THE LONDON SCHOOL OF ECONOMICS AND
POLITICAL SCIENCE

PUBLIC PROCUREMENT AUCTIONS IN BRAZIL

DIMITRI SZERMAN

30th July 2012

A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy, London, July 2012

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 32,374 words.

Statement of conjoint work

I confirm that Chapter 2 was jointly co-authored with Can Celiktemur and I contributed 60% of this work.

Statement of conjoint work

I confirm that Chapter 4 was jointly co-authored with Daniel Silva-Junior and Fabio Miessi Sanchez and I contributed 60% of this work.

Abstract

This thesis provides an empirical analysis of data generated by ComprasNet, the online procurement bidding platform developed and used by the Brazilian federal government. ComprasNet is a large bidding platform used since 2001 by more than 2200 public purchasing units who list around one million lots each year. Over 70,000 unique bidders have participated in these auctions. In 2010, 46 percent of all procurement for the federal government was conducted through ComprasNet, totalling R$ 27 billion, or 0.7 percent of Brazil's GDP. In short, these auctions represent a large share of federal tenders and a substantial amount is contracted through them each year. Chapter 1 provides an overview of ComprasNet. After reviewing the literature on various topics which this dissertation contributes to, I describe the institutional background surrounding ComprasNet. I then present the baseline data used throughout the remainder of this dissertation. Chapter 2 addresses one important aspect of designing an online ascending auction, namely how to end the auction. ComprasNet varied its ending rules over time, providing an unique opportunity to test theories of bidder behaviour, as well as assessing the impact of ending rules on auction outcomes. Chapter 3 analyses a two-stage auction format which ComprasNet uses. Two-stage designs have long been proposed by the theoretical literature, but there are virtually no empirical works apart from experimental studies. Finally, chapter 4 analyses a bid preference programme targeted at small and micro enterprises (SMEs). The programme consists of setting aside eligible lots for SMEs. We first use eligibility rules as a source of exogenous variation in the treatment assignment to estimate the effects of the programme on auction outcomes. We then set up an open auction model with endogenous entry and asymmetric bidders and estimate the model's primitives. In particular, we estimate entry costs, which we interpret as red tape costs.

Acknowledgements

I would like to thank Professor Martin Pesendorfer for the support throughout the PhD. I have benefited from conversations with a number of colleagues and faculty members, and I will forget some names with probability one. First, I would like to thank my co-authors, Can Celiktemur, Daniel Silva-Junior and Fabio Sanches for embarking in this project and putting up with me. Working with Michael Best in a related project indirectly contributed to the work presented here. I had extremely useful conversations with Eduardo Azevedo, John Bardear, Patrick Blanchenay, Konrad Buchardi, Pedro Carvalho, Francisco Costa, Joachim Groeger, James Hansen, Luke Miner, Francesco Nava, João Pessoa, Andrea Pratt, Rudi Rocha, Stephan Seiler, Mark Schankerman, Matt Skellern and Gabriel Ulyssea. I first heard of ComprasNet in a conversation with Claudio Ferraz. Eduardo Fiuza was my first point of contact with the data set. Joachim Krebs provided invaluable help with scrapping the data from ComprasNet. I gained important insights into public procurement in conversations with several Brazilian civil servants. It goes without saying that all remaining errors are mine.

Preface

This thesis provides an empirical analysis of data generated by ComprasNet, the online bidding platform developed and used by the Brazilian federal government for procurement of various goods and services. ComprasNet is a large bidding platform used since 2001 by more than 2200 public purchasing units who list approximately one million lots each year. Over 70,000 unique bidders have participated in these auctions to date. In 2010, 46% of all procurement for the federal government was done through ComprasNet, totalling R$ 27 billion, or 0.7% of Brazil's GDP. In short, these auctions represent a large share of federal tenders and a substantial amount is contracted through them every year.

The ultimate goal of this thesis is to improve our understanding about the functioning of online bidding platforms. To the best of my knowledge, the data analysed has not yet been explored.

Chapter 1 provides an overview of ComprasNet. After reviewing the literature on topics addressed in the subsequent chapters, I describe the institutional background surrounding ComprasNet. I then present a descriptive analysis of the baseline data used throughout the remainder of this thesis. The chapter closes with some considerations for future research based on the data presented.

Chapter 2 addresses one important aspect of designing an online ascending auction, namely how to end the auction. While traditional English auctions end when no bidder is willing to outbid the current bid, time limits are sometimes used to close the auction, particularly in online auctions. eBay, for example, has a fixed and known ending time, or a *hard close*. In contrast, ComprasNet auctions have a *random close*: auction durations are drawn from a distribution, but the realizations remain unknown to bidders until the auction closes. We first document a number of empirical regularities under random close. We find, for example, that a random close is not enough to prevent late-bidding. We then build on the work of Ockenfels and Roth (2006) to offer a stylised model in order to rationalise observed bidding behaviour. Finally, we close the chapter with some considerations about the efficiency and revenues under a random close. We conjecture that a random close may harm revenues and efficiency when entry is held constant. On the other hand, it gives weak bidders better chances of winning the auction, thus encouraging entry. This increase in participation should mitigate the post-entry negative effect on revenues.

Chapter 3 analyses offline auctions held in ComprasNet. Offline auctions were common in the early days of ComprasNet, when the online bidding software was not fully developed. These offline auctions have a two-stage design. Two-stage designs such as this have long been proposed by the theoretical literature, however there are virtually no empirical works analysing data from such designs apart from experimental studies. I first present a stylised model of bidding in a two-stage auction similar to that used by ComprasNet. I then explore bidding behaviour in offline auctions and confront the predictions of the model with the data. The model is currently unsatisfactory, as it fails to capture some key features of the data. I discuss possible modifications of the model in order to rationalise observed bidding behaviour. I then compare online and offline auctions in terms of their outcomes, taking advantage of a change in regulation that required purchasing units to use online auctions. This change introduces exogenous variation on the choice of the auction format, which can be used to identify the effects of auction rules on outcomes. I find that offline auctions attract substantially fewer bidders than online auctions and that this results in higher procurement costs. The analysis however cannot disentangle the effects of auction rules from those of the way the auction is held (online vs. offline).

In chapter 4, we analyse the effects of a bid preference programme targeting small and micro enterprises (SMEs) in ComprasNet auctions. The programme consists of *setting aside* eligible lots for favoured firms by restricting the participation by non-favoured firms. We first provide reduced-form evidence on the effects of the programme, taking advantage of the criteria used to restrict participation. These criteria are based on lots' reserve prices and provide discontinuities in the probability of treatment, allowing us to make use of them as a source of exogenous variation to identify the effects of the programme. We find that restricting participation of large firms has little effect on prices, while it increases participation of small firms. This finding is consistent with a model of bidder asymmetry and costly participation. In such a model, restricting participation by large, strong bidders increases the incentives of small, weak bidders to participate thus mitigating the adverse effects of the restriction on prices. We then set up a structural model to estimate entry costs and simulate the effects of using different criteria to set aside lots. We interpret entry costs in the context of ComprasNet as red-tape costs.

Contents

List of Figures

List of Tables

Chapter 1

An Overview of ComprasNet

Abstract

This chapter provides an overview of ComprasNet. After reviewing the literature on topics addressed in the subsequent chapters, I describe the institutional background surrounding ComprasNet. I then present a descriptive analysis of the baseline data used throughout the remainder of this thesis. The chapter closes with some considerations for future research based on the data presented.

1.1 Introduction

Online auction markets have attracted substantial attention from economists. Besides being interesting in their own right, online bidding platforms generate useful data for researchers to test economic theory and understand consumer and firm behaviour. Both the amount and type of data generated in these marketplaces are different from what economists have been used to. For example, more than 600 million items are listed every year in eBay (Bajari and Hortaçsu, 2004), a number several orders of magnitude greater than conventional datasets. Also, due to reduced costs in changing auction parameters, experimental variation abound in data generated by these marketplaces (Varian (2010)). In fact, some of early field experiments in the economics literature were conducted in online bidding platforms (Levitt and List (2009)).

Online bidding platforms for procurement have gained popularity among practitioners in the private and public sector alike. For example, General Electric claims to have saved USD 600 million in 2001 alone by using reverse auctions instead of other procurement methods. FreeMarkets.com provides a platform for companies like Quaker Oats and GlaxoSmithKline to procure billions of dollars every year (Ellison and Ellison (2005)). Yet, business-to-consumer bidding platforms–eBay in particular–have been the focus of most empirical research on online auctions. The importance of eBay notwithstanding, it is perhaps surprising then that researchers have not used data generated by online procurement platforms.

This thesis provides an empirical analysis of data generated by ComprasNet, the online procurement bidding platform developed and used by the Brazilian federal government. ComprasNet is a large bidding platform used by more than 2200 public purchasing units since 2001 who list around 1 million lots every year. Over 70,000 unique bidders have participated in these auctions. In 2010, 46 percent of all procurement for the federal government was conducted through ComprasNet, totalling R$ 27 billion, or 0.7 percent of Brazil's GDP. In short, these auctions represent a large share of federal tenders and a substantial amount is contracted through them every year. While this thesis focuses on the federal ComprasNet platform, many states, municipalities and government-owned companies run their own platforms, which are on the whole identical to ComprasNet. Thus, the actual importance of ComprasNet-style auctions is likely to be even greater than what the figures above suggest.

The ultimate goal of this thesis is to improve our understanding about the functioning of online bidding platforms. To the best of my knowledge, the data I analyse has not been explored before. In fact, I am unaware of any studies that analyse data generated by government-to-business online procurement platforms such as ComprasNet.

It is useful to take eBay as a benchmark to describe ComprasNet. Like eBay, ComprasNet generates data for a wide range of products in which many unique sellers and an even greater number of unique buyers participate. Auctioneers can also vary many auction parameters, although they typically have less discretion than eBay sellers. There are important differences, though: on both sides of the platform, stakes are higher. Rather than consumers, bidders are firms, many of which see ComprasNet as an important source of business. Moreover, the data generated is essentially administrative. Bidders are identified by their tax revenue numbers, allowing researchers to not only track bidders across auctions, but also to obtain bidders' characteristics such as location, firm size and revenues, sector of activity, etc. Sellers are public bodies, which display variation in observable characteristics such as location, governance, the incentive structure within the organisation and personnel qualification.

One key motivation for this dissertation stems from the fact that the bulk of the empirical literature on auctions has focused on a restricted set of industries and cases. In particular, works using data from OCS auctions, timber auctions, highway construction and, more recently, eBay have formed the core of this literature. To be sure, there are many works studying settings as varied as used cars, treasury bonds, fish, art, wine and school milk. However justified, such narrow focus has the drawback of limiting researchers' ability to assess their finding's external validity. This dissertation sees ComprasNet as a candidate to expand the set of empirical application in the auctions literature.

The remainder of this chapter is organized as follows. Section 1.2 reviews the literature on topics addressed in the next chapters. Section 1.3 gives a brief background of ComprasNet. Section 1.4 describes the baseline sample used throughout the remainder of the thesis. As I present the data, I discuss selected features of ComprasNet. Section 1.5 concludes by briefly discussing possibilities for future research.

1.2 Literature Review

In this section, I review the literature to which this dissertation contributes. Before focusing on the specific topics addressed in each chapter, it is worth considering the big picture. Most of the literature analysing online auctions data focuses on business-to-consumer platforms, specially eBay–by far the most popular platform amongst researches (see, for example, the review in Hasker and Sickles (2010) for a long list of studies using eBay data). Studies analysing data generated by procurement platforms (business-to-business or government-to-business) are less abundant. One exception is the growing literature on internet advertisement (Edelman and Ostrovsky (2007), Ostrovsky and Schwarz (2009)). These are "sponsored search" auctions, where bidders (advertisers) bid for keywords in search mechanisms, most notable Google and Yahoo!.

When analysing auction data, especially from internet marketplaces, researchers typically take the approach of focusing on a narrowly defined product. For example, Bajari and Hortaçsu (2003) study collectible coins, Ely and Hossain (2009) study selected DVD titles, and Lucking-Reiley (1999) studies trading cards. When items are not identical, one typically tries to collect as many observable characteristics as possible, hoping that, once those are controlled for, the products can be considered homogeneous.

Einav et al. (2011) propose a different approach to analysing data from eBay, which takes advantage of the vast heterogeneity of products and as such can be valuable for other large platforms. The authors define a *seller experiment* in eBay as a group of auctions for the same item by the same seller. The idea is that, if a seller places multiple listings for identical items while varying auction parameters–say, the reserve price or the buy-it-now option–, then the authors take this variation in auction parameters as being as good as random. Applying their definition of a seller experiment, they identify over 240,000 experiments, with considerable variation in the many auction parameters. They can then estimate average treatment effects for each of the parameters, and let those treatment effects to vary by some observables (e.g., product categories). Though one may object treating variation within a seller experiment as being as good as random, their approach does improve the trade-off between a credible identification strategy and external validity. Moreover, it illustrates how researches can use new approaches to analyse data from internet marketplaces.

1.2.1 Chapter 2: Ending Rules, Duration, and (Last-Minute) Bidding

Chapter 2 analyses ending rules used in online ComprasNet auctions. This section reviews the evidence on the relationship between auction ending rules and duration on one hand, and bidding behaviour and auction outcomes on the other.

There is robust evidence that bidders use two "strategies" in eBay: late bidding (or sniping) and incremental bidding. For example, Bajari and Hortaçsu (2003) note that 25% of winning bids in their sample of collectible coins arrive after 99.8% of the auction has elapsed. Ockenfels and Roth (2006) note that 38% of bidders in their sample submit more than one bid per auction. Although researches have used different measures to gauge the extent of sniping and incremental bidding, a number of studies have found similar patterns with larger samples of eBay auctions for multiple product categories (see Bajari and Hortaçsu (2004), Ockenfels et al. (2006) and Hasker and Sickles (2010) for a comprehensive review of these studies). At first sight, it is hard to rationalise both patterns. eBay's proxy-bidding mechanism makes it resemble a sealed-bid second-price auction. Under the standard independent private values (IPV) framework, the dominant strategy is to submit one's valuation (willingness to pay). The timing of bids should not play a role in bidders' strategies.

Besides detecting the practice of late bidding, Roth and Ockenfels (2002) and Ockenfels and Roth (2006) observe that Amazon.com auctions have less late bidding than eBay. Amazon uses a "soft close" ending rule, whereby the auction is automatically extended if there are any bids within 10 minutes of the scheduled end time. The authors interpret this as evidence that eBay's fixed ending time, or *hard close* ending rule, is a key ingredient for late-biding. One concern with these studies is the possibility of confounding factors, such as bidders' self-selection into the platforms. To address these concerns, Ariely et al. (2005) perform lab experiments in which only ending rules are varied. They find that a hard close rule leads to more late bidding than a soft close in their controlled environment, confirming previous findings. To explain the practice of sniping, and its link to ending rules, many hypotheses have been put forward and tested by the literature. I arrange them into three groups.

First, and perhaps most notably, Ockenfels and Roth (2006) argue that sniping may be a rational (best) response to bidders using an incremental bidding strategy, whereby they place multiple bids. This theory implies that, from the bidder's standpoint, sniping must dominate incremental bidding. Gray and Reiley (2007) and Ely and Hossain (2009) conduct field experiments in eBay to measure benefits to the bidder. Gray and Reiley (2007) select 70 pairs of identical listings of various product categories. For one listing of each pair, they submit an early bid while in the other, they snipe by submitting the same bid ten seconds before the auction closed. Sniping led to 2.54% lower winning bids than early bidding, but that effect is not statistically significant. Ely and Hossain (2009) bid for identical DVDs while randomising between a sniping strategy and an early-biding strategy. To measure bidder surplus gains, they fix their hypothetical valuation at different levels, and bid their valuations using the proxy system. They find that sniping increases the probability of winning the item by 12.7%, but has virtually no effect on the price they pay conditional on winning. Moreover, the effect on the winning probability decreases as their hypothetical valuation increases, since a bidder with high valuation is likely to win the item regardless of the timing of her bid.

If sniping dominates incremental bidding, then one should wonder why some bidders use the dominated strategy. Three answers emerge: inexperienced, irrational and shill bidders. Using proxies for bidders' experience, Ockenfels and Roth (2006) show that inexperienced bidders tend to use incremental bidding. This explanation gained much support from the literature. Srinivasan and Wang (2010) use entire bidding histories, and support this hypothesis. Further, they show that novice bidders learn relatively fast that sniping dominates multiple bidding. Ely and Hossain (2009) provide indirect evidence that experienced bidders are more more likely to use a sniping strategy. They find that transaction prices are significantly lower in auctions where snipping was assigned, and show that sniping encourages entry from other potential bidders (what they term a *competition effect*), but provokes those who enter into bidding less aggressively (an *escalation effect*).[1] Since sniping reduces sellers' revenues, they conclude that the escalation dominates the competition effect. Experienced bidders are more likely to understand these effects, and therefore to place their bids late in the auction. On the irrational front, Heyman et al. (2004) argue that there is a quasi-endowment effect, whereby bidders' willingness to pay increase during the auction. They find support for this hypothesis using laboratory experiments. However, more recent work using observational data with reasonable identification assumptions by Einav et al. (2011) reject the hypothesis that bidders overbid in eBay. Finally, part of the observed incremental bidding might actually be shill biding. Engelberg and Williams (2009) estimates that 1.39% of the incremental bidding observed in eBay comes sellers themselves trying to drive prices up.

Second, Roth and Ockenfels (2002) argue that due to network traffic, late bids

[1] They interpret the competition effect as a consequence of the substitutability that bidders with unit demand typically face in eBay.

in eBay have a positive probability of not going through. Therefore, if all bidders place their bids late some bids will be suppressed with a positive probability– hence lowering the winning bid and raising the surplus of the winner. Because of this effect on transaction prices, Roth and Ockenfels (2002) interpret this story as tacit collusion among bidders. Ockenfels and Roth (2006) show that bidding late can arise as an equilibrium even in an IPV framework. The key parameter in their model of eBay bidding is the probability with which a late bid is successfully transmitted. This theory implies that if late bids were guaranteed to go through, then the incentives for sniping are reduced. In their lab experiment, Ariely et al. (2005) varied the probability that late bids are successfully transmitted. Under the scenario in which bids are guaranteed to go through, the amount of late bidding actually increases, leading them to conclude that there is little support for the tacit collusion hypothesis. Another implication of the tacit collusion hypotheses is that late bidding should soften competition, leading to lower prices. The findings of Ely and Hossain (2009) therefore also provide support for this hypothesis. Other works (Bajari and Hortaçsu (2003) and Wintr (2008)) reject this hypothesis based on tests for differences in the distribution of bids in the presence or absence of late bidding. However, these studies use observational data and lack credible identification strategies. Summing up, there is inconclusive evidence for the hypotheses of tacit collusion.

A third group of hypotheses relates to the informational structure of the auction game departing from IPV. The common feature of these explanations is that some bidders are better informed than others about the value of object. Better-informed bidders choose to withhold their private information and snipe. Bajari and Hortaçsu (2003) formalise this idea and show that in a model of eBay bidding with common values bidding late is a symmetric Bayesian Nash equilibrium. In the same vein, Rasmusen (2006) presents a model in which bidders are uncertain about their private values, and would like to learn it during the auction, and would benefit from knowing other bidders' values before bidding. Thus, their opponents would be unwilling to reveal this information during the auction, leading to sniping. One problem with this type of argument is that it cannot explain late bidding on eBay at large, unless one is willing to assume that all eBay auctions have a common component strong enough to trigger such bidding behaviour.

A related, but less explored, issue is that of auction duration. In eBay, sellers can choose between 1, 3, 5, 7 and 10 days for a listing to be open for bidding. eBay charges sellers a small fee for 10-day durations, indicating that longer duration may make the listing more profitable, perhaps by making it more visible to buyers. Haruvy and Leszczyc (2010) conduct a field experiment in eBay in which they list pairs of identical items. They randomly assign the duration for 1 day to one of the items in each pair, and 3 days to the other. They find that increasing duration led to prices 11 percent higher. One issue with this finding is there may be diminishing marginal returns from extending a listing's duration. Einav et al. (2011) report that the average listing in eBay last 5.6 days. The average (population) effect of randomly extending listings' duration is probably smaller. In fact, using their "seller experiment" approach, Einav et al. (2011) report that find small effects of auction duration on transaction prices, and that those effects are statistically significant only when interacted with the buy-it-now option.

1.2.2 Chapter 3: Two-round Auctions; Online vs Offline auctions

Chapter 3 analyses an offline auction format used by ComprasNet, and compares it to the more popular online format. Offline auctions were common in the early days of ComprasNet, when the online bidding software was not fully developed. The offline format is interesting for two reasons. First, it uses a two-stage design– a feature that has received increasing attention from the literature. Second, it is commonly argued that online auctions impose less participation costs on bidders. The fact that we observe similar items being auctioned off under the two formats gives us a unique opportunity to compare the two.

There is a growing literature on two-stage auctions that combine features of different standard formats. Two things motivate this literature. First, some authors note that two-stage designs are already in use, but little is known about them. For example, Perry et al. (2000) and Dutra and Menezes (2002) observe that a number of privatization processes of state-owned companies in Italy and Brazil were auctioned off using two-stage designs, and Ye (2007) documents widespread use of two-stage auctions in the context of high-valued assets in the US electricity generation industry. Second, other authors argue that two-stage auctions can improve on standard designs in certain contexts. Klemperer (1998) proposed a two-stage design in the context of airwaves auctions, where standard auction formats faced the trade-off between efficient network formation and maximizing revenue.

Klemperer (1998) proposes the following two-stage auction game, which he names the Anglo-Dutch auction. In the first stage, the auctioneer runs an ascending auction in which the price is risen continuously until all but a predetermined number of bidders have dropped out. The remaining bidders qualify for the second round, which is a first-price sealed bid auction with reserve price equal to the first round's highest bid. The first round is therefore an English auction, and the second round is equivalent to a Dutch auction. The justification for this design lies on ex-ante bidder asymmetries and participation costs. If participation is costly, weak bidders have little incentives to participate in the auction, specially in an open (English) format. A first-price, sealed-bid design however, is inefficient in the presence of asymmetries. Thus, Klemeperer argues, combining the two formats gives enough incentives to weak bidders to participate, while eliminating inefficiencies. Levin and Ye (2008) formally analyse this model and corroborate Klemperer's insights.

Reversing the order of the stages, Perry et al. (2000) consider a Dutch-Anglo format under affiliated values. In the first stage, bidders submit binding sealed bids. The two best bidders continue to the second stage, and rejected bids are publicly announced. The second stage is a sealed-bid second-price auction augmented by a reserve price, which equals the highest first-stage bid. Perry et al. (2000) find a continuum of symmetric equilibria in first-stage bidding function. Also, they show that their auction is revenue-equivalent to an English (ascending) auction, but defend the two-stage design on the basis that it is less susceptible to collusion, while keeping an easy implementation.

Dutra and Menezes (2002) also consider a Dutch-Anglo auction, but allow the game to end in the first stage if the highest bid exceeds the second-highest bid by more than a predetermined amount. The second stage takes place if and only if there are bids close enough to the highest bid (i.e., their difference from the highest bid does not exceed that predetermined amount). In this case, all bidders who submitted bids sufficiently close to the highest bid play a second-price sealed-bid auction with reserve price equal to the first stage's highest bid. Dutra and Menezes (2002) analyse the case with three risk-neutral bidders, and bidders' valuations are discrete and have both a private and a common component. They show that in this setting, their two-stage format generates more revenue than standard auction formats.

Ye (2007) considers two-stage auctions in which first-stage bids are not binding. First-stage bids serve only as an opportunity for bidders to indicate their interest in participating before paying the entry costs.[2] The first stage is a sealed-bid auction. The n highest bidders qualify to the second stage and learn the highest rejected bid. In the second stage, bidders play a first-price, sealed bid auction. Ye (2007) notes that this indicative biding model has no symmetric increasing equilibrium in the first stage. This is a problem for efficient entry, since the mechanism does not ensure that the bidder with the highest valuation enters. She proposes an alternative qualifying rule, in which the n highest first-stage bidders must pay a fee to enter the bid in the second stage. The fee equals the highest loosing first-stage bid. Kagel et al. (2008) provides experimental evidence models proposed by Ye (2007). They find that, despite the theoretical lack of increasing equilibria, first-stage bids reflect first-stage valuations. As a result, the indicative bidding model performs well in terms of efficiency.

[2] Ye (2007) motivates this setting by observing that in the context of merges and takeovers, the auctioneer first collects non-binding bids; then, selected bidders go through the costly process of due diligence and learn their valuation.

1.2.3 Chapter 4: Set Asides

Chapter 4 analyses a preference programme targeting small and micro enterprises (SMEs). In the context of procurement auctions, preference programmes are typically of one of two forms: bid subsidies or set asides. Bid subsidies give favoured firms an advantage by scaling their bids by a factor, whereas set asides restrict participation in certain auctions to favoured firms. Two issues are particularly relevant in the context of preference programmes: bidder asymmetry and entry costs. Preference programmes are typically justified based on bidder asymmetry: the stated goal of these programmes is to give disadvantaged (or weak) bidders a better chance to win the auction. Apart from its empirical relevance, entry costs are important because they create an important channel through which the preference programme operates. In the presence of entry costs, bidders have reduced incentives to participate in an auction. The preference programme gives extra incentives for favoured bidders to participate by increasing their post-entry prospects. In this section, I summarize the empirical findings on these policies, and briefly review the literature on the estimating (static) models of endogenous participation in auctions.

Endogenous participation in auctions is generally associated with entry costs– under free entry, bidders cannot do worse than participating, and all potential bidders should participate. Entry decisions are typically modelled as a two-stage game. In the first stage, bidders decided whether or not to pay a fixed entry cost, which enables them to then submit bids in the second stage. Following the works of Levin and Smith (1994) and Samuelson (1985), researchers typically make one of the following two assumptions regarding the timing that bidders learn their values. In the model of Levin and Smith (1994), bidders incur the entry cost before learning their private values, whereas in the model of Samuelson (1985), bidders incur the entry cost after learning their private values. This difference in the timing of the game has important implications for players' strategies in the entry stage: in the Levin-Smith model, there is a symmetric mixed-strategies equilibrium (and many asymmetric pure strategies equilibria), while in Samuelson's model the entry strategy is a cut-off-type strategy in which only bidders whose valuation are above an equilibrium cut-off enter. Roberts and Sweeting (2010) note that these are two extreme assumptions: bidders either have no information, or have full information prior to paying the entry cost. Roberts and Sweeting (2010) propose an alternative model, in which bidders have imperfect private signals prior to paying the entry cost. Upon paying the entry cost, they then learn their values perfectly. It is not clear however which parameters can be identified in their model.

Which entry model to use depends on the specificities of the application at hand. For example, in the context of timber auctions Athey et al. (2011b) use the Levin-Smith model, arguing that bidders must perform a (costly) cruise in the tract to assess its potential. Krasnokutskaya and Seim (2011) interpret entry costs as bidding preparation costs in highway procurement auctions, and modify the Levin-Smith model to allow heterogeneity in entry costs, as opposed to having all bidders incurring the same fixed cost. Krasnokutskaya and Seim (2011) also "purify" the Levin-Smith model by assuming that entry costs are private information, which yields pure-strategy equilibria in the entry stage, whereby only bidders with an entry cost below a certain cut-off participate in the auction.

The fundamental methodological problem with estimating models of market entry is the existence of multiple equilibria. Works estimating such models have adopted a variety of approaches to this problem. The early literature (Bresnahan and Reiss (1990, 1991)) placed restriction in the players' payoffs and assumed away mixed strategy equilibria. The set of pure strategies equilibria gives a unique mapping to the the total number of entrants. Other works proceed by making some kind of equilibrium selection assumption. For example, Athey et al. (2011a) assume all potential strong bidders use a pure strategy entering with certainty, whereas weak bidders use a (non-degenerated) mixed strategy.

Turning to empirical results on bid preference programmes, the only two studies we are aware of that use data from set asides are those of Brannman and Froeb (2000) and Athey et al. (2011a). Both studies use data from the US Forest Service timber auctions. According to the programme, in a fraction of the auctions only small mills or loggers are allowed to participate. Brannman and Froeb (2000) use data from 51 open auctions and maintain the "button auction" assumption and the IPV framework, parametrizing the distribution of private information. Without accounting for endogenous participation, Brannman and Froeb (2000) find that eliminating the set aside programme would increase government revenues by 15 percent. Athey et al. (2011a) use data from 381 first-price sealed bid auctions which were not set aside to estimate a model of endogenous participation and ex-ante asymmetric bidders. Athey et al. (2011a) find that the set aside programme induces losses both in terms of revenue (5%) and efficiency (17%). They also find that entry choice of small firms mitigate the effects of the set aside programme: small firms did not increase their participation, revenue and efficiency losses would be much larger, at 30 and 28 percent, respectively.

Other studies analyse bid subsidy programmes. Marion (2007) and Krasnokutskaya and Seim (2011), use data from the highway construction procurement from the California Department of Transportation (Caltrans). Caltrans gives a 5% bid preference for small businesses. Marion (2007) takes advantage from the fact that there is no preferential treatment in auctions for projects using federal funds. On source of concern with this identification strategy is that state- and federal-funded project may differ along unobservable characteristics. In a reduced-form analysis, he finds that preferential treatment leads to an increase in winning bids of 3.8 percent. He also shows that, despite the fact that large firms bid more aggressively on preference auctions conditional on participation, they also participate less often in those auctions. Thus, entry choices from unfavoured firms exacerbate the costs of the programme. Using a structural model, he is able to assess efficiency losses from the programme, and finds that the marginal dollar allocated to favoured firm leads to a 20-cent loss in efficiency. Again, most of this loss comes from reduced participation from large firms. Krasnokutskaya and Seim (2011) estimate their model using data only from auctions using bid preferences, and allow for heterogeneous effects of the bid subsidy policy across different projects. In particular, they show that the programme is more effective on large projects–i.e., projects where small firms are less likely to participate. Overall, they find that the subsidy programme increases costs of procurement by 1.5 percent.

To sum up, the existing literature on bid preference programmes is increasing, but there are few studies analysing data from set asides in particular. Moreover, these studies analyse data from the same industry–auctions for timber harvest in the U.S. They also face empirical challenges which can be improved on.

1.3 Institutional Background

Public procurement in Brazil uses different methods depending on the value and nature of the goods being purchased. Reverse auction (the *pregão*) is a method used for procurement of off-the-shelf goods, regardless of their value. Broadly, an off-the-shelf good has three features. First, it can only be auxiliary to the end-activity of the public body procuring them. Second, the auctioneer must be able to fix the item specifications in a precise and concise way, so that bids can be compared solely based on the price dimension. Finally, engineering projects are not considered an off-the-shelf good. Although the legislation does not provide a clear-cut definition of an "engineering project", it is known, for example, to include entire road resurfacing works. On the other hand, reverse auction are sometimes used to procure small demolition works.[3]

These auctions are known as ComprasNet auctions, after the one-stop internet portal that hosts the bidding platform and listings. Listings are posted by purchasing units (PUs) of federal public bodies (PBs). During our sample period, the Brazilian federal administration had 278 PBs; some PBs, such as the Army, have hundreds of PUs, while others have only one PU.[4] PUs have varying degrees of

[3]Federal Law 8666/93 regulates public procurement. Federal Law 10520/2002 and Federal Decree 5450 are specific to procurement auctions. For a detailed description of public procurement in Brazil, see World Bank (2004).

[4]A public body as defined here is a *Unidade Orçamentária* (UO), and a purchasing unit is a

discretion inside the PB. The majority of listings are posted by PUs belonging to either an education institution (45%) or the military (39%).

CcomprasNet also centralizes a list of qualified suppliers. To participate in CompraNet auction, bidders must be pre-qualified. Pre-qualification requirements for these auctions are kept to a minimum, although they have changed over time and may change from auction to auction. Being up-to-date with tax obligations is the most important condition for a firm to be allowed to sign up for participation in an auction.

Figure 1.1 shows the monthly number of listings appearing in ComprasNet. ComprasNet auctions started in 2001, when PUs could choose whether to use ComprasNet auctions or other procurement methods. As of July 2005, it is mandatory for PUs to use ComprasNet auctions to procure off-the-shelf goods. Since 2007, ComprasNet auctions has accounted for 49% of all procurement done at the federal level.

Figure 1.1 Monthly Evolution of ComprasNet usage: 2001-2010

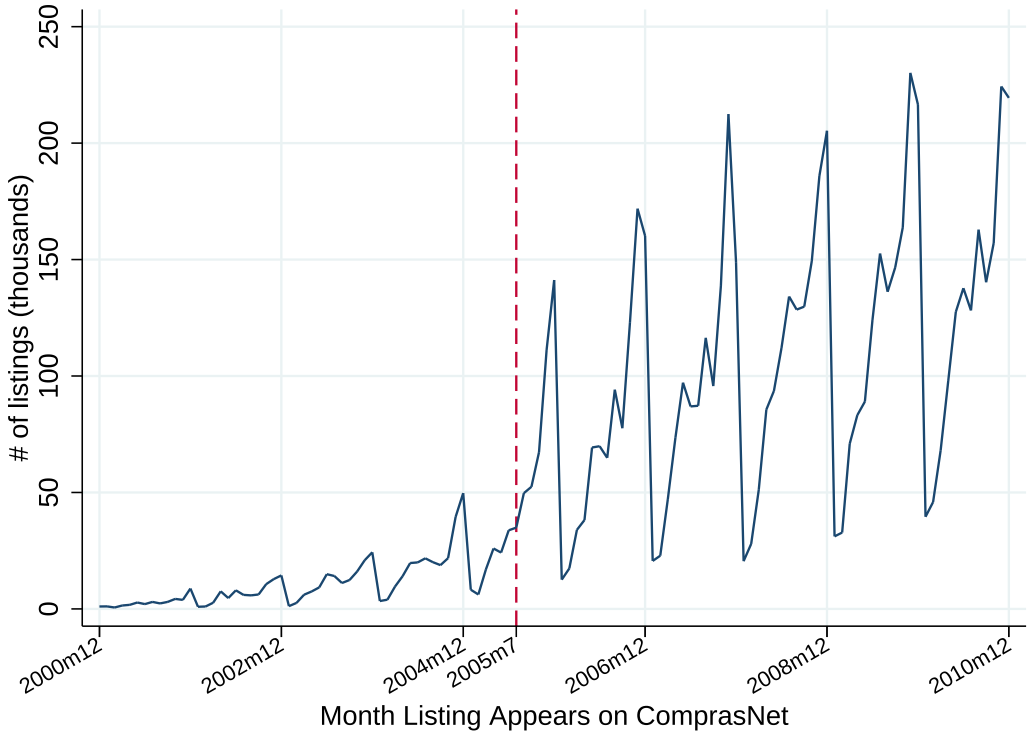

1.3.1 A Typical ComprasNet Auction: Sequence of Events

Pre-bidding A typical ComprasNet auction starts with a PU defining lots it needs to procure. A lot consists of some indivisible quantity of an off-the-shelf

Unidade Adminstrativa de Serviços Gerais (UASG).

good or service.[5] Several lots can be procured at the same letting session; such group of lots is called a *batch*. Next, the PU must provide a reservation price for each lot. The reservation price is calculated as the average of at least three quotes obtained through market research conducted by the PU, and is meant to capture the retail price of the lot. Finally, the PU advertises the tender at least 8 working days before the letting session and publishes a tender document. The tender document contains a detailed description of each lot, the date of the letting session, reservation prices and the contract's terms and conditions. It is free to download anonymously from ComprasNet.

Bidding For each lot, interested bidders must submit qualifying (opening) bids before a deadline specified in the tender document.[6] There is no minimum bid, nor a minimum bid increment.[7] When the auction starts, the low bid is announced and bidders engage in a descending auction, where bidders can place new bids.[8] A bidder's new bids must be strictly lower than her own previous bids.[9] Throughout the auction, bidders only learn the currently low bid, but neither the identity of the bidders nor the history of bids. Bidding closes at random up to 30 minutes after the auctioneer gives a warning. Chapter 2 discusses and analyses the ending rules used in ComprasNet.

Post-Bidding After bidding closes, the auctioneer checks if the best bid is below the reservation price. If it is, the best bidder is requested to submit supporting documentation. Required documents vary across lots, but are detailed in the tender announcement.[10] If the documentation is accepted, the lot is adjudicated. Otherwise, the bid is disqualified and the auctioneer may request the documentation of the second-best bidder, and proceed that way until a valid bid is found. The auctioneer may, at any point, cancel the auction. If the best bid is above the reserve price, the auctioneer tries to negotiate a better price. If the bidder is unwilling to meet the reservation price, the auctioneer has three options. First, she can declare the bid invalid and proceed to negotiate with the second-best bidder, and so on. Second, she may cancel the auction. Finally, she may adjudicate the lot at a price higher than the reservation price. This is rarely done, and when it

[5] In principle, auctioneers may allow bidders to bid for fractions of the lot. In practice, this is very rarely done. In the data, we noted 724 lots (out of more than 6 million) in which two or more bidders were awarded fractions of the lot.

[6] Such requirement is common in some open auctions–e.g., in the U.S. Forest Service auctions, see Haile (2001).

[7] To be precise, the minimum bid increment is R$0,01.

[8] Note that, unlike eBay, no proxy bidding is available.

[9] Bidders can, however, submit bids higher than other bidders' previous bids. This is to avoid a situation in which that typos (deliberate or otherwise) prevent bidders from placing new bids.

[10] Documents typically concern firms' tax duties, but may include, for example, a cost breakdown when the lot is a service, or sample items if the lot is a good.

is, the tender has a higher chance of being externally audited and the auctioneer must justify her decision–e.g., reservation prices were calculated with dated market research.

1.4 Data Description

This section describes the baseline sample used throughout the remainder of the thesis. The sample described in this section contains 6,510,541 million lots auctioned off by federal public bodies between 2004 and 2010 in ComprasNet. I use two sources of administrative data. First, I use publicly available data automatically recorded by the ComprasNet platform.

For each lot, ComprasNet collects the following information: which firms bid and all bids placed by each firm; the timing of the bids; the opening and closing dates and times of lots; the purchasing unit running the auction. Second, I complement these with internal data from the Ministry of Planning, Budget and Management. These data contain information on lots, bidders, and purchasing units. On lots, there is a paragraph-long description of the item along with classification codes following the United States' Federal Supply Codes (FSC) for materials and U.N. Central Product Classification for services. These classification schemes define product categories by 2-digit codes, and sub-categories by 4-digit codes[11]. There are also finer 6-digit codes which are created by purchasing units on a rolling basis. On bidders, the data contain the following information: whether they are registered as a small or micro enterprise (SME); their geographical location and their industry, as defined by the International Standard Industrial Classification (ISIC). Finally, the internal data contains the geographical location of purchasing units, as well as their place within the government's organisational structure.

Table 1.1 reports statistics for the 20 most frequent product categories in the sample. As the categories header suggests, various types of goods and services from different industries are procured through ComprasNet auctions. Categories range from books, to pharmaceuticals, to building materials. Moreover, items auctioned are primarily goods; only one service category (Maintenance & Installation Services) makes it into the top 20. Overall, services make up 5 percent of the number of lots (not shown in the table). Columns 1 and 2 give the total and relative frequencies of each category. The top 6 categories account for more than 50 percent of the total number of lots.

Columns 3 and 4 give the number of unique 4-digit and 6-digit codes within each product category. Some product categories are divided in up to 26 subcategories (Electrical and Eletronic Equipment Components), while other are di-

[11] The Federal Supply Codes are available at http://www.dlis.dla.mil/H2/search.aspx.

Table 1.1: 20 Most Frequent Product Categories in ComprasNet: 2004-2010

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	# of Lots	% of Total	# of 4-digit	# of 6-digit	% Canceled	% Offline	% Set Aside	% Price Regist
Books, Maps, Other Publications	771,633	12	7	185	6.5	.18	8.9	11
Medical & Veterinary Equip	758,963	12	12	42,764	11	5.9	6.7	62
Laboratory Equipment	656,605	10	19	5,747	12	3.9	7.4	40
Office Supplies and Devices	429,570	6.6	4	11,311	5.3	4.3	16	52
Subsistence (Food)	379,220	5.8	15	3,546	4.7	8	10	69
IT E&S	348,475	5.4	11	8,689	7.4	5.2	12	45
Vehicular Equipment Components	327,116	5	5	1,150	3.5	6.2	13	85
Electrical/Electronic Equip Com	238,824	3.7	26	7,995	8.9	3.1	13	47
Construction & Building Materia	177,283	2.7	8	1,504	6.3	3.3	11	75
Chemicals and Chemical Products	146,998	2.3	5	4,894	11	6.7	6.6	37
Hardware and Abrasives	126,065	1.9	16	3,588	6.8	3.9	12	58
Hand Tools	122,771	1.9	7	2,347	7.4	2.7	14	52
Pipe, Tubing, Hose, Fittings	121,256	1.9	3	3,272	5.9	4.6	12	63
Brushes, Paints & Sealers	117,813	1.8	4	1,823	6.5	4.3	12	59
Furniture	111,532	1.7	4	3,645	8.4	5.1	11	41
Maintenance & Installation Serv	101,914	1.6	16	269	4.8	6.1	29	64
Electric Wire & Power Equipment	87,745	1.3	13	3,389	8.8	3.6	14	49
Cleaning E&S	87,354	1.3	3	836	6.8	3.7	15	62
Food Preparation E&S	86,080	1.3	6	2,330	10	3.6	15	51
Nonmetallic Fabricated Materials	70,116	1.1	6	2,798	8.9	3.2	15	49
Total	6,510,541	100	719	143,202	8	4.4	11	49

Notes: Table reports statistics for all listings in the 20 most frequent categories in ComprasNet between June 2004 and December 2010. The last row shows total for all categories in the sample, not only the ones showed in the table. Categories are 2-digit codes defined by the U.S. Federal Supply Classification for goods and the U.N. Central Product Classification for services. Column (1) shows the number of lots in each category. Column (2) shows the percentage each category represents. Column (3) shows the number of subcategories in each category. Subcategories are defined by 4-digit codes of the aforementioned classification schemes. Column (4) shows the number of 6-digit codes in each category. 6-digit codes are created by procurement officers in ComprasNet. Column (5) shows the fraction of lots in each category that are cancelled. Column (6) shows the fraction of lots in each category that are held offline. Column (7) shows the fraction of lots in each category for which participation is exclusive to small businesses. Column (8) shows the fraction of lots in each category that are held under the price-registration system.

vided in only 3 subcategories (e.g., Cleaning Equipment and Supplies). Codes at the 6-digit level, which are created on-the-go by purchasing units, display even larger disparities. For example, Medical and Veterinary Equipment and Supplies, a category that includes pharmaceuticals, is divided up in more than 42,000 products at the 6-digit level. Books, Maps and Other Publications, on the other hand, are described by 185 unique products.

Cancellation Column 5 of table 1.1 shows the percentage of lots not awarded to any bidder–that is, lots that were cancelled. Cancellations may happen either before or after bidding, for several reasons. First, there might be problems in the tender document or in the online listing. For instance, the description of the lot might contain mistakes, or the online listing is not in conformance with the tender document. Second, cancellations are likely to happen if the best bid is higher than the reservation price, and the best bidder is unwilling to negotiate. Finally, bids might be rejected for non-price reasons. For instance, a bidder might have submitted a bid for an item which does not meet the required specifications, or might fail to submitted required documentation. Overall, 8% of lots were cancelled, and there is considerable variation across categories.

Offline auctions Column 6 shows the fraction of lots that are auctioned off using offline auctions. The offline format was more common prior to 2005, when the online platform was not fully developed and PUs could choose between the online and offline format. In June 2005, the same legislation that mandates the usage of ComprasNet auctions for off-the-shelf goods, also required PUs to use the online format. The differences between the formats go beyond what their names suggest, and Chapter 3 discusses and analyses these differences in depth. In total, 4% of our sample is of offline ComprasNet auctions.

Set Asides Column 7 reports the fraction of lots that are auctioned off exclusively to small and micro enterprises (SMEs).[12] These lots are part of a preferential programme in federal procurement introduced in 2007. Lots whose reserve prices are below R$80,000 are eligible to be included in the programme, at discretion of the purchasing unit. Table 1.1 shows that 11% of all lots are included in programme, and that this fraction varies across products categories. The set-aside programme is the subject of Chapter 4.

Price registration Column 8 reports the fraction of lots that held under the *price registration system* (PRS). The PRS was created to solve two problems. First, there are cases in which PUs cannot know ex-ante either quantities to be demanded, or

[12]To qualify as a micro (small) enterprise, a firm must not have gross revenues larger than R$ 360,000 (3,600,000). Prior to 2009, this threshold was R$ 240,000 (2,400,000).

delivery dates, or both. In this case, the quantity announced in the tender is an upper bound of how much of that item the PU will demand in the next, say, 12 months. Second, the PRS helps to circumvent the requirement that public bodies have available funds for any purchases at the time of the tender announcement. Due to uncertainty in the public bodies' budgets, auctioneers may find it beneficial to hold listings under the PRS which does not require funds to be available. To sum up, the PRS gives auctioneers the flexibility of acquiring quantities up to the amount originally announced in the tender, without creating the obligation to acquire any of that quantity. Around 49% of the listings are held under the PRS. This is an interesting parameter for it may introduce uncertainty to bidders, who do not know how much, if any, of the quantity announced will be purchased. Moreover, some bidders may be better informed than others if they have experience in supplying a particular public body.

To highlight some key features of these auctions, Table 1.2 presents summary statistics of selected variables.

Reservation Prices and Winning Bids The third row of table 1.2 reveals that, conditional on being accepted, winning bids are on average 69 percent of the lot's reservation prices. There is also sizeable variation on winning bids: in one quarter of lots, the winning bid is below 51 percent, and in another quarter, they are above 93 percent of the reserve price. As noted in section 1.3, reservation prices are calculated based on market surveys and are meant to capture the retail price of the lot. In fact, the reservation price is commonly called "estimated" or "reference" price in the ComprasNet jargon. Such feature has a few interesting implications. First, reservation prices should capture many unobservable (to the researcher) characteristics of the lot. Second, identical lots display a lot more variation in reservation prices when compared to other settings.[13] Second, reservations prices of large lots are likely to ignore possible quantity discounts. This helps to explain the magnitude of discrepancy between reservation prices and winning bids.

Tender Notice, Auction Duration and Bids Placed Table 1.2 reveals that tenders are announced on average 14.4 days prior to the letting day. During this period, bidders may ask for clarifications about the lot's specifications. Bidding itself takes place over the course of 65 minutes on average. In this respect, ComprasNet is very different from eBay, where listings are open for bidding for periods from 3 up to 10 days. ComprasNet does not offer a proxy bidding system, which helps explaining the number of bids per bidder placed in an auction.

[13]For example, in the US Forest Service, auctioneers set reservation prices at one of three pre-established levels.

Table 1.2: Sample Descriptive Statistics: All Listings 2004–2010

	Mean	Std. Dev.	25th pctile	75th pctile
Log Reserve Price[a,c]	6.55	2.30	4.91	8.01
Log Winbid[a,b]	6.23	2.31	4.60	7.69
Winbid/Reserve	0.69	0.25	0.51	0.93
$\frac{Ranked2 - Ranked1}{Ranked1}$	0.10	0.25	0.00	0.08
# of Bidders	5.84	5.11	2.00	8.00
# of Small Bidders	4.85	4.67	2.00	7.00
# of Bidders in Same State	2.71	2.86	1.00	4.00
Tender Notice (days)[d]	14.46	7.53	12.00	14.00
Auction Duration (minutes)[e]	65.64	55.66	34.53	75.28
# of bids per bidder	4.16	6.55	1.25	4.22
Listing Canceled	0.08	0.27	0.00	0.00
Geographic Region				
North	0.12	0.33	0.00	0.00
Northeast	0.21	0.40	0.00	0.00
Southeast	0.34	0.47	0.00	1.00
South	0.17	0.37	0.00	0.00
Central-West	0.16	0.37	0.00	0.00
# of Lots	6,510,541			

Notes: Table shows summary statistics for all listings posted by federal purchasing units between June 2004 and December 2010. (a) Monetary values are measured in nominal R$. (b) No winning bid is observed in 875,287 observations due to cancelation or no bids being submitted. (c) No reserve price is observed in 71,475 the observations. (d) Tender notice is the number of days between tender announcement in ComprasNet and bidding day. (e) Only for Online auctions; no duration is observed in 814,880 observations due to no bids being submitted, or the way data is sometimes recorded.

Bundles Participation in one lot is not tied to participation in other lots, except in cases where lots are *bundled*. Bundled lots are awarded to the bidder making the best combined bid in all lots of the bundle. Bundles were introduced for two reasons. First, some items may require some kind of uniformity among them (different pieces of furniture to refurbish an office space, for example). Second, there are gains of scale in contract monitoring for similar, or complementary items. For example, air conditioners may be bundled with their respective installation services. In total, 662,760 items are bundled. In the next chapters, bundled lots are excluded from the data.

Batches As noted in section 1.3, lots are typically auctioned off in *batches*. Although batches are common in many settings– think, for instance, of an evening of art auctions at Sothebys'–I emphasize this feature here for two reasons. First, lots belonging to the same batch share certain characteristics. In other words, certain variables in my dataset have variation only across batches. For example, the tender notice is the same for all lots of a batch. Second, lots of a batch are typically similar, and bidders usually participate in multiple lots in a batch. There may complementarities in lots of the same batch–for example, economies of scale in transportation costs or fixed costs of participation at the batch level. To give a sense of how much is auctioned off in a batch, Table 1.3 gives statistics of the batches' total reserve prices.

Table 1.3: Total Batch Reserve Prices

		Total Batch Reserve Prices (R$)			
# of Lots	# of Batches	Mean	10th pctile	Median	90th pctile
1	65,572	396,080	7,700	48,585	512,000
2-9	60,621	524,982	8,019	44,313	524,500
10-49	47,063	726,763	10,525	65,537	765,046
50-100	15,110	1,198,542	15,077	103,106	1,542,010
100+	16,520	2,846,299	29,957	257,170	4,083,161
Total	204,886	774,113	9,346	60,000	810,948

Notes: the first column shows the number of batches in the sample with lots given by the row headers. The remaining columns give statistics of the total reserve price of batches with lots given by the row headers. Reserve price are measured in nominal R$.

Bidders Over 70,500 unique firms have participated in these auctions. Table 1.2 sheds some light on the nature of these firms. In a typical auction, 5 out of 6 firms are registered as a small or micro enterprise with the tax authorities.[14]

[14]The set aside programme can partially explain this figure, as 11% of all auctions were restricted to small or micro enterprises.

One feature of these data is the remarkable variation in the number of bidders across auctions. I explore the determinants of entry in an auction with a Poisson regression of the number of bidders on various covariates.

The results are presented in Table 1.4. I find that the number of bidders has increased over the years, but displays seasonal variation within a year. For example, the coefficients on the yearly dummies imply that the typical auction in 2009 had 27.7 percent more bidders than a typical auction in 2004. But auctions in the last quarter have on average 34.6 percent *fewer* bidders than auctions in the first quarter. I offer the following explanation for these patterns. Looking back at Figure 1.1, we can see that the number of auctions has increased over time, and displays wide seasonal variation. The fact that the number of bidder has increased in the long run is probably related to the increase in the supply-side of ComprasNet: by relying more on ComprasNet, the government attracted more bidders to the platform. However, the large spike in the number of auctions at the end of each years creates a congestion effect in the short run, whereby PUs compete for a relatively inelastic pool of bidders.

1.5 Concluding Remarks

This chapter provides an overview of ComprasNet. ComprasNet is a large online marketplace which accounts for a sizeable share of public procurement in Brazil. The next chapters analyse some interesting features of ComprasNet. There are a number of other interesting topics to be addressed in future research with these data, two of which I outline here.

First, these auctions–and federal procurement in general–display tremendous seasonality; 25% of the procurement activity is conducted in the four weeks preceding the end of the fiscal year. After talking to procurement officials, the following story emerged. Public bodies have lapsing budgets, that is, any left-overs from one fiscal year are lost and cannot be spent on the subsequent year. In addition, there is uncertainty both on the supply of funds and on the demand side. As a result, procurement managers prioritize spending throughout the year, and rush to spend any unused resources before the budget expires. Liebman and Mahoney (2010) document the same phenomenon in US public procurement. In fact, many organisations, both private and public, operate their budget under similar rules.

Such seasonality in procurement activity is negatively correlated with the average number of bidders, as we saw in table 1.4. This suggests that bidders have participation constraints. There are alternative explanations, such as seasonality in other business in which bidders also participate. Nevertheless, it is worthwhile exploring further price patterns, to gauge the extent to which budget rules have

Table 1.4: Determinants of Entry: Number of Bidders

	OLS		Poisson	
Variable	Estimate	Std Error	Estimate	Std Error
Log Reserve Price	0.552	(0.001)	0.092	(0.000)
PRS	0.908	(0.004)	0.149	(0.001)
Offline	-2.661	(0.008)	-0.566	(0.002)
Log # Items in Batch	-0.245	(0.002)	-0.041	(0.000)
Tender Notice	0.009	(0.000)	0.001	(0.000)
Quarter of Year (1st omitted)				
2 Quarter	-0.714	(0.009)	-0.105	(0.001)
3 Quarter	-0.913	(0.008)	-0.135	(0.001)
4 Quarter	-2.101	(0.008)	-0.347	(0.001)
Year (2004 omitted)				
2005	0.377	(0.016)	0.070	(0.004)
2006	0.619	(0.016)	0.124	(0.003)
2007	0.995	(0.016)	0.200	(0.003)
2008	1.105	(0.016)	0.219	(0.003)
2009	1.464	(0.016)	0.276	(0.003)
2010	1.396	(0.016)	0.264	(0.003)
Geographic Region (North omitted)				
Northeast	0.627	(0.006)	0.109	(0.001)
Southeast	0.475	(0.006)	0.080	(0.001)
South	0.024	(0.006)	0.005	(0.001)
Central	0.766	(0.006)	0.121	(0.001)
4-digit code dummy	Yes		Yes	
N	6,439,066		6,439,066	
R^2	0.297		0.171	

Notes: Dependent variable is the number of bidders in the auction. Robust standard errors in parenthesis. 71,475 observations are dropped from the sample due to missing reservation prices.

knock-on effects on auction outcomes. Transaction-level data can help assessing whether these rules lead to wasteful spending in the sense of Bandiera et al. (2009).

Second, corruption and collusion are common phenomena in public procurement, both in developed and developing economies. For example, Ferraz and Finan (2011) list some observed practices of corruption in public procurement in Brazil. It is therefore intriguing that after more than 10 years of ComprasNet operation, reports of corruption and/or collusion in these auctions have been, at best, anecdotal. This fact can be interpreted in two opposite, but non-excludable ways. On one hand, it could mean that these practices are, in fact, rare on ComprasNet. A World Bank report, for example, found that the procedures in ComprasNet, although not foolproof, are amongst best procurement practices. Tran (2008) finds that replacing traditional procurement methods with online reverse auctions, reduced bribing in one developing country's public tenders. On the other hand, it may be that the detection, but not the practice, of corruption and collusion in ComprasNet is low. The truth probably lies in between the two extremes; given the absence of rigorous studies using these data, a thorough analysis is called for.

Chapter 2

Ending Rules in Online Auctions

Abstract

One important aspect of designing an online ascending auction is how to end the auction. While traditional English auctions end when no bidder is willing to outbid the previous bid, time limits are sometimes used to terminate an auction. eBay, for example, has a fixed and known ending time, or a *hard close*. In contrast, ComprasNet auctions have a *random close*: auction durations are drawn from a distribution, but the realisations remain unknown to bidders until the auction closes. We first document a number of empirical regularities under random close in ComprasNet. We find that (i) bidders defer bidding to the end phase of the auction; (ii) a sizeable fraction of auctions is resolved early in the auction; (iii) bid increments are typically small; (iv) large bid increments are more likely to occur early in the auction. We then build on the work of Ockenfels and Roth (2006) to offer a stylised model in order to rationalise observed bidding behaviour. Finally, we close the chapter with some considerations about the efficiency and revenues under a random close. We conjecture that a random close may harm revenues and efficiency when entry is held constant. On the other hand, it gives weak bidders better chances to win the auction, thus encouraging entry. This increase in participation should mitigate the post-entry negative effect on revenues.

2.1 Introduction

One important aspect of designing an ascending auction is how to end the auction. While traditional English auctions end when no bidder is willing to outbid the previous bid, time limits are frequently used to terminate an auction. Two prominent examples are eBay, which uses a fixed and known ending time, known as a *hard close*, and Amazon, where auctions are automatically extended if a bid is placed close to the current scheduled ending time–a *soft close*. Such differences can have significant effects on bidding behaviour: for example, bidders on eBay have an incentive to place last-minute bids, a strategy known as sniping (Ockenfels and Roth (2006)). The effects of ending rules on auction outcomes (revenues and participation) are less well understood.

ComprasNet is, for a number of reasons, an interesting case study of the effects of ending rules on bidding behaviour and auction outcomes. First, ComprasNet has varied its ending rules over time, a feature rarely observed in online platforms.[1] Evidence based on cross-platforms comparisons suffer from a series of confounding factors, such as buyer and seller self-selection, and platform reputation. Analysing data from a single platform eliminates many of these concerns. In addition, the data generated by ComprasNet allows us to track bidders over time and observe them in different regimes. We can therefore control for bidder idiosyncrasies while at the same time assessing the influence of experience in bidding patterns.

Second, ComprasNet uses different ending rules from those observed in eBay and Amazon, providing us with evidence on new auction designs. In one design, auctions close at random – bidders do not know and cannot influence the ending times. We call this ending rule *random close*. In another design, the auction duration is extended for each submitted bid, but the maximum duration is capped. Bidders are not aware of the fact that new bids extend the duration, although they certainly can infer that the auction duration is skewed to the left. We show that in practice this ending rule amounts to a *quasi-hard close* because, despite some random variation in the ending time, there is considerable mass at one short interval of the support of auction duration.

The fact that auctions with a random close are rarely observed outside of ComprasNet might be partially due to IBM's US patent on this kind of auction (IBM calls them "smooth-finish" auctions).[2] As Klemperer (2004) points out, the fact that IBM was granted this patent in 2003 is perhaps surprising, as auctions incorporating a random ending time have a long history. Cassady (1967) describes variants of a *candle auction*, which was commonly used in Great Britain in the

[1] One exception is Yahoo, which used to offer sellers the option to list items under a hard close (eBay-style auctions) and a soft close (Amazon-style).

[2] Available at http://patft.uspto.gov/.

seventeenth century. In one variant, the auctioneer inserts a pin close to the top of a candle, and bids can only be placed as long as the pin is held to the candle. The other variation occurred in some auctions held outdoors. Bids could only be placed so long as the flame was alive. In essence, both variants are random close auctions, since bidders cannot tell precisely when the auction is going to end.

Both Cassady's account and IBM's patent suggest that auctions with a random ending time may have attractive features for auctioneers. More generally, Cassady's writings also demonstrate that the use of time limits for closing ascending auctions long predates, and is not restricted to, online auctions.[3] It is not surprising then that eBay's hard close ending rule and its implications for bidding behaviour, has generated a large literature on the topic (see for instance the references in Hasker and Sickles (2010)).

The objectives of this chapter are twofold. First, we document stylised facts of bidding behaviour under the different ending rules used by ComprasNet. We find that, under both ending rules (i) bidders wait until the end phase of the auction to bid; (ii) a sizeable fraction of auctions is resolved early; (iii) bid increments are typically small; (iv) large bid increments are more likely to occur early in the auction. Among other implications, these facts imply that a random close does not necessarily rule out late bidding. Second, we build on the work Ockenfels and Roth (2006) to develop a simple model that captures the key elements of bidding behaviour in auctions with a random ending time. It is important to note that ascending auctions are complex dynamic games, and generally cannot be fully analysed with available methods. Such complexity has led researchers to rely on abstractions for tractability, modelling ascending auctions as a "button auctions", as in Milgrom and Weber (1982). Time limits bring additional difficulties in modelling ascending auctions by augmenting bidders' strategy space. As a result, any attempt to model ascending auctions with time limits is bound to be a drastic simplification of reality, aimed at capturing only the most salient features of the auction game.

English auctions have many desirable features (Cramton (1998)), and so it is not clear what the arguments are for modifying its ending rules and include time limits. To the best of our knowledge, the literature has not offered any answers to this question from a theoretical perspective. From a practical point of view, bidding can be time-consuming and costly for auction participants. Time limits can therefore be seen as a way of limiting the escalation of these costs.[4] In the case

[3] Füllbrunn and Sadrieh (2012) describe various instances of items sold through random close (candle) auctions. Cassady (1967) describes an auction for toll rights conducted in 1932 in which a time limit was set by an hour glass. Cassady points out that no bids were placed until the sand had nearly run out; it seems that bidders on eBay did not invent the "sniping" strategy after all.

[4] Cassady (1967) puts forward an alternative rationale for time limits. Bidders may need time to assess the item's value; an English auction without time limits forces bidders to keep on bidding. Of course, this argument relies on the fact that bidders are not well informed of the item's value.

of ComprasNet, such practical considerations were probably in play, but there were two other reasons for the choice for a random close. First, it was meant to preclude procurement officials giving favoured bidders an unfair advantage by passing on information about the ending time of the auction. Second, it was meant to curb the practice of sniping, widely observed in hard close auctions in eBays.[5] Procurement officials seem to believe that sniping is detrimental to auction revenues, a view that finds some support in the empirical literature, see Ely and Hossain (2009).

In our model, we show that, in auctions with a random ending time, bidders with valuations above a certain threshold may have an incentive to bid more aggressively early in the game, but also that, as the game enters the ending phase, all bidders use an incremental bidding strategy. We conjecture, but do not show formally, that this behaviour will lead low-valuation bidders who are still willing to bid in the final stage to win the auction with higher probability than under a traditional English auction. This, in turn, has a direct negative effect on revenues and efficiency, but also has a concomitant positive effect on entry; this additional entry positively affects revenues, thus counteracting the negative direct effect.

This chapter contributes to the literature on ascending auctions with time limits. In particular, this chapter is close in spirit to the works of Roth and Ockenfels (2002) and Ockenfels and Roth (2006). We are unaware of any existing work analysing, either theoretically or empirically, the auction institutions studied in this paper. The only exception is Füllbrunn and Sadrieh (2012), who attempt to develop a model of these auctions and perform lab experiments. The authors attempt to model a random close auction, but they model it as a second-price auction. As a result, they find an equilibrium in which bidders bid their true valuation. The experimental evidence they provide however rejects the prediction of their model.

The remainder of this chapter is organized as follows. Section 2.2 describes in detail the bidding timeline in ComprasNet, the ending rules used and the changes that have occurred to those rules. In Section 2.3, we present the data and describe some key facts of bidding under the different ending rules. Section 2.4 presents a model of auctions with a random ending time. Section 2.5 concludes.

2.2 Environment and Data

This section describes the timing of an online ComprasNet auction, as well as the natural experiments that have occurred.

It also ignores other possible arrangements that give bidders enough time to revise their bids without imposing time pressure on them. For example, in the UK 3G spectrum auction bidding went for as many rounds as needed, and the number of rounds in any given day was limited.

[5]This argument is also made by Trevathan and Read (2011) and Malaga et al. (2010).

Pre-bidding Listings appear on ComprasNet at least 8 working days before the letting session. The tender document contains a detailed description of each lot, the date of the letting session, reservation prices and the contract's terms and conditions. It is free to download anonymously from ComprasNet.

Bidding Interested bidders must submit a sealed bid before a pre-specified deadline, after which no bidder may enter the auction. Figure 2.1 depicts the bidding timeline after this point. At t_0 sealed bids are open, and bidders learn the low bid. Bidders can now place as many new bids as they want. A bidder can only place bids strictly lower than her own previous bids.[6] t_1 is unknown to bidders. At t_1 the auctioneer announces t_2, the start of the ending phase. At t_2, a number d_C is calculated according to an algorithm described below. The realisation of d_C remains unknown to bidders (and the auctioneer). At $t_3 = t_2 + d_C$ the auction ends. The low bidder at t_3 wins and pays her bid.

Figure 2.1 Bidding Timeline

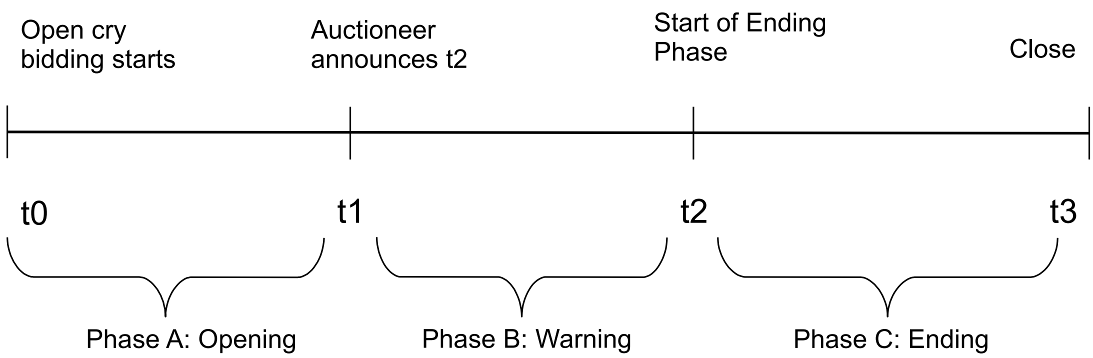

Ending Rules An ending rule is an algorithm that determines the duration of Phase C. Two ending rules have been in used in ComprasNet. Under a *random close* ending rule, d_C is drawn from the uniform distribution in the [0,30] minutes interval. In a *quasi-hard close* ending rule, d_C is the sum of a random draw from the uniform [5,30] and one random draw from the uniform [0,2] for each bid placed in the auction, but it remains capped at 30 minutes. The precise *quasi-hard close* algorithm is given in the appendix. In practice, a *quasi-hard close* puts a lot of mass on the interval between 28 and 29 minutes, as we will see in the next section.

[6]Bidders can, however, submit bids higher than other bidders' previous bids. This is to avoid a situation in which that typos (deliberate or otherwise) prevent bidders from placing new bids. The platform software uses an algorithm to spot this sort of typos, but might take a few seconds to exclude mistaken bids, or might simply fail to spot typos.

Figure 2.2 Changes in Ending Rules for Online Auctions

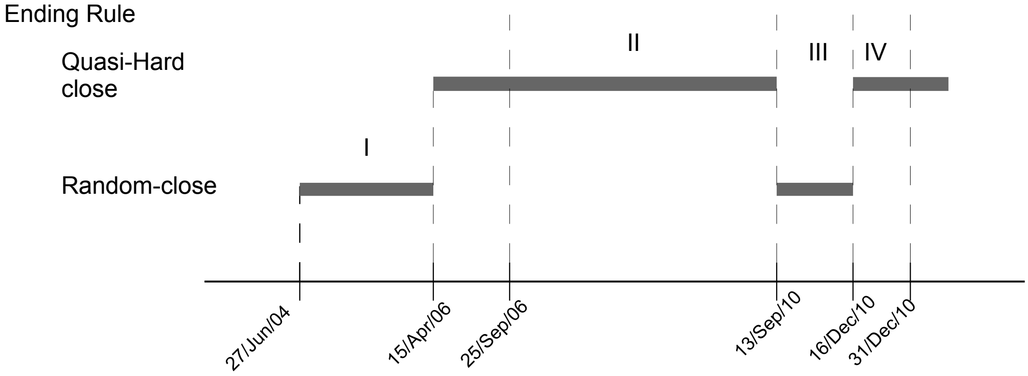

Notes: All changes between 15/Apr/2006 and 16/Dec/2010 were unannounced both to bidders and auctioneers. The period between 15/Apr/2006 and 25/Sep/2006 was similar to the quasi-hard close algorithm described in the paper, but had minor differences. We therefore do not use data from that period.

Changes to Ending Rules There have been changes in the ending rule over time, as illustrated by Figure 2.2. Between June 2004 and April 2006 (Period I), the random close was used.[7] In April 2006, ComprasNet changed the ending algorithm to the quasi-hard close. After speaking to procurement officials, the following story emerged. The ending rule algorithm was changed after repeated complaints from (loosing) bidders that they did not have enough time to place their bids. ComprasNet developers then changed the ending rule to the quasi-hard close, effectively allowing for more bidding time in Phase C. It seems that during the period between April 2006 and September 2006, different algorithms were used, but we were not able to recover the precise rules and therefore did not use data from this period.[8] Between September 2006 and September 2010 (Period II), the quasi-hard close was unchanged. It is unclear why ComprasNet switched back to the random close between September and 13 December 2010 (Period III), and back again to quasi-hard close after 14 December 2010 (Period IV).[9] It is important to note that bidders did not have knowledge of the exact functioning of the quasi-hard algorithm. In particular, bidders were unaware of the fact that the software extended the duration of Phase C with the number of

[7]It is unclear how to classify the the way ComprasNet ended auctions prior to June 2004. It seems that auctioneers had control over the ending time, but would not reveal it precisely to bidders. Furthermore, duration data was not embedded in the HTML files containing the records of each auction, so we cannot collect this information.

[8]However, we noted that all the algorithms used between April 2006 and September 2006 had in common a minimum of 5 minutes for bidding as well as automatic extensions when bids were placed. The duration of Phase C was kept capped at 30 minutes for most of this period, although we observed a few cases in which Phase C lasted for up to 35 minutes.

[9]In fact, procurement officials do not seem to be aware of those changes at all. In March 2011, an external auditing body ruled that the ending rule was to be reversed to a random close. Our data however end in December 2010.

bids; they simply observe changes in the duration of Phase C.[10]

Figure 2.3 Distribution of Phase C Duration

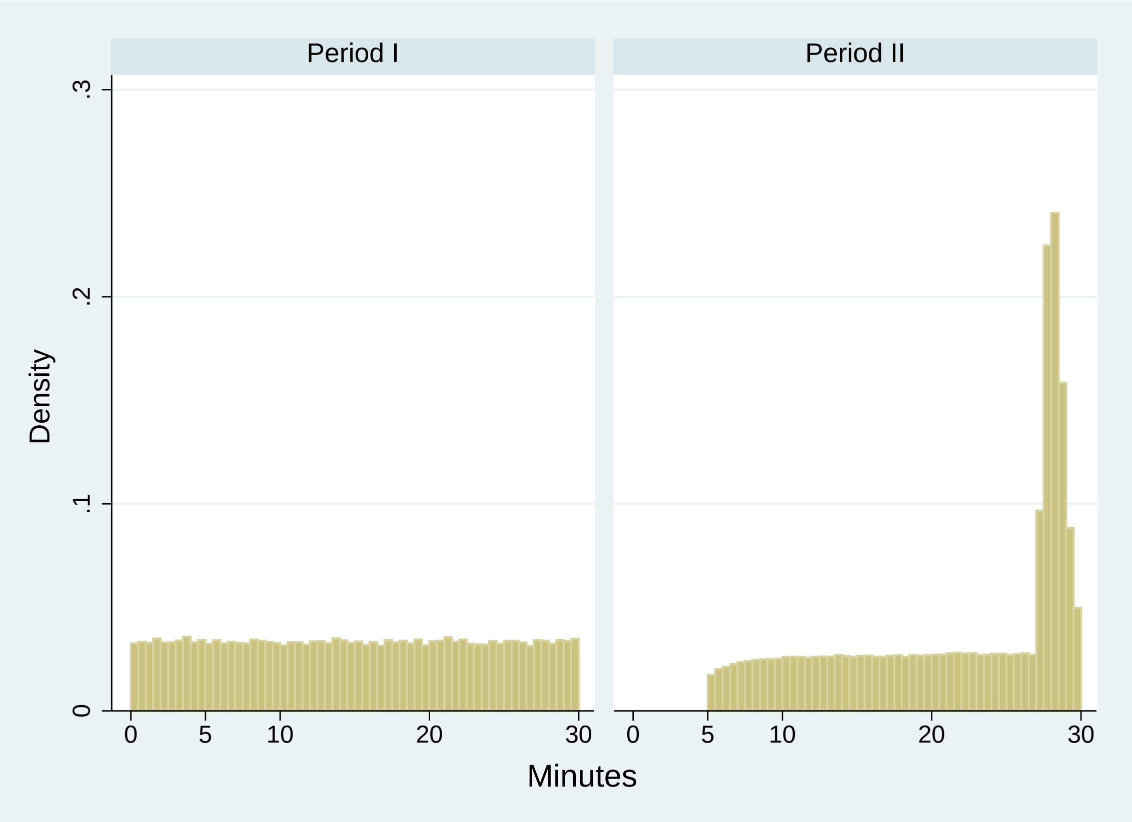

Notes: The Figure shows the distribution of Phase C in Periods I (random close) and II (quasi-hard close). Bins represent 30-second intervals.

2.2.1 Data

For the purposes of this chapter we restrict the sample to auctions held in Periods I and II.[11] We also require that auctions in our sample have at least one bidder. We select data for eight commonly purchased products: books, A4 paper, ink cartridges, printers, air conditioners, coffee powder, computers and bottled cooking gas. Table 2.1 presents descriptive statistics for each period in the sample. The table shows detailed information on the duration of each auction phase. As expected, the average duration of Phase C in Period I is 15 minutes. More details on the distribution of of Phase C duration are given in Figure 2.3.

Table 2.1 also reports statistics for the number of bids per bidder placed. In a typical auction in Period I, a bidder places on average 1.95 bids. Of these, 1.36

[10]In fact, bidders (and auctioneers) were not officially warned of any of these changes. Any changes in bidders' behaviour due to change in the ending rules therefore must have come through learning.

[11]Analysing periods III and IV would certainly be interesting, but adds little value given the current purposes of this chapter.

bids (or 70 percent) are placed in Phase C. Note however, that in a typical auction, the duration of Phase C accounts for 19 percent of the total auction duration. This is the first sign that bidders delay bidding to Phase C, a point we will return to in the next section. This pattern is even stronger in Period II, where Phase C accounts for 38 percent of the auction duration and for 85 percent of the bids.

Table 2.1: Summary Statistics by Period

	Period 1		Period 2	
	Mean	Std. Dev.	Mean	Std. Dev.
Log Win Bid	5.21	1.74	5.46	1.88
Log Reserve	5.60	1.80	5.87	1.94
Win Bid/Reserve	0.71	0.19	0.70	0.20
# Bidders	6.61	6.10	7.76	7.58
Duration of Auction Phases (minutes)				
Phase A (d_A)	55.75	68.53	28.74	45.43
Phase B (d_B)	8.79	10.91	6.32	8.84
Phase C (d_C)	15.02	8.67	21.73	7.49
Total Auction Dur.	79.14	71.59	56.77	47.94
# of Bids Per Bidder				
Phase C	1.36	3.48	2.37	5.59
All Phases	1.95	4.15	2.77	5.94
Product Subcategory				
Air Conditioners	0.02		0.02	
Book	0.73		0.74	
Coffee	0.00		0.01	
Computer	0.02		0.02	
GLP	0.00		0.01	
Ink Cartridges	0.20		0.17	
Paper	0.02		0.01	
Printer	0.01		0.02	
Observations	65,593		336,624	

Notes: Period I has a random close and Period II a quasi-hard close. The definitions of periods are given in Figure 2.2. Duration of Phase A is missing in 43,323 observations in Period I and 7,095 observations in Period II due to data recording issues. In those cases, overall Auction Duration is also missing.

2.3 Stylised Facts

In this section we document some key features of bidding behaviour in random close auctions.

2.3.1 Bid Increments and Jump Bidding

Since we observe all bids placed, we can gauge the extent of incremental and jump bidding. In each auction, we calculate bid increments. Specifically, we define the k-th increment in auction t as $i_{k,t}$ as follows:

$$i_{k,t} = \frac{j_{k,t} - j_{k-1,t}}{r_t}, \quad k = 1, ..., K, \ t = 1, ...T, \tag{2.1}$$

where $j_{k,t}$ is the k-th bid placed in auction t and r_t is the reservation price of auction t. When $k = 1$, we take the best sealed bid as $k = 0$. Since these are procurement auctions, bid increments are negative.[12] A large negative number indicates bid jumping, whereas a number close to zero indicates a small increment.

Table 2.2 shows the distribution of bid increments. In Period I, the average and the median increment are respectively 0.35 and 0.05 percent of the reserve price. These numbers are not very different in Period II: 0.56 and 0.04 for the average and median, respectively. However, bid increments display variation across auction phases. In Phases A and B of Period I, the average increment was of 0.44 and 0.43 percent, whereas in Phase C this figure was 0.31 percent (p-values against Phase A and B both $<$0.001). The same pattern can be seen in Period II. Overall, these numbers suggest that (i) there is little jump bidding in both periods, and that (ii) increments tend to be larger in Phases A and B than in Phase C.

Table 2.2: Distribution of Bid Increments

	Mean	Std. Dev.	25th Pctile	50th Pctile	# of Obs.
			Period I		
Phase A	-0.0044	0.0272	-0.0020	-0.0003	179,681
Phase B	-0.0043	0.0173	-0.0023	-0.0005	80,869
Phase C	-0.0031	0.0126	-0.0021	-0.0005	524,917
Total	-0.0035	0.0175	-0.0021	-0.0005	785,467
			Period II		
Phase A	-0.0109	0.0477	-0.0034	-0.0004	476,987
Phase B	-0.0117	0.2098	-0.0040	-0.0005	254,630
Phase C	-0.0047	0.0206	-0.0019	-0.0004	4,415,186
Total	-0.0056	0.0525	-0.0021	-0.0004	5,146,803

Notes: Table shows statistics for increments between consecutive bids normalised by the reserve price. In Period I, a total of 897,394 bids were placed in 65,593 auctions. In Period II, a total of 6,713,923 bids were placed in 336,624 auctions. Increments can sometimes be positive, as bidders may submit bids that have already been outbid. Positive increments are not counted, leaving 785,467 observed increments in Period I and 5,146,803 in Period II.

2.3.2 Timing of Bids

In section 2.2, we documented the fact that bidders place more bids in Phase C of the auction. To better capture this behaviour, Figure 2.4 plots the empirical cumulative distribution functions of the number of bids per minute, for the three auction phases. The distribution of bids per minute in Phase C dominates the

[12] In practice, increments can sometimes be positive, since bidders may submit bids that have already been outbid. In what follows, positive increments are ignored.

those of Phases B and C. This indicates that bidders delay placing their bids to Phase C.

Figure 2.4 Distribution of Number of Bids Per Minute by Phase

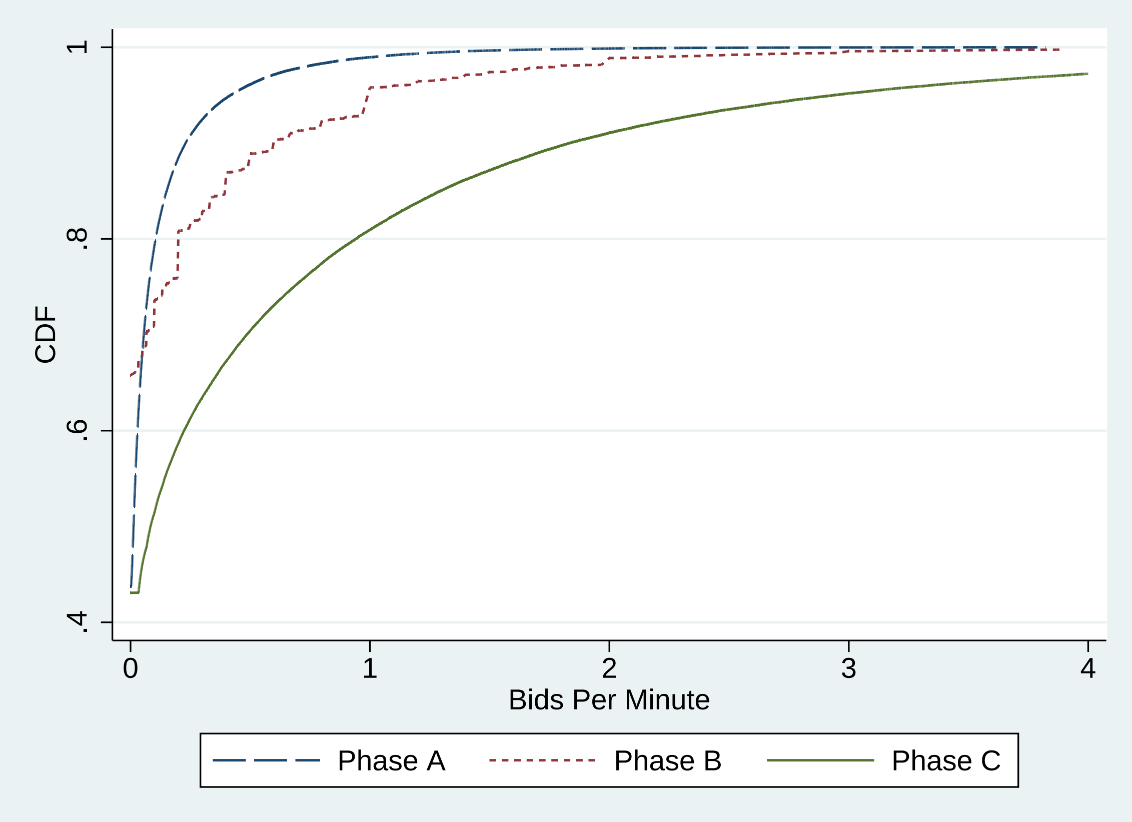

We look at this phenomenon from a different angle in Figure 2.5. For auctions in Period I, we calculate the time distance of each bidders' last bids to the end of the auction. We then plot the cumulative distribution of this distance for auctions whose duration was in the (10,11], (15,16] and (20,21] minutes intervals in Figure 2.5. All three distributions display a kink at beginning of Phase C. The figure reveals, for example, that when Phase C lasts between 15 and 16 minutes, 20 percent of bidders' last bids arrive in the last 5 minutes. Figure 2.5 also reveals that, despite the increased activity in Phase C, many bidders place their last bids in Phases A or B. For example, when Phase C lasts between 20 and 21 minutes, 40 percent of last bids arrive before Phase C starts.

We look at these numbers more closely in Table 2.3. The table reveals that, in Period I, 17 percent of auctions were resolved in Phases A or B, and 53 percent had bids placed in Phase C. The remaining 30 percent of auctions were resolved without a single bid being placed in the ascending stage–i.e., these auctions were won by the bidder submitting the highest sealed bid. The numbers for Period II are similar to those for Period I.

We now turn to the importance of the timing of bids to auction outcomes. Specifically, we look at the relationship between auction duration and winning

Figure 2.5 Cumulative Distribution over time of Bidders' Last Bids

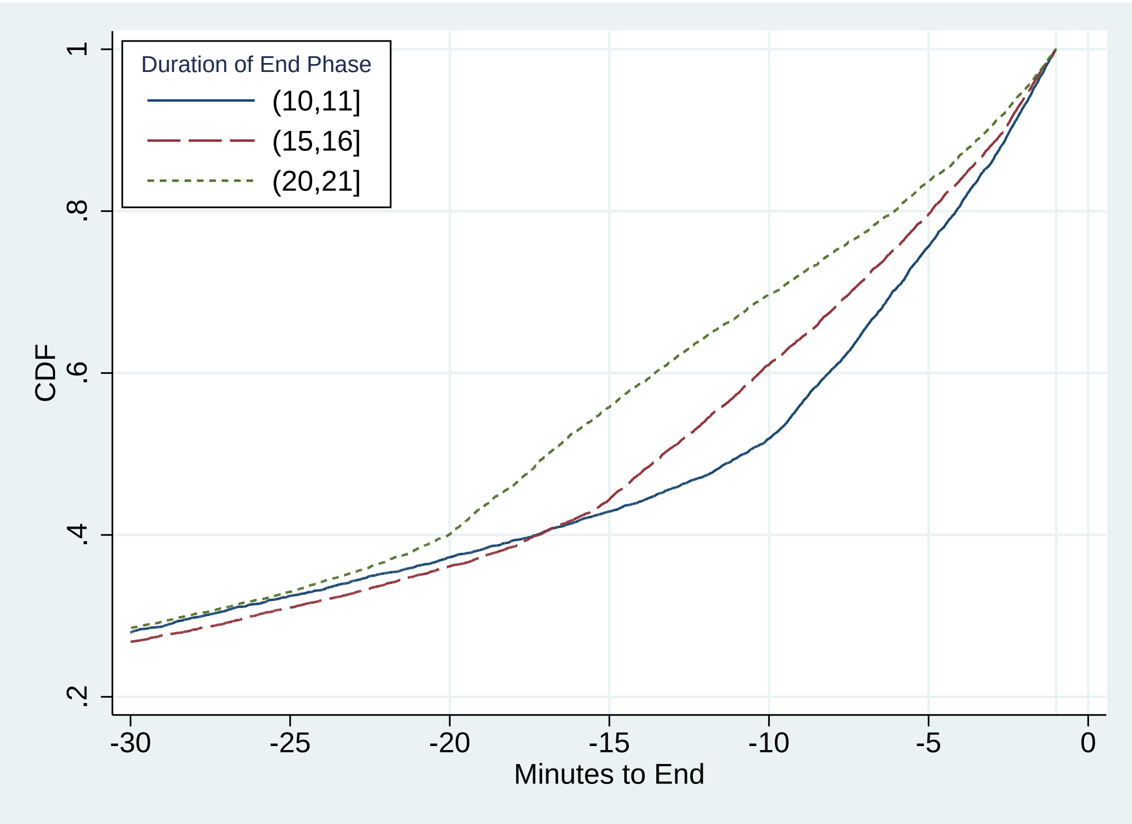

Notes: The Figure uses data from Period I only.

bids. If bidders actively use Phase C to place their final bids, then auctions with longer Phase C duration should have smaller winning bids. To capture this effect, we run regressions of the form:

$$p_t = X_t\beta + \gamma f(d_{Ct}) + \epsilon_t, \tag{2.2}$$

where p_t is the (normalized) winning bid of auction t; X_t is a vector of auction controls including, geographical and time-specific covariates, and the length of the opening and warning phases; and $f(d_{Ct})$ is a polynomial function of the

Table 2.3: Fraction of Auction's Last Bids in...

	Period 1	Period 2
Phases A or B	0.17	0.11
Phase C	0.53	0.68
Last 10 Seconds	0.08	0.09
Last 1 Minute	0.23	0.18
Last 5 Minutes	0.38	0.35
Last 10 Minutes	0.47	0.50

length of the ending phase.

Table 2.4: Effects of duration on Price

	Period 1		Period 2	
	(1)	(2)	(3)	(4)
$d_C/10$	-0.022	-0.016	0.015	-0.002
	(0.003)	(0.006)	(0.003)	(0.003)
$(d_C/10)^2$	0.004	0.001	-0.015	-0.006
	(0.001)	(0.002)	(0.001)	(0.001)
$d_B/10$		-0.032		0.007
		(0.003)		(0.001)
$(d_B/10)^2$		0.005		-0.001
		(0.001)		(0.000)
$d_A/10$		0.005		0.001
		(0.001)		(0.000)
$(d_A/10)^2$		-0.000		-0.000
		(0.000)		(0.000)
Year Dummies		✓		✓
Quarter Dummies		✓		✓
Good Dummies		✓		✓
State Dummies		✓		✓
Observation	65,593	22,270	336,624	329,529
R^2	0.003	0.180	0.023	0.185

Notes: Robust standard errors in parenthesis. Variations in the sample size within a Period are due to missing data for the duration of Phase A, as explained in section 2.2.

Table 2.4 reports the results. For each Period, we run 2 specifications. In Period I, the winning bid is decreasing and concave on the duration of Phase C (d_C). The coefficients in column (1) imply that the winning bid drops, on average, by 1.4 percent in the first 10 minutes, and by 0.6 percent in the next 10 minutes of Phase C. The introduction of other auction covariates in column (2), as expected, does not alter this result. In Period II, the winning bid is decreasing and convex on the d_C: the coefficients in column (3) imply that in the first 10 minutes the winning bid is expected to drop by 1.5 percent, and by an extra 4.5 percent in the next 10 minutes. One possible explanation for this change is that, in Period II, d_C is positively correlated with the number of bidders in the auction, and the number of bidders in the auctions is negatively correlated with the winning bid. In other words, the coefficients on d_C in columns (3) and (4) are likely to be picking up the effect of the increased number of bidders over time in ComprasNet.

To sum up, we find that (i) bidders defer bidding to the end phase of the auction; (ii) a sizeable fraction of auctions is resolved early; (iii) bid increments are typically small; (iv) large increments are more likely to occur early on the auction. In the next section, we developed a model to rationalise these observed patterns in the data.

2.4 A model of Auctions with Random Ending Time

In this section, we present a basic theoretic framework for ascending auctions with a random ending time. We closely follow the formulation presented in Ockenfels and Roth (2006) for auctions with a hard close and adjust the game for the random close case. We also perform some comparative statics exercises to highlight the key differences between hard close and a random close.

In what follows, we assume that the highest bidder wins and pays the second highest bid. This would be the case in ascending auctions with a proxy bidding system, like eBay. In absence of a proxy mechanism, as it is the case of ComprasNet, this assumption is an abstraction necessary for tractability (Milgrom and Weber (1982)). Below we present the strategic structure of an auction with random ending time.

- There is a single seller auctioning a single indivisible object.
- There are $n \geq 2$ buyers (bidders), denoted by $N = \{1, 2, \ldots, n\}$.
- Each bidder has private valuation $v_i \in [0, \bar{v}]$ drawn identically and independently according to some distribution $F(v)$.
- The minimum initial bid equals 0. (i.e. there is no reservation price for the seller[13].)
- A player can place a single bid b_i^t at any time $t \in T = \{0\} \cup \{t_1(m) = \frac{m}{m+1}; m = 1, 2, \ldots\} \cup \{1\} \cup \{t_2(m) = \frac{2m+1}{m+1}; m = 1, 2, \ldots\} \cup \{2\}$. If a player $i \in N$ at some time $t \in T$ does not bid, then we denote her bid as $b_i^t = \emptyset$. This formulation states that the auction game has four periods, two of which are divided in an infinite and countable number of subperiods. This resembles the bidding timeline of ComprasNet, see Figure 2.1.
- Every new bid of a player has to be higher than her last nonempty bid, i.e. $b_i^{t'} > b_i^t$ if $t' > t$ for $b_i^{t'} \neq \emptyset$ and $b_i^t \neq \emptyset$.
- At any given time $t \in T$, players can submit bids simultaneously without knowing what other bids are placed.
- The bid history at some time t lists all the bids placed up to that time along with the identities of the bidders.
- The auction ends either at the end of time $t = 1$ with probability $h \in (0, 1)$ or at the end of time $t = 2$ with the remaining probability, $1 - h$.

[13]For any seller's reservation price $r \in (0, \bar{v})$, the results do not change.

- Depending on the ending time realization $\bar{t}$, the highest bidder wins the auction paying the highest submitted bid from another bidder according to the bid history and the last bids placed at $\bar{t}$ (if any).

- A bidder who wins the auction at some price p earns $v_i - p$, a bidder who does not win earns 0.

- A player has time to react to another player's bid at any time $t \in T \setminus \{1,2\}$, however the reaction can not be instantaneous. Any reaction $b_i^{t'}$ to a bid b_j^t for $i \neq j$ can arrive earliest at $t' = t_1(1)$ if $t = 0$, or at $t' = t_1(m+1)$ if $t = t_1(m)$ for some m, or at $t' = t_2(m'+1)$ if $t = t_2(m')$ for some m'.

- Equal bids from different bidders are resolved by order of arrival (first bidder to submit has priority) or, if they were simultaneously submitted, at random with equal probability.

- Any bid submitted at $t \in T \setminus \{1,2\}$ is transmitted with certainty.

- A bid submitted at time $t = 2$ is successfully transmitted with probability $0 < q < 1$, where q is an exogenously given probability.

- Similarly, a bid submitted at time $t = 1$ arrives at the end of $t = 1$ with probability $q \in (0,1]$. If the auction does not end at $t = 1$, and the bid does not arrive at $t = 1$, then it arrives at $t_2(1)$[14]. This is a crucial difference between the last-minute bid at $t = 2$ and (possibly) a last-minute bid at $t = 1$.

These games are fairly different than the standard second price sealed bid auction. In particular, similar to the second price eBay-model auction, there are no dominant strategies in these games[15]. Furthermore, in these games the set of undominated strategies are rather large. In particular, a strategy with bidding above own valuation is not necessarily dominated.

Lemma 1. *A strategy where player i bids $b_i^t > v_i$ at some $t \in T$ may not be dominated.*

Proof. Let $n = 2$ (if $n > 2$, assume all additional bidders do not bid at any time). It is sufficient to show that for some types of i, bidding $b_i^t > v_i$ is better than bidding $b_i^t \leq v_i$.

Suppose that bidder j's strategy is to bid once at $t = t_1(1)$ either $b_j^{t_1(1)} = v_j$ if $b_i^0 < \tilde{v}$ for some $\tilde{v} \in (0, \bar{v})$ or to bid $b_j^{t_1(1)} = 0$ otherwise. In that case, for all types

[14] This implies that the other bidders have opportunity to reply to a bid placed at $t = 1$ with probability $1 - h$, i.e. the probability that the auction ends at $t = 2$.

[15] The proof of the claim is analogous to that of Theorem (no dominant strategies) on page 301 of Ockenfels and Roth (2006), and hence is omitted here.

bidding any amount $b_i^0 \geq \tilde{v}$ yields v_i while bidding below $\tilde{v}$ yields $v_i - v_j$ only when $v_i > v_j$. Thus bidding $b_i^0 = \bar{v}$ is weakly better than any other bid which is a strategy that results in bidding above valuation v_i, almost surely. $\square$

In what follows, we search for equilibria in strategies where bidders do not bid above their private valuations.

Lemma 2. *Every bidder $i \in N$ bidding according to the following strategy is an equilibrium:*

There exists $t_i \in \{0\} \cup \{t_1(m); m = 1, 2, \dots\}$ such that $b_i^{t_i} = v_i$ and never updates after t_i, i.e. $b_i^{t'} = \emptyset$ for all $t' > t_i$.

Proof. Consider a bidder $i \in N$. By bidding $b_i^t > v_i$ at any t, she incurs losses whenever the highest valuation among the other bidders, $v_{-i} = \max_{j \in N \setminus \{i\}}\{v_j\}$ satisfies $b_i^t > v_{-i} > v_i$. These losses would be avoided if i were never to bid higher than v_i. Next denote by $\hat{t}_{-i} = \max_{j \in N \setminus \{i\}}\{t_j\}$. This is the time of the last submitted bid from bidders other than i such that $b_j = v_j$. At $\hat{t}_{-i}$, bidder i is faced with one of two situations. She may see that there exists some bidder $j \neq i$ with value $v_j > v_i$, in which case, bidding v_i at some time $t_i < 1$ is optimal (by indifference). Alternatively, i may see that she has the highest value, in which case it is optimal for her to bid v_i at some time $t_i < 1$[16]. $\square$

Although the equilibria constructed above includes the standard equilibrium in weakly dominant strategies of second price sealed bid auctions[17], we can see that there are far more equilibria in our game. They differ in the timing of announcement of the true valuation, however they are all payoff-equivalent to the standard Vickrey equilibria. The crucial point is that all players bid their true valuations at some point $t_i < 1$.

Next, we search for equilibria that display late bidding, and more importantly sniping. Sniping will be referred to as a situation in which a player places a bid that can not be retaliated against. In our game, conditional on the game ending at $t = 2$, a bid successfully placed at $t = 2$ is a snipe. Similarly a bid successfully placed at $t = 1$ may be a snipe with probability h. We show by construction that there may exist equilibria with sniping at $t = 2$ in the presence of a random ending time. Before we state the main result consider the subgame at $t = 2$.

Lemma 3. *If a bidder has not already placed a bid equal to her true valuation before $t = 2$, i.e. $\lim_{m \to \infty} b_i^{t_2(m)} < v_i$, then it is a weakly dominant strategy to bid $b_i^2 = v_i$.*

[16] Although any bid $b_i^{t_i} = p \in (p_{-i}, v_i]$ where $p_{-i} = \max_{j \in N \setminus \{i\}}\{b_j^{t_j}\}$ works, by indifference we pick $b_i^{t_i} = v_i$.

[17] Namely, all bidders bid $b_i^0 = v_i$ and never update at any $t > 0$.

Proof. At $t = 2$, it is common knowledge that the game ends. Given valuations are private, then under the assumption that all bidders bid up to their respective valuations, it is a weakly dominant strategy to bid $b_i^2 = v_i$. The reasoning follows that of a standard second price sealed bid auction. $\square$

Observe that, despite the fact that a bid is transmitted with q probability at $t = 2$, it is nevertheless optimal to bid your true valuation. The reason is that with non-zero probability what you bid matters. In cases where your bid does not go through, a player is indifferent to what she bids. However, when her bid goes through, by bidding below her valuation she is giving up expected rents regardless, when entering $t = 2$, whether she is the highest bidder or not[18]. On the other hand, since no other bidder bids above their true valuation, bidding above own valuation yields expected losses whenever the overbid satisfies $v_{-i} > b > v_i$ where $v_{-i} = \max_{j \in N \setminus \{i\}} \{v_j\}$.

The main result is as follows.

Theorem 2.1. *There may exist symmetric perfect Bayesian equilibria where bidders with valuation above a threshold $p \in [0, \bar{v})$ snipe each other mutually at $t = 2$ and do not place any bids in $(0, 2)$.*

Proof. Without loss of generality, consider the case of 2 bidders. Then the following strategies comprise the equilibrium profile. There exists a cutoff bid $p \in [0, \bar{v})$ such that:

S1. If $v_i \leq p$, then she bids $b_i^0 = v_i$ and never updates after time $t = 0$, i.e. $b_i^t = \emptyset$ for all $t > 0$.

S2. If $v_i \geq p$, then she bids $b_i^0 = p$ at time $t = 0$. If the opponent has bid $b_j^0 < p$ then she bids $b_i^{t_1(1)} = v_i$. Otherwise, she does not update until $t = 2$, i.e. $b_i^t = \emptyset$ at any $t \in (0, 2)$. At $t = 2$ she bids $b_i^2 = v_i$. If i observes that $b_j^0 > p$ or $b_j^t \neq \emptyset$ at some $t \in (0, 2)$, then she bids $b_i^{t'} = v_i$ at the next possible period.

We can also define the beliefs as follows. At time $t = 0$ beliefs are given by the ex-ante probability distribution $F(v)$. Then the beliefs are updated at $t = t_1(1)$. If both bidders have bid below p or both have bid p at $t = 0$ then beliefs remain the same for the rest of the game in all times $t \in [t_1(1), 2]$ on the equilibrium path. Therefore, after the first bids are announced, a bidder i believes that the other bidder j has a value that is;

- Equal to her bid b_j^0 if $b_j^0 < p$ with certainty, i.e. $\mathbb{P}(v_j = b_j^0 | b_j^0 < p) = 1$,
- Distributed according to the conditional distribution between p and $\bar{v}$ if $b_j^0 = p$, i.e. $\mathbb{P}(p \leq v_j \leq v | b_j^0 = p) = \frac{F(v) - F(p)}{1 - F(p)}$ for some $v \in [p, \bar{v}]$.

[18]To see why, observe that with nonzero probability all players' last-minute bids go through.

If i has bid some $b_i^0 < p$ and j had bid $b_j^0 = p$, then on the equilibrium path, $b_j^{t_1(1)} = v_j$. Therefore, in that case i updates her belief from the conditional distribution to a belief where the other bidder j has a value that is equal to her bid $b_j^{t_1(1)}$ with certainty, i.e. $\mathbb{P}(v_j = b_j^{t_1(1)} | b_i^0 < p \text{ and } b_j^0 = p) = 1$.

If there are any deviations, i.e. a bid at $t \in (0,2)$ or a bid $b_j^0 > p$ at $t = 0$, then beliefs at those out-of-equilibrium nodes put full probability to the deviation; $\mathbb{P}(v_j = \tilde{b} | \text{ if } b_j^0 = \tilde{b} > p \text{ or } b_j^t = \tilde{b} \neq \emptyset \text{ for } 0 < t < 2) = 1$.

We will verify that the above strategy profile is subgame perfect by backwards induction. The main point is to show why it is optimal for any bidder i with value $v_i > p$ to not deviate at any any point if she sees that her opponent has bid p at $t = 0$.

Let us start by considering bidder i's strategy when she has $v_i \leq p$. She believes that with $F(p)$ probability the other bidder has valuation $v_j \leq p$, in which case bidder j bids her true valuation at $t = 0$. Similarly, with $1 - F(p)$ probability the opponent has valuation $v_j > p$ and bids p at $t - 0$. For i, bidding above v_i is weakly dominated and creates additional loss making. If she bids below her valuation, then she could update at any $t = t_1(m)$ before 1 and receive the same payoff. In other words, bidding below her valuation at $t = 0$ can not yield her higher payoff at any $t > 0$. By indifference it is therefore optimal for her to bid true valuation at the first period, which in turn yields her an expected payoff of:

$$\Pi_i^0(v_i \leq p) = F(p) \left[\left(v_i - \int_0^{v_i} \frac{vf(v)}{F(v_i)} dv \right) \frac{F(v_i)}{F(p)} \right] + (1 - F(p)) \cdot 0 = \int_0^{v_i} F(v) dv$$

Next consider the strategy of bidder i when $v_i > p$. On the equilibrium path, there are two possibilities. Either i will be faced with an opponent having a valuation below p, which she can infer at the beginning of $t = t_1(1)$. In that case, by indifference it is optimal for i to bid her true valuation at $t_1(1)$, since it guarantees her a payoff of $v_i - v_j$ where $v_j \leq p$. Alternatively, i may be faced with an opponent who also bids p at $t = 0$ which leads i to believe that v_j is distributed in the $[p, \bar{v}]$ interval with distribution $\frac{F(v) - F(p)}{1 - F(p)}$. Let us evaluate under what conditions it is indeed optimal for i to not update her initial bid of p until the last instance $t = 2$.

We know from Lemma 3 that it is optimal for i to bid her true valuation at $t = 2$. Since both bidders have bid p at the beginning of the game, one of them is randomly chosen to be the high bidder[19]. Therefore, by bidding $b_i^2 = v_i$ at the last instance in a subgame where both i and j's bids are at p, the expected payoffs

[19] We will refer to the randomly chosen highest bidder as the *leader* and the other bidder will be referred to as the *follower*.

to i when she is the leader or the follower are equal to, respectively:

$$\Pi_i^2(v_i, v_j > p)^L = q^2 S(v_i|v_j > p) + (1-q)(v_i - p)$$
$$\Pi_i^2(v_i, v_j > p)^F = q^2 S(v_i|v_j > p) + (1-q)q(v_i - p)$$

where $S(v_i|v_j > p)$ is the expected information rents conditional on the opponent also having valuation above p and is equal to:

$$\begin{aligned} S(v_i|v_j > p) &= \left[v_i - \int_p^{v_i} \frac{vf(v)}{F(v_i) - F(p)} dv\right] \frac{F(v_i) - F(p)}{1 - F(p)} \\ &= \int_p^{v_i} \frac{F(v)}{1 - F(p)} dv - (v_i - p)\frac{F(p)}{1 - F(p)} \end{aligned}$$

Firstly, observe that the expected payoff to the leader, $\Pi_i^2(v_i, v_j > p)^L$, is strictly larger than that of the follower. Secondly, in both of the expected payoffs, we account for the friction of probabilistic bid transmission. In that regard, the expected payoff has the standard information rents appear only with q^2 probability, since it is the probability with which both bidders can successfully update their bids. Similarly, with $(1-q)q$ probability only one of the bidders' bids goes through. If indeed only i's bid goes through, then i wins with certainty having to pay only p. On the other hand, if it is only j whose bid was updated, then i loses with certainty earning 0. Finally with $(1-q)^2$ probability, neither bidder's bid is updated from p. In that case i gets to keep the whole rent of $(v_i - p)$ to herself only if she is the leader.

Let us now go back one step further and consider a subgame beginning at time $t = t_2(m)$ for some m. We check whether it is possible to sustain an equilibrium in which no bids are placed before $t = 2$, whenever the auction has reached some period $t_2(m)$ at price p. If i sticks to the suggested strategy profile then, she gets an expected payoff given by $\Pi_i^2(v_i, v_j > p)^L$ or $\Pi_i^2(v_i, v_j > p)^F$ depending on whether she is the leader or the follower, respectively. If i were to deviate and place a bid (above p) at some $t_2(m)$, then j retaliates by bidding her true valuation. Hence, i's (weakly) most profitable deviation at $t_2(m)$ is to bid her valuation[20]. By bidding $b_i^{t_2(m)} = v_i$, the expected payoff i gets is equal to $S(v_i|v_j > p)$, since j retaliates at $t_2(m+1)$ by bidding her valuation hence yielding i her expected information rent with certainty. One important point to observe is that the deviation and the reply arrives without any frictions. Then i would rather delay bidding to $t = 2$ than deviating if $\Pi_i^2(v_i, v_j > p)$ is weakly larger than $S(v_i|v_j > p)$. Since i's incentives to deviate are higher if she is the follower, it is

[20]To see why, if she deviates by bidding above v_i, after j bids her true valuation, i might find herself making a loss. Any deviation to some bid between p and v_i could eventually be updated to v_i at some later time $t_2(m')$ for $m' > m$. In that case, by indifference, we consider a single deviation to v_i at $t_2(m)$.

enough to check only the follower's inequality:

$$\begin{aligned} q^2 S(v_i|v_j > p) + (1-q)q(v_i - p) &\geq S(v_i|v_j > p) \\ \iff (v_i - p)\left[\frac{F(p)+q}{1+q}\right] &\geq \int_p^{v_i} F(v)dv \end{aligned}$$

The last line gives a condition which should be satisfied for all $v_i > p$. Observe that the left hand side is linear in v_i while the right hand side is strictly convex in v_i.[21] Since both sides are positive valued for all $v_i \in (p, \bar{v}]$ and the right hand side of the inequality convex, it is sufficient to check whether the inequality is satisfied for $v_i = \bar{v}$. This gives us our first condition:

$$(\bar{v} - p)\left[\frac{F(p)+q}{1+q}\right] \geq \int_p^{\bar{v}} F(v)dv \tag{2.3}$$

Next consider bidder i's incentives at time $t = 1$. Remember that, the game might end at $t = 1$ with some probability h or continue on until $t = 2$ with the remaining probability. If she sticks to the equilibrium strategy, then as a leader and as a follower she expects to receive, respectively:

$$\begin{aligned} \Pi_i^1(v_i, v_j > p)^L &= h \cdot (v_i - p) + (1-h)\left[q^2 S(v_i|v_j > p) + (1-q)(v_i - p)\right] \\ \Pi_i^1(v_i, v_j > p)^F &= h \cdot 0 + (1-h)\left[q^2 S(v_i|v_j > p) + (1-q)q(v_i - p)\right] \end{aligned}$$

Observe that, again, $\Pi_i^1(v_i, v_j > p)^L$ is strictly larger than $\Pi_i^1(v_i, v_j > p)$. If i deviates at $t = 1$ then her bid goes through with probability q successfully at $t = 1$, while with remaining probability bid goes through at $t_2(1)$. In case the game does not end at $t = 1$ and i deviates, then j will learn this with certainty by the end of $t_2(1)$ and retaliate at $t_2(2)$ by bidding v_j. Now let us evaluate the expected payoff i accrues from deviating. By deviating at $t = 1$, i takes two risks: first, her bid may not go through; second, the auction may not end at $t = 1$. Given that bidding true valuation is the most profitable deviation due to the same reasons as before, then the deviation yields an expected payoff equal to $h \cdot q(v_i - p) + (1-h)S(v_i|v_j > p)$.

In order i not to deviate, her equilibrium path expected payoff must be larger than or equal to the deviation expected payoff. Again, the incentives to deviate for the follower are larger as her equilibrium path payoff is smaller. Evaluating

[21] This is because the first derivative at any $v_i > p$ equals $F(v_i) \geq 0$ and the second derivative equals $f(v_i) > 0, \forall v_i \in [0, \bar{v}]$.

the inequality for the follower yields the following condition:

$$(1-h)\Big[q^2 S(v_i|v_j>p)+(1-q)q(v_i-p)\Big] \geq h\cdot q(v_i-p)+(1-h)S(v_i|v_j>p)$$
$$\Longleftrightarrow \quad (v_i-p)\left[\frac{F(p)+q}{1+q}-\frac{h\cdot q}{1-h}\frac{(1-F(p))}{(1-q^2)}\right] \geq \int_p^{v_i} F(v)dv$$

Since the right hand side of the last inequality is strictly convex while the left hand side is linear in v_i and that both sides are always positive, it is sufficient to check that inequality is satisfied for $\bar{v}$ to imply it holds at all $v_i \in [p,\bar{v}]$. This gives us the following second condition:

$$(\bar{v}-p)\left[\frac{F(p)+q}{1+q}-\frac{h\cdot q}{1-h}\frac{(1-F(p))}{(1-q^2)}\right] \geq \int_p^{\bar{v}} F(v)dv \tag{2.4}$$

It is immediate to observe that condition (2.3) is implied by the latter condition (2.4). This is simply because the constant term in the brackets is smaller, making it is easier for the inequality to be satisfied.

Going further back, at any time $t=t_1(m)$, if i sticks to the equilibrium strategy, then her expected payoff $\Pi_i^{t_1(m)}(v,v_j>p)$ is equivalent to $\Pi_i^1(v_i,v_j>p)^L$ or $\Pi_i^1(v_i,v_j>p)^F$ depending on whether she is the leader or the follower. On the other hand, deviating by bidding $b_i^{t_1(m)}$ at some $t_1(m)$ yields her an expected payoff of $S(v_i|v_j>p)$ with certainty. This is different than at $t=1$, simply because at any $t_1(m)$ there exists the next period $t_1(m+1)$ where the opponent retaliates with certainty. Concentrating on the follower's incentives, she never finds it optimal to deviate when the following inequality is satisfied:

$$(1-h)\Big[q^2 S(v_i|v_j>p)+(1-q)q(v_i-p)\Big] \geq S(v_i|v_j>p)$$
$$\Longleftrightarrow \quad (v_i-p)\left[1-\frac{(1-(1-h)q)}{(1-(1-h)q^2)}(1-F(p))\right] \geq \int_p^{v_i} F(v)dv$$

By evaluating the above condition at $v_i=\bar{v}$ we get the following condition:

$$(\bar{v}-p)\left[1-\frac{(1-(1-h)q_2)}{(1-(1-h)q_2^2)}(1-F(p))\right] \geq \int_p^{\bar{v}} F(v)dv \tag{2.5}$$

Simple algebra yields that the constant on the left hand side from the last condition (2.5) is larger than that in condition (2.4). Hence the previous condition implies condition (2.5) stated above.

Finally let us consider i's options at $t=0$. If she bids p, then with probability $F(p)$ the opponent has valuation below p in which case she gets v_i-v_j for sure. With remaining probability $(1-F(p))$, opponent also bids p and from there on

i receives the continuation payoffs[22]. Combining all of the above, we get that i's expected payoff from equilibrium path play at $t = 0$ is given as follows:

$$\Pi_i^0(v_i > p) = \left[\left(\frac{v_i - p}{2}\right)[1 + F(p)] - \int_p^{v_i} F(v)dv\right](1 - (1-h)q^2) + \int_0^{v_i} F(v)dv$$

If i were to deviate at $t = 0$ by bidding above p, and in particular bidding v_i, then i would get the standard expected information rent, denoted $S(v_i)$, where $S(v_i) = \int_0^{v_i} F(v)dv$.[23]

On the other hand, if i were to deviate by bidding some $b_i^0 = \tilde{b} < p$, then the opponent would bid v_j latest by $t_1(1)$. To be more precise, if $v_j < p$ then $b_j^0 = v_j$, anyways, and i has no gain in deviating by underbidding at $t = 0$. If $v_j > p$, on the other hand, then j initially bids $b_j^0 = p$. After observing i's deviation to $\tilde{b} < p$ at $t = 0$, the opponent would bid $b_j^{t_1(1)} = v_j$. In any case, underbidding would result in i receiving an expected deviation payoff of $S(v_i)$, as well. Altogether, bidder i bids $b_i^0 = p$ only when her equilibrium path expected payoff is larger than or equal to $S(v_i)$. This inequality is shown below:

$$\begin{aligned} \Pi_i^0(v_i > p) &\geq \int_0^{v_i} F(v)dv \\ \iff (v_i - p)\left[\frac{1 + F(p)}{2}\right] &\geq \int_p^{v_i} F(v)dv \end{aligned}$$

Since the constant in brackets on the left hand side is even larger than $\frac{F(p)+q}{1+q}$, condition (2.4) also implies this last condition, similar to the previous cases.

Therefore the strategy profile described by S1 and S2 along with the beliefs form a perfect Bayesian equilibrium, whenever the parameters of the model satisfy condition (2.4), which is shown below:

$$(\bar{v} - p)\left[\frac{F(p) + q}{1 + q} - \frac{h \cdot q}{1 - h}\frac{(1 - F(p))}{(1 - q^2)}\right] \geq \int_p^{\bar{v}} F(v)dv$$

□

In Theorem 2.1, we have proved that for a given set of parameters q, h and a given probability distribution $F(v)$, it is possible to have equilibria with a threshold $p \in [0, \bar{v})$ where bidders with valuations above the threshold mutually snipe each other at the end of the game. We provided a sufficient condition to have such equilibria. It was not surprising that the most relevant condition comes from the no deviation check at period $t = 1$. The reason is that, a bidder with high valu-

[22]Note that, if both v_i and v_j is larger than p, then i becomes the leader or the follower with equal probability after $t = 0$, which is accounted for in the continuation payoffs.

[23]We can think of $S(v_i)$ as a special case of $S(v_i | v_j > p)$ where $p = 0$. Remember that $S(v_i | v_j > p)$ is the expected information rent conditional on the opponent having valuation higher than some p. When $p = 0$ it is simply equivalent to the unconditional expected rent.

ation (higher than p) has the most incentives to deviate at $t = 1$ when he is faced with another bidder who has bid p at $t = 0$ and has not updated her bid. A deviation at $t = 1$, if succesful, will generate a big surplus whenever $t = 1$ is the terminal node. Deviation at any other period, however, will be retaliated against with certainty. Therefore, it is sufficient to make sure that high valuation bidders will not want to deviate even at $t = 1$ which is equivalent to the condition (2.4) being satisfied.

Several remarks are in order regarding the equilibrium threshold bid, p. Firstly, observe that in the characterization of sniping equilibria, we exclude $\bar{v}$ from the set of possible threshold values. The reason is that for any set of parameters and any probability distribution, the threshold $p = \bar{v}$ is always an equilibrium. However there is no sniping occuring in those equilibria as they are simply those where all bidders bid their true valuations at $t = 0$[24]. This is not surprising, as it is a special case of those equilibria characterized in Lemma 2, where $t_i = 0$, and hence $b_i^0 = v_i$, for all $i \in N$.

Secondly, it may be the case that there are multiple values p's that satisfy the equilibrium condition (2.4). Let us denote by $\mathcal{P} = \{p \in [0, \bar{v}] | (2.4) \text{ is satisfied.}\}$. Then, it is ambiguous which of those multiple equilibria would be played. One intuitive selection argument would be to pick $p^* \in \mathcal{P}$ which maximizes the expected ex-ante payoff of a bidder[25]. Let us denote by Π_i, the ex-ante expected payoff to i which is equal to the following:

$$
\begin{aligned}
\Pi_i =& F(p)\left[\int_0^p \Pi_i^0(v_i < p)\frac{f(v_i)}{F(p)}dv_i\right] + (1 - F(p))\left[\int_p^{\bar{v}} \Pi_i^0(v_i > p)\frac{f(v_i)}{1 - F(p)}dv_i\right] \\
=& \int_0^{\bar{v}} \left[\int_0^{v_i} F(v)dv\right] dF(v_i) + \\
&+ \int_p^{\bar{v}} \left[\left[\frac{v_i - p}{2}[1 + F(p)] - \int_p^{v_i} F(v)dv\right]\left(1 - (1 - h)q^2\right)\right] dF(v_i)
\end{aligned}
$$

In order to maximize Π_i over p, let us consider the first derivative with respect to it:

$$
\begin{aligned}
\frac{d\Pi_i}{dp} &= \int_p^{\bar{v}} \left[f(p)\left(\frac{v_i - p}{2}\right) - \frac{1 - F(p)}{2}\right]\left(1 - (1 - h)q^2\right)dF(v_i) \\
&= \kappa\left[f(p)\left[(\bar{v} - p) - \int_p^{\bar{v}} F(v_i)dv_i\right] - \left[1 - F(p)\right]^2\right], \qquad (2.6)
\end{aligned}
$$

where $\kappa = \frac{\left(1 - (1 - h)q^2\right)}{2}$. Any $p \in \mathcal{P}$ that satisfies the first order condition is a local

[24] Only bidders with valuation $v_i \leq p = \bar{v}$ bid p at $t = 0$ and then snipe each other mutually at $t = 2$. However probability of are having a bidder with valuation above $\bar{v}$ is 0 which implies that there is no sniping occuring in such an equilibrium.

[25] Ex-ante is used in the sense of before $t = 0$, i.e. before bidders learn their true valuation.

maximizer so long as the second derivative, shown below, is negative:

$$\frac{d^2\Pi_i}{dp^2} = \kappa \left[f'(p) \left[(\bar{v} - p) - \int_p^{\bar{v}} F(v_i)dv_i \right] + f(p)\left[1 - F(p)\right] \right]$$

Below we provide a numeric example where we have equilibria as per Theorem 2.1:

Example 1. *Consider the private valuations to be drawn over the support $[0,1]$ (i.e. $\bar{v} = 1$) according to the distribution $F(v) = v^2$. Let $q\frac{3}{4}$ and $h = \frac{1}{25}$. Then we can only have sniping equilibria with threshold bid $p \in [0,1]$ whenever condition (2.4) is satisfied:*

$$\begin{aligned} & (\bar{v} - p)\left[\frac{F(p)+q}{1+q} - \frac{h \cdot q}{1-h}\frac{(1-F(p))}{(1-q^2)}\right] \geq \int_p^{\bar{v}} F(v)dv \\ \iff \quad & (1-p)\left[\frac{p^2+(3/4)}{1+(3/4)} - \frac{(1/25)(3/4)}{(24/25)}\frac{(1-p^2)}{(1-(9/16))}\right] \geq \int_p^{\bar{v}} F(v)dv \\ \iff \quad & (1-p)(1-13p) \geq 0 \end{aligned}$$

Hence $\mathcal{P} = [0, 1/13]$. In other words, there may be a sniping equilibrium for any $p \in [0, 1/7]$. If we consider equilibrium selection, then we can see the first derivative of Π_i for the particular parameters and distribution is equal to:

$$\begin{aligned} \frac{d\Pi_i}{dp} &= \frac{23}{100}\left[2p\left[1 - p - \frac{1-p^3}{3}\right] - (1-p^2)^2\right] \\ &= \underbrace{-\frac{23}{100}(1-p)^2\left[(1+p)^2 + 2\right]}_{<0 \text{ for all } p\in[0,1) \text{ and } =0 \text{ when } p=1} \end{aligned}$$

Since the FOC holds only at $p = 1$, evaluating the second derivative, we get $\frac{d^2\Pi}{dp^2}|_{p=1} = 0$ which implies that $p = 1$ is not a local maximum. On the other hand the first derivative of Π_i is negative for the remaining range of p values. Among all other candidate thresholds $p \in \mathcal{P} = [0, 1/7]$, the ex-ante expected payoff is at a maximum when $p = 0$. In other words, it is best for a bidder from an ex-ante point of view, that all the bidders "collude" at lowest price and snipe each other mutually at the end of the game.

2.4.1 Comparative Statics

We now analyse the robustness of the sniping equilibria with respect to the parameters of our model, namely q and h. We have already established in Theorem 2.1, that in our model existence of sniping equilibria depends on a particular condition given by (2.4). For a given probability distribution $F(v)$ and threshold price p, we investigate whether it is more or less likely for sniping equilibria to exist, by varying the parameters q and h. In the exercises below, we assume that

$F(v)$ and p are such that they satisfy the condition (2.4) for some values of the parameters and we investigate how the inequality in the condition changes in the parameters.

The two parameters h and q matter for existence of sniping equilibria. When the value of parameter h varies, the only term in condition (2.4) that changes is the constant inside the brackets. Taking the first derivative of the term in brackets with respect to h, we get the following expression:

$$\frac{\partial}{\partial h}\left[\frac{F(p)+q}{1+q}-\frac{h\cdot q}{1-h}\frac{(1-F(p))}{(1-q^2)}\right]=-\frac{q\big(1-F(p)\big)}{(1-h)^2(1-q^2)}<0$$

Hence as h increases, the term in brackets from condition (2.4) gets smaller, in turn making it harder for the sniping equilibrium to be satisfied, all else equal. The intuition is clear; the higher is h, the more likely it becomes that the game ends at the end of period $t = 1$. Therefore, the more tempting it becomes to deviate and bid at some time $t \in (0, 1]$. In particular, remember that we evaluated the condition using the expected payoffs of a follower on the equilibrium path and a follower receives a payoff of 0 if the game were to end at $t = 1$. Then it is very sensible that a follower would experience a rise in incentives to deviating when probability of game ending at $t = 1$ becomes larger.

Another important point is that the condition is easiest to satisfy when $h = 0$ at which value the term in the brackets takes the largest value. Note that the special case in which $h = 0$ represents a hard close ending. Therefore, we can argue that introducing random ending time lowers the incentives of collusive sniping equilibria, merely from the fact that it is harder for those equilibria to be sustained when $h > 0$.

Next, let us consider the parameter q, which represents the sniping friction at periods $t = 1$ and $t = 2$. When the value of parameter q changes, we have a different picture. Taking the first derivative of the bracket term yields the following expression:

$$\frac{\partial}{\partial q}\left[\frac{F(p)+q}{1+q}-\frac{h\cdot q}{1-h}\frac{(1-F(p))}{(1-q^2)}\right]=\frac{\big(1-F(p)\big)}{(1+q)^2}\left[1-\frac{h}{(1-h)}\frac{1-3q^2}{(1-q)^2(1-q^2)}\right]$$

We can observe that the derivative may take both positive and negative values. Furthermore, if $h < \frac{1}{2}$ or that it is more likely for the auction to end at $t = 2$, then the constant in brackets from Condition (2.4) achieves a maximum at some $q^* \in (0, 1)$ or an intermediary value[26]. This implies that the impact of q is non-monotonic.

We can summarise the intuition as follows. When q is too low (below q^*), the

[26]To see why, the constant evaluated at $q = 0$ is positive, and it has a strictly positive derivative. Also as q approaches 1, the constant in brackets approaches $-\infty$.

last minute mutual snipes are very imprecise. A bidder who enters $t = 2$ as a follower has to sacrifice some of the surplus in order to maintain lower prices. The less precise snipes are, the more are the potential surplus losses. On the other hand, if q is too high (above q^*), then the contrary is true. Since the snipes are very precise, there are less gains from waiting[27]. Therefore, it becomes harder to sustain sniping equilibria all else equal.

2.5 Concluding Remarks

This chapter analysed auctions with a random ending time. We first established the following stylised facts about bidding behaviour in ComprasNet auctions with a random close: (i) bidders defer bidding to the end phase of the auction; (ii) a sizeable fraction of auctions is resolved "early"; (iii) bid increments are typically small; and (iv) large increments are more likely to occur early in the auction. We then developed a model to rationalise these observed patterns in the data.

Our findings imply that, contrary to the intentions of those who designed the ComprasNet rules, ending the auction at random does not prevent late-bidding. Bidders actively use the last minutes of the auction, outbidding previous bids by tiny amounts. We observe that, when the auction duration is randomly extended, the winning bid decreases significantly. This in turn may explain why ComprasNet decided to change its algorithm to allow for more bidding time. However, by changing the ending rule, ComprasNet also changed bidders' time incentives, who again use the last minutes of the auction to bid. We now conclude with some conjectures and directions for future research.

One likely consequence of ending an auction at random is that it is not ex-post efficient: the allocation is essentially decided randomly among bidders actively engaged in late bidding. Taking entry and bid history as given, this is also a source of concern for revenues: if given more time, bidders would carry on outbidding each other, thus decreasing price. However, the introduction of randomness in the final allocation might give weak bidders better chances of winning the auction. If entry is costly, such uncertainty encourages weak bidders to participate in the auction. The positive effect of increased participation thus mitigates, and may even outweigh, the negative effect on revenues conditional on entry.

One alternative policy is to introduce a proxy-bidding system, similar to that of eBay. Again, it is not clear whether bidders would place their best bids early on the auction, as the eBay experience teaches us. However it would allow bidders the option of placing their best bid at the outset without having to worry about over-bidding. In essence, a proxy-bidding system would reduce monitoring costs

[27] Remember that the gain for a high value bidder who enters $t = 2$ as a follower in a sniping equilibrium is the possibility of only her bid going through at $t = 2$ which guarantees her the purchase at low price.

for bidders – that is, the cost of time spent on closely following the auction until it closes.

These conjectures call for further inquiry. Experimental evidence can be particularly useful for deciding on potential changes to the design of the bidding platform, such as the introduction of a proxy-bidding system.

Chapter 3

A comparison between offline and online auctions

Abstract

This chapter analyses offline auctions held in ComprasNet. Offline auctions were common in the early days of ComprasNet, when the online bidding software was not fully developed. These offline auctions have a two-stage design. Two-stage designs have long been proposed by the theoretical literature, but there is virtually no empirical work analysing data from such designs apart from experimental studies. I first present a stylised model of bidding in a two-stage auction similar to that used by ComprasNet. I then explore bidding behaviour in offline auctions and confront the predictions of the model with the data. The model is currently unsatisfactory, as it fails to capture some key features of the data. I discuss possible modifications of the model to rationalize observed bidding behaviour. I then compare online and offline auctions in terms of their outcomes, taking advantage of a change in regulation that required purchasing units to use online auctions. This change introduces exogenous variation on the choice of the auction format, which can be used to identify the effects of auction rules on outcomes. I find that offline auctions attract substantially fewer bidders than online auctions and that this results in higher procurement costs. The analysis however cannot disentangle the effects of auction rules from those of the way the auction is held (online vs. offline).

3.1 Introduction

There is a growing literature on two-stage auctions that combine features of different standard formats (e.g., Ye (2007), Hörner and Sahuguet (2007), Perry and Reny (2002)). Two-stage designs have a wide range of applications. For example, Perry et al. (2000) and Dutra and Menezes (2002) observe that a number of privatisation processes of state-owned companies in Italy and Brazil were auctioned off using two-stage designs, while Ye (2007) documents widespread use of two-stage auctions in the context of high-valued assets in the US electricity generation industry. Moreover, some authors argue that two-stage auctions can improve on standard designs in certain contexts. For example, Klemperer (1998) proposed a two-stage design in the context of airwaves auctions, where standard auction formats faced the trade-off between efficient network formation and maximizing revenue.

Some ComprasNet auctions use a two-stage auction design as an alternative to the standard online format discussed in the previous two chapters. In the first stage, bidders submit (binding) sealed bids. A qualifying rule selects the best bids from the first stage stage. In the second stage, qualified bidders engage in an ascending auction in which the minimum bid is set at the best bid submitted in the first stage. These auctions are held offline and were mostly used when the online bidding platform was not yet ready for large-scale utilisation. The offline, two-stage format uses a unique design closely related to the Dutch-Anglo auction proposed by Klemperer (1998). To the best of my knowledge, there are no studies using data for this type of auction.

The objectives of this chapter are twofold. First, I present a stylised model of bidding in a two-stage auction similar to the one used in ComprasNet. I then explore bidding behaviour in offline auctions, and confront the predictions of the model with the data. The model is currently unsatisfactory, as it fails to capture some key features of the data. I discuss possible modifications of the model to rationalise observed bidding behaviour.

Second, I estimate the impact that the two auction formats have on auction outcomes. I take advantage of a change in regulation that required purchasing units to use online auctions. This change introduces exogenous variation in the choice of the auction format, which can be used to identify the effects of auction rules on outcomes. I find that offline auctions attract substantially fewer bidders than online auctions and that this results in higher procurement costs. One key challenge when comparing the online and offline formats is to disentangle the effects of the medium–online or offline–from those of the auction rules.

This chapter contributes to the literature on two-stage auction designs. Two-stage designs have received considerable attention from the theoretical literature (e.g., Ye (2007), Hörner and Sahuguet (2007), Perry and Reny (2002)), but empirical analysis has been restricted to experimental studies (Kagel et al. (2008)). To the best of my knowledge, this is the first study using field data generated by a two-stage auction design.

The remainder of this chapter is organized as follows. Section 3.2 describes the mechanism used in offline auctions and present a model of these auctions. Section 3.3 gives some background on the usage of offline auctions in ComprasNet, and describes the data used. Section 3.4 presents results of reduced form regressions that test for differences in auction outcomes between online and offline auctions. Section 3.5 concludes.

3.2 The Offline Auction Mechanism

In this section I describe how offline auctions work in ComprasNet. The online format was extensively described in Chapter 2 and will be referred to here for comparison purposes only. I also focus on the auction game itself, as the environment surrounding the auction game was described in section 1.3.

- Stage 1: Bidders submit a (binding) sealed bid before a pre-specified deadline, after which no bidders may enter the auction.[1] After receiving bids from all participants, the auctioneer opens the sealed bids and orders them. All bids become common knowledge to all participants.

- Qualifying Rule: All bids no greater than 10% of the best (i.e., lowest) bid qualify for Round 2; if there are less than 3 bids in this situation, the best 3 bids qualify for Round 2.

- Stage 2: Qualified bidders can now verbally call new bids. Bidding starts from the highest (i.e., worst) qualified bid, followed by the 2nd highest and so forth. There is an *activity rule*: if a bidder skips her turn, she withdraws from the auction and cannot return. All bids must be strictly smaller than any previous bids. In particular, the first bid placed in Round 2 must be smaller than the low bid of Round 1. Bidding ends when there is only one bidder who has not quit. The winner pays her last bid. If no bidders wish to place new bids in Round 2, the low bidder in Round 1 wins the auction and pays her bid.

Example Figure 3.1 illustrates one typical physical auction. Six bidders entered this auction, as indicated in Round 1. The auctioneer opens the envelopes and ranks the bids. The low bid in Round 1 came from Bidder C, $\$63,000$. According to the qualifying rule, all bids up to $\$63,000 \times 1.1 = \$69,300$ qualify for Round

[1] Sometimes, the deadline coincides with bidders arriving at the bidding room.

2. There are no bids in this range, so the bottom 3 bids (i.e., bidders A, B and C) qualify. Bidder A placed the highest Round 1 bid among the qualified bids, so she is the first one to bid in Round 2. She must either bid lower than \$63,000 or quit. She bids \$62,000. Next, it's B's turn. The auction goes on this way until C decides to withdraw. The next one to withdraw is bidder A, so B wins and pays her last bid, \$51,900.

Figure 3.1 A typical offline auction

Round	Bidder	Bid	Qualify
1	Z	119.88	
1	Y	118.80	
1	X	116.64	
1	A	90.00	✓
1	B	75.24	✓
1	C	63.00	✓
2	A	62.00	
2	B	61.00	
2	C	60.00	
⋮	⋮	⋮	⋮
2	C	52.50	
2	A	52.49	
2	B	52.00	
2	C	Gives up	
2	A	51.99	
2	B	51.90	
2	A	Gives up	

Discussion In both online and offline auctions bidders submit a sealed bid and then engage in a descending auction. But the two formats differ in three key respects, each of which can have implications for bidding behaviour and auction outcomes. First, all bidders must be physically present at a public meeting in offline auctions. Participation costs are thus likely to be greater. More information is revealed to bidders, since bidders learn all the bids placed and the respective identity of other bidders. In contrast, bidders in online auctions only learn the current low bid. Under the IPV framework this difference should not matter for bidding behaviour, but combined with other information structures such difference may translate into different bidding behaviours across formats. Second, offline auctions have a qualifying rule which is absent in online auctions. This qualifying rule is what makes the offline format a two-stage auction game. It creates a trade-off between the probability that a bidder qualifies for Stage 2 and the price she will have to pay in case she wins the auction. As a result, bidders must be "careful" with their opening bids, as they may not have a chance to revise them. In contrast, online auctions the opening bid can always be revised. Third, offline auctions also have an activity rule. The activity rule reveals the drop out points of each bidder, up to bid increments. This has implications to how the data on loosing bids can be interpreted.

3.2.1 A Model of Offline Auctions

Consider the following auction game. N risk-neutral bidders bid for a single object. Bidders have independent private values: bidders' i valuation is $v_i \sim F(v)$, with support $[v^L, v^H]$, and v_i is independent from $\{v_j\}_{i \neq j}$. All this is common knowledge. The auction has two stages. The rules are as follows:

- Stage 1: Bidders submit a sealed bid. Denote a bid submitted in Stage 1 by b_i, and the highest bid by $b^{(1)}$, the second highest by $b^{(2)}$, and so on. Let $k \in (0, 1)$. All bids in $[kb^{(1)}, b^{(1)}]$ qualify to Stage 2. If there are less than 3 bids $[kb^{(1)}, b^{(1)}]$, the best 3 bids qualify to Stage 2.

- Stage 2: All bids from Stage 1 are revealed. Qualified bidders play an English "button" auction; they can choose to update their bids, or to maintain their Stage 1 bids. If they choose to update, their bid must be no lesser than the high bid from Stage 1. If no bidders choose to update their bids, the high bidder from Round 1 wins the auction and pays his bid.

The next two propositions show that there is no symmetric, strictly increasing equilibrium bidding function Stage 1. I assume throughout that in Stage 2 players use the weakly dominant strategy of bidding up to their valuations.

Proposition 3.1. *Suppose b is an equilibrium bidding function in Stage 1, and that $b' > 0$. The Revenue Equivalence Theorem must hold.*

Proof. Since $b' > 0$, the high-valuation bidder goes to Stage 2 and wins the auction with probability 1; a bidder with valuation $\underline{v}$ has an expected payoff of zero, since she only wins when all her opponents have valuation $\underline{v}$, and in that case she must bid $\underline{v}$. Therefore the Revenue Equivalence Theorem holds. □

Proposition 3.2. *Suppose b is an equilibrium bidding function in Stage 1, and that $b' > 0$. If $b(x) > \underline{v}$ for any x, the Revenue Equivalence Theorem cannot hold*

Proof. Without loss of generality, denote bidder 1's valuation by x and the highest valuation of the remaining $N-1$ bidders by Y_1^{N-1}. Since $b' > 0$, the high-valuation bidder goes to Stage 2 and wins the auction. So bidder 1 wins if and only if $x > Y_1^{N-1}$. If she wins, she pays the maximum between her Stage 1 bid and the highest valuation among her opponents. Hence we can write bidder 1's expected payment as

$$
\begin{aligned}
m(x) =& E[\max b(x), Y_1 | x \geq Y_1] \times \Pr[x \geq Y_1] \\
=& b(x) \Pr[b(x) \geq Y_1 | x \geq Y_1] \times \Pr[x \geq Y_1] \\
& + E[Y_1 | x \geq Y_1, b(x) < Y_1] \times \Pr[b(x) < Y_1 | x \geq Y_1] \times \Pr[x \geq Y_1] \\
=& b(x) G_1(b(x)) + \int_{b(x)}^{x} z g_1(z) dz, \qquad (3.1)
\end{aligned}
$$

where $G_1(t) = \Pr[Y_1 \leq t]$ and $g_1 = G_1'$. By Proposition 3.1, the Revenue Equivalence Theorem holds so we can equate $m(x)$ in (3.1) to the expected payment in a Vickrey auction:

$$
\begin{aligned}
& b(x) G_1(b(x)) + \int_{b(x)}^{x} z g_1(z) dz = \int_{0}^{x} z g_1(z) dz, \quad \forall x \\
\iff & b'(x) G_1(b(x)) + b(x) g_1(b(x)) b'(x) + \\
& + x g_1(x) - b(x) g_1(b(x)) b'(x) = x g_1(x), \quad \forall x \\
\iff & b'(x) G_1(b(x)) = 0 \quad \forall x
\end{aligned}
$$

The first line makes it clear that for the Revenue Equivalence Theorem to hold, we must have $b(x) < x$, i.e., bidders cannot bid more than their valuation in Stage 1. The second line is obtained by differentiating both side of the preceding line. The last line says that for the Revenue Equivalence Theorem to hold, we must have either that $b' = 0$, which contradicts the assumption that b is monotonic increasing, or $G_1(b(x)) = 0$, which means $b(x) < \underline{v}$. □

To sum up, $b' > 0$ and $b(x) \geq \underline{v}$ together imply a contradiction (the Revenue Equivalence Theorem must hold and cannot hold). Therefore we conclude that there is no symmetric, strictly increasing equilibrium bidding function Stage 1. However, the next proposition demonstrates that there can be symmetric equilibria so long as bidders bid below $\underline{v}$.

Proposition 3.3. *Any $b \leq \underline{v}$ is an equilibrium bidding function in Stage 1.*

Proof. If all bidders bid the same amount, they all go to Stage 2. The outcome of the auction is the same of the "button" auction of Milgrom and Weber (1982). A downward deviation in Stage 1 does not increase the probability of winning the auction, nor reduces the expected payment. An upward deviation only increases the expected payment, and does not increase the probability of winning. Therefore it does not pay to deviate. □

Remark 1 The symmetric equilibria in which all bidders in Stage 1 bid below the low valuation are fragile in two respects. First, although trembling hand perfectness is not formally defined in this game, there is a sense in which bidding the low valuation (say, zero) in the first stage does not satisfy the concept of trembling hand perfectness. This is especially relevant in the context of ComprasNet because "bidding zero" is not defined: ComprasNet auctions are *descending auctions* and there is no maximum bid permitted. Second, while these equilibria survive deviations from one bidder, they do not survive deviations from more than two bidders. That is, a coalition of bidders would have strong incentives to deviate by a little amount in Stage 1, and exclude all other bidders from competing in Stage 2.

Remark 2 One way of getting a symmetric increasing equilibrium in Stage 1 is to assume that updating the bids between the two stages is costly: in a limiting case where updating the bids is infinitely costly, no bidder will update and the game reduces to a first-price auction. However, it is not clear how to interpret such costs in our empirical setting. It is worth noting that other researchers (Decarolis (2010), Ye (2007)) have documented the use of auction mechanisms that also lack symmetric increasing equilibrium bidding strategies.

3.3 Background and Data

3.3.1 Background

Since its inception in 2001, ComprasNet had the long-term view of utilizing communication technology to improve the government's purchasing system. To that end, the government commissioned an online bidding platform. An alpha version of the platform was put in operation in 2001, but it needed fixes and was not ready for large-scale utilisation. As a result, few purchasing units used it; most purchasing units using reverse auctions chose the offline format. In 2005, as the online platform was ready for large-scale utilisation, purchasing units were mandated to use ComprasNet auctions. They were also required to used the online format, although the offline format remained available. Figure 3.2 illustrates the impacts of the 2005 regulation on the relative use of online and offline auctions.

As described in Section 3.2 offline ComprasNet auctions follow a two-stage design. Why such a special design? There are two answers to this question. First, this design was borrowed from the experiences of one particular autonomous government agency which used this design since the 1990s.[2]. Second, government officials were concerned with the possibility of collusion in offline auctions. The fact that bidders physically gather in a room increases the chances of communication among them. The first-round, they argue, was meant to mitigate the risk of collusion by giving bidders an incentive to deviate from a collusive agreement.

[2]The *Agência Nacional de Telecomunicações*, or ANATEL, had autonomy to use its own procurement methods.

Figure 3.2 Online and Offline Lots in ComprasNet: Jun/2004-Jun/2006

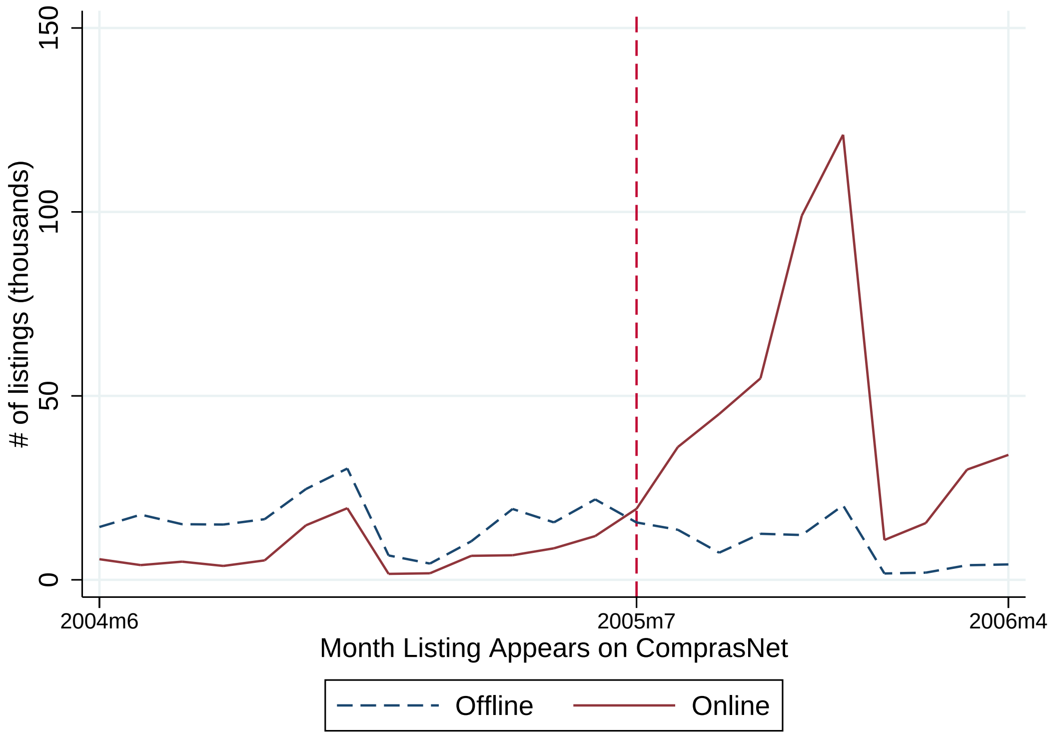

In online auctions, government officials perceived such risks to be reduced, and decided to abolish the qualifying rule.

The choice of auction format can therefore be linked to these concerns. If it is true that offline auctions may benefit certain groups of firms, it is also likely that corrupt procurement officials may choose this format to favour those firms. Moreover, there may be cases in which the choice of format is not related to collusion and/or corruption concerns, but rather to some unobservable characteristic (to the researcher) that make one auction format better suited. Because of these concerns, I take advantage from the fact that the timing of regulation requiring the use of the online format was driven by the technical conditions of the online bidding platform. While the offline format did not cease to be used, the regulation had a significant (negative) impact on the probability that the offline format is used, as evidenced by Figure 3.2.

3.3.2 Data

I focus on auctions between December 2004 and April 2006. The choice is based on the following four arguments. First, ending rules for online auctions have changed over time (see section 2.2). Restricting the sample avoids dealing with different flavours of online auctions. Second, unobserved factors may drive the choice for the offline format years after the 2005 legislation. Selecting auctions around the date of the policy change minimises such concerns. Third, other policy changes after 2006 were more likely to affect online auctions, thus introducing confounding factors in the comparison between the formats. Finally, because of data collection issues, there are many incomplete data points prior to December 2004, especially for offline auctions. I therefore avoid ending with data spanning different periods for the different auction formats by disregarding data from before December 2004.

Table 3.1 reports descriptive statistics for the sample used in this chapter. Overall, the table suggests that there are substantial differences between the two auction institutions. Normalised winning bids are 1.3 percentage points higher in offline auctions (p-value$<$ 0.001) and offline auctions attract less bidders than online auctions (3.7 vs 5.1, p-value$<$ 0.001). Finally, the third row shows what is known as the "money left on the table"–the percentage difference between the best (winning) bid and the runner up bid. This measure captures the dispersion in private information among bidders. Offline auctions display on average a much higher "money left on the table" than online auctions.

The last two rows of table 3.1 report statistics for two bid-level variables. The second-to-last row reports normalised opening bids, defined as bidder' stage 1 bids normalised by the lot's reserve price. The table shows that bidders tend to place more aggressive opening bids in offline auctions: the average opening bid in offline auctions is 14.3 percent below the reserve price, whereas in online auctions it is 1.4 percent *above* the reserve price. The last row shows by how much bidders update their bids between their opening and final bids. Over the course the auction, bidders revise their bids substantially more in the online than offline (9.8 percent and 4.3 percent).

3.4 Results

3.4.1 Bidding in Two Stages

The qualifying rule in offline auctions creates a trade-off between the probability that a bidder qualifies for Stage 2 and the price she will have to pay in case she wins the auction. In contrast, online auctions do not have this qualifying rule, and bidders have nothing to fear by submitting higher opening bids. If there are no costs in updating bids, one should expect opening bids in online auctions to be lower than opening bids in offline auctions. This also implies that the difference between a bidder's opening bid and her final bid should be greater in online auctions. To capture the extent to which the qualifying rule affects the opening bid, I run regressions of the form

$$y_{it} = \beta \text{Offline}_t + \mathbf{x_t}\gamma + \alpha_i + \varepsilon_{it}, \tag{3.2}$$

Table 3.1: Summary Statistics

	Online			Offline		
	Mean	Std. Dev	# of Obs.	Mean	Std. Dev	# of Obs.
Auction-level Variables						
Log Winbid	5.742	2.215	353,906	6.800	2.414	68,199
Log Reserve	6.065	2.209	388,243	7.152	2.376	74,769
WinBid/Reserve	0.704	0.238	353,906	0.717	0.244	68,199
$\frac{Ranked2 - Ranked1}{Ranked1}^{(a)}$	0.111	0.246	345,576	0.174	0.285	65,160
# of Bidders	5.079	4.545	388,243	3.684	3.247	74,769
Bid-level Variables						
Stage1/Reserve	1.014	0.467	1,806,030	0.867	0.390	251,150
$\frac{(Stage1 - Stage2)}{Stage1}^{(b)}$	0.098	0.176	1,804,433	0.043	0.095	163,270

Notes: table reports statistics for auctions held between December 2004 and April 2006. Auctions for which the reservation price is not observed are dropped from the sample. (a) If only one bidder enters the auction or if the second best bid is above the reserve price, the reserve price is used instead of the second best bid. (b) In offline auctions, bid update is only defined for bidders who qualified for stage 2.

where y_{it} is either the (normalised) first-stage bid of bidder i in auction t or her update between her opening and final bids; offline equals 1 if auction t is held offline, $\mathbf{x_t}$ is a vector of auction characteristics, α_i is a bidder-specific effect, and ε_{it} is an error term satisfying the usual assumptions. I experiment with treating α_i as a bidder fixed-effect or random-effect, as well as different auction controls.

The results are in Table 3.2. Columns (1) and (2) imply that the bidders submit more aggressive opening bids in offline auctions. On average, opening bids are 13 percent lower in offline auctions. After controlling for bidder fixed-effects, this figure remains at 9.7 percent. Columns (3) and (4) present results for the normalised difference of bidders' opening and final bids. The results suggest that bidders are less aggressive in offline auctions when updating their bids.

3.4.2 Revenues and Participation

The theoretical model presented in section 3.2 gives little predictions regarding differences in participation and revenues across the online and offline formats. As argued in the literature, it is likely offline auction have higher entry costs, reducing participation. In a independent private value framework, this should lead to lower revenues (higher procurement costs) on average. If bidders incur those entry costs after they learn their valuation, as in Samuelson (1985), then negative effect of reduced entry is mitigated by the fact that entry selects bidders with higher valuations.

The ideal experiment to assess the impact of auction formats on outcomes,

Table 3.2: Effect of Auction Format on Stage 1 bids

	(1) Opening Bid	(2) Opening Bid	(3) Update	(4) Update
Offline	-.131***	-.097***	-.075***	-.057***
	(.001)	(.002)	(.000)	(.001)
Log Reserve	-.030***	-.035***	.003***	.004***
	(.000)	(.000)	(.000)	(.000)
Bidder F.E.	No	Yes	No	Yes
Observations	1,987,468	1,336,452	2,169,316	1,470,803
R^2	.043	.025	.028	.017

Notes: opening bids are normalized by the reserve price. Update is the difference between a bidder's opening and final bid, normalized by the opening bid. Specifications with bidder fixed-effects include only bidders who bid in both online and offline auctions. All specifications include controls for year; quarter of year; state; and subcategory (4-digit code). Robust standard errors in parentheses.

would be to randomly assign formats to auctions, and then measure the differences in procurement costs and participation across the formats. I am unaware of any purchasing units following a similar strategy, although this may be possible.[3] To illustrate the empirical challenge we face, consider estimating the following equation:

$$y_t = \beta \text{Offline}_t + \mathbf{x}_t \gamma + \varepsilon_t, \tag{3.3}$$

where y_t is the outcome of interest (winning bid or number of bidders) in auction t and Offline_t indicates whether auction t is held offline. The vector $\mathbf{x}_t$ contains auction characteristics which determine outcomes (geographic location, lots' reserve price, etc). The term ε_t represents unobserved auction characteristics that determine auction outcomes. If the choice of auction format is correlated with any of such characteristics, $Cov(\varepsilon_t, \text{Offline}_t) \neq 0$, making the OLS estimator for β inconsistent.

To overcome these concerns, I take advantage of the policy change occurring in July 2005 that required purchasing units to use the online format. As argued in section 3.3, the timing of the change was driven by the suitability of the online platform. I therefore take the variation in the format choice induced by the policy as exogenous. In other words, I use a dummy for the period starting in July 2005 as an instrument for the choice of auction format.

[3]It is important to note that even such an experiment would not enable us to disentangle the effects of participation costs and differences in auction rules, as these two features are perfectly confounded with the auction format. To isolate the effect of the auction rules, the experiment would announce an auction, and ask interested bidders to submit two bids–one for each auction format. Then, randomize the auction format and consider only the appropriate bids.

The validity of this strategy relies on the assumption that there are no confounding factors over time that affect procurement costs and participation. The fact that ComprasNet has expanded over time might make it more attractive for suppliers. As a result, there may be an increase in participation over time that is not directly related to the choice of auction format. I use time trends to alleviate these concerns.

Table 3.3: Results

	Offline	Norm. Winning Bid		# of Bidders	
		OLS	2SLS	OLS	2SLS
AfterJuly05	-.391***				
	(.002)				
Offline		.022***	.082***	-2.351***	-4.907***
		(.001)	(.003)	(.016)	(.050)
N	463,012	422,105	422,105	463,012	463,012
R^2	.285	.082	.075	.305	.268

Notes: the first column shows first-stage regression results. The other columns show OLS and 2SLS results for each dependent variable. All regressions include controls for the log of reserve price, a quadratic time trend, and dummies for 4-digit codes and states.

Table 3.3 presents the results. The first column reports the results of a linear probability model for the choice of auction format. This is the first-stage regression of our IV strategy. As expected, the probability that the offline format is chosen drops significantly (39 percentage points) after July 2005. The third column in implies that the offline format increases procurement costs by 8.2 percent. The OLS results displayed in the second column however put this figure at 2.2 percent. One interpretation for this result is that purchasing units base their choice of auction format on unobservable factors (to the researcher) that decrease procurement costs. The fourth and fifth columns report the results for the number of bidders entering the auctions. As expected, offline auctions attract fewer bidders than the online format. The IV estimates imply that offline auctions have 5 fewer bidders than online auctions.

3.5 Concluding remarks

In this chapter, I analysed an offline version of ComprasNet auctions. The offline format follows a two-stage design. In the first stage, bidders submit sealed bids. A qualifying rule selects the best bids from Stage 1. In Stage 2, qualified bidders play an open-cry auction. I find differences in bid behaviour vis-à-vis online auctions, which do not have this qualifying rule. As expected, opening bids are more aggressive in offline auctions. However, final bids are less aggressive. Taking advantage of a change in regulation that required purchasing units to use online auctions, I find that, on average, offline auctions result in prices 8 percent higher than online auctions and that offline auctions attract, on average 5 fewer bidders than online auctions. The effect on the number of participants is consistent with the conjecture that offline auctions have higher participation costs.

The analysis in this chapter can be improved and extended in a number of ways. First, the analysis of section 3.4.2 could be improved by using cross-section variation in the identification of effects of auction format. For example, if the pace of adoption of the online format varied across cities or Public Bodies, and if that variation is due to factors not related to auction outcomes directly, then one could use this extra variation for estimating the desired effects. Useful variation however depends on institutional detail, and may not be availability in the present setting.

Second, the model presented in section 3.2 lacks an increasing equilibrium in Stage 1 bidding strategies. The trade-off created by the qualifying rule however, suggest that bidders should weight the probability that they will not qualify for Stage 2. Perhaps more importantly, the data suggests that Stage 1 bids are informative about Stage 2 bids. The existence of increasing equilibria is not desirable from a theoretical perspective, but also important for structural analysis of auction games. By appropriately modifying the theoretical model, future research will be able to carry a structural estimation of the offline auction game, thus recovering the primitives of the model.

There are other avenues that future research can explore. The distinction between online and offline, as well as the role of geography, is underexplored in the current version of this chapter. Economists have debated about the impacts of electronic commerce. It is commonly argued that web-based technologies reduce trade frictions (Brown and Goolsbee (2002)) and product differentiation through geographic location (Ellison and Ellison (2005)), leading to a more competitive environment. For example, Brown and Goolsbee (2002) find that price comparison websites helped to reduce insurance prices and price dispersion, and conclude that the internet had significant impacts on market competition in the life insurance industry. On the other hand, Hortaçsu et al. (2009) show that while eBay increased the geographical distance of transactions, it has not eliminated the "home bias". Thus, local firms (sellers) still enjoy some market power.

In ComprasNet, geographic location is arguably the only source of differentiation across bidders, since competition is strictly on prices. While bidders closely located to buyers will still enjoy costs advantages on online auctions, such advantages are certainly greater in offline auctions. Future research could investigate whether price dispersion across different geographic locations has fallen with the introduction of online auctions, contributing to the debate about the impact of information technologies on trade.

Chapter 4

Set Asides

Abstract

In this chapter, we analyse the effects of a bid preference programme targeting small and micro enterprises (SMEs) in ComprasNet auctions. The programme consists of *setting aside* eligible lots for favoured firms by restricting the participation by non-favoured firms. We first provide reduced-form evidence on the effects of the programme, taking advantage of the criteria used to restrict participation. These criteria are based on lots' reserve prices and provide discontinuities in the probability of treatment, allowing us to make use of them as a source of exogenous variation to identify the effects of the programme. We find that restricting participation of large firms has little effect on prices, while it increases participation of small firms. This finding is consistent with a model of bidder asymmetry and costly participation. In such a model, restricting participation by large, strong bidders increases the incentives of small, weak bidders to participate thus mitigating the adverse effects of the restriction on prices. We then set up a structural model to estimate entry costs and simulate the effects of using different criteria to set aside lots. We interpret entry costs in the context of ComprasNet as red-tape costs.

4.1 Introduction

Bid preference programmes have attracted economists' attention for at least three reasons. First, such programmes are ubiquitous in public procurement, and as such they potentially have large distributional effects. According to Krasnokutskaya and Seim (2011) 20 percent of US federal government procurement dollars in 2006 went to favoured firms. Second, preference programmes, and affirmative actions more generally, provide a good testing ground for economic theory. The theory predicts that both favoured and unfavored agents' incentives change. With the appropriate data, researchers can test the extent to which those incentives change, and the implications of those changes. Finally, affirmative action is a controversial political issue due to its distributional impacts, and unknown overall effects. Pinning down the costs and benefits of such policies is important to guide public debate and inform policy making.

In procurement auctions, bid preference programmes usually take one of two forms. The most common approach is to give targeted firms a bid subsidy. For example, a subsidy of 5 percent equates a $105 bid from a favoured bidder to a $100 bid from an unfavored bidder for purposes of determining the wining bid. The winner, however, is paid the amount of her bid. The other alternative is to set aside some lots by restricting participation to favoured firms. The effects of both bid subsidies and set-asides depend on the auction rules, the extent to which bidders are asymmetric and auction participation costs.

At first glance, both policies would appear to increase procurement costs as they tend to select weaker bidders. But handicapping a group of bidders may be optimal for the auctioneer. For example, in a first-price auction with asymmetric bidders, McAfee and McMillan (1989) show that there is an optimal bid subsidy to weak bidders that make unfavored (strong) bidders bid more aggressively, minimizing procurement costs to the auctioneer. In a set-aside, restricting participation of strong bidders gives weak bidders extra incentives to participate. Procurement costs may therefore decrease if the absence of the strong bidders is compensated by an increase in the number of favoured bidders (Athey et al. (2011a)).

This paper analyses the effects of a set aside programme targeting small and micro enterprises (SMEs) in Brazil.[1] We use data from ComprasNet, the online bidding platform used by the Brazilian federal government for procurement of various goods and services. In 2010, 50% of all procurement for the federal government was conducted through ComprasNet, totalling R$ 27 billion, or 0.7% of Brazil's GDP. Since the start of the set-aside programme in 2007, the government restricted participation to SMEs in 17% of the lots procured in ComprasNet, cor-

[1]We deliberately use the acronym SME for small and *micro* enterprises, instead of the usual "small and *medium* enterprises".

responding to 3% of procurement dollars.

We first provide reduced-form evidence on the effects of the programme. We find that restricting participation of large firms has little effect on prices, while it increases participation of small firms. This finding is consistent with a model of bidder asymmetry and costly participation. In such a model, restricting participation of large, strong bidders increases the incentives of small, weak bidders to participate thus mitigating the adverse effects of the restriction on prices. To estimate entry costs, we then set up a structural model of open auctions with asymmetric bidders and costly participation. The models' estimates also allows us to simulate the effects of using different criteria to set aside lots.

Interpretation of entry costs vary across empirical works on endogenous participation in auctions. In the context of ComprasNet, we argue that the opportunity cost of participating in the auction and, more importantly, red-tape costs are the main components of entry costs. Despite the government's efforts to reduce red-tape in public procurement, interested bidders still have to fill forms, provide documentation and fulfil a number of formal procedures. In other empirical settings, such features may be negligible when compared to other components of entry costs. For example, in timber auctions bidders must perform (costly) cruises to assess the value of a tract (see Athey et al., 2011b), and in road construction auctions, bid preparation is a costly activity that requires hours of work from skilled personnel (see Krasnokutskaya and Seim, 2011).[2] In a typical auction in ComprasNet, bidders are retailers or wholesalers who have good knowledge of their private costs before entering the auction. Bidders are unlikely to spend large amounts of time reading lots' descriptions, as these are short due to the standardisation of products being traded. We therefore interpret the time and effort that bidder put in preparing their bids in as a cost imposed by the formal requirements technological limitations of ComprasNet.

The data we use has at least two features that are distinct from previously used data. First, we observe auctions both before and after the programme was introduced. Pre-intervention data helps in the identification of the programme impact on auction outcomes. Secondly, the data spans various industries with different pools of potential bidders. Previous studies have focused almost exclusively on road construction (Krasnokutskaya and Seim (2011), Marion (2007)) and timber harvest (Athey et al. (2011a), Brannman and Froeb (2000)). Our setting allows to see if previous findings generalize and improves our understanding about such policies.

Our work contributes to the empirical literature on endogenous participation in auctions, e.g. Li and Zheng (2009), Krasnokutskaya and Seim (2011), Athey

[2]Bid preparation costs are also relevant in other settings. Bajari and Hortaçsu (2003) argue that in eBay auctions for collectible coins bidders must spend time collecting book values and other information on the lots being sold.

et al. (2011b), Groeger (2011). The fundamental methodological problem with estimating models of market entry is the existence of multiple equilibria. Recent works typically proceed by making some kind of equilibrium selection assumption.[3] For example, Athey et al. (2011a) consider an equilibrium in which large firms participate with probability one, and small firms use a mixed strategy. We build on the work of Bajari et al. (2010) to deal with multiple equilibria in the participation stage. Differently from the approaches currently used in the literature, Bajari et al. (2010) propose an equilibrium selection mechanism that can be estimated together with the model primitives. The main advantage of this idea is that it allow us to estimate games with multiple equilibria without imposing restrictive assumptions on how equilibrium actions are played.

In our structural estimation, we also face challenges due to the open nature of the auctions we analyse. Point identification of open auction models requires two restrictive assumptions, namely that bidders' valuations are independent private costs and that bidders play a button auction (Athey and Haile (2002)). We take an alternative approach and consider only information on bidders' entry decisions to estimate a reduced-form specification for bidders' utilities, as in Bresnahan and Reiss (1991).

Empirically, few studies have addressed the effects of set aside programmes. To the best of our knowledge, the only two studies using data from set aside auctions are those of Brannman and Froeb (2000) and Athey et al. (2011a).[4] Both studies analyse the set aside programme held by US Forest Service timber auctions – though using different datasets. According to the programme, in a fraction of the auctions only small mills or loggers are allowed to participate. Without considering endogenous entry, Brannman and Froeb (2000) find that eliminating the set aside programme would increase government revenues by 15%. Athey et al. (2011a) find that the set aside programme induces losses both in terms of revenue (5%) and efficiency (17%). However, by assuming endogenous entry Athey et al. (2011a) show that the losses are mitigated in set aside auctions due to the entry of small firms – without this effect the losses would be 30% in terms of revenue and 28% in terms of efficiency. Athey et al. (2011a) also use the model to estimate the effect of changing the set aside programme to a bid subsidy programme and find that with a 6% bid subsidy small firms would win the same proportion of auctions, the price would be 4% higher, and the efficiency would increase by 2%.[5]

[3] The early literature (Bresnahan and Reiss (1991, 1990)) placed restriction on the players' payoffs and assumed away mixed strategy equilibria. The set of pure strategy equilibria gives a unique mapping to the the total number of entrants.

[4] Ayres and Cramton (1996) provide a case study of bid subsidies in the US Federal Communications Commission (FCC) auctions.

[5] There are other two important works analysing the effects of bid subsidy programmes. Marion (2007) studies the effect of bid subsidies in California highway procurement auction. He finds that bid subsidy programme decrease the participation of big firms, and make them bid

The remainder of this paper is organized as follows. Section 4.2 discusses the institutional setting of ComprasNet auctions. Section 4.3 presents the data and provides some descriptive analysis. In Section 4.4 we provide reduced-form evidence of the effects of the set-aside programme on auction outcomes. In Section 4.5 we present the structural model we take to the data. Section 4.6 discusses the estimation procedure.

4.2 Background: ComprasNet Auctions and the Set-Aside programme

ComprasNet Auctions The Brazilian public administration has used reverse auctions as a procurement method for various types of off-the-shelf goods since 2001. Such auctions are commonly referred to as *ComprasNet* auctions, after the government-run internet portal that hosts the online bidding platform and listings for public procurement auctions.[6] Listings are posted by purchasing units (PUs) of federal public bodies (PBs).[7] In our sample, we observe 203 PBs with a total of 1,730 PUs located across 318 municipalities.[8] As of 2005, it is mandatory for PUs to use ComprasNet auctions to procure off-the-shelf goods.[9] In 2010, 50% of all procurement for the federal government was conducted through ComprasNet, totalling R$ 27 billion, or 0.7% of Brazil's GDP. In short, these auctions represent a large share of federal tenders and a substantial amount is contracted every year through them.

A ComprasNet auction starts with a PU defining lots it needs to purchase. Typically, several lots are procured at the same letting session; such lots are said to form a *batch*.[10] Upon defining lots' characteristics, the PU conducts market research, whereby at least three different firms give quotes for each lot. A reserve price for each lot is calculated as the average of these quotes, and it is meant to capture the retail price of the lot. The PU then advertises the tender at least 8

more aggressively. The overall effect is an increase the price by 3.8%. Krasnokutskaya and Seim (2011), also working with California highway procurement auction, conclude that bid subsidies increases the procurement costs by less than 1%.

[6]www.comprasnet.gov.br.

[7]In the Brazilian federal administration, a purchasing unit is called a *Unidade de Adminstração de Serviços Gerais* (UASG).

[8]The sample used is described in detail in Section 4.3.

[9]ComprasNet auctions cannot be used for procuring one-of-a-kind goods, engineering works or any project in which bids have multiple dimensions (e.g., technique and price). For each of these cases, different procurement methods are available.

[10]We emphasize such feature here because it will matter when we explain the set-aside programme and participation costs. Batches are a common feature of many auction settings; for examples relating to different industries, see Hendricks et al. (1987) for oil leases, Groeger (2011) for road construction works and Ashenfelter and Graddy (2003) for arts and wine auctions. In the Comprasnet auctions we analyse, the timing in which lots are open for bidding typically overlaps. We therefore assume that lots in a batch are auctioned off simultaneously.

working days before the letting session and publishes a tender document. The tender document contains a detailed description of each lot, including reserve prices, and is free to download from ComprasNet.

We focus on online ComprasNet auctions, the predominant procurement method since 2005.[11] Bidders must submit binding sealed bids before a deadline specified in the tender document to qualify to participate in the auction. When the auction starts, the low bid is announced and bidders engage in an open auction. Throughout the auction bidders only know what the current low bid is; they don't learn how many bidders are bidding, let alone their identities. The winner is the firm making the lowest bid.

Set Aside Programme In October 2007, the federal government introduced a bid preference programme for small and micro enterprises (SME). The programme consists of *setting aside* eligible lots for SMEs.[12] Due to ambiguous wording in the regulation implementing the set-aside programme, there are two competing eligibility criteria:

(C1) the lot's reserve price must be below R$ 80,000;

(C2) the batch's total reserve price (i.e., the sum of the reserve prices for all lots in the batch) must be below R$ 80,000.

Criterion C2 is therefore more stringent than C1. Procurement officials have considerable discretion in carrying out these auctions. Besides deciding which eligibility criterion to use, they may opt not to apply restrictions to eligible lots. In fact, the set aside programme explicitly instructs officials not to apply restrictions when they are expected to result in substantially higher procurement costs, or in fewer than three bidders participating in the auction.[13]

4.3 Data

We analyse data of auctions held between October 2007 and December 2010. We use two sources of data. First, we use publicly available data automatically re-

[11]Before 2005, it was common for PUs to use an offline version of the ComprasNet auction analysed in this paper.

[12]See *Lei Complementar* 123/2006 (14 December 2006) and *Decreto* 6204 (05 September 2007). Among other requirements, to qualify as a micro (small) enterprise, a firm must not have gross revenues larger than R$ 240,000 (2,400,000). This threshold was changed in 2009 to R$ 360,000 (3,600,000).

[13]Procurement officers seem to have different views on when to set aside eligible lots. Approximately 25% of the public organisations never restrict participation, while another 25% only restrict when criterion C2 is met. Among the organisations which restrict participation in eligible lots, there is considerable variation in the proportion of eligible lots which end up being restricted. In other words, public bodies which are similar in terms of observables seem have different policies. In future versions of this paper, we plan to use this cross-section variation to help identifying the impacts of the set aside programme.

corded by the ComprasNet platform. For each lot, ComprasNet collects the following information: which firms bid and all bids placed by each firm; whether the firm is registered as SME with tax authorities; the size of the lot; the purchasing unit running the auction. Second, we complement these with internal data from the Ministry of Planning, Budget and Management. This data contain information on lots, bidders, and purchasing units. On lots, there is a paragraph-long description of the item along with classification codes following the United States' Federal Supply Codes (FSC) for materials and the U.N. Central Product Classification for services. These classification schemes define product categories by 2-digit codes, and sub-categories by 4-digit. There are also finer 6-digit codes which are created by purchasing units if needed. On bidders, we have their geographical locations and industry defined by the International Standard Industrial Classification (ISIC). Finally, the internal data contains the geographical location of purchasing units, as well as their place within the government's organizational structure.

We restrict the sample to seven commonly purchased products, namely ink cartridges, A4 paper, printers, air conditioning, coffee powder, computers, and bottled cooking gas. We select these products for two reasons. First, these are commonly purchased products with reasonable variation in treatment. Second, by focusing on a narrower set of products we can extract lots' characteristics from text data using a piece of computer software. We can ensure, for example, that all lots of the same product category are measured in the same units.[14] The appendix gives further detail of the sample composition.

Table 4.1 reports sample descriptive statistics for restricted and unrestricted lots. Unrestricted lots have significantly higher reservation prices. On a typical unrestricted lot there are 2 large bidders participating, and 15 small bidders. This represents a proportion of small bidders of 88 percent in a typical unrestricted lot. Small firms however win 82 percent of the lots, suggesting that there are asymmetries between small and large bidders.

Figure 4.1 illustrates the effects of criteria C1 and C2 on the probability that a lot is restricted. This is the variation in the data we use to identify the effects of the set aside programme on auction outcomes. The figure plots the probability that a lot is restricted against the lot's reservation price and weather or not the lot belongs in a batch whose overall reservation price is below R$80,000. The figure shows that 14 percent of lots meeting criterion C1 but not C2 are treated, and that this proportion increases to 40 percent for lots meeting both criteria.

[14] We plan to expand the set of products analysed. This is a time-consuming task though, as it requires searching for product-specific regular expressions in the text data.

Table 4.1: Sample Descriptive Statistics

	Unrestricted				Restricted			
	Mean	Std. Dev.	p25	p75	Mean	Std. Dev.	p25	p75
Price (000's R$)	41.39	1000.26	.67	10.00	4.64	9.04	.36	4.43
Reserve (000's R$)	62.68	1447.42	1.30	17.00	6.98	12.06	.70	7.20
# Big Firms	1.97	1.88	1.00	3.00	.00	.00	.00	.00
# Small Firms	15.43	9.81	8.00	21.00	14.66	9.41	8.00	19.00
Prop. Small Firms	.88	.13	.83	.96	1.00	.00	1.00	1.00
Small Firm Wins	.82	.38	1.00	1.00	1.00	.00	1.00	1.00
Observations	57623				16639			
Num of Batches	10274				3588			

Notes: The table reports the mean, standard deviation, 25th and 75th percentiles for ComprasNet online auctions held by federal purchasing units between October 2007 and December 2010. As explained in the text, the sample is restricted to lots of ink cartridges, A4 paper, printers, computers, coffee powder, and bottled cooking gas (GLP).

Figure 4.1 Discontinuity in the Probability of Treatment

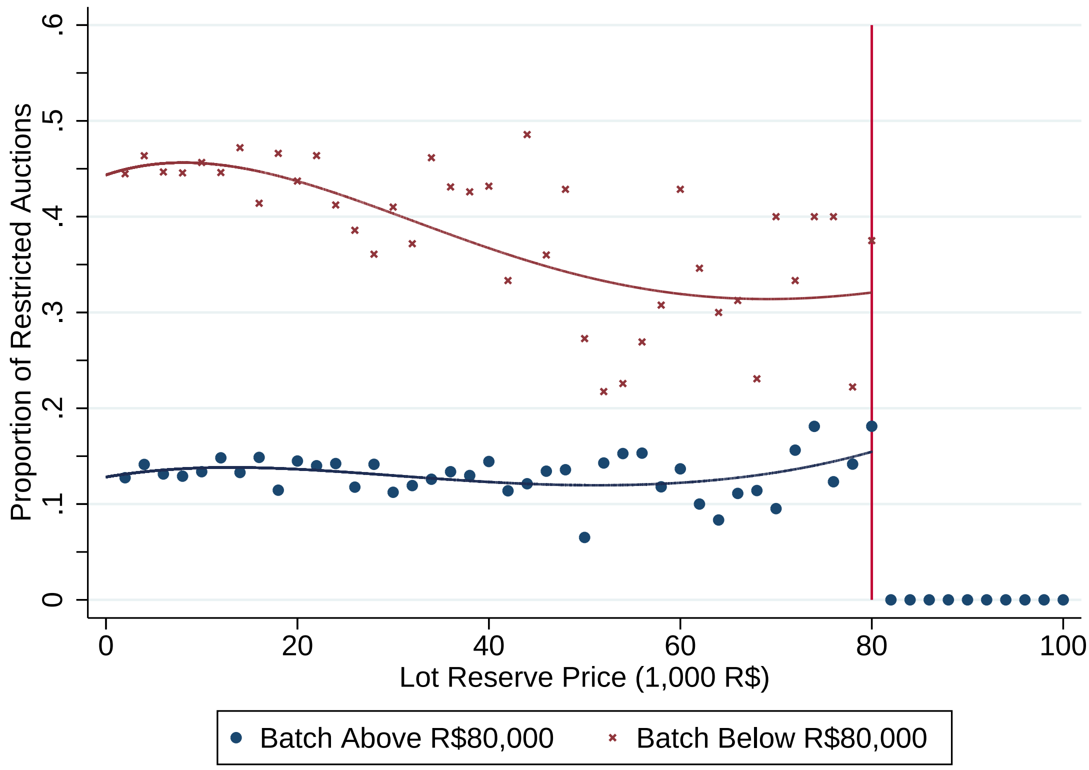

Notes: Dots represent local averages in R$ 2,000 bins. Fitted lines are from a probit regression of the treatment indicator on a 4th-order polynomial of the running variable (lot's reserve price), B_i and interactions of the polynomial and B_i. The regression uses only lots with reserve price below R$ 80,000.

4.4 Reduced Form Analysis

The ideal experiment The theory gives ambiguous predictions regarding the effects of a set-aside programme on both revenues and participation. The objective of this section is to estimate the impact of the ComprasNet set-aside programme on auction outcomes imposing minimal assumptions to the data. The ideal experiment in such setting would be to randomly restrict participation in some auctions, and then measure the differences in procurement costs and participation across restricted and unrestricted auctions. However feasible, we are unaware of such an experiment being carried out. From the discussion in Section 4.2, it is unlikely that our sample of restricted auctions has been randomly selected from the pool of eligible auctions. In particular, procurement officers are instructed to base their decisions on expected auction outcomes. While we, as econometricians, can observe some variables influencing officers' decisions (e.g., geographical region), we cannot be certain to observe all such relevant variables. A simple comparison between restricted and unrestricted auctions is therefore likely to give inconsistent estimates of the desired effects.

To illustrate the empirical challenge we face, consider estimating the following equation:

$$y_i = \beta SAS_i + \mathbf{X}_i\gamma + \lambda_t + \varepsilon_i, \tag{4.1}$$

where y_i is the outcome of interest (winning bid or number of bidders) in auction i and SAS_i indicates whether auction i is restricted to SMEs. The vector $\mathbf{X}_i$ contains auction characteristics which determine outcomes. For example, it may contain the geographic location of delivery, and the lots' reserve price. Note that $\mathbf{X}_i$ may contain variables used by procurement officers when deciding whether or not to restrict participation in auction i. The term ε_i represents unobserved auction characteristics that determine auction outcomes. As argued above, ε_i is likely to contain relevant information for the decision-making of whether to restrict participation. Hence, $Cov(\varepsilon_i, SAS_i) \neq 0$, making the OLS estimator for β inconsistent.

To overcome these concerns, we explore the discontinuity created by criteria C1 and C2 to identify β in equation (4.1). Because not all lots meeting either C1 or C2 are restricted, we end up with a fuzzy discontinuity around the cutoff values used by each criterion. To be more precise, let $L_i = 1$ if the reserve price of lot i is below R\$ 80,000 and 0 otherwise; and $B_i = 1$ if the reserve price of the batch in which lot i belongs is below R$ 80,000, and 0 otherwise. Now, let

$$Pr[SAS_i = 1|r_i] = \begin{cases} 0, & \text{if } L_i = 0; \\ g_2(r_i), & \text{if } L_i = 1, B_i = 0; \\ g_3(r_i), & \text{if } B_i = 1. \end{cases} \tag{4.2}$$

We assume that $g_2(r_i)$ and $g_3(r_i)$ can be described as p-order polynomials on r_i, so that

$$\begin{aligned} E[SAS_i = 1|r_i] &= \left[\delta_{20} + \delta_{21} r_i + \delta_{22} r_i^2 + ... + \delta_{2p} r_i^p\right] [(L_i)(1 - B_i)] + \\ &\quad + \left[\delta_{30} + \delta_{31} r_i + \delta_{32} r_i^2 + ... + \delta_{3p} r_i^p\right] (B_i) \\ &= \left(\delta_{21} r_i L_i + \delta_{22} r_i^2 L_i + ... + \delta_{2p} r_i^p L_i\right) \\ &\quad + \left(\alpha_1 r_i B_i + \alpha_2 r_i^2 B_i + ... + \alpha_p r_i^p B_i\right) \\ &\quad + \delta_{20} L_i + \alpha_0 B_i, \end{aligned} \tag{4.3}$$

where $\alpha_j = \delta_{3j} - \delta_{2j}$, $j = 0, ..., p$. The second equality uses the fact that $L_i B_i = B_i$. This last equation tells us that L_i and B_i, as well as the interaction terms $r_i L_i, ..., r_i^p L_i$, $r_i B_i, ..., r_i^p B_i$ can be used as instruments for SAS_i, since none of these terms belongs in (4.1).

This strategy assumes that the reserve prices are not manipulated in order to cross the eligibility threshold in any way. The fact that procurement officials are not mandated to restrict participation in auctions mitigates this risk. Moreover, using pre-intervention data we are able to test this assumption. We find no evidence that reserve prices have have a disproportionate chance to fall on either side of the eligibility threshold during the treatment period vis-a-vis the pre-treatment period.

4.4.1 Results

Table 4.2 presents OLS and IV results from estimating equation 4.1. We look at the effects of the set aside programme on the winning bid (conditional there being a winner) and the number of bidders. The specifications in columns (1) and (2) include dummies for each product category interacted with year dummies, reservation price and size of the lot. We also include state and quarter dummies. Column (3) restricts the sample to lots with reserve price around the discontinuity of R$80,000.

The results in the first row suggest that restricting participation to SMEs has very little effect, on average, on the winning bid (conditional on the lot being awarded). Both the OLS and two-stage least squares estimates are negative, indicating that winning bids actually *fall*, but these results are not statistically signi-

Table 4.2: Results: IV estimates

	OLS	2SLS	
		Full Sample	Discontinuity
Dep. Var.	(1)	(2)	(3)
Log of Price	-.010	-.014	.019
	(.005)	(.018)	(.026)
Number of bidders	-1.070***	2.436***	.019
	(.070)	(.231)	(.332)
Instruments			
B. Interaction (Quad)		yes	
N	74,009	74,009	3,954

Notes: Columns (1) and (2) include as controls dummies for each product category interacted with year dummies, reservation price and size of the lot. We also include state and quarter dummies. Column (2) uses B_i, $B_i r_i$ and $B_i r_i^2$ as instruments for SAS_i. Column (3) restricts the sample to lots with reserve price between R$69,000 and R$91,000. *** indicates significance at the 0.01% level.

ficant at the 5 percent level. The estimate in the discontinuity sample is positive, but again not significant at the usual levels.

The results on the number of entrants, presented in the second row, are ambiguous. The OLS estimate is negative, indicating that the number of participants decrease by one bidder when participation is restricted. The two-stage least squares estimate, on the other hand, suggest that the number of entrants *increase* by 2 bidders. To reconcile these results, recall that the IV estimator recovers the local-average treatment effect, that is, the effect of the set aside programme on the sub-population of compliers. In the present setting, the compliers are those lots which are restricted if and only if (i) the reserve price is below R$80,000 and (ii) the restriction is unlikely to have adverse effects on auction outcomes (i.e., winning bid and participation).

Overall, the evidence suggests that the SAS programme had positive effects on entry, but no effects on prices. This is consistent with the following story. The SAS programme induces the entry of SMEs, which more than compensates the reduction in the number of large bidders. The increase in the number of SMEs however does not translate into lower prices because, possibly, SMEs have higher costs in comparison to large bidders. This will motivate the structural model that we use to estimate the primitives of the auction game.

4.4.2 Testing Instrument Validity

Our identification strategy uses the discontinuity created by eligibility rules. This assumes reserve prices are not manipulated in order to cross the eligibility threshold in any way. Although reserve prices are the result of market research and in prin-

Table 4.3: TSLS and LIML estimates

Dependent Variable	Method	(1)	(2)	(3)	(4)	(5)	(6)
Log of Price	2SLS	.001	-.023	-.018	-.029	.006	-.014
		(.018)	(.017)	(.017)	(.017)	(.018)	(.018)
	LIML	.001	-.023	-.018	-.031	.006	-.014
		(.018)	(.017)	(.017)	(.018)	(.018)	(.018)
Number of bidders	2SLS	2.363***	2.678***	2.601***	2.714***	2.328***	2.436***
		(.232)	(.227)	(.227)	(.226)	(.232)	(.231)
	LIML	2.363***	2.697***	2.657***	2.814***	2.331***	2.455***
		(.232)	(.228)	(.228)	(.229)	(.232)	(.232)
Instruments							
B.		yes					
L. & B.			yes				
L. & B. Interaction				yes			
L. & B. Interaction (Quad)					yes		
B. Interaction						yes	
B. Interaction (Quad)							yes
N		74009	74009	74009	74009	74009	74009
Partial F stat		7186	3799	1907	1291	3607	2429

Notes: The table shows two-stage least squares (2SLS) and limited-information maximum likelihood (LIML) results when instruments are varied. Column (1) uses only an indicator for batch below R$80,000 ($B_i$). Column (2) adds an indicator for lot below R$80,000 ($L_i$). The other columns interact L_i and B_i with polynomials on the reserve price. L_i alone turns out to be a weak instrument, due to low variation in the data. We therefore omit the specification using only L_i as the instrument. *** indicates significance at the 0.01% level.

ciple are not under the control of the PU, there are a number of ways it could be manipulated by procurement officials. For example, the size of the lot could be changed. The fact that procurement officials are not mandated to restrict participation in auctions should mitigate the risk of inflating reserve prices. Still, using pre-intervention data we are able to to provide evidence in support of this assumption.

As a test to validate our identifying assumption, we first restrict the sample to auctions with reserve prices around the discontinuity. We then test whether auctions after the policy change have increased chances of lying on either side of the cutoff. We experiment with various windows widths and various "placebo" cutoffs, to allow for the possibility that the distribution of reserve prices changes over time. That is, we search for changes at the cut-off of R$ 80,000 which are unusual vis-à-vis other cutoffs. For example, there might be bunching around multiples of R$ 10,000; if that is the case, then looking only at the R$ 80,000 would wrongly suggest that reserve prices are tweaked to meet (or unmeet) eligibility criteria.[15]

Formally, for various cutoffs c and window widths w, we restrict the sample to auctions with $r_i \in (c - w, c + w)$, and define $below_i = 1$ if $r_i \in (c - w, c]$, and zero otherwise. That is, $below_i$ is an indicator of whether the reserve price is auction i is just below the cut-off c, where "just below" is defined by window width w. We then run the regression

$$below_i = \alpha + \beta AfterOct07_i + \varepsilon_i, \tag{4.4}$$

where $AfterOct07_i = 1$ if the auction is held after the SAS programme started. The parameter β gives the differential chance that auctions in the treatment period fall just below the cut-off.

We plot $\widehat{\beta}$ along with 95% confidence intervals in Figure 4.2 for various cutoffs and windows of R$500 and R$5,000. The top row of Figure 4.2 reveals that lots' reserve prices have a slightly higher change of falling just below the R$80,000 cutoff in the treatment period. This effect is not statistically significant at the 5% level when the window is R$500 and is just significant when we expand the window to R$5,000. More importantly, the effect at R$80,000 is not different (statistically nor visually) from the placebo cutoffs. The same can be said about the plots at the bottom of Figure 4.2, which use batches' reserve prices. As opposed to lots, batches have a decreased chance of laying at the bottom half of the window after October 2007. Again, this effect is not different from those at other placebo cut-offs. Overall, taking the pre-treatment period as a baseline, there is no evidence that reserve prices are manipulated around the cutoff of R$80,000 during the

[15] We could also perform the test suggested by McCrary (2008). This test should be included in future versions of this paper. McCrary (2008) however does not use pre-intervention in his test.

Figure 4.2 Change in Probability of Reserve Price

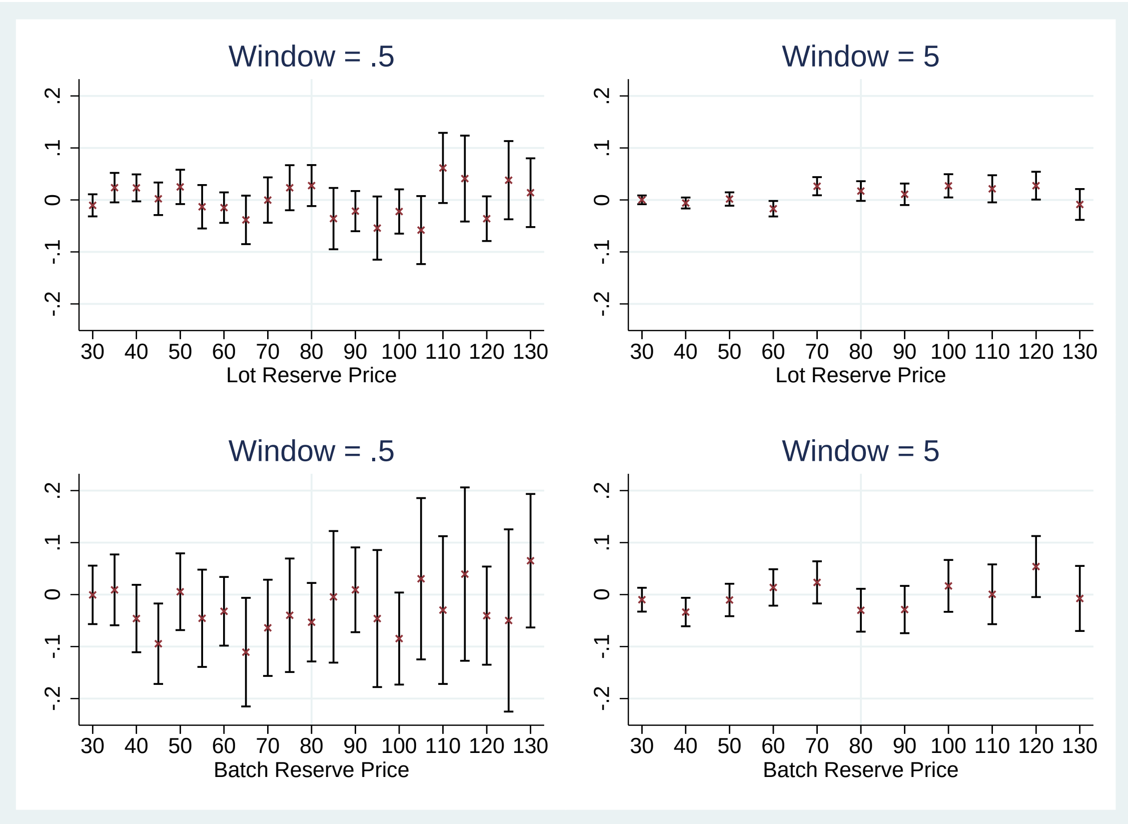

Notes: each panel shows point estimates and 95% confidence intervals of estimates of β in equation (4.4). The top left panel consider windows of R$500 on lots' reserve price. The bottom-right panel considers windows of R$5,000 on batches' overall reserve prices.

treatment period.

4.5 Structural Model

In this section we develop a structural model to study the effects of the set aside programme and alternative counterfactual policies on procurement costs. We first define an entry model with type-symmetric players that maps observed auction characteristics and entry costs on the distribution of the number of entrants. We then use this distribution to calculate expected procurement costs.

4.5.1 Entry Model

We focus on a model with two types of players, small and large. We assume that players are type-symmetric: all players from the same type draw their private information from the same distribution and use the same entry strategies. Symmetry is a standard assumption and avoids computational issues involved in the solution and the estimation of entry models with a large number of players. In what follows we specify the main assumptions behind the structural model.

Basic Assumptions

Players We consider a procurement auction with N^s small potential bidders and N^l large potential bidders. From this pool of potential bidders $n^s \leq N^s$ and $n^l \leq N^l$ denotes the number of small and large firms participating in the auction. Bidders are type-symmetric, meaning that the subscript j identify any player $i \in j, i \leq N^j, j \in \{s, l\}$.

Entry costs Bidders have type-specific entry costs $K^j + \epsilon^j$, where K^j is a type-specific parameter and ϵ^j is a type-specific shock. We assume that players observe both K^j and ϵ^j, whereas ϵ^j is not observed by the econometrician. We are therefore assuming that the entry stage is a complete information game, as in Bresnahan and Reiss (1991, 1990).

Payoffs Conditional on entry, bidders of type j have *ex-ante* expected utilities conditional on winning the auction given by $u^j\left(n^s, n^l, x\right)$ where x is a vector of auction characteristics. The vector of auction characteristics, x, is observed by all the bidders and by the econometrician.

Sequence of actions Formally, the game has two stages. In the first stage, players simultaneously choose either to enter ($a^j = 1$) or to stay out of the auction ($a^j = 0$). In the second stage, bidders learn about their completion costs and submit bids in the open auction. The procurement contract is awarded to the low bidder.

Solution of the entry model

To solve the model, we restrict attention to (subgame perfect) Nash Equilibria in pure strategies in the $\left(n^s, n^l\right)$ space. Solving for equilibria in non-degenerated mixed-strategies is not feasible in our setting due to the increased number of potential entrants.[16]In the next paragraphs we use the hypotheses outlined in the last section and this equilibrium concept to solve the model. The solution of this model maps the set of observables and parameters into the joint probability distribution of $\left(n^s, n^l\right)$.

First stage: pure strategy entry In the first stage bidders know the number of small and large potential bidders, the specification of the contract and entry costs. They estimate the expected profit conditional on winning the auction,

[16]One possible solution for the feasibility problem is to restrict attention to a set of bidders–say, regular bidders–instead of restricting attention to a set of equilibria as we currently do. We are working on such a solution as an alternative. In principle, one could accommodate in this framework other equilibrium concepts. In future versions of this paper, we plan to estimate a model of joint profit maximization as in Bajari et al. (2010).

$u^j\left(n^s, n^l, x\right)$, and choose simultaneously either $a^j = 1$ or $a^j = 0$. Formally, the expected utility from $a^j = 1$ or $a^j = 0$ is given by:

$$U^j\left(n^s, n^l, x, \epsilon^j, a^j; K^j\right) = \left(u^j\left(n^s, n^l, x\right) + K^j + \epsilon^j\right) \cdot I\left(a^j = 1\right), \tag{4.5}$$

where $I(.)$ is an indicator function that assumes 1 if the condition inside the parenthesis is satisfied.

Second stage: bidding In the second stage, entrants learn their completion costs and simultaneously submit their bids. To analyse bidding decisions we must specify the format of our auction and how completion costs are distributed across bidders. It turns out that point identification of the distribution of bidders' valuations in an open auction requires the independent private value paradigm as well as the assumption that bidders play a "button' auction (Athey and Haile (2002)). These are strong assumptions, and recent works have taken the approach of working with incomplete models of open auctions and estimate bounds on the distribution of valuations (Haile and Tamer (2003)) or seller's revenues (Aradillas-López et al. (2011)). We take a different approach.[17] As already hinted by our notation, we do not use bidding information to identify entry costs. We are thus agnostic about the auction format and the distribution of private values. We impose a parametric structure on $u^j\left(n^s, n^l, x\right)$, and assume it is a known function of the number of small and large participants and of auction characteristics. In practical terms, this means that, as in Bresnahan and Reiss (1991), the function $u^j\left(n^s, n^l, x\right)$ will be estimated only using information on the number of small and large participants as well as auction covariates.

Equilibria Using the set of hypothesis discussed above we find the subgame perfect entry strategies of small and large players. The following proposition establishes this result.

Proposition 4.1. *Suppose* $\left(N^s, N^l\right) > 0$. *For any* $\left\{K^j, x, \epsilon^j, u^j(.); j \in \{s, l\}\right\}$, *where* $u^j(.)$ *is any function defined in* $\mathbb{R}$, *the pair* $\left(n^s, n^l\right) \leq \left(N^s, N^l\right)$ *is a* Nash Equilibrium

[17] It is not clear that the bounds approach is feasible in our setting. Unlike our paper, the literature taking the bounds approach do not consider endogenous participation and entry costs. The challenge is to identify entry costs K^j with bounds on $u^j\left(n^s, n^l, x\right)$ in (4.5).

in Pure Strategies *if and only if it satisfies the following system of inequalities:*

$$
\begin{aligned}
\left\{u^s\left(n^s,n^l,x\right)-K^s-\epsilon^s\geq 0\right\}I\left(0<n^s\leq N^s\right)\\
\left\{u^s\left(n^s+1,n^l,x\right)-K^s-\epsilon^s< 0\right\}I\left(0\leq n^s< N^s\right)\\
\left\{u^l\left(n^s,n^l,x\right)-K^l-\epsilon^l\geq 0\right\}I\left(0<n^l\leq N^l\right)\\
\left\{u^l\left(n^s,n^l+1,x\right)-K^l-\epsilon^l< 0\right\}I\left(0\leq n^l< N^l\right)
\end{aligned}
$$

Proof. See Appendix. □

This proposition establishes the conditions under which any pair $\left(n^s,n^l\right)\leq\left(N^s,N^l\right)$ is a *Nash Equilibrium in Pure Strategies* for the entry game. It should be noted that (i) neither the existence nor the uniqueness of equilibrium in pure strategies is guaranteed and (ii) the inexistence and the multiplicity of equilibria implies that the mapping from the set of observables and parameters into the joint probability distribution of $\left(n^s,n^l\right)$ will be ill defined. This in turn impedes the estimation of the model, see Bajari et al. (2010) and Bresnahan and Reiss (1990).

Existence Bresnahan and Reiss (1990) solve the inexistence problem restricting the format of $u^j(.)$. In our setting it is easy to see that if $u^j(.)$ is non increasing in n^j then, for any $\left\{K^j,x,\epsilon^j,u^j(.);j\in\{s,l\}\right\}$ there will be at least one pair $0<\left(n^s,n^l\right)\leq\left(N^s,N^l\right)$ that solves the system of inequalities in proposition 4.1. Here we followed the same approach proposed in Bresnahan and Reiss (1990) and assumed that $u^j(.)$ is non increasing in n^j.

Multiplicity To address the multiplicity problem we construct an equilibrium selection mechanism using the idea proposed by Bajari et al. (2010). More specifically, for any $\left\{K^j,x,\epsilon^j,u^j(.);j\in\{s,l\}\right\}$, let $\mathcal{E}(K,x,\epsilon,u(.))$ be the set of all pairs $\left(n^s,n^l\right)$ solving the system in proposition 4.1 and $e(K,x,\epsilon,u(.))$ a particular element of $\mathcal{E}(K,x,\epsilon,u(.))$. Now define the following dummy variables:

$$
\begin{aligned}
y_1(e(K,x,\epsilon,u(.)))=1 \;\; if\; n^s>n^l;\; 0\; otherwise\\
y_2(e(K,x,\epsilon,u(.)))=1 \;\; if\; n^s<n^l;\; 0\; otherwise\\
y_3(e(K,x,\epsilon,u(.)))=1 \;\; if\; n^s=n^l;\; 0\; otherwise,
\end{aligned}
$$

which holds for every $e(K,x,\epsilon,u(.))\in\mathcal{E}(K,x,\epsilon,u(.))$ and every possible information set, $\left\{K^j,x,\epsilon^j,u^j(.);j\in\{s,l\}\right\}$. Now define the following cdf over the set of equilibria $\mathcal{E}(K, x, \epsilon, u(.))$:

$$\lambda\left[e\left(K, x, \epsilon, u\left(.\right)\right); \theta\right] = \frac{exp\left\{\sum_{p=1}^{3} y_p\left(e\left(K, x, \epsilon, u\left(.\right)\right)\right)\theta_p\right\}}{\sum_{\forall e'\left(K,x,\epsilon,u(.)\right)\in\mathcal{E}\left(K,x,\epsilon,u(.)\right)} exp\left\{\sum_{p=1}^{3} y_p\left(e'\left(K, x, \epsilon, u\left(.\right)\right)\right)\theta_p\right\}}, \tag{4.6}$$

where $\theta = \{\theta_p\}_{p=1}^{3}$ is a set of auxiliary parameters weighting the probability of any equilibrium in $\mathcal{E}(K, x, \epsilon, u(.))$ according to the equilibrium characteristics defined by the set of dummy variables, y.

Entry probabilities Based on the structure used to address multiplicity and inexistence of equilibrium in pure strategies we are finally able to define a probability distribution for (n^s, n^l) conditional on the set of parameters and observed auction characteristics. The following proposition brings this result.

Proposition 4.2. *Suppose $\epsilon = (\epsilon^s, \epsilon^l)$ is an iid vector with cdf $F_\epsilon(.,.)$. Then the distribution of (n^s, n^l) conditional on $\{N^j, K^j, x, \theta, u^j(.); j \in \{s, l\}\}$ is given by:*

$$\begin{aligned} &p\left(\left(n^s, n^l\right) | N, K, x, \theta, u\left(.\right)\right) = \\ &\int I\left(\left(n^s, n^l\right) | \epsilon, N, K, x, \theta, u\left(.\right)\right) \lambda\left[e\left(\epsilon, K, x, u\left(.\right)\right) = \left(n^s, n^l\right); \theta\right] dF_\epsilon\left(\epsilon\right), \end{aligned} \tag{4.7}$$

where $I((n^s, n^l) | \epsilon, N, K, x, \theta, u(.))$ is an indicator function that assumes 1 if (n^s, n^l) satisfies the system of inequalities in proposition 4.1.

Proof. See Appendix. □

Proposition 4.2 establishes a mapping from $\{N^j, K^j, x, \theta, u^j(.); j \in \{s, l\}\}$ into the distribution of the number of players. Intuitively, based on this map one can understand the effects of changes in auction characteristics, number of potential players, entry costs, etc. on participation decisions of small and large players. In the next subsection we use this distribution to define expected procurement costs.

4.5.2 Procurement Costs

Auction structure For a given number of bidders, $n = n^s + n^l$, participating in a given procurement auction write the order statistics of the completion costs as $V_{n:n}^{n^s,n^l} \leq V_{n-1:n}^{n^s,n^l} \leq ... \leq V_{1:n}^{n^s,n^l}$, and as $F_{k:n}^{n^s,n^l}(\cdot)$ the distribution of $V_{k:n}^{n^s,n^l}$. Notice firstly that because the order statistics depends not only on the total number of bidders in the uction, n, but also on the composition of the pool, n^s, n^l, we indexed our order statistics and the respective distributions on n^s, n^l.

Notice as well that we are not making any assumption on the distribution of valuations. As in Aradillas-López et al. (2011) we assume that in any auction with $n = n^s + n^l$ bidders the transaction price is the minimum between the reserve price and the second lowest cost of the bidders participating in that auction.

Entry We assume that, *ex-ante*, the government does not observe the number of bidders participating in the auction but only the distribution of $\left(n^s, n^l\right)$ conditional on the information set $\left\{N^j, K^j, x, \theta, u^j\left(.\right); j \in \{s, l\}\right\}$. All the elements in this information set are observed by the government. The distribution over $\left(n^s, n^l\right)$ was defined in proposition 4.2.

Reserve prices To calculate procurement expected costs we used the fact that in our auctions estimated reserve prices, r, are always equal to the cost estimated by the government engineers, π_0. The main advantage of this hypothesis is that it allows us to calculate expected procurement costs as a function of second highest distributions only, implying that expected procurement costs can be estimated directly from our data without the necessity of applying any procedure based on the estimation of bounds–see Haile and Tamer (2003) and Aradillas-López et al. (2011).

Procurement costs In the next proposition we define a closed form expression for *ex-ante* expected procurement costs.

Proposition 4.3. *Suppose that the following set of assumption holds:*

1. *Government observes the tuple* $\left\{N^j, K^j, x, \theta, u^j\left(.\right); j \in \{s, l\}\right\}$;
2. $\left(n^s, n^l\right)$ *is distributed according to the probability function in proposition (4.2);*
3. *The winning bid is the highest between the reserve price and the second lowest cost; and,*
4. *Reserve prices, r, are always equal to the cost estimated by the government engineers, π_0.*

Then ex-ante *expected procurement costs are given by:*

$$\pi\left(N, x; K, \theta, u\left(.\right)\right) =$$
$$\sum_{n^s, n^l} \left[\int_0^r v dF_{n-1:n}^{n^s, n^l}\left(v|x\right)\right] p\left(\left(n^s, n^l\right)|N, K, x, \theta, u\left(.\right)\right) - \pi_0 \tag{4.8}$$

where $F_{n-1:n}^{n^s, n^l}\left(v|x\right)$ is the distribution of the second lowest costs conditional on the vector of auction characteristics x.

Proof. See Appendix. □

4.6 Estimation Procedure

In this section we design an estimation procedure that allows us to estimate the entry probabilities in proposition 4.2 and subsequently the procurement costs, described in equation (4.8). The estimation procedure is divided in two blocks. In the first part we set up a minimum distance estimator that allows us to consistently estimate the parameters of the entry model, $\{K, \theta, u(.)\}$. With these elements we can define entry probabilities. In the second part we estimate the distributions of completion costs. With the distribution of completion costs and the entry probabilities we can finally calculate expected procurement costs.

4.6.1 Estimation of the Entry Model

To estimate the entry model we propose an Asymptotic Least Squares Estimator (hereafter ALSE) similar to that used by Pesendorfer and Schmidt-Dengler (2008) (see also Gourieroux and Monfort (1995)). Before describing the details of this procedure we invoke a new set assumptions that will be used in the empirical model.

Data Suppose that the econometrician observes a set of $\mathcal{T} = 1, 2, ..., T$ auctions. For this set of auctions she has $\{n_t^s, n_t^l, N_t^s, N_t^l, x_t; \forall t \in \mathcal{T}\}$.

Unobservables The distribution of the vector of shocks, $\epsilon = (\epsilon^s, \epsilon^l)$, is known with cdf denoted by $F_\epsilon(.,.)$.

Utility Bidders utility conditional on winning the auction, $u^j(n^s, n^l, x)$, is a parametric function of (n^s, n^l, x). From now on we assume that this function is known and we denote it by $u^j(n^s, n^l, x; \beta)$, where β is a vector of parameters to be estimated.

The estimator Pesendorfer and Schmidt-Dengler (2008) propose an ALSE to estimate dynamic games. In what follows we adapt this idea to the present framework. Generically, the ALSE is defined directly from the theoretical constraint developed in proposition 4.2. The estimation is developed in two steps. In the first step we estimate the vector of entry probabilities in the LHS of (4.7) using the available data. The RHS of the same equation is simulated using the equilibrium conditions and the distribution of the unobservables. In the second step we choose the vector of parameters, $\{K, \theta, \beta\}$, that minimizes the (weighted) difference between the LHS and the RHS of (4.7). Formally, the LHS of (4.7) can be easily estimated using the dataset $\{n_t^s, n_t^l, N_t^s, N_t^l, x_t; \forall t \in \mathcal{T}\}$. For this one can define a non parametric estimator for the probabilities of $\left(n^s, n^l, x\right)$ conditional on the set of auction observables. For example, for any x,

$$\hat{p}\left(\left(n^s, n^l\right)|x\right) = \frac{\sum_{l=1}^{L} I\left(n_t^s = n^s, n_t^l = n^l\right) K\left(\frac{x_t - x}{h_t}\right)}{\sum_{l=1}^{L} K\left(\frac{x_t - x}{h_t}\right)} \tag{4.9}$$

is a consistent estimator of the LHS of (4.7). In this expression, $K(\cdot)$ is a kernel function and h_t is a given bandwidth. For the RHS of the equilibrium condition we fix (N, K, x, θ, β) and take $d = 1, 2, ..., D$ draws from the distribution of ϵ. Denoting each draw as ϵ_d, a consistent estimator of the RHS of the equilibrium condition is:

$$\begin{aligned} &\hat{p}\left(\left(n^s, n^l\right)|N, K, x, \theta, \beta\right) = \\ &\frac{\sum_{d=1}^{D} I\left(\left(n^s, n^l\right)|\epsilon_d, N, K, x, \theta, \beta\right) \lambda\left[e\left(\epsilon_d, K, x, \beta\right) = \left(n^s, n^l\right); \theta\right]}{D} \end{aligned} \tag{4.10}$$

In practice, for any auction $t \in \mathcal{T}$ one can calculate consistent estimates of the LHS and the RHS of (4.7) for all equilibrium pairs $\left(n^s, n^l\right)$ (conditional on x_t), using the non parametric approach in expression (4.9) and the procedure summarized by equation (4.10). By stacking the difference between the LHS and the RHS of (4.7) for all equilibrium pairs and for all auctions, the ALS estimator for the model parameters solves:

$$\min_{K,\theta,\beta} \left(\hat{p} - \hat{p}(K, \theta, \beta)\right)' W \left(\hat{p} - \hat{p}(K, \theta, \beta)\right),$$

where $\hat{p}$ and $\hat{p}(K, \theta, \beta)$ are column vectors with dimension corresponding to the number of possible equilibrium pairs in each auction times the number of observed auctions and W is a weight matrix. Pesendorfer and Schmidt-Dengler (2008) derive the asymptotic distribution of the efficient ALSE building on Gourieroux and Monfort (1995) that, in turn, show that the ALSE is asymptotically normal. Given, however, the complexities involved in the estimation of the asymptotic matrix of the ALSE we estimate the variance-covariance matrix of the estimated parameters using a bootstrap procedure and use the asymptotic normality result for inference.

4.6.2 Estimation of Procurement Costs

Now we turn to the estimation of expected procurement costs. To do this, we firstly estimate the term inside the square brackets in equation (4.8), that gives us the expected procurement cost when auction l has $\left(n_t^s, n_t^l\right)$ small and large bidders. Following Aradillas-López et al. (2011) this element can be non parametrically estimated by:

$$\hat{T}_{n^s,n^l}(x) = \frac{\sum_{l=1}^{L} B_t \cdot K\left(\frac{x_t - x}{h_t}\right) \cdot I\left(n_t^s = n^s, n_t^l = n^l\right)}{\sum_{l=1}^{L} K\left(\frac{x_t - x}{h_t}\right) I\left(n_t^s = n^s, n_t^l = n^l\right)},$$

where x is a given vector of covariates, $K(\cdot)$ is a kernel function, h_t is a given bandwidth and B_t is the winning bid in auction l. With this element it is straightforward to see that a consistent estimator of expected procurement costs are:

$$\hat{\pi}\left(N, x; \hat{K}, \hat{\theta}, \hat{\beta}\right) = \sum_{n^s,n^l} \hat{T}_{n^s,n^l}(x)\, \hat{p}\left(\left(n^s, n^l\right) | N, x; \hat{K}, \hat{\theta}, \hat{\beta}\right) - \pi_0$$

With this expression we close our estimation procedure. In the next section we estimate the parameters of this model and provide some "goodness of fit" analysis. Subsequently we discuss how we can use this model to calculate the costs of the set aside policy and to produce counterfactuals.

4.7 Conclusions

In this paper we analyse the effects of a set aside programme targeting small and micro enterprises (SMEs) in Brazil. We use data from ComprasNet, the online bidding platform used by the Brazilian federal government for procurement of various goods and services.

We first provided reduced-form evidence on the effects of the programme using a fuzzy RDD approach. We found that restricting participation of large firms has little effect on prices, while it increases participation of small firms. This finding is consistent with a model of bidder asymmetry and costly participation. In such a model, restricting participation of large, strong bidders increases the incentives of small, weak bidders to participate thus mitigating the adverse effects of the restriction on prices. The increase in the number of SMEs however does not translate into lower prices because, possibly, SMEs have higher costs in comparison to large bidders.

Based on the reduced form evidence, we then set up a structural model to estimate entry costs and simulate the effects of alternative policy rules on participation and procurement costs. Our model has two interesting features. First, differently from the exiting literature, the model and estimation approach proposed here explicitly deals with the multiple equilibria problem arising from entry games with asymmetric players. In particular, we addressed the multiplicity problem using an equilibrium selection mechanism in the spirit of Bajari et al. (2010).

Second, because we observe open auctions and because point identification of open auction models requires very restrictive assumptions - namely that bidders' valuations are independent private costs and that bidders play a button auction (Athey and Haile (2002)) - we considered only information on bidders' entry decisions to estimate a reduced-form specification for bidders utilities, as in Bresnahan and Reiss (1991). Based on this approach we are able to point estimate bidders' entry costs and subsequently government procurement costs using weaker assumptions than those normally used in the literature that studies the estimation of open auctions – Aradillas-López et al. (2011) and Haile and Tamer (2003).

Another interesting feature of the ComprasNet setup is that it allow us to interpret estimated entry costs as red-tape costs. In ComprasNet, bidders are retailers or wholesalers who have good knowledge of their private costs before entering the auction. Bidders are unlikely to spend large amounts of time reading lots' descriptions, as these are short due to the standardisation of products being traded. It means, therefore, that entry costs captures opportunity costs of filling forms, providing documentation and abiding to certain formal procedures. With the structural model we can estimate the magnitude of these costs and analyse how the reduction in bureaucratic procedures affects participation of small and large bidders and procurement costs.

Appendices

Appendix A

Appendix to Chapter 2

This appendix describes the quasi-hard close algorithm described in section 2.1.

```
1. At the start of Phase C, draw X from Unif[5,30] minutes.
2. d=X.
3. While d<28:
4.   If a bid is placed, draw e from Unif[0,2].
5.   d=X+e.
6. end while
7. Return d.
```

Appendix B

Appendix to Chapter 4

B.1 Proofs

Proof of Proposition 4.1 Suppose the system of inequalities in proposition 4.1 holds. Then, no firms (small or big) have an incentive to leave the auction as the outside option is equal to zero. Likewise, no outside firms (small or big) have an incentives to enter as it would mean negative payoff, instead of zero. Therefore, a pair $\left(n^s, n^l\right)$ satisfying the system of equations in proposition 4.1 is a Nash equilibrium in pure strategies.

Conversely, suppose there is pair $\left(n^s, n^l\right) \leq \left(N^s, N^l\right)$ that is a Nash equilibrium in pure strategies and does not satisfy the set of equation in proposition 4.1. Suppose first that:

$$\left\{u^j\left(n^s, n^l, x\right) - K^j - \epsilon^j < 0\right\} I\left(0 < n^j \leq N^j\right) \tag{B.1}$$

For $j = s$ and/or $j = l$. Then, the firms in the group that are making negative profits will leave the auction as they can get zero outside. Now, suppose:

$$\left\{u^j\left(n^s + 1, n^l, x\right) - K^j - \epsilon^j \geq 0\right\} I\left(0 \leq n^s < N^s\right) \tag{B.2}$$

For $j = s$ and/or $j = b$. Then, at least one firm has incentive to enter in the market. Therefore $\left(n^s, n^l\right) \leq \left(N^s, N^l\right)$ cannot be a Nash equilibrium in pure strategies.

Proof of Proposition 4.2 Notice that:

$$\sum_{e(K,x,\epsilon,u(.)) \in \mathcal{E}(K,x,\epsilon,u(.))} \lambda\left[e\left(K, x, \epsilon, u\left(.\right)\right); \theta\right] = 1 \tag{B.3}$$

Therefore, the probability of observing a given outcome $\left(n^s, n^l\right) \leq \left(N^s, N^l\right)$ is equal to the integral of $\lambda\left[e\left(K, x, \epsilon, u\left(.\right)\right); \theta\right]$ over all the possible values of the vector $\left(\epsilon^s, \epsilon^l\right)$.

Proof of Proposition 4.3 Under the hypotheses above, for a given (known) number of small and big players $\pi\left(N, x; K, \theta, u\left(.\right)|n^s, n^l\right) = \int_0^r v dF_{n-1:n}^{n^s, n^l}\left(v|x\right) - \pi_0$ follows readily from Aradillas-López et al. (2011). To get the final expression we weight this term by the distribution probability designed in proposition (4.2).

B.2 Tables

Table B.1: Sample Composition

	NumLots	LotBelow80	BatchBelow80	SetAside	SetAsideC2	SetAsideNotC2
Ink	48689	47406	15749	11802	7373	4429
Paper	4438	3980	1772	1103	731	372
Printer	6853	6449	1815	1248	740	508
AirCo	7774	6946	2219	1623	972	651
Coffee	2111	1965	981	631	496	135
Computer	7622	6074	1790	1050	673	377
GLP	3102	2761	1281	573	428	145
Total	80589	75581	25607	18030	11413	6617

Table reports the composition of the sample described in section 4.3. Column (1) shows the total number of lots for each product category in the sample. Column (2) shows the number of lots satisfying criterion C1 (reserve price below R$ 80,000). Column (3) shows the number of lots meeting criterion C2 (batches with reserve price below R$ 80,000). Column (4) shows the number of lots that were set-aside in each product category. Column (5) shows the number of lots that were set-aside and satisfied criterion C2. Column (5) shows the number of lots that were set-aside and satisfied only criterion C1.

Table B.2: Sample Descriptive Statistics

	Ink		Paper		Printer		AirCo		Coffee		Computer		GLP	
Set Aside =	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Reserve (R$/unit)	211.4	194.5	24.47	29.29	1882.8	1418.7	3095.7	2787.6	13.20	14.91	3341.3	3038.3	76.19	82.04
	(367.2)	(209.4)	(34.9)	(41.2)	(2971.6)	(1964.1)	(5757.2)	(2139.9)	(39.3)	(43.9)	(3125.7)	(2704.9)	(71.8)	(65.7)
Price (R$/unit)	110.9	106.4	18.91	23.24	1374.1	1069.6	2369.7	2464.5	9.693	12.92	2606.4	2572.9	69.64	79.88
	(158.2)	(138.8)	(27.8)	(33.3)	(2369.8)	(1470.9)	(3664.1)	(6791.1)	(31.2)	(47.2)	(2583.4)	(2794.4)	(80.5)	(60.7)
# Big Firms	2.061	0	2.089	0	2.101	0	1.874	0	1.755	0	1.884	0	.919	0
	(1.71)	(0)	(2.38)	(0)	(2.46)	(0)	(1.78)	(0)	(1.98)	(0)	(2.21)	(0)	(.87)	(0)
# Small Firms	17.98	16.71	8.988	8.966	14.29	13.47	13.01	10.70	8.368	8.557	13.20	13.35	1.915	2.274
	(10.1)	(9.66)	(5.31)	(5.45)	(8.17)	(8.68)	(7.75)	(6.25)	(4.46)	(4.20)	(7.46)	(8.24)	(1.59)	(1.20)
Prop. Small Firms	.896	1	.843	1	.885	1	.872	1	.849	1	.880	1	.625	1
	(.085)	(0)	(.15)	(0)	(.12)	(0)	(.12)	(0)	(.14)	(0)	(.13)	(0)	(.36)	(0)
Small Firm Wins	.838	1	.750	1	.847	1	.867	1	.808	1	.822	1	.559	1
	(.37)	(0)	(.43)	(0)	(.36)	(0)	(.34)	(0)	(.39)	(0)	(.38)	(0)	(.50)	(0)
N	35070	11130	3202	1041	4998	1094	5768	1475	1369	571	5877	916	2086	412

Table reports the mean and standard deviation for selected variables by each product in the sample. Standard deviations are in parentheses. A4 paper is measured in 500-sheet packs; coffee is measured in kilos; GLP is measure in 13-kilo bottles.

Bibliography

Aradillas-López, A., Gandhi, A., and Quint, D. (2011). Identification and inference in ascending auctions with correlated private values.

Ariely, D., Ockenfels, A., and Roth, A. E. (2005). An experimental analysis of ending rules in internet auctions. *RAND Journal of Economics*, 36(4):890–907.

Ashenfelter, O. and Graddy, K. (2003). Auctions and the price of art. *Journal of Economic Literature*, 41(3):763–786.

Athey, S., Coey, D., and Levin, J. (2011a). Set-asides and subsidies in auctions. Mimeo.

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

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

Ayres, I. and Cramton, P. (1996). Deficit reduction through diversity: How affirmative action at the fcc increased auction competition.

Bajari, P., Hong, H., and Ryan, S. P. (2010). Identification and estimation of a discrete game of complete information. *Econometrica*, 78(5):1529–1568.

Bajari, P. and Hortaçsu, A. (2003). The winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions. *RAND Journal of Economics*, 34(2):329–355.

Bajari, P. and Hortaçsu, A. (2004). Economic insights from internet auctions. *Journal of Economic Literature*, 42(2):457–486.

Bandiera, O., Prat, A., and Valletti, T. (2009). Active and passive waste in government spending: Evidence from a policy experiment. *American Economic Review*, 99(4):1278–1308.

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

Bresnahan, T. F. and Reiss, P. C. (1990). Entry in monopoly markets. *Review of Economic Studies*, 57(4):531–553.

Bresnahan, T. F. and Reiss, P. C. (1991). Entry and competition in concentrated markets. *Journal of Political Economy*, 99(5):977–1009.

Brown, J. R. and Goolsbee, A. (2002). Does the internet make markets more competitive? evidence from the life insurance industry. *Journal of Political Economy*, 110(3):481–507.

Cassady, R. (1967). *Auctions and Auctioneering*. University of California Press.

Cramton, P. (1998). Ascending auctions. *European Economic Review*, 42(3-5):745–756.

Decarolis, F. (2010). When the highest bidder loses the auction: Theory and evidence from public procurement. Mimeo.

Dutra, J. C. and Menezes, F. M. (2002). Hybrid auctions. *Economics Letters*, 77(3):301–307.

Edelman, B. and Ostrovsky, M. (2007). Strategic bidder behavior in sponsored search auctions. *Decision Support Systems*, 43(1):192–198.

Einav, L., Kuchler, T., Levin, J., and Sundaresan, N. (2011). Learning from seller experiments in online markets. NBER Working Paper w17385.

Ellison, G. and Ellison, S. F. (2005). Lessons about markets from the internet. *The Journal of Economic Perspectives*, 19(2):139–158.

Ely, J. C. and Hossain, T. (2009). Sniping and squatting in auction markets. *American Economic Journal: Microeconomics*, 1(2):68–94.

Engelberg, J. and Williams, J. (2009). eBay's proxy bidding: A license to shill. *Journal of Economic Behavior & Organization*, 72(1):509–526.

Ferraz, C. and Finan, F. (2011). Electoral accountability and corruption: Evidence from the audits of local governments. *American Economic Review*, 101(June):1274–1311.

Füllbrunn, S. and Sadrieh, A. (2012). Sudden termination auctions: An experimental study. *Journal of Economics & Management Strategy*, 21(2):519–540.

Gourieroux, C. and Monfort, A. (1995). *Statistics and Econometric Models*. Cambridge University Press.

Gray, S. and Reiley, D. H. (2007). Measuring the benefits to sniping on eBay: Evidence from a field experiment. Mimeo.

Groeger, J. (2011). A study of participation in dynamic auctions. Mimeo.

Haile, P. A. (2001). Auctions with resale markets: an application u.s. forest service timber sales. *American Economic Review*, 91(3):399–432.

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

Haruvy, E. and Leszczyc, P. T. L. P. (2010). The impact of online auction duration. *Decision Analysis*, 7(1):99–106.

Hasker, K. and Sickles, R. (2010). eBay in the economic literature: Analysis of an auction marketplace. *Review of Industrial Organization*, 37(1):3–42.

Hendricks, K., Porter, R. H., and Boudreau, B. (1987). Information, returns, and bidding behavior in ocs auctions: 1954-1969. *The Journal of Industrial Economics*, 35(4):517–542.

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

Hortaçsu, A., Martinez-Jerez, F. A., and Douglas, J. (2009). The geography of trade in online transactions: Evidence from eBay and mercadolibre. *American Economic Journal: Microeconomics*, 1(1):53–74.

Hörner, J. and Sahuguet, N. (2007). Costly signalling in auctions. *Review of Economic Studies*, 74(1):173–206.

Kagel, J., Pevnitskaya, S., and Ye, L. (2008). Indicative bidding: An experimental analysis. *Games and Economic Behavior*, 62(2):697–721.

Klemperer, P. (1998). Auctions with almost common values: The 'wallet game' and its applications. *European Economic Review*, 42(3-5):757–769.

Klemperer, P. (2004). America's patent protection has gone too far. Financial Times.

Krasnokutskaya, E. and Seim, K. (2011). Bid preference programs and participation in highway procurement auctions. *American Economic Review*, 101(6):2653–2686.

Levin, D. and Smith, J. L. (1994). Equilibrium in auctions with entry. *American Economic Review*, 84(3):585–99.

Levin, D. and Ye, L. (2008). Hybrid auctions revisited. *Economics Letters*, 99(3):591–594.

Levitt, S. D. and List, J. A. (2009). Field experiments in economics: The past, the present, and the future. *European Economic Review*, 53(1):1–18.

Li, T. and Zheng, X. (2009). Entry and competition effects in first-price auctions: Theory and evidence from procurement auctions. *Review of Economic Studies*, 76(4):1397–1429.

Liebman, J. B. and Mahoney, N. (2010). Do expiring budgets lead to wasteful year-end spending? evidence from federal procurement. Mimeo.

Lucking-Reiley, D. (1999). Using field experiments to test equivalence between auction formats: Magic on the internet. *American Economic Review*, 89(5):1063–1080.

Malaga, R., Porter, D., Ord, K., and Montano, B. (2010). A new end-of-auction model for curbing sniping. *Journal of the Operational Research Society*, 61:1265–1272.

Marion, J. (2007). Are bid preferences benign? the effect of small business subsidies in highway procurement auctions. *Journal of Public Economics*, 91(7-8):1591–1624.

McAfee, R. P. and McMillan, J. (1989). Government procurement and international trade. *Journal of International Economics*, 26:291–308.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714.

Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122.

Ockenfels, A., Reiley, D. H., and Sadrieh, A. (2006). Online auctions. NBER Working Paper 12785.

Ockenfels, A. and Roth, A. E. (2006). Late and multiple bidding in second price internet auctions: Theory and evidence concerning different rules for ending an auction. *Games and Economic Behavior*, 55(2):297–320.

Ostrovsky, M. and Schwarz, M. (2009). Reserve prices in internet advertising auctions: A field experiment. Mimeo.

Perry, M. and Reny, P. J. (2002). An efficient auction. *Econometrica*, 70(3):1199–1212.

Perry, M., Wolfstetter, E., and Zamir, S. (2000). A sealed-bid auction that matches the english auction. *Games and Economic Behavior*, 33(2):265–273.

Pesendorfer, M. and Schmidt-Dengler, P. (2008). Asymptotic least squares estimators for dynamic games. *Review of Economic Studies*, 75(3):901–928.

Rasmusen, E. B. (2006). Strategic implications of uncertainty over one's own private value in auctions. *BE Press Journal*, 6(1):1534–5963.

Roberts, J. W. and Sweeting, A. (2010). Entry and selection in auctions. Mimeo.

Roth, A. E. and Ockenfels, A. (2002). Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and amazon auctions on the internet. *American Economic Review*, 92(4):1093–1103.

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

Srinivasan, K. and Wang, X. (2010). Bidders' experience and learning in online auctions: Issues and implications. *Marketing Science*, 29(6):988–993.

Tran, A. (2008). Can procurement auctions reduce corruption? evidence from the internal records of a bribe-paying firm. Mimeo.

Trevathan, J. and Read, W. (2011). Disarming the bid sniper. *Journal of Electronic Commerce Research*, 12(3):176–186.

Varian, H. R. (2010). Computer mediated transactions. *American Economic Review*, 100(2):1–10.

Wintr, L. (2008). Some evidence on late bidding in ebay auctions. *Economic Inquiry*, 46(3):369–379.

World Bank (2004). Brazil country procurement assessment report. Technical Report 28446, World Bank.

Ye, L. (2007). Indicative bidding and a theory of two-stage auctions. *Games and Economic Behavior*, 58(1):181–207.