BOUNDING BEST-RESPONSE VIOLATIONS IN DISCRIMINATORY AUCTIONS WITH PRIVATE VALUES*

James T.E. Chapman
Dept. of Banking Operations
Bank of Canada

David McAdams
Sloan School of Management
MIT

Harry J. Paarsch[†]
Dept. of Economics
University of Iowa

First Draft: September 2005
This Version: July 2006

Abstract

Deciding whether bidders at auctions are playing a best response is perhaps one of the most fundamental questions faced by empirical workers employing the structural econometric approach. Investigations of this question depend heavily on whether the true valuations of bidders are observed. We develop an approach to bound best-response violations when bidders' private values are unobserved in multi-unit discriminatory auctions under the assumption of non-increasing marginal valuations. We also derive a bound on the expected profits lost when a best response is not adopted. Subsequently, we implement our framework using data from auctions of Receiver General term deposits conducted by the central bank in Canada, finding that best-response violations are frequent. For most bidders, however, the median lower bound on the economic distance between the estimated best-response and the bid data is small.

JEL Classification Numbers: C14, D44, E4, E5, L1.

Keywords: multi-unit auctions; discriminatory auctions; pay-as-bid auctions; testing best response; conditionally-independent private values; cash management.

[†]Corresponding author: 108 PBB; 21 East Market Street; Iowa City, Iowa 52242-1944. Telephone: (319) 335-0936; Facsimile: (319) 335-1956; E-mail: `hjp@paarsch.biz.uiowa.edu`

1. Motivation and Introduction

During the last five decades, economists have made considerable progress in understanding the theoretical structure of strategic behaviour in markets with small numbers of participants, such as auction markets; see Krishna (2002) for a comprehensive book-length survey of progress.

During the last two decades, several researchers have brought these theoretical models to data. Most structural econometric research has been devoted to investigating equilibrium behaviour at single-object auctions (SOAs) within the symmetric independent private-values paradigm (IPVP) where each potential bidder gets an independent and identically distributed valuation draw for the object on sale. Examples include Paarsch (1992,1997); Donald and Paarsch (1993,1996,2002); Laffont, Ossard, and Vuong (1995); Guerre, Perrigne, and Vuong (2000); Haile and Tamer (2003); and Li (2005). Paarsch and Hong (2006) provide an introduction to this literature.

In reality, however, most auctions involve the sale of several units of the same good or of several different goods. Economic theorists often make a distinction between multi-object and multi-unit auctions. At multi-unit auctions (MUAs), the objects for sale are identical, so each bidder only cares about how many units he wins. At multi-object auctions (MOAs), on the other hand, the objects for sale are different and it matters to a bidder which specific objects he wins. Thus, an example of a MOA would involve the sale of an apple, an orange, and a pear, while an example of a MUA would involve the sale of three identical apples. Examples of structural econometric research involving MUAs include Donald, Paarsch, and Robert (forthcoming); Jofre-Bonet and Pesendorfer (2004); Wolak (2003); Hortaçsu (2002a,b); Brendstrup (2002); and Brendstrup and Paarsch (2004a,2006). Brendstrup and Paarsch (2004b) have investigated MOAs, considering bundling, while Cantillon and Pesendorfer (2006a,b) have investigated combinatorial auctions in procurement.

Empirical researchers studying auctions have chosen to investigate a variety of different questions using data collected in laboratory experiments or gathered in the field. For the structural econometric approach to be valid, however, each bidder must

be playing a strategy that is a best response to the strategies of others. Do bidders play best responses? When the true valuations are observed at SOAs, as is the case with data from laboratory experiments, a researcher can test for best response in a straightforward way by examining whether the observed bid strategy, conditional on the valuation generated under the experiment, is in fact a best response to the strategies of others; see, for example, Paarsch and Robert (2003). When the true costs are observed, Hortaçsu and Puller (2004) have generalized the approach of Guerre, Perrigne, and Vuong (2000) to test for a best response at MUAs under the uniform-pricing rule. To implement their procedure, Hortaçsu and Puller gathered accurate estimates of producer costs to use as proxies for the true costs.

When field data are used and the true valuations are unobserved, as is more typically the case, testing for a best response may be impossible. For example, consider a single-object, first-price, sealed-bid auction. When the monotone likelihood-ratio property (MLRP) holds for observed bids, there always exists a latent distribution of values that can rationalize actual bidding behaviour. Thus, in this case, one cannot possibly reject the hypothesis that each bidder always plays a best response; see Guerre *et al.* as well as Athey and Haile (forthcoming) for thoughtful discussions.

One might suspect the same to be true at MUAs. Indeed, if the true values (or costs) are unobserved and a condition that generalizes the MLRP holds, then there always exists a latent distribution of values that can rationalize observed bidding behaviour. When values are unobserved, it would seem that rejecting the best-response hypothesis is equally hopeless in MUAs. Even when values are unobserved, however, a researcher may be willing to impose restrictions on the space of possible values. For example, in some applications, it may be reasonable to assume that bidders have non-increasing marginal values; *i.e.*, each bidder's value for a first unit is weakly greater than his value for a second unit, and so forth.

In this paper, we describe a strategy to bound the extent of best-response violations when bidders' private values are unobserved and when one is willing to assume non-increasing marginal values. That is, on a bid-by-bid basis, we establish

2

a lower bound on each bidder's lost expected profit relative to a best response. No matter what values the bidder has, he must be able to increase his expected profit by at least this much with another bid.

We implement our strategy using data from auctions of Receiver General term deposits (hereafter RG auctions) conducted by the central bank in Canada, the Bank of Canada. The RG auctions we investigate are multi-unit, simultaneous, sealed-bid auctions where bidding financial institutions pay what they have tendered for each unit sold, a pricing rule often referred to as *discriminatory*. The format of these auctions was designed by the Bank of Canada to manage Canadian federal government cash balances by lending them out to the highest bidders from a small number of financial institutions who are members of the Large Value Transfer System. When modelling the behaviour of the financial institutions who bid at these auctions, we impose the (natural) assumption that each has non-increasing marginal values for cash.

Why would members of the Large Value Transfer System have different private values for cash? Imagine a financial institution that has inflows and outflows of cash as well as a variety of different instruments of different maturities in its portfolio. In each period, the financial institution solves a large programming problem concerning how to manage its investments. The solution to this programming problem will have Lagrange multipliers associated with constraints involving the various instruments of different maturities. These indicate to managers the "value" of instruments of different maturities. For example, an institution may need a lot of cash tomorrow, but have a portfolio in which very few of its assets are maturing tomorrow. Thus, it has an individual-specific need for one-day cash. It is in this sense that we imagine that cash has different private values for different quantities at different financial institutions.

We confront the observable implications of our model with data from a sample of auctions held between 1 October 2001 and 31 December 2003. We find that best-response violations are frequent. However, for most bidders, depending on our method

of estimation, the median lower bound on the economic distance (in terms of expected profit lost) between the estimated best-response and the bid data is small, often less than the price of a cup of coffee when the implicit values at risk are in the tens of thousands of dollars.

Our paper is in six more parts. In the next section, we outline a theoretical model, while in section 3 we present a summary of some relevant implications of the theory. In section 4, we describe our application, cash-management auctions held on behalf of the Canadian federal government by the Bank of Canada, while in section 5 we present our empirical results. We conclude the paper in section 6 and, in an appendix to the paper, we present a detailed description concerning the construction of the data set used.

2.   Model

Below, we develop a theoretical model that respects the institutional features common to many multi-unit auctions. We divide this section into five parts: 1) the informational environment; 2) the auction rules; 3) a discussion of beliefs; 4) assumptions concerning valuations; and 5) an example.

2.1. Informational Environment

We consider behaviour within the private values paradigm where each bidder has an individual-specific valuation of every unit for sale. Thus, the unobserved heterogeneity across bidders is multi-dimensional.

2.2. Auction Rules

We begin by making the following assumptions: First, supply is known to all bidders; next, bids are discrete and must be chosen from a finite set of alternatives; third, quantities are discrete and must be chosen from a finite set of alternatives; fourth, only a small number of (quantity,bid) pairs can be chosen; and, finally, we focus on discriminatory auctions where each winning bidder pays what he has bid for the units he has won.

We adopt the following notation: At any auction, the total supply is denoted S, while the set of bidders is $\mathcal{N} = \{1, 2, \ldots, N\}$. Permissible prices lie on a grid $\mathcal{P}$ which equals $(p_{\text{out}}, \underline{p}, p_2, \ldots, \bar{p})$, while permissible quantities lie on a grid $\mathcal{Q}$ which equals $(\underline{q}, q_2, \ldots, \bar{q})$. Here, p_{out} denotes unwillingness to buy that unit at any price. Note that, in principle, $\bar{q}$ can vary from participant to participant, but we suppress this dependence for notational parsimony.

We denote by $b_{i,q}$ bidder i's unit-bid on quantity q where $b_{i,q} \in \mathcal{P}$ for each $q \in \mathcal{Q}$. Bidder i's unit-bids are collected in the vector $\boldsymbol{b}_i$, while the profile of bids for all of i's opponents is denoted $\boldsymbol{b}_{-i}$, and the profile of bids for all participants is $\boldsymbol{b}$. The same pattern of subscripts applies to marginal values, bidder strategies, and residual supply. To wit, the first subscript denotes a bidder or set of bidders, the second subscript denotes a quantity or set of quantities, while the absence of subscripts denotes *all* of the participants or quantities. Hence, $v_{i,q}$ denotes bidder i's marginal value for quantity q, while $b_{i,q}(\cdot)$ denotes the projection of bidder i's pure strategy $\boldsymbol{b}_i(\cdot)$ onto quantity q.[1] Also, $b_{i,q}(\boldsymbol{v}_i)$ denotes bidder i's unit-bid on quantity q given his values $\boldsymbol{v}_i$ and strategy $\boldsymbol{b}_i(\cdot)$. Here, $s_{i,q}(\boldsymbol{b}_{-i})$ denotes bidder i's *residual supply* of quantity q. In words, when opponents tender bids $\boldsymbol{b}_{-i}$, bidder i wins less than q units whenever his unit-bid on the q^{th} unit is less than $s_{i,q}(\boldsymbol{b}_{-i})$, or wins at least q units whenever this unit-bid is greater than $s_{i,q}(\boldsymbol{b}_{-i})$. Under this notation, $s_{i,q}[\boldsymbol{b}_{-i}(\boldsymbol{v}_{-i})]$ denotes bidder i's residual supply when his opponents have values $\boldsymbol{v}_{-i}$. Let $q_i(\boldsymbol{b}_i, \boldsymbol{s}_i)$ denote the quantity that bidder i wins. We introduce the following definition:

Definition. A *step* in bid $\boldsymbol{b}_i$ is a maximal interval $[q, q']$ of quantities having the property that

$$b_{i,q} = \cdots = b_{i,q'}.$$

For any step $Q \subset \mathcal{Q}$, let $b_{i,Q}$ denote bidder i's step-bid on those quantities. We denote by $\mathcal{Q}(\boldsymbol{b}_i)$ the partition of $\mathcal{Q}$ into steps induced by the bid vector $\boldsymbol{b}_i$. Each bid

---

[1] McAdams (2003) has proven that a monotone pure-strategy equilibrium exists, while McAdams (forthcoming) proves that every mixed-strategy equilibrium is outcome equivalent to a monotone pure-strategy equilibrium. Thus, it is without loss of generality that we restrict attention to equilibria in pure strategies.

Figure 2.1

Example of Bids Inducing Steps
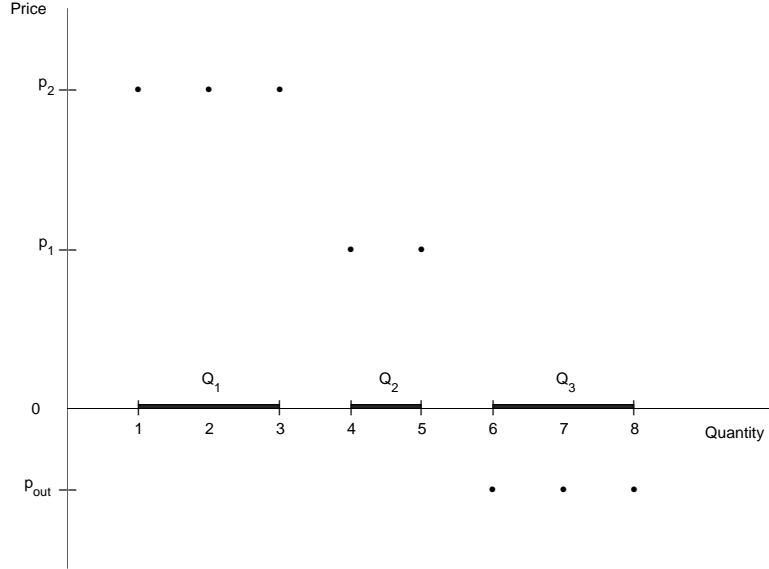
induces a step-partitition of $\{\underline{q}, \ldots, \bar{q}\}$. We illustrate the partition induced by a bid vector having three steps in Figure 2.1, when $\underline{q}$ is one and $\bar{q}$ is eight.

2.3. Beliefs

We ignore any dynamic features of the game and focus, instead, on short-run best-responses. If all bidders are following a best response, then they are in a static Bayes–Nash equilibrium.

We denote by $\mathcal{B}_i$ the (correct) *beliefs* of bidder i concerning the distribution of $\boldsymbol{s}_i(\boldsymbol{b}_{-i})$, given the strategies of others and the distribution of $\boldsymbol{V}_{-i}$. We denote by $\Pi_i(\boldsymbol{b}_i, \boldsymbol{v}_i, \mathcal{B}_i)$ bidder i's expected payoff, given his own values $\boldsymbol{v}_i$, his bid vector $\boldsymbol{b}_i$, and his beliefs $\mathcal{B}_i$ about the distribution of others' bids.

2.4. Assumptions concerning Valuations

We make the following assumptions concerning the valuations of bidders:

A1. $(\boldsymbol{V}_1, \ldots, \boldsymbol{V}_N)$ are independent.

6

A2. For every bidder $i \in \mathcal{N}$, $(V_{i,\underline{q}}, \ldots, V_{i,\bar{q}})$ are drawn from the same joint distribution at each auction.

A3. Each bidder has non-increasing marginal values,

$$V_{i,\underline{q}} \geq V_{i,2} \geq \ldots \geq V_{i,\bar{q}} \geq 0.$$

Other than non-increasing marginal values, we have not made any other restrictive assumptions concerning the joint distribution of $\boldsymbol{V}_i$.

In equilibrium, each bidder i tenders a bid that maximizes his expected payoff given his values $\boldsymbol{v}_i$ and correct beliefs $\mathcal{B}_i$ concerning the distribution of his residual supply. Computing an equilibrium is difficult, even if we know the distribution of bidder values. However, inferring possible bidder values from observed bids is much less difficult than calculating an equilibrium, provided we *assume* that observed bids are consistent with equilibrium.

An observed bid can only be a best response if the bidder does not prefer to perturb that bid in any way. (Global deviations provide additional restrictions.) This provides us with a (huge) system of equalities and inequalities involving $\boldsymbol{v}_i$ that must be satisfied, corresponding to all the perturbations that must be unprofitable.

2.5. An Example

To see this, consider the following example: Suppose there are three units for sale and we observe bidder i making the bid ($100, $100, $50). For this bid to be a best response, given a minimum price increment of $1, bidder i must at least weakly prefer

7

it to *all* of the following bids:

$$(\$101, \$101, \$50), (\$101, \$100, \$50), (\$100, \$99, \$50),$$

$$(\$101, \$99, \$50), (\$99, \$99, \$50), (\$101, \$101, \$51),$$

$$(\$101, \$100, \$51), (\$100, \$100, \$51), (\$100, \$99, \$51),$$

$$(\$101, \$99, \$51), (\$99, \$99, \$51), (\$101, \$101, \$49), \tag{2.1}$$

$$(\$101, \$100, \$49), (\$100, \$100, \$49), (\$100, \$99, \$49),$$

$$(\$101, \$99, \$49), (\$99, \$99, \$49).$$

Fortunately, others have shown that each bidder's expected payoff is additively-separable in his unit-bid on the first unit, the second unit, and so on; see Hortaçsu (2002b). In particular, expected payoffs can be expressed as

$$\Pi_i(\boldsymbol{b}_i, \boldsymbol{v}_i; \mathcal{B}_i) = \sum_{q \in \mathcal{Q}} \Pi_{i,q}(b_{i,q}, v_{i,q}; \mathcal{B}_i) \tag{2.2}$$

where

$$\Pi_{i,q}(b_{i,q}, v_{i,q}; \mathcal{B}_i) = \Pr(b_{i,q} > s_{i,q}|\mathcal{B}_i)(v_{i,q} - b_{i,q}). \tag{2.3}$$

Note that bidder i's expected unit-payoff for quantity q does not depend on his marginal value or unit-bid for other units. Furthermore, it only depends on the marginal distribution of bidder i's residual supply on the q^{th} unit $s_{i,q}$, rather than the joint distribution of $\boldsymbol{s}_i$.[2]

Thus, all perturbations of (2.1) are unprofitable as soon as we show that all perturbations in a smaller set

$$(\$101, \$101, \$50), (\$101, \$100, \$50), (\$100, \$99, \$50),$$

$$(\$99, \$99, \$50), (\$100, \$100, \$51), (\$100, \$100, \$49)$$

---

[2] Ties may be important theoretically, but in our application below they are unimportant empirically, so for parsimony we omit discussing them here.

are unprofitable. For example, if ($101, $99, $51) is a profitable deviation, then either ($101, $100, $50), or ($100, $99, $50), or ($100, $100, $51) must also be a profitable deviation.

If $\mathcal{P}$ were a continuum, so we could take derivatives, then these requirements would reduce to

$$\frac{d\Pi_{i,1}(b_{i,1}, v_{i,1}; \mathcal{B}_i)}{db_{i,1}} \leq 0$$

$$\frac{d\Pi_{i,1}(b_{i,1}, v_{i,1}; \mathcal{B}_i)}{db_{i,1}} + \frac{d\Pi_{i,2}(b_{i,2}, v_{i,2}; \mathcal{B}_i)}{db_{i,2}} = 0$$

$$\frac{d\Pi_{i,2}(b_{i,2}, v_{i,2}; \mathcal{B}_i)}{db_{i,2}} \geq 0$$

$$\frac{d\Pi_{i,3}(b_{i,3}, v_{i,3}; \mathcal{B}_i)}{db_{i,3}} = 0.$$

Notice that the first three conditions do not depend on $v_{i,3}$, while the fourth condition does not depend on $(v_{i,1}, v_{i,2})$. When identifying the values $\boldsymbol{v}_i$ for which the observed bid is a best response, it suffices to consider the two step-bids separately. The pair $(v_{i,1}, v_{i,2})$ is identified by the step-bid $(b_{i,1}, b_{i,2})$, which equals ($100, $100), while $v_{i,3}$ is identified by the step-bid $b_{i,3}$ equal $50. The identification problem for $v_{i,3}$ is much like at a single-object, first-price, sealed-bid auction: We find the $v_{i,3}$ that rationalizes $b_{i,3}$ by inverting the fourth condition above. The identification problem for $(v_{i,1}, v_{i,2})$ can be broken into two parts. First, replace the first three conditions with the following two equalities:

$$\frac{d\Pi_{i,1}(b_{i,1}, v_{i,1}; \mathcal{B}_i)}{db_{i,1}} = 0 \tag{2.4}$$

and

$$\frac{d\Pi_{i,2}(b_{i,2}, v_{i,2}; \mathcal{B}_i)}{db_{i,2}} = 0. \tag{2.5}$$

For now, use the notation $(\tilde{v}_{i,1}, \tilde{v}_{i,2})$ for values that solve the system of equalities defined by equations (2.4) and (2.5). Second, identify all other $(v_{i,1}, v_{i,2})$ pairs that satisfy

$$\frac{d\Pi_{i,1}(b_{i,1}, v_{i,1}; \mathcal{B}_i)}{db_{i,1}} < 0, \tag{2.6}$$

9

$$\frac{d\Pi_{i,2}(b_{i,2}, v_{i,2}; \mathcal{B}_i)}{db_{i,2}} > 0, \tag{2.7}$$

and

$$\frac{d\Pi_{i,1}(b_{i,1}, v_{i,1}; \mathcal{B}_i)}{db_{i,1}} + \frac{d\Pi_{i,2}(b_{i,2}, v_{i,2}; \mathcal{B}_i)}{db_{i,2}} = 0. \tag{2.8}$$

Overall, there will typically be a *set* of values that can rationalize the observed step-bid. We depict this set graphically as the line segment between A and B in Figure 2.2.[3] In that figure, the element marked A satisfies equations (2.4) and (2.5), while all others between A and B satisfy conditions (2.6)–(2.8). Note that, in this example, all points on the line segment between A and B correspond to *increasing* marginal values. For instance, point B corresponds to the extreme case in which the bidder requires two units to get any value from winning. Given this bid, we conclude that *either* bidder i has increasing marginal values *or* ($\$100, \$100, \$50$) is not a best response. Now, under the assumption that the bidder has non-increasing marginal values, the bidder type corresponding to some point like C will be "closest" to finding the bid ($\$100, \$100, \$50$) to be a best response. We devise methods to estimate the economic distance between points like A and points like C in terms of lost expected profits relative to a best response.

3. Implications of the Theory

In this section, building on McAdams (2006), we develop the relevant theoretical implications for our investigation of the best-response hypothesis. First, as in the study of first-price SOAs by Guerre *et al.*, a bid in a discriminatory MUA model cannot possibly be a best response unless a MLRP-like property is satisfied by the distribution of observed bids. Unlike in models of first-price SOAs, however, one can still reject the best-response hypothesis in discriminatory MUA models when a MLRP-like property is satisfied. In particular, given non-increasing marginal values, we develop methods to detect and to bound the magnitude of inferred best-response violations.

---

[3] If negative marginal values are admitted, then the set would correspond to the ray from A passing through B.

Figure 2.2

Best-Response Violation given Non-Increasing Marginal Valuations
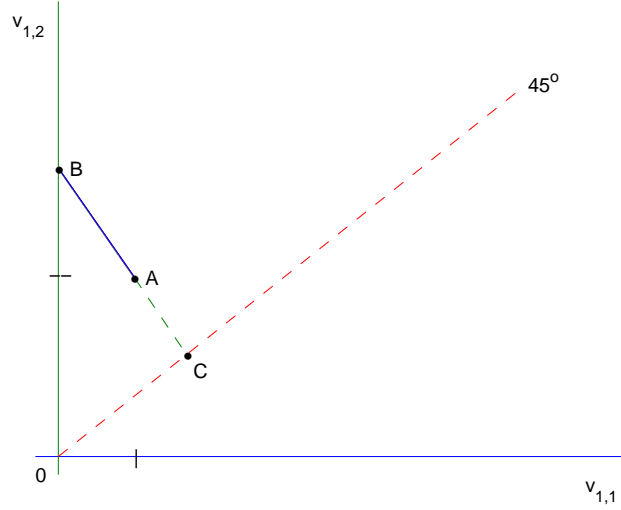
3.1. Restrictions on Bids

For each bidder i and each quantity q, introduce

$$\lambda_{i,q}(b) \equiv \Pr(b_{i,q} > s_{i,q}|\mathcal{B}_i).$$

From section 2, we then know that bidder i's expected payoff from bid $\boldsymbol{b}_i$, given values $\boldsymbol{v}_i$ and belief $\mathcal{B}_i$, takes the form

$$\Pi_i(\boldsymbol{b}_i, \boldsymbol{v}_i; \mathcal{B}_i) = \sum_{q\in\mathcal{Q}} \lambda_{i,q}(b_{i,q})(v_{i,q} - b_{i,q}).$$

3.1.1. Monotone Likelihood-Ratio Property

For the moment, consider the first-price SOA, a special case of the discriminatory MUA when only one unit is sold. Here, $\lambda_{i,1}(b)$ is the probability that bidder i wins one unit with bid b. Suppose that bidder i submits a bid b_i that exceeds the reserve

price. Bidder i could have instead submitted (b_i-1) or (b_i+1). By revealed preference, it must have been that bidder i preferred submitting b_i over these two other bids:

$$\lambda_{i,1}(b_i)(v_i - b_i) \geq \lambda_{i,1}(b_i - 1)(v_i - b_i + 1)$$

$$\lambda_{i,1}(b_i)(v_i - b_i) \geq \lambda_{i,1}(b_i + 1)(v_i - b_i - 1).$$

Re-arranging these two inequalities yields the well-known MLRP of equilibrium bids at a first-price SOA: For bid b_i to be a best response for bidder i, it must be the case that

$$\frac{\lambda_{i,1}(b_i - 1)}{[\lambda_{i,1}(b_i) - \lambda_{i,1}(b_i - 1)]} \leq v_i - b_i \leq \frac{\lambda_{i,1}(b_i + 1)}{[\lambda_{i,1}(b_i + 1) - \lambda_{i,1}(b_i)]}. \tag{3.1}$$

If condition (3.1) fails, then bidder i must prefer *either* to lower his bid to $(b_i - 1)$ *or* to raise his bid to $(b_i + 1)$, regardless of his value; see, for example, the discussion in Athey and Haile (forthcoming).

In practice, when several units are sold at a discriminatory auction, bidders typically submit bids having steps that contain several units. Let $b_{i,Q}$ be one step-bid in an observed bid $\boldsymbol{b}_i$. Bidder i could have instead submitted $(b_{i,Q} - 1)$ or $(b_{i,Q} + 1)$ on all units in the step Q. By revealed preference, it must have been that bidder i prefers submitting $b_{i,Q}$ on these units:

$$\sum_{q \in Q} \lambda_{i,q}(b_{i,Q})(v_{i,q} - b_{i,Q}) \geq \sum_{q \in Q} \lambda_{i,q}(b_{i,Q} - 1)(v_{i,q} - b_{i,Q} + 1) \tag{3.2}$$

$$\sum_{q \in Q} \lambda_{i,q}(b_{i,Q})(v_{i,q} - b_{i,Q}) \geq \sum_{q \in Q} \lambda_{i,q}(b_{i,Q} + 1)(v_{i,q} - b_{i,Q} - 1). \tag{3.3}$$

Re-arranging conditions (3.2) and (3.3) yields

$$\sum_{q \in Q} \lambda_{i,q}(b_{i,Q} + 1) \geq \sum_{q \in Q} (v_{i,q} - b_{i,Q}) \left[\lambda_{i,q}(b_{i,Q} + 1) - \lambda_{i,q}(b_{i,Q}) \right]$$

$$\sum_{q \in Q} \lambda_{i,q}(b_{i,Q} - 1) \leq \sum_{q \in Q} (v_{i,q} - b_{i,Q}) \left[\lambda_{i,q}(b_{i,Q}) - \lambda_{i,q}(b_{i,Q} - 1) \right],$$

which requires

$$\sum_{q \in Q} (v_{i,q} - b_{i,Q}) \frac{\left[\lambda_{i,q}(b_{i,Q}) - \lambda_{i,q}(b_{i,Q} - 1) \right]}{\sum_{q \in Q} \lambda_{i,q}(b_{i,Q} - 1)} \geq 1 \geq$$
$$\sum_{q \in Q} (v_{i,q} - b_{i,Q}) \frac{\left[\lambda_{i,q}(b_{i,Q} + 1) - \lambda_{i,q}(b_{i,Q}) \right]}{\sum_{q \in Q} \lambda_{i,q}(b_{i,Q} + 1)}. \tag{3.4}$$

Table 3.1
Summary Statistics: Step Sizes versus Bids
Millions of Dollars

Summary Statistic	All Steps	1$^{\text{st}}$ Steps	2$^{\text{nd}}$ Steps	3$^{\text{rd}}$ Steps	4$^{\text{th}}$ Steps
Average	304	374	102	42	32
First Decile	23	25	20	20	10
Lower Quartile	33	50	25	20	20
Median	99	225	25	25	20
Upper Quartile	632	716	100	27	25

When Q is $\{1\}$, condition (3.4) reduces to condition (3.1) and fails to hold whenever

$$\frac{\lambda_{i,1}(b_{i,1} - 1)}{[\lambda_{i,1}(b_{i,1}) - \lambda_{i,1}(b_{i,1} - 1)]} > \frac{\lambda_{i,1}(b_{i,1} + 1)}{[\lambda_{i,1}(b_{i,1} + 1) - \lambda_{i,1}(b_{i,1})]}.$$

More generally, condition (3.4) fails whenever

$$\frac{\lambda_{i,q}(b_{i,Q} - 1)}{[\lambda_{i,q}(b_{i,Q}) - \lambda_{i,q}(b_{i,Q} - 1)]} > \frac{\lambda_{i,q}(b_{i,Q} + 1)}{[\lambda_{i,q}(b_{i,Q} + 1) - \lambda_{i,q}(b_{i,Q})]} \quad \text{for } all \; q \in Q. \quad (3.5)$$

As long as condition (3.5) does not fail for *some* $\tilde{q} \in Q$, however, one can find marginal values (with relatively large marginal value for the $\tilde{q}^{\text{th}}$ unit), so the inequalities (3.4) are satisfied for step Q. We summarize this discussion in the following theorem:

Theorem 1: *An observed bid $\boldsymbol{b}_i$ cannot be a best response at a discriminatory MUA unless, for every step $Q \in \mathcal{Q}(\boldsymbol{b}_i)$, there exists $q \in Q$ such that*

$$\frac{[\lambda_{i,q}(b_{i,Q}) - \lambda_{i,q}(b_{i,Q} - 1)]}{\sum_{q \in Q} \lambda_{i,q}(b_{i,Q} - 1)} \geq \frac{[\lambda_{i,q}(b_{i,Q} + 1) - \lambda_{i,q}(b_{i,Q})]}{\sum_{q \in Q} \lambda_{i,q}(b_{i,Q} + 1)}.$$

Note that a best response can only be rejected by our generalization of the MLRP when $\#(Q)$ separate conditions hold simultaneously for some $Q \in \mathcal{Q}(\boldsymbol{b})$ where $\#(Q)$ denotes the number of quantities in step Q. For our application, we tabulated in Table 3.1 the average step-size as well as some of the quantiles for all of the bids as well as the first, second, third, and fourth steps. Note that ninety percent of the steps span a range of more than twenty million dollars. In our application, after the first step, which is five million dollars, the minimum step size is one million dollars. Thus, almost all observed steps contain *many* quantities. Not surprisingly, we *never* fail to reject the best-response hypothesis due to a failure of MLRP in our application.

3.1.2. Non-Increasing Marginal Values

In this section, we derive novel testable restrictions, provided one is willing to assume that bidders have non-increasing marginal values. We summarize these in the following theorem:

Theorem 2: *Suppose that bidder i has non-increasing marginal values. An observed bid $\boldsymbol{b}_i$ cannot be a best response at a discriminatory MUA unless, for every step $Q \in \mathcal{Q}(\boldsymbol{b}_i)$ and every $q \in Q$,*

$$\frac{\sum_{\tilde{q} \in [\min Q, q]} \lambda_{i,\tilde{q}}(b_{i,Q} + 1)}{\sum_{\tilde{q} \in [\min Q, q]} \left[\lambda_{i,\tilde{q}}(b_{i,Q} + 1) - \lambda_{i,\tilde{q}}(b_{i,Q})\right]} \geq \frac{\sum_{\tilde{q} \in [q, \max Q]} \lambda_{i,\tilde{q}}(b_{i,Q} - 1)}{\sum_{\tilde{q} \in [q, \max Q]} \left[\lambda_{i,\tilde{q}}(b_{i,Q}) - \lambda_{i,\tilde{q}}(b_{i,Q} - 1)\right]}. \tag{3.6}$$

Proof: Consider any step $Q \in \mathcal{Q}(\boldsymbol{b}_i)$ and any $q \in Q$. One feasible deviation for bidder i is to raise his unit-bids on units $[\min Q, q] \subset Q$ from $b_{i,Q}$ to $b_{i,Q}+1$. Another feasible deviation is to lower his unit-bids on units $[q, \max Q] \subset Q$ from $b_{i,Q}$ to $b_{i,Q} - 1$.[4] For $\boldsymbol{b}_i$ to be a best response, both of these deviations must be unprofitable. In particular, the following two inequalities must hold:

$$\sum_{\tilde{q} \in [\min Q, q]} \lambda_{i,\tilde{q}}(b_{i,Q})(v_{i,\tilde{q}} - b_{i,Q}) \geq \sum_{\tilde{q} \in [\min Q, q]} \lambda_{i,\tilde{q}}(b_{i,Q} + 1)(v_{i,\tilde{q}} - b_{i,Q} - 1) \tag{3.7}$$

$$\sum_{\tilde{q} \in [q, \max Q]} \lambda_{i,\tilde{q}}(b_{i,Q})(v_{i,\tilde{q}} - b_{i,Q}) \geq \sum_{\tilde{q} \in [q, \max Q]} \lambda_{i,\tilde{q}}(b_{i,Q} - 1)(v_{i,\tilde{q}} - b_{i,Q} + 1) \tag{3.8}$$

Re-arranging conditions (3.7) yields

$$\sum_{\tilde{q} \in [\min Q, q]} \lambda_{i,\tilde{q}}(b_{i,Q} + 1) \geq \sum_{\tilde{q} \in [\min Q, q]} \left[\lambda_{i,\tilde{q}}(b_{i,Q} + 1) - \lambda_{i,\tilde{q}}(b_{i,Q})\right] (v_{i,\tilde{q}} - b_{i,Q})$$

$$\geq (v_{i,q} - b_{i,Q}) \sum_{\tilde{q} \in [\min Q, q]} \left[\lambda_{i,\tilde{q}}(b_{i,Q} + 1) - \lambda_{i,\tilde{q}}(b_{i,Q})\right].$$

---

[4] In our application, around 2.54 percent of our unique bid pairs contain four steps, the maximal number of steps allowed. Thus, the maximal number of bids allowed rarely appears to be binding, at least in our application.

where the second inequality follows from the assumption of non-increasing marginal values since $v_{i,\tilde{q}}$ is less than or equal to $v_{i,q}$ for all $\tilde{q}$ less than or equal to q. Similarly, re-arranging condition (3.8) and using the fact that $v_{i,\tilde{q}}$ is less than or equal to $v_{i,q}$ for all $\tilde{q}$ greater than or equal to q, yields

$$
\sum_{\tilde{q}\in[q,\max Q]} \lambda_{i,\tilde{q}}(b_{i,Q}-1) \leq \sum_{\tilde{q}\in[q,\max Q]} \left[\lambda_{i,\tilde{q}}(b_{i,Q}) - \lambda_{i,\tilde{q}}(b_{i,Q}-1)\right] (v_{i,\tilde{q}} - b_{i,Q})
$$
$$
\leq (v_{i,q} - b_{i,Q}) \sum_{\tilde{q}\in[q,\max Q]} \left[\lambda_{i,\tilde{q}}(b_{i,Q}) - \lambda_{i,\tilde{q}}(b_{i,Q}-1)\right].
$$

Altogether then

$$
\frac{\sum_{\tilde{q}\in[\min Q,q]} \lambda_{i,\tilde{q}}(b_{i,Q}+1)}{\sum_{\tilde{q}\in[\min Q,q]} \left[\lambda_{i,\tilde{q}}(b_{i,Q}+1) - \lambda_{i,\tilde{q}}(b_{i,Q})\right]} \geq v_{i,q} - b_{i,q} \geq
$$
$$
\frac{\sum_{\tilde{q}\in[q,\max Q]} \lambda_{i,\tilde{q}}(b_{i,Q}-1)}{\sum_{\tilde{q}\in[q,\max Q]} \left[\lambda_{i,\tilde{q}}(b_{i,Q}) - \lambda_{i,\tilde{q}}(b_{i,Q}-1)\right]}.
$$

If (3.6) is violated for any quantity q, then bidder i must find one of these two deviations to be profitable, *regardless of his values*. ∎

3.2. Discussion

When observed bids contain steps having many quantities, the restrictions imposed by Theorem 1 become very weak, while those imposed by Theorem 2 remain relatively strong. Consider an example in which an observed bid $\boldsymbol{b}_i$ for two hundred units has two steps, each of length one hundred. To conclude that $\boldsymbol{b}_i$ could not be a best response, Theorem 1 requires that either one hundred conditions hold simultaneously on the functions $[\lambda_{i,1}(\cdot); ...; \lambda_{i,100}(\cdot)]$ or one hundred conditions hold simultaneously on the functions $[\lambda_{i,101}(\cdot); ...; \lambda_{i,200}(\cdot)]$. If one is willing to assume that bidders have non-increasing marginal values, then Theorem 2 provides many more restrictions. In particular, we may conclude that $\boldsymbol{b}_i$ is not a best response if any of two hundred conditions hold. One hundred of these conditions depend on $[\lambda_{i,1}(\cdot); ...; \lambda_{i,100}(\cdot)]$, while the other one hundred depend on $[\lambda_{i,101}(\cdot); ...; \lambda_{i,200}(\cdot)]$.

In our empirical application, a typical step contains dozens if not hundreds of quantities. Not surprisingly, the testable restrictions of Theorem 1 have no bite. More interesting is the fact that many observed bids violate at least one of the restrictions imposed by Theorem 2. This allows us to conclude that *either* some bidders often fail to have non-increasing marginal values *or* some bidders often fail to play a best response.

Kastl (2005) has provided an alternative explanation for why bidders submit bids having steps with many quantities. In his model, bidders prefer to submit bids having fewer steps, because they must pay a cost that is increasing in the number of steps that they use. Could such bidding complexity costs account for observed bidding behaviour? Fortunately, our analysis can be adapted to test this alternative hypothesis as well. From any given bid, each bidder has a large set of feasible deviations available to him, including many that do not increase the number of steps. A bidder who faces complexity costs must not prefer to deviate in any way that *does not increase* the number of steps.

Consider any observed bid $\boldsymbol{b}_i$ having *at least three* steps, and consider any step in this bid that is neither the first nor the last step; *i.e.*, $Q \in \mathcal{Q}(\boldsymbol{b}_i)$ and $Q \not\ni \underline{q}, \bar{q}$. Let $\underline{Q}, \overline{Q} \in \mathcal{Q}(\boldsymbol{b}_i)$ be the steps that are adjacent below and above Q; *i.e.*, $\underline{Q} \ni \min Q - 1$ and $\overline{Q} \ni \max Q + 1$. For every $q \in Q$, bidder i could raise all unit-bids on quantities $[\min Q, q]$ from $b_{i,Q}$ to $b_{i,\underline{Q}}$ *or* lower all unit-bids on quantities $[q, \max Q]$ from $b_{i,Q}$ to $b_{i,\overline{Q}}$. Each of these deviations does not increase the number of steps. Unless an inequality very similar to condition (3.6) is satisfied, one of these two deviations must be profitable for bidder i.

Theorem 3: *Suppose that bidder i has non-increasing marginal values. An observed bid $\boldsymbol{b}_i$ having three or more steps cannot be a best response (even given bidding complexity costs) in the discriminatory auction unless, for every triple of consecutive steps $(\underline{Q}, Q, \overline{Q}) \subset \mathcal{Q}(\boldsymbol{b}_i)$ and every $q \in Q$,*

$$\frac{\sum_{\tilde{q} \in [\min Q, q]} \lambda_{i,\tilde{q}}(b_{i,\underline{Q}})}{\frac{\sum_{\tilde{q} \in [\min Q, q]} (\lambda_{i,\tilde{q}}(b_{i,\underline{Q}}) - \lambda_{i,\tilde{q}}(b_{i,Q}))}{b_{i,\underline{Q}} - b_{i,Q}}} \geq \frac{\sum_{\tilde{q} \in [q, \max Q]} \lambda_{i,\tilde{q}}(b_{i,\overline{Q}})}{\frac{\sum_{\tilde{q} \in [q, \max Q]} (\lambda_{i,\tilde{q}}(b_{i,Q}) - \lambda_{i,\tilde{q}}(b_{i,\overline{Q}}))}{b_{i,Q} - b_{i,\overline{Q}}}}.$$

The proof of Theorem 3 is almost identical to that of Theorem 2, and omitted to save space. To be applicable, Theorem 3 requires that bidders submit bids having at least three steps. Unfortunately, in our data set, bidders rarely submit bids having three or more steps. For this reason, we must wait for future research to test whether bidding complexity costs alone can rationalize observed bids.

3.3. Calibrating the Magnitude of the Best-Response Violations

Let $\boldsymbol{b}_i(\boldsymbol{v}_i)$ denote a best response for bidder i given values $\boldsymbol{v}_i$. When an observed bid vector $\boldsymbol{b}_i$ fails condition (3.6), we may conclude that

$$\Pi_i(\boldsymbol{b}_i, \boldsymbol{v}_i, \mathcal{B}_i) < \Pi_i[\boldsymbol{b}_i(\boldsymbol{v}_i), \boldsymbol{v}_i, \mathcal{B}_i] \quad \text{for all non-increasing} \ \ \boldsymbol{v}_i.$$

Although $\boldsymbol{b}_i$ may not to be a best response, given non-increasing marginal values, it need not be "far" from being a best response. For each bid that violates condition (3.6), we compute a lower bound of the bidder's lost expected profit, relative to a best response:

$$\mathsf{Loss}_i(\boldsymbol{b}_i) = \min_{\boldsymbol{v}_i} \ \{\Pi_i[\boldsymbol{b}_i(\boldsymbol{v}_i), \boldsymbol{v}_i, \mathcal{B}_i] - \Pi_i(\boldsymbol{b}_i, \boldsymbol{v}_i, \mathcal{B}_i)\}$$

$$\text{for all non-increasing} \ \ \boldsymbol{v}_i. \tag{3.9}$$

4.   Empirical Application: Auctions of Receiver General Term Deposits

As the fiscal agent of the Canadian federal government, the Bank of Canada advises the government on its profile of borrowing as well as manages its day-to-day cash and foreign exchange reserves.[5] The RG auctions, which are conducted in both the morning and the afternoon, are the main instruments through which the Bank of Canada conducts cash management. The Bank of Canada uses the morning MUAs to invest excess government cash in term-deposits held at various financial institutions. The afternoon MUAs are used to sterilize the effect that government cash flows

---

[5] The majority of the material in this section was derived from Merrett, Boisvert, and Côté (1995).

have on the net-supply of reserves funds among Canadian financial institutions, thus ensuring that the Bank of Canada's monetary-policy objectives are met.

Short-term repurchases of central-bank funds, also referred to as *repos*, are the main traded instruments in the overnight market. These repos are exchanged among major investment dealers in Canada. While the maturities of the repos can vary from overnight to one week, the bulk involves overnight funds. The market-makers are direct clearers in the main Canadian payment and settlement system, the Large Value Transfer System (LVTS). Firms permitted to use the LVTS include the seven chartered banks of Canada as well as several large credit unions and some foreign banks.[6]

The LVTS, which began operations in February 1999, is the payment and settlement system through which large-value or time-critical transactions are processed. Only financial institutions that, among other technical requirements, are members of the Canadian Payments Association (CPA) *and* have an account at the Bank of Canada may use the LVTS. All LVTS payments are settled through the accounts of participants held at Bank of Canada. Negative balances held at the end of trading need to be financed through loans from the Bank of Canada at the Bank rate, which the Bank of Canada sets as part of its monetary-policy target.[7] Positive balances held at the end of trading accrue interest at the Bank rate minus 50 basis points; *i.e.*, at one-half of one percent below the Bank rate. No participant is required to hold a specific level of reserves at the Bank of Canada. This is essentially a tunnel system of monetary policy implementation as described in Whitesell (forthcoming) and Chapman (2006).

The end-of-day net-supply of reserve funds within LVTS is targeted by the Bank

---

[6] In November 2002, for example, the LVTS participants included: Alberta Treasury Branches, Bank of America National Association, Bank of Montreal, Bank of Nova Scotia, BNP Paribas (Canada), La Caisse Centrale Desjardins du Québec, Canadian Imperial Bank of Commerce, Credit Union Central of Canada, HSBC Bank Canada, Laurentian Bank of Canada, National Bank of Canada, Royal Bank of Canada, and Toronto Dominion Bank.

[7] Prior to 2001, this was the Bank rate; after 2001, it is the Bank of Canada's target rate, which is the Bank rate minus 25 basis points.

of Canada as part of an effort to reduce the effects of uncertainty concerning public-sector flows (such as federal government disbursements and receipts of funds) have on the overnight rate. Cash managers at the Bank of Canada determine the target for end-of-day net-supply in the afternoon prior to the afternoon RG auction. At the inception of LVTS, this targeted level was a zero net-supply of reserves. In the latter half of 2000, this target was changed to a small positive amount. Actual net-supply in the LVTS may not equal the target amount because of unexpected changes that occur after the afternoon RG auction. In general, these differences are small and unforecastable by the Bank of Canada.

The constant level of net-supply of reserve funds as well as the penalty rates that the Bank of Canada imposes on non-zero account balances in the LVTS effectively mean that the overnight rate remains inside a 50-basis point band, with the Bank rate being the upper bound and the target rate being the middle.

At RG auctions, cash in excess of that needed by the Canadian federal government in its daily operations is deposited with direct clearers of the LVTS. The interest rate on these deposits is determined at a sequence of MUAs. These auctions can be thought of as auctions at which the Bank of Canada seeks the highest interest rate for its deposits. The size of an RG auction is based on the daily operational needs of the Canadian federal government as well as daily monetary-policy operations; see Merrett *et al.* (1995) as well as Howard (1998). Thus, it is reasonable to assume that the supply of funds is exogenous.

The auctions conducted by the Bank of Canada are multi-unit, simultaneous, dependent auctions. The term *multi-unit* refers to the fact that billions of dollars of cash are at auction, while the term *simultaneous* refers to the fact that all participants in the LVTS must submit bids by a certain deadline, 9:15 a.m. for the morning RG auctions and 4:15 p.m. for the afternoon RG auctions. The bids are then processed at the same time. The term *dependent* refers to the fact that a bidder's action determines both the price he pays and the number of units he wins. At the RG auctions,

19

participants submit (size,rate) pairs for term-deposits of different maturities.[8] For example, in the market for overnight deposits, the most important market, a financial institution might submit a tender of 3.15 percent (an annualized interest rate) for $100 million, another of 3.00 percent for $100 million, and 2.90 percent for $200 million. At the RG auctions, participants can submit up to four (size,rate) pairs. The Bank of Canada then aggregates all of these bid vectors and allocates the available units to the highest bidders.

The prices that bidders pay at auction can be determined in at least two different ways. At *uniform-price* auctions, all bidders pay the same per-unit price, which typically equals the lowest bid that won some amount of the good. The United States Treasury sells Treasury bills according to this pricing rule as does the Bank of Canada when selling inflation-protected bonds. On the other hand, at discriminatory auctions, each bidder pays the amount he bid for each unit he has won. The Bank of Canada uses the discriminatory pricing rule at RG auctions.

At both the morning and the afternoon RG auctions, the term to maturity for the funds being auctioned can vary from overnight to a month, but the bulk is overnight. Multiple deposits with different maturities may be offered simultaneously. Each distinct deposit being offered at auction is known as a *tranche*. Auctions are conducted through the Bank of Canada's Communication, Auction and Reporting System (CARS). Results are settled using the LVTS, directly through accounts maintained at the Bank of Canada.

Bids at RG auctions are subject to the following restrictions: First, only participants in the LVTS system are eligible to bid; second, each bidder may submit up to four bids for any tranche of deposits being auctioned; third, bids take the form of (size,rate) pairs, where the size equals the amount the participant is willing to accept as a deposit and the rate is the specified interest rate on this deposit; fourth, bidders pay their bid; fifth, all rates bid are in the form of an annualized interest rate specified

---

[8] Equivalently, one can think of the (size,rate) pairs as (quantity,price) pairs.

up to a single basis point;[9] and, sixth, each quantity bid must be in an increment of $1 million CAD, starting from a lower bound of $5 million CAD and going up to a bidder-specific limit.

The bidding limit for each participant is based on its CPA ratio. Essentially, this ratio is the participant's market share of Canadian dollar deposits; see Bank of Canada (2000). The formula used to determine the bidding limits is the following:

$$\ell_i = \rho_i \times 2.5 \max \left(2 \times 10^9, S\right) \tag{4.1}$$

where S is the auction size, ρ_i is bidder i's CPA ratio, and ℓ_i is bidder i's bidding limit. Therefore, the aggregate bidding limit equals two and one-half times the amount being auctioned or two billion dollars, whichever is greater. Each bidder is assigned a share of this aggregate in proportion to its CPA ratio. A limit is placed on individual bidding in an effort to avoid a market squeeze on available reserves. The two and one-half multiple is set to ensure that the auction has adequate *cover*, which refers to the ratio of bids to total amount being auctioned. A MUA at which inadequate cover exists may result in the failure to place all the units being sold.

5. Econometric Specification and Empirical Results

Below, we implement the theoretical framework of sections 2 and 3 to interpret the data described in the appendix using the institutional rules summarized in section 4.

5.1. Econometric Specification

We believe that the following are important empirical features of the RG auctions that must be incorporated in our empirical model:

1) Each LVTS member appears to have a private value determined by its own cash-reserve needs.

---

[9] A basis point is 0.01 of one percent.

2) The bids are constrained to be on a lattice of 51 points because bids can only be quoted to four decimal points of the interest rate and must be within a band of the target rate plus or minus 25 basis points.[10]

3) For bidder i, quantities are constrained to be on a lattice of points starting at \$5 million and increasing by increments of \$1 million up to the nearest million of ℓ_i.

4) There is a common-value component, which Chapman (2006) has isolated to be the target rate and for which we have data.

We denote by θ the common part to the values of all bidders.

5.1.1. Additional Assumptions concerning Valuations

We assume that bidder i' private component $\boldsymbol{V}_i$ and the common component θ satisfy:

A4. Bidder i's marginal value for quantity q is $(V_{i,q} + \theta)$ for all q.

That is, the common component θ affects all marginal values in the same way. We also assume

A5. θ varies from auction to auction, but is commonly known to the bidders.

Note that, under these assumptions, in equilibrium, all bids will rise and fall, one-for-one, with θ. An implication of this is that, in equilibrium, participant i's bid vector will be of the form $[\boldsymbol{b}_i(\boldsymbol{v}_i) + \theta]$.

Because the good being sold is cash, the assumption of non-increasing marginal valuations seems natural. Assumption A4 imbeds the implicit notion that the distribution of the demand for cash in Canada by LVTS members, after normalizing for changes in the target rate, was stable over the sample period 1 October 2001 through 31 December 2003.

---

[10] This is not strictly true. For while CARS will only accept a bid within this band, in some circumstances, an institution can fax the Bank of Canada with a bid outside this range. In our data, this happened only four times, all in the four days before the fixed action date of November 2001.

Figure 5.1

Difference between the Overnight Rate and the Target Rate
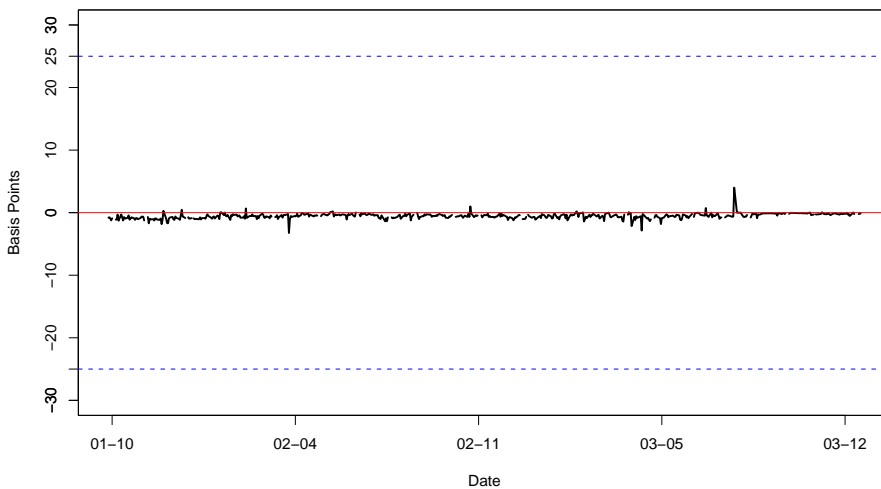
Although we do not directly observe θ, we have an *ex ante* proxy for it, the target rate. One *ex post* measure of θ_t would be the actual interest rate determined in the overnight market on date t. In Figure 5.1, we present a graph of the difference between the actual θ_t, the overnight rate, and our proxy of θ_t, the target rate. This difference has a sample mean very close to and not significantly different from zero. Also, most of the deviations are very small. This evidence suggests that the target rate is a very accurate proxy for θ_t.

5.2. Empirical Results

As mentioned above, in section 4, each auction can have several different tranches for sale. Thus, the RG auctions have both multi-unit and multi-object features: Within a tranche, millions of units (dollars) of cash are available, but funds of different maturities are effectively different objects. In our data, tranches are defined by days to maturity. The most common tranche is for one-day maturities, overnight paper. Nearly two-thirds of all maturities are overnight. Even some of the maturities between two and eight days are effectively overnight: Two- or three-day maturities sold on

Fridays ensure that a bank gets to Monday, or Tuesday in the case of a long weekend. Four-day maturities sold on the Thursday before Good Friday, a statutory holiday in Canada, just get a bank to the Tuesday following Easter Monday, which is also a statutory holiday. Only a small fraction of tranches are for longer-term maturities; *i.e.*, those over three days. For these reasons, we focus exclusively on the overnight market.

In addition to holding RG auctions in both the morning (AM) and the afternoon (PM), the Bank of Canada operated under different policy regimes, one before "fixed action dates" (FAD), which were instituted in the third quarter of 2001, and the other after. Basically, under fixed action dates, the Bank of Canada committed to making changes only on pre-specified "action" dates. We focus on data from the Post-FAD regime.

In order to reduce heterogeneity across auctions, both observed and unobserved, we focused on a subset of our data concerning the LVTS—the post-FAD period, auctions taking place between 1 October 2001 and 31 December 2003. In addition, we focused on the shortest-maturity instrument in the morning auctions that involved a total supply S of between $800 and $1,200 million CAD. As a short-hand, we refer to this sample as the post-FAD, TDAM auctions. We chose this sample for two reasons. First, since bidders know the volume coming into the auction, their bidding behaviour will depend on the volume. It is important that volumes not vary too much. While it would be desireable to focus solely on volumes of, say, $1,200 million CAD, not enough of these existed in our sample. Thus, we expanded the window that met a second consideration: Since the $\lambda_{i,q}$s described in section 3 are to be estimated nonparametrically, we need a relatively large sample of data. Under the $800–$1,200 million CAD window, we used about one-half of the data set, 46 percent. Note, however, that even though the total amount for sale might be $1,200 million CAD, winning bids are typically only ever observed up to $500–$600 million CAD. This is because the auctions are sufficiently well covered so no participant who bids all the way out to his bid limit ever wins at those large quantities. Thus, if a step bid goes

to, say, $1,000 million CAD, then we truncate it to what we can use to estimate the $\lambda_{i,q}$s. In the post-FAD TDAM sample, this is $575 million CAD.

5.2.1. Estimating the $\boldsymbol{\lambda}$-Probabilities

The $\boldsymbol{\lambda}$-probabilities are estimated using a quantity-price grid of observed inverse residual-supply curves. This grid was constructed as follows: For each bidder, a quantity-price grid was constructed; for each auction, the inverse residual-supply curve that the bidder encountered was traced out on this grid of points and one was added to each of the cells in which it fell. This left, for each quantity-bidder combination, a frequency distribution concerning the number of times that the residual supply at a given quantity was observed for a given normalized interest rate. Here, normalized means that we adjusted bids by subtracting out the target rate. In other words, for quantity q, the normalized interest rate r and bidder j, we observed the number of times $n_{q,r,j}$ out of the T that the residual supply for quantity q was observed to be r for bidder j. For a typical quantity-price grid, in this case for quantity 150 ($150 million CAD) and bidder having identification number 40, we obtained the relative frequency distribution presented in Figure 5.2.

Under this grid, the $\boldsymbol{\lambda}$-probabilities were constructed nonparametrically in two related ways—first, using the empirical distribution function (EDF) and, second, by nonparametric smoothing. Using the EDF, $\lambda_{i,q}(b)$ is estimated by

$$\hat{\lambda}_{i,q}(b) = \frac{\sum_{r \leq b} n_{q,r,i}}{\sum_{r=-25bp}^{25bp} n_{q,r,i}} \tag{5.1}$$

where bp denotes *basis point*. In the smoothing case, a discrete version of the Gaussian kernel was employed where we experimented with different bandwidth parameters between 0.001 and 0.08.

The discrete version of the Gaussian kernel was constructed as follows. First, a Gaussian-kernel nonparametric smoother was used in the usual way to create a nonparametric smoothed probability density function of the winning values. Second, the probabilities at the lattice points we were using (*i.e.,* $\pm 25bp$) were kept and all

25

Figure 5.2

Histogram of Inverse Residual-Supply Curve

$q = 150$; **Bidder 40**

Figure 5.3

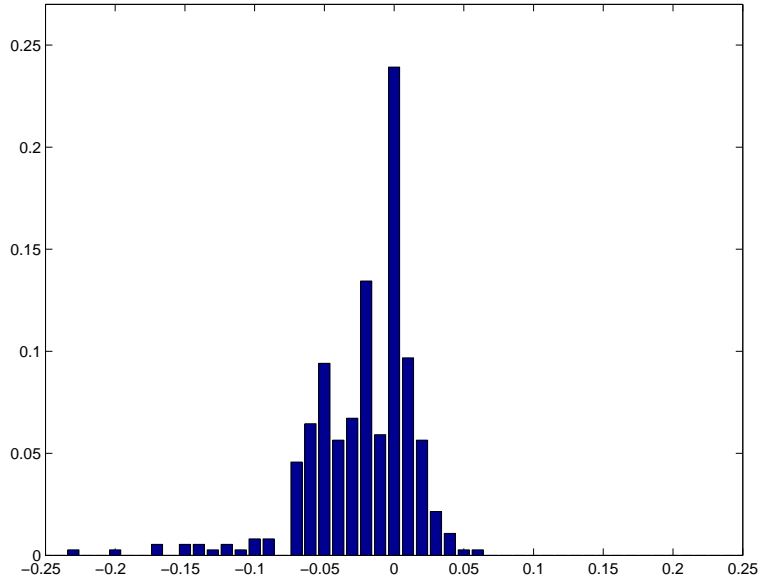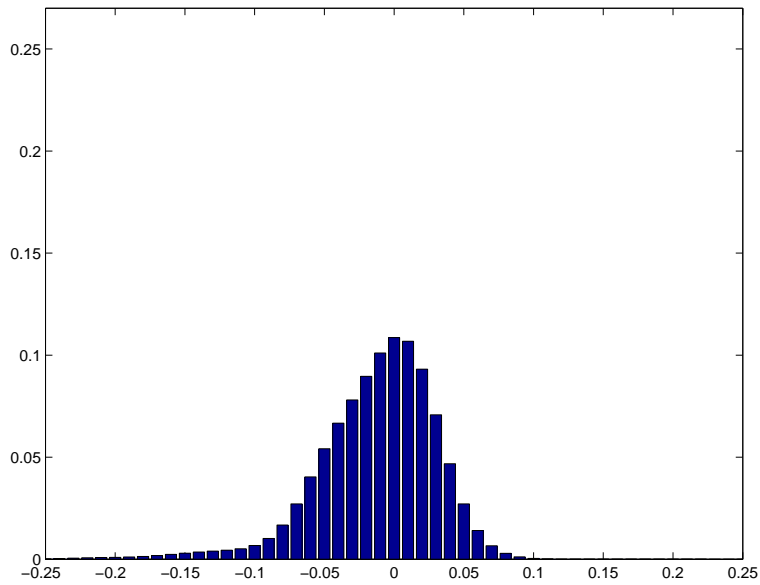**Smoothed Histogram of Inverse Residual-Supply Curve**

$q = 150$; **Bidder 40; Bandwidth = 0.01**

other values were discarded. Third, these were reweighted, by dividing them by their sum. Thus, these 51 values summed to one. Fourth, one minus the cumulative sum of these values was then used as the probability of winning the given quantity. The above weights are weakly positive and sum to one and are, therefore, a valid kernel to conduct this discrete kernel smoothing; see Santer and Duffy (1989). For a typical quantity-price grid, in this case for quantity 150 ($150 million CAD) and bidder having identification number 40, we obtained the "smoothed" relative frequency distribution presented in Figure 5.3.

We chose to conduct our analysis using both a bandwidth of zero (the EDF) and a bandwidth of 0.01. We chose the second bandwidth because it was small enough to introduce only a small amount of bias and because this degree of smoothing eliminated the "flat" parts of the probability of winning which appeared to be the source of the majority of best-response violations. As a check that the bandwidth 0.01 was reasonable, Silverman's rule was calculated for the each of the sample bidder-quantity probability distributions.[11] The average choice of Silverman's rule was 0.0081 with a standard deviation of 0.0019. The number of best-response violations did not change appreciably with a bandwidth of 0.0081.

5.2.2. Sampling Variability

To calculate estimates of the sampling variability of various quantities, we employed the method called *subsampling*; see Politis, Romano, and Wolf (1999). We chose not to use the *bootstrap* because in auction models within the private values paradigm the support of the distribution of bids typically depends on the underlying distribution, violating one of the regularity conditions required to use the bootstrap. Also, we chose not to use the *jackknife* because Efron (1982) has shown that jackknife estimates of quantiles typically do not work well, presumably because the smoothness regularity

---

[11] Note that Silverman's rule, which is derived for continuous random variables, is not necessarily valid theoretically with discrete bids. We include this discussion here because others have expressed interest concerning what Silverman's rule would have indicated and how that would have worked in this case.

condition is violated; see Shao and Tu (1995). Our estimates of the inverse residual-supply functions are based on quantiles. Because subsampling works under mild regularity conditions, when an asymptotic distribution exists, we felt that it was the most robust for our purposes. We chose a block size T_1 equal $T^{\frac{3}{4}}$ where T is 372, and rounded down to 84. $T^{\frac{3}{4}}$ is the appropriate block size when the underlying estimators are $\sqrt{T}$ consistent, which seems reasonable given that our estimates are based on quantiles and sums.

Thus, having calculated the $\hat{\boldsymbol{\lambda}}_T$ using all T observations, we then drew subsamples of size T_1. For each subsample, we then carried out the empirical analysis and then obtained K different estimates $\{\hat{\boldsymbol{\lambda}}_{T_1}^k\}_{k=1}^K$ where the superscript k and the T_1 subscript highlight the fact that this k^{th} estimated quantity is based on a subsample of size T_1 drawn from the observed sample of size T. Based on these subsampled estimates, we then calculated, among other quantities, the losses that obtained when a best response was not used. From these subsamples, we then estimated the sampling variability in our point estimates; these are reported in terms of confidence intervals in Tables 5.6 and 5.7.[12]

5.2.3. Highlights of Empirical Results

One of the most striking features of the data from the LVTS is that participants rarely submitted the four (size,rate) pairs permitted under the rules of the auctions. In fact, the bulk of the participants used just one (size,rate) pair. We illustrate this point by reporting the number of unique pairs submitted in our data set in Table 5.1. Using only one (size,rate) pair does not appear on the surface to be a best response. We should like to investigate how far from a best response submitting one, two, or three bid pairs is.

The most striking feature of Table 5.2 is the large proportion of best-response violations—an average over all bidders of 42 percent, with a range of between 23 and 71 percent, depending on the bidder. Note, however, that the vast majority of these

---

[12] The confidence intervals reported in these tables are one-sided, to be used to test the hypothesis that the mean (median) loss is greater than zero against the null hypothesis that it equals zero.

Table 5.1

Number of Unique Bid Pairs

Pairs	1 Day	2–8 Days
1	3222	1218
2	344	146
3	264	93
4	91	47

violations appear to arise from the jagged nature of non-smoothed estimates of the λs derived from data illustrated by the example in Figure 5.2. Only a small amount of smoothing caused the proportion of best-response violations to fall to between 5 and 34 percent, again depending on the bidder, with an average of about 9 percent. Note, too, there does not appear to be a systematic pattern between the proportion of best-response violations and the step on which they obtained as illustrated in Tables 5.4 and 5.5.

From Tables 5.6 and 5.7, we note that the mean and median lower bound for estimated losses are relatively small. We estimate that the implicit value at risk for an average winning bid to be in tens of thousands of dollars, when calibrated at a daily rate; the comparable average loss, when calibrated at the daily rate, is between \$1.00 and \$826.00 when the EDF estimates are used, but only between \$0.00 and \$35.19 when the smoothed EDF estimates are used. The median estimates are smaller for both methods of estimation; in the smoothed case all of the estimated median losses are zero. The ninety-five percent, one-sided confidence intervals that we have calculated using subsampling are quite small for the EDF estimates, but sometimes large for the smoothed EDF estimates of the mean losses: Smoothing appears to admit outliers that are quite influential in determining our estimated mean losses. Because the median is less influenced by extreme values, our ninety-five percent one-sided confidence intervals for the smoothed EDF median losses are quite small.

6. Summary and Conclusions

Deciding whether bidders at auctions are playing a best response is perhaps one

Table 5.2

Proportion of Violations by Bidder: EDF Estimates

Bidder Number	Percent of Bids
38	0.4696
40	0.2391
41	0.2588
42	0.3699
43	0.5031
52	0.7061
510	0.4886
50	0.4588
296	0.4007
39	0.3214
51	0.6466
46	0.5000
573	0.2813
574	0.2857
350	0.2375
344	0.2632
431	0.4444
Average	0.4191

Table 5.3

Proportion of Violations by Bidder: Smoothed EDF Estimates

Bidder Number	Percent of Bids
38	0.1768
40	0.0516
41	0.1779
42	0.0116
43	0.0497
52	0.0571
510	0.0183
50	0.0118
296	0.0941
39	0.1071
51	0.0733
46	0.1111
573	0.1563
574	0.1429
350	0.0625
344	0.2632
431	0.3333
Average	0.0852

of the most fundamental questions faced by empirical workers employing the structural econometric approach. We have investigated this question in a model of a multi-unit, discriminatory auction, assuming that the true valuations of bidders are unobserved, but that marginal valuations are non-increasing. When we implemented our framework using data from RG auctions, we found that best-response violations are frequent, but for most bidders the median lower bound of the economic distance between the estimated best-response and the bid data is small.

Our main conclusions are the following: While the bidders in this market are extremely skilled, they make small mistakes relatively frequently. Only occassionally, however, do they make a costly mistake. The majority of the mistakes that we measure are small, often less than the price of a cup of coffee when the implicit values at risk are in tens of thousands of dollars. Thus, equilibrium play appears to be a reasonable *approximation* of actual bidder behaviour.

Table 5.4

Violations of Best Response by Step: EDF Estimates

Bidder Number	Step One	Step Two	Step Three	Step Four
38	0.4696	NA*	NA	NA
40	0.2255	0.5000	NA	NA
41	0.2507	0.2500	0.1429	0.0000
42	0.3526	0.4286	NA	NA
43	0.4658	0.1944	0.5000	NA
52	0.4735	0.2966	0.3645	0.3571
510	0.4703	0.3000	0.0000	1.0000
50	0.4118	0.2500	0.5000	NA
296	0.4007	NA	NA	NA
39	0.1786	0.8000	NA	NA
51	0.4310	0.2671	0.2667	0.1395
46	0.4444	0.3333	NA	NA
573	0.1875	0.1667	1.0000	NA
574	0.1429	0.3333	1.0000	NA
350	0.1875	0.2333	0.0000	NA
344	0.1579	0.4000	0.0000	NA
431	0.2222	0.5000	0.0000	0.0000
Average	0.3596	0.2860	0.3115	0.1875

* Here, NA denotes not applicable.

Table 5.5

Violations of Best Response by Step: Smoothed EDF Estimates

Bidder Number	Step One	Step Two	Step Three	Step Four
38	0.1768	NA*	NA	NA
40	0.0489	0.1667	NA	NA
41	0.1698	0.2500	0.1429	0.0000
42	0.0116	0	NA	NA
43	0.0373	0.0556	0.0000	NA
52	0.0245	0.0276	0.0561	0.0000
510	0.0091	0.1000	0.0000	1.0000
50	0.0118	0.0000	0.0000	NA
296	0.0941	NA	NA	NA
39	0.1071	0.0000	NA	NA
51	0.0000	0.0497	0.0762	0.1163
46	0.0556	0.3333	NA	NA
573	0.0938	0.1667	0.4000	NA
574	0.0000	0.3333	1.0000	NA
350	0.0125	0.1333	0.0000	NA
344	0.2105	0.2000	1.0000	NA
431	0.1111	0.5000	0.0000	0.0000
Average	0.0674	0.0665	0.0779	0.0938

* Here, NA denotes not applicable.

Table 5.6

Statistics on Estimated Losses: EDF Estimates

Bidder Number	Non-rationalizable/ Number of Auctions	Mean	95% CI	Median	95% CI
38	85/181	24.21	(0.00, 468.01)	0.00	(0.00, 33.77)
40	88/368	42.97	(0.00, 127.16)	0.00	(0.00, 0.00)
41	96/371	85.26	(0.00, 220.20)	0.00	(0.00, 30.12)
42	64/173	5.89	(0.00, 51.03)	0.00	(0.00, 5.03)
43	81/161	222.84	(0.00, 678.19)	0.00	(0.00, 13.28)
52	173/245	8.84	(0.00, 70.17)	0.24	(0.00, 12.24)
510	107/219	3.22	(0.00, 24.39)	1.05	(0.00, 4.84)
50	39/85	1.62	(0.00, 42.55)	0.00	(0.00, 2.22)
296	115/287	3.08	(0.00, 7.48)	4.61	(0.00, 1.40)
39	9/28	24.64	(0.00, 375.51)	0.00	(0.00, 22.04)
51	150/232	3.96	(0.00, 45.82)	2.08	(0.00, 7.53)
46	9/18	13.57	(0.00, 65.04)	7.23	(0.00, 16.95)
573	9/32	126.79	(0.00, 206.09)	27.38	(0.00, 2.32)
574	4/14	157.15	(0.00, 404.97)	16.81	(0.00, 289.95)
350	19/80	826.00	(0.00, 875.01)	32.70	(0.00, 5.56)
344	5/19	413.67	(0.00, 760.38)	73.21	(0.00, 40.16)
431	4/9	5.87	(0.00, 117.80)	4.84	(0.00, 146.11)
Overall		50.54	(0.00, 427.64)	0.00	(0.00, 16.11)

Table 5.7

Statistics on Estimated Losses: Smoothed EDF Estimates

Bidder Number	Non-rationalizable/ Number of Auctions	Mean	95% CI	Median	95% CI
38	32/181	0.88	(0.00, 152.04)	0.00	(0.00, 0.00)
40	19/368	0.37	(0.00, 2702.41)	0.00	(0.00, 0.00)
41	66/371	35.19	(0.00, 103253.69)	0.00	(0.00, 0.00)
42	2/173	0.00	(0.00, 1.96)	0.00	(0.00, 0.00)
43	8/161	0.12	(0.00, 245.87)	0.00	(0.00, 0.00)
52	14/245	0.02	(0.00, 956.72)	0.00	(0.00, 0.00)
510	4/219	0.01	(0.00, 0.28)	0.00	(0.00, 0.00)
50	1/85	0.00	(0.00, 0.05)	0.00	(0.00, 0.00)
296	27/287	0.18	(0.00, 19.43)	0.00	(0.00, 0.00)
39	3/28	0.39	(0.00, 310.90)	0.00	(0.00, 0.00)
51	17/232	0.31	(0.00, 185.81)	0.00	(0.00, 0.00)
46	2/18	0.40	(0.00, 1.07)	0.00	(0.00, 0.01)
573	5/32	0.81	(0.00, 3224.08)	0.00	(0.00, 0.00)
574	2/14	0.34	(0.00, 1.49)	0.00	(0.00, 0.45)
350	5/80	10.83	(0.00, 25.25)	0.00	(0.00, 0.00)
344	5/19	16.81	(0.00, 2040.92)	0.00	(0.00, 27.75)
431	3/9	0.25	(0.00, 0.54)	0.00	(0.00, 0.52)
Overall		5.47	(0.00, 427.64)	0.00	(0.00, 16.11)

A. Appendix

In this appendix, we describe the construction of the data set used. Our *entire* data set contains 2409 auctions occurring twice a day at the Bank of Canada from 7 February 1999 to 1 January 2004.[13] There are 1222 morning auctions and 1187 afternoon auctions: No afternoon auctions occur on the day before holidays. However, we did not use all of these data.

Of the 2409 auctions, 807 morning auctions and 785 afternoon auctions occurred during the Fixed Action Day (FAD) regime. Of this subset, 326 morning auctions and 321 auctions occurred during the period when the number of bidders was increased from LVTS participants to other bidders who have pledged collateral to the Bank of Canada and the Department of Finance.

A.1. Auction Data

Data concerning the RG auctions were supplied by the Bank of Canada. The format of the data set is especially complex, requiring that we use a relational database management system (RDBMS) to organize it. We used the open-source RDBMS *MySQL* to store and to manipulate the data.

The data concerning the morning and afternoon auctions on a given day are organized as several tables in the *MySQL* database.[14] These tables are in an hierarchy, with multiple entries being associated with a single coarser unit of data. The following simplified example may help illustrate this relationship.

For a given auction, the `tender` table contains an identification number of the tender, the identity of the bidder, the time-stamp of when the bidder submitted the bid, and whether the bid was accepted. The `bid` table contains the identity of the bidder and a tender number, uniquely associating with one entry in the `tender` table. In addition to these identification entries, the `bid` table contains the yield of the bid, the amount of the bid, and the amount that the bidder won. For a given entry in the `tender` table, up to four entries exist in the `bid` table, each uniquely associated with the `tender` table entry, one per submitted (size,rate) pair. Below, we explain in detail the entries in each table.

---

[13] A note concerning vocabulary is warranted. Auction managers at the Bank of Canada define an *auction* to be a time and a place when multiple term-deposits are sold. They define a *tranche* as one of these sales. Therefore, a sale of overnight deposits conducted simultaneously with a sale of three-day term-deposits would be referred to as one auction with two tranches by these managers, but would be referred to as two simultaneous auctions by an economic theorist. In the descriptions that follow, to avoid confusion, we adopt the vocabulary of managers at the Bank of Canada.

[14] A table in a RDBMS may be thought of as a single spreadsheet with a column or combination of columns containing unique data points for all rows. This column or these columns are known as the *index* or *key* of the table. The process through which this index is constructed is known as *normalization*. Information across tables of the relational database is linked by the key.

A.2. Auction Tables

These tables contain an identification number which uniquely identifies the group of auctions as well as data concerning the auctions (*e.g.*, the type of auction—morning collateralized, morning uncollateralized, or afternoon auctions) and the deadline for the submission of bids. The primary key of this table is the `auction id` variable, which is a sequential increasing number. Therefore, any information in later tables that refers to a specific auction will contain the `auction id` key.

A.3. Tranche

The `tranche` table contains the information pertaining to a specific tranche of an auction: The associated `auction id`; the issue date of the term deposit; the maturity date of the term deposit; the type of term, either `1D` for one business day deposits, `2-8D` for deposits less than eight days, or `8+D` for longer deposits; the issue amount; the total amount allotted to bidders; the total amount bid; the lowest, highest, and value-weighted average yields bid; the cutoff or clearing yield; the total amount allotted and bid at the cutoff; and the ratio of allotted and bid amount at the cutoff.

A.4. Tender

The `tender` table contains data corresponding to a specific bidder at a specific auction. A tender is the *package* of (size,rate) pairs that a bidder submits to the Bank of Canada at an auction. An entry in the `tranche` table contains the `auction id` key for the associated auction, the maturity and issue date for the associated tranche, the identifying number for the bidder, the tender number of the bid, the time-stamp of when the bid was submitted, and a status code.

In fact, two tranche tables exist. One table contains the official final bids of the participants; *i.e.*, those bids with a status code `S`, for sold. These are the bids which are used to allocate the deposits to the various participants and which are legally binding. A second table contains all the tenders submitted for the auction. Thus, this latter table, in addition to containing all the official tenders, contains previous bids that bidders entered and withdrew as well as bids which were deemed to be in error or late. For example, if bidder A were to submit a bid at 9:01 a.m. for an auction which closes at 9:15 a.m. and later changes his bid at 9:10 a.m., then the second bid and its time-stamp will be recorded in both the `official tender` table and the `all tender` table, while the former bid would be recorded in the `all tender` table. In another example, suppose bidder A has submitted a bid at 9:01 a.m. for an auction that closed at 9:15 a.m. and later realized that this bid contained an error in the yield bid. Bidder A could then resubmit the bid at 9:10 a.m. changing the error. The latter bid would be recorded in both tables, while the former would be recorded only in the `all tender` table with a status code of `E`, for error. Similarly, cancelled bids are only recorded in the `all tender` table with a status code of `C`, for cancelled.

A.5. Bids

The `bid` table contains information concerning a specific (size,rate) pair of a particular tender at a given auction. An entry in the `bid` table contains the `auction id` of the associated auction, the issue date and maturity date of the associated tranche, the bidder

identification and tender number of the associated tender, these pieces of information along with the bid yield serve to uniquely identify an entry in the `bid` table. In addition, an entry in the `bid` table also contains the amount of the bid in Canadian dollars, the amount that was allotted to this bid, and the amount of the allotted amount that was uncollateralized.

A.6. Bidding Limits

The `bidding limit` table contains the maximum amount that bidders may bid at an auction. These bidding limits are based on the bidders CPA ratio, essentially a bidder's market share of retail bank deposits, and per auction are determined by the following formula

$$\ell_i = \rho_i \times 2.5 \max(2 \times 10^9, S)$$

where ℓ_i is bidder i's bidding limit, ρ_i is their CPA ratio and S is the amount being auctioned. The Bank of Canada sets the total bidding limit to equal either five billion dollars or two and one-half times the amount being auctioned.

A.7. Interest Rates

The interest-rate series used were downloaded from the Bank of Canada's website and consist of the target rate, which is the instrument used by the Bank of Canada in monetary policy, and the overnight rate, which is the average yield in the overnight market weighted by transaction volume.

A.8. Bidders

The eligible bidders at RG auctions consist of the direct clearers in the LVTS and, after September 2002, additional participants who had signed appropriate legal agreements with the Bank of Canada and the Department of Finance. The total number of distinct bidders is sixteen before September 2002 and twenty-four after September 2002.

B. Bibliography

Athey, S. and P. Haile. "Nonparametric Approaches to Auctions." In *Handbook of Econometrics*, Volume 6, edited by J. Heckman, and E. Leamer. Amsterdam: Elsevier, forthcoming.

Bank of Canada. "Proposed Revisions to the Rules Pertaining to Auctions of Receiver General Term Depositis." Bank of Canada, Research Department, discussion paper, 2000.

Brendstrup, B.. "Nonparametric Estimation of Sequential English Auctions." University of Aarhus, Department of Economics, typescript, 2002.

Brendstrup, B. and H. Paarsch. "Nonparametric Estimation in Models of Multi-Unit, Sequential, Dutch Auctions with Asymmetric Bidders." University of Iowa, Department of Economics, typescript, 2004a.

Brendstrup, B. and H. Paarsch. "Semiparametric Estimation of Multi-Object, Sequential, English Auctions." University of Iowa, Department of Economics, typescript, 2004b.

Brendstrup, B. and H. Paarsch. "Identification and Estimation in Sequential, Asymmetric English Auctions." *Journal of Econometrics*, 134 (2006), 69–94.

Cantillon, E. and M. Pesendorfer. "Auctioning Bus Routes: The London Experience." In *Combinatorial Auctions*, edited by P. Cramton, Y. Shoham, and R. Steinberg. Cambridge, MA: MIT Press, 2006a.

Cantillon, E. and M. Pesendorfer. "Combination Bidding in Multi-Unit Auctions." London School of Economics, typescript, 2006b.

Chapman, J.. "The Market Microstructure of the Canadian Overnight Market." University of Iowa. Department of Economics, typescript, 2006.

Donald, S. and H. Paarsch. "Piecewise Pseudo-Maximum Likelihood Estimation in Empirical Models of Auctions." *International Economic Review*, 34 (1993), 121–148.

Donald, S. and H. Paarsch. "Identification, Estimation, and Testing in Parametric Empirical Models of Auctions within the Independent Private Values Paradigm." *Econometric Theory*, 12 (1996), 517–567.

Donald, S. and H. Paarsch. "Superconsistent Estimation and Inference in Structural Econometric Models Using Extreme Order Statistics." *Journal of Econometrics*, 109 (2002), 305–340.

Donald, S., H. Paarsch, and J. Robert. "An Empirical Model of the Multi-Unit, Sequential, Clock Auction." *Journal of Applied Econometrics*, forthcoming.

Efron, B.. *The Jackknife, the Bootstrap, and Other Resampling Plans*. Philadelphia: SIAM, 1982.

Guerre, E., I. Perrigne, and Q. Vuong. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica*, 68 (2000), 525–574.

Haile, P. and E. Tamer. "Inference with an Incomplete Model of English Auctions." *Journal of Political Economy*, 111 (2003), 1–51.

Hortaçsu, A.. "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market." University of Chicago, Department of Economics, typescript, 2002a.

Hortaçsu, A.. "Bidding Behavior in Divisible Good Auctions: Theory and Evidence from the Turkish Treasury Auction Market." University of Chicago, Department of Economics, typescript, 2002b.

Hortaçsu, A. and S. Puller. "Testing Strategic Models of Bidding in Deregulated Electricity Markets: A Case Study of ERCOT." University of Chicago, Department of Economics, typescript, 2004.

Howard, D.. "A Primer on the Implementation of Monetary Policy in the LTVS Environment." *Bank of Canada Review*, (1998), 57–66.

Jofre-Bonet, M. and M. Pesendorfer. "Estimation of a Dynamic Auction Game." *Econometrica*, 71 (2004), 1443–1489.

Kastl, J.. "Discrete Bids and Empirical Inference in Divisible Good Auctions." Northwestern University, Department of Economics, typescript, 2005.

Krishna, V.. *Auction Theory.* San Diego: Academic Press, 2002.

Laffont, J.-J., H. Ossard, and Q. Vuong. "Econometrics of First-Price Auctions." *Econometrica*, 63 (1995), 953–980.

Li, T.. "Econometrics of First-Price Auctions with Binding Reservation Prices." *Journal of Econometrics*, 126 (2005), 173–200.

McAdams, D.. "Isotone Equilibrium in Games of Incomplete Information." *Econometrica*, 71 (2003), 1191–1214.

McAdams, D.. "Monotone Equilibrium in Multi-Unit Auctions." *Review of Economic Studies*, forthcoming.

McAdams, D.. "Partial Identification in Multi-Unit Auctions." MIT, Sloan School, typescript, 2006.

Merrett, D., S. Boisvert, and P. Côté. "Managing the Federal Governent's Cash Balances: A Technical Note." *Bank of Canada Review*, (1995), 55–69.

Paarsch, H.. "Deciding between the Common and Private Value Paradigms in Empirical Models of Auctions." *Journal of Econometrics*, 51 (1992), 191–215.

Paarsch, H.. "Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales." *Journal of Econometrics*, 78 (1997), 333–357.

Paarsch, H. and H. Hong. *An Introduction to the Structural Econometrics of Auction Data*. Cambridge, MA: MIT Press, 2006.

Paarsch, H. and J. Robert. "Testing Equilibrium Behaviour at First-Price, Sealed-Bid Auctions with Discrete Bid Increments." University of Iowa, Department of Economics, typescript, 2003.

Politis, D., J. Romano, and M. Wolf. *Subsampling*. Berlin: Springer–Verlag, 1999.

Santner, T. and D. Duffy. *The Statistical Analysis of Discrete Data*. Berlin: Springer–Verlag, 1989.

Shao, J. and D. Tu. *The Jackknife and Bootstrap*. Berlin: Springer–Verlag, 1995.

Whitesell, W.. "Interest Rate Corridors and Reserves," *Journal of Monetary Economics*, (forthcoming).

Wolak, F.. "Identification and Estimation of Cost Functions using Observed Bid Data: An Application to Electricity Markets." In *Advances in Economics and Econometrics: Theory and Applications*, Volume II, edited by L. Hansen, and S. Turnovsky. New York: Cambridge University Press, 2003, 133–169.