

Information Transparency in B2B Auction Markets: The Role of Winner Identity Disclosure

Yixin Lu (yixinlu@gwu.edu)¹, Alok Gupta (alok@umn.edu)²,
Wolfgang Ketter (ketter@wiso.uni-koeln.de)^{3,4}, and Eric van Heck (eheck@rsm.nl)³

¹George Washington University - School of Business, United States

²University of Minnesota - Carlson School of Management, United States

³Erasmus University - Rotterdam School of Management, The Netherlands

⁴University of Cologne - Faculty of Management, Economics, and Social Sciences, Germany

Abstract

We study the impact of information transparency in B2B auctions. Specifically, we measure the effect of concealing winners' identities on auction outcomes using a large-scale, quasi-natural field experiment. Contrary to the conventional wisdom that "the more information, the better," we find that concealing winners' identities leads to a significant increase in price by approximately 6%, and such effect holds true across both online and offline channels as well as different types of bidders. We further explore the mechanism that drives the observed effect. The empirical analysis suggests that the price increase may primarily stem from the disruption of imitative bidding which relies on the identification of fellow competitors. Our findings have important implications for the design of auction markets, especially multi-channel B2B markets.

Key words: Auction design, field experiment, information transparency, identity disclosure, sequential auctions

1 Introduction

One of the central issues in auction design is how much information should be disclosed to bidders (Milgrom and Weber 1982). The choice arises in both the public and private sectors, ranging from allocating natural resources such as timber, offshore oil rights, and spectrum licenses to sales of artwork, real estate, and various commodities. While a large body of theoretical work has identified conditions under which one disclosure policy outperforms the others on efficiency or revenue, empirical evidence is rather limited, partly because many markets tend to "operate under a given set of rules rather than experimenting with alternative designs" (Athey et al. 2011, p.1).

Undoubtedly, the choice of information disclosure policies reflects a variety of considerations (e.g., efficiency, privacy, etc.). For example, in the U.S., due to “Freedom of Information Act”, procurement auctions in the public sector is often subject to strict transparency requirements that require full disclosure of bidders’ identities as well as their bids. By contrast, auction houses such as Christie’s and Sotheby’s typically preserve the anonymity of winning bidders. Further, in many real markets, the information disclosure problem can be very subtle: sellers may not be able to observe and disclose information directly, or they may not have full control of bidders’ perceived content of the information (Abraham et al. 2013).

Recognizing these real-world complications, in this paper, we study the information disclosure problem in the context of the worlds’ largest wholesale market of cut flowers, namely, the Dutch Flower Auctions (DFA). The DFA features a dynamic, complex B2B market where trades between sellers and buyers are facilitated through multi-unit, sequential Dutch auctions. 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 practice has conflicting effects on the bidding dynamics. On the one hand, it may induce more aggressive bidding behavior in the presence of negative identity-dependent externalities (Varma 2002), as bidders can track their competitors’ successful purchases in the past auctions. On the other hand, the public disclosure of winners’ identities may also facilitate bidders’ learning in sequential rounds (Jeitschko 1998) and result in considerable bid shading. As such, whether or not to disclose winners’ identities is a critical design choice for this B2B auction market.

We conducted a large-scale, quasi-natural field experiment to examine the impact of withholding winners’ identities on auction outcomes. Specifically, we manipulated the information disclosure policy at a major auction site from mid-November to early December 2012 by removing winner’s identities from the screen of auction clocks while keeping all other aspects unchanged. To establish a causal inference of the effect of the policy change, we constructed a matched sample that contains transaction data from both the experiment site and a comparison site from late October till the end of December 2012. Using the difference-in-differences estimation approach, we find that bidders pay significantly higher prices when winners’ identities were *concealed* from public view. Such finding holds up to various specifications and robustness tests.

Several potential mechanisms that may result in the price increase when winners’ identities are withheld rather than publicly disclosed. Drawing upon 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 differential treatment effects and changes in bidders’ dynamic interaction pattern, the explanation that

is most consistently supported for why withholding winners’ identities leads to the price increase is that bidders tend to imitate their *fellow* competitors and shade their bids accordingly. Once the winners’ identities are concealed, it becomes difficult to identify the fellow competitors as well as their bids, thus bidders are forced to use the most recent bids as the reference points (“anchors”) when making their bidding decisions.

Our paper makes several contributions to both 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 literature of market design, the identities of participants are typically assumed to play a minimal role in determining the outcome of the exchanges. However, with the proliferation of online markets, this is not the case anymore: due to the lack of face-to-face interactions, participants in online markets have to establish themselves in a formal manner as legitimate (Smith 2007). To the best of our knowledge, this is the first that has systematically investigated the identity disclosure problem in sequential auctions. Second, by exploring the underlying mechanisms which drive the differences in bidding dynamics and outcomes under different disclosure policies, we also shed new light on the understanding of the information transmission and learning in sequential auctions. Third, unlike most of the previous studies, we empirically measure the impact of alternative disclosure policies in a complex, dynamic B2B market. Thus our findings offer more actionable insights into the practical design. Specifically, contrary to the common wisdom that “the more transparent, the better” in market design, our results show 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. Finally, it is worthy to note that although we have primarily focused 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 and preliminary analysis. 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, and reflect on the limitations and point out the directions for future work.

2 Literature Review

The current research draws upon the interdisciplinary research of auction design, particularly the growing literature on the effects of information disclosure policies on market outcomes, and prior works on the bidding dynamics in sequential auctions.

2.1 Information Disclosure in Auctions

One of the most significant findings on information disclosure in auctions is the *linkage principle*, which states that sellers can raise the expected revenue by increasing information transparency and allowing bidders to learn from each others' bids (Milgrom and Weber 1982). Due to the challenge of comparing alternative information disclosure policies in the field, the majority of the empirical work related to the linkage principle is based on controlled lab experiments, and the results from these studies are mixed. Kagel and Levin (1986) and Goeree and Offerman (2002) have shown that providing more information can reduce bidders' value uncertainty and thus increase the expected revenue, whereas Haruvy and Katok (2013) find that sealed-bid format generates higher buyer surplus than the dynamic open-bid format in procurement auctions.

Among the few studies that use field data, Jap (2007) reports that in procurement auctions, the less visible the bids are, the more positive the impact is on the interorganizational relationship. Such finding is echoed by Athey et al. (2011), where the authors show that sealed bid auctions can generate higher revenues than open auctions. By contrast, Cho et al. (2014) and Tadelis and Zettelmeyer (2015) find that disclosing more information results in a significant increase in seller's revenue, which is consistent with the linkage principle. Table 1 provides an overview of the mixed findings from the empirical studies discussed above, which, regardless of the research design and the findings, has exclusively focused on comparing the performance of open and sealed bid auctions.

The linkage principle does not hold, in general, beyond single-unit auctions. For example, Perry and Reny (1999) have shown that in a two-unit Vickrey auction, the release of information has a downward effect (on average) on winning bids and the linkage principle breaks down. Tu (2007) has compared the revenue properties of a number of auction formats and disclosure policies in various two-period sequential auctions. He finds that in the first-price auction, announcing the winning bid alone yields higher revenue than announcing both the winning and losing bids, which again contradicts the prediction of the linkage principle. Bergemann and Hörner (2010) have also reported that in repeated first-price auctions, the minimal disclosure policy – each bidder only learns whether he wins or loss privately at the end of each round – generates the highest expected revenue. Cason et al. (2011) provide additional insights regarding the observed conflicting results from prior studies by showing that whether or not a full disclosure policy (i.e., both winning and losing bids get revealed) yields higher buyer surplus in procurement auctions depends on the *perceived* competitiveness in the market.

Our current paper differs from the research mentioned above in at least two aspects. First, the prior studies have exclusively focused on the auctions where bidders have single-

Table 1: Summary of prior empirical work on information disclosure in auctions

Title and Author	Data Source	Consistent with Linkage Principle?
The Winner’s Curse and Public Information in Common Value Auctions (Kagel and Levin 1986)	Lab Experiment	Yes
Efficiency in Auctions with Private and Common Values: An Experimental Study (Goeree and Offerman 2002)	Lab Experiment	Yes
The Impact of Online Reverse Auction Design on Buyer-Supplier Relationships (Jap 2007)	Field	No
Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions (Athey et al. 2011)	Field	No
An Experimental Study of Information Revelation Policies in Sequential Auctions (Cason et al. 2011)	Lab Experiment	No
Increasing Revenue by Decreasing Information in Procurement Auctions (Haruvy and Katok 2013)	Lab Experiment	No
Is the ‘Linkage Principle’ Valid? Evidence from the Field (Cho et al. 2014)	Field Experiment	Yes
Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment (Tadelis and Zettelmeyer 2015)	Field Experiment	Yes

unit demand (thus demand reduction is not a concern) whereas we study competitive bidding in sequential auctions where bidders have multi-unit demand in each round. Second, most of the existing research have chosen to address the information disclosure problem by comparing the market performance under complete and partial disclosure of bids. We, on the other hand, focus on a more granular level of the problem by examining alternative disclosure policies within the partial disclosure regime, i.e., the disclosure of bids with and without identity information.

2.1.1 The Effects of Identity Disclosure in Auctions.

Over the past decade, the role of identity disclosure in auction markets has attracted increasing attention from auction theorists. Jehiel and Moldovanu (1996) have documented that the presence of identity-dependent externalities may create intricate strategic effects on bidders’ willingness to participate in the market process. Specifically, when bidders experience considerable non-reciprocal identity-dependent externalities¹, disclosing bidders’ identities can increase seller’s revenue (Varma 2002).

Nevertheless, one may think of numerous circumstances where the disclosure of bidders’

¹In the presence of non-reciprocal externalities, each bidder’s equilibrium willingness-to-pay is determined by the identities of his opponents.

identity information has a negative impact on the auction outcome. In fact, in recent years, many online consumer auction markets (e.g., eBay) have moved toward a less transparent setting where bidders' identities are no longer publicly disclosed. Besides privacy concerns, the public disclosure of bidders' identities may also create a screening effect (Snir and Hitt 2003) which discourages participation. Additionally, it is widely accepted that the disclosure of bidders' identities is that it increases an auction's susceptibility to collusion (Bajari and Yeo 2009, Cramton and Schwartz 2000, Marshall and Marx 2009) and thereby leads to poor market performance (e.g., reduced revenue). 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.

Given the existence of these conflicting forces, whether it is beneficial to disclose bidder's identity is an empirical question. In light of this, we conduct a large-scale field experiment where we exogenously vary the visibility of bidder's identity information. Compared to prior works on identity disclosure, our empirical setting features a complex B2B market with highly experienced bidders, and this allows us to develop a better understanding about the effects of identity disclosure on bidding dynamics and market performance.

2.2 Learning in Sequential Auctions

Compared to single-unit auctions, sequential auctions allow bidders to learn about the market trend from previous rounds of competition. This also makes the analysis of these auctions much more complicated. Within the symmetric independent private value (IPV) paradigm, when bidders have single-unit demand, Weber (1983) shows that the equilibrium price path under the four standard auction formats (English, Dutch, first-price sealed bid, second-price sealed bid) follows a martingale; when bidders have multi-unit Poisson demand, Donald et al. (2006) demonstrate that the equilibrium price path follows a super-martingale, i.e., on average, the equilibrium price increases over time.

Unfortunately, such neat theoretical results are not supported by empirical findings. Specifically, McAfee and Vincent (1993) have detected price declines in sequential rounds in wine auctions; Beggs and Graddy (1997) have observed a similar trend in art auctions; Van den Berg et al. (2001) have provided evidence of price declines in the flower auctions. 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 direct and indirect effect on bidders’ strategies. The direct effect refers to the belief updating about the probability that another bidder has the same valuation, and the indirect effect being the trade-off between the benefits associated with winning in early rounds and the benefits of losing but learning more about their competitor’s valuations.

In the current paper, we draw upon these learning aspects in exploring the underlying mechanisms of the observed differences under different information disclosure policies. However, unlike prior studies which often impose restrictive assumptions about bidders’ dynamic interaction in sequential rounds, we make inferences about bidders’ strategies based on the observed bidding outcomes. In this regard, our findings also contribute to the growing literature of behavioral game theory (Camerer 2003).

3 Research Setting

3.1 Research Context

We examine the identity disclosure problem in the context of the Dutch Flower Auctions (DFA). The DFA account for more than 60% of the global flower trade and generate an annual turnover of over €4 billion. More than 6,000 global suppliers involved in the day-to-day operation of this market². On weekdays, 38 auction clocks which are strategically located at six sites across the country run simultaneously between 6 AM to 10 AM. On average, there are 125,000 daily transactions.

The DFA use a sequential Dutch auction mechanism. They are implemented using fast-paced auction clocks displayed on an electronic board. Apart from the current asking price, each clock also contains information about the setup of the current auction (for example, monetary unit, minimum purchase units as well as bundling properties). Further, bidders can also see the information of the product under auction (for example, the name of the product, the identity of the grower, and 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 which corresponds to a high price of the product 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

²More details can be found from <https://www.royalfloraholland.com/en/>.




Figure 1: Illustration of the auction screen. The setup of the current auction is shown on the clock whereas the product information and the upcoming schedules are shown on the right and top-left of the screen, respectively.

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 the products get sold, or the price falls below the seller's reserve price³, 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, and reserve prices) as well as the information disseminated during the sequential rounds.

Recent years, as more and more bidders have adopted the remote bidding application and participate in the auctions via the online channel, the information disclosure decision becomes 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 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 "who bought what at which clock" confidential, that is, removing the winner's identity from the auction clock). However, it is unclear whether and

³Currently, the reserve price is fixed for the entire year, regardless of the auction site and flower types.

how such change would affect market performance.

3.2 Research Design

The data for this study were collected through a large-scale, quasi-natural field experiment (Shadish et al. 2002) during the last quarter of 2012 in the DFA. The treatment site was chosen randomly among the six sites, and the policy change with respect to winner’s identity disclosure was implemented at a clock which auctioned chrysanthemums, the flowers in season. The experiment lasted from November 19 to December 7, 2012, during which the winner’s identity was removed from the clock screen⁴. It should be noted that while none of the bidders except the winner knew who had won in each round, we as researchers could see who were the winners because such 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 due to systematic changes in supply or demand over time. Fortunately, we were able to obtain data from a *comparison* site where the same type of flower was auctioned and bidders could always observe winners’ identities throughout the eight-week study period. Figure 2 provides an overview of our research design.

The treatment site and comparison site are approximately 60 Kilometers away; each serves a large buyer population (both with three auction halls which can accommodate a total of 450 bidders at the same time). Further, 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 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 centers⁵. This observation helps to alleviate the concern of selection bias associated with our quasi-experimental 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. 2013).

⁴The 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.

⁵Based on our interview with the wholesalers, such location preference has to do with the logistics cost.




Figure 2: Overview of the research design. In addition to data from the experiment period (Nov. 19 - Dec. 7), we also obtained data before (Oct. 29 - Nov. 16) and after (Dec. 10 - Dec. 21).

4 Data

Our dataset 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 (for example, 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); (7) clock specification (for example, clock stand and currency unit).

Table 2 gives a stylized example of a sequence of transactions from our dataset. Due to space constraint, we do not include all the 22 attributes but a set of representative attributes. In this example, a lot⁶ containing 18 units gets sold. Note that the sales prices are not monotonically decreasing or increasing⁷. Also, unlike the 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 a lot.

⁶A lot is a bundle of homogeneous products.

⁷Van den Berg et al. (2001) show empirical evidence for declining price anomaly in the flower auctions; however, if we look at individual auctions, the price trend is inconclusive.

Table 2: A Sample Entry in a Logbook (Auctioneer’s Decision Variables are Italicized).

Transaction Time	Seller ID	Flower ID	Stems		<i>Minimum</i>	<i>Starting</i>	Buyer ID	Purchase Units	Price (cent)	Online
			Per Unit	Available Units	<i>Purchase Units</i>	<i>Price</i> (cent)				
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

To control for product heterogeneity, we selected the product group with the highest transaction amount, namely, Chrysanthemum spray white/yellow GP. Since 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 comparison site.

We performed preliminary analysis to examine the market-level characteristics of both auction sites. The results are reported in Table 3. For brevity, we use Pre-, Exp- and Post- to denote the pre-, during- and post-experiment period. The average number of auctioned lots (per week) was quite stable at both sites. In terms of the number of unique (winning) bidders, there was a small decrease during the post-experiment period, which might have to do with the fact that there were only two weeks during this period (note that both pre- and during-experiment period consisted of three weeks). 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.

Table 3: Market-Level Characteristics at Treatment Site and Comparison Site.

	Treatment Site			Comparison Site		
	Pre-	Exp-	Post-	Pre-	Exp-	Post-
Average Number of Auctions (Per Week)	273	283	278	290	286	296
Total Number of Winning Bidders	323	321	303	389	381	360
Usage of Online Channel (%)	92	93	92	75	78	76

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-value < 0.05) during the experiment period. Specifically, the average price increased by approximately 18% at the treatment site and 11% at the comparison site. Table 4 and 5 summarize the auction-level descriptive statistics from the treatment site and comparison site, respectively.

As we mentioned above, theoretically speaking, bidders are free to enter any auction held

Table 4: Descriptive Statistics at Auction Level from Treatment Site.

Statistics	No. of Win. Bidders			Price (cent)			Purchase units		
	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
Std.dev.	4.2	4.1	3.9	9.3	7.1	9.7	19.8	16.7	15.9
Min.	1	1	1	6	10	5	1	1	1
Max.	31	32	19	62	66	78	347	264	180

Table 5: Descriptive Statistics at Auction Level from Comparison Site.

Statistics	No. of Win. Bidders			Price (cent)			Purchase units		
	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
Std.dev.	4.9	5.3	4.5	10.6	8.4	13.0	18.7	17.3	19.1
Min.	1	1	1	5	5	5	1	1	1
Max.	33	34	27	70	87	81	371	384	396

at the six sites in person (via offline channel) or remotely (via online channel). Thus we would like to know whether bidders had switched between the two auction sites during the study period, or perhaps they were simultaneously bidding across both sites. Note that 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: (*i*) one company may own several IDs across different auction sites; (*ii*) the same ID number across different sites may refer to different companies. In light of this, we requested the list of the registered buyers (companies) as well as the latest allocation information of bidder IDs at the treatment and comparison 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 was participating 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 (thus unobservable in our data) might have participated in the auctions across the two sites. Nevertheless, the presence of these bidders is less of 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 use a difference-in-differences (DID) approach (Bertrand et al. 2004) to estimate the causal effect of the policy change in identity disclosure. By measuring the difference in differences between the treatment site and control site over time, it allows us to control for the characteristics that were 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 between November 19 to December 7, 2012. 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 due to market trends at both sites throughout the study period.

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

$$\begin{aligned} \ln(\text{Price}_{i,j,t}) = & \beta_0 + \beta_1 \text{Treat}_{ij} \times \text{Experiment}_t + \beta_2 \text{Treat}_{ij} \times \text{Post}_t \\ & + \beta_3 \text{Treat}_{ij} + \theta_1 \text{Experiment}_t + \theta_2 \text{Post}_t + \gamma \mathbf{X}_{ijt} + \epsilon_{ijt} \end{aligned} \quad (1)$$

In Equation 1, $\ln(\text{Price})$ is the natural log of the winning price. i indexes each transaction, j indexes each auction, t indexes the time (week) in the study period. Treat_{ij} , Experiment_t and Post_t are dummy variables: Treat_{ij} equals one if the transaction was from the treatment site; Experiment_t equals one during the experiment period and zero otherwise; Post_t equals one during the post-experiment period and zero otherwise. \mathbf{X}_{ijt} a vector of control variables. Specifically, to account for product heterogeneity, we include the key product characteristics *Stem Length*, *Bundling Condition*, *Blooming Stage* and 17 grower dummies. To account for variation from the supply side, we include *Lot Size* and *Minimum Purchase Quantity*. In addition, we control for the day of the week for each transaction, which may have implications for the market demand. The error term ϵ_{ijt} 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 before and after the policy change at the treatment site.

Note that we can also replace the 5th and 6th terms in Equation 1, (i.e., Experiment_t and Post_t) by a set of more detailed week fixed effects θ_t . Essentially, these week fixed effects serve as non-parametric 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):

$$\begin{aligned} \ln(\text{Price}_{i,j,t}) = & \beta_0 + \beta_1 \text{Treat}_{ij} \times \text{Experiment}_t + \beta_2 \text{Treat}_{ij} \times \text{Post}_t \\ & + \beta_3 \text{Treat}_{ij} + \theta_t + \gamma \mathbf{X}_{ijt} + \epsilon_{ijt} \end{aligned} \quad (2)$$


5.1.1 The Parallel-Trend Assumption.

It is worth mentioning that one of the key underlying assumptions of the DID approach is that the difference between the treatment and comparison site would remain constant over time in the absence of the treatment (Abadie 2005). If something other than the treatment changes at one site but not the other, such parallel-trend assumption would be violated, in which case we have no guarantee that the above DID estimator is unbiased. With this in mind, we first check whether the auction price exhibits parallel trends before the policy change at the treatment site and comparison site. Following Autor (2003), we use a leads and lags model to explore the pre-experiment time trends at the two sites. Specifically, we create 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 estimate a model similar to Equation 2 that replaces Experiment_t and Post_t with the two lead and three lag indicator variables. In this case, the reference group is the week immediately before the experiment period (i.e., Nov.12 - Nov.16). We find the coefficients corresponding to the two lead indicator variables are not statistically significant (p-value > 0.2). This suggests that there are no differential trends across the two sites before the experiment. Figure 3 also provide qualitative evidence of parallel trends between the unadjusted prices at the treatment site and comparison site during the three weeks before the policy change at the treatment site.

Before proceeding to the estimation results, we would like to discuss the assumption about the variance-covariance matrix of the error term ϵ_{ijt} in Equation 1 and 2. One possibility would be assuming these errors are independent and identically distributed (IID). Unfortunately, a Breusch-Pagan test rejects the hypothesis that errors are homoskedastic across auctions (p-value < 0.001), and a Breusch-Godfrey test rejects the hypothesis of no first-order autocorrelation (p-value < 0.001). To address the serial correlation problem (Bertrand et al. 2004), we use heteroskedastic-robust standard errors clustered at 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.



(a) Price trend by week.



(b) Price trend by day.

Figure 3: Price trends based on raw data that are unadjusted for covariates or fixed effects. Data are collapsed to week and day level. The vertical bars correspond to the standard errors of means.

5.2 Effect of Withholding Winner’s Identity on Auction Price

Table 6 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 (i.e., $Treatment \times Experiment$) is positive and significant (coefficient = 0.065, p-value = 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 comparison site is no different to 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-value = 0.592). Further, we note that all the control variables except for *LotSize* have a significant impact on the winning price. Specifically, we find 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 (3.7% to 11.2%) than Monday. Finally, although the results from the preliminary analysis (see Table 4 and 5) suggests there might be a difference in average winning price between the treatment and comparison site, such 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 effects model (Column 2) are qualitatively similar with those from the baseline model. Most importantly, the average winning price increased significantly (coefficient = 0.062, p-value = 0.001) at the treatment site during the experiment period. By contrast, the difference in winning prices between the treatment and comparison site during

the post-experiment period was no different relative to the pre-experiment period.

Table 6: Average Effect of Withholding Winner’s Identity on Winning Price

Variable	(1) Baseline	(2) Time Fixed Effects
<i>Treatment</i> × <i>Experiment</i>	0.065 ** (0.001)	0.062 ** (0.001)
<i>Treatment</i> × <i>Post</i>	0.015 (0.592)	0.016 (0.565)
<i>Treatment</i>	0.043 (0.412)	0.031 (0.551)
<i>Experiment</i>	0.137 *** (0.000)	
<i>Post</i>	-0.091 *** (0.000)	
<i>LotSize</i>	0.002 (0.164)	0.002 (0.259)
<i>MinimumPurchaseQuantity</i>	-0.025 *** (0.000)	-0.024 *** (0.000)
<i>StemLength</i>	-0.018 * (0.037)	-0.024 ** (0.005)
<i>BundlingCondition</i>	0.003 ** (0.003)	0.002 ** (0.007)
<i>BloomingStage</i>	0.081 *** (0.000)	0.084 *** (0.000)
<i>Tuesday</i>	-0.037 * (0.011)	-0.037 * (0.011)
<i>Wednesday</i>	-0.082 *** (0.000)	-0.082 *** (0.000)
<i>Thursday</i>	-0.112 *** (0.000)	-0.114 *** (0.000)
<i>Friday</i>	-0.052 *** (0.000)	-0.056 *** (0.000)
<i>Week Fixed Effects</i>		Yes
<i>Grower Fixed Effects</i>	Yes	Yes
Observations	31,848	31,848
Adjusted R-squared	0.337	0.409

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

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

$$\begin{aligned} \ln(AvgPrice_{j,t}) = & \beta_0 + \beta_1 Treat_j \times Experiment_t + \beta_2 Treat_j \times Post_t \\ & + \beta_3 Treat_j + \theta_t + \gamma \mathbf{X}_{jt} + \epsilon_{jt}, \end{aligned} \tag{3}$$

$$\begin{aligned} \ln(\textit{WeightedAvgPrice}_{j,t}) = & \beta_0 + \beta_1 \textit{Treat}_j \times \textit{Experiment}_t + \beta_2 \textit{Treat}_j \times \textit{Post}_t \\ & + \beta_3 \textit{Treat}_j + \theta_t + \gamma \mathbf{X}_{jt} + \epsilon_{jt}, \end{aligned} \quad (4)$$

where $\textit{AvgPrice}_{j,t}$ is the simple average of the winning prices in auction j at time t , and $\textit{WeightedAvgPrice}_{j,t}$ is the weighted average which accounts for the relative purchase amount associated with each transaction. To address the concern about serial correlation, we use heteroskedastic-robust standard errors clustered at day level.

Table 7 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 6, except for *LotSize* and *StemLength*. The former exhibits significant effect on the winning price based on the alternative model specifications, although such effect is of negligible magnitude (0.3%); the latter has no significant effect on the (weighted) average price at auction level.

5.3 Robustness Checks

We conducted a series of checks to assess the robustness of the observed effect. To start with, we performed a falsification test to examine the possibility that the observed effect of policy change is due to spurious correlations. To do so, we randomly indicated which transactions were from the treatment site and re-estimated Equation 2. Because \textit{Treat}_{ij} is replaced by a new dummy variable which does not reflect the reality of whether transaction i was from a treatment or comparison site, we expect the coefficient which captures the true treatment effect to be insignificant. We ran 100 falsification regressions using the above randomization procedure and found in no instance that the treatment effect is statistically significant. This helps alleviate the concern that the observed treatment effect from Table 6 is pure coincidence.

Further, we note that not all the bidders participated in the auctions every day during the study 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 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*

Table 7: Estimation Results with Alternative Model Specifications

Variable	(1) Average Price	(2) Weighted Average Price
<i>Treatment</i> × <i>Experiment</i>	0.070 *** (0.000)	0.073 *** (0.000)
<i>Treatment</i> × <i>Post</i>	0.024 (0.311)	0.027 (0.264)
<i>Treatment</i>	0.033 (0.458)	0.039 (0.398)
<i>LotSize</i>	0.003 * (0.042)	0.003 * (0.029)
<i>StemLength</i>	0.004 (0.479)	0.004 (0.479)
<i>BundlingCondition</i>	0.010 *** (0.000)	0.010 *** (0.000)
<i>BloomingStage</i>	0.055 *** (0.000)	0.055 *** (0.000)
<i>Tuesday</i>	-0.034 ** (0.009)	-0.040 ** (0.002)
<i>Wednesday</i>	-0.068 *** (0.000)	-0.071 *** (0.000)
<i>Thursday</i>	-0.106 *** (0.000)	-0.106 *** (0.000)
<i>Friday</i>	-0.062 *** (0.000)	-0.064 *** (0.000)
<i>Week</i> Fixed Effects	Yes	Yes
<i>Grower</i> Fixed Effects	Yes	Yes
Observations	4,545	4,545
Adjusted R-squared	0.438	0.438

Notes. All the coefficients are estimated at the auction level. Column 1 and 2 summarize the estimation results from Model (3) and Model (4), respectively. Standard errors (in parentheses) are heteroskedasticity robust and clustered by day. *p-value <0.05, **p-value <0.01, ***p-value < 0.001.

effect or novelty effect (Adair 1984). To investigate these alternative explanations to our main finding, we created two sub-samples which include only the transactions from bidders who had participated in the auctions throughout the eight-week study period at the treatment site and comparison site, respectively. As we did not find any bidders who had switched from one auction site to the other during the study period (see discussion in the preliminary analysis of the data in Section 4), we then merged the two sub-samples and re-estimated Equation 1 and 2 on this combined sample. If the price change were mainly due to private shocks or Hawthorne effect, we would expect to see a significant reduction in the magnitude of the treatment effect. The results presented in Table 8 show that the treatment effect remains positive and significant, and the estimated coefficients are qualitatively similar to the ones shown in Table 6. Such observation provides evidence against the aforementioned

alternative explanations.

Table 8: Treatment Effect on Active Bidders

Variable	(1) Baseline	(2) Time Fixed Effects
<i>Treatment</i> × <i>Experiment</i>	0.072 *** (0.000)	0.067 *** (0.000)
<i>Treatment</i> × <i>Post</i>	0.011 (0.706)	0.011 (0.692)
<i>Treatment</i>	0.024 (0.646)	0.008 (0.867)
<i>Experiment</i>	0.128 (0.000)	
<i>Post</i>	-0.087 (0.000)	
<i>LotSize</i>	0.002 (0.284)	0.001 (0.456)
<i>MinimumPurchaseQuantity</i>	-0.025 *** (0.000)	-0.024 *** (0.000)
<i>StemLength</i>	-0.025 ** (0.006)	-0.032 *** (0.000)
<i>BundlingCondition</i>	0.001 (0.114)	0.001 (0.171)
<i>BloomingStage</i>	0.077 *** (0.000)	0.080 *** (0.000)
<i>Tuesday</i>	-0.043 ** (0.004)	-0.043 ** (0.004)
<i>Wednesday</i>	-0.084 *** (0.000)	-0.084 *** (0.000)
<i>Thursday</i>	-0.115 *** (0.000)	-0.116 *** (0.000)
<i>Friday</i>	-0.055 *** (0.000)	-0.060 *** (0.000)
<i>Week Fixed Effects</i>		Yes
<i>Grower Fixed Effects</i>	Yes	Yes
Observations	26,524	26,524
Adjusted R-squared	0.338	0.409

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

Finally, we also investigated whether the treatment effect varies across the market channels. Note that compared to online bidders, offline (onsite) bidders can also acquire market state information (Koppius 2002) through implicit or explicit verbal communications in the auction hall. As such, they may have an informational advantage in the bidding competition. If this were the case, one would expect to see a significant difference in the magnitude of the treatment effect between offline and online bidders. With this in mind, we re-estimated

Equation 2 for transactions⁸ via the offline (onsite) and online channel, separately. The results are shown in Table 9, and they are largely consistent with the full sample estimates⁹ presented in Table 6.

Table 9: Treatment Effect on Bidders from Different Channels

Variable	(1) Offline	(2) Online
<i>Treatment</i> × <i>Experiment</i>	0.058 * (0.041)	0.062 ** (0.001)
<i>Treatment</i> × <i>Post</i>	0.021 (0.571)	0.013 (0.635)
<i>Treatment</i>	0.003 (0.967)	0.028 (0.598)
<i>LotSize</i>	0.001 (0.539)	0.002 (0.332)
<i>MinimumPurchaseQuantity</i>	-0.027 *** (0.000)	-0.024 *** (0.000)
<i>StemLength</i>	-0.015 (0.228)	-0.027 ** (0.002)
<i>BundlingCondition</i>	0.002 (0.148)	0.003 ** (0.004)
<i>BloomingStage</i>	0.083 *** (0.000)	0.082 *** (0.000)
<i>Tuesday</i>	-0.027 (0.189)	-0.038 ** (0.007)
<i>Wednesday</i>	-0.062 ** (0.002)	-0.086 *** (0.000)
<i>Thursday</i>	-0.109 *** (0.000)	-0.114 *** (0.000)
<i>Friday</i>	-0.089 *** (0.000)	-0.050 *** (0.000)
<i>Week Fixed Effects</i>	Yes	Yes
<i>Grower Fixed Effects</i>	Yes	Yes
Observations	5,074	26,752
Adjusted R-squared	0.413	0.410

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

⁸Among the 31,848 transactions within the original sample, there were 22 transactions where the channel usage information was missing. Thus we excluded them for the estimation.

⁹To examine whether the treatment has differential effect across the two channels, we also estimated a model by including the three-way interactive term $Treat \times Experiment \times Online$ as well as the two-way interactive terms $Treat \times Online$ and $Experiment \times Online$ to Equation 2. The estimated coefficient for the three-way interactive term is not significant.

5.4 Effect of Withholding Winner’s Identity on Price Dynamics

Given the conflicting results on price trends in sequential auctions (see Section 2.2), we would like to find out whether withholding winner’s identity has any impact on the price dynamics. To begin with, we plotted the price trend during pre-, during- and post-experiment period, 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 4 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 2nd round, the rank number is 2. The vertical bar denotes one standard deviation of the normalized prices in each round. Overall, we can see that the winning price exhibits a declining trend in all the three time periods. However, such trend seems to get 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 period. Further, we can find that the variances of the normalized prices decrease considerably during the experiment period. Note that although Figure 4 provides qualitative evidence that withholding winner’s identity can mitigate the price declining trend in sequential rounds, such result could be misleading since we have not controlled for any potential confounding factors.




Figure 4: Comparison of price trends. The vertical bars correspond to the standard errors of means.

To measure the treatment effect on price dynamics, we adapt the model from the seminal

work of Van den Berg et al. (2001) which also looks at the sequential sales at the DFA:

$$\begin{aligned} \log \frac{P_{k,j,t}}{P_{k-1,j,t}} = & \mu_0 + \mu_1 \textit{Experiment}_t \times (\textit{Available}_{k,j} - 2) + \mu_2 \textit{Experiment}_t \times (k - 2) \\ & + \mu_3 \textit{Experiment}_t + \mu_4 (\textit{Available}_{k,j} - 2) + \mu_5 (k - 2) + \epsilon_{k,j}. \end{aligned} \quad (5)$$

In Equation 5, j indexes each auction, k indexes the rank within auction j , and t indexes the time. $\textit{Available}_{k,j}$ stands for the available units before the k -th transaction at auction j . $\textit{Experiment}_t$ is a dummy variable which equals one during the experiment period and zero otherwise. Our coefficient of interest is μ_3 , as it captures the main effect of the policy change during the experiment period on the price dynamics in sequential rounds. The dependent variable is the difference of log prices in consecutive rounds. As discussed in 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. To account for the potential curvilinear relationship between the rank number and the dependent variable, we also estimate an alternative model as follows:

$$\begin{aligned} \log \frac{P_{k,j,t}}{P_{k-1,j,t}} = & \mu_0 + \mu_1 \textit{Experiment}_t \times (\textit{Available}_{k,j} - 2) + \mu_2 \textit{Experiment}_t \times (k - 2) \\ & + \mu_3 \textit{Experiment}_t + \mu_4 (\textit{Available}_{k,j} - 2) + \mu_5 (k - 2) + \mu_6 (k - 2)^2 + \epsilon_{k,j}. \end{aligned} \quad (6)$$

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 the variation of prices in sequential rounds. We estimated Model (5) and Model (6) using data from the treatment site¹⁰. The results are presented in Table 10. The estimates of the coefficients largely confirm our observations from Figure 4. Specifically, under both specifications, the estimated coefficient of the intercept is negative and significant (coefficient ≈ -0.025 , p-value < 0.001), suggesting that there is indeed a declining trend in price. The coefficient corresponding to the treatment effect (i.e., $\textit{Experiment}$) is positive and significant (coefficient = 0.011, p-value < 0.001), indicating that withholding winner's identity has a mitigation effect on the price declining trend. Further, we note that the magnitude of the price decline tends to decrease over time (i.e., the coefficient for \textit{Rank} is positive and significant), and the treatment does not impact

¹⁰We have also tried estimating another model by adding a square root term, $(k - 2)^{1/2}$, to Equation 6. The results are qualitatively similar to those shown in Table 10.

such relationship (i.e., the coefficient for $Experiment \times Rank$ is not statistically significant).

Table 10: Effect of Identity Withholding on Price Dynamics

Variables	(1) Baseline	(2) Alternative
(Intercept)	-0.024 *** (0.000)	-0.028 *** (0.000)
$Experiment \times Available$	0.000 (0.710)	0.000 (0.718)
$Experiment \times Rank$	0.000 (0.061)	0.000 (0.060)
$Experiment$	0.011 *** (0.000)	0.011 *** (0.000)
$Available$	0.000 (0.651)	0.000 (0.850)
$Rank$	0.001 *** (0.000)	0.003 *** (0.000)
$Rank^2$		0.000 *** (0.000)

Notes. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions. *p-value < 0.05, **p-value < 0.01, ***p-value < 0.001.

Before moving to the investigation of the potential mechanisms that may have led to our findings, we would like to discuss the economic significance of the observed treatment effects briefly. First, we consider the effects on the expected revenue. The difference-in-differences analyses above reveal that all else being equal, withholding winner’s identity increases the winning price at the auction level by approximately 7% (see Column 2 in Table 7). For chrysanthemums alone, for which the current annual turnover is about €300 million, such increase implies an extra €20 million in the expected revenue. Another way to benchmark the observed effect is to compare our estimated effects with results from the prior literature on the optimal design of auctions. Specifically, it is interesting to see how much difference a good information disclosure policy can make relative to setting optimal pricing mechanisms in similar environments. Using counterfactual simulations based on advertiser bidding data from Microsoft Advertising Exchange, Celis et al. (2014) demonstrate a hybrid mechanism, *buy-it-now or take-a-chance*, can increase the expected revenue by 4.4%. This indicates that the revenue effect that we find in the experiment are comparable to the potential benefits of optimizing the current auction mechanism. Finally, it is worthwhile to mention 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 ¹¹.

6 Understanding the Effects of Identity Withholding

In the previous section, we used the DID approach to show that the policy change (i.e., withholding winner’s identity) could increase the winning price in the sequential auctions of the DFA by approximately 6% at the transaction level. In this section, we complement these analyses by exploring why withholding winner’s identity would lead to such price increase.

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

A second potential mechanism is that withholding winner’s identity disrupts bidders’ *imitation heuristic* and deters strategic bid shading (Zeithammer 2007). Imitation is an attractive heuristic when decision makers – for example, bidders in the DFA – have little information about the strategic environment but can observe others’ success. When using such heuristic, a bidder would refer to a subset of his competitors, i.e., *fellow* competitors, and imitate the strategy of the most successful one (Schlag 1998, Selten et al. 2005). Note that public disclosure of 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 his reference group. As a result, bidders are more likely to blindly follow the crowd, which could partly explain the increased price stability observed in Section 5.4.

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 of them may serve as a better explanation. Specifically, if the leading explanation is the mitigation of tacit collusion, we would see a differential increase in the average price paid by the *collusive* bidders. By contrast, if the leading explanation is the disruption of the imitation heuristic, we would see a differential increase in the average price paid by the

¹¹See <https://www.royalfloraholland.com/en/> for more details about the current initiatives taken by the market makers in the DFA

imitating bidders.

6.1 Mitigation of Tacit Collusion

To examine whether withholding winner’s identity has differential effect on collusive and non-collusive bidders, we first need to differentiate these two types of bidders. Unfortunately, current literature on tacit collusion does not provide clear guidance to distinguish collusive bids from non-collusive bids (Bajari and Yeo 2009). In light of this, we propose a descriptive method to detect potential bid rotation (i.e., one of the cartel members is bidding on behalf of others) at the treatment site by examining the variation of bidder’s behavioral patterns (reflected in the transaction data) between the pre-experiment and experiment period. As such, we focus on the bidders who were actively bidding throughout this 6-week time frame, i.e., 288 bidders at the treatment site and 334 bidders at the comparison site.

For both the treatment and control site, we first calculated the co-occurrence frequency of any two active bidders – the frequency that the two won in the same auction – during the pre-experiment and experiment period, respectively. If two bidders at the treatment site were colluding via bid rotation during the pre-experiment period, we would expect to observe an increase of their co-occurrence frequency during the experiment period, as the treatment – withholding winner’s identity – makes it difficult or impossible to identify defectors. Following this rationale, we select the pairs which exhibited substantial increase in their co-occurrence frequency from the pre-experiment to experiment period. This results in 59 suspected colluders at the treatment site. For the comparison site, we also selected the bidder pairs which exhibited the most significant increase in the co-occurrence frequency from the pre-experiment to experiment period. However, unlike the treatment site, the 66 bidders found in these pairs were unlikely to have engaged in bid rotation, because if they had, the co-occurrence frequency associated with these bidders would remain low (the winner’s ID was visible and thereby could still be used to maintain the cartel).

We then estimate the following fixed-effect log-linear model at the transaction level on the sub-sample which consists transactions from active bidders during pre-experiment and experiment period:

$$\begin{aligned} \ln(\text{Price}_{i,j,t}) = & \lambda_0 + \lambda_1 \text{Treat}_{ij} \times \text{Experiment}_t \times \text{CoOccurrencePlus}_{ij} + \lambda_2 \text{Treat}_{ij} \times \text{Experiment}_t \\ & + \lambda_3 \text{Experiment}_t \times \text{CoOccurrencePlus}_{ij} + \lambda_4 \text{Treatment}_{ij} \times \text{CoOccurrencePlus}_{ij} \\ & + \lambda_5 \text{CoOccurrencePlus}_{ij} + \lambda_6 \text{Treatment}_{ij} + \theta_t + \gamma \mathbf{X}_{ijt} + \epsilon_{ijt} \end{aligned} \quad (7)$$

In Equation 7, $\text{CoOccurrencePlus}_{ij}$ is a dummy variable which equals one if the transaction was made by a bidder among one of those 125 (= 59 + 66) bidders mentioned above

(i.e., bidders involved in pairs which associated with a substantial increase of co-occurrence frequency) and zero otherwise. All other variables remain the same as in Equation 2. Table 11 presents the estimation results. The coefficient of the three-way interaction term ($Treatment \times Experiment \times CoOccurrencePlus$) is not significant (p-value = 0.264), whereas the coefficient for the main treatment effect ($Treatment \times Experiment$) remains positive and significant (coefficient = 0.06, p-value = 0.001). We note that if the treatment effect were mainly due to the mitigation of tacit collusion, we would expect the coefficient for the three-way interaction term to be positive and significant. However, this is not the case.

Table 11: Differential Treatment Effect on Collusive and Non-Collusive Bidders

Variable	Estimates
$Treatment \times Experiment \times CoOccurrencePlus$	-0.045 (0.264)
$Treatment \times Experiment$	0.060 (0.001) **
$Experiment \times CoOccurrencePlus$	0.056 (0.071)
$Treatment \times CoOccurrencePlus$	0.046 (0.094)
$CoOccurrencePlus$	-0.137 (0.000) ***
$Treatment$	0.068 (0.268)
Observations	24,165
Adjusted R-squared	0.451

Notes. All the coefficients are estimated at the transaction level. For brevity, we have omitted the estimates for control variables and time fixed effects as the estimates are consistent with the ones in Table 1 and 2. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions. *p-value < 0.05, **p-value < 0.01, ***p-value < 0.001.

6.2 Disruption of Imitation Heuristic

Similar as above, to assess whether withholding winner’s identity has a differential effect on the price paid by a bidder who initially relied on imitation heuristic and one who did not, we first need to identify the imitating bidders – bidders who repeatedly followed a small subset of their competitors. Intuitively, if two bidders, a and b , frequently won in the same auction (not necessarily in consecutive rounds), and a always won after b , it is likely that a was imitating b . Further, we note that an essential prerequisite for the imitation heuristic is that bidders can identify their *fellow* competitors – the ones worthwhile to imitate. Hence we would expect to find a significant difference in the bidding pattern of an imitating bidder between the pre-experiment and experiment period. Specifically, while the dyadic, follower-follower pairs associated with any imitating bidder (i.e., “follower”) would involve a small, fixed subset of competitors (i.e., “followees”) during the pre-experiment period, such dyadic pairs would include a random subset of competitors during the experiment period. In other words, the *dispersion* of the occurrence of the directed dyads, which is measured by the

ratio of the standard deviation and mean of the occurrence of these dyads, would be smaller during the experiment period than the pre-experiment period.

Following this rationale, we calculated the decrease in the dispersion from the pre-experiment to experiment period for the active bidders during this 6-week time window, and ranked these bidders with respect to this measure (i.e., decrease in the dispersion). Figure 5 illustrates the distribution of active bidders at the treatment site where a positive value in the x-axis indicates a decrease in the dispersion whereas a negative value an increase in the dispersion (from the pre-experiment to experiment period). For both the treatment and comparison site, we picked the top 20% bidders associated with the highest decrease in the dispersion, 58 from the treatment site and 67 from the comparison site. It should be noted, as in the above discussion of tacit collusion, that while the 58 bidders from the treatment site were likely to be imitating bidders, the 67 bidders from the comparison site were unlikely to have followed such heuristic.

We estimate a model similar to Equation 7 that replaces the dummy variable *CoOccurrencePlus* with a new dummy variable *DispersionMinus* which equals one if the transaction was made by one of the 125 bidders (i.t., those who were associated with a substantial decrease of dispersion) and zero otherwise. The estimation results are presented in Table 12. The coefficient of the three-way interaction term (*Treatment* \times *Experiment* \times *DispersionMinus*) is positive (0.047) and significant (p-value=0.001), suggesting that bidders who relied on imitative heuristics were paying significantly more than others under the treatment condition. Further, we note that the same bidders from the treatment site were, all else being equal, paying less prior to the experiment, suggesting that the imitative heuristic worked well for them at that time. Such finding lends support to the hypothesis that the price increase may primarily stem from the disruption of imitation heuristic.

Table 12: Differential Treatment Effects on Imitative and Non-Imitative Bidders

Variable	Estimates
<i>Treatment</i> \times <i>Experiment</i> \times <i>DispersionMinus</i>	0.047 (0.001) **
<i>Treatment</i> \times <i>Experiment</i>	0.049 (0.010) *
<i>Experiment</i> \times <i>DispersionMinus</i>	-0.006 (0.484)
<i>Treatment</i> \times <i>DispersionMinus</i>	-0.044 (0.000) ***
<i>DispersionMinus</i>	0.018 (0.006) **
<i>Treatment</i>	0.080 (0.194)
Observations	24,165
Adjusted R-squared	0.448

Notes. All the coefficients are estimated at the transaction level. For brevity, we have omitted the estimates for control variables and time fixed effects as the results are consistent with the ones in Table 1 and 2. Standard errors (in parentheses) are heteroskedasticity robust and clustered by auctions. *p-value <0.05, **p-value <0.01, ***p-value < 0.001.

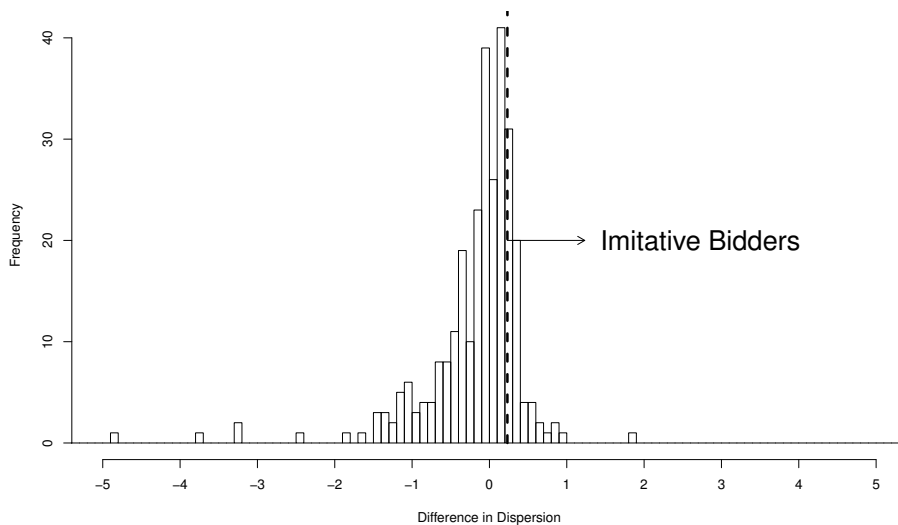


Figure 5: Distribution of the change in dispersion at the treatment site. The occurrence of the directed dyads for bidders on the right side is less dispersed (more random) during the experiment period than the pre-/post-experiment period.

7 Discussion

Since Milgrom and Weber (1982)’s seminal paper, the linkage principle – sellers can increase the expected revenue by disclosing more information – has long been used as the guide to auction design. However, prior studies from both lab and field provide mixed results. In this study, we empirically show that sellers do not necessarily benefit from the commitment to disclosing all the available information and thereby challenges the common wisdom that the more information, the better. Specifically, we show that withholding winner’s identity in sequential auctions can significant increase the average price and thereby raise the seller’s revenue.

Previous studies examining the information disclosure problem are largely based on game-theoretical models with strong assumptions (e.g., bidders are fully rational) or controlled lab experiments where bidders do not possess much experience. We, on the other hand, bring the richness of real-world operation environments while maintaining a high level of control through a quasi-natural field experiment design, combined with a difference-in-differences identification strategy.

While our primary goal of this study is to examine the effects of the policy change in identity closure on market performance, we also provide a theoretical explanation 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, i.e., imitating the successful winners from certain reference groups (Selten et al. 2005) in previous rounds and strategically shade their bids. Withholding winner's identity makes it difficult for bidders to keep track of the members from their reference groups and thereby disrupt such 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, most of the existing research has restricted attention to 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 the revealed winners' information. In other words, the price system does facilitate the information aggregation in the market. However, it does this imperfectly.

Second, we conducted a large-scale field experiment to examine the performance of alternative information disclosure policies in a dynamic, complex B2B market where exchanges of goods are facilitated through sequential Dutch auctions. While the importance of field work has been well acknowledged (Athey et al. 2011), current research has primarily focused on single-unit auctions (Cho et al. 2014, Tadelis and Zettelmeyer 2015). To the best of our knowledge, this is the first study which examines the information disclosure problem in sequential auctions using field data.

Third, we find that withholding winner's identity can significantly mitigate the declining price trend in sequential rounds. Further, 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 new light on the well-known declining price anomaly in sequential auctions (Van den Berg et al. 2001).

Finally, our study adds to the literature on behavioral biases in complex decision environments. While 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 upon the recent literature on behavioral

economics (Gigerenzer and Selten 2002, Weibull 1997), we find that bidders are likely to use imitation heuristic in sequential auctions when winner’s identity is publicly disclosed.

7.2 Implication

Our study also provides useful implications for practitioners. First, our finding confirms that information feedback in repeated competition has a significant impact on the outcome (Ockenfels and Selten 2005). In particular, given our finding that withholding winner’s identity can raise 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. Further, given that humans tend to rely on heuristics (e.g., 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 the technological developments, the auction house still embraces traditions going back more than a century, and its methods have changed little. The results from the current study suggest the current auction rules on and information disclosure has ample room for improvement. Specifically, we show that by removing the winner’s identity from the auction clock, auctioneers can increase the average price by approximately 6%. Such policy change can also effectively mitigates the declining price and thereby increase the 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 as well as the market processes.

To this point, we would like to point out that although the current study mainly focuses on the DFA, the results also provide insights to practitioners in other B2B markets. Specifically, for many procurement auctions, the disclosure of competitors’ identities also play a critical role. Thus the current 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. 1). The empirical findings from the current study suggest that there are important dynamics at play in the 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-drive approach in 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 causal inferences with the current data. If there were a demand shock of another product which 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 has focused 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 the post-auction trades. For example, when winners’ identities are not publicly disclosed, it becomes much more difficult for the customers in the down-stream 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 his profit margin and thereby affect his bidding strategies in the auctions. An integrated model which takes into account the post-auction competition in the down-stream market can be very helpful in understanding the impact of different disclosure policies in the whole supply chain.

Third, in the current study, we have explored two alternative 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. While 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.

References

- Abadie, A. 2005. Semiparametric difference-in-differences estimators. *The Review of Economic Studies* **72**(1) 1–19.
- Abraham, I., S. Athey, M. Babaioff, M. Grubb. 2013. Peaches, lemons, and cookies: designing auction markets with dispersed information. *ACM Conference on Electronic Commerce*. 7–8.
- Adair, J. G. 1984. The hawthorne effect: A reconsideration of the methodological artifact. *Journal of applied psychology* **69**(2) 334–345.
- Angrist, J. D., J. Pischke. 2008. *Mostly harmless econometrics: An empiricist’s companion*. Princeton university press.
- Athey, S., J. Levin, E. Seira. 2011. Comparing open and sealed bid auctions: Evidence from timber auctions. *The Quarterly Journal of Economics* **126** 207–257.
- Autor, D. H. 2003. Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of Labor Economics* **21**(1) 1–42.
- Bajari, P., J. Yeo. 2009. Auction design and tacit collusion in fcc spectrum auctions. *Information Economics and Policy* **21**(2) 90–100.
- Beggs, A., K. Graddy. 1997. Declining values and the afternoon effect: evidence from art auctions. *RAND Journal of Economic* **28**(3) 544–565.
- Bergemann, D., J. Hörner. 2010. Should auctions be transparent? Discussion paper, Cowles Foundation Discussion Paper 1764.
- Bertrand, M., E. Duflo, S. Mullainathan. 2004. How much should we trust differences-in-differences estimates? *The Quarterly journal of economics* **119**(1) 249–275.
- Bichler, M., A. Gupta, W. Ketter. 2010. Designing smart markets. *Information Systems Research* **21**(4) 688–699.
- Camerer, C. 2003. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- Cason, T. N., K. N. Kannan, R. Siebert. 2011. An experimental study of information revelation policies in sequential auctions. *Management Science* **57**(4) 667–688.
- Celis, E., G. Lewis, M. Mobius, H. Nazerzadeh. 2014. Buy-it-now or take-a-chance: Price discrimination through randomized auctions. *Management Science* **60**(12) 2927–2948.
- Cho, S., H. Paarsch, J. Rust. 2014. Is the “linkage principle” valid? evidence from the field. *The Journal of Industrial Economics* **62**(2) 346–375.
- Cramton, P., J. Schwartz. 2000. Collusive bidding: Lessons from the fcc spectrum auctions. *Journal of regulatory Economics* **17**(3) 229–252.
- Donald, S. G., H. J. Paarsch, J. Robert. 2006. An empirical model of the multi-unit, sequential, clock auction. *Journal of Applied Econometrics* **21**(8) 1221–1247.

- Gigerenzer, G., R. Selten. 2002. *Bounded rationality: The adaptive toolbox*. MIT press.
- Goeree, J., T. Offerman. 2002. Efficiency in auctions with private and common values: An experimental study. *American Economic Review* **92**(3) 625–643.
- Harrington, J. E. 2012. A theory of tacit collusion. Tech. rep., The Johns Hopkins University, Department of Economics.
- Haruvy, E., E Katok. 2013. Increasing revenue by decreasing information in procurement auctions. *Production and Operations Management* **22**(1) 19–35.
- Imbens, G. W., J. M. Wooldridge. 2009. Recent developments in the econometrics of program evaluation. *Journal of economic literature* **47**(1) 5–86.
- Jap, S. 2007. The impact of online reverse auction design on buyer-supplier relationships. *Journal of Marketing* **71**(1) 146–159.
- Jehiel, P., B. Moldovanu. 1996. Strategic nonparticipation. *The Rand Journal of Economics* 84–98.
- Jeitschko, T. D. 1998. Learning in sequential auctions. *Southern Economic Journal* **65**(1) 98–112.
- Kagel, J. H., D. Levin. 1986. The winner’s curse and public information in common value auctions. *The American Economic Review* **76**(5) 894–920.
- Koppius, O. R. 2002. Information architecture and electronic market performance. Ph.D. thesis, Erasmus University, Rotterdam.
- Lacetera, N., B. J Larsen, D. G. Pope, J. R. Sydnor. 2013. Bid takers or market makers? the effect of auctioneers on auction outcomes. Tech. rep., National Bureau of Economic Research.
- Lu, Y., A. Gupta, W. Ketter, E. Van Heck. 2016. Exploring bidder heterogeneity in multichannel sequential b2b auctions: Evidence from the dutch flower auctions. *MIS Quarterly* **40**(3) 645–662.
- Marshall, R., L. Marx. 2009. The vulnerability of auctions to bidder collusion. *The Quarterly Journal of Economics* **124**(2) 883–910.
- McAfee, R., D. Vincent. 1993. The declining price anomaly. *Journal of Economic Theory* **60** 191–212.
- Milgrom, Paul, Robert J Weber. 1982. The value of information in a sealed-bid auction. *Journal of Mathematical Economics* **10**(1) 105–114.
- Ockenfels, A., R. Selten. 2005. Impulse balance equilibrium and feedback in first price auctions. *Games and Economic Behavior* **51**(1) 155–170.
- Perry, M., P. Reny. 1999. On the failure of the linkage principle in multi-unit auctions. *Econometrica* **67**(4) 895–900.
- Roth, A. E. 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. *Journal of Economic Theory* **78**(1) 130–156.

- Selten, R., K. Abbink, R. Cox. 2005. Learning direction theory and the winner's curse. *Experimental Economics* **8**(1) 5–20.
- Shadish, W. R., T. D. Cook, D. T. Campbell. 2002. *Experimental and quasi-experimental designs for generalized causal inference*. Houghton and Mifflin, Boston.
- Sherstyuk, K., J. Dulatre. 2008. Market performance and collusion in sequential and simultaneous multi-object auctions: evidence from an ascending auctions experiment. *International Journal of Industrial Organization* **26**(2) 557–572.
- Smith, C. 2007. Markets as definitional practices. *The Canadian Journal of Sociology* **32**(1) 1–39.
- Snir, E., L. Hitt. 2003. Costly bidding in online markets for it services. *Management Science* **49**(11) 1504–1520.
- Tadelis, S., F. Zettelmeyer. 2015. Information disclosure as a matching mechanism: Theory and evidence from a field experiment. *American Economic Review* **105**(2) 886–905.
- Tu, Z. 2007. Why do we use the dutch auction to sell flowers: Information disclosure in sequential auctions. Discussion Paper, University of Pittsburg.
- Van den Berg, G. J., J. Van Ours, M. Pradhan. 2001. The declining price anomaly in dutch dutch rose auctions. *The American Economic Review* **91**(4) 1055–1062.
- Varma, G. 2002. Standard auctions with identity-dependent externalities. *RAND Journal of Economics* **33**(4) 689–708.
- Weber, R. 1983. Multiple-object auctions. R. Engelbrecht-Wiggans, M. Shubik, R. Stark, eds., *Auctions, Bidding, and Contracting: Uses and Theory*. New York University Press, New York, NY, 165–191.
- Weibull, J. W. 1997. *Evolutionary game theory*. MIT press.
- Zeithammer, R. 2007. Research note—strategic bid-shading and sequential auctioning with learning from past prices. *Management Science* **53**(9) 1510–1519.