


An Auction-Based Test of Private
Information in an Interdealer FX
Market

By: Pietro Bonaldi
Mauricio Villamizar-Villegas

Borradores de ECONOMÍA



No. 1049
2018



tá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Colombia - Bogotá - Col



An Auction-Based Test of Private Information in an Interdealer FX Market*

Pietro Bonaldi[†]

Mauricio Villamizar-Villegas[‡]

The opinions contained in this document are the sole responsibility of the authors and do not commit Banco de la República nor its Board of Directors

Abstract

There are several financial markets where dealers trade a large share of total volume, while also having access to periodic auctions of the same asset conducted by a third party. For such a market, we derive a test of private information about the value of the asset that combines data on both bidding behavior and market trades. Our approach is to test for private versus common values, as defined in auction theory. We use changes in trading prices of extreme bidders before and after the auction to test the null hypothesis of private values (no private information) against the alternative of common values (private information). Additionally, we use a regression discontinuity design where we compare the behavior of dealers bidding right below and right above the auction's cutoff price to control for inventory effects, understood here as decreasing marginal valuations as functions of inventory. Our case study are foreign exchange auctions conducted by the Central Bank of Colombia during the period 2008-2014, and the corresponding interdealer market for the Colombian peso against the US dollar. Overall, the data does not reject the null hypothesis of private values. Specifically, information about other bidders' valuations has no significant effect on trading prices, not even shortly after the auction takes place.

JEL Codes: C57, D44, F31, G14

Key Words: Auctions, Common Values, Private Values, Private Information, Foreign Exchange Market, Regression Discontinuity Design

*We are thankful to Burton Hollifield, Diego Rojas, Thomas Ruchti, and seminar participants at Banco de la República for valuable comments. We also thank Felipe García, Joselo Peña, and Maria Alejandra Ruiz for their research assistance.

[†]Carnegie Mellon University, Tepper School of Business, email: pbonaldi@andrew.cmu.edu

[‡]Banco de la República Colombia, email: mvillavi@banrep.gov.co

Una Prueba de Información Privada Basada en Subastas en un Mercado Cambiario

Pietro Bonaldi

Mauricio Villamizar-Villegas

Las opiniones contenidas en el presente documento son responsabilidad exclusiva de los autores y no comprometen al Banco de la República ni a su Junta Directiva

Abstract

Existen varios mercados financieros en los que algunos intermediarios tranzan una gran parte del volumen total, además de tener acceso a subastas periódicas del mismo activo realizadas por un tercero. Para dicho mercado, derivamos una prueba de información privada sobre el valor del activo que combina datos sobre el comportamiento de las pujas y las transacciones de mercado. Nuestro enfoque es probar valores privados versus valores comunes, tal como se define en la teoría de subastas. Utilizamos cambios en los precios de mercado de los postores extremos antes y después de la subasta para probar la hipótesis nula de valores privados (sin información privada) frente a la alternativa de valores comunes (información privada). Además, utilizamos un diseño de regresión discontinua en el que comparamos el comportamiento de los postores que ofertan justo debajo y encima del precio de corte de la subasta para controlar por posibles efectos de inventario. Nuestro caso de estudio son las subastas de divisas realizadas por el Banco de la República de Colombia durante el periodo 2008-2014, y el mercado cambiario correspondiente de pesos-dólar. En general, nuestra prueba no rechaza la hipótesis nula de valores privados. En particular, la información revelada a los postores sobre sus valoraciones relativas no tiene un efecto sobre los precios de mercado, ni siquiera justo después del momento de la subasta.

Códigos JEL: C57, D44, F31, G14

Palabras Clave: Subastas, Valores Comunes, Valores Privados, Información Privada, Mercado Cambiario, Regresión Discontinua

1 Introduction

In this paper we derive an empirical test to shed new light on an old but ongoing debate in finance: do market participants (dealers, traders) have private information that is not fully reflected in market prices? In brief, we test for private information as defined in Fama (1991). We revisit this question in the context of the interdealer market for a specific currency, namely, the Colombian peso (COP). We start from the observation that the recent widespread use of electronic trading platforms in foreign exchange (FX) markets has increased transparency and reduced trading costs (King et al., 2013), which could in principle alter the results of previous research that has provided empirical evidence of private information in what used to be highly opaque over-the-counter (OTC) FX markets

As pointed out by Ito et al. (1998), in principle public information could be expected to fully determine prices in foreign exchange markets since, in contrast with equity markets, there is no such thing as inside information. And yet these authors, together with Evans and Lyons (2002), and Evans (2002) among others, provide empirical evidence of price-relevant private information in FX markets. However, they all analyze data from highly opaque OTC markets, where the details of any trade (price, amount, quotes, etc.) are only observed by the two counterparts involved. Evans (2002) argues that this lack of transparency is precisely the reason why there is a price distribution in equilibrium reflecting heterogeneous information, and yet, arbitrage opportunities do not arise.

Although the FX market is still the largest OTC market worldwide, with a daily average turnover of \$5.1 trillion dollars (Bank of International Settlements, 2016), trades are now predominantly settled at lower costs through electronic trading platforms, and the details of such trades are known in real-time by any dealer trading on those platforms. King et al. (2013) and Rime and Schrimpf (2013) document how recent technological changes have decreased search costs and increased the velocity of trading.¹ Given such increase in transparency and subsequent reduction in trading costs, the question that we address in this paper is whether evidence of private information remains in FX markets. We focus on a close-to-centralized market, where a single electronic platform covers roughly 95% of all interdealer spot trades, in volume, for a specific currency pair.

Our test exploits the availability of detailed data on market transactions as well as bidding behavior in an auction for the same asset, and could be applied to any other asset besides currency. Following standard definitions from auction theory, our analysis can also be described as a test of

¹Institutions today may trade with dealers electronically through Bloomberg Tradebook, Reuters, or through multi-bank platforms and electronic brokerage systems such as Electronic Broking Services (EBS) and Reuters Matching. Note that EBS and Reuters Matching were once exclusive to interdealer trading but opened up to individual costumers in 2004 in order to compete with other multi-bank platforms such as FXall, Currenex, or Hotspot.

private versus common values. In a private values model (PV) with incomplete information, bidders know how much the asset is worth to them, i.e. their valuation, but ignore the valuations of other bidders. We emphasize that PV allows for a common component in all valuations (e.g. the asset’s fundamental value or its resale price), but in such case PV imposes the strong assumption that all bidders share the same information about it. Hence, all variation in valuations across bidders is purely idiosyncratic. Alternatively, in a common values model (CV), bidders are uncertain about the common value. Instead, they form different expectations based on privately known information. In sum, the main difference between the two models, and the one that we exploit to derive the test, is whether information about the common component is shared by all bidders (PV) or dispersed across them (CV).

We note that a particular case of the PV model is when all bidders’ expected value of the common component coincides with an observable market price. On the contrary, under the CV model information is dispersed (usually in the form of privately observed signals), hence not all information available to market participants is already contained in the price. If it were, there would be no disagreement among rational bidders who know the price. Namely, CV is equivalent to private information. Moreover, private information rejects the strong version of the efficient market hypothesis (EMH) which states that prices incorporate all relevant information, including that which is not publicly available. As expressed in Grossman (1976), a market where the price system fails to perfectly aggregate information dispersed across several traders is not “informationally efficient”. Therefore, our test could also be interpreted as a test of market efficiency.

A key feature of the PV model is that knowledge of other bidders’ valuations should not affect a bidder’s own valuation. Intuitively, PV assumes that bidders value the asset differently, not because they have different information about its common value, but because their valuation of the asset is determined by characteristics that might be distinct among bidders. In contrast, in a CV model bidders update their valuations whenever they receive information about other bidders’ valuations. For instance, learning that an opponent expects the common value to be higher increases a bidder’s own valuation. An ideal test of PV versus CV would then compare bidders’ expected valuations before and after they learn about other bidders’ valuations. Keeping everything else constant, any change in a bidder’s valuation after receiving such information would reject PV. However, valuations are generally not observable to researchers (except for experimental settings), and are usually recovered structurally from a model that imposes assumptions on whether values are common or private.

Instead, and with the use of a simple model of trading, we derive a test that requires data on market transactions taking place right before and after an auction of the same asset, and involving extreme bidders —those placing either the lowest or highest bids—. Specifically, for a given dealer bidding in a given auction, we estimate the causal effect of learning that she is an extreme bidder

on her price adjustment thereafter. Under the PV model, we expect no systematic price change, since this information should have no effect on her idiosyncratic private valuation. Contrarily, under CV, we expect a positive (negative) price adjustment after she learns that she submitted the lowest (highest) bid. The rationale for the latter is the same argument explaining the winner's curse in common value auctions. Intuitively, bidders who learn that all other bids are lower (higher) than theirs should realize that their ex-ante valuation is too high (low), and update it accordingly. Such changes in valuations should then be reflected in the prices at which these bidders are willing to buy or sell the asset, which in turn are at least partially revealed by their actual trading prices. All else equal, higher valuations imply higher prices. Since valuations are essentially unobservable, we focus instead on changes in prices.

Our paper is closely related to the empirical auction literature. Hortaçsu and Kastl (2012) test for private versus interdependent values in multi-unit auctions of Canadian treasury bills. Their test is based on the fact that dealers are allowed to modify their bids at a given auction after observing their customers' bids. Using a model of bidding, Hortaçsu and Kastl estimate the dealers' values under the null hypothesis of private values and conclude that the change in dealers' bids after observing their customers' is not attributable to dealers updating their beliefs about the fundamental value of the asset. That is, similar to our findings, their test does not reject the null hypothesis of private values. Haile et al. (2003) also develop nonparametric tests for private versus common values, focusing on first-price sealed-bid auctions. Their test only requires observing the bids, and exploiting variation in the number of bidders that participate at each auction. Based on previous work by Bajari and Hortaçsu (2003) and Hendricks et al. (2003), they test for bidding behavior that is consistent with bidders anticipating the winner's curse, since the latter only arises in common value auctions. McAfee et al. (1999) use resale prices as ex-post measures of values and compare them to winning bids to test their auction model which implies that, on average, highest winning bids must be higher than values. The data reject such implication, however the authors conclude that it does not provide enough evidence to reject a private values null hypothesis against common values. As McAfee et al. acknowledge, testing the model is tantamount to a joint hypotheses test, thus it is hard to further assess which hypotheses fails when the model is rejected.

Our test differs from the previous ones in several respects. To begin with, the test we propose does not depend on a specific model of bidding, hence it does not require additional assumptions regarding the bidders payoff functions (e.g. risk neutrality) or the independence of their idiosyncratic values. Moreover, it is generally agnostic with respect to the bidders' equilibrium beliefs and strategies. Therefore, we do not face the joint hypothesis problem in McAfee et al. (1999). The only restriction we impose on the auction equilibrium for the baseline specification of our test is the assumption that the lowest bidder always has a lower expected valuation of the asset than the highest bidder. Several models of bidding imply this as a result. As opposed to Hortaçsu and Kastl

(2012), our data does not allow us to observe bidders modifying their bids, nor learning about other bids before submitting their final bid. Instead, we analyze changes in pricing decisions that might result from learning about other bidders' valuations once the results of the auction are announced to the bidders. Interestingly, we find similar null results regarding whether dealers update their valuations, although in an entirely different market where previous literature has highlighted the role of private information.

Ideally, our test should depend exclusively on price changes induced by knowledge of the auction results, from where bidders learn about their relative valuations. This is challenging because prices might be changing concurrently for several other reasons. For instance, new information about the common value of the asset might become available to the bidders, or other market conditions might be subject to exogenous shocks (e.g. shocks to liquidity or to external demand for the asset). To control for these common factors likely affecting all bidders, we choose a difference-in-differences estimator for the baseline specification of our test. More precisely, we compare the change in the price at which the lowest bidder trades before and after the auction, with that of the highest bidder. The underlying (parallel trends) assumption, necessary for identification of the causal effect of the information released on prices, is that both extreme bidders are, on average across all auctions, equally affected by any common shocks.

A more worrying confounding factor would be that the auction itself affected prices through channels other than the information released about relative valuations, but yet in different ways for different bidders. Primarily, the auction reallocates the asset among bidders and the auctioneer. That is, winners at the auction consequently change their inventory or holdings of the asset while losers keep theirs constant. Such a change might have a direct effect on their marginal valuation if the latter decreases with inventory. Since the auction affects inventories in different ways for the two extreme bidders, this could imply a violation of the parallel trends assumption in our baseline specification, even when extreme bidders do not respond differently to common shocks.²

To address this concern, we use a Regression Discontinuity Design (RDD). We have data from procurement auctions where all bidders bidding at or below a cutoff price are winners, in the sense that the auctioneer buys from them some or all of their offered amounts. Correspondingly, bidders right above the cutoff are losers, that is, the auctioneer buys nothing from them. This introduces a discontinuity in the change in the asset inventory as a function of the price bid. Two bidders with almost identical price bids, one at the cutoff and the other right above it, change their holdings of the asset quite differently as a result of the auction, even though their bids reveal their intention of

²Standard models of multi-unit auctions assume non-increasing marginal valuations (see McAdams, 2008 for a discussion on this assumption). Moreover, inventory effects have been extensively documented in the market microstructure literature. In particular, Lyons (1995) provides a model and empirical evidence of such effects in FX markets.

selling the asset at almost the same price. Such discontinuity produces variation in asset holdings that is exogenous to the bidders' asset valuations. In fact, it is as good as randomly assigned at the limit as bids approach the cutoff price. We exploit such variation to identify and estimate changes in interdealer market prices induced by changes in inventories. In particular, we use a fuzzy RDD to estimate the slope of the marginal valuation as a function of asset holdings. We use this estimate to control for potential inventory effects in the main specification of our test. Roughly, our test statistic is comprised of the difference between price changes of the extreme bidders induced by the auction, net of any inventory effect.

We note that our proposed estimation of inventory effects can be interpreted as a test of the decreasing marginal valuations assumption common in multi-unit auctions. Moreover, it allows us to identify the causal effect of marginal variation in quantities on a variable of interest (in our case, prices in the interdealer market). Hence, it constitutes an additional (and stand-alone) contribution of our work. Very few previous studies have used an RDD approach to exploit the quasi-experimental variation induced by an auction. Among the few, Kawai and Nakabayashi (2015) use this approach to detect collusion in consecutive procurement auctions, when bids at the first auction fail to meet a secret reserve price. The authors focus on failed auctions where the bids of the lowest and the second lowest bidders are very close to each other, and test whether they preserve their order at the second auction. Kong (2017) uses a similar strategy to distinguish affiliation and synergy across sequential auctions of similar or related objects (contracts). She compares the bids at a second auction of bidders who were marginal winners and losers at the first auction, to determine whether observed correlation in bids is explain by affiliation (roughly, value correlation) or synergy (roughly, complementarities in winning both auctions). In these two previous studies the treatment assigned is a binary status (i.e., winning or losing the auction). We extend this methodology to a multi-unit auction, where the quantity sold to, or purchased from the winners is a continuous variable.

Our case study centers on foreign exchange auctions conducted by the Central Bank of Colombia during the period of 2008-2014. For each multi-unit uniform clearing price auction (1,098 in total), we observe all bids along with the bidders' identities, the clearing (cutoff) price, and the US dollar amount purchased by the Central Bank. We also have data on over 2 million interdealer tic-by-tic FX transactions taken from *SET-ICAP FX S.A.*, an electronic trading platform in charge of administrating the largest Colombian peso against the US dollar spot market. A back-of-the-envelope calculation based on the entirety of the COP-USD market, as reported in Pérez et al. (2015), suggests that the platform providing our data covers over 95% of the total spot market share. Moreover, we have unique identification numbers (NIT) for all dealers and bidders in our sample, which allow us to unambiguously match all bidders to their corresponding transactions in the interdealer market. After restricting the sample to fit our research design, we are left with 180 - 465 auctions (depending on the particular specification considered), corresponding to those

dates when we observe all extreme bidders trading in time windows of fixed duration (namely, 30, 60, 90 or 120 minutes) both before and after the auction.

Our main findings indicate that dealers do not update their valuations after receiving information on other dealers' valuations. That is, our test does not reject the null hypothesis of private values (no private information). When focusing on trading occurring 60 minutes before and after the auction, and comparing the trading prices of the lowest and highest bidders only, we estimate differences in price changes of 0.11 and -0.22 COP/USD for sales and purchases of US dollars, respectively. These magnitudes are rather small (the average exchange rate in our sample is 1,867 COP/USD), and none of them are statistically different from zero. In contrast, the average difference between the highest and lowest bids across all auctions in the sample is 1.9 COP/USD. We find very similar results when using time windows of 30, 90 and 120 minutes. In all but one case, the test does not reject the null hypothesis at the 5% significance level, and the point estimates of price changes range from -0.22 to 0.18 COP/USD. The only specification of the test where the null hypothesis is rejected is one where we do not control for inventory effects despite the fact that the first stage RDD yields a negative and statistically significant estimate of the slope of the marginal valuation as a function of inventory.

Overall, these null results contrast with our RDD estimation of inventory effects. When looking at changes in buying prices of dealers bidding right below the cutoff price compared to those bidding right above, we find a statistically significant difference of 0.31 COP/USD. Hence, our data provides enough statistical power to identify (admittedly small) effects of the auctions on buying prices. Our test of private values strongly suggests that such difference is not attributable to dealers updating their valuations after learning the auction results.

Prima facie, our results seem to be at odds with previous literature. Menkhoff et al. (2016) find that the order flow of long-term demand side investment managers (mutual and pension funds) predicts persistent shifts in exchange rates out-of-sample, when looking at data from 2001 to 2011 across several currency pairs. A likely explanation that has been extensively explored in the market microstructure literature is that order flow contains private information about fundamentals (e.g. Evans and Lyons (2005), Evans and Lyons (2008) and Chinn and Moore (2011)). We attribute the seeming discrepancy to the centralization of the COP-USD spot market. Transparency accelerates information aggregation, hence, even if some customers' order flow provide private information to their dealers, this information is rapidly reflected in the interdealer market price. Lyons (1997) and Evans and Lyons (2002), among others, argue that interdealer order flow aggregates information dispersed across dealers, and that interdealer trades are at least partially informative to counterparties. In a centralized market, all trades are observable to all dealers, hence we find that all remaining variation in dealers' valuations at any point in time is purely idiosyncratic, i.e. it does not reflect different information about fundamentals. Put differently, in a centralized FX market we

find evidence that the cross-sectional variation in dealers' valuations is not explained by private information. This does not imply an absence of information in customers' order flow. Potentially, the market might be aggregating such information efficiently.

Our paper proceeds as follows. In Section 2 we derive our test of private information and explicitly state its underlying assumptions. Further, we provide a procedure through which we control for potential confounding effects with an RD design. Section 3 describes our data and the contextual characteristics of auctions conducted by the Central Bank of Colombia. Finally, Section 4 presents the results of our test when applied to the Colombian peso against the US dollar interdealer market.

2 Test

2.1 Test of private information

In this section we present a stylized model of marginal valuations and market prices and we use it to derive our test of private information.

For any given dealer, we define net inventory as the sum of current holdings of the asset (FX) and signed pending orders from customers. While customer orders to sell increase a dealer's inventory, orders to buy decrease it. That way, if a dealer buys or sells some amount of the asset on behalf of a customer (to execute a customer's order), we still interpret the transaction as having an effect on the dealers' inventory. This definition allows us to model dealers' marginal valuations as functions of both holdings and orders from clients in a simple way, as follows:

For any dealer $i \in \{1, \dots, n\}$ with inventory x , we assume her marginal valuation at any time t is given by

$$V_{i,t}(x) = \beta_i(x) + C_t + \epsilon_{i,t} \quad (1)$$

where $\beta_i(\cdot)$ is a non-increasing function capturing how marginal valuations change with inventory, C_t is common to all dealers, and $\epsilon_{i,t}$ is purely idiosyncratic, has mean zero, and is independent of C_t . We assume bidder i observes $\epsilon_{i,t}$ directly, but not C_t . Instead, i has private information about C_t , consisting of a scalar signal $s_{i,t}$ that is independent of the idiosyncratic valuation. Hence,

$$E[V_{i,t}(x) | s_{i,t}] = \beta_i(x) + E[C_t | s_{i,t}] + \epsilon_{i,t} \quad (2)$$

Our goal is to test whether dealers form different expectations of the common component C_t . More

precisely, we want to test the null hypothesis

$$H_0 : E [C_t | s_{i,t}] = P_t, \text{ for all } i \in \mathcal{D} \quad (3)$$

which states that, regardless of dealers' private information, they all agree on their expectation of the common component. A particular case is when such expectation coincides with an observable market price (P_t). This is consistent with the absence of private information, since all relevant information would already be contained in the price. In the auction literature, H_0 is commonly referred to as a private values (PV). The key feature of such a PV model is that if the signals of other dealers were revealed to bidder i , such information would not change i 's valuation. In contrast, bidders having common values (CV) would update their valuations if they knew other bidders' private signals. Correspondingly, we consider the alternative hypothesis:

$$H_1 : E [C_t | s_{1,t}, \dots, s_{n,t}] \text{ is strictly increasing in } s_{j,t}, \text{ for each } j \in \{1, \dots, n\} \quad (4)$$

where H_1 allows for disagreement in bidders' expectations $E [C_t | s_{i,t}]$, provided some of them do not share the same information (their signals differ). Moreover, it assumes that all signals are informative, since higher signals are assumed to imply higher expected C_t and, thus, higher valuations. The key feature of the CV model is that learning that an opponent has a higher signal increases i 's valuation.³

2.2 A hypothetical experiment


Let us start by describing a hypothetical experiment to test H_0 against H_1 . Let \mathcal{B} be the set of n bidders that participate in the experiment. In the first stage, the experimenter draws a common value C from a known distribution, a signal S_i and an idiosyncratic value ε_i , for each bidder $i \in \mathcal{B}$. The distributions have common support across bidders, but these supports are allowed to be different for S_i and ε_i .⁴ All distributions are common knowledge. Further, all idiosyncratic values have zero mean, are mutually independent, and are also independent of each signal and C . Signals are mutually independent, conditional on C . Also, $E [S_i | C = c] = c$. Bidder i privately observes her private signal s_i , her idiosyncratic value ε_i , and all publicly available information about the common component C . Hence, each bidder i privately learns $v_i = c_i + \varepsilon_i$, where $v_i = E [V_i | S_i = s_i]$ and $c_i = E [C | S_i = s_i]$. Also, for simplicity, we assume for now that inventories are held constant. In the second stage, the experimenter uses an incentive-compatible mechanism to learn v_i , from all $i \in \mathcal{B}$. Then, in the third stage, the experimenter reveals the lowest valuation $v^{(1)} = \min \{v_i : i \in \mathcal{B}\}$ to all

³We follow the definition of private and common values for auction models in Athey and Haile (2002) and Athey and Haile (2007).

⁴Hereafter, we use S_i and ε_i to denote random variables, with respective realizations s_i and ε_i .

bidders. In the fourth stage, bidders update their valuations after observing the extreme valuation. Finally, in the last step, the experimenter learns the updated valuations $\hat{v}_i = \hat{c}_i + \epsilon_i$ of all bidders. These steps are outlined in Figure 1.

Figure 1: Stages of Experiment: Propagation of Information



Under the null hypothesis, $v_i = \hat{v}_i$ for all i . No bidder should change her valuation because the signals of other bidders are not informative of how valuable the asset is for her. However, under the alternative hypothesis, bidders update their valuations. Moreover, if the joint distribution of valuations is absolutely continuous with respect to the Lebesgue measure, $\hat{v}_L > v_L$ for the lowest bidder L , with probability one under H_1 . The underlying logic is similar to that of the winner's curse in common value auctions. If a bidder learns that she had the lowest valuation, she rationally updates her expectation upwards. Given the common support assumption, the addition of idiosyncratic values does not change the conclusion. We include a proof of this statement in Appendix A.

Now let us suppose that in the third stage the experimenter reveals more information, for instance, any function of the vector of all valuations $g(v_1, \dots, v_n)$. In such case, H_1 no longer implies that $\hat{v}_L > v_L$. For instance, suppose that ϵ_L is very low (although it is drawn from a distribution with zero mean), but s_L is higher than the mean of the unconditional distribution of C . If the lowest bidder learns that all other bidders have valuations below this mean, she might update her own downwards with positive probability, despite knowing that she has the lowest valuation. Intuitively, she might attribute having the ex-ante lowest valuation to a low ϵ , rather than to a low signal of C , relative to the other bidders' private values and signals. Whether this is the case, depends on the underlying distributions and the function of observed valuations.

However, for arbitrarily many independent repetitions of the latter version of the experiment, it is still the case that, on average, $\hat{v}_L > v_L$ under H_1 . Proposition 1 states this claim formally. Notice that now L can also be seen as a random variable since in every repetition new signals and private values are independently drawn for each i , hence the bidder with the lowest valuation is determined randomly. We prove Proposition 1 under the mild condition that $E[C|L] = E[C]$ (the

proof is left for Appendix A). Several stronger yet still reasonable conditions imply the former one. One such case is the standard assumption in the empirical auction literature that the distributions of the signals s_i conditional on C are the same for all $i \in \mathcal{B}$.

Proposition 1. *Let L be a discrete random variable taking values in the set \mathcal{B} . For any $i \in \mathcal{B}$, $L = i$ denotes that bidder i has the lowest valuation v_i among all bidders. Also, let $\mathbf{V} = (V_1, \dots, V_n)$ be the vector of all bidders' valuations, and $g(\cdot)$ be an arbitrary function of \mathbf{V} . If $E[C|L] = E[C]$, then,*

$$E\left[E[C|S_i, \varepsilon_i, L = i, g(\mathbf{V})] - E[C|S_i, \varepsilon_i] \mid L = i\right] > 0 \quad (5)$$

2.3 Auction

We focus on a multi-unit, uniform price, reverse (procurement) auction, with only one bid per bidder, because that is precisely the format used in our case study, as we will describe in Section 3.2. However, our test could be easily adapted to fit other auction formats.

The auction provides a setting similar to the hypothetical experiment just described. Namely, there is a set of potential bidders \mathcal{B} that are allowed to bid in every auction. Let $\mathcal{B}_t \subset \mathcal{B}$ be the subset of bidders that participate in auction t . Each $i \in \mathcal{B}_t$ submits a bid $(b_{i,t}, q_{i,t})$, consisting of a number of units to be sold, $q_{i,t}$, and a price per unit $b_{i,t}$. The auctioneer sorts the bids by price in ascending order: $b_t^{(1)} \leq \dots \leq b_t^{(n_t)}$, and purchases the units offered by the lowest bidders, until it reaches a previously announced quantity Q_t .

Right after the auction, the auctioneer reveals to all bidders the lowest and highest bids, $b_t^{(1)}$ and $b_t^{(n_t)}$ respectively, and the cutoff price –the highest bid among winners–. Thus, while one bidder learns that she submitted the lowest bid, another one learns that he submitted the highest one (bidders also learn additional information about the distribution of all bids). Henceforth, we denote these two bidders by L (lowest) and H (highest).⁵ Let $\hat{v}_{L,t}$ and $\hat{v}_{H,t}$ denote the valuations of these two bidders immediately after this information is released. Ceteris paribus, H_0 implies that $v_{L,t} = \hat{v}_{L,t}$ and $v_{H,t} = \hat{v}_{H,t}$. Again, under PV, for any given bidder, information about other dealers's bids, values, or signals, has no effect on their own valuation.

We apply Proposition 1 here to conclude that, under the alternative hypothesis, L and H update their valuations upwards and downwards, respectively. However, we would need to impose a restrictive structure on the auction's equilibrium to guarantee that L and H have not only the

⁵The lowest bidder is always a winner, that is, she sells all her units to the auctioneer. The highest bidder could be either a winner, a partial winner, or a loser. That is, the highest bidder could also sell some or all his offered units $(q_{H,t})$, if total demand Q_t is just below, equal, or higher than total supply $\sum_{i \in \mathcal{B}_t} q_{i,t}$. Otherwise, if $Q_t \leq \sum_{i \in \mathcal{B}_t} q_{i,t} - q_{H,t}$, the auctioneer does not buy any units from the highest bidder.

lowest and highest bids but also the lowest and highest valuations.⁶ Instead of focusing on a specific such equilibrium, we propose an assumption that seems weak enough to encompass a large set of plausible equilibria. Assumption 1 simply states that the valuation of the lowest bidder must be lower than that of the highest bidder. Besides, in any Bayes-Nash equilibrium of the auction all bidders know each others' strategies, and thus can infer information about valuations from the revealed bids.

Assumption 1. *Let L and H be the lowest and highest bidders in the auction described in section 3.1. Let $v_L(q)$ and $v_H(q)$ be their respective expected marginal valuations, that is, their valuations of the q -th unit they bid for at the auction. Then, $v_L(q) < v_H(q)$.*

Under H_1 , Assumption 1 guarantees that, on average across multiple auctions, $v_L(q) < \hat{v}_L(q)$ and $v_H(q) > \hat{v}_H(q)$. The proof is the same as in Proposition 1, but with a minor modification where we let $L = i$ denote that bidder i has the lowest bid, instead of the lowest valuation. Under Assumption 1, bidder i learns that her valuation is lower than the highest bidder's and thus, on average, updates her valuation upwards. Similarly, the highest bidder updates, on average, downwards.

In the ideal experiment previously described, the experimenter tests the null hypothesis by directly comparing valuations before and after the information is revealed. However, valuations are generally not observable. Instead, we see all bidders, including L and H , trading in the interdealer market both before and after the auction. The prices of those transactions provide noisy measures of valuations.

More precisely, we assume that at time t , dealer i 's selling (buying) price is

$$p_{i,t} = E[V_{i,t}(x_{i,t}) | s_{i,t}] + \xi_{i,t} \tag{6}$$

where $\xi_{i,t} \geq 0$ ($\xi_{i,t} \leq 0$) is the seller's (buyer's) marginal surplus. That is, at any moment in time, a dealer sells (buys) some amount only if she gets a price above (below) her current valuation. Any marginal surplus $\xi_{i,t}$ is determined by market conditions at the time of the transaction that cannot be fully anticipated.

Right after the auction, results are revealed to all bidders containing information \mathcal{A}_t . Dealer i 's valuation when submitting her bid is $v_{i,t} = E[V_{i,t}(x_{i,t}) | s_t]$. Once the auction results are released, i 's valuation potentially changes to $\hat{v}_{i,t} = E[V_{i,t}(\hat{x}_{i,t}) | s_{i,t}, \mathcal{A}_t]$. The difference, if any, between $v_{i,t}$

⁶The are still several open questions regarding equilibrium existence and multiplicity in uniform-price multi-unit auctions. In particular, McAdams (2007) shows that if bidders are not risk neutral the existence of an equilibrium in monotone strategies is not guaranteed. In that case, it is not generally true that the highest (lowest) bidder has the highest (lowest) valuation among all bidders.

and $\hat{v}_{i,t}$ could be the result of learning new information, but also the effect of changes in inventory, from $x_{i,t}$ to $\hat{x}_{i,t}$, if dealer i is among the winners. We want to test if some of the average difference $\frac{1}{T} \sum_{t=1}^T \hat{v}_{i,t} - v_{i,t}$ is explained by obtaining information from the auction results. In what follows, we focus on selling prices, but an analogous reasoning holds for buying prices.

Let $p_{i,t}^0$ and $p_{i,t}^1$ be observable per-unit prices at which bidder i sells units in the market $\tau_{i,t}^0$ and $\tau_{i,t}^1$ minutes before and after auction t , respectively. All other variables with a 0 or 1 superscript are similarly defined. For now, we fix an arbitrary auction t and drop the index for ease of notation. The change in price $\Delta p_i = p_i^1 - p_i^0$ can be written as

$$\Delta p_i = (p_i^1 - \hat{v}_i) + (\hat{v}_i - v_i) + (v_i - p_i^0) \quad (7)$$

Moreover,

$$\begin{aligned} p_i^1 - \hat{v}_i &= E [V_i^1(x_i^1) | s_i^1, \mathcal{A}] + \xi_i^1 - E [V_i(\hat{x}_i) | s_i, \mathcal{A}] \\ &= \beta_i(x_i^1) - \beta_i(\hat{x}_i) + E [C^1 | s_i^1, \mathcal{A}] - E [C | s_i, \mathcal{A}] + \epsilon_i^1 - \epsilon_i + \xi_i^1 \end{aligned}$$

Similarly,

$$v_i - p_i^0 = \beta_i(x_i) - \beta_i(x_i^0) + E [C | s_i] - E [C^0 | s_i^0] + \epsilon_i - \epsilon_i^0 - \xi_i^0 \quad (8)$$

Putting these two terms together,

$$\begin{aligned} \Delta p_i &= \beta_i(x_i^1) - \beta_i(x_i^0) \\ &\quad + E [C^1 | s_i^1, \mathcal{A}] - E [C | s_i, \mathcal{A}] + E [C | s_i] - E [C^0 | s_i^0] \\ &\quad + \epsilon_i^1 - \epsilon_i^0 + \xi_i^1 - \xi_i^0 \end{aligned}$$

In order to be able to test the null by comparing prices rather than valuations, we need a few more assumptions.

Assumption 2. Let $W_i \in \{l, h, n\}$ be a random variable that equals l or h if bidder i submits the lowest bid or highest bid, respectively, and n otherwise. Also, let $\eta_i = \epsilon_i^1 - \epsilon_i^0 + \xi_i^1 - \xi_i^0$. For all i , $E[\eta_i | W_i] = 0$ and $E[\eta_i^2 | W_i] = \sigma_\eta^2$.

Assumption 2 requires that the direction of changes in idiosyncratic values and trader's surplus, when comparing transactions right before and after the auction, is not predictable even if we know the auction results. For instance, the highest bidders might have, on average, higher idiosyncratic values at the time of the auction, but under assumption 2, that is not enough to predict the sign of $\epsilon_i^1 - \epsilon_i^0$. Notice that this is weaker than assuming that the distributions of ϵ_i^0 and ϵ_i^1 , and those of ξ_i^0 and ξ_i^1 , are the same, respectively, even after conditioning on W_i .

Assumption 3. For all i , $E[C_t|s_{i,t}]$ and $E[C_t|s_{i,t}, \mathcal{A}]$ follow a random walk.

Assumption 3 implies that new information about the common component is independent of previous shocks and follows a normal distribution. We rely on independence to rule out auto-correlated innovations, however normality is stronger than what we need to obtain the asymptotic distribution of our test statistic. Notice that, under the null hypothesis, Assumption 3 is consistent with currency prices following a random walk (for the specific case where $E[C_t|s_{i,t}]$ is equal to the market price, for all bidders).

Assumption 4. For all i , $\beta_i(\cdot)$ is a linear function of inventory x , with slope $\beta \geq 0$.

Assumption 4 imposes restrictions on how marginal values change with inventory, but allows us to provide a simple estimate of the corresponding function based on a fuzzy regression discontinuity design, as we explain below.

To save on notation, let $c_{i,t} = E[C_t|s_{i,t}]$ and $\hat{c}_{i,t} = E[C_t|s_{i,t}, \mathcal{A}_t]$. Given Assumptions 2 - 4:

$$\begin{aligned} \Delta p_{i,t} &= (\hat{c}_{i,t} - c_{i,t}) + \beta \Delta x_{i,t} + \Pi(\tau_{i,t}) + \eta_{i,t} \\ &= \Delta c_{i,t} + \beta \Delta x_{i,t} + \nu_{i,t} \end{aligned} \tag{9}$$

where, $\Delta x_{i,t} = x_{i,t}^1 - x_{i,t}^0$, $\tau_{i,t} = \tau_{i,t}^0 + \tau_{i,t}^1$ and $\Pi(\tau) \sim \mathcal{N}(0, \tau \sigma_\pi^2)$.

Under the null hypothesis, $\Delta c_{i,t} = \hat{c}_{i,t} - c_{i,t} = 0$ and the composite error, $\nu_{i,t} = \Pi(\tau_{i,t}) + \eta_{i,t}$, has mean zero and variance $\sigma_\eta^2 + \tau_{i,t} \sigma_\pi^2$. It is tempting to regress $\Delta p_{i,t}$ linearly on a constant and $\Delta x_{i,t}$, and test whether the coefficient on the constant (β_0) is statistically different from zero. But such an strategy would fail to test the null for several reasons. First, $\Delta x_{i,t}$ is most likely endogenous, since it is at least partially determined by the dealer based on her knowledge of $\nu_{i,t}$, resulting in a biased estimate of β_0 . Second, $\hat{c}_{i,t} - c_{i,t}$ may have different signs for different i and t (different dealers may updated their expected C_t in different directions), and thus β_0 could be statistically indistinguishable from zero even when some dealers are updating their expectations. In such a case, the test would fail to reject a false null. Additionally, the composite error term $\nu_{i,t}$ is likely heteroscedastic, and this should be accounted for to perform inference correctly. Below we explain with detail how we address each of these concerns.

As we explain with more detail below in section 2.4, to estimate β consistently we use a regression discontinuity design. Intuitively, the auction introduces a discontinuity in inventory as a function of bids. All bidders with bids lower or equal to the cutoff price sell some positive amount to the auctioneer and hence reduce their inventories. In contrast, all bidders above the cutoff keep theirs constant. For those bidders submitting bids approaching the cutoff from both

sides, the auction provides variation in inventory that is exogenous to their valuations. Let $\hat{\beta}_{RD}$ be the corresponding estimate of the slope of the marginal valuation function.

In the main specification of the test, we focus on the extreme bidders L_t and H_t . We model these as discrete random variables with support \mathcal{B} (the set of all potential bidders). That is, for a given dealer i , $L_t = i$ means that i submits the lowest bid at auction t , which is an uncertain outcome for the dealers when submitting their bids. Now let $\Delta\tilde{p}_{i,t} = \Delta p_{i,t} - \beta\Delta x_{i,t}$, since equation (9) holds for all bidders, $\Delta\tilde{p}_L = \Delta c_L + \nu_L$, $\Delta\tilde{p}_H = \Delta c_H + \nu_H$ and then

$$\Delta\tilde{p}_L - \Delta\tilde{p}_H = \Delta c_L - \Delta c_H + \nu_L - \nu_H \quad (10)$$

Under the null $\Delta c_L = \Delta c_H = 0$, with probability one, then obviously $E[\Delta c_L - \Delta c_H] = 0$. Moreover, under the alternative hypothesis of common values, Proposition 1 implies that $E[\Delta c_L - \Delta c_H] > 0$. Let $\mu_c = E[\Delta c_L - \Delta c_H]$ and $\gamma_t = \Delta c_{L_t} - \Delta c_{H_t} - \mu_c$. The test of private vs common values we propose is just

$$\tilde{H}_0 : \mu_c = 0 \text{ against } \tilde{H}_1 : \mu_c > 0. \quad (11)$$

We will now derive a test statistic for (11). Under the null, $\tilde{\gamma}_t = \gamma_t + \nu_{L_t} - \nu_{H_t}$ is a mean zero heteroscedastic error term, and

$$\Delta\tilde{p}_{L_t} - \Delta\tilde{p}_{H_t} = \mu_c + \tilde{\gamma}_t \quad (12)$$

Notice, though, that $\Delta\tilde{p}_{L_t}$ and $\Delta\tilde{p}_{H_t}$ are not directly observable because they depend on β , but they can be estimated by $\hat{\Delta}\tilde{p}_{L_t} = \Delta p_{L_t} - \hat{\beta}_{RD}\Delta x_{L_t}$, and similarly for H_t . This adds an additional error term to the right hand side of equation 12, since in any finite sample $\hat{\beta}_{RD}$ is estimated with error. However, this error term vanishes asymptotically, provided $\hat{\beta}_{RD}$ is a consistent estimator of β . Moreover, its asymptotic distribution is known, as long as we know the distribution of $\hat{\beta}_{RD}$.

It follows that μ_c is the mean of the asymptotic distribution of $\hat{\Delta}\tilde{p}_{L_t} - \hat{\Delta}\tilde{p}_{H_t}$. Therefore, given a random sample $\left\{ \hat{\Delta}\tilde{p}_{L_t} - \hat{\Delta}\tilde{p}_{H_t} \right\}_{t=1}^T$, a natural test statistic for 11 is simply the sample average. However, to derive the asymptotic distribution of this test statistic we need to account for error in the estimation of β .

More precisely, Let $\xi_T = \hat{\beta}_{RD} - \beta$ be the estimation error for a sample of size T . Then,

$$\hat{\Delta}\tilde{p}_{L_t} - \hat{\Delta}\tilde{p}_{H_t} = \mu_c + \tilde{\gamma}_t - \xi_T(\Delta x_{L_t} - \Delta x_{H_t}) \quad (13)$$

Non-parametric estimators of β converge at a rate slower than \sqrt{T} .⁷ In consequence,

⁷For a derivation of the asymptotic distribution of the local-linear RD estimator under different choices of bandwidth see Hahn et al. (2001), Porter (2003), Imbens and Kalyanaraman (2011) and Calonico et al. (2014).

$T^{-1} \sum_{t=1}^T \Delta \hat{p}_{L_t} - \Delta \hat{p}_{H_t}$ converges at the same rate and to the same distribution as ξ_T (up to a multiplicative constant). We state this result precisely in Proposition 2. The proof is included in Appendix A.

Proposition 2. *Let T^r be the rate of convergence of $\hat{\beta}_{RD}$, with $0 < r < \frac{1}{2}$. Then,*

$$T^r \left(\frac{1}{T} \sum_{t=1}^T \Delta \hat{p}_{L_t} - \Delta \hat{p}_{H_t} - \mu_c \right) \xrightarrow{d} \xi E[\Delta x_{L_t} - \Delta x_{H_t}] \quad (14)$$

where ξ is a random variable with the same asymptotic distribution of $T^r(\hat{\beta}_{RD} - \beta)$.

2.4 Addressing confounding effects with RDD

The auctioneer only purchases units from bidders with price bids lower or equal than the cutoff price. Hence, the auction cutoff price introduces a discontinuity in the amount purchased from bidders as a function of their bids. Those bidders situated at or below the cutoff price receive a payment of domestic currency (COP) in exchange for foreign currency (USD), as a direct result of winning the auction, while those marginally above it keep their inventories of foreign exchange constant. By design, there is always at least one bid at the threshold, since the cutoff price is defined as the highest bid among winners. Crucially, the auction results reveal essentially the same information to all bidders within a neighborhood of the threshold about their relative bids, i.e. they all learn that their bids are close or equal to the cutoff. Therefore, subsequent differences in the behavior of marginal winners compared to marginal losers cannot be attributed to such information.

These features pertaining to the auction enable us to identify the causal effects of exogenous variation in inventories (induced by the auction) on subsequent trading prices, by comparing the trading behavior of marginal winners and marginal losers. Intuitively, the treatment assigned by being at or below the cutoff price (inventory change) is as good as randomly assigned when narrowing locally at the threshold. RDD requires the assumption that conditional expectations of potential outcomes are continuous at the cutoff. Namely, the two potential outcomes of interest are the changes in trading prices of bidders if hypothetically they won or lost the auction, regardless of their actual treatment status. Hence, we are implicitly assuming that conditional expected prices do not jump discontinuously as bids cross the cutoff price if, counterfactually, the auction result (who wins and who loses) is held constant.⁸

More precisely, since we are interested in the marginal effect of inventory on prices, we use the auction cutoff price b_t^c to instrument variation in inventory, and identify the effect using the Wald

⁸Hahn et al. (2001) show that, in the RDD context, continuity of the conditional expectations of potential outcomes at the threshold is sufficient for identification of the (local) average treatment effect.

estimator corresponding to a fuzzy RD design, as described in Lee and Lemieux (2010):

$$\beta_{RD} = \frac{\lim_{\delta \uparrow 0} E[\Delta p_{i,t} | b_{i,t} = b_t^c + \delta] - \lim_{\delta \downarrow 0} E[\Delta p_{i,t} | b_{i,t} = b_t^c + \delta]}{\lim_{\delta \uparrow 0} E[\Delta x_{i,t} | b_{i,t} = b_t^c + \delta] - \lim_{\delta \downarrow 0} E[\Delta x_{i,t} | b_{i,t} = b_t^c + \delta]} \quad (15)$$

Under the assumption that $\beta_i(\cdot)$ in equation (1) is a linear function of inventory x , with slope β (for all bidders), β_{RD} identifies β , and can be used to derive a consistent estimator. For estimation, we use the fuzzy local-linear RD with robust confidence intervals of Calonico et al. (2014).⁹

3 Data and Institutional Environment

From June 2008 until December 2014, the Central Bank of Colombia (**CBoC henceforth**) intervened periodically in the FX market with multi-unit uniform price auctions in order to accumulate reserves denominated in US dollars. Through these auctions, the CBoC purchased a total of \$23.9 billion dollars, the largest of all its dollar purchases (\$41.3 billion in total) and sales (\$2.9 billion) since 1999.¹⁰ Nonetheless, the amount purchased by the central bank is still small compared to the roughly \$1 billion dollars that are traded daily in the Colombian interdealer market.

3.1 Data

We obtained the auction data from the Market Operations and Development Department at the CBoC (*Departamento de Operaciones y Desarrollo de Mercados -Mesa de Dinero*). The data contain the specific date and time of each auction (1,099 auctions in total), the resulting bids (prices and amounts) with the identity of each bidder, the quota or upper bound on the amount to be purchased announced by the CBoC, the actual amount purchased after receiving all bids, and the resulting cutoff clearing price.

Our second data source is *SET-ICAP FX S.A.*, a financial institution that administrates the largest Colombian electronic FX market. We note that this platform reports over 90% of the total USD-COP market volume, since offshore trading of USD-COP is restricted by regulation.¹¹ We confirm this with a back-of-the-envelope calculation based on the entirety of the COP-USD market,

⁹Local-polynomial kernel-based RD estimators require the selection of a bandwidth. Conventional bandwidth choice based on cross-validation or mean squared error minimization is likely to result in confidence intervals that are too large and hence lead to substantial over-rejection of the null hypothesis of zero treatment effect. Calonico et al. (2014) improve on these previous inference methods by using a bias-corrected estimator to construct confidence intervals with better coverage.

¹⁰Aside from the multi-unit uniform price auctions studied in this investigation, the CBoC conducted FX trades either directly in the electronic spot market (2004-2007) or through the auctioning of FX options (1999-2008). For a more detailed description of all FX intervention carried out by the CBoC, see the editorial note found in Uribe (2016).

¹¹Regulatory details were obtained from the following central bank's official decree: *Resolución Externa* (2018).

which, as reported in Pérez et al. (2015) averaged a daily trade volume of \$1 billion dollars in 2012. Since the daily average in our data set is \$950 million for 2012, we conclude that it covers roughly 95% of the total market. Data from this source include 2 million tick-by-tick foreign exchange transactions. On average, USD \$784,000 are traded at 1,867 COP/USD in a single transaction. Crucially, the exchange trading data also include the identities of both counterparts (financial institutions) involved in every transaction. A subset of these institutions also bid at the CBoC auctions as primary dealers, and we possess unique identification numbers that allow us to match them unequivocally.¹²

3.2 Colombian FX Auctions

The CBoC preselects a set of FX dealers (mostly private banks) to participate in the USD procurement auctions.¹³ Although the list of potential bidders is publicly available, at any given auction bidding is voluntary and anonymous. Only the CBoC knows the bidders' identities, and this information is never released to the public, nor to the bidders.

Each auction takes place at some point during the trading session, from 8:00am to 1:00pm on business days. The CBoC makes an official announcement on the electronic trading platforms calling for an auction, only two minutes before it starts. The announcement includes information on the maximum amount of foreign currency to be purchased (the quota). However, the CBoC holds the right to buy a lower amount or even render the auction null and void if deemed necessary. Thus, the exact total demand at any given auction is uncertain to the dealers.¹⁴

Each auction lasts for three minutes during which participants can present and modify only one bid, consisting of both a single price (COP/USD) and a total amount of dollars to be sold, as long as it does not exceed 80% of the quota. The minimum amount allowed per bid is \$1 million dollars and all bids have to be submitted in multiples of \$100,000 USD. Once a bid is entered, the CBoC's electronic platform notifies the dealer whether the bid is momentarily: (i) “*in*”, meaning that the price is below the temporary cutoff price, (ii) “*out*”, meaning that the bid is above the cutoff price, or (iii) “*partially in*”, meaning that the bid is equal to the cutoff price.

At the end of each auction, all bids are sorted out in ascending order by price. The CBoC

¹²An overlapping data set was previously used by Kuersteiner et. al (2016a, 2016b) to analyze the effect of sterilized foreign exchange intervention by the CBoC. They include detailed descriptions of the data and the corresponding market. The question we address in this paper is only tangentially related to theirs.

¹³To grant participation in auctions, the CBoC uses criteria based on balance sheet information and the existence of foreign currency deposits within the CBoC. The detailed criteria are found in this regulatory document: *Circular Reglamentaria Externa DODM-143*.

¹⁴Out of a total of 1,099 auctions in our sample, there was only one declared void, on July 11, 2008. We thus analyze the remaining 1,098 valid auctions. Information about the auction's format was obtained from the official central bank's regulatory documents, in particular: *Circular Reglamentaria Externa DODM-143*.

then buys the amount offered by each bidder, starting from the one bidding the lowest price, until the quota is filled. The resulting cutoff price is the highest bid among all dealers from whom the CBoC buys a positive amount. Hence, bidders with a price lower than the cutoff price are the resulting winners, and those above the cutoff are the resulting losers (they do not trade with the CBoC). We refer to bidders with a price equal to the cutoff price as partial winners, since the CBoC only purchases a fraction of the total amount they offer (to avoid exceeding the quota). When there are ties at the cutoff price, rationing is pro-rata on the margin, that is, the amount sold by each bidder is proportional to the amount offered. The auction is uniform, hence the CBoC pays the cutoff price to all winners regardless of their bid.

One and a half minutes after the end of the auction, the central bank communicates to the bidders: (i) the bids with the lowest and highest prices, and (ii) the cutoff price. Note that the almost immediate release of this information is key to our identification strategy, as we rely on extreme bidders learning that their valuations are higher or lower than others. Finally, the auction results are announced later that day through the CBoC's website and other electronic platforms. The announcement includes information on total demand, prices of both extreme bids, total amount purchased, cutoff price and number of bidders. While we observe the identities of all bidders, this information is kept confidential from them.

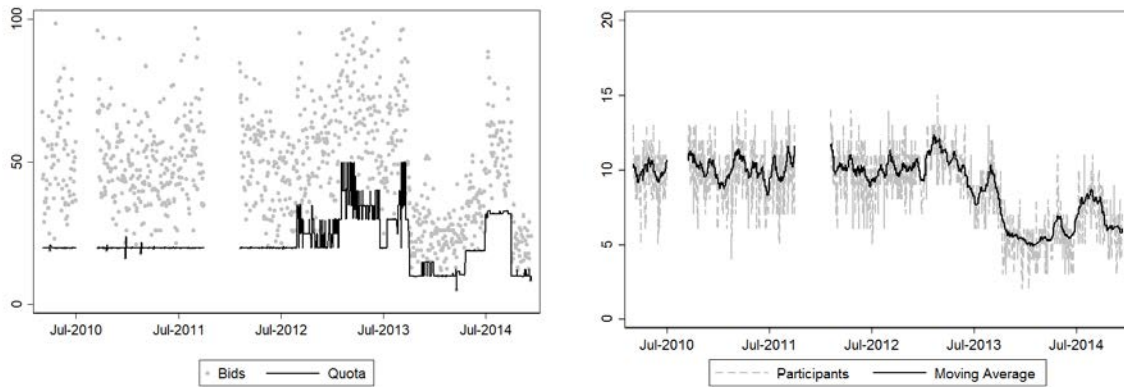
As shown in panel (a) of Figure 2, daily dollar purchases through the CBoC auctions remained stable around 20 million during June 2008 - July 2012. Purchases then ranged from 20-50 million during August 2012 - September 2013, and from 10-33 million during October 2013 - December 2014. With very few exceptions, the total amount offered (to sell) by all bidders is larger than the amount purchased by the CBoC, and the difference exhibits substantial variation across auctions. Panel (b) shows the number of financial institutions bidding at each auction. Of the 17 institutions that participated in at least one auction throughout our sample, there was an average of 10 bidders in every auction; roughly six winners and four losers per auction.

Across all 1,098 auctions in our sample, the median absolute difference between the most extreme bid per auction and the cutoff price is 0.3 COP/USD, and the 99th percentile is 5 COP/USD. However, in a few anomalous cases we see bids that differ from the cutoff price for more than 50 or even 100 COP/USD. Since there is no plausible rationalization of such large deviations, we attribute them to misreports, introduced either when placing the bids or when compiling the dataset. In consequence, we drop 42 auctions containing bids that deviate for more than 10 COP/USD from the cutoff. The resulting sample data thus have 1,056 auctions.

To perform our test of private information, we match each auction with the interdealer market data by bidder and time of each auction. More precisely, an auction is successfully matched to market data only if each extreme bidder purchases or sells US dollars in the interdealer market

within a time window of fixed duration (30, 60, 90 or 120 minutes, respectively), both before and after the auction. As a result, we are left with 161 - 461 auctions, depending on the duration of the time window, and whether we focus on bidders' purchases or sales. Table 1 presents summary statistics for all auctions in the sample (1,056), as well as for auctions successfully matched to dollar sales by extreme bidders 60 minutes before and after the auction (328), and for auctions matched to purchases in time windows of the same duration (335).

Figure 2: Foreign Exchange Intervention through FX Auctions



(a) Amount Auctioned (in USD million)

(b) Number of Participants

Panel (a) shows the total amount offered by all bidders (solid dots) and the amount purchased by the CBoC at each auction. Panel (b) shows the number of bidders at each auction. The solid line is a 1-month moving average.


Table 1: Auction summary statistics

	Mean	Std. Dev	Percentile 25%	Percentile 50%	Percentile 75%
All auctions (1056)					
Max - Min bid ^(b)	1.94	1.46	1	1.5	2.25
Quota ^(a)	22.97	7.87	20	20.1	25.2
Amount offered ^(a)	51.76	19.61	38	50.5	63
Amount purchased ^(a)	22.83	7.84	19.9	20	25.2
Cutoff price ^(b)	1864.32	90.12	1793.6	1839.5	1913.3
Number of bidders	5.19	2.99	3	5	7
Auctions matched to sales (328)					
Max - Min bid ^(b)	1.52	1.2	0.8	1.2	1.77
Quota ^(a)	22.79	10.4	10.5	20.1	30.25
Amount offered ^(a)	47.29	22.03	27.5	46.75	63.25
Amount purchased ^(a)	22.68	10.3	10.4	20	30.25
Cutoff price ^(b)	1893.24	94.35	1813.7	1886.65	1936.7
Number of bidders	9	4.9	5	9	13
Auctions matched to purchases (335)					
Max - Min bid ^(b)	1.58	1.26	0.9	1.25	1.75
Quota ^(a)	23.05	10.28	11.5	20.1	30.2
Amount offered ^(a)	47.95	22.17	29	46.5	63.5
Amount purchased ^(a)	22.88	10.18	11.5	20	30.2
Cutoff price ^(b)	1890.25	96.48	1809.4	1884.2	1933.9
Number of bidders	9	4.9	5	9	13

Authors' calculations: An auction is successfully matched to market data only if we observe each extreme bidder selling (purchasing) US dollars in the interdealer market in the last 60 minutes before the auction announcement and also in the first 60 minutes after the auction results are disclosed to the bidders. We build to separate samples for sales (auctions matched to sales) and purchases (auctions matched to purchases).^(a) Amounts are measured in million USD.^(b) Prices are measured in COP/USD.

Figure 3 shows the distribution of the difference between the highest and the lowest price bids placed at each auction. The mean difference is 1.9 COP/USD across 1,056 auctions. Table 1 reports slightly smaller averages for the restricted samples, after matching the auctions to market data. Overall, this indicates substantial cross-sectional variation in bidders' valuations. In fact, these magnitudes are comparable to the average difference between the highest and lowest trading prices in the interdealer market in any given 10-minute window, which is close to 2 COP/USD (See Table 4 in Appendix B). This is despite the fact that at least some of the time series variation in market prices, even at high frequencies, should be explained by the arrival of new common information. Hence, we take these statistics as further motivation for our test. In particular, it relates to the question of whether there is statistical evidence that the cross-sectional heterogeneity in bids reflects private information.

Figure 3: Distribution of difference between extreme bids



The mean difference between the highest and the lowest price bid across 1,056 auctions is 1.9 COP/USD, with a standard deviation of 1.47 and a median of 1.5.

Table 2 shows different participation and bidding patterns among bidders. For confidentiality reasons we cannot disclose their names. There is substantial variation in participation. The average dealer bids in approximately half of all auctions, while two dealers bid in less than 10%, and two others participate in more than 90%. Conditional on participating, there are also large differences in how likely a bidder is to win or lose the auction, and also how likely it is to place an extreme bid (either highest or lowest). Since our test is based on whether the auction results affect the

valuations of extreme bidders, we also show how frequently a bidder places an extreme bid as a fraction of the total number of auctions. The percentages in the last two columns of Table 2 show that all dealers are effectively included in our test, although disproportionately so.

It is worth emphasizing that our test does not require dealers to draw their valuations from the same distribution (ex-ante homogeneity) or even to bid according to symmetric equilibrium strategies (symmetric equilibrium), hence bidder heterogeneity in valuations or strategies is less of a concern for our test than for a test based on estimating bidders' valuations given a specific auction model. That said, the results in Table 2 might still raise some concerns that the results of our test are mostly driven by a small subset of all dealers. In an untabulated exercise, for each bidder who participated in the auctions, we perform the test of private information after dropping all auctions where such dealer placed the highest (lowest) bid. Overall, the main conclusion of the test is robust to these exclusions. The null hypothesis of no private information is not rejected. This alleviates concerns that the results might be driven by a few dealers, specially those who place extreme bids in a large share of all auctions.

Table 2: Participation and Bidding Patterns at the CBoC Auctions

Dealer	Participated	Won ^(a)	Lost ^(a)	Partially Won ^(a)	Lowest Bid ^(a)	Highest Bid ^(a)	Lowest Bid ^(b)	Highest Bid ^(b)
9	93%	40%	26%	34%	16%	6%	15%	6%
17	91%	44%	32%	25%	8%	5%	7%	4%
20	76%	28%	56%	16%	2%	9%	2%	7%
14	71%	43%	36%	21%	11%	8%	8%	5%
27	65%	74%	12%	14%	36%	3%	24%	2%
23	64%	72%	14%	14%	30%	1%	19%	1%
12	60%	16%	58%	27%	6%	34%	4%	20%
15	51%	37%	40%	22%	15%	11%	7%	5%
16	49%	13%	67%	20%	3%	29%	1%	14%
35	47%	23%	46%	31%	6%	9%	3%	4%
46	46%	28%	56%	16%	4%	11%	2%	5%
55	45%	14%	70%	17%	2%	23%	1%	10%
13	45%	34%	48%	18%	9%	10%	4%	4%
48	42%	28%	62%	10%	5%	17%	2%	7%
26	14%	24%	50%	27%	6%	7%	1%	1%
53	6%	12%	77%	12%	5%	17%	0%	1%
22	5%	24%	70%	6%	4%	22%	0%	1%

Notes: Dealers are sorted by participation. The numbers assigned to each dealer are arbitrary, but match the numeration used elsewhere in the paper. On average, a dealer participated in 540 out of 1,056 auctions.

^(a) Measured as a share of all auctions where the bank participated.

^(b) Measured as a share of all auctions in the sample.

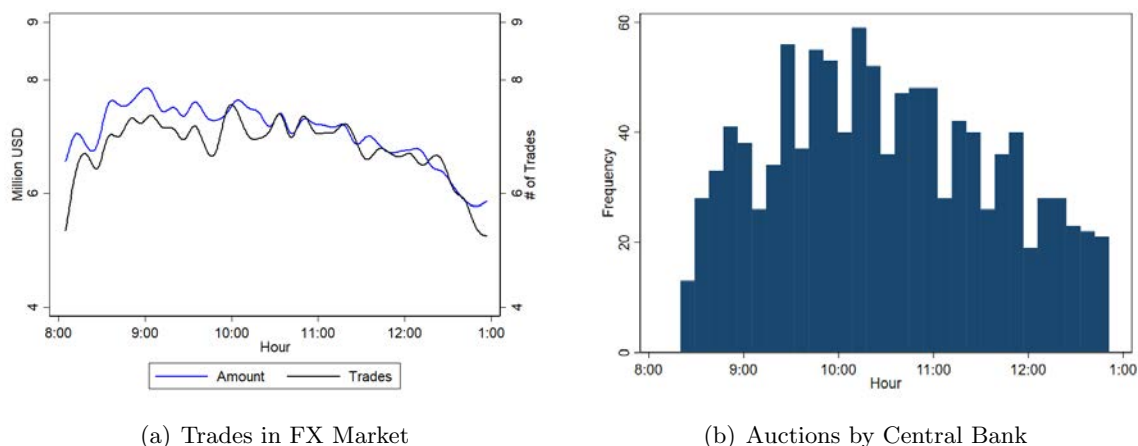
3.3 Trading Patterns

We next turn our attention to the behavior of the electronic COP-USD market. We first note that, as shown in panel (a) of Figure 4, number of trades and amounts traded are roughly uniformly

distributed throughout the typical trading day, from 8:00am - 1:00pm.¹⁵ Similarly, a histogram of the time of day when auctions are conducted by the CBoC (panel (b) of Figure 4) suggests that the time of a specific auction is hardly predictable. This helps ruling out potential trading patterns that might arise in anticipation of an auction, and that could bias the test if they were present.

We explore this potential source of bias further in a Placebo test reported in Table 5 in Appendix B. We examine the behavior of extreme bidders when trading in the centralized market around arbitrary thresholds at either 30 or 60 minutes before each auction (avoiding any overlap with the actual auction). For instance, if an auction is announced at noon on a given day, then a placebo diff-in-diff estimator within a 30 minute window would consider trades conducted between 11:00am-11:30am and compare them with trades between 11:30am-noon. We report differences in volume traded, prices, and standard deviations (of prices) for different time windows. Further, we break down statistics by purchases and sales of US dollars. We find no significant effects that would otherwise suggest asymmetric trading patterns for different types of extreme bidders before the auction.

Figure 4: Trades in centralized FX market and Auctions by Central Bank



The daily distribution of trades and amounts in the centralized market are shown using a Gaussian Kernel. Auctions issued by the Central Bank are shown using a 300-bin histogram (each bin corresponding to one minute).

4 Results

In this section we present the results of the main specifications of our test. As explained in Section 2.3, we compare the change in selling and buying prices before and after the auction, for the highest

¹⁵Table 4 in Appendix B presents summary statistics of trading behavior for 10 minute intervals within the trading day.

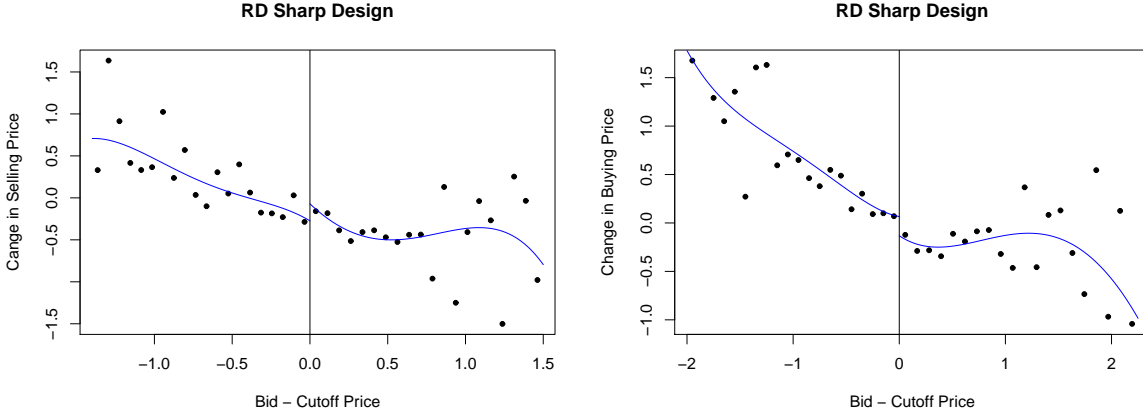
and the lowest bidders. Under PV, there should be no systematic differences between the change in prices of these two bidders, since they are already perfectly informed about their private values, and hence any information inferred from the auction about other bidders' valuations should have no effect on their own. Another way to put this is that under the null hypothesis, there is no discrepancy about the common component, any remaining heterogeneity in valuations would be explained only by differences in the idiosyncratic values. In contrast, if bidders had different expectations of the common value, this would explain some of the difference in their bids. As shown in Section 2.3, under common values, both extreme bidders would update their beliefs about the common value, although in opposite directions, after learning that their bids are extreme among the set of all bidders.

Any changes in prices induced by winning or losing the auction, for reasons other than the information it conveys about the relative valuations would potentially bias the test. Specifically, given that winning the auction implies variation in inventories, we must control for changes in prices that might result from such variation, when testing the PV hypothesis. As described in section 2.4, we address this concern using a fuzzy RDD to estimate inventory effects, understood here as decreasing marginal valuations as a function of dollar holdings.

We describe first the RD correction. A unit of observation is a pair bidder-auction. The difference between the price bid and the auction cutoff price is the forcing variable in all our RDDs. The outcome of interest is the change in a bidder's trading price, from the last transaction before the auction is announced (two minutes before it starts) to the first transaction after the auction results are disclosed to the dealers (one and a half minutes after the end of the auction). We focus on either selling or buying prices, separately. We report here the results for time windows of 60 minutes before and after the auction. The results for windows of other duration are included in Table 6 in Appendix B.

We first report the result of a sharp design (which is also the numerator of the Wald estimator in our fuzzy design). The bias-corrected local-polynomial RD estimate (Calonico et al. (2014)) is an average of 0.03 COP/USD lower selling prices for the winners, which is quite small (the average exchange rate in our sample is 1,867 COP/USD) and not significantly different from zero (with a p-value of 0.41). Figure 5, plots the results of this sharp design. For clarity, the support of the bids is partitioned in 20 disjoint bins at each side of the threshold. Every point in the graph is the sample mean of the outcome variable at the corresponding bin. Surprisingly, when we use changes in buying prices as the outcome variable the results change. The estimated treatment effect at the threshold (the auction cutoff price) is -0.31 COP/USD, and it is statistically significant at the 0.01 level. Figure 5 illustrates this result. Therefore, when focusing on purchases rather than sales, not controlling for inventory effects could result in a biased test.

Figure 5: RD Sharp Design



In our fuzzy design specification, we let the change in inventory due to the auction (the amount sold) be the endogenous treatment variable. We keep the distance between bids and the cutoff as the running variable, and the changes in selling (buying) prices as the outcome variable. When we focus on changes in selling prices at 60 minute long windows around the auction, the bias-corrected estimate is a non-significant decrease of 0.01 COP/USD for each one million dollars sold at the auction. However, if instead we use changes in buying prices as the outcome, the RD estimate is -0.07 COP/USD for each one million dollars sold, and it is statistically significant at the 0.01 level. On average, bidders sell 4.7 million dollars at an auction, conditional on winning. Thus, marginal winners (bidding the cutoff price) pay 0.33 COP less than marginal losers for every dollar they buy after reducing their inventories at the auction.

A standard assumption when modeling multi-unit auctions is that marginal valuations are either non-increasing or strictly decreasing. The estimations just reported could be interpreted as a test of this assumption. Unfortunately, the results are not conclusive since only the test based on purchases rejects the null hypothesis of constant marginal valuations (no inventory effects) in favor of decreasing marginal valuations. The test based on sales yields non-significant results. That said, our main focus here is to correct any bias from such effects, if any, in our test of private information. Since we only obtain a negative and statistically significant estimate of the slope of the marginal valuation when we use buying prices as the outcome of interest, we only consider such correction meaningful when looking at buying prices. For completeness, we still report the results of the test with and without correcting for inventory effects.

Our private vs common values test is derived directly from the model in section 2.3. As shown there,

$$(\Delta p_{L_t} - \beta \Delta x_{L_t}) - (\Delta p_{H_t} - \beta \Delta x_{H_t}) = \mu_c + \gamma_t \tag{16}$$

Table 3: Private vs Common Values: Main Test

Time window (mins)	Change in Selling Prices				Change in Buying Prices			
	30	60	90	120	30	60	90	120
$\hat{\mu}_c$	0.10	-0.03	-0.06	0.04	-0.08	-0.22	-0.19	-0.12
	(0.13)	(0.13)	(0.11)	(0.13)	(0.17)	(0.11)	(0.12)	(0.12)
$\hat{\mu}_c$ (No correction)	0.07	0.00	-0.02	0.03	0.11	0.10	0.13	0.18*
	(0.12)	(0.10)	(0.10)	(0.10)	(0.13)	(0.10)	(0.09)	(0.10)
N Obs.	180	335	427	465	161	328	404	445

Notes: This table reports the test statistic for each time window considered, and for purchases and sales by extreme bidders separately, both with and without correcting for inventory effects. Standard errors are reported in parenthesis. When correcting for inventory effects, standard errors are computed using the asymptotic distribution of the test statistic in Proposition 2. Otherwise, we report HC standard errors.

* (One-sided) p-value < 0.05.

where $\mu_c = E[\Delta c_{L_t} - \Delta c_{H_t}]$, and under the null hypothesis $\tilde{\gamma}_t$ is a mean zero heteroscedastic error term. To correct for potential inventory effects of the form $\beta \Delta x$, we use our RD point estimate, $\hat{\beta}_{RD}$. Otherwise we set $\hat{\beta}_{RD} = 0$. We then test $H_0 : \mu_c \leq 0$ v.s. $H_1 : \mu_c > 0$. When correcting for inventory effects, we use the asymptotic distribution of the test statistic in Proposition 2 for statistical inference. Otherwise, we report heteroscedastic-consistent (HC) standard errors.

We report the results of the test in Table 3. We only include transactions in the interdealer market that occurred in time windows of 30, 60, 90, or 120 minutes before and after the auction, and we compute the test statistic separately for selling and buying prices, and for each time window. For all this cases, we report the test statistic and corresponding standard error with and without correcting for inventory effects. All but one of these sixteen specifications do not reject the null hypothesis of no private information at a 0.05 significance level. When considering a 60-minute window, our test yields a price change of COP -\$0.22 to \$0.11 COP, depending on the specification considered. These magnitudes are rather small when compared to the average COP/USD exchange rate of of \$1,867 COP/USD, and are not statistically greater than zero. The results for the other time windows considered are very similar. Only when we focus on purchases by extreme bidders in time windows of 120 minutes around the auction, and we do not correct for inventory effects, we find significant differences in price changes. However, the fuzzy RD estimate of the inventory effect ($\hat{\beta}_{RD}$) using the same sample (purchases in 120 minutes windows) is a statistically significant drop in price of 0.07 COP/USD for each one million dollars sold. This indicates that for this sample the unbiased test statistic is the one that includes the corresponding correction. When included, the statistic is -0.12 COP/USD, which is obviously not statistically greater than zero.

5 Conclusion

The once opaque and fairly decentralized nature of the largest financial market in the world, the foreign exchange market, has only until recently undergone a profound renovation largely due to technological advances and the widespread use of electronic trading platforms. In fact, in 2004 many electronic brokerage systems opened up to individual costumers, providing real-time trading information at significantly lowered costs.

In this paper we empirically revisit the question of private information, but in a close-to-centralized market, where a single electronic platform covers roughly 95% of all interdealer trades for the US dollar-Colombian peso currency pair. Our main findings indicate that dealers in this market do not update their valuations after receiving information on other dealers' valuations. Specifically, we estimate that extreme bidders (participating in central bank auctions) adjust their trading prices in 0.11 and -0.22 COP/USD, for sales and purchases of US dollars, respectively, which are not statistically significant and small compared to the average exchange rate price of \$1,867 COP/USD. Our results are robust across time windows of 30, 60, 90 and 120 minutes before and after each auction.

Additionally, our results hold after correcting for inventory effects, understood as decreasing marginal valuations as functions of inventory. This correction, which we consider as a stand-alone contribution of our investigation, compares bidders within a close vicinity of the auctions' cutoff-price (barely winners and losers) with the use of a a fuzzy regression discontinuity design approach.

Overall, our results support the idea that market transparency accelerates information aggregation, hence, even if some customers' order flow provide private information to their dealers, this information is rapidly reflected in the market price. Potentially, the market might be aggregating such information efficiently.

6 Bibliography

- Athey, S. and Haile, P. (2007). Nonparametric approaches to auctions. In Heckman, J. and Leamer, E., editors, *Handbook of Econometrics*, volume 6A, chapter 60. Elsevier, 1 edition.
- Athey, S. and Haile, P. A. (2002). Identification of standard auction models. *Econometrica*, 70(6):2107–2140.
- Bajari, P. and Hortaçsu, A. (2003). The winner’s curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. *The RAND Journal of Economics*, 34(2):329–355.
- Bank of International Settlements (2016). Triennial central bank survey: Foreign exchange turnover in april 2016.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.
- Chinn, M. D. and Moore, M. J. (2011). Order flow and the monetary model of exchange rates: Evidence from a novel data set. *Journal of Money, Credit and Banking*, 43(8):1599–1624.
- Evans, M. D. (2002). Fx trading and exchange rate dynamics. *The Journal of Finance*, 57(6):2405 – 2447.
- Evans, M. D. and Lyons, R. K. (2005). Meese-rogooff redux: Micro-based exchange-rate forecasting. *The American Economic Review*, 95(2):405–414.
- Evans, M. D. and Lyons, R. K. (2008). How is macro news transmitted to exchange rates? *Journal of Financial Economics*, 88(1):26 – 50.
- Evans, M. D. D. and Lyons, R. K. (2002). Order Flow and Exchange Rate Dynamics. *Journal of Political Economy*, 110(1):170–180.
- Fama, E. F. (1991). Efficient capital markets: Ii. *The Journal of Finance*, 46(5):1575–1617.
- Grossman, S. (1976). On the efficiency of competitive stock markets where trades have diverse information. *The Journal of Finance*, 31(2):573–585.
- Hahn, J., Todd, P., and Van der Klaauw, W. (2001). Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. *Econometrica*, 69(1):201–09.
- Haile, P. A., Hong, H., and Shum, M. (2003). Nonparametric tests for common values at first-price sealed-bid auctions. Working Paper 10105, National Bureau of Economic Research.
- Hendricks, K., Pinkse, J., and Porter, R. H. (2003). Empirical implications of equilibrium bidding in first-price, symmetric, common value auctions. *The Review of Economic Studies*, 70(1):115–145.
- Hortaçsu, A. and Kastl, J. (2012). Valuing dealers’ informational advantage: A study of canadian treasury auctions. *Econometrica*, 80(6):2511–2542.
- Imbens, G. and Kalyanaraman, K. (2011). Optimal Bandwidth Choice for the Regression Discontinuity Estimator. *The Review of Economic Studies*, 79(3):933–959.
- Ito, T., Lyons, R. K., and Melvin, M. T. (1998). Is there private information in the fx market? the tokyo experiment. *The Journal of Finance*, 53(3):1111–1130.
- Kawai, K. and Nakabayashi, J. (2015). Detecting large-scale collusion in procurement auctions. Mimeo.
- King, M., Osler, C. L., and Rime, D. (2013). The market microstructure approach to foreign exchange: Looking back and looking forward. *Journal of International Money and Finance*, 38(C):95–119.
- Kong, Y. (2017). Sequential auctions with synergy and affiliation across auctions. Mimeo.

- Kuersteiner, G. M., Phillips, D. C., and Villamizar-Villegas, M. (2016). Supplementary Material for “The Effects of Foreign Exchange Intervention: Evidence from a Rule-Based Policy in Colombia”. Borradores de Economía 965, Banco de la Republica de Colombia.
- Kuersteiner, G. M., Phillips, D. C., and Villamizar-Villegas, M. (2018). Effective sterilized foreign exchange intervention? Evidence from a rule-based policy. *Journal of International Economics*, 113(C):118–138.
- Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2):281–355.
- Lyons, R. K. (1995). Tests of microstructural hypotheses in the foreign exchange market. *Journal of Financial Economics*, 39(2):321 – 351.
- Lyons, R. K. (1997). A simultaneous trade model of the foreign exchange hot potato. *Journal of International Economics*, 42(3):275 – 298.
- McAdams, D. (2007). On the failure of monotonicity in uniform-price auctions. *Journal of Economic Theory*, 137(1):729 – 732.
- McAdams, D. (2008). Partial identification and testable restrictions in multi-unit auctions. *Journal of Econometrics*, 146(1):74 – 85.
- McAfee, R. P., Takacs, W., and Vincent, D. R. (1999). Tariffing auctions. *The RAND Journal of Economics*, 30(1):158–179.
- Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2016). Information flows in foreign exchange markets: Dissecting customer currency trades. *The Journal of Finance*, 71(2):601–634.
- Pérez, J., León, C., and Marino, R. (2015). Aproximación a la estructura del mercado cambiario colombiano desde el análisis de redes. BORRADORES DE ECONOMIA 012583, Banco de la República.
- Porter, J. (2003). Estimation in the regression discontinuity model. Department of Economics, University of Wisconsin.
- Rime, D. and Schrimpf, A. (2013). The anatomy of the global fx market through the lens of the 2013 triennial survey. *BIS Quarterly Review*, pages 27–43.
- Uribe, J. D. (2016). Revista Banco de la Republica. Volumen LXXXV. Nota Editorial Agosto 2016.

Appendix A

In this section we present the proofs of Proposition 1 and Proposition 2.

As motivation for Proposition 1 we will first show that under the alternative hypothesis a bidder than learns that she has the lowest valuation updates her expectation upwards with probability one, despite the fact that all bidders' valuations also depend on idiosyncratic components, when no additional information about the valuations is disclosed.

Proposition. *Let $\mathbf{V} = (V_1, \dots, V_n)$ be the random vector of all bidders' valuations, with $V_i = E[C|S_i] + \varepsilon_i$. Under the alternative hypothesis H_1 , for all for all $\varepsilon_i > \underline{\varepsilon}$ and $S_i > \underline{s}$,*

$$E[C|S_i = s_i, \varepsilon_i = \varepsilon_i, V_i = \min\{V_1, \dots, V_n\}] > E[C|S_i = s_i, \varepsilon_i = \varepsilon_i] \quad (\text{A1})$$

Proof. Let, $\varepsilon_i > \underline{\varepsilon}$ and $s_i > \underline{s}$. First notice that $E[C|S_i = s_i, \varepsilon_i = \varepsilon_i]$ can be written as

$$\begin{aligned} & E[C|s_i, \varepsilon_i, V_i = \min\{V_1, \dots, V_n\}] P(V_i = \min\{V_1, \dots, V_n\} | s_i, \varepsilon_i) \\ & + E[C|s_i, \varepsilon_i, V_i > \min\{V_1, \dots, V_n\}] (1 - P(V_i = \min\{V_1, \dots, V_n\} | s_i, \varepsilon_i)) \end{aligned}$$

Given the common support assumption, it follows that the random variables (V_1, \dots, V_n) also have a common support. Therefore, $P(V_i = \min\{V_1, \dots, V_n\} | s_i, \varepsilon_i) < 1$. Moreover,

$$E[C|s_i, \varepsilon_i, V_i = \min\{V_1, \dots, V_n\}] > E[C|s_i, \varepsilon_i, V_i > \min\{V_1, \dots, V_n\}] \quad (\text{A2})$$

hence $E[C|s_i, \varepsilon_i] < E[C|s_i, \varepsilon_i, V_i = \min\{V_1, \dots, V_n\}]$. \square

Informally, Proposition 1 states that the former proposition still holds on expectation, when the bidders with the lowest valuation learn additional information about other bidders' valuations.

Proposition 1. *Let L be a discrete random variable taking values in the set \mathcal{B} . For any $i \in \mathcal{B}$, $L = i$ denotes that bidder i has the lowest valuation v_i among all bidders. Also, let $\mathbf{V} = (V_1, \dots, V_n)$ be the vector of all bidders' valuations, and $g(\cdot)$ be an arbitrary function of \mathbf{V} . If $E[C|L] = E[C]$, then,*

$$E\left[E[C|S_i, \varepsilon_i, L = i, g(\mathbf{V})] - E[C|S_i, \varepsilon_i] \mid L = i\right] > 0 \quad (\text{A3})$$

Proof. The Law of Iterated Expectations implies that $E[E[C|S_i, \varepsilon_i, L = i, g(\mathbf{V})] | L = i] = E[C|L = i]$. By assumption, $E[C|L = i] = E[C]$. This unconditional expectation can also be written as $E[C] = E[E[C|S_i, \varepsilon_i]]$, hence $E[E[C|S_i, \varepsilon_i, L = i, g(\mathbf{V})] | L = i] = E[E[C|S_i, \varepsilon_i]]$.

Bayes Theorem implies that

$$P(S_i \leq s | L = i) = \frac{P(L = i | S_i \leq s) P(S_i \leq s)}{P(L = i)}, \quad (\text{A4})$$

and for all s in the interior of the support of S_i , $P(L = i | S_i \leq s) > P(L = i)$. Therefore, $P(S_i \leq s | L = i) \geq P(S_i \leq s)$ for all s in the support of S_i , and $P(S_i \leq s | L = i) > P(S_i \leq s)$ for at least one such s . It follows that the distribution of S_i first-order stochastically dominates the

distribution of S_i conditional on $L = i$. Hence, since $E[C|S_i = s]$ is an increasing function of s , $E[E[C|S_i]] > E[E[C|S_i]|L = i]$.

Finally, $E[C|S_i] = E[C|S_i, \varepsilon_i]$, since C and ε_i are independent, and then $E[E[C|S_i, \varepsilon_i]] > E[E[C|S_i, \varepsilon_i]|L = i]$. The conclusion follows immediately. \square

We now derive the asymptotic distribution of the test statistic in Section 2.

Proposition 2. *Let T^r be the rate of convergence of $\hat{\beta}_{RD}$, with $0 < r < \frac{1}{2}$. Then,*

$$T^r \left(\frac{1}{T} \sum_{t=1}^T \Delta \hat{p}_{L_t} - \Delta \hat{p}_{H_t} - \mu_c \right) \xrightarrow{d} \xi E[\Delta x_{L_t} - \Delta x_{H_t}] \quad (\text{A5})$$

where ξ is a random variable with the same asymptotic distribution of $T^r(\hat{\beta}_{RD} - \beta)$.

Proof. From equation (13)

$$\Delta \hat{p}_{L_t} - \Delta \hat{p}_{H_t} - \mu_c = \tilde{\gamma}_t - \xi_T(\Delta x_{L_t} - \Delta x_{H_t}) \quad (\text{A6})$$

and, under the null hypothesis, $\tilde{\gamma}_t = \nu_{L_t} - \nu_{H_t}$.

Assumptions 2 and 3 imply that $E[\tilde{\gamma}_t] = 0$ and $E[\tilde{\gamma}_t^2]$ is finite. Hence, given an arbitrarily large random sample $\{\tilde{\gamma}_t\}_{t=1}^T$, the Central Limit Theorem implies that $\sqrt{T}\tilde{\gamma}_T \xrightarrow{d} \mathcal{N}(0, \sigma_\gamma^2)$, where $\tilde{\gamma}_T = \frac{1}{T} \sum_{t=1}^T \tilde{\gamma}_t$. Since $0 < r < \frac{1}{2}$, it follows that $T^r \tilde{\gamma}_T \xrightarrow{p} 0$.

$\xi_T = \hat{\beta}_{RD} - \beta$ and, by assumption, $T^r \xi_T \xrightarrow{d} \xi$. The Weak Law of Large Numbers implies that

$$\frac{1}{T} \sum_{t=1}^T \Delta x_{L_t} - \Delta x_{H_t} \xrightarrow{p} E[\Delta x_{L_t} - \Delta x_{H_t}] \quad (\text{A7})$$

Since convergence in probability implies convergence in distribution, the desired result follows from Slutsky's theorem. \square

Appendix B

Table 4: COP - USD Electronic Market Behavior

Hour	Average amount	Std. dev. amount	Average price Std. dev.	Average price max-min	Std. dev. price max-min
8:00	7.051416	17.65664	0.9924135	1.224043	2.415219
8:10	30.52977	18.88439	1.187962	4.993953	3.1003
8:20	33.89107	22.42347	0.8293839	3.331473	2.170445
8:30	38.09209	23.67539	0.7743075	3.045565	2.143136
8:40	41.91714	27.61976	0.7152025	2.719604	1.845875
8:50	44.15775	32.15321	0.6537201	2.51878	1.820704
9:00	38.16455	23.92161	0.5735334	2.216949	2.96128
9:10	51.13174	33.25654	0.6591426	2.517846	1.819816
9:20	43.89756	31.01556	0.5563857	2.132327	1.524664
9:30	44.43125	27.54898	0.5493024	2.106112	1.575111
9:40	46.16465	31.50749	0.5431035	2.071766	1.450401
9:50	43.62982	26.34276	0.5294372	2.045486	1.724843
10:00	39.57581	26.84598	0.443804	1.651268	0.957841
10:10	53.50715	35.6701	0.5759913	2.244664	2.005767
10:20	47.11869	32.44186	0.5138739	2.058123	1.771786
10:30	44.87132	27.06443	0.463302	1.783935	1.540969
10:40	46.05784	29.03926	0.5217314	2.026117	1.665303
10:50	47.44162	29.00567	0.5138695	1.978849	1.771306
11:00	43.92082	30.21631	0.4920517	1.929836	1.381215
11:10	50.69081	30.77058	0.5061741	1.972403	1.266016
11:20	44.66331	25.16726	0.4805873	1.946937	1.880413
11:30	45.54335	31.3349	0.4862873	1.968654	2.426538
11:40	43.06268	25.66601	0.4999638	2.059937	2.623638
11:50	44.7196	28.65876	0.4691281	1.846012	1.305219
12:00	42.587	29.09794	0.4667986	1.808533	1.323775
12:10	48.85782	28.71108	0.536661	2.159461	1.951353
12:20	44.9808	29.31306	0.5182727	1.99316	1.424898
12:30	50.23669	29.32336	0.5576244	2.423052	5.546124
12:40	44.02139	25.04799	0.4623142	1.815756	1.493834
12:50	44.17594	28.578	0.5161457	2.110002	3.191997
13:00	46.9609	25.48808	0.5253813	2.39617	2.035466

Notes: Summary statistics of trading patterns for 10 minutes intervals within the trading day. The sample is comprised of the universe of all trades executed and reported in the largest electronic COP - USD market (Approx. 90% of total market volume). All amounts are measured in million USD, and prices in COP/USD.

Table 5: Trading behavior of extreme bidders around arbitrary thresholds of 30 and 60 minutes before the auction

	Sales Diff in diff	Purchases Diff in diff
Volume	0.003	0.004
	0.004	0.004
Std Dev price	0.000	0.000
	0.001	0.000
Average Price	0.002	0.001
	0.001	0.001
Weighted Average Price	0.002	0.001
	0.001	0.001

Notes: The Diff-in-Diff estimator compares mean differences between extreme bidders, before and after each placebo threshold arbitrarily placed either 30 or 60 minutes before the auction announcement. The time windows considered do not overlap with the auction or its announcement. None of the differences reported are statistically different from zero at the 0.05 significance level. Volume is measured in million USD, Prices are in COP/USD. Weighted average prices are weighted by volume.

Table 6: Inventory effects: fuzzy regression discontinuity results

Time window (mins)	Change in Selling Prices				Change in Buying Prices			
	30	60	90	120	30	60	90	120
Inventory effect $\hat{\beta}$	0.005	-0.007	-0.009	0.003	-0.032	-0.071**	-0.064**	-0.070**
	(0.025)	(0.027)	(0.026)	(0.027)	(0.028)	(0.025)	(0.025)	(0.027)

Notes: This table reports the estimation of the inventory effect (the slope of the marginal valuation as a function inventory) for each time window considered, and for purchases and sales separately. $\hat{\beta}$ is measured in COP/USD for each 1 million USD. Robust standard errors, based on Calonico et al. (2014), are reported in parenthesis.

* $p < 0.05$, ** $p < 0.01$.

