bidder would have bid higher than $B^{n-1:n-1,n}$ by Assumption 2(ii). It is then also true that $B^{n-1:n-1,n} \geq B^{n-2:n} - \tau$ by Assumption 2(i). Therefore, the drop in revenue from removing a random bidder, $\Delta(n)$, is bounded weakly above by $\frac{2}{n}E\left[B^{n:n} - (B^{n-2:n} - \tau)\right]$. □

Asymmetric Bidders

We give sufficient conditions for valuations to be independent of N with asymmetric bidders and private values, following the setup of Coey et al. (2017). Let $\mathbb{N}$ be the full set of potential bidders. Let $\mathcal{P}$ be a random vector representing the identities or types of bidders participating in an auction, with realizations $P \subset \mathbb{N}$. Let N be a random variable representing the number of bidders participating in an auction, with realizations $n \in \mathbb{N}$. When necessary to clarify the number of bidders in a set of participating bidders, we let P_n denote an arbitrary set of n participating bidders. Define F^P to be the joint distribution of $(V_i)_{i\in P}$ when P is the set of participating bidders.[30] As before, F^n represents the joint distribution of values conditional on there being n entrants, but unconditional on the set of participants. Therefore, $F^n(v_1 \dots v_n) = \sum_{P_n \subset \mathbb{N}} \Pr(\mathcal{P} = P_n | N = n) F^{P_n}(v_1 \dots v_n)$. For $P' \subset P$, let $F^{P'|P}$ denote the joint distribution of $(V_i)_{i\in P'}$ in auctions where P is the set of participants. Let F^P_m denote the joint distribution of values of m bidders drawn uniformly at random without replacement from P, when the set P enters, and let $F^n_m(v_1 \dots v_n) = \sum_{P_n \subset \mathbb{N}} \Pr(\mathcal{P} = P_n | N = n) F^{P_n}_m(v_1 \dots v_n)$.

We consider a subset of bidders to be of the same *type* if they are exchangeable, in the sense that $F^n(v_1, \dots, v_n) = F^n(v_{\sigma(1)}, \dots, v_{\sigma(n)})$ for any permutation $\sigma : \{1, \dots, n\} \to \{1, \dots, n\}$ and any $(v_1, \dots, v_n)$. Let $\Pr(P_n | P_{n+1})$ denote the probability that P_n would be obtained by dropping a bidder at random from P_{n+1}.[31]

Definition 1. ***Valuations are independent of supersets*** *if for all* $P' \subset P$, $F^{P'|P} = F^{P'}$.

Definition 2. ***Bidder types are independent of*** N *if, for all* P_n, $\Pr(\mathcal{P} = P_n | N = n) = \sum_{P_{n+1} \supset P_n} \Pr(P_n | P_{n+1}) \Pr(\mathcal{P} = P_{n+1} | N = n+1)$.

These definitions describe different kinds of exogeneity. Definition 1 requires that conditional on some set of bidders participating, those bidders' values are independent of which other bidders participate (what Athey and Haile (2002) refer to as exogenous participation). Definition 2 requires that the distribution of participating bidder types in n bidder auctions is just like the distribution of participating bidder types in $n+1$ bidder auctions, with one bidder randomly removed. It restricts who participates, but not what their values are. Coey

[30] We adopt the convention that bidders are ordered according to their identities, i.e. if $P = \{2, 5, 12\}$ then F^P is the joint distribution of (V_2, V_5, V_{12}), rather than, for example, the joint distribution of (V_5, V_2, V_{12}).

[31] For example, consider a case with two types, H and L. Then $\Pr(\{2H, 2L\}|\{3H, 2L\}) = \frac{3}{5}$, $\Pr(\{3H, 2L\}|\{3H, 3L\}) = \frac{1}{2}$, etc. If instead each bidder is a distinct type, then for any n, $\Pr(P_n|P_{n+1}) = \frac{1}{n+1}$ for all n. To see this, fix P_n and note that for each $P_{n+1} \supset P_n$, P_n is obtained by dropping the bidder $P_{n+1} \backslash P_n$ from P_{n+1}. When bidders are dropped uniformly at random, this occurs with probability $\frac{1}{n+1}$.

et al. (2017) demonstrate that one immediate implication of this condition is that for any bidder type τ the expected fraction of bidders that are of type τ should be constant across realizations of N. The following proposition shows that together these conditions imply that valuations are independent of N. Consequently, evidence of dependence of valuations and N suggests either that valuations are not independent of supersets, or that bidder types are not independent of N.

Proposition 6. *If valuations are independent of supersets and bidder types are independent of N, then valuations are independent of N.*

Proof. The proof follows Coey et al. (2017), Lemma 4. It suffices to prove that $F_m^n = F_m^{n+1}$ for any $n \geq m$.

$$\begin{aligned}
F_m^n(v) &= \sum_{P_n} \Pr(\mathcal{P} = P_n | N = n) F_m^{P_n}(v) \\
&= \sum_{P_n} \sum_{P_{n+1} \supset P_n} \Pr(P_n | P_{n+1}) \Pr(\mathcal{P} = P_{n+1} | N = n+1) F_m^{P_n}(v) \\
&= \sum_{P_{n+1}} \sum_{P_n \subset P_{n+1}} \Pr(P_n | P_{n+1}) \Pr(\mathcal{P} = P_{n+1} | N = n+1) F_m^{P_n}(v) \\
&= \sum_{P_{n+1}} \sum_{P_n \subset P_{n+1}} \Pr(P_n | P_{n+1}) \Pr(\mathcal{P} = P_{n+1} | N = n+1) F_m^{P_n | P_{n+1}}(v) \\
&= \sum_{P_{n+1}} \Pr(\mathcal{P} = P_{n+1} | N = n+1) F_m^{P_{n+1}}(v) \\
&= F_m^{n+1}(v)
\end{aligned}$$

The second equality follows because bidder types are independent of N. The fourth equality follows because $F^{P_n} = F^{P_n | P_{n+1}}$, as valuations are independent of supersets. The fifth equality follows because randomly selecting m bidders from $n+1$ bidders is the same as randomly selecting n bidders from $n+1$ bidders, and then randomly selecting m bidders from those n bidders. □

Monte Carlo Power Simulations

For some evidence on how powerful our test is, we compare it to another test, which simply compares bidders' mean values in n and $n+1$ bidder auctions. This latter test requires the econometrician to observe all bidders' values. Relative to our test based on the bidder exclusion effect, it requires more data, and does not allow for low bidding. Furthermore, this mean comparison test is not actually feasible in ascending auctions in practice given that the highest bid is never observed.

In our simulation, there are 10 potential bidders, who have iid lognormal private values drawn from $\ln N(\theta, 1)$, where θ is itself a random variable. All potential bidders see a common signal $\delta = \theta + \epsilon$, and Bayes update on the value of θ given their observation of δ. The random variables $(\delta, \theta, \epsilon)$ are jointly normally distributed:

$$\begin{pmatrix} \delta \\ \theta \\ \epsilon \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1+\sigma_\epsilon^2, & 1, & \sigma_\epsilon^2 \\ 1, & 1, & 0 \\ \sigma_\epsilon^2, & 0, & \sigma_\epsilon^2 \end{pmatrix} \right). \tag{32}$$

As σ_ϵ increases, the ratio of noise to signal increases, and the variable δ becomes less informative about the variable θ. To learn their value and bid in the ascending auction, potential bidders must pay an entry cost of 0.5. They play mixed entry strategies, entering with a probability p that depends on δ. Thus, this is a model of conditionally independent private values, where bidders do not know their own valuation until after paying an entry fee. In the limit as $\sigma_\epsilon \to \infty$, the signal δ is uninformative about θ and the entry probability p no longer varies with δ. This limiting case corresponds to the entry model of Levin and Smith (1994).

For each $\sigma_\epsilon \in \{1, 1.25, 1.5, \ldots, 7\}$, and for $n \in \{3, 4\}$, we generate 1,000 datasets with auctions in which n or $n+1$ bidders choose to enter. Each dataset contains 500 n bidder auctions and 500 $n+1$ bidder auctions. We calculate the probability of rejecting the null hypothesis of no dependence betweeen valuations and the number of bidders at the 5% level over the 1,000 datasets, for each value of σ_ϵ and n, and for both bidder exclusion test, and the comparison of means test. Figure 2 shows the rejection probabilities as a function of σ_ϵ. The comparison of means uses more data (in the case of Panel (A), all three bids from $n = 3$ auctions and all four bids from $n = 4$ auctions; and, in the case of Panel (B), all four bids from $n = 4$ auctions and all five bids from $n = 5$ auctions), and is more powerful. The simulation results suggest that when not all bidders' values are observed and the comparison of means test is infeasible (as in ascending auctions), the bidder exclusion based test is a reasonably powerful alternative, especially when the dependence between valuations and the number of bidders is stronger (corresponding in this model to low values of σ_ϵ.)

FIGURE 2. Monte Carlo Power Comparison

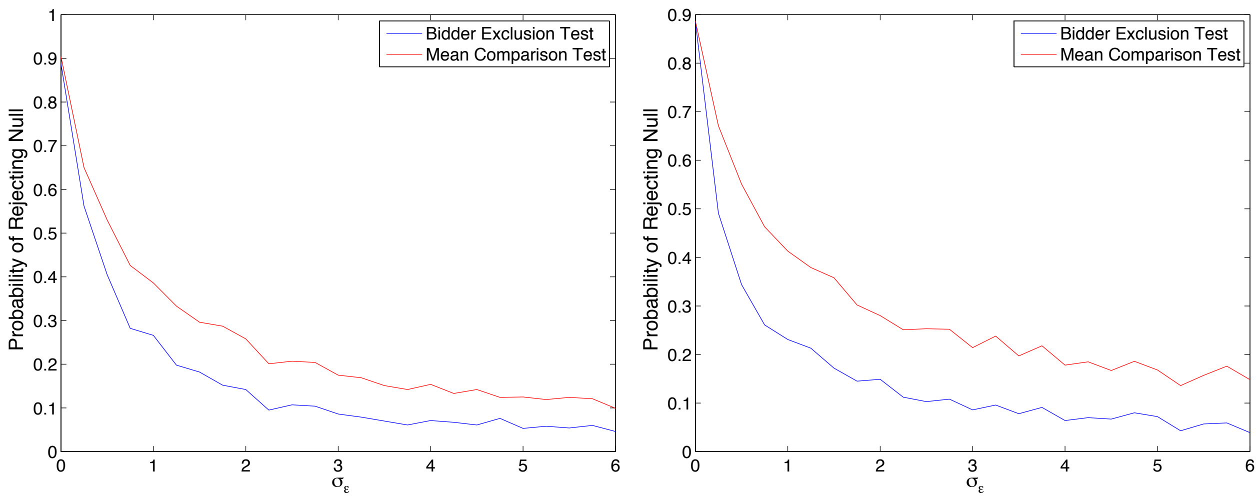

(A) 3 and 4 Bidder Auctions (B) 4 and 5 Bidder Auctions

Notes: Figures show the simulated probability of rejecting the null hypothesis of no selective entry for various levels of entry selectiveness (with greater σ_ϵ corresponding to less selective entry), for two tests: one based on a comparison of means between n and $n+1$ bidder auctions, and one based on the bidder exclusion effect computed on n and $n+1$ bidder auctions. The left panel shows the case of $n = 3$, and the right panel shows the case of $n = 4$.