Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Information Transparency in Business-to-Business Auction Markets: The Role of Winner Identity Disclosure

Yixin Lu, Alok Gupta, Wolfgang Ketter, Eric van Heck

informs
http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–19
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Information Transparency in Business-to-Business Auction Markets: The Role of Winner Identity Disclosure

Yixin Lu,[a] Alok Gupta,[b] Wolfgang Ketter,[c,d] Eric van Heck[d]

[a] School of Business, George Washington University, Washington, District of Columbia 20052; [b] Carlson School of Management, University of Minnesota, Minneapolis, Minnesota 55455; [c] Faculty of Management, Economics and Social Sciences, University of Cologne, 50931 Cologne, Germany; [d] Rotterdam School of Management, Erasmus University, Rotterdam, 3062 PA, Netherlands

Contact: yixinlu@gwu.edu, http://orcid.org/0000-0001-6515-1940 (YL); alok@umn.edu, http://orcid.org/0000-0002-2097-1643 (AG); ketter@wiso.uni-koeln.de, http://orcid.org/0000-0001-9008-142X (WK); evanheck@rsm.nl, https://orcid.org/0000-0002-8058-6640 (EvH)

Received: September 30, 2015
Revised: January 29, 2017; November 22, 2017; April 9, 2018
Accepted: May 28, 2018
Published Online in Articles in Advance: May 1, 2019

https://doi.org/10.1287/mnsc.2018.3143

Abstract. One of the central issues in auction design is how much information should be disclosed to bidders. In this paper, we examine bidder's identity disclosure in sequential business-to-business (B2B) auctions. Specifically, we compare two information disclosure policies, one that publicly discloses winners' identities (the status quo) and an alternative policy that conceals winners' identities. Using a large-scale field experiment in the Dutch flower auction market, we find that concealing winners' identities can significantly increase the average winning price and thereby raise the seller's revenue. We further explore the underlying mechanism that drives the observed effect. The empirical analysis of bidding behavior in these auctions suggests that bidders tend to imitate some of their competitors who have won in previous rounds of auctions and shade their bids accordingly. Concealing winners' identities can disrupt such imitation heuristic, which in turn mitigates the price-declining trend in sequential rounds. Our findings have important implications for the design of information disclosure policies in B2B auction markets.

History: Accepted by Chris Forman, information systems.

Keywords: auction design • bidding heuristics • field experiment • information transparency • identity disclosure • sequential auctions • tacit collusion

1. Introduction

Auctions account for an enormous volume of economic activities in both the public and private sectors, including the sales of mineral rights, Treasury bills, spectrum licenses, artwork, flowers, and real estate, and the procurement of construction contracts, office equipment, food supplies, and transportation services (Klemperer 1999). One of the central issues in auction design is how much bid information—for example, bidding prices (Arora et al. 2007), bidders' identities (Marshall and Marx 2009), and bid states (Adomavicius et al. 2012)—to disclose to bidders. Undoubtedly, the information disclosure policy for any real-world auction market must account for a variety of considerations such as efficiency, revenue, susceptibility to collusion, and risk of corruption. For example, the Federal Communications Commission (FCC) has introduced anonymous bidding to mitigate tacit collusion (Cramton and Schwartz 2000). In a similar vein, many European Union countries have adopted sealed bidding in government procurement auctions (Haberbush 2000).

In this paper, we study the information disclosure problem in the context of the world's largest wholesale market of cut flowers—namely, the Dutch flower auctions (DFAs). The DFA features a dynamic, complex business-to-business (B2B) market where trades between sellers (growers) and buyers are facilitated through multi-unit, sequential Dutch auctions. Buyers in these auctions could be retailers (e.g., florists), regional wholesalers, or multinational wholesalers. Whereas retailers are typically serving distinct market segments (and thereby do not compete directly), wholesalers often compete among themselves in the (postauction) resale market.

Traditionally, each bidder in the DFA is assigned a unique identity, which is often used by the same bidder for years, and the winner's identity is publicly disclosed after each sale. Such a policy has two conflicting effects on bidding outcomes. On one hand, the public disclosure of winners' identities may lead to lower prices, as it increases the auctions' susceptibility to collusion (Bajari and Yeo 2009) and reduces the level of competition. On the other hand, disclosing winners' identities may lead to higher prices through the reinforcement of the opponent effect,[1] which drives losing bidders to bid more aggressively in subsequent rounds. Whether the status quo policy in the DFA should be maintained is a question that has been debated for years in the DFA consortium.

To address this question, we conducted a large-scale field experiment to compare the performance of the status quo policy, which discloses winners' identities publicly, with an alternative policy that conceals identities. Using the difference-in-differences (DID) estimation approach (Bertrand et al. 2004), we are able to quantify the effect of the policy change on auction outcomes. Specifically, we find that the average winning prices increased by more than 6% when winners' identities were hidden from public view. This finding holds for various model specifications and robustness tests. Furthermore, concealing winners' identities also increased the stability of prices.

Several potential mechanisms may drive the price up when winners' identities are withheld. Drawing on the B2B nature of these auctions, we consider two mechanisms that may explain the effect: the mitigation of tacit collusion and the disruption of imitative bidding. Based on our analyses of bidders' dynamic interaction patterns, the explanation that is most consistently supported for why withholding winners' identities leads to a price increase is that bidders tend to imitate some of their competitors who have won in previous rounds and shade their bids accordingly. Once winners' identities are concealed, it becomes difficult for bidders to identify the competitors they chose to imitate. As a result, bidders are forced to use the most recent bids as reference points ("anchors") when making their bidding decisions.

Our paper makes several contributions to both the theory and practice of auction design. First, we contribute to the literature on information disclosure by examining the role of identity disclosure in sequential auctions. In the traditional market design literature, participants' identities are typically assumed to play a minimal role in determining the outcome of the exchanges. However, with the proliferation of online markets, this is no longer the case: as a result of the lack of face-to-face contact, participants in online markets have to establish themselves in a formal manner as legitimate (Smith 2007). As such, participants' identities will inevitably shape the strategic interactions and the market outcome (Varma 2002b). To the best of our knowledge, this is the first empirical study that systematically examines the problem of identity disclosure in sequential auctions. Second, the exploration of the underlying mechanism that drives the differences in bidding dynamics and outcomes under the policy change sheds light on the information transmission and learning aspects (Jeitschko 1998) in sequential auctions and thus improves our understanding of the bidding dynamics in these complex auctions. Third, unlike prior studies that rely on stylized models or laboratory experiments with inexperienced bidders (see Arora et al. 2007 and Cason et al. 2011, for example), we empirically measure the impact of alternative disclosure policies in a complex, dynamic B2B market through a large-scale field experiment, which allows us to maintain a high level of control while accounting for the richness of the real-world operation environment. Therefore, our findings provide actionable insights into the practical design. Specifically, contrary to the popular view "the more information the better" in market design, our results suggest that sellers do not necessarily benefit from the commitment to disclosing all the available information and highlight the behavioral aspects of different disclosure policies in shaping bidders' dynamic interactions. To this point, it is worth noting that, although we primarily focus on the DFA, our findings are also useful for policy makers in other B2B markets.

The rest of this paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the empirical setting. Section 4 describes the data used in the empirical analysis and provides descriptive statistics. Section 5 provides the details of our empirical analyses and results. Section 6 discusses the potential mechanisms behind the empirical findings and offers an explanation to the observed effect. Finally, Section 7 discusses the contributions and implications, reflects on the limitations, and suggests directions for future work.

2. Related Literature

Our paper draws on two streams of literature: economic implications of identity disclosure and bidding dynamics in sequential auctions.

2.1. Identity Disclosure

The public disclosure of bidders' identities has conflicting effects on bidding dynamics and outcomes in auction markets. On one hand, identity disclosure may lead to lower prices and revenue. Specifically, prior literature has shown that identity disclosure increases the risk of collusion among bidders (Cramton and Schwartz 2000, Bajari and Yeo 2009, Marshall and Marx 2009) and thereby reduces competition. The rationale is as follows: if bidders' identities are publicly observable during an auction, any deviation from the collusive agreement among cartel members can be easily identified, and the punishment can be directed to the defecting member. In addition, disclosing bidders' identities may also create a screening effect (Snir and Hitt 2003, Hong et al. 2016), which discourages participation and lower prices. This may explain why some of the popular online consumer auction markets such as eBay have moved to a less transparent setting where bidders' identities are no longer publicly disclosed. On the other hand, identity disclosure may increase bidding prices (and thus benefit sellers) by reinforcing the opponent effect (Heyman et al. 2004). In particular, disclosing winners' identities can increase losing bidders' frustration by making the social component in the bidding competition more salient, thereby causing them to overbid in future auctions (Delgado et al. 2008).

When auctions are followed by interactions in the downstream market, bidders may experience identity dependent externalities; that is, conditional on losing the auction, a bidder's payoff is dependent on the identity of the winner (Varma 2002a). Identity-dependent externalities arise in a variety of economic contexts such as the sale of patents, the change of ownership (e.g., mergers and acquisitions, privatizations), and the sale of spectrum licenses (Jehiel and Moldovanu 1996). In general, when bidders experience identity-dependent externalities, the price implication of identity disclosure is ambiguous (Aseff and Chade 2008). Consider the case of the DFA, for example. Because wholesalers often compete in the downstream retail market after the auction, a losing wholesaler would bid more aggressively after learning that the current winner is her direct competitor. By contrast, if the current winner is a retailer, wholesalers may bid less aggressively in subsequent auctions. Nevertheless, Varma (2002b) has shown that for a large set of standard auction mechanisms, disclosing bidders' identities can increase bidding prices by helping bidders form a more accurate estimate of their ex post willingness to pay.

At the outset, the result from Varma (2002b) shares the same spirit of the linkage principle (Milgrom and Weber 1982): sellers can raise the expected revenue by disclosing all available information to bidders. Despite its general acceptance as a guide to auction design, the linkage principle typically does not hold beyond single-unit auctions (Perry and Reny 1999). Specifically, depending on the underlying assumptions, the linkage principle may or may not hold in sequential auctions. For example, Arora et al. (2007) find that a complete revelation policy that announces both winning and losing bids generates higher buyer surplus than a partial revelation policy that announces winning bids alone in a two-stage procurement setting where bidders face market-structure uncertainty. By contrast, Tu (2005) reports that announcing winning bids alone yields higher revenue than announcing both winning and losing bids in a two-period first-price auction with two risk-neutral bidders. Bergemann and Hörner (2017) extend Tu's study and show that the minimal disclosure policy (i.e., where each bidder only learns whether he wins or loses privately at the end of each round) generates higher expected revenue than the complete revelation policy and the partial revelation policy in repeated first-price auctions with a fixed number of bidders. The mixed findings from the existing literature suggest that the choice of a disclosure policy for a real-world auction market must be tailored to the empirical details of the environment.

2.2. Learning in Sequential Auctions

Compared with single-unit auctions, sequential auctions allow bidders to learn about the market trend from previous rounds of competition (Goes et al. 2010). This also makes the analysis of these auctions much more complicated (Klemperer 1999, Overby and Kannan 2015). Within the symmetric independent private value (IPV) paradigm, Weber (1983) and Donald et al. (2006) demonstrate that the equilibrium price is nondecreasing over time. Unfortunately, such theoretical results are not supported by empirical findings.

Specifically, prior studies have detected price declines in sequential auctions of various products including artworks, flowers, and wine (McAfee and Vincent 1993, Beggs and Graddy 1997, van den Berg et al. 2001). So far, researchers have offered different explanations for the declining price trend in sequential auctions, which can be broadly cast into two strands. The first strand of work attributes the price decline to bidder heterogeneity regarding risk profiles (McAfee and Vincent 1993). The second strand emphasizes the information transmission and learning in sequential rounds. Specifically, Jeitschko (1998) points out that if a bidder thinks there is a positive probability that another bidder has the same valuation for the object under auction, the information transmission has both a direct and an indirect effect on bidders' strategies. The direct effect refers to the belief updating about the probability that another bidder has the same valuation. The indirect effect refers to the trade-off between the benefits associated with winning in early rounds and the benefits of losing but learning more about competitors' valuations.

In the current paper, we draw on the learning aspects in exploring the underlying mechanism that drives the observed differences between the two information disclosure policies. In this regard, our paper is related to Cason et al. (2011), where the authors conduct a laboratory experiment to study the trade–off between the bidders' desire to learn (i.e., learning effect) and their desire to prevent their opponents from learning (i.e., deceptive effect). Cason et al. show that the perceived degree of competition can moderate the impact of different information revelation policies on the market outcome when bidders account for such a learning-related trade-off. Our paper differs from Cason et al. (2011) in two important ways. First, we use a field experiment to explore the learning aspects under different policies without imposing any restrictive assumptions about bidders' valuations (e.g., whether bidders' valuations are affiliated) or market dynamics (e.g., whether winners drop out with certainty or not), whereas Cason et al. consider a two-stage game where bidders' private costs are drawn from a discrete distribution of two types and winners do not drop out.

Second, our empirical setting features an information-rich and time-critical B2B market with highly experienced bidders. This allows us to provide more actionable insights to practitioners regarding the choice of information disclosure policies.

3. Research Setting

3.1. Research Context

We examine the identity disclosure problem in the context of the DFA. The DFA serve as efficient trading centers for cut flowers and potted plants, generating an annual turnover of over 4 billion euros.[2] The DFA use a sequential Dutch auction mechanism, which is implemented using a single-handed clock displayed on an electronic board. The clock initially points to a high price and then quickly ticks down in a counterclockwise direction until a bidder accepts the current price by stopping the clock and purchasing all or part of the lot[3] under auction. In case part of the lot is unsold, the clock will restart at a high price and the process is repeated.

Apart from the current asking price, each clock also contains information about the setup of the current auction (e.g., monetary unit, minimum purchase units, bundling properties as well). Furthermore, bidders can see the information of the product under auction (e.g., the name of the product, the identity of the grower, a representative picture of the product) from the electronic board. Figure 1 provides an illustration of the clock interface.

Auctioneers in the DFA represent the growers. Thus, the primary goal of their work is to maximize the revenue. At the beginning of an auction, the auctioneer decides the starting position of the clock and sets the clock in motion. As the clock ticks down counterclockwise, each bidder can stop the clock by pressing a button indicating that she is willing to accept the price corresponding to the current clock position. The first bidder who makes a bid wins. The winning bidder, whose identity is displayed on the clock screen, can select the purchase quantity (which must exceed the minimum required amount). If the winning bidder does not select the entire amount available, the clock ticks backward and restarts at a high price, and the auction continues. This process repeats until all of the products are sold or the price falls below the seller's reserve price,[4] in which case any unsold goods in that lot are destroyed. Auctioneers can influence the dynamic competition of these auctions by controlling the key auction parameters (e.g., starting prices, minimum purchase quantities, reserve prices) as well as the information disseminated during the sequential rounds.

In recent years, as more and more bidders have chosen to participate in the auctions via the online channel, the information disclosure decision has become increasingly important in the operationalization of these auctions. On one hand, the online channel allows more bidders to participate in the auctions and significantly increases the market-level uncertainty. On the other hand, the increased transparency resulting from the high adoption rate of the online channel raises several concerns from both suppliers and buyers, the most salient of which is that it enables bidders to easily track their competitors' bids and adjust their bidding strategies accordingly. A straightforward way to address this concern is making the information about

Figure 1. (Color online) Illustration of the Auction Screen

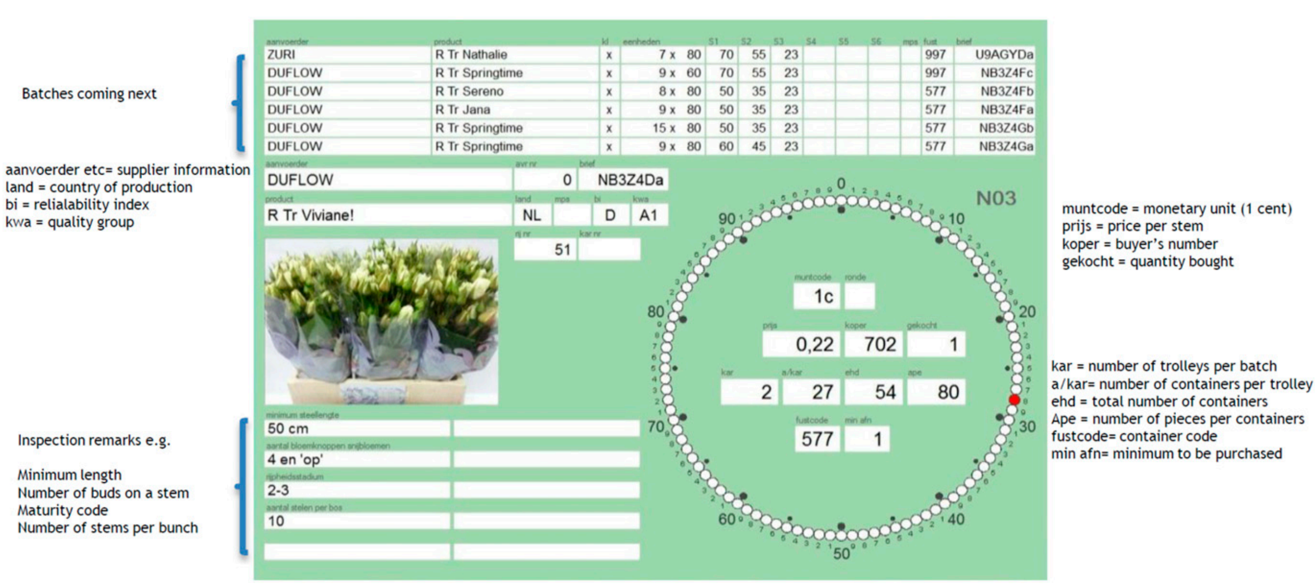

Note. The setup of the current auction is shown on the clock, the product information is shown to the left of the clock, and the upcoming schedules are shown on the top of the screen.

"who bought what at which clock" confidential—that is, removing the winner's identity from the auction clock. However, it is unclear whether and how such a change would affect the market outcome. Specifically, although concealing winners' identities can discourage the potential collusion among bidders, it may also lower bidders' willingness to pay by preventing them from forming an accurate estimate of the expected payoff in the presence of identity-dependent externalities (Varma 2002b) or mitigating the opponent effect (Heyman et al. 2004). The existence of these conflicting forces calls for an empirical examination to reveal whether a less transparent policy (i.e., a policy of identity concealment) should be adopted.

3.2. Research Design

The data for this study were collected through a large-scale field experiment[5] during the last quarter of 2012 in the DFA. The treatment site was chosen randomly among the four major auction sites, and the policy change with respect to winner's identity disclosure was implemented at a clock that auctioned chrysanthemums, the flowers in season. The experiment lasted from November 19 to December 7, during which time the winner's identity was removed from the clock screen.[6] It should be noted that whereas none of the bidders except the winner knew who had won in each round, we as researchers could see the winners because that information was registered in the auctioning system and recorded in the logbook.

The ability to attribute potential changes in market performance to the policy change of identity disclosure requires more than a simple before–after research design. In particular, it is necessary to demonstrate that the changes did not happen simply as a result of systematic changes in supply or demand over time. Fortunately, we were able to obtain data from a control site where the same type of flower was auctioned and bidders could observe winners' identities throughout the study period. Figure 2 provides an overview of our research design. The inclusion of data from the post-experiment period allows us to show whether the treatment effect dissipated once the treatment was switched off (Kumar and Tan 2015).

The treatment site and control site are approximately 60 kilometers apart; each serves a large buyer population (both with three auction halls that can accommodate a total of 450 bidders at the same time). Furthermore, the two sites have three key commonalities. First, they use the same auction format and payment rules, and the suppliers of products (cut flowers and potted plants) are largely the same as well. Second, both sites provide high-quality transport and delivery services. Theoretically speaking, bidders can

Figure 2. (Color online) Overview of the Research Design

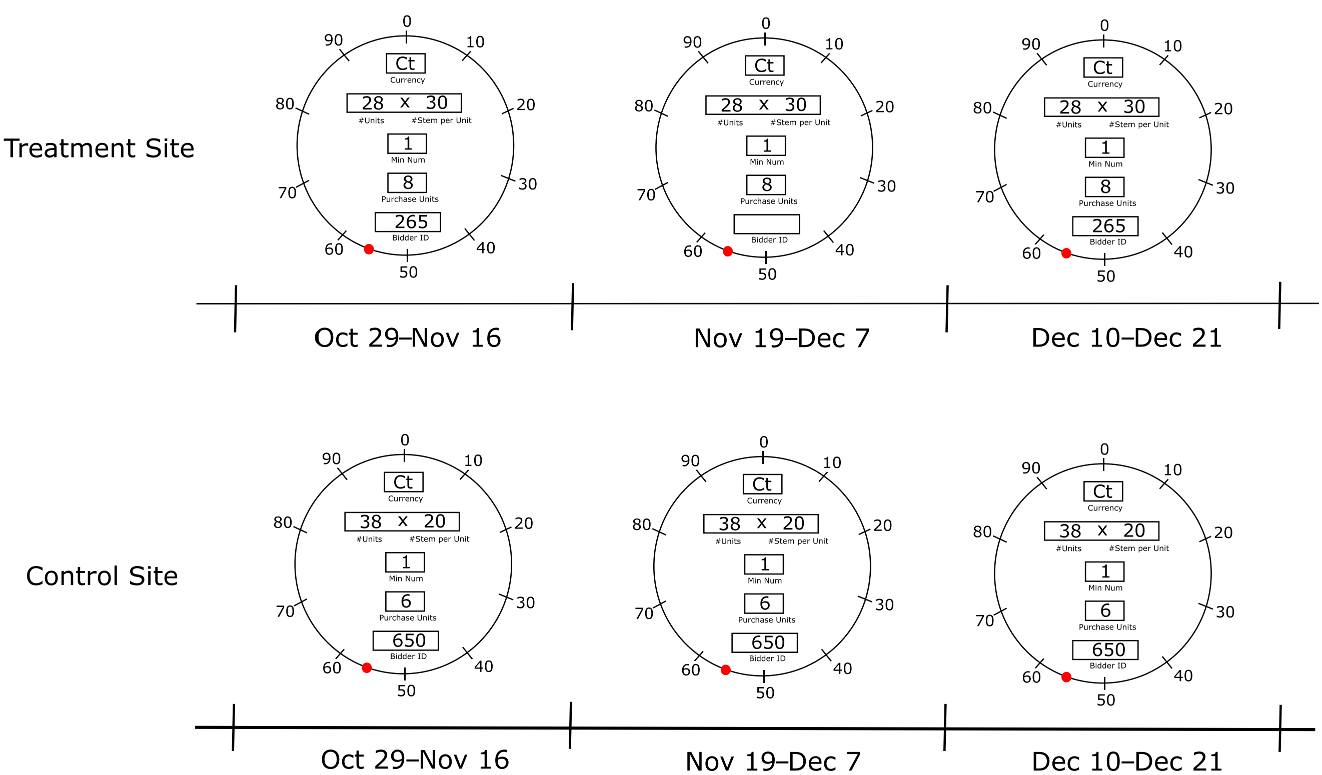

Note. In addition to data from the experiment period (November 19–December 7), we also obtained data before (October 29–November 16) and after (December 10–December 21) the experiment period.

Table 1. A Sample Entry in a Logbook

Transaction time	Seller ID	Flower ID	Stems per unit	Available units	*Minimum Purchase Units*	*Starting Price* (¢)	Buyer ID	Purchase units	Price (¢)	Online
07:10:54	5644	182	50	18	1	100	439	1	30	Yes
07:10:56	5644	182	50	17	3	42	395	5	29	No
07:10:57	5644	182	50	12	4	41	601	8	26	Yes
07:10:59	5644	182	50	4	4	38	563	4	29	Yes

Note. Auctioneer's decision variables are italicized.

make a purchase from any auction site. However, in practice, bidders, especially the large wholesalers, often choose to buy from the auction site closest to their distribution centers.[7] This observation helps to alleviate the concern of selection bias associated with our quasiexperimental design. Third, the auctioneers at both sites have extensive experience in conducting these auctions, and there was no replacement or new hire during our study period. This enables us to disentangle the effect of the policy change from potential auctioneer effect (Lacetera et al. 2016).

4. Data

Our data set consists of 22 attributes, two of which are bidders' real-time decision variables: price and quantity. The rest can be classified into seven broad categories: (1) product characteristics (e.g., product type, stem length, bundling size, blooming scale, quality); (2) transaction timing (date and time); (3) supply-side information, which includes lot size and minimum purchase quantity; (4) the precise market actors (seller identity and buyer identity); (5) logistics (stems per unit, units per trolley, and number of trolleys); (6) bidding channel (online or offline); and (7) clock specification (e.g., clock stand, currency unit).

Table 1 provides a stylized example of a sequence of transactions from our data set. Because of space constraints, we do not include all 22 attributes but a set of representative attributes. In this example, a lot containing 18 units gets sold. Note that the sales prices are not monotonically decreasing or increasing.[8] Also, unlike existing studies, which focus on the situation where only one unit gets sold in each round, in our case the purchase quantity in each round can vary significantly.

To control for product heterogeneity, we selected the product group with the highest transaction amount—namely, chrysanthemum spray white/yellow. Because there were many growers for this product, we only included lots from 18 of them who were selling the particular flower at both sites every day during the study period. This left us a total of 31,848 transactions, with 14,570 from the treatment site and 17,278 from the control site. It is worth mentioning that all products in this sample were rated as of the highest quality (level "A").

We performed a preliminary analysis to examine the market-level characteristics of both auction sites. The results are reported in Table 2. For brevity, we use *Pre*, *Exp*, and *Post* to denote the pre-, during-, and post-experiment period. Both the average number of auctioned lots (per week) and the average number of winning bidders (per week) were quite stable at the two sites. Additionally, we can see that the treatment site had a higher rate of online transactions (i.e., the transactions where the corresponding winning bidders participated in the auctions via the online channel) during the study period.

At the auction level, we found that the average number of winning bidders per auction was quite stable at both sites throughout the study period. By contrast, the mean and standard deviation of the winning price varied significantly ($p < 0.05$) during the experiment period. Specifically, the average price increased by approximately 18% at the treatment site and 11% at the control site. Tables 3 and 4 summarize the auction-level descriptive statistics from the treatment site and control site, respectively.

As we mentioned above, bidders are free to enter any auction in person (via the offline channel) or remotely (via the online channel). Thus we would like to know whether bidders had switched between the two auction sites during the study period or whether they were simultaneously bidding across both sites. Note that

Table 2. Market-Level Characteristics at Treatment Site and Control Site

	Treatment site			Control site		
	Pre	*Exp*	*Post*	*Pre*	*Exp*	*Post*
Average number of auctions (per week)	273	283	278	290	286	296
Average number of winning bidders (per week)	265	254	250	296	285	286
Usage of online channel (%)	92	93	92	75	78	76

Table 3. Descriptive Statistics at Auction Level from the Treatment Site

	No. of winning bidders			Price (¢)			Purchase units		
Statistics	*Pre*	*Exp*	*Post*	*Pre*	*Exp*	*Post*	*Pre*	*Exp*	*Post*
Mean	6.5	6.4	6.1	26.7	31.6	25.1	11.8	11.0	11.3
Median	6.0	6.0	6.0	25.0	31.0	23.0	5.0	5.0	5.0
Standard deviation	4.2	4.1	3.9	9.3	7.1	9.7	19.8	16.7	15.9
Minimum	1	1	1	1	10	1	1	1	1
Maximum	31	32	19	62	66	78	347	264	180

Table 4. Descriptive Statistics at Auction Level from the Control Site

	No. of winning bidders			Price (¢)			Purchase units		
Statistics	*Pre*	*Exp*	*Post*	*Pre*	*Exp*	*Post*	*Pre*	*Exp*	*Post*
Mean	7.0	7.5	7.0	29.6	33.0	28.2	10.1	9.4	10.0
Median	6.0	7.0	7.0	29.0	33.0	25.0	4.0	4.0	4.0
Standard deviation	4.9	5.3	4.5	10.6	8.4	13.0	18.7	17.3	19.1
Minimum	1	1	1	5	5	5	1	1	1
Maximum	33	34	27	70	87	81	371	384	396

although bidder IDs are unique in the sense that each ID at one site is typically owned and used by the same company over many years, we cannot directly compare them across different sites, because (1) one company may own several IDs across different auction sites, and (2) the same ID number across different sites may refer to different companies. In light of this, we requested the list of registered buyers (companies) as well as the latest allocation information of bidder IDs at the treatment and control site from the market maker. This allows us to map the bidder IDs observed in the transaction data from both sites to the companies. After cross-checking the winning identities and their corresponding companies, we did not find any company that had participated at both sites during the study period. However, given that only winning bids are observable in a Dutch auction, we cannot rule out the possibility that some losing bidders might have participated in the auctions across the two sites. Nevertheless, the presence of these bidders is not a concern regarding the data-generating process (Donald et al. 2006).

At the outset, these aggregate-level results suggest that withholding winners' identities has a positive impact on the auction prices. However, as we explained in the research design, such before–after comparison does not control for any potential systematic changes in the auction market. For example, it may be that there was a higher demand during the experiment period.

5. Empirical Analyses and Results

5.1. The Empirical Strategy

We used a difference-in-differences (DID) approach (Bertrand et al. 2004) to estimate the effect of the policy change in identity disclosure. By measuring the difference in differences between the treatment site and the control site over time, we can control for the characteristics that are unobservable to researchers but may impact the market processes and outcomes (Imbens and Wooldridge 2009). Note that the policy change was introduced as an exogenous shock at the treatment site from November 19 to December 7. The DID approach accounts for differences in baseline levels of market performance across the two auction sites and adjusts for any potential differences that may arise as a result of market trends at both sites during the study period.

To quantify the change in auction price when winners' identities were withheld as opposed to publicly disclosed, we first estimated a fixed effect log-linear model at the transaction level (baseline) as follows:

$$\begin{aligned} \ln(Price_{i,j,t}) = {} & \beta_0 + \beta_1 Treatment_{i,j} \times Experiment_t \\ & + \beta_2 Treatment_{i,j} \times Post_t + \beta_3 Treatment_{i,j} \\ & + \theta_1 Experiment_t + \theta_2 Post_t + \gamma \mathbf{X}_{i,j,t} + \epsilon_{i,j,t}. \end{aligned} \tag{1}$$

In Equation (1), $\ln(Price)$ is the natural log of the winning price; i indexes each transaction, j indexes each auction, and t indexes the time (week) in the study period. The variables $Treatment_{i,j}$, $Experiment_t$, and $Post_t$ are dummy variables: $Treatment_{i,j}$ equals 1 if the transaction was from the treatment site, $Experiment_t$ equals 1 during the experiment period and 0 otherwise, and $Post_t$ equals 1 during the post-experiment period and 0 otherwise; $\mathbf{X}_{i,j,t}$ is a vector of control variables. Specifically, to account for product heterogeneity, we included the key product characteristics *StemLength*, *BundlingCondition*, *BloomingStage*, and 17 grower dummies. To account for variation from the supply side, we included *LotSize* and *MinimumPurchaseQuantity*. In addition, we controlled for the day of the week for each transaction, which may have implications for the market demand. The error term $\epsilon_{i,j,t}$ reflects the idiosyncratic variation in potential outcomes (i.e., log prices) that varies across transactions, auction lots, and time. Our coefficient of interest is β_1. It captures the difference in the log winning price between the treatment site and control site during the policy change at the treatment site.

Note that we can also replace the fifth and sixth terms in Equation (1) (i.e., $Experiment_t$ and $Post_t$) by a set of more detailed week fixed effects θ_t. In essence, these week fixed effects serve as nonparametric controls for the temporal variation in winning prices that is common across both sites. This results in the following alternative model (time fixed effect model):

$$\ln(Price_{i,j,t}) = \beta_0 + \beta_1 Treatment_{i,j} \times Experiment_t + \beta_2 Treatment_{i,j} \times Post_t + \beta_3 Treatment_{i,j} + \theta_t + \gamma \mathbf{X}_{i,j,t} + \epsilon_{i,j,t}. \tag{2}$$

Before proceeding to the estimation results, we would like to discuss the assumption about the variance–covariance matrix of the error term $\epsilon_{i,j,t}$ in Equations (1) and (2). One possibility would be to assume these errors are independent and identically distributed (IID). Unfortunately, a Breusch–Pagan test rejects the hypothesis that errors are homoskedastic across auctions ($p < 0.001$), and a Breusch–Godfrey test rejects the hypothesis of no first-order autocorrelation ($p < 0.001$). To address the serial correlation problem (Bertrand et al. 2004), we used heteroskedastic-robust standard errors clustered at the auction level (Angrist and Pischke 2008). Clustering at the auction level allows for heteroskedastic errors across auctions as well as arbitrary correlation of errors within auctions.

5.2. Effect of Withholding Winner's Identity on Auction Price

Table 5 summarizes our main results. For the baseline model (column (1)), we can see that the coefficient of the interaction between the treatment site and the experiment period is positive and significant (coefficient = 0.065, $p = 0.001$). This suggests that the average winning price was indeed higher when winners' identities were concealed from public view at the treatment site. When the policy change was revoked, the difference between the average winning price from the treatment site and control site is no different from what it was before the policy change. This is indicated by the coefficient of the interaction between the treatment site and the post-experiment period (i.e., $Treatment \times Post$), which is statistically insignificant ($p = 0.592$).

Furthermore, we can see that all control variables except *LotSize* have a significant impact on the winning price. Specifically, the minimum purchase quantity has a negative effect. This is consistent with the findings of Lu et al. (2016). Also, the day-of-the-week effect is quite salient: on average, the winning prices from Tuesday to Friday were significantly lower than those on Monday. Finally, although the results from the preliminary analysis (see Tables 3 and 4) suggests there might be a difference in the average winning price between the treatment and control site, such a difference turns out to be statistically insignificant after we control for the product heterogeneity and day-of-the-week effect. The results from the time fixed effect model (column (2)) are

Table 5. Average Effect of Withholding Winner's Identity on Winning Price

Variable	(1) Baseline	(2) Time fixed effect
Treatment × *Experiment*	0.065**	0.062**
	(0.021)	(0.019)
Treatment × *Post*	0.015	0.016
	(0.029)	(0.028)
Treatment	0.043	0.031
	(0.053)	(0.052)
Experiment	0.137***	
	(0.015)	
Post	−0.091***	
	(0.021)	
LotSize	0.002	0.002
	(0.002)	(0.002)
MinimumPurchaseQuantity	−0.025***	−0.024***
	(0.001)	(0.001)
StemLength	−0.018*	−0.024**
	(0.009)	(0.008)
BundlingCondition	0.003**	0.003**
	(0.001)	(0.001)
BloomingStage	0.081***	0.084***
	(0.009)	(0.010)
Tuesday	−0.037*	−0.037*
	(0.014)	(0.014)
Wednesday	−0.082***	−0.082***
	(0.015)	(0.014)
Thursday	−0.112***	−0.114***
	(0.015)	(0.014)
Friday	−0.052***	−0.056***
	(0.014)	(0.014)
Grower fixed effects	Yes	Yes
Week fixed effects		Yes
Observations	31,848	31,848
Adjusted R^2	0.337	0.409

Notes. All coefficients are estimated at the transaction level. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions.
$^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

qualitatively similar to those from the baseline model. Most important, the average winning price increased significantly (coefficient = 0.062, $p = 0.001$) at the treatment site during the experiment period. By contrast, the difference in winning prices between the treatment and control site during the post-experiment period was no different relative to the pre-experiment period.

We also considered two alternative specifications, both at the auction level, to examine the effect of policy change on the winning price:

$$\ln(AvgPrice_{j,t}) = \beta_0 + \beta_1 Treatment_j \times Experiment_t + \beta_2 Treatment_j \times Post_t + \beta_3 Treatment_j + \theta_t + \gamma \mathbf{X}_{j,t} + \epsilon_{j,t}, \quad (3)$$

$$\ln(WeightedAvgPrice_{j,t}) = \beta_0 + \beta_1 Treatment_j \times Experiment_t + \beta_2 Treatment_j \times Post_t + \beta_3 Treatment_j + \theta_t + \gamma \mathbf{X}_{j,t} + \epsilon_{j,t}, \quad (4)$$

where $AvgPrice_{j,t}$ is the simple average of the winning prices in auction j at time t, and $WeightedAvgPrice_{j,t}$ is the weighted average that accounts for the relative purchase amount associated with each transaction. Table 6 summarizes the estimation results of these alternative model specifications. To begin with, we can see that the treatment effect remains positive and significant (i.e., the average winning price increased by approximately 7% at the treatment site during the experiment period). The effects of the control variables are qualitatively similar to those observed from Table 5, except for *LotSize* and *StemLength*. The former exhibits a significant effect on the winning price based on the alternative model specifications, although such an effect is of negligible magnitude (0.3%); the latter has no significant effect on the (weighted) average price at auction level.

Table 6. Estimation Results with Alternative Model Specifications

	Auction-level time fixed-effects model	
Variable	(1) Average price	(2) Weighted average price
Treatment × *Experiment*	0.070***	0.073***
	(0.018)	(0.018)
Treatment × *Post*	0.024	0.027
	(0.024)	(0.024)
Treatment	0.033	0.039
	(0.045)	(0.046)
LotSize	0.003*	0.003*
	(0.001)	(0.001)
StemLength	0.004	0.004
	(0.006)	(0.006)
BundlingCondition	0.010***	0.010***
	(0.001)	(0.001)
BloomingStage	0.055***	0.055***
	(0.006)	(0.006)
Tuesday	−0.034**	−0.040**
	(0.013)	(0.013)
Wednesday	−0.068***	−0.071***
	(0.013)	(0.013)
Thursday	−0.106***	−0.106***
	(0.013)	(0.013)
Friday	−0.062***	−0.064***
	(0.013)	(0.013)
Grower fixed effects	Yes	Yes
Week fixed effects	Yes	Yes
Observations	4,545	4,545
Adjusted R^2	0.438	0.438

Notes. All coefficients are estimated at the auction level. Columns (1) and (2) summarize the estimation results from model (3) and model (4), respectively. Standard errors are shown in parentheses.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

5.2.1. The Parallel-Trend Assumption.

A critical assumption underlying the DID approach is that the difference between the treatment and control site would remain constant over time in the absence of the treatment (Abadie 2005, Angrist and Pischke 2008). If something other than the treatment changes at one site but not the other, the parallel-trend assumption would be violated, in which case we have no guarantee that the DID estimator described above is unbiased. With this in mind, we checked whether there were differential trends before the policy change at the treatment site and the control site.

Following Autor (2003), we used a leads and lags model—also referred to as a relative time mode (Greenwood and Wattal 2017)—to explore the pre-experiment time trends at the two sites. Specifically, we created two lead indicator variables corresponding to the first two weeks during the pre-experiment period and three lag indicator variables corresponding to the three weeks during the experiment period. We estimated a model similar to Equation (2) except that $Experiment_t$ and $Post_t$ are replaced by the lead and lag indicator variables. In this case, the reference group is the week immediately before the experiment period (i.e., November 12–16). Table 7 summarizes the estimation results. Here, *Lead*(−3) and *Lead*(−2) denote the two lead indicators (i.e., week 1 and week 2 in the pre-experiment period), and *Lag*(0), *Lag*(+1), and *Lag*(+2) stand for the three lag indicators (i.e., weeks 1, 2, and 3 during the experiment period). We can see that the coefficients corresponding to the two lead indicators are not statistically significant, whereas the coefficients for the three lag indicators are all significant at the level of 0.001. These results indicate that the parallel trend assumption for our DID model is not violated.[9]

5.2.2. Robustness Checks.

We conducted a series of checks to assess the robustness of the observed effect. To start with, we noticed that not all bidders participated in the auctions every day during the study

Table 7. Estimation Results of the Pre-experiment Trend

Variable	Relative time model
Treatment × *Lead*(−3)	−0.003 (0.036)
Treatment × *Lead*(−2)	−0.012 (0.035)
Treatment × *Lag*(0)	0.060** (0.020)
Treatment × *Lag*(+1)	0.074** (0.022)
Treatment × *Lag*(+2)	0.067** (0.021)
Treatment	0.075 (0.062)
LotSize	0.004* (0.002)
MinimumPurchaseQuantity	−0.021*** (0.002)
StemLength	0.004 (0.009)
BundlingCondition	0.004*** (0.001)
BloomingStage	0.079*** (0.007)
Tuesday	−0.040** (0.014)
Wednesday	−0.086*** (0.015)
Thursday	−0.124*** (0.015)
Friday	−0.061*** (0.015)
Grower fixed effects	Yes
Week fixed effects	Yes
Observations	24,165
Adjusted R^2	0.448

Notes. All coefficients are estimated at the transaction level. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions.

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

period. This observation leads to the concern that the observed change in the winning price may be due to bidder heterogeneity across the three different periods. For example, if the bidders participating in the auctions at the treatment site during the experiment period happened to have much higher valuations than those participating in the auctions before or after the experiment period, the change in the winning price might be due to the idiosyncratic shocks in bidders' private valuations rather than the withholding of the winner's identity. Alternatively, if the bidders exposed to the treatment condition varied from week to week, the observed difference in the winning price may be due to the so-called Hawthorne effect or novelty effect (Adair 1984).

To investigate these alternative explanations to our main finding, we created two subsamples, which include only the transactions from bidders who had participated in the auctions throughout the eight-week study period at the treatment site and control site. As we did not find any bidders who switched from one auction site to the other during the study period (see the discussion in the preliminary analysis of the data in Section 4), we then merged the two subsamples and reestimated the baseline model in Equation (1), the time fixed effect model in Equation (2), and an augmented model by including the bidder fixed effect to Equation (2). If the price change was primarily driven by the Hawthorne effect or changes in the relative proportion of the bidder population (i.e., bidders with higher valuation purchased more during the experiment period at the treatment site), we would observe a significant reduction in the magnitude of the treatment effect. However, this is not the case: according to Table 8, the treatment effect remains positive and significant, and the estimated coefficients are qualitatively similar to the ones shown in Table 5. This finding rules out the aforementioned alternative explanations.

Next, because bidders in the online channel are likely to pursue different strategies as opposed to those bidding onsite (Lu et al. 2016), it is interesting to see whether the treatment effect varies across the market channels. Specifically, prior research argues that compared with online bidders, those bidding onsite can acquire additional market state information through implicit or explicit verbal communications in the auction hall (Koppius 2002). As such, they may have an informational advantage in the bidding competition. If this was the case, one would expect to see a significant difference in the magnitude of the treatment effect between onsite (offline) and online bidders. With this in mind, we reestimated Equation (2) on the subsamples that consist of transactions[10] from the offline and online channels. Table 9 summarizes the estimation results.

We can see that the treatment effect remains positive and significant, and the magnitude of the effect does not differ between the two channels. This finding indicates that, from the bidders' perspective, the loss of identity information in their decision-making process cannot be compensated by the additional market state information from the offline channel. The rest of the coefficient estimates for both subsamples are also largely consistent with the full sample estimates presented in Table 5.

As the third robustness check, we changed the benchmark from the pre-experiment period to the post-experiment period—the dummy variable $Post_t$ corresponding to the post-experiment period is replaced by a dummy variable Pre_t corresponding to the pre-experiment period—and replicated the DID estimations described above. If the observed price increase was mainly due to market shocks that coincided with the policy change rather than the policy change itself, the coefficient of the interaction between the treatment site

Table 8. Treatment Effect on Active Bidders

Variable	(1) Baseline	(2) Time fixed effect	(3) Time and bidder fixed effect
Treatment × *Experiment*	0.072***	0.067***	0.066***
	(0.020)	(0.019)	(0.018)
Treatment × *Post*	0.011	0.011	0.008
	(0.030)	(0.028)	(0.028)
Treatment	0.024	0.008	
	(0.052)	(0.051)	
Experiment	0.128***		
	(0.015)		
Post	−0.087***		
	(0.021)		
LotSize	0.002	0.001	0.001
	(0.002)	(0.002)	(0.002)
MinimumPurchaseQuantity	−0.025***	−0.024***	−0.022***
	(0.002)	(0.002)	(0.002)
StemLength	−0.025**	−0.032***	−0.034***
	(0.009)	(0.009)	(0.009)
BundlingCondition	0.001	0.001	0.0001
	(0.001)	(0.001)	(0.001)
BloomingStage	0.077***	0.080***	0.072***
	(0.011)	(0.011)	(0.011)
Tuesday	−0.043**	−0.043**	−0.045**
	(0.015)	(0.015)	(0.014)
Wednesday	−0.084***	−0.084***	−0.087***
	(0.015)	(0.015)	(0.014)
Thursday	−0.115***	−0.116***	−0.113***
	(0.015)	(0.015)	(0.014)
Friday	−0.055***	−0.060***	−0.058***
	(0.014)	(0.014)	(0.014)
Grower fixed effects	Yes	Yes	Yes
Week fixed effects		Yes	Yes
Bidder fixed effects			Yes
Observations	26,524	26,524	26,524
Adjusted R^2	0.338	0.409	0.437

Notes. All coefficients are estimated at the transaction level. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

and the experiment period would become insignificant. Table 10 presents the results corresponding to Equations (1) and (2) under the alternative benchmark. We can see that the coefficient for the interaction term, *Treatment* × *Experiment*, remains positive and significant. This observation provides additional evidence that the price increase is driven by the policy change.

Finally, for each of the above robustness checks, we also performed a similar analysis using the relative time model as in Section 5.2.1. The results show no violation of the parallel trend assumption (the p-values for the coefficients of the lead indicators are all above 0.20).

5.3. Effect of Withholding Winner's Identity on Price Dynamics

Drawing on prior research on information aggregation and rational learning in sequential auctions (see Section 2.2), we would like to find out whether withholding the winner's identity has any impact on the price dynamics. To begin with, we plotted the price trend at the treatment site in pre-, during-, and post-experiment periods, respectively. Given the heterogeneity across different auctions, we normalized the prices from the same auction with respect to the price in the first round of that auction. Figure 3 depicts the price trends. The rank number in the horizontal axis denotes the rank of a transaction. For example, if a transaction was made in the second round, the rank number is 2. The vertical bar denotes one standard error of the normalized prices in each round. Overall, we can see that the winning price exhibits a declining trend in all three time periods. However, this trend seems to be mitigated during the experiment period. Specifically, the mean values of the normalized prices in the subsequent rounds are higher during the experiment period than those observed during the pre- and post-experiment

Table 9. Treatment Effect on Bids from Different Channels

	Sample	
Variable	Offline bids	Online bids
Treatment × *Experiment*	0.058*	0.062**
	(0.028)	(0.019)
Treatment × *Post*	0.021	0.013
	(0.037)	(0.028)
Treatment	0.003	0.028
	(0.075)	(0.053)
LotSize	0.001	0.002
	(0.002)	(0.002)
MinimumPurchaseQuantity	−0.027***	−0.024***
	(0.004)	(0.002)
StemLength	−0.015	−0.027**
	(0.012)	(0.009)
BundlingCondition	0.002	0.003**
	(0.001)	(0.001)
BloomingStage	0.083***	0.082***
	(0.015)	(0.009)
Tuesday	−0.027	−0.038**
	(0.021)	(0.014)
Wednesday	−0.062**	−0.086***
	(0.020)	(0.014)
Thursday	−0.109***	−0.114***
	(0.020)	(0.014)
Friday	−0.089***	−0.050***
	(0.019)	(0.014)
Grower fixed effects	Yes	Yes
Week fixed effects	Yes	Yes
Observations	5,074	26,752
Adjusted R^2	0.413	0.410

Notes. All coefficients are estimated at the transaction level. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

periods. Furthermore, we find that the variances of the normalized prices decrease considerably during the experiment period. Note that although Figure 3 provides qualitative evidence that withholding the winner's identity can mitigate the price-declining trend in sequential rounds, such a result could be misleading because we have not controlled for any potential confounding factors.

To measure the treatment effect on price dynamics, we adapted the model used in van den Berg et al. (2001), which also looks at the sequential sale at the DFA:

$$\ln\frac{Price_{j,k,t}}{Price_{j,k-1,t}} = \mu_0 + \mu_1 Experiment_t \times (Available_{j,k-1} - 2) + \mu_2 Experiment_t \times (Rank_{j,k} - 2) + \mu_3 Experiment_t + \mu_4 (Available_{j,k-1} - 2) + \mu_5 (Rank_{j,k} - 2) + \epsilon_{j,k,t}. \tag{5}$$

In Equation (5), j indexes each auction, k indexes the rank within auction j, and t indexes the time; $Available_{j,k-1}$ denotes the available units at the beginning of the $(k-1)$th transaction at auction j and $Rank_{j,k}$ denotes the rank of the transaction at auction j; and $Experiment_t$ is a dummy variable that equals 1 during the experiment period and 0 otherwise. We used the difference of log prices in consecutive rounds as the dependent variable. According to van den Berg et al. (2001), this comes with two benefits. First, it controls for potential confounding factors that influence the length (the maximum of the rank number) of an auction and the transaction prices simultaneously. Second, it addresses the potential correlation of prices within a given auction, as well as the observed or unobserved heterogeneity across different auctions, both of which may result in biased estimation. Note that

Table 10. Treatment Effect Under an Alternative Benchmark

Variable	(1) Baseline	(2) Time fixed effect
Treatment × *Experiment*	0.060**	0.060**
	(0.014)	(0.015)
Treatment × *Pre*	−0.015	−0.016
	(0.029)	(0.028)
Treatment	0.059	0.047
	(0.057)	(0.057)
Experiment	0.229***	
	(0.019)	
Pre	0.091***	
	(0.021)	
LotSize	0.002	0.002
	(0.002)	(0.002)
MinimumPurchaseQuantity	−0.025***	−0.024***
	(0.001)	(0.001)
StemLength	−0.019*	−0.025**
	(0.009)	(0.009)
BundlingCondition	0.003**	0.003**
	(0.001)	(0.001)
BloomingStage	0.081***	0.084***
	(0.009)	(0.009)
Tuesday	−0.035*	−0.037*
	(0.014)	(0.014)
Wednesday	−0.085***	−0.082***
	(0.015)	(0.014)
Thursday	−0.124***	−0.114***
	(0.015)	(0.014)
Friday	−0.055***	−0.056***
	(0.014)	(0.014)
Grower fixed effects	Yes	Yes
Week fixed effects		Yes
Observations	31,848	31,848
Adjusted R^2	0.338	0.409

Notes. All coefficients are estimated at the transaction level. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

Figure 3. Comparison of Price Trends

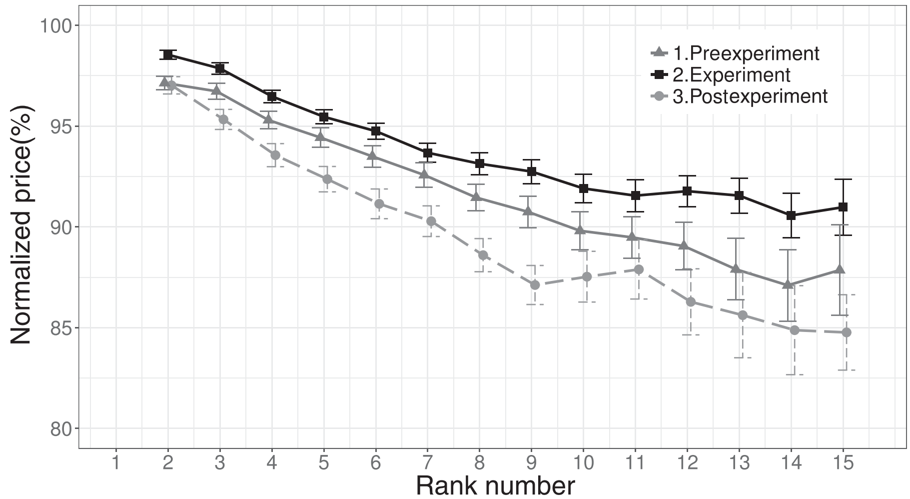

Notes. The *x* axis denotes the rank number. The vertical bars correspond to one standard errors of the means. The curve corresponding to the pre-experiment (post-experiment) period is left (right) shifted to avoid the overlapping of the vertical bars.

Equation (5) is not well defined for single-unit or single-round auctions. Nevertheless, these auctions are irrelevant to understanding the price dynamics in sequential rounds. To distinguish between the effect of the available units and the effect of the rank, we used the reference case of a two-unit lot that is auctioned in two rounds (i.e., subtracting $Available_{j,k-1}$ and $Rank_{j,k}$ by 2).

To account for the potential curvilinear relationship between the rank number and the dependent variable, we considered an alternative model as follows:

$$\begin{aligned}\ln\frac{Price_{j,k,t}}{Price_{j,k-1,t}} &= \mu_0 + \mu_1 Experiment_t \times (Available_{j,k-1} - 2)\\ &+ \mu_2 Experiment_t \times (Rank_{j,k} - 2)\\ &+ \mu_3 Experiment_t + \mu_4 (Available_{j,k-1} - 2)\\ &+ \mu_5 (Rank_{j,k} - 2) + \mu_6 (Rank_{j,k} - 2)^2 + \epsilon_{j,k,t}.\end{aligned} \quad (6)$$

In both specifications (with and without the quadratic term), our coefficient of interest is μ_3, as it captures the impact of the policy change on price dynamics in sequential rounds. We estimated models (5) and (6) using data from the treatment site. The results[11] are presented in Table 11. The estimates of the coefficients largely confirm our observations from Figure 3. Specifically, under both specifications, the estimated coefficient of the intercept is negative and significant ($p < 0.001$), suggesting that there is indeed a declining price trend. The coefficient corresponding to the treatment effect (i.e., *Experiment*) is positive and significant (coefficient = 0.011, $p < 0.001$), indicating that withholding the winner's identity has a mitigation effect on the price-declining trend. Furthermore, we can see that both the coefficient for *Rank* and the coefficient for *Experiment* × *Rank* are positive and significant, indicating that the magnitude of the price decline tends to decrease over time, and the treatment makes it decrease even faster.

Before moving to the investigation of the potential mechanisms that may have led to our findings, we would like to briefly discuss the economic significance of the observed treatment effects. First, we consider the effects on the expected revenue. The DID analyses above reveal that, all else being equal, withholding the

Table 11. Effect of Identity Withholding on Price Dynamics

Variable	(1)	(2)
(Intercept)	−0.024***	−0.028***
	(0.001)	(0.001)
Experiment × *Available*	0.000007	0.000007
	(0.00002)	(0.00002)
Experiment × *Rank*	0.0006*	0.0006*
	(0.0003)	(0.0003)
Experiment	0.011***	0.011***
	(0.002)	(0.002)
Available	0.000005	0.000002
	(0.00001)	(0.00001)
Rank	0.002***	0.003***
	(0.0002)	(0.0003)
$Rank^2$		0.00007***
		(0.00002)
Observations	12,348	12,348
Adjusted R^2	0.010	0.011

Notes. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions. Column (1) contains the estimation result using the baseline model (Equation (5)), and column (2) contains the estimation result using the alternative model by including a quadratic term of the rank number (Equation (6)).
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

winner's identity increases the winning price by more than 6%. For chrysanthemums alone, for which the current annual turnover is about 300 million euros, such increase implies an extra 20 million euros in the expected revenue. Furthermore, it is worth mentioning that the increased stability of price in sequential rounds is good news for both suppliers and buyers in the market. In fact, one of the primary goals for the market maker behind these auctions is to have a stable price-setting mechanism.[12]

6. Understanding the Effects of Identity Withholding

So far, our DID analyses have consistently shown that the policy change (i.e., withholding the winner's identity) could increase the winning price in the sequential auctions of the DFA by more than 6%. In this section, we conduct additional analyses to uncover the mechanism that may lead to such a price increase.

Drawing on the literature on information disclosure and sequential auctions, we consider two potential mechanisms. First, because bidders have been competing in these auctions repeatedly for a long period of time, some bidders, especially the large ones, may have an incentive to implicitly coordinate their bids and engage in tacit collusion (Sherstyuk and Dulatre 2008,Bajari and Yeo 2009). Public disclosure of the winner's identity could serve as an effective coordinating tool on a collusive outcome, as it allows cartel members to identify the defecting ones (Harrington 2012). Therefore, the price increase may be attributed to the mitigation of tacit collusion enabled by the identity withholding policy.

Another potential mechanism is that withholding the winner's identity disrupts bidders' *imitation heuristic* and deters strategic bid shading (Zeithammer 2007). Imitation is an attractive heuristic when decision makers have little information about the strategic environment but can observe others' success. In the context of the DFA, bidders face not only high uncertainty but also extreme time pressure. As a result, bidders tend to rely on their own experiences and pursue different heuristic strategies (Lu et al. 2016). When using the imitation heuristic, a bidder would refer to a subset of her competitors and imitate the strategy of the most successful ones (Schlag 1998, Selten et al. 2005). Note that public disclosure of the winner's identity information is critical to the use of the imitation heuristic: without such information, bidders are not able to tell whether a winner from the previous round is from their reference group. As a result, bidders are more likely to blindly follow the crowd, which could partially explain the increased price stability observed in Section 5.3.

Note that empirically separating the two explanations is a challenging task given that they are not mutually exclusive of one another. With this in mind, we attempt to find out which mechanism may serve as a leading explanation to the observed effect.

6.1. Mitigation of Tacit Collusion

To examine whether withholding the winner's identity mitigates the potential tacit collusion, we first need to identify the collusive bids of cartel members. Unfortunately, current literature on tacit collusion does not provide clear guidance to distinguish collusive bids from noncollusive bids (Bajari and Yeo 2009). In light of this, we chose to explore the potential bid rotation (i.e., one of the cartel members is bidding on behalf of the others in an effort to reduce competition) by examining the bidding patterns. Note that, by definition, cartel members who were pursuing a bid rotation scheme would not actively bid in the same auction. As a result, they were unlikely to be winners from the same auction.

Following this rationale, we constructed two bidder networks for the 288 bidders who were actively bidding at the treatment site throughout the three-week pre-experiment period and the three-week experiment period. The nodes in the networks represent the bidders and the edges correspond to the dyadic constituents of bid rotation: if two bidders never showed up in the same auction during a given (i.e., pre-experiment or experiment) period, they would form a collusive tie represented by an edge (see Figure 4 for an illustration). Similarly, we also constructed two bidder networks for the 334 active bidders at the control site. Given the nature of the edges in these networks, we can use the closeness centrality[13] to measure bidders' collusive tendency (i.e., bidders with a higher closeness centrality were more likely to be involved in bid rotation).

To examine the structural changes in the bidding pattern associated with bid rotation, we estimated the following baseline DID model at the bidder level:

$$\ln(Close_{i,t}) = \lambda_0 + \lambda_1 Treat_i \times Experiment_t + \lambda_2 Treat_i + \lambda_3 Experiment_t + \epsilon_{i,t}, \quad (7)$$

Figure 4. Illustration of the Collusive Ties Among Five Bidders

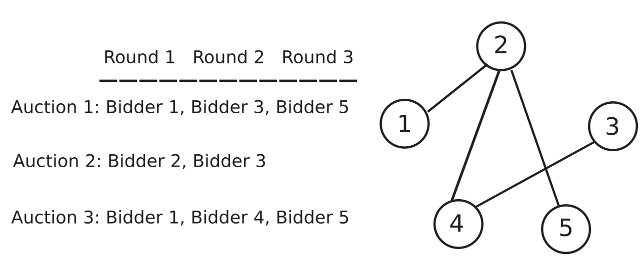

Notes. Bidder 2 never showed up in the same auction with Bidders 1, 4, and 5; thus we generated three edges between Bidder 2 and Bidders 1, 4, and 5, respectively. Similarly, Bidders 3 and 4 are connected by an edge as they did not win the same auction.

where $Close_{i,t}$ denotes the closeness centrality of bidder i during time t, $Experiment_t$ is a dummy variable that equals 1 during the experiment period, and $Treat_i$ is a dummy variable that equals 1 if bidder i was from the treatment site. If the price increase was mainly driven by the mitigation of bid rotation, we would expect to see a significant decrease in the closeness centrality, which is measured by the coefficient λ_1. To account for the potential confounding factors that affect a bidder's collusive tendency (e.g., large bidders may have more bargaining power or benefit more from bid rotation), we also considered two alternative model specifications where we include bidders' average purchase quantity per week (*AveragePurchaseQuantWeek*) and average purchase quantity per transaction (*AveragePurchaseQuantTransact*) as a control variable, respectively.

Table 12 presents the estimation results.[14] The coefficient of the interaction term (i.e., *Treatment* × *Experiment*) is not significant ($p > 0.2$), suggesting that the mitigation of tacit collusion (particularly, bid rotation) is not the most likely explanation to the price increase during the policy change.

6.2. Disruption of Imitation Heuristic

Intuitively, if two bidders, i and j, frequently won in the same auction (not necessarily in consecutive rounds) and i always won after j, it is likely that i was imitating j. Furthermore, an essential prerequisite for the imitation heuristic is that bidders can identify the competitors that they want to imitate. Given this consideration, we performed network analysis similar to that above to examine whether withholding winners' identities deters the imitation heuristic. However, unlike bid rotation where the colluding members are in a symmetric, reciprocal relationship, imitation between bidders is directed. As a result, we had to reconstruct the bidder networks.

To do so, we kept the 288 (bidder) nodes for the networks at the treatment site and the 334 (bidder) nodes for those at the control site, but we generated edges based on the dyadic relationship of imitation (see Figure 5 for an illustration). Specifically, if bidder i *always* won after bidder $j, j \neq i$, during a given (i.e., pre-experiment or experiment) period, they would form an imitation tie represented by an edge. On the other hand, if bidder i won before bidder j in some auctions but after in other auctions in a given period, the two nodes corresponding to i and j would not form an imitation tie.

We then estimated the baseline DID model specified in Equation (7) as well as the alternative models based on these imitation networks. The results[15] are reported in Table 13. The coefficient of the interaction term is negative and highly significant ($p < 0.001$). Furthermore, we can also see that, prior to the policy intervention, bidders at the treatment site are more likely to use the imitative heuristic. Such a finding provides compelling evidence that the observed price increase was primarily due to the disruption of the imitation heuristic.

7. Discussion

In this paper, we examine the role of identity disclosure in shaping bidders' strategic interactions in sequential Dutch auctions. Using a large-scale field experiment, we compare two alternative identity disclosure policies in a dynamic, complex B2B market with experienced bidders and find that concealing winners' identities can significantly increase the average winning price and thereby raise the seller's revenue. It is worth noting that whereas the importance of field work has been well acknowledged (Athey et al. 2011), prior research has largely focused on single-unit auctions (Cho et al. 2014, Tadelis and Zettelmeyer 2015). To the best of our

Table 12. Treatment Effect on Tacit Collusion

Variable	(1)	(2)	(3)
Treatment × *Experiment*	0.018 (0.015)	0.018 (0.015)	0.018 (0.015)
Treatment	0.142*** (0.011)	0.142*** (0.011)	0.142*** (0.011)
Experiment	−0.006 (0.011)	−0.006 (0.010)	−0.007 (0.011)
AveragePurchaseQuantWeek		−0.00006** (0.00002)	
AveragePurchaseQuantTransact			−0.0002 (0.0004)
Observations	1,092	1,092	1,092
Adjusted R^2	0.264	0.269	0.264

Notes. This table reports the estimation results regarding the treatment effect on bidders' collusive tendency (measured by the closeness centrality). Column (1) shows the results from the baseline model; columns (2) and (3) show the results from the alternative models with the additional control variables *AveragePurchase QuantWeek* and *AveragePurchaseQuantTransact*, respectively. Standard errors are shown in parentheses.

*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

Figure 5. Illustration of the Imitation Ties

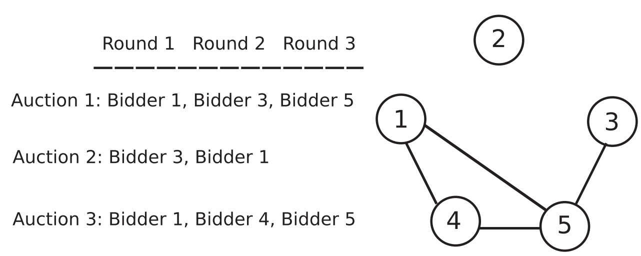

Note. Bidder 5 always won after Bidders 1, 3, and 4, and Bidder 4 won after Bidder 1; thus we generated four edges corresponding to these relationships.

Table 13. Treatment Effect on Imitation

Variable	(1)	(2)	(3)
Treatment × *Experiment*	−0.481***	−0.479***	−0.480***
	(0.022)	(0.021)	(0.022)
Treatment	0.571***	0.568***	0.569***
	(0.016)	(0.015)	(0.016)
Experiment	0.476***	0.476***	0.477***
	(0.015)	(0.014)	(0.015)
AveragePurchaseQuantWeek		0.0004***	
		(0.00003)	
AveragePurchaseQuantTransact			0.001*
			(0.0006)
Observations	1,092	1,092	1,092
Adjusted R^2	0.629	0.679	0.630

Notes. This table reports the estimation results regarding the treatment effect on bidders' imitation (measured by the closeness centrality). Column (1) shows the results from the baseline model; columns (2) and (3) show the results from the alternative models with the additional control variables *AveragePurchaseQuantWeek* and *AveragePurchase QuantTransact*, respectively. Standard errors are shown in parentheses.
*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

knowledge, this is the first study that examines the information disclosure problem in sequential auctions using field data.

Although the primary goal of this study is to examine the effects of the policy change in identity disclosure on market performance, we also provide a theoretical explanation as to why bidders would pay higher prices when winners' identities were concealed. Specifically, given the complexity of the bidding problem and the extreme time pressure, the public disclosure of winners' identities allows bidders to follow a simple heuristic—that is, imitating successful winners from certain reference groups (Selten et al. 2005) in previous rounds and strategically shading their bids. Withholding the winner's identity makes it difficult for bidders to keep track of the members from their reference groups and thereby disrupts the imitation heuristic.

7.1. Contribution

The current paper makes several contributions to the literature on information disclosure in complex markets. First, despite the growing interest in the information disclosure problem in online auction markets, prior research has largely focused on the choice of bid visibility by comparing the performance of open and sealed-bid auctions. By examining the role of competitors' identity information on bidding processes and outcomes, our study complements the existing literature and improves our understanding of the way that different information elements could jointly shape the dynamic interactions in auction markets. Specifically, our empirical results suggest that in real-world sequential auctions, bidders' inferences of the market trend are not merely based on the revealed prices in previous rounds but also based on the revealed winners' information. In other words, the price system does facilitate information aggregation in the market. However, it does this imperfectly.

Second, we find that withholding the winner's identity can significantly mitigate the declining price trend in sequential rounds. Furthermore, by exploring the bidding heuristics in the field setting with professional bidders, we provide empirical evidence that the declining trend can be attributed, to a large extent, to bidders' adaptive learning (Jeitschko 1998). These results shed light on the well-known declining price anomaly in sequential auctions (van den Berg et al. 2001) and improves our understanding of the dynamic decision making in these complex auctions.

Third, our study adds to the literature on behavioral biases in complex decision environments. Whereas extensive research has shown that bidders' behaviors systematically deviate from theoretical predictions and they are not making the best responses (Bichler et al. 2010), there is limited understanding about how bidders handle different levels of complexity and time pressure in real-world environments. Drawing on the literature on behavioral economics (Weibull 1997,Gigerenzer and Selten 2002), we find that bidders are likely to use the imitation heuristic in sequential auctions when the winner's identity is publicly disclosed.

7.2. Implication

Our study also provides useful implications for practitioners. First, our findings confirm that information feedback in repeated competition has a significant impact on the outcome (Ockenfels and Selten 2005). In particular, given our finding that withholding the winner's identity can raise the seller's expected revenue in sequential auctions, we suggest that market designers should be cautious in following the conventional wisdom about information disclosure. Instead, any decision should be grounded on a systematic investigation of the potential impact on the market processes, especially the dynamic interaction between different market participants.

Second, our results suggest that adjustments or changes in the information disclosure policy can have different effects on market participants' strategic behaviors, even if they are equally experienced and knowledgeable. As a result, market makers (e.g., auctioneers) need to closely monitor the composition of the participant population and adapt the information disclosure policy accordingly. Furthermore, given that humans tend to rely on heuristics (e.g., the imitation heuristic) instead of the theoretically sound optimal strategies in complex environments, it is important to account for the various behavioral biases when making policy changes.

Third, in the empirical setting of the DFA, despite the various initiatives enabled by technological developments, the auction house still embraces traditions going back more than a century, and its methods have changed little. The results from our study suggest the current auction rules on information disclosure have ample room for improvement. Specifically, we show that by removing the winner's identity from the auction clock, auctioneers can increase the average winning price by more than 6%. Such a policy change can also effectively mitigate the declining price trend and thereby increase price stability in the market. In addition, given the observed effects of the control variables—for example, day-of-the-week and minimum purchase quantity—auctioneers could leverage the rich historical data from daily transactions as well as well-designed experiments to optimize the pricing rules and the market processes.

To this point, we would like to point out that although the current study mainly focuses on the DFA, the results provide insights to practitioners in other B2B markets. Specifically, for many government procurement auctions, the disclosure of competitors' identities plays a critical role, given the high uncertainty about the expected cost and the potential risk of collusion. In this respect, the results from our study can serve as a useful starting point for future work in investigating information disclosure issues in these markets.

Finally, as Nobel laureate Alvin Roth writes, "Market design involves a responsibility for detail, a need to deal with all of a market's complications, not just its principle features" (Roth 2002, p. 1341). The empirical findings from the current study suggest that there are important dynamics in sequential auctions that are not captured by the classic framework of auction design. From the market design perspective, this highlights the need for experimentation or, more broadly, a data-driven approach to optimizing the design of real-world auctions.

7.3. Limitation and Future Work

The current study bears several limitations and offers opportunities for future work. First, we could not account for the potential complementarity or substitutability of different products in the empirical analysis with the current data. If there had been a demand shock of another product that serves as a complement to the product chosen in our analyses, it would have led to an overestimation of the actual effect of the policy change. Fortunately, we were reassured by the market makers that there was no such demand shock during our study period. Therefore, this is not a big concern in this paper. Nevertheless, future work can take transaction data from multiple types of products to account for the potential complementarity and verify the robustness of our results.

Second, our empirical analysis focuses on the role of different identity disclosure policies on the bidding competition in the auction market. Given that these are B2B auctions, it is also interesting to see how the policy change impacts post-auction trades. For example, when winners' identities are not publicly disclosed, it becomes much more difficult for customers in the downstream market to track the original purchasing prices in the auction market even if they have online access to view the real-time auctions. This may allow bidders to increase profit margins and thereby affect bidding strategies in the auctions. An integrated model that takes into account the postauction competition in the downstream market can be very helpful in understanding the impact of different disclosure policies in the whole supply chain.

Third, our current study only looks at two special identity disclosure policies (i.e., complete revealing and complete hiding). Drawing upon recent work on information disclosure in online advertising exchanges (e.g., Sun et al. 2016), it would be interesting to explore the continuum of disclosure policies by examining other important information elements (e.g., the real-time market structure information), which may affect bidders' decision making.

Finally, we consider two mechanisms that may serve as the leading explanations to the observed treatment effect. However, there might be other mechanisms that could drive the observed effect. Although it is challenging to identify the exact mechanism or disentangle different potential mechanisms with our current data, we do believe that a full treatment of this subject would be a very promising direction for future work.

Acknowledgments

The authors thank Chris Forman (department editor), the associate editor, and three anonymous reviewers for valuable suggestions. This research would not have been possible without the support of Royal FloraHolland via the "Veilen met Advies" project. The authors also thank the constructive feedback from the participants in presentations at the 2014 Conference on Information Systems and Technology (CIST 2014) and the 2014 International Conference on Information Systems (ICIS 2014).

Endnotes

[1] The opponent effect refers to "an increase in the subjective value of winning the auction when the behavior of the other bidders in the auction is perceived to be competitive" (Heyman et al. 2004, p. 9).

[2] More details can be found from https://www.royalfloraholland.com/en/ (last accessed July 1, 2017).

[3] A *lot* is a bundle of homogeneous products (flowers) from a grower.

[4] Currently, the reserve price is fixed for the entire year, regardless of the auction site or flower type.

[5] Strictly speaking, our study used a quasi-experiment design, as the treatment (i.e., the policy change) is not randomly assigned to the bidders or auctions (Shadish et al. 2002).

[6] The specific time frame was determined in consultation with the managing team of the auction market, and bidders at the treatment site were also informed about the policy change during this time frame.

[7] Based on our interview with the wholesalers, such location preference has to do with the logistics costs.

[8] The paper by van den Berg et al. (2001) shows empirical evidence for declining price anomaly in the flower auctions; however, if we look at individual auctions, the price trend is inconclusive.

[9] We performed similar analysis to examine the parallel trend assumption underlying the alternative specifications in Equation (3) and Equation (4) (i.e., we replaced $Experiment_t$ and $Post_t$ in the two equations by the lead and lag dummies and estimated the corresponding relative time models). Again, we find that the coefficients corresponding to the lead terms are not statistically significant. Specifically, the p-values for the coefficient estimates of the two lead terms are 0.74 and 0.51 for the simple average price model and 0.71 and 0.53 for the weighted average price model.

[10] Among the 31,848 transactions within the original sample, there were 22 transactions where the channel usage information was missing. Thus we excluded them from the estimation.

[11] We also estimated an alternative model by including a square root term, $(k-2)^{1/2}$, to Equation (6). The results are qualitatively similar to those shown in Table 11.

[12] See https://www.royalfloraholland.com/en/ (last accessed July 1, 2017) for more details about the current initiatives taken by the market makers in the DFA.

[13] A node's closeness centrality is the average minimal distance between the node and any other nodes in the network. Formally, for a node i, the closeness centrality is defined as $(n-1)/\sum_{j,j\neq i} d(i,j)$, where n is the total number of nodes in the network and $d(i,j)$ denotes the shortest path between node i and node j. If node j is not reachable from node i, $d(i,j)$ is defined as the total number of nodes in the network (Wasserman and Faust 1994).

[14] We also estimated the relative time models corresponding to the baseline specification in Equation (7) and the two alternative specifications by constructing six bidder networks corresponding to the six weeks at each site. We did not find evidence indicating violation of the parallel trend assumption (i.e., the coefficients for the lead terms are not significant). Furthermore, the estimation results are quite consistent with those from Table 12.

[15] As before, we also estimated the three relative time models based on the bidder networks generated under the new dyadic relationship. Again, we did not find evidence indicating violation of the parallel trend assumption (i.e., the coefficients for the lead terms are not significant).

References

Abadie A (2005) Semiparametric difference-in-differences estimators. *Rev. Econom. Stud.* 72(1):1–19.

Adair JG (1984) The Hawthorne effect: A reconsideration of the methodological artifact. *J. Appl. Psych.* 69(2):334–345.

Adomavicius G, Curley S, Gupta A, Sanyal P (2012) Effect of information feedback on bidder behavior in continuous combinatorial auctions. *Management Sci.* 58(4):811–830.

Angrist JD, Pischke J (2008) *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton University Press, Princeton, NJ).

Arora A, Greenwald A, Kannan K, Krishnan R (2007) Effects of information-revelation policies under market-structure uncertainty. *Management Sci.* 53(8):1234–1248.

Aseff J, Chade H (2008) An optimal auction with identity-dependent externalities. *RAND J. Econom.* 39(3):731–746.

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

Autor DH (2003) Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *J. Labor Econom.* 21(1):1–42.

Bajari P, Yeo J (2009) Auction design and tacit collusion in FCC spectrum auctions. *Inform. Econom. Policy* 21(2):90–100.

Beggs A, Graddy K (1997) Declining values and the afternoon effect: Evidence from art auctions. *RAND J. Econom.* 28(3): 544–565.

Bergemann D, Hörner J (2017) Should auctions be transparent? Cowles Foundation Discussion Paper No. 1764R2, Cowles Foundation for Research in Economics, Yale University, New Haven, CT.

Bertrand M, Duflo E, Mullainathan S (2004) How much should we trust differences-in-differences estimates? *Quart. J. Econom.* 119(1):249–275.

Bichler M, Gupta A, Ketter W (2010) Designing smart markets. *Inform. Systems Res.* 21(4):688–699.

Cason TN, Kannan KN, Siebert R (2011) An experimental study of information revelation policies in sequential auctions. *Management Sci.* 57(4):667–688.

Cho S, Paarsch H, Rust J (2014) Is the "linkage principle" valid? Evidence from the field. *J. Indust. Econom.* 62(2):346–375.

Cramton P, Schwartz J (2000) Collusive bidding: Lessons from the FCC spectrum auctions. *J. Regul. Econom.* 17(3):229–252.

Delgado MR, Schotter A, Ozbay EY, Phelps EA (2008) Understanding overbidding: Using the neural circuitry of reward to design economic auctions. *Science* 321(5897):1849–1852.

Donald SG, Paarsch HJ, Robert J (2006) An empirical model of the multi-unit, sequential, clock auction. *J. Appl. Econometrics* 21(8): 1221–1247.

Gigerenzer G, Selten R (2002) *Bounded Rationality: The Adaptive Toolbox* (MIT Press, Cambridge, MA).

Goes P, Karuga G, Tripathi A (2010) Understanding willingness to pay formation of repeat bidders in online sequential auction. *Inform. Systems Res.* 21(4):907–924.

Greenwood B, Wattal S (2017) Show me the way to go home: An empirical investigation of ride-sharing and alcohol related motor vehicle fatalities. *MIS Quart.* 41(1):163–187.

Haberbush KL (2000) Limiting the government's exposure to bid rigging schemes: A critical look at the sealed bidding regime. *Public Contract Law J.* 30(1):97–122.

Harrington JE (2012) A theory of tacit collusion. Working paper No. 588, Department of Economics, Johns Hopkins University,Baltimore.

Heyman JE, Orhun Y, Ariely D (2004) Auction fever: The effect of opponents and quasi-endowment on product valuations. *J. Interactive Marketing* 18(4):7–21.

Hong Y, Wang C, Pavlou P (2016) Comparing open and sealed bid auctions: Evidence from online labor markets. *Inform. Systems Res.* 27(1):49–69.

Imbens GW, Wooldridge JM (2009) Recent developments in the econometrics of program evaluation. *J. Econom. Lit.* 47(1):5–86.

Jehiel P, Moldovanu B (1996) Strategic nonparticipation. *RAND J. Econom.* 27(1):84–98.

Jeitschko TD (1998) Learning in sequential auctions. *Southern Econom.* J. 65(1):98–112.

Klemperer P (1999) Auction theory: A guide to the literature. *J. Econom. Surveys* 13(3):227–286.

Koppius OR (2002) Information architecture and electronic market performance. PhD thesis, Erasmus University Rotterdam, Rotterdam, Netherlands.

Kumar A, Tan Y (2015) The demand effects of joint product advertising in online videos. *Management Sci.* 61(8):1921–1937.

Lacetera N, Larsen BJ, Pope DG, Sydnor JR (2016) Bid takers or market makers? The effect of auctioneers on auction outcome. *Amer. Econom. J. Microeconom.* 8(4):195–229.

Lu Y, Gupta A, Ketter W, van Heck E (2016) Exploring bidder heterogeneity in multichannel sequential B2B auctions. *MIS Quart.* 40(3):645–662.

Marshall R, Marx L (2009) The vulnerability of auctions to bidder collusion. *Quart. J. Econom.* 124(2):883–910.

McAfee R, Vincent D (1993) The declining price anomaly. *J. Econom. Theory* 60(1):191–212.

Milgrom P, Weber RJ (1982) The value of information in a sealed-bid auction. *J. Math. Econom.* 10(1):105–114.

Ockenfels A, Selten R (2005) Impulse balance equilibrium and feedback in first price auctions. *Games Econom. Behav.* 51(1):155–170.

Overby E, Kannan K (2015) How reduced search costs and the distribution of bidder participation affect auction prices. *Management Sci.* 61(6):1398–1420.

Perry M, Reny P (1999) On the failure of the linkage principle in multi-unit auctions. *Econometrica* 67(4):895–900.

Roth AE (2002) The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica* 70(4):1341–1378.

Schlag K (1998) Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *J. Econom. Theory* 78(1): 130–156.

Selten R, Abbink K, Cox R (2005) Learning direction theory and the winner's curse. *Experiment. Econom.* 8(1):5–20.

Shadish WR, Cook TD, Campbell DT (2002) *Experimental and Quasi-Experimental Designs for Generalized Causal Inference* (Houghton Mifflin Company, Boston).

Sherstyuk K, Dulatre J (2008) Market performance and collusion in sequential and simultaneous multi-object auctions: Evidence from an ascending auctions experiment. *Internat. J. Indust. Organ.* 26(2): 557–572.

Smith C (2007) Markets as definitional practices. *Canadian J. Sociol.* 32(1):1–39.

Snir E, Hitt L (2003) Costly bidding in online markets for it services. *Management Sci.* 49(11):1504–1520.

Sun Z, Dawande M, Janakiraman G, Mookerjee V (2016) The making of a good impression: Information hiding in ad exchanges. *MIS Quart.* 40(3):717–739.

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

Tu Z (2005) Why do we use the Dutch auction to sell flowers? Information disclosure in sequential auctions. Discussion paper, University of Pittsburgh, Pittsburgh.

van den Berg GJ, van Ours J, Pradhan M (2001) The declining price anomaly in Dutch Dutch rose auctions. *Amer. Econom. Rev.* 91(4): 1055–1062.

Varma GD (2002a) Standard auctions with identity-dependent externalities. *RAND J. Econom.* 33(4):689–708.

Varma GD (2002b) Who else is bidding? The Pareto optimality of disclosing bidder identities. *Rev. Econom. Design* 7(2):155–171.

Wasserman S, Faust K (1994) *Social Network Analysis: Methods and Applications*, vol. 8 (Cambridge University Press, Cambridge, UK).

Weber R (1983) Multiple-object auctions. Engelbrecht-Wiggans R, Shubik M, Stark R, eds. *Auctions, Bidding, and Contracting: Uses and Theory* (New York University Press, New York), 165–191.

Weibull JW (1997) *Evolutionary Game Theory* (MIT Press, Cambridge, MA).

Zeithammer R (2007) Strategic bid-shading and sequential auctioning with learning from past prices. *Management Sci.* 53(9): 1510–1519.