NBER WORKING PAPER SERIES

MARKET POWER AND CAPITAL CONSTRAINTS

Milena Wittwer
Jason Allen

Working Paper 34645
http://www.nber.org/papers/w34645

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026

The views presented are those of the authors and not necessarily those of the Bank of Canada. Any errors are our own. We thank the Triton team at the Bank of Canada, Rui Albuquerque, Eric Budish, David Cimon, Robert Clark, Darrell Duffie, Benjamin Hébert, Amy Huber, Ali Hortaçsu, Mark Egan, Mahyar Kargar, Leonid Kogan, Liran Einav, Matthew Gentzkow, Katja Seim, Daniel Smith, Or Shachar, Monika Piazzesi, Marzena Rostek, Fabricius Somogyi, Xavier Vives, Ji Hee Yoon, Bart Zhou Yueshen, Anthony Zhang, in addition to all participants of talks at Boston College, Chicago Booth, London School of Economics, University College London, University of Maryland, Stanford University, the Annual Meeting of Women in Microstructure, Boston Conference on Markets and Competition, EARIE, EFA, EWMES, FIRS, NBER Market Design, NYU IO day, and SFS Cavalcade. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Market Power and Capital Constraints
Milena Wittwer and Jason Allen
NBER Working Paper No. 34645
January 2026
JEL No. D40, D44, G12, G18, G20, L10

ABSTRACT

We explore how traders' equity capitalization influences asset prices in a framework that accounts for market power. In our model, traders with capital constraints engage in transactions in an imperfectly competitive market. We demonstrate that looser capital constraints elevate both asset prices and price impact, the latter diminishing market liquidity. Using Canadian Treasury auction data, we illustrate how to apply our model to quantify these effects. We estimate the shadow costs of capital constraints by leveraging a temporary policy exemption during 2020-2021. We show that while these constraints are only infrequently binding, their relative impact when activated can be sizable.

Milena Wittwer
Columbia University
Columbia Business School
and NBER
mw3941@columbia.edu

Jason Allen
Bank of Canada
jallen@bankofcanada.ca

Market Power and Capital Constraints*

Milena Wittwer[§] and Jason Allen[+]

December 27, 2025

Abstract

We explore how traders' equity capitalization influences asset prices in a framework that accounts for market power. In our model, traders with capital constraints engage in transactions in an imperfectly competitive market. We demonstrate that looser capital constraints elevate both asset prices and price impact, the latter diminishing market liquidity. Using Canadian Treasury auction data, we illustrate how to apply our model to quantify these effects. We estimate the shadow costs of capital constraints by leveraging a temporary policy exemption during 2020-2021. We show that while these constraints are only infrequently binding, their relative impact when activated can be sizable.

Keywords: Financial intermediaries, market power, price impact, asset demand, asset pricing, government bonds, Basel III, capital requirements, leverage ratios

JEL: G12, G18, G20, D40, D44, L10

1 Introduction

The intermediary asset pricing literature suggests that prices depend on the equity capitalization of financial intermediaries who invest and trade the assets (e.g., Gromb and Vayanos

*Previous versions of the paper were circulated under the titles "Intermediary asset pricing: Capital constraints and market power", and "Intermediary Marker Power and Capital Constraints". The views presented are those of the authors and not necessarily those of the Bank of Canada. Any errors are our own. We thank the Triton team at the Bank of Canada, Rui Albuquerque, Eric Budish, David Cimon, Robert Clark, Darrell Duffie, Benjamin Hébert, Amy Huber, Ali Hortaçsu, Mark Egan, Mahyar Kargar, Leonid Kogan, Liran Einav, Matthew Gentzkow, Katja Seim, Daniel Smith, Or Shachar, Monika Piazzesi, Marzena Rostek, Fabricius Somogyi, Xavier Vives, Ji Hee Yoon, Bart Zhou Yueshen, Anthony Zhang, in addition to all participants of talks at Boston College, Chicago Booth, London School of Economics, University College London, University of Maryland, Stanford University, the Annual Meeting of Women in Microstructure, Boston Conference on Markets and Competition, EARIE, EFA, EWMES, FIRS, NBER Market Design, NYU IO day, and SFS Cavalcade. Correspondence to: [§]Milena Wittwer (Columbia University): mw3941@columbia.edu, and [+]Jason Allen (Bank of Canada and University of Wisconsin School of Business): jallen39@wisc.edu

(2002); Brunnermeier and Pedersen (2009); He and Krishnamurthy (2012, 2013); Brunnermeier and Sannikov (2014)). In this literature, intermediaries typically face funding or capital constraints and execute trades in perfectly competitive markets. In practice, however, intermediaries and other large traders enjoy market power—as documented for various trade settings, including Treasury, repo, foreign exchange, mortgage-backed securities, and equity securities lending markets (e.g., Hortaçsu et al. (2018); Allen and Wittwer (2023); An and Song (2023); Chen et al. (2023); Huber (2023); Pinter and Üslü (2023); Wallen (2023)).

Our contribution is to study how equity capitalization affects asset prices and to quantify the effects in a framework that allows for market power (as in Wilson (1979); Klemperer and Meyer (1989); Kyle (1989); Vives (2011); Rostek and Weretka (2012); Rostek and Yoon (2021); among others).[1] We introduce a model in which capital-constrained traders buy or trade an asset in an imperfectly competitive market, and we estimate it with data on Canadian Treasury auctions. Differing from much of the existing asset pricing literature, but in line with the banking literature, our emphasis lies on capital constraints that prevent equity from falling below a specified ratio to the assets on the balance sheet, in accordance with Basel III standards. Since the last financial crisis, this type of constraint has risen in significance, becoming a primary concern for traders, as evidenced by surveys of market participants and regulatory reports (e.g., CGFS (2014); CGFS (2016); ESRB (2016); Group of Thirty (2021); ISDA (2024)), and public commentary (e.g., Baer (2020); Popowicz (2021); Duffie (2023)).[2]

In the model, presented in Section 2, traders compete to buy or trade multiple units of an asset of uncertain supply. They are risk averse and subject to a capital constraint, which depends on the auction outcome and the amount of capital they will need to hold post-auction. They have private information about the asset's return. The market is imperfectly competitive, so that each trader has price impact; it clears via one of two auction formats,

[1]Capital requirements aim to strengthen the risk management of banks and avoid the build-up of systemic risks. Our analysis does not incorporate how these risks change when relaxing constraints.

[2]Chart 15 in ESRB (2016) lists capital requirements as the main friction affecting market-making, above other types of constraints that have achieved more attention in the academic literature, for example, value-at-risk constraints. Table 1 in CGFS (2016) provides a summary of private sector views about the costs of regulation, including Basel III. Group of Thirty (2021) recommends modifications of the leverage ratio to improve resiliency of the U.S. Treasury market.

which represent different financial markets, including primary auctions and exchanges. In the benchmark model, dealers submit decreasing net-demand functions that specify how much they are willing to pay for different units of the asset; the market clears at the price at which aggregate dealer net-demand meets supply, and each dealer wins the amount they asked for at that price (uniform price auction). In the extended model, winning dealers pay the prices they bid (discriminatory price auction).

Solving for an equilibrium in this environment is challenging because point-wise maximization —a common approach in the literature—does not work when bidders face outcome-dependent constraints. Instead, we must consider all feasible net-demand functions. By doing so, we derive necessary conditions for symmetric Bayesian Nash Equilibria (hereinafter referred to as equilibria) for all auction environments we consider. We show that traders who face capital constraints behave as if they were bidding in an auction without constraints, only that their willingness to pay is discounted by the shadow cost of the capital constraint. For the uniform price auction, which is the standard auction format in the market-microstructure literature, we prove that there is a unique symmetric equilibrium with linear net-demand functions, and derive its functional form.

We show that when capital constraints are eased, not only does the asset price increase, but so does the price impact across all traders—an effect that does not arise in models with perfect competition. It occurs because traders discount their willingness to pay (which declines with quantity due to risk aversion) by the shadow cost of the constraint. Relaxing the constraint lowers the shadow cost, which not only increases the overall willingness to pay but also steepens the curve. To understand how this affects price impact, consider the decision of a single trader, assuming all other traders naively submit discounted willingness to pay curves. As these curves steepen, the possible market-clearing points—determined by how much supply is left given aggregate competitor demand at each price—lie along a steeper (residual supply) curve. This means that it takes a smaller trade amount to change the clearing price by one dollar: price impact rises and liquidity declines. In equilibrium, all traders account for their own price impact, which introduces an additional effect from strategic demand reduction (as in Ausubel et al. (2014)). In uniform price auctions, this strategic component reinforces the direct effect, whereas in discriminatory price auctions, it

may partially offset it.

To demonstrate how to quantify the price and price impact effect of capital constraints with our framework, we use data on Canadian Treasury auctions in Section 3. Our data combine bidding details from all Canadian government bond auctions between January 2019 and February 2022 with balance sheet information from the eight largest bidders—primary dealers. We can identify each bidder thanks to unique identifiers, and observe all submitted bids, in addition to the quarterly Basel III Leverage Ratio (LR) at the holding company level (following He et al. (2017)). The LR is the ratio between a bank's capital and assets. It must be above a regulatory threshold and is considered a relevant capital constraint when trading government bonds (CGFS (2016)). Lastly, we collect data on all secondary market trades by dealers to assess the volatility of returns from buying bonds at auction and selling them in the secondary market.

We estimate the key parameters of the model: dealer risk aversion and the shadow costs of the capital constraint. To accomplish this, we employ estimation techniques from the auctions literature (introduced by Guerre et al. (2000); Hortaçsu and McAdams (2010); Kastl (2011)) to estimate each bidder's willingness to pay at a discrete number of points. Then we fit the model-implied functional form for the willingness to pay through these points. Finally, we take advantage of a temporary exemption of domestic government bonds from the leverage ratio during the COVID-19 pandemic to identify the degree of dealer risk aversion and their shadow costs of the capital constraint by analyzing how the willingness to pay varies around the policy change.

Our findings reveal that dealers exhibit a moderate level of risk aversion, and that the shadow costs of constraints vary significantly across auctions. In a typical auction the constraint is not binding. The distribution of shadow costs, however, demonstrates pronounced tails, indicating that while constraints are only occasionally binding, their impact when activated is severe. This observation underscores the importance of examining the entire distribution of shadow costs, aligning with the notion that violations of these constraints are rare but high-cost events.

A back-of-the-envelope calculation indicates that a 1% decrease in the shadow cost of capital constraints might lead to an average increase in both market prices and price impact

by approximately 0.19%. Therefore, completely removing these constraints could raise the average price and price impact by 19%, due to rare, costly events. Given that Canadian Treasury auctions are liquid, independent of capital constraints, the absolute impact of modifying the constraint on price impact is small. In less liquid markets, even a slight change in price impact can result in an economically meaningful absolute change. The effect becomes more pronounced with fewer competing traders and heightened levels of risk aversion or return volatility.

Given the generality of our framework, it offers a tool for analyzing the impact of Basel-III style capital constraints on trading, pricing, and market liquidity across various markets characterized by a finite number of sophisticated traders. Examples include markets for foreign exchange, repo, mortgage-backed securities, equity securities lending, and government asset purchase programs. Understanding the influence of capital constraints on prices and their interaction with market power across diverse market settings is crucial for ongoing policy discussions, particularly regarding potential adjustments to Basel III requirements following the events of the COVID-19 pandemic (Group of Thirty (2021); ISDA (2024)). More broadly, the significance of market power is inherently an empirical inquiry, vital for policymakers concerned about the unintended consequences of policy interventions, such as relaxing capital constraints.

Related literature. By analyzing how capital constraints affect asset prices and price impact or market liquidity, we contribute to several strands of the theoretic and empirical literature.

A substantial body of the theory literature, surveyed by Vayanos and Wang (2012), examines the effects of various market frictions—such as imperfect competition and funding constraints—on market prices and liquidity. Studied in isolation, the typical result is that individual frictions tend to raise price impact and lower liquidity. We find that funding constraints can actually decrease price impact and enhance liquidity when traders have market power. This outcome arises not just from examining the interaction between two frictions, but also because we investigate a different type of constraint. While the market microstructure literature focuses on margin or value-at-risk constraints (following Gromb and Vayanos (2002); Brunnermeier and Pedersen (2009); Pedersen and Gârleanu (2011)), we

examine capital constraints, motivated by Basel III standards. In this regard we are closer to a large banking literature (e.g., Begenau (2020); Benetton (2021); Jamilov (2021); Corbae and D'Erasmo (2021); Begenau and Landvoigt (2022); Wang et al. (2022)) which analyses how capital constraints affect different aspects of the economy, rather than analyzing its effect on trading.

Our theory also relates to the ample intermediary asset pricing literature that examines the impact of capitalization of financial intermediaries, who trade and invest on behalf of households, on asset price behavior due to constraints on debt (e.g., Brunnermeier and Pedersen (2009)), or constraints on equity (e.g., He and Krishnamurthy (2013, 2012); Brunnermeier and Sannikov (2014)). The key difference relative to these macroeconomic models is that we zoom in on the market in which large traders, such as intermediaries, interact, and allow them to impact prices as a result of market power.

The market clears via a multi-unit auction following Wilson (1979), Kyle (1989), and Klemperer and Meyer (1989), and more recent literature, overviewed by Rostek and Yoon (2020). Our innovation is the introduction of bidder constraints that are dependent on the auction outcome. While we focus on capital constraints, our methods to characterize equilibria generalize to auctions with other types of constraints, including budget constraints, and quantity or price caps.

Our empirical analysis adds to the growing empirical literature that investigates the impact of different costs or constraints faced by traders and intermediaries on asset prices and market liquidity (e.g., Adrian and Shin (2010); Adrian et al. (2014); He et al. (2017, 2022); Du et al. (2018, 2023a,b); Haddad and Muir (2021); Siriwardane et al. (2025)). There is increasing evidence that Basel III capital requirements are having a negative impact on bank-intermediated trades (e.g., Boyarchenko et al. (2020); Du et al. (2018); Cenedese et al. (2021); Wallen (2023)). Our paper highlights a positive effect on liquidity of capital requirements, which adds to the mixed evidence on how constraints affect liquidity (e.g., Adrian et al. (2017); Anderson and Stulz (2017); Trebbi and Xiao (2017); Bessembinder et al. (2018); Breckenfelder and Ivashina (2021)). Existing studies rely on proxy variables to capture intermediary costs or constraints, such as spreads or VIX, while we directly estimate the shadow cost of the capital constraint.

For estimation, we adopt techniques from the literature on multi-unit auctions, developed by Guerre et al. (2000), Hortaçsu and McAdams (2010), and Kastl (2011) and extended by Hortaçsu and Kastl (2012) and Allen et al. (2020, 2024), among others. This literature commonly assumes that financial institutions are risk-neutral. This stands in contrast to the related market microstructure literature which builds on Kyle (1989) and assumes that financial institutions are risk averse. We follow this literature and estimate risk aversion. To do this, we impose a functional form on preferences to circumvent the impossibility result by Guerre et al. (2009) that one cannot non-parametrically identify risk aversion (in first-price auctions)—similar to a handful of papers that estimate risk aversion in auctions for procurement, timber, and other non-financial goods (e.g., Campo et al. (2011); Bolotnyy and Vasserman (2023); Häfner (2023); Luo and Takahashi (2023)).[3] The auction approach complements the common macroeconomic practice of calibrating risk aversion for households using Euler equations. Given the significant role that intermediary risk aversion plays in intermediary asset pricing models, our estimates offer valuable insights for calibrating these models.

2 Model

Our goal is to study how prices and price impact change when capital constraints are relaxed (or tightened) and traders have market power. Consistent with the theoretical literature, we model market clearing via a uniform price auction, in which winning traders pay the market clearing price. In Appendix A, we extend the model to discriminatory price auctions, which is more complicated, but relevant for our empirical analysis.

The market may be one-sided, meaning that traders buy but not sell, or double-sided, so that traders buy and sell. In practice, primary markets, for instance, are typically one-sided, while trading on an exchange can be approximated via a double-sided uniform price auction, where packages of limit orders form net-demand schedules (e.g., Kyle (1989)). In order to facilitate the comparison with the empirical analysis, we present our framework using a one-sided market but explain how to adjust it to represent a double-sided market in Appendix

[3] Gupta and Lamba (2017) highlights the importance of risk aversion in Treasury auctions, but does not estimate it.

A. Proofs are in Appendix C. Random variables are highlighted in **bold**.

2.1 Benchmark

There are $N > 2$ traders who compete for units of an asset in an auction. When there are finitely many traders, each one has some market power in that they can impact the market clearing price. When $N \to \infty$ the market is perfectly competitive.

Total supply $\boldsymbol{A}$ is random; it is drawn from some continuous distribution with support $(0, \overline{A}]$ where $\overline{A} \in \mathbb{R}^+$, and has a strictly positive density, $\phi(\cdot)$. In our empirical application, supply is random because traders don't know the issuance size when they compete. In other settings, the supply might be random due to noise traders.

Each trader seeks to maximize the payoff they expect to earn from buying in the auction. Ex-post, when the market clears at price P^c and the trader wins amount a_i^c, the payoff is the gross utility they achieve from the asset, $V(a_i^c, \epsilon_i)$, minus the amount they have to pay, $P^c a_i^c$. Beyond the amount of the asset, gross utility depends on the trader's private signal, ϵ_i. The signal is independently drawn across traders from a distribution with bounded support, $[\underline{\epsilon}, \overline{\epsilon}]$, and strictly positive density, $\psi(\cdot)$. When $\underline{\epsilon} = \overline{\epsilon} = \epsilon_i = 0$ with $\psi(\epsilon_i) = 1$ for all traders i, all of them observe the same signal, generating an auction environment without private information.

For most of our analysis we follow Vives (2011), Rostek and Weretka (2012), Chen and Duffie (2020), Rostek and Yoon (2021), among others, and assume that utility is linear-quadratic:

$$V(a_i^c, \epsilon_i) = (\mu + \epsilon_i) a_i^c - \frac{\rho\sigma^2}{2}[a_i^c]^2, \tag{1}$$

with commonly known parameters $\mu > 0, \sigma > 0$, and $\rho > 0$. This enables us to derive an equilibrium with linear demand functions in closed-form, because the trader's willingness to pay for amount a, $v(a, \epsilon_i) = \frac{\partial V(a,\epsilon_i)}{\partial a}$, is linear:

$$v(a, \epsilon_i) = (\mu + \epsilon_i) - \rho\sigma^2 a. \tag{2}$$

Imposing utility (1) is similar to postulating mean-variance preferences (c.f., Wang and Zender (2002)). Moreover, when assuming that the asset pays an unknown gross return

of $\boldsymbol{R_i} \sim N(\mu + \epsilon_i, \sigma^2)$ to trader i, assuming utility function (1) is typically equivalent to imposing utility with constant absolute risk aversion (CARA) with risk-aversion $\rho > 0$ (as in Kyle (1989)).[4] To highlight the connection, we refer to σ^2 as return variance, and ρ as risk aversion—short hand for the trader's overall risk-bearing capacity, which could be shaped by a variety of factors.

To prepare for the moment at which they receive signal ϵ_i, each trader determines demand schedules, $p_i(\cdot, \epsilon_i) : \mathbb{R}^+ \to \mathbb{R}^+$, for all possible signals they might observe. Each demand schedule specifies the price, $p_i(a, \epsilon_i)$, a trader is willing to pay for amount a, given signal ϵ_i. Function $p_i(\cdot, \cdot)$ is strictly decreasing and twice continuous in quantity, and continuous and bounded in the signal. We denote the set of functions with these properties by $\mathcal{B}$, and the inverse with respect to quantity by $a_i(p, \epsilon_i) = p_i^{-1}(p, \epsilon_i)$.

When determining their demand schedules, each trader faces a capital constraint, according to which banks must hold sufficient equity capital relative to their total balance sheet exposure:

$$\mathbb{E}[\boldsymbol{\theta_i} - \kappa \boldsymbol{P^c a_i^c}] \geq 0. \tag{3}$$

Here $\boldsymbol{\theta_i}$ is trader i's equity capital position. It is drawn iid from some continuous distribution with strictly positive density on bounded support. $\boldsymbol{P^c a_i^c}$ is the total balance sheet exposure of the trader, where $\boldsymbol{P^c}$ is the market clearing price, and $\boldsymbol{a_i^c}$ is the amount the trader wins. Capital threshold $\kappa > 0$ is commonly known, for example, 3% according to Basel III.

The constraint must hold in expectation, rather than for any specific signal draw, ϵ_i. This captures the feature that it is imposed by a regulator who has less information about the trader's beliefs and demand than the trader themselves. Furthermore, the constraint

[4]The equivalence holds in the standard setting, where an equilibrium can be found via pointwise maximization, for each fixed price P^c. Specifically, maximizing expected CARA utility, $\mathbb{E}[1 - \exp(-\rho \boldsymbol{\omega}(a_i^c, P^c))]$ with respect to a_i^c over future wealth $\boldsymbol{\omega}(a_i^c, P^c) = (\boldsymbol{R_i} - P^c)a_i^c$ that is generated by holding amount a_i^c at price P^c is equivalent to maximizing $V(a_i^c, \epsilon_i)$, given Normally distributed returns $\boldsymbol{R_i} \sim N(\mu + \epsilon_i, \sigma^2)$. In our model, the mathematical equivalence breaks down when constraints bind. The equivalence would continue to hold if, instead of modeling the actual constraint, we imposed an exogenous cost term, $\lambda \kappa a^c P^c$, with $\lambda \kappa \geq 0$. We avoid this approach because an exogenous cost term cannot capture the endogenous nature of the shadow cost, including its dependence on equity capital and the capital threshold.

is imposed ex-ante, prior to market clearing, and therefore may be violated ex-post. This reflects the fact that, in reality, capital requirements must be met on average over a longer time horizon than one instantaneous moment of market clearing—the trading game is played many times. The expected equity and balance sheet exposure of a trader reflect their average positions over this longer time horizon.

When all traders have observed their signals, ϵ_i, they submit the demand schedule that corresponds to the observed signal, $p_i(\cdot, \epsilon_i)$. The timing assumption—that a trader first selects their demand function for all potential signals before observing their signal—reflects the reality that traders aim to include the most recent private information in their trading decisions, while calculating optimal trading strategies is intricate and time-consuming.[5] Technically, the timing assumption is similar to a common assumption in Bayesian persuasion games (Kamenica and Gentzkow (2011)) where players on one side of the market commit to a signal realization space and a corresponding profile distribution over that space, prior to observing their private signal. It makes the model tractable.[6] In the special case in which all traders are equally informed ($\epsilon_i = \epsilon$ for all i), the timing assumption is without bite. In that case, we could equivalently assume that all bidders first observe their shared signal, and then choose their bidding strategies.

When all traders have submitted their demand, the auction clears at the price, P^c, such that aggregate demand meets total supply: $P^c : \sum_i a_i(P^c, \epsilon_i) = A$. Each trader pays the market clearing price, $P^c = p_i(a_i^c, \epsilon_i)$, for the amount won, $a_i^c = a_i(P^c, \epsilon_i)$ at that price. To highlight equilibria, we refer to the equilibrium market clearing price by P^* and the winning amount by a_i^*.

[5]In our empirical application, Treasury auctions, traders observe customer orders placed close to the auction closing time. To incorporate this information, traders predefine bidding strategies for all contingencies and adjust their bids accordingly when the orders arrive. This approach is also common in high-frequency trading environments where information arrives from various sources.

[6]In the absence of a capital constraint, selecting $p_i(\cdot, \cdot)$ to maximize the expected payoff $\mathbb{E}[V(\boldsymbol{a_i^c}, \boldsymbol{\epsilon_i}) - \boldsymbol{P^c a_i^c}] = \mathbb{E}[\mathbb{E}[V(\boldsymbol{a_i^c}, \boldsymbol{\epsilon_i}) - \boldsymbol{P^c a_i^c}|\boldsymbol{\epsilon_i}]]$, and then submitting the demand function $p_i(\cdot, \epsilon_i)$ based on the observed signal ϵ_i, is equivalent to maximizing the expected payoff conditional on the signal, $\mathbb{E}[V(\boldsymbol{a_i^c}, \epsilon_i) - \boldsymbol{P^c a_i^c}|\epsilon_i]$, with respect to $p_i(\cdot, \epsilon_i)$. This aligns with the standard timing assumption in the auction literature in the absence of constraints. However, when accounting for the constraint, the trader does not condition their bidding strategy on any particular signal draw to correctly account for the fact that the capital constraint is imposed in expectation over signals, rather than conditional on any particular signal draw.

Summarizing, the timing of events is as follows. First, each trader chooses a demand schedule $p_i(\cdot,\cdot)$ for all possible realizations of signals by maximizing their expected payoff at market clearing subject to the capital constraint:

$$\max_{p_i(\cdot,\cdot)\in\mathcal{B}} \mathbb{E}\big[V(\boldsymbol{a_i^c},\boldsymbol{\epsilon_i}) - \boldsymbol{P^c a_i^c}\big] \text{ subject to}$$
$$\text{the capital constraint:} \qquad \mathbb{E}\big[\boldsymbol{\theta_i} - \kappa \boldsymbol{P^c a_i^c}\big] \geq 0,$$
$$\text{and market clearing:} \qquad \boldsymbol{P^c} = p_i(\boldsymbol{a_i^c},\boldsymbol{\epsilon_i}). \tag{4}$$

Second, each trader observes their signal and submits the corresponding demand, $p_i(\cdot,\epsilon_i)$. Last, supply and equity positions realize, the market clears, and all transactions take place.

We focus on symmetric equilibria since traders are ex-ante identical. Later, we restrict attention to equilibria with linear demand functions, which is common in the related literature (recently reviewed by Rostek and Yoon (2020)). We will establish that, within this restricted class of symmetric linear equilibria, the equilibrium is unique. We do not make claims about uniqueness within the broader class of symmetric equilibria or the full set of possible equilibria.[7]

Definition 1. *A symmetric Bayesian Nash equilibrium is a collection of demand functions $p^*(\cdot,\cdot)$ that for each trader maximizes expected surplus subject to the capital constraint and market clearing. An equilibrium is linear if $\frac{\partial p^*(a,\epsilon_i)}{\partial a}$ is constant at all a.*

To derive equilibrium conditions, take the perspective of trader i—who solves optimization problem (4)—and assume that all other traders follow the equilibrium strategy. Substituting the market-clearing constraints, $\boldsymbol{P^c} = p_i(\boldsymbol{a_i^c},\boldsymbol{\epsilon_i})$, into the maximization problem, we obtain the following Lagrangian with Lagrange multiplier $\lambda \geq 0$:

$$\max_{p_i(\cdot,\cdot)\in\mathcal{B}} \mathbb{E}\big[V(\boldsymbol{a_i^c},\boldsymbol{\epsilon_i}) - (1+\lambda\kappa) p_i(\boldsymbol{a_i^c},\boldsymbol{\epsilon_i})\boldsymbol{a_i^c}\big].$$

This optimization problem mirrors the trader's unconstrained problem, with the key difference being an additional cost term, $\lambda\kappa p_i(\boldsymbol{a_i^c},\boldsymbol{\epsilon_i})\boldsymbol{a_i^c}$, which equals $\lambda\kappa \boldsymbol{P^c a_i^c}$ by market clearing.

[7] In some auction environments without private information stronger uniqueness results are possible—for instance, uniqueness within the full class of symmetric equilibria (see Glebkin et al. (2023).

This term implies that the effective price the trader faces per unit consists of the market price plus a shadow cost associated with the capital constraint—similar to an ad valorem tax, $(1 + \lambda\kappa)\boldsymbol{P^c}$.

Proposition 1 characterizes how this cost term influences bidding behavior. To emphasize the generality of the result, we establish it for a broader class of utility functions than the linear-quadratic form (1).

Proposition 1. *Let gross utility $V(a, \epsilon_i)$ be measurable, bounded, continuously differentiable and strictly concave in a for all ϵ_i, and bounded in ϵ_i for all a.*

In any symmetric equilibrium, traders behave as if they were bidding in an auction without capital constraints, in which their willingness to pay is

$$\tilde{v}(a, \epsilon_i) = \frac{v(a, \epsilon_i)}{1 + \lambda\kappa} \tag{5}$$

instead of $v(a, \epsilon_i)$, where $\lambda \geq 0$ is the Lagrange multiplier of the capital constraint. Trader i submits demand function, $p^*(\cdot, \cdot)$, that for all p, a, and ϵ_i satisfies

$$p = \tilde{v}(a, \epsilon_i) - \underbrace{a \left(\frac{\partial G(a, p|\epsilon_i)}{\partial a} \Big/ \frac{\partial G(a, p|\epsilon_i)}{\partial p} \right) (-1)}_{shading}, \tag{6}$$

where $G(a, p|\epsilon_i) = \Pr(\boldsymbol{a_i^c} \leq a|\epsilon_i)$ with $\boldsymbol{a_i^c} = \boldsymbol{A} - \sum_{j \neq i} a^*(p, \boldsymbol{\epsilon_j})$ is the probability that trader i, who bids price $p = p^*(a, \epsilon_i)$, wins less than a at market clearance, given that the other traders play the equilibrium strategy.

The main insight of Proposition 1 is that the trader bids as if participating in a standard uniform price auction without capital constraints, where the marginal utility, $v(a, \epsilon_i)$, is discounted by the shadow cost of the capital constraint, $\lambda\kappa$, which is strictly positive when the capital constraint is binding, and zero otherwise.[8] As in auctions without constraints, strate-

[8]The shadow cost is common across all traders (within an auction). In our empirical application, the premise likely holds because traders, acting as dealers, have access to a relatively frictionless inter-dealer market that allows them to engage in arbitrage. This logic also extends to other scenarios where traders can mitigate their exposures in a separate market than the one we model, or at another time. While our model does not explicitly simulate the arbitrage process, it generates shadow costs consistent with this concept.

gic traders shade their willingness to pay to reduce payments—unless the market is perfectly competitive—with the extent of shading depending on the auction rules, as discussed by Clark et al. (2021).

To explain what determines shading in uniform price auctions, we solve for the equilibrium in closed-form. This relies on the assumption that gross utility takes a linear-quadratic functional form (1).

Proposition 2. *Let gross utility $V(a, \epsilon_i)$ be linear-quadratic, i.e., given by expression (1), and $\tilde{v}(a, \epsilon_i)$ be given by expression (5).*

There exists a unique symmetric linear equilibrium in which each trader submits

$$p(a, \epsilon_i) = \tilde{v}(a, \epsilon_i) - a\Lambda, \ \textit{with} \tag{7}$$

$$\Lambda = \frac{1}{N-2} \frac{\rho\sigma^2}{1+\lambda\kappa}, \ \textit{where} \tag{8}$$

$$\lambda = \begin{cases} 0 & \textit{if } \mathbb{E}[\boldsymbol{\theta_i}] \geq \tilde{\theta} \\ \frac{1}{\kappa}\left(\frac{\tilde{\theta}}{\mathbb{E}[\boldsymbol{\theta_i}]} - 1\right) > 0 & \textit{otherwise,} \end{cases} \quad \textit{where } \tilde{\theta} = \frac{\kappa}{N}\left(\mathbb{E}[\boldsymbol{A}](\mu + \mathbb{E}[\boldsymbol{\epsilon_i}]) - \rho\sigma^2\left(\frac{N-1}{N-2}\right)\frac{\mathbb{E}[\boldsymbol{A^2}]}{N}\right).$$

We know from Proposition (1) that in equilibrium the trader submits a bid p for amount a that equals their discounted willingness to pay minus a shading factor. In uniform price auctions the shading factor is determined by the bid's price impact, Λ—commonly referred to as Kyle's lambda, which is constant in linear equilibria. It captures the change in the market-clearing price when trader i changes their demand marginally. To see why, suppose the trader bids their full discounted willingness to pay at quantity a. If they shade their bid slightly, they may win slightly less, but they also decrease the clearing price by Λ. Since the clearing price applies to all units won, the total cost savings are $a\Lambda$.

The inverse of price impact, $1/\Lambda$, is a standard measure of liquidity in exchange markets, as it represents the order size required to move the market price by one dollar (Kyle (1985); Vayanos and Wang (2013)). When the market becomes perfectly competitive ($N \to \infty$), price impact converges to zero and liquidity infinite, i.e., perfect. The key feature in our setting is that price impact, Λ, and therefore liquidity depends not only on the number of traders, N, and the curvature of gross utility—captured by $\rho\sigma^2$ in the linear-quadratic

case—but also on the shadow cost of the capital constraint, $\lambda\kappa$.

2.2 How capital constraints affect the price and price impact

The main prediction of the model is about what happens to prices and price impact when the capital constraint is relaxed, for instance, because the minimal capital threshold decreases.

Corollary 1. *When the capital constraint is relaxed, price and price impact increase.*

To build an intuition for Corollary 1, assume for a moment that the market is perfectly competitive ($N \to \infty$). This market is perfectly liquid, and no trader has price impact ($\Lambda \to 0$), implying that all traders are price-takers who bid their discounted willingness to pay. This does not change when relaxing the constraint. Relaxing the constraint, however, still has an effect on the price. Since the capital constraint depends on the nominal (rather than real) value of the amount the trader wins at market clearance, the effective per-unit price each trader faces, $(1 + \lambda\kappa)p$, decreases as the shadow cost $\lambda\kappa$ declines. Consequently, the trader's discounted willingness to pay (5) becomes steeper. This leads to an increase in the market-clearing price—unless offset by a supply response—as illustrated in Figure 1a.[9]

Now consider an imperfectly competitive market ($N < \infty$), and take the perspective of trader i. To decompose the effects at play, assume that all other traders are naive; they do not internalize that their own bid impacts the price (since $\Lambda > 0$), and therefore falsely submit their full discounted willingness to pay. Naive traders behave like price-takers, and therefore submit steeper demand curves as the constraint is relaxed. For trader i, this means that the prices at which the market can clear align along a steeper curve: the residual supply curve $RS_i^{-1}(\cdot)$, as shown in Figure 1b. This curve reflects the available supply net of the aggregate demand from all other traders for a given draw of signals and supply. In the price-quantity space it is defined as:

$$RS_i(p) = A - \sum_{j \neq i} a^*(p, \epsilon_j). \tag{9}$$

[9]This prediction is in line with He and Krishnamurthy (2012, 2013) and Brunnermeier and Sannikov (2014). In their models, a positive shock to a trader's net worth, i.e., equity capital, increases its risk-bearing capacity, which leads to higher asset prices. In our model, risk aversion is constant.

Figure 1: Change in demand and price impact when capital constraints are relaxed

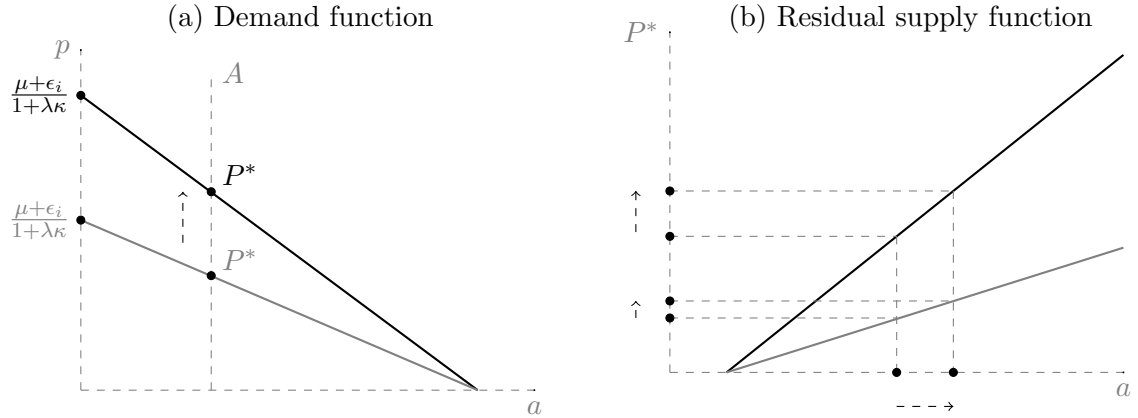

Figure 1a illustrates how the trader's demand function, $p^*(\cdot, \epsilon_i)$—which coincides with their discounted willingness to pay, $\tilde{v}(\cdot, \epsilon_i)$, under perfect competition—shifts when capital constraints are relaxed. Figure 1b shows how the residual supply curve, $RS_i^{-1}(\cdot)$, for a given realization of supply and signals, changes with looser capital constraints. The slope of the residual supply curve reflects the price impact. In both panels, the initial curves are shown in gray, and the curves under relaxed constraints appear in black.

In the quantity-price space, a steeper residual supply curve, $RS_i^{-1}(\cdot)$, implies that a smaller adjustment in trader i's demand—corresponding to movement along a steeper residual supply curve—is sufficient to shift the market price by one dollar. In other words, price impact, which is defined as the amount by which the market price moves when the trader demands marginally more, $\Lambda = \frac{\partial RS_i^{-1}(a)}{\partial a}$, increases, while liquidity, $1/\Lambda$, declines.

When all traders internalize price impact there is an additional effect coming from the bidder's incentive to shade their bids. In uniform-price auctions, bidders strategically shade their bids on subsequent units to decrease the clearing price, and save money on the purchase of earlier units. This increases the total price impact of each trader. To see this formally, we decompose the total price impact (8) into two parts. First, the price impact that would arise if traders were naive and therefore submitted their actual discounted willingness to pay (5). Since its slope is $\left(\frac{\rho\sigma^2}{1+\lambda\kappa}\right)$ the implied price impact is $\left(\frac{1}{N-1}\right)\left(\frac{\rho\sigma^2}{1+\lambda\kappa}\right)$. We refer to this as the direct effect. Second, the effect coming from demand reduction—the strategic effect:

$$\Lambda = \underbrace{\left(\frac{1}{N-1}\right)\left(\frac{\rho\sigma^2}{1+\lambda\kappa}\right)}_{\text{Direct effect}} + \underbrace{\left(\frac{1}{(N-1)(N-2)}\right)\left(\frac{\rho\sigma^2}{1+\lambda\kappa}\right)}_{\text{Strategic effect}} > 0. \tag{10}$$

When $\lambda\kappa$ decreases, both the price and price impact increase by more when traders are strategic than if they were naive.

Corollary 2. *A 1% decrease in the shadow cost of the capital constraint, $\lambda\kappa$, leads to an increase in the price, and the price impact, equal to $\eta = \left|\frac{1}{1+\lambda\kappa} - 1\right|\%$.*

Since increasing supply by 1% decreases the market price by $\frac{N-1}{N}\Lambda\%$, Corollary 2 implies that the market price reacts more strongly to a change in supply when capital constraints are relaxed.

The result that relaxing constraints increases price impact, and thus decreases liquidity, is surprising in light of the existing literature. The conventional view, summarized by Vayanos and Wang (2012) and references therein, is that weaker constraints makes traders less risk averse because they are less afraid of hitting the constraint.[10] Less risk averse traders can respond to price changes more easily, so that the market becomes more liquid.[11]

Our framework differs in two key dimensions. First, we focus on a different type of constraint. Unlike value-at-risk or margin constraints, our constraint captures requirements on bank leverage, which have become increasingly important for traders since the last financial crisis (e.g., CGFS (2016); ESRB (2016); Group of Thirty (2021); Boyarchenko et al. (2020); Duffie (2023)). Second, we keep risk aversion constant to underscore a distinct and opposing channel through which capital constraints influence asset prices and price impact, separate from risk aversion. Our identified effect arises solely from changes in the shadow cost of the constraint, which affects traders' willingness to pay and thus their strategic behavior. Notably, this shadow cost effect persists even when we eliminate private information in our model, indicating it does not stem from private or asymmetric information among traders.

[10]Sometimes, this relationship is directly built into the model, for instance, by imposing logarithmic utility, which implies that absolute risk aversion decreases in wealth (e.g., Kyle and Xiong (2001); Xiong (2001)).

[11]Two exceptions are Glebkin et al. (2023), who find that the result can flip if payoffs are non-Gaussian as in the case of derivatives, and Fardeau (2024), who extends Gromb and Vayanos (2002) to illustrate that tighter value-at-risk constraints can improve liquidity when arbitrageurs compete a la Cournot.

2.3 Model extensions

We consider five model extensions. First, in our benchmark model, the trader does not hold any inventory going into the auction. In Appendix A.1, we explain how to extend the model to account for privately observed inventory positions. These positions affect both the traders gross utility from winning at auction, and the capital constraint, since it affects the traders total exposure.

Second, in our benchmark model, asset supply is exogenous and all traders buy from the auctioneer. It is straightforward to adjust the setting to approximate centralized double-sided markets, such as an exchange. In this case, demand, $p_i(\cdot, \epsilon_i) : \mathbb{R} \to \mathbb{R}$, represents the trader's demand net of supply. In Appendix A.2, we consider two trade environments, one with symmetric and one with asymmetric market conditions for buyers versus sellers. We show that our main prediction, that the price and price impact increase when constraints are relaxed, can generalize to trade settings.

Third, we consider other types of constraints in Appendix A.3. We start by explaining how to adjust our framework to study the effects of wealth or budget constraints, price caps, quantity-capacity constraints, and other Basel III regulatory requirements (such as risk-weighted leverage ratios, and liquidity coverage ratios).[12] In fact, any linear constraint of the form: $\mathbb{E}[h(\boldsymbol{P^c}, \boldsymbol{a_i^c}] \geq 0$, where $h(\boldsymbol{P^c}, \boldsymbol{a_i^c}) = \Gamma + \Xi a_i^c + \Omega P^c + \Upsilon P^c a_i^c$ with $\Gamma, \Xi, \Omega, \Upsilon \in \mathbb{R}$, gives rise to a linear equilibrium, which can be determined analogously to our equilibrium in Proposition 2. Furthermore, we illustrate what happens when imposing that the constraint must hold ex-post for all states of the world (see Proposition 6, Corollaries 3 and 4).

Fourth, in our benchmark model, the auction format is uniform price. Another popular format for selling or trading multiple units of a good is the discriminatory price auction. In this format, winning bidders pay the prices they bid, rather than the market clearing price. We define price impact, Λ, like in a uniform price auction, as the inverse slope of a bidder's residual supply curve. At the marginal bid (which is the last winning bid, and equals to the market price), $1/\Lambda$ approximates the amount necessary to change the market price by 1

[12]Risk-weighted leverage ratios assign different weights to assets based on their riskiness, while the liquidity cover ratio is a quantity-based funding constraint that depends primarily on banks' holdings of high-quality liquid assets, such as Treasuries. Both are cornerstones in the Basel III regulation.

dollar, similar to the uniform price auction.

In Appendix A.4, we characterize equilibrium conditions for discriminatory price auctions (Proposition 7), and provide conditions under which we can solve for a linear equilibrium (Proposition 8), analogously to Proposition 1 and 2. This allows us to show that the main result of the paper, that the price and price impact increase when capital constraints are relaxed, generalizes to this auction format when the equilibrium is linear.

There is one difference to the uniform price auction, which is that bid-shading in discriminatory price auctions need not be increasing in quantity. Ausubel et al. (2014) show that it can be optimal to shade bids for small quantities more strongly compared to bids for large quantities, depending on the bidder's beliefs about market clearing, and ultimately the distribution of winning quantities. Price impact is now:

$$\Lambda^{DPA} = \underbrace{\left(\frac{1}{N-1}\right)\left(\frac{\rho\sigma^2}{1+\lambda\kappa}\right)}_{\text{Direct Effect}} + \underbrace{\left(\frac{N\xi}{(N-1)(N(1-\xi)-1)}\right)\left(\frac{\rho\sigma^2}{1+\lambda\kappa}\right)}_{\text{Strategic Effect}} > 0, \qquad (10')$$

where $\xi \in \mathbb{R}, \xi \neq 1$ is a shape parameter of the distribution of winning quantities, which, in any linear equilibrium, follow a Generalized Pareto Distribution. When $\xi < 0$, the hazard rate of the amount the bidder wins at market clearing is increasing in quantity, which generates incentives for bid shading to decrease in quantity, and pushes down price impact. The reverse is true when $\xi > 0$.

Finally, in our benchmark model, demand is continuous in quantity. Working with continuous demand functions is common in the related theory literature in order to achieve tractability, even though in practice demand functions are often discrete (in quantity). For example, bidders must submit step functions in most Treasury auctions. Therefore, we also provide equilibrium conditions for step functions in Proposition 3.

Take away. Summarizing, our model helps explain how capital constraints affect asset prices, price impact, and, with that, liquidity. When constraints are relaxed, both the market price and price impact increase. For primary markets, this highlights that relaxing capital constraints increases auction revenues at an implicit cost of larger price impact when issuing more supply. In the context of exchange markets, higher price impact indicates reduced market liquidity.

3 Quantification

Leveraging data on Canadian Treasury auctions, we illustrate how to use our framework to map changes in capital requirements into changes in price and price impact. Our interest lies in the relative change in prices and price impact due to altered constraints, illustrating theoretical mechanisms in action within the data. We expect price effects to be small in absolute terms, in light of the previous literature that indicates that Treasury auctions are in most countries highly liquid, and therefore characterized by relatively flat bidding functions and small average bid shading (e.g., Kang and Puller (2008); Kastl (2011); Hortaçsu et al. (2018); Allen et al. (2020)).

3.1 Institutional setting and data

Canadian primary market. Governments issue bonds of different maturities in the primary market via regularly held uniform price or discriminatory price auctions. In Canada, regular auctions are discriminatory price. Each bidder submits a step function with at most seven steps, which specifies how much a bidder offers to pay for specific amounts of the asset for sale. Anyone may participate, but eight dealers purchase the majority of the Treasury supply.[13] These dealers are federally regulated deposit-taking banks who face Basel III requirements. They dominate the Canadian Treasury market, hold a substantial amount of Government of Canada bonds on their own balance sheets (as shown in Appendix Figure A1), and intermediate the vast majority of the daily trade volume in government bonds (Berger-Soucy et al. (2018)). More broadly, these banks dominate the Canadian banking sector and hold over 90% of the sector's assets.

According to a survey of market participants, the Basel III leverage ratio represents the most relevant capital constraint when trading government bonds (CGFS (2016)). This regulatory requirement came into effect in September 2014 to reduce systematic risk—a

[13]In total there are eleven primary dealers. One of these dealers is provincially regulated, and two are private securities dealers. They face different capital regulation than the eight dealers we study. We do not observe any balance sheet information for these players. Technically, two of the eight banks have multiple dealers. For example, the Bank of Montreal has two dealers (Bank of Montreal and BMO Nesbitt Burns) who attend different Treasury auctions and therefore do not compete or share information within an auction. We treat these as a single dealer.

benefit which we do not consider in this paper. We focus on the cost side of the constraint, which was emphasized by Duffie (2018), He et al. (2022), among others. Formally, the LR measures a bank's Tier 1 capital relative to its total leverage exposure, and must be at least 3%:

$$\text{LR}_{iq} = \frac{\text{regulatory Tier 1 capital of bank } i \text{ in quarter } q}{\text{total leverage exposure of } i \text{ in } q}.$$

Tier 1 capital consists primarily of common stock and disclosed reserves (or retained earnings), but may also include non-redeemable non-cumulative preferred stock; the leverage exposure includes the total notional value of all cash and repo transactions of all securities, including government bonds, regardless of which securities are used as collateral (for more details, see OSFI (2023)).

In reality, banks refrain from getting close to the Basel III threshold. One reason is that banks tend to hold a conservation buffer for Tier 1 capital so as to avoid punishment in the form of restricted distributions (including dividends, share buybacks, and discretionary payments, such as bonuses). Through the lens of the model, this implies that the relevant threshold of the capital constraint is unobservable to the econometrician

To separately identify shadow costs of the capital constraint and risk aversion, we rely on a regulatory change that temporarily eliminated the capital constraint for Treasuries. When dealers failed to absorb the extraordinary supply of government bonds in March 2020, government bonds, central bank reserves, and sovereign-issued securities that qualify as high-quality liquid assets (HQLA) were temporarily exempted from the LR constraint—starting on April 9, 2020.[14] As a result, the LR spiked upward, moving away from the constraint, as shown in Figure 2b. Other regulatory requirements, such as risk-weighted capital ratios, or liquidity coverage ratios, remained in place. The exemption of government bonds and HQLA ended on December 31, 2021, while reserves continued to be excluded.

[14]The announcement to start and subsequently end the temporary exemption can be found here: `https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/unwinding-certain-temporary-exclusions-leverage-ratio-exposure-measures`, accessed on 05/01/2024. Exposures related to the US Government Payment Protection Program (PPP), which are minor in the case of Canadian banks, were also temporarily exempted.

Figure 2: The effect of the exemption on Treasury positions and the LR

(a) Aggregated positions in Treasuries

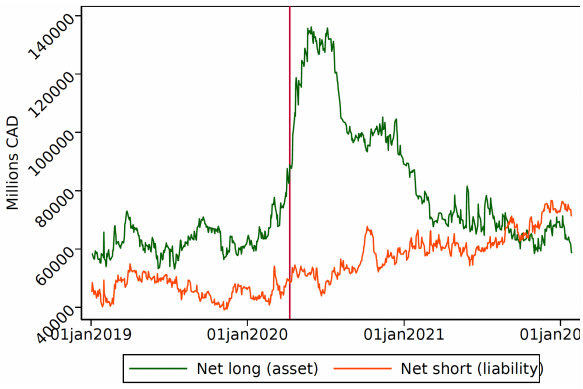

(b) Time series of LR for an average bank

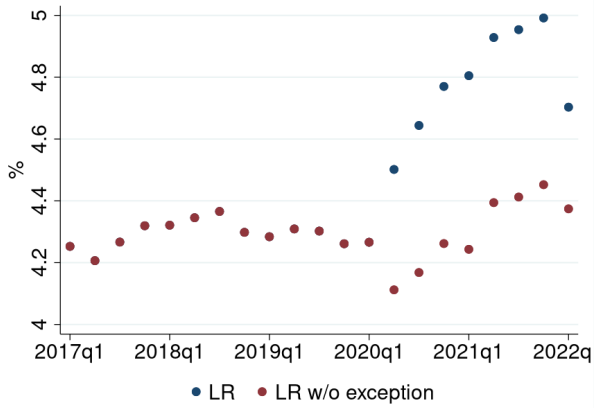

Figure 2a shows the aggregated amount of Canadian government bonds that the biggest six Canadian banks hold in long (in green) and short (in red) positions in millions of C$ from January 2019 until February 2022. The vertical line is April 9, 2020, when the exemption period began. Figure 2b shows the time series of the actual LR (in %) of an average bank in blue. In red is the counterfactual LR that the average bank would have had in absence of the exemption. In 2022q1, the LR does not get back to its original level, partially because central bank reserves are still exempted.

Data. We combine multiple data sources. First, we obtain bidding data of all regular government bond auctions between January 1, 2019, and February 1, 2022, from the Bank of Canada. We see how much is issued of which security, and the maturity category, of which there are five (2Y, 3Y, 5Y, 10Y and 30Y). We also observe who bids (identified by a legal entity identifier) and all winning and losing bids at auction closure. For consistency, we restrict attention to bids of the eight dealers who are deposit-taking.

Second, we collect balance sheet information for these eight dealers at the holding company level. We obtain the quarterly LRs of each dealer from 2015q1 until 2022q1 from the Leverage Requirements Return, in addition to the daily aggregated long and short positions in government bonds of the six largest dealers from the Collateral and Pledging Report (H4).

Third, we gather information on the volatility of the return, i.e., the price, that a dealer expects to obtain from selling government bonds in the secondary market. For this we leverage the fact that dealers start selling bonds that are about to be issued at auction when the tender call opens, which happens one week before the auction closes. This means that dealers already observe the distribution of prices at which they can sell a particular

bond, which gives them a precise idea about the return volatility. To also observe this price distribution, we obtain prices (and yields) of essentially all trades with Canadian government bonds from January 1, 2019, until February 1, 2022. These data are collected by the Industry Regulatory Organization of Canada in the Debt Securities Transaction Reporting System and are made available for research with a time lag.

Fourth, we collect the Implied Volatility Index for Canadian Treasuries over the same time period. This index measures the expected volatility of the market over the next 30 days and is based on option prices on short-term interest rate futures (Chang and Feunou (2014)). It is similar to the Merrill Lynch Option Volatility Estimate (MOVE) for U.S. Treasuries. Different to the U.S., the Canadian option market isn't always liquid, which creates volatility in the index. Therefore, we don't rely on this index in our estimation.

An overview of the main variables is presented in Table 1. We express bond prices and values in yields-to-maturity to make the value of bonds that have different maturities and coupon payments more comparable. In line with this convention, we compute the auction-specific return volatilities as standard deviation of yields (expressed in %) at which a dealer sells a bond that is to be auctioned during the week preceding the auction. To avoid our estimates being driven by the absolute magnitude of the volatility, we normalize the return volatility by its average. Figure 3 shows that the resulting return volatility is similar, yet not identical, to the implied volatility index for Canadian Treasuries (in %).

3.2 Identifying shadow costs of capital and dealer risk aversion

To quantify by how much the price and price impact change when capital constraints are relaxed or tightened, we adjust the benchmark model to better fit the data-generating process. Appendix A.4 provides technical details.

Model adjustments. Consider an auction t that issues a bond of maturity m. In line with the institutional setting, the auction is discriminatory price, and bidders submit step functions with $K(\epsilon_{ti})$ steps for each signal, ϵ_{ti}. As before, a dealer with private signal ϵ_{ti} is willing to pay

$$v_t(a_{tik}, \epsilon_{ti}) = (\mu + \epsilon_{ti}) - \rho_m \sigma_t^2 a_{tik} \tag{2}$$

Table 1: Summary statistics

	Mean	Median	Std	Min	Max
Supply (in bn C$)	4.12	4.00	1.23	1.40	7.00
Average bid yield (in %)	1.04	1.09	0.58	0.20	2.18
Years to maturity	8.40	5.03	9.62	2.0	32.62
Number of deposit-taking dealers	8	8	0	8	8
Number of steps in demand curve	4.80	5	1.43	1	7
Maximal amount demanded (in % of supply)	21.32	21.17	10.70	0.08	50.00
Amount dealer won (in % of supply)	6.55	4.98	6.22	0	44.22
Quarterly LR (in %)	4.41	4.36	0.28	-	-
Return volatility (normalized)	1	0.74	0.91	0	8.14

Table 1 shows the average, median, standard deviation, minimum, and maximum of key variables in our sample. Our auction data goes from January 1, 2019, until February 1, 2022, and counts 176 bond auctions. The LR data goes from 2015q1 until 2022q1. The min and max LR are empty because we cannot disclose this information.

Figure 3: Return volatility

(a) Across maturities

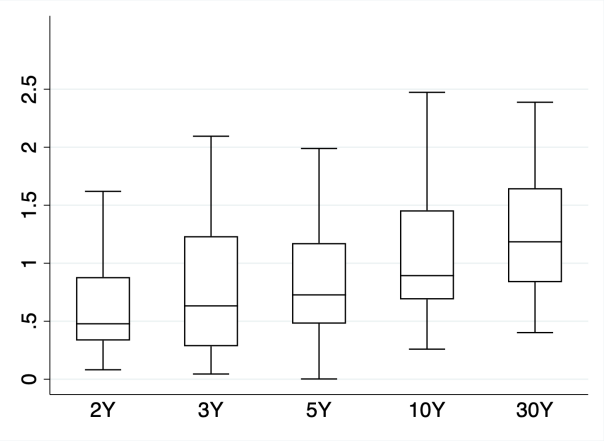

(b) Over time

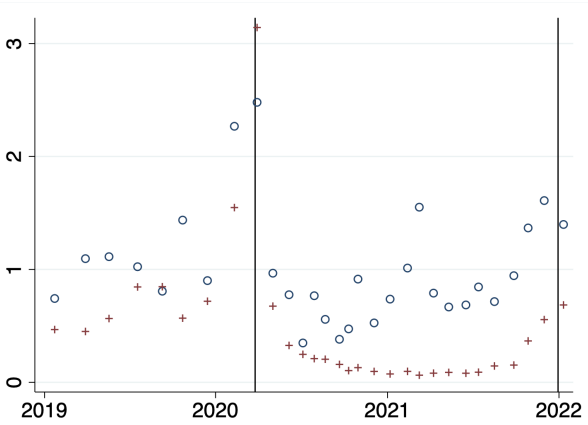

Figure 3a shows the distribution of the normalized return volatility for each maturity category, excluding outliers. Figure 3b shows a binned scatter plot of the return volatility (in circles) and the implied volatility index in % (in pluses) across time. The correlation between these two volatility indices is 0.3. The black lines mark the beginning (09 April 2020) and end (01 January 2022) of the exemption period.

for amount a_{tik} in auction t for maturity m with return volatility σ_t^2 when the constraint does not bind, i.e., $\lambda\kappa_t = 0$.[15] Parameter $\rho_m \geq 0$ measures the degree of risk aversion for a bond with maturity m. Thus, we allow, but do not impose, risk aversion to vary in the bond's maturity to reflect the fact that longer bonds may be riskier to hold than shorter bonds.[16] Parameters $\lambda\kappa_t \geq 0$ represent the shadow costs of the capital constraint. Note that we estimate the product of the Lagrange multiplier of the constraint and the capital threshold to avoid having to specify a capital threshold. To highlight this, we relabel the shadow costs $\lambda\kappa_t$ instead of $\lambda_t\kappa$.

Proposition 3 characterizes how bidders bid in this environment. The main difference to the conditions of our benchmark model (Proposition 1) is that bidders submit step functions instead of continuous functions. This changes the distribution of the market clearing price and the amount each bidder wins, and therefore the shading factor. The economic insights, however, carry over.

Proposition 3. *Consider a discriminatory price auction t with step-function demand , and denote bidder i's step-function by $\{p_{tk}(\epsilon_{ti}), a_{tik}\}_{k=1}^{K_t(\epsilon_{it})}$, where we abbreviate $a_{tik} = a_{tk}(\epsilon_{ti})$. Denote the discounted willingness to pay by*

$$\tilde{v}_t(a_{tik}, \epsilon_{ti}) = \frac{v_t(a_{tik}, \epsilon_{ti})}{1 + \lambda\kappa_t}. \tag{5}$$

The equilibrium demand function satisfies for all ϵ_i

$$p_{tk}(\epsilon_{ti}) = \tilde{v}_t(a_{tik}, \epsilon_{ti}) - \frac{\Pr(p_{tk+1}(\epsilon_{ti}) \geq \boldsymbol{P_t^*} | \epsilon_{ti})}{\Pr(p_{tk}(\epsilon_{ti}) > \boldsymbol{P_t^*} > p_{tk+1}(\epsilon_{ti}) | \epsilon_{ti})} \tag{11}$$

[15]The functional form of the intercept $\mu + \epsilon_{ti}$ is not relevant, as long as the private signal, ϵ_{ti}, only affects the intercept, and not the slope. Moreover, we could rely on our extended model, where bidders have private information about their inventory position. This would give rise to the same conditions and estimates. In both model specifications bidders are ex-ante identical. We think that this is a valid assumption because the eight deposit-taking dealers we focus on in the paper play similar roles in the secondary market, and face the same regulation. In Allen et al. (2020), we provide evidence to support this idea. Specifically, in that paper we estimate a more general auction model in which each dealer may have a latent business type that affects their true willingness to pay. Our findings suggest that the eight largest dealers share the same business type, and therefore the same preferences.

[16]An alternative would be to impose that risk aversion is constant across maturities—this is rejected by the data, as shown in Figure 5.

at every step but the last one; at the last step, the trader bids truthfully. The Lagrange multiplier, and the implied shadow cost, $\lambda\kappa_t$, are pinned down by the capital constraint.

Identification. Before explaining how we estimate the shadow costs of the capital constraint, $\lambda\kappa_t$, and the dealer's degree of risk aversion, ρ_m, we discuss what variation in the data helps to identify which parameter under what assumptions.

For this, recall the main insight from Proposition 1, which generalizes to the empirical environment (Proposition 3): In auction t, dealer i bids for step k like in an auction without constraints in which their willingness to pay is discounted by the shadow cost, as given in expression (5). When the unconstrained willingness to pay for maturity m follows expression (2), this implies that the slope in the dealer's discounted willingness to pay curve only depends on their degree of risk aversion, ρ_m, shadow costs, $\lambda\kappa_t$, and bond return-volatility, σ_t^2:

$$\frac{\partial \tilde{v}_t(a_{tik}, \epsilon_{ti})}{\partial a} = -\beta_t \sigma_t^2 \text{ with } \beta_t = \frac{\rho_m}{1 + \lambda\kappa_t}. \tag{12}$$

Thus, if we observed the discounted willingness to pay curves, $\tilde{v}_t(\cdot, \epsilon_{ti})$, and bond return-volatility, we would be able to identify the shadow cost, $\lambda\kappa_t$, and risk-aversion, ρ_m, by comparing the bidders' willingness to pay during the exemption period of the capital constraint—where $\lambda\kappa_t = 0$—to periods when the constraint is in effect. Specifically, risk aversion, ρ_m, is identified from the average slope of the willingness to pay, $\tilde{v}_t(\cdot, \epsilon_{ti})$, when the constraint is exempt ($\lambda\kappa_t = 0$) because the slope of the unconstrained willingness to pay, $\rho_m\sigma_t^2$, is assumed to vary across auctions solely due to observable changes in observed return volatility, σ_t^2. Shadow costs, $\lambda\kappa_t \geq 0$, are identified from the slopes under the capital constraint.

So far, we have assumed that we observe how much bidders are willing to pay. However, Proposition 3 makes clear that this is not the case in equilibrium. As in standard auctions without constraints, dealers shade their bids, with the degree of shading determined by the auction rules and the bidder's belief about the market-clearing price. This feature is not specific to our setting, which allows us to follow the existing literature in separating bid shading from the underlying willingness to pay.

Building on Hortaçsu and McAdams (2010) and Kastl (2012), we identify the shading factor from variation in the submitted bids for given auction rules that determine the way the

market clears. The main idea is to take the perspective of one bidder, fix their submitted bidding step-function, and simulate market clearing many times to back out the bidder's belief about the market clearing price, i.e., the $\Pr(...|\epsilon_{ti})$ terms in condition (11).[17] With the shading factors, we identify the unique values, $\tilde{v}_t(a_{tik}, \epsilon_{ti})$, that rationalizes the observed bid in each auction t of each dealer i and submitted step k from the equilibrium conditions (11), under the assumption that bidders play the equilibrium strategy (Kastl (2011)).

Estimation. We estimate bidder values, $\tilde{v}_{tik} = \tilde{v}_t(a_{tik}, \epsilon_{ti})$, following the resampling approach in Allen et al. (2024), which takes institutional details of Canadian Treasury auctions, that are omitted in our theoretic model, into account. For example, it adjusts for the fact that there are not only dealers, but also customers who bid via dealers. Importantly, these details only affect the way we estimate the dealers' beliefs about where the market will clear, but not the equilibrium condition itself for a given price distribution.

Using value functions with at least two steps, which represent 99% of all functions, we then estimate parameters $\{\rho_m, \lambda\kappa_t\}$ for all m, t, by fitting

$$\tilde{v}_{tik} = \zeta_{ti} - \sum_t \beta_t \mathbb{I}(auction = t)\sigma_t^2 a_{tik} + \epsilon_{tik}, \text{ with } \beta_t = \frac{\rho_m}{1 + \lambda\kappa_t} \tag{13}$$

such that $\lambda\kappa_t = 0$ for all auctions t when Treasuries are exempt or otherwise $\lambda\kappa_t \geq 0$, and $\rho_m \geq 0$. Here ζ_{ti} is a dealer-auction fixed effect, σ_t^2 is the return volatility plotted in Figure (3), and ϵ_{tik} represents finite sample measurement error in the values, $\tilde{v}_{tik}$.

We estimate two separate sets of parameters, one for when the exemption period started in 2020, and one when it ended in 2021. Given that capital requirements must be fulfilled quarterly, we use data from auctions that took place within one quarter around each policy change, i.e., 2020q1–2020q2 and 2021q4–2022q1.

We express bid and values in percentages of yields, and quantities in percentages of auction supply to avoid that changes in the supply, which increased substantially during the

[17] To simulate market clearing, bids from other competitors in the auction are randomly drawn with replacement to construct a residual supply curve. This curve is then intersected with the fixed bidder's step-function to determine a potential market-clearing price. By repeating this simulation many times, one obtains an unbiased estimate of the bidder's belief about the market-clearing price, under the assumption that bond values are independent across bidders within the auction.

COVID-19 pandemic, affect our estimates. Therefore, β_t measures by how many percentage points the dealer's willingness to pay decreases on average when demand increases by 1% of auctions supply in an auction with average return volatility ($\sigma_t^2 = 1$).

Estimation findings. Before identifying our parameters of interest, we analyze the slope coefficients, β_t, from regression (13), when estimated without constraints on shadow costs or risk aversion, and using data from all auctions in our sample. We do this using estimated values and observed bids to visualize the key variation in the data that pins down our parameters of interest.

Our theory predicts that if constraints bind then the willingness to pay, expressed in equation (5), becomes steeper when Treasuries are exempt. In that case, slope coefficients, β_t, should be larger during the exemption period than during regular times. In contrast, a model without capital constraints, where the degree of risk aversion increases when the constraint is more binding, would predict that slope coefficients, β_t, are smaller during the exemption period than in regular times. In line with our model, Figure 4 shows that β_t is larger during the exemption period than in regular times for both events—the start of the exemption period, and the end, where we are less concerned about COVID-related confounding factors. A t-test we conduct in Appendix B.1 confirms that the difference in slopes during versus outside of the exemption period is statistically significant at the 5% level.

Next, we separate the degree of risk aversion from the shadow costs by estimating regression (13) with constraints using data of auctions around the policy changes. We find that risk aversion is relatively low for all bond types with no clear pattern with respect to maturity (see Figure 5). The median (mean) degree of risk aversion is 0.005 (0.004). This implies that a typical dealer is willing to pay 0.5 bps less for 1% more of the auction supply in an auction with average return volatility. If dealers were risk neutral, their willingness to pay would be perfectly flat.

In comparison, the existing auction literature estimates risk aversion of similar, yet typically larger, magnitudes in non-financial settings, and given CARA preferences. Most papers consider single-unit auctions. For instance, Bolotnyy and Vasserman (2023) estimate a median degree of risk aversion of firms in procurement auctions to be 0.08. One exception is

Figure 4: Slope coefficients of estimated value and observed bids

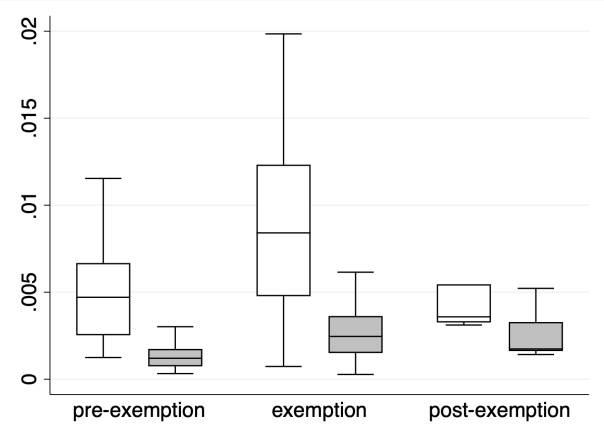

The white box plots of Figure 4 show the distribution of the estimated slopes coefficient of the average dealer's willingness to pay in auction t of regression (13) without imposing restrictions on ρ or $\lambda\kappa_t$ for three time periods: before the exemption of Treasuries from the LR (2019q1–2020q1), during the exemption period (2020q1–2021q4), and after the exemption (2022q1). The gray box plots show the analogue when using bids instead of estimated values. Dealer values and bids are in %; quantities are in % of auction supply.

Häfner (2023), who analyzes discriminatory price auctions for Swiss tariff-rate quotas. He finds that the majority of bidders exhibit a risk aversion parameter of 0.007.

Shadow costs, as shown in Figure 6a, vary substantially across auctions; they are positive in auctions (outside of the exemption period) in which the average willingness to pay curve is relatively flat, but the auction-specific return volatility is high. Intuitively, whether capital constraints bind depends on the dealer bank's entire balance sheet, which likely vary over time due to a wide range of factors, including conditions in multiple financial markets and the banks' broader business activities. In line with this idea, we find that the auctions with binding constraints occurred on different dates—spanning beyond the peak of the COVID crisis in March 2020—and across the maturity spectrum. The long tail in the distribution of shadow costs suggests that there are some auctions in which dealers expect to take large losses. Due to these outliers, the average shadow cost is large (53%), with the exact number depending on the specification we consider.

However, in a typical auction the constraint is not binding—the median shadow cost is zero across all regression specifications we consider. This finding aligns with extensive empirical evidence indicating that banks typically avoid operating too close to regulatory

Figure 5: Risk aversion bond type

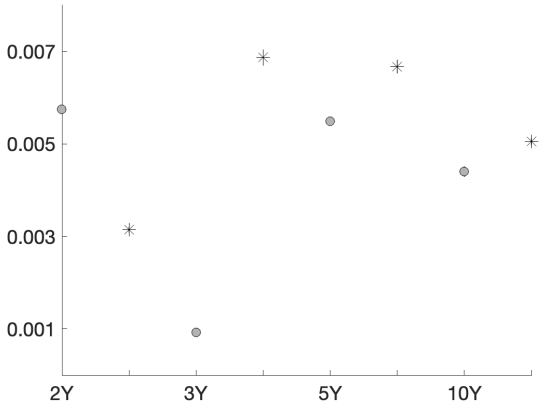

Figure 5 shows the risk aversion estimates, ρ_m, for $m = \{2Y, 3Y, 5Y, 10Y\}$ around both policy changes in 2020q1–202q2 (in circles) and in 2021q4–2022q1 (in stars). We exclude 30Y bonds because they were not issued in all four quarters we consider. The graph also plots the 95% confidence intervals for 2021q4–2022q1, but given that these intervals are very tight, they are not visible. To compute standard errors and these intervals, we fit equation (13) for each bootstrapped estimate of values. Each coefficient is measured in % of yield relative to % of supply.

constraints (e.g., Barth et al. (2005); Berger et al. (2008); Brewer et al. (2008); Corbae and D'Erasmo (2021)).[18] It supports the view that banks face substantial costs when constraints bind. It also echoes evidence from Duffie et al. (2023), who provide evidence in line with occasionally binding constraints on dealers' intermediation capacity in the U.S. Treasury market. More broadly, our estimates contribute to a literature that quantifies costs associated with constraints in financial markets, using diverse data and methodologies.[19]

We conduct a series of robustness checks in Appendix B.2. We explain what happens

[18]In particular, Corbae and D'Erasmo (2021) document U.S. banks holding Tier 1 capital that is two or three times above the regulatory minimum. In their model, the capital constraint is only occasionally binding since banks hold precautionary savings to protect their sizable franchise values and hitting the constraint is costly.

[19]For instance, Du et al. (2018) use the overnight spread between the interest rates on excess reserves paid by the Federal Reserve and the Fed Funds rate as proxy for the average shadow costs of bank's balance sheets, which is a couple of basis points. Adrian et al. (2014) fit an augmented Fama-French factor model using quarterly balance sheet data from U.S. security broker-dealers from 1968q1 to 2009q4 and U.S. stock returns. They compute a price of leverage (which proxies for the funding constraint of Brunnermeier and Pedersen (2009), among others) of roughly 10% per year.

Figure 6: Shadow costs and their effect on price and price impact

(a) Estimates and confidence intervals (b) Price and price impact elasticity

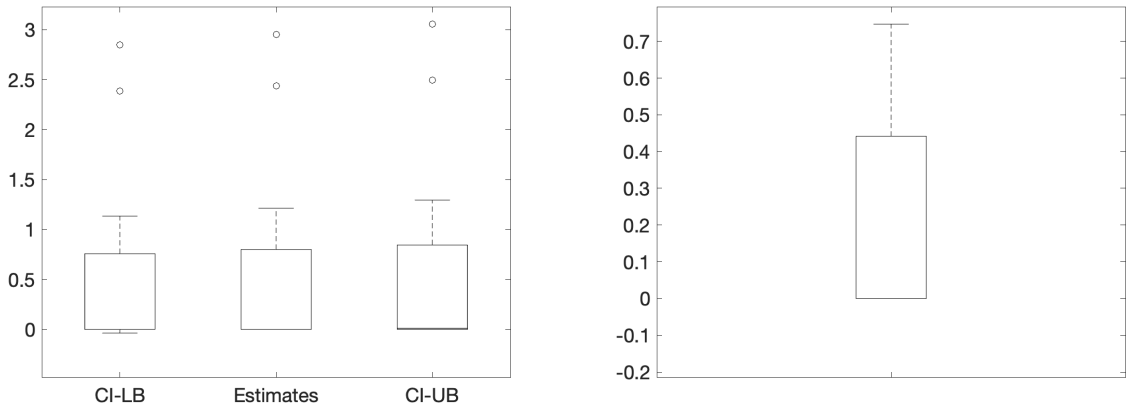

Figure 6a shows the distribution of the lower and upper bounds of the 95% confidence intervals for shadow costs (CI-LB and CI-UB, respectively), in addition to the distribution of the point estimates (Estimates) for all auctions in quarters around the policy changes, excluding outliers. The confidence intervals of the shadow costs are bootstrapped, analogous to those of the degree of risk aversion. Figure 6b displays the distribution of elasticity $\eta_t = |\frac{1}{1+\lambda\kappa_t} - 1|$ of Corollary 2 using the shadow cost point-estimates for auctions in the first quarters of 2020 and 2022.

when we rely on different volatility indices, when expressing quantities in absolute terms rather than in percentages of supply, and when including different samples of bidding functions in the estimation. Our main finding, that risk aversion is moderate (around 0.005), and that constraints are typically not binding, but when they do, it is costly, is robust across all specifications.

Discussion. To separate shadow costs from risk aversion, we rely on several assumptions that enable identification. We outline each assumption, discuss its motivation, and consider the implications of possible violations.

We assume that dealers behave according to the equilibrium strategy of our empirical auction model. This assumption is standard in the empirical auction literature and enables us to recover willingness to pay from observed bids. It aligns with the institutional context, where dealers are both sophisticated and experienced in Treasury auctions. As a robustness check, we also estimate the model under the alternative assumption that dealers bid their true willingness to pay plus a random error. This specification yields lower estimates of

Figure 7: Shading

(a) Per step

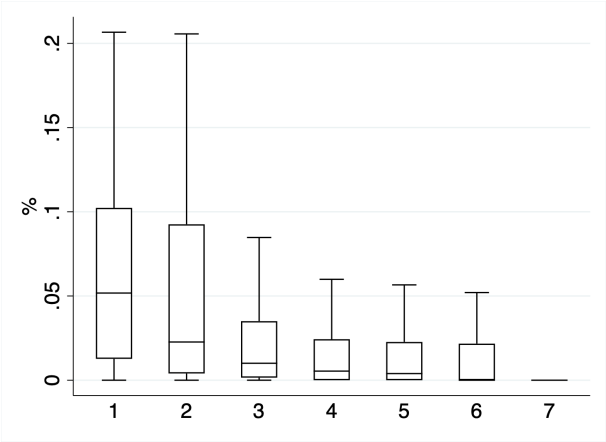

(b) Average per period

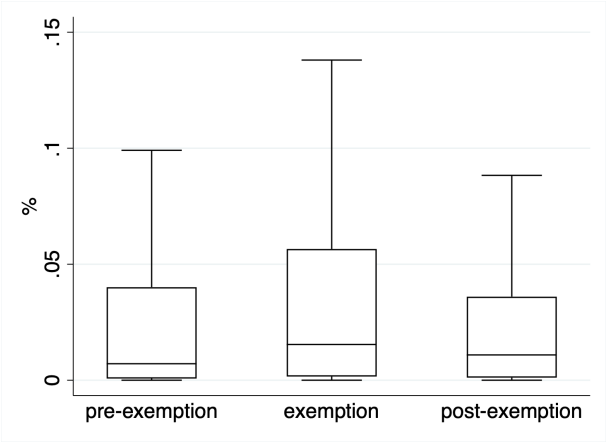

Figure 7a shows box plots of how much dealers shade their bids at each step. It is the difference between the submitted bid and the estimated value, both in percentage. The distribution for each step is taken over dealers and auctions. Shading factors are small in absolute terms, and comparable to those in the literature. Figure 7b shows the distribution of shading across auctions, dealers, and steps in 2020q1 (pre-exemption), 2020q2 and 2021q4 (exemption), and 2022q1 (post-exemption).

risk aversion (median of 0.002), since bid shading decreases in quantity (Figure 7a); yet the average and median shadow costs (43% and zero, respectively) remain similar to the baseline.

Moreover, we assume that the unconstrained willingness to pay follows the linear functional form in expression (2), which arises, for example, under CARA preferences. While this assumption is common in the theoretical auction literature (as it ensures tractability), it is less frequently adopted in the empirical auction literature, which typically aims for non-parametric identification. We cannot directly test our parametric assumption, but we can provide evidence suggesting that it is a reasonable first-order approximation: Appendix Table A1 shows that submitted bids are approximately linear. Therefore, it would be surprising if the underlying willingness to pay was highly nonlinear.[20]

The slope of our willingness to pay curve during the exemption period of the capital constraint is $\rho_m \sigma_t^2$, where σ_t^2 captures auction-specific return volatility. This would not be the

[20] Ausubel et al. (2014) show that a linear equilibrium arises in discriminatory price auctions without private information (where bidders submit continuous bidding functions) if and only if willingness to pay is linear and supply follows a Generalized Pareto distribution.

case if banks faced balance sheet constraints about the nominal value of assets beyond those stemming from (LR) capital constraints. If this was the case, what we would be identifying is the variation in total nominal balance sheet costs as regulatory conditions shift. The median slope coefficients depicted in Figure 4 of the different time periods indicate that the total nominal balance sheet costs (which show up in the slope's denominator) were lower during the exemption period. In light of recent research on the LR constraint, e.g., Boyarchenko et al. (2020), we expect most of the effect stems from changes in the shadow cost of that constraint.[21]

The slope may also vary across auctions in a manner not fully captured by changes in the auction-specific return volatility, σ_t^2, or due to differences across dealers. To investigate heterogeneity, we flexibly estimate the slope coefficient in each bidder's willingness to pay curve per auction, β_{ti}, by leveraging variation in value estimates across steps within a step function for each bidder and auction.[22] This identifies the bidder's effective risk aversion after controlling for return volatility, and provides precise estimates in auctions in which bidders submit sufficiently many steps, for instance, in Turkish Treasury auctions (c.f., Hortaçsu and McAdams (2010)). In our setting, where bidding functions have only four steps on average, these estimates are less precise, though the observed pattern aligns with Figure 4. The defining feature is that the β_{ti}, like the β_t estimates, are typically higher during the exemption period than outside of it. To account for the median differences in β_{ti} estimates between these periods, either risk aversion (or factors driving it beyond market volatility) must have increased by approximately 200% during the exemption period, or the LR constraint was binding in some auctions outside of the exemption period.

Finally, we assume that we observe the return volatility, σ_t^2, that a dealer expects to earn from buying bonds at auction and selling them in the secondary market. In practice,

[21]The LR constraint was the primary limitation, because dealers were buying more Treasuries during the exemption period. Given that Treasuries are safe assets with zero risk-weight, this implies that risk-weighted capital ratios and liquidity coverage ratios improved (see Appendix Figure A6). Therefore, neither risk-weighted capital constraints nor liquidity coverage ratios were binding. Moreover, value-at-risk limits were not binding during that period, as is shown, for example in Figure 8 of Lu and Wallen (2024).

[22]We estimate regression (13), but replace β_t with β_{ti} and include an indicator variable for each dealer i, $\mathbb{I}(dealer = i)$, which multiplies the auction-indicator, $\mathbb{I}(auction = t)$.

dealers form such knowledge by actively trading the bond in anticipation of the auction. With finite data, the observed volatility measure may contain measurement error, which could bias our slope estimates downward. If return volatility were unobserved, we would need to impose additional structure on the data—for example, by specifying a model for the secondary market price process.

3.3 Approximating the price and price impact effect

Ideally, we would use our empirical model to quantify the impact of counterfactual capital requirements on prices and price impact. Recent advances by Richert (2023) proposes one way to do this for multi-unit auctions in which bidders do not face constraints. Given the complexity of our problem, the fact that our estimates of bid-shading are small, and bids are approximately linear, we take a more modest approach. Specifically, we rely on our theoretic framework outlined in Section 2 to approximate the magnitude of all effects, setting the number of bidders to eight.

As a starting point, we rely on Corollary 2, which holds for uniform-price and discriminatory price auctions, and tells us that the market price increases and the price impact in auction t increases by $\eta_t = |\frac{1}{1+\lambda\kappa_t} - 1|\%$ when the shadow cost of the capital constraint decreases by 1%.[23] Like shadow costs, this elasticity varies across auctions with a median of zero and a mean of 0.19%. This implies that eliminating the constraint—implying a 100% reduction in the shadow cost of the capital constraint—doesn't change the price or price impact in a typical auction, but results in a significant increase of 19% on average due to rare occasions when it is extremely costly to satisfy the constraint.

Next, in Figure 8, we decompose the price impact in Canadian auctions into the direct and strategic effect that comes from demand reduction. For this, we rely on expression (10'), which depends on the distribution of the amount bidders win at market clearing, estimated separately.[24] In line with the existing literature, we estimate a small median price impact of

[23]This statement hinges on the assumption that volatility is independent of $\lambda\kappa$. In practice, this might not always be the case (e.g., Du et al. (2023a)). Our calculation, therefore, neglects the possible indirect effect that a change in $\lambda\kappa$ has on the price and the markup via a change in volatility.

[24]We first simulate market clearing 1,000 times for each bidder, following the same procedure that we adopt to estimate values. This gives us 1,000 realizations of winning quantities for each

Figure 8: Price impact and decomposition

(a) Discriminatory price auction (b) Uniform price auction

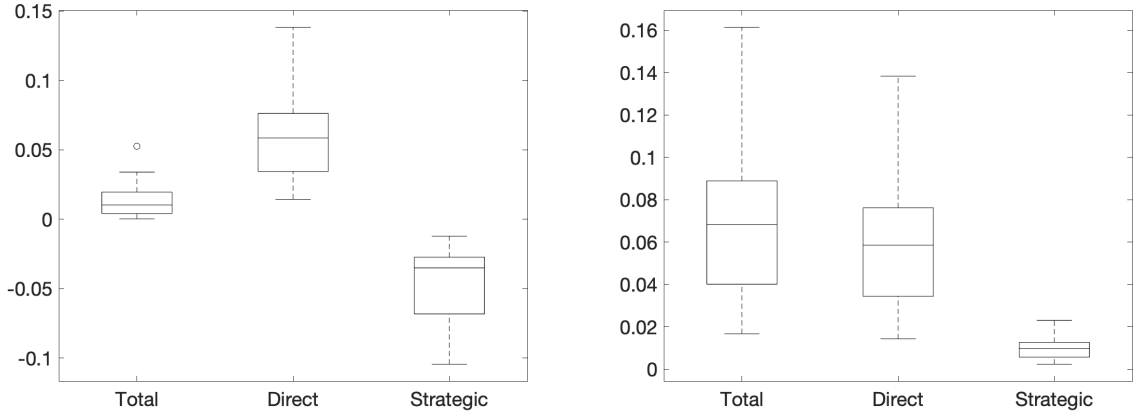

Figure 8 displays the distribution of the total price impact, the direct and strategic effect for both auction formats according to expressions (10) and (10'), using risk aversion and shadow cost estimates for auctions in the first quarter in 2020 and the first quarter in 2022. Units are bps.

0.01 bps. This means that the auction price would drop by less than 0.01 bps if the Canadian government issued 1% more of the bond (assuming a face value of C$100).

If Canada sold their debt using a uniform price auction (like the U.S.), the price impact would be slightly larger with a median of 0.06 bps. The reason for the difference is that bid shading is decreasing in quantity in Canadian Treasury auctions under the status quo, as shown in Figure 7a, while it is increasing in quantity under uniform price auctions. This implies that the strategic effect goes in the opposite direction than the direct effect, reducing the total price impact with discriminatory price auctions.

Given that the absolute value of the price impact is small under both auction formats, eliminating capital constraints results in small absolute price impact changes (see Figure 9). Even the largest change in price impact is relatively small—0.15 bps in the discriminatory price auction and 0.39 bps in the uniform price auction. This finding is in line with the insignificant change in bid shading we observe when the policy changed—see Figure 7b.

Taken together, these calculations suggest that the Canadian regulator did not face a

bidder. Then, we compute the mean and variance of those realizations for each bidder. Finally, we match these "observed" moments to the mean and the bidder-specific variances of the Generalized Pareto Distribution, which is specified in Proposition 8, using GMM.

Figure 9: Change in price impact when eliminating the capital constraint

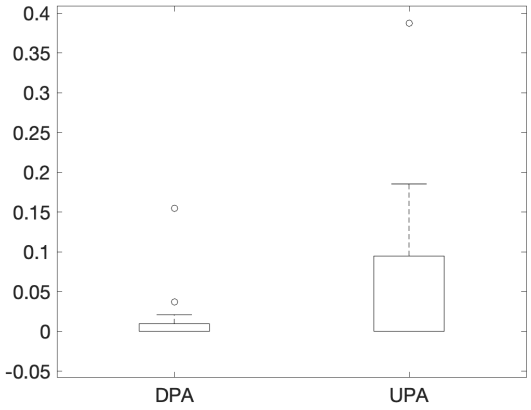

Figure 9 shows the distribution of the change in total price impact when setting all shadow costs to zero, using risk aversion and shadow cost estimates for auctions in the first quarter in 2020 and the first quarter in 2022. For the uniform price auction we compute these effects using expression (10), and for the discriminatory price auction using expression (10').

quantitatively meaningful trade-off in 2020 and 2022 when deciding whether to relax and tighten capital constraints—in addition to the way the LR affects trading in the secondary market and concerns about systematic risk. However, this might not always be the case. The fewer bidders competing in the auctions (as illustrated in Appendix Figure A2), the higher return volatility and risk aversion, the larger the change in price impact in absolute value.[25]

4 Future research

There is room for future research branching into multiple directions. Our theoretic analysis relies on the assumption that traders are ex-ante identical. In some real-world applications this might not be the case. For example, different traders might have different degrees of risk aversion, or face different constraints. These heterogeneities create possibly intriguing asym-

[25]Formally, price impact (10') decreases in the number of bidders, and increases in the return volatility and risk aversion when $\xi < 0$ which is the case for all auctions in our data. This is also true for price impact (10) in uniform price auctions.

metries.[26] Furthermore, we focus on linear equilibria. In the absence of linearity it becomes challenging to prove that a demand function that fulfills necessary equilibrium conditions is a global optimum. This is due to the fact that sufficient conditions for maximization over functions typically hold only locally, as explained, for example, by Elsgolts (1977) in Chapter 8. Addressing this challenge could open the door to exploring potential non-linear effects on prices and price impact.

To maintain the linearity of equilibrium, we concentrate on a particular category of constraints, detailed in Appendix A.3. This category excludes value-at-risk constraints, which have been extensively explored in related literature following Gromb and Vayanos (2002). Given that traders often operate under several constraints simultaneously, exploring how these various constraints interact would provide valuable insights. Likewise, examining how changes in risk aversion due to variations in equity or wealth influence this dynamic—both through the direct effect on risk aversion and indirectly through the constraints—would offer further understanding of these complex interactions.

Finally, it would be valuable to examine how regulatory constraints shape traders' willingness to participate in the market. Such an analysis would complement empirical findings and anecdotal evidence, which suggest that traders—traditionally acting as market makers and who are subject to regulatory constraints—are exiting these markets (e.g., Geddie (2016); Favara et al. (2022); Allen et al. (2024)).

Beyond motivating future theoretical research, our framework could also be used to empirically quantify how capital constraints affect trading, pricing, and liquidity in markets with a limited number of sophisticated traders. This requires access to micro-level data, which has become increasingly available. Since we model market clearing via an auction, the framework is best suited to centralized markets—both one-sided auctions (e.g., for mortgage-backed securities, credit events, or quantitative easing) and double-sided auctions (e.g., exchanges, where packages of limit orders form net demand curves). Applying the model across differ-

[26]For instance, each trader could have an individual expectation over how much capital it holds, and this is common knowledge. Then, a reduction in the expected equity position of one trader might cause this trader to demand more. This increases the market price and reduces the amount that the other traders win. The constraints of the other traders may tighten or weaken, depending on whether the price or quantity effect dominates. The effect on price impact seems ambiguous.

ent market settings can help identify the conditions under which constraints bind and when traders' risk aversion intensifies. An especially promising direction would be to explore potential feedback loops between capital constraints and market volatility, as suggested by Du et al. (2023a), to better understand the dynamics of market behavior.

5 Conclusion

We develop a framework to analyze the interplay between trader market power and capital constraints. We illustrate how to use our framework to quantify the effect of softer capital constraints on prices and price impact using data from Canadian Treasury auctions. Our findings contribute to the discourse on adjusting capital requirements, such as those under Basel III, in the aftermath of the COVID-19 pandemic.

References

Adrian, T., Etula, E., and Muir, T. (2014). Financial intermediaries and the cross-section of asset returns. *Journal of Finance*, 69(6):2557–2596.

Adrian, T., Fleming, M., Shachar, O., and Vogt, E. (2017). Market liquidity after the financial crisis. Federal Reserve Bank of New York Staff Reports, No. 796.

Adrian, T. and Shin, H. S. (2010). Liquidity and leverage. *Journal of Financial Intermediation*, 19:418–437.

Allen, J., Hortaçsu, A., Richert, E., and Wittwer, M. (2024). Exit and entry in Treasury auctions. Working paper.

Allen, J., Kastl, J., and Wittwer, M. (2020). Estimating demand systems with bidding data. Working paper.

Allen, J. and Wittwer, M. (2023). Centralizing over-the-counter markets? *Journal of Political Economy*, 131:3310–3450.

An, Y. and Song, Z. (2023). Does the Federal Reserve obtain competitive and appropriate prices in monetary policy implementation? *Review of Financial Studies*, 36(1-45):4113–4157.

Anderson, M. and Stulz, R. M. (2017). Is post-crisis bond liquidity lower? Working paper.

Ausubel, L. M., Cramton, P., Pycia, M., Rostek, M., and Weretka, M. (2014). Demand reduction and inefficiency in multi-unit auctions. *Review of Economic Studies*, 81:1366–1400.

Baer, J. (2020). The day coronavirus nearly broke the financial markets. `https://www.wsj.com/articles/the-day-coronavirus-nearly-broke-the-financial-markets-11589982288`, accessed 2024-04-08.

Barth, J., Caprio, G., and Levine, R. (2005). *Rethinking Bank Regulation: Till Angels Govern*. Cambridge University Press.

Begenau, J. (2020). Capital requirements, risk choice, and liquidity provision in a business cycle model. *Journal of Financial Economics*, 136:355–378.

Begenau, J. and Landvoigt, T. (2022). Financial regulation in a quantitative model of the modern banking system. *Review of Economic Studies*, 89:1748–1784.

Benetton, M. (2021). Leverage regulation and market structure: A structural model of the UK mortgage market. *Journal of Finance*, 76(6):2997–3053.

Berger, A. N., DeYoung, R., Flannery, M. J., Lee, D., and Öztekin, Ö. (2008). How do large banking organizations manage their capital ratios? *Journal of Financial Services Research*, 34:123–149.

Berger-Soucy, L., Garriott, C., and Usche, A. (2018). Government of Canada fixed-income market ecology. Bank of Canada discussion paper 2018-10.

Bessembinder, H., Jacobsen, S., Maxwell, W., and Venkataraman, K. (2018). Capital commitment and illiquidity in corporate bonds. *Journal of Finance*, 73(4):1615–1661.

Bolotnyy, V. and Vasserman, S. (2023). Scaling auctions as insurance: A case study in infrastructure procurement. *Econometrica*, 91:1205–1259.

Boyarchenko, N., Eisenbach, T., Gupta, P., Scachar, O., and Tassel, P. V. (2020). Bank-intermediated arbitrage. New York Federal Reserve Staff Report No. 858.

Breckenfelder, J. and Ivashina, V. (2021). Bank balance sheet constraints and bond liquidity. Working paper.

Brewer, E., Kaufman, G. G., and Wall, L. D. (2008). Bank capital ratios across countries: Why do they vary? *Journal of Financial Services Research*, 34:177–201.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238.

Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

Campo, S., Guerre, E., Perrigne, I., and Vuong, Q. (2011). Semiparametric estimation of first-price auctions with risk-averse bidders. *Review of Economic Studies*, 78:112–147.

Cenedese, G., della Corte, P., and Wang, T. (2021). Currency mispricing and dealer balance sheets. *Journal of Finance*, 76(6):2763–2803.

CGFS (2014). CGFS papers no 52: Market-making and proprietary trading: Industry trends, drivers and policy implications. Bank of International Settlements.

CGFS (2016). CGFS papers no 55: Fixed income market liquidity. Bank of International Settlements.

Chang, B. and Feunou, B. (2014). Measuring uncertainty in monetary policy using implied volatility and realized volatility. *Bank of Canada Review*, pages 32–41.

Chen, D. and Duffie, D. (2020). Market fragmentation. NBER working paper No. 26828.

Chen, S., Kaniel, R., and Opp, C. C. (2023). Market power in the securities lending market. Working paper.

Clark, R., Houde, J.-F., and Kastl, J. (2021). The industrial organization of financial markets. In: The Handbook of Industrial Organization, K. Ho, A. Hortaçsu, and A. Lizzeri (eds.), North-Holland Press, Volume 5, Chapter 15.

Corbae, D. and D'Erasmo, P. (2021). Capital buffers in a quantitative model of banking industry dynamics. *Econometrica*, 86(6):2975–3023.

Du, W., Hébert, B., and Hubert, A. W. (2023a). Are intermediary constraints priced? *Review of Financial Studies*, 36:1464–1507.

Du, W., Hébert, B., and Li, W. (2023b). Intermediary balance sheets and the Treasury yield curve. *Journal of Financial Economics*, 150(3):1–19.

Du, W., Tepper, A., and Verdelhan, A. (2018). Deviations from covered interest rate parity. *Journal of Finance*, 73(3):915–957.

Duffie, D. (2018). Financial regulatory reform after the crisis: An assessment. *Management Science*, 64:4835–4857.

Duffie, D. (2023). Resilience redux in the US Treasury market. Jackson Hole.

Duffie, D., Fleming, M., Keane, F., Nelson, C., Shachar, O., and Tassel, P. V. (2023). Dealer capacity and U.S. Treasury market functionality. New York Federal Reserve staff report No. 1070.

Elsgolts, L. (1977). *Differential Equations and the Calculus of Variation*. MIR Publishers.

ESRB (2016). Market liquidity and market-making. European Systemic Risk Board.

Ewerhart, C., Cassola, N., and Valla, N. (2010). Declining valuations and equilibrium bidding in central bank refinancing operations. *International Journal of Industrial Organization*, 28:30–43.

Fardeau, V. (2024). Arbitrage with financial constraints and market power. *Journal of Economic Theory*, 217:1–32.

Favara, G., Infante, S., and Rezende, M. (2022). Leverage regulation and Treasury market participation: Evidence from credit line drawdowns. Working paper.

Geddie, J. (2016). Squeezed bank dealers quit European government bond markets. `https://www.reuters.com/article/idUSL8N14R0GW/`, accessed 2024-04-08.

Glebkin, S., Malamud, S., and Teguia, A. (2023). Illiquidity and higher cumulants. *Review of Financial Studies*, 36(5):2131–2173.

Gromb, D. and Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66:361–407.

Group of Thirty (2021). US Treasury markets: Steps toward increased resilience. G30 Working Group on Treasury Market Liquidity, Group of 30, Washington, DC.

Guerre, E., Perrigne, I., and Vong, Q. (2009). Nonparametric identification of risk aversion in first-price auctions under exclusion restrictions. *Econometrica*, 77(4):1193–1227.

Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574.

Gupta, S. and Lamba, R. (2017). Treasury auctions during a crisis. Working paper.

Haddad, V. and Muir, T. (2021). Do intermediaries matter for aggregate asset prices? *Journal of Finance*, 76(6):2719–2761.

Häfner, S. (2023). Risk aversion in share auctions: Estimating import rents from TRQs in Switzerland. *Quantitative Economics*, 14:419–470.

He, Z., Kelly, B., and Manela, A. (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35.

He, Z. and Krishnamurthy, A. (2012). A model of capital and crises. *Review of Economic Studies*, 79(2):735–777.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103(2):732–70.

He, Z., Nagel, S., and Song, Z. (2022). Treasury inconvenience yields during the COVID-19 crisis. *Journal of Financial Economics*, 143(1):57–79.

Holmberg, P. (2009). Supply function equilibria of pay-as-bid auctions. *Journal of Regulatory Economics*, 36(2):154–177.

Hortaçsu, A. and Kastl, J. (2012). Valuing dealers' informational advantage: A study of Canadian Treasury auctions. *Econometrica*, 80(6):2511–2542.

Hortaçsu, A., Kastl, J., and Zhang, A. (2018). Bid shading and bidder surplus in the U.S. Treasury auction. *American Economic Review*, 108(1):147–69.

Hortaçsu, A. and McAdams, D. (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the Turkish Treasury auction market. *Journal of Political Economy*, 118(5):833–865.

Huber, A. W. (2023). Market power in wholesale funding: A structural perspective from the triparty repo market. *Journal of Financial Economics*, 149(2):235–259.

ISDA (2024). Letter to U.S. regulators and organizations regarding the SLR Reform - U.S. Treasuries. Available at `https://www.isda.org/a/h3sgE/ISDA-Submits-Letter-to-US-Agencies-on-SLR-Reform.pdf`.

Jamilov, R. (2021). A macroeconomic model with heterogeneous banks. Working paper.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kang, B.-S. and Puller, S. L. (2008). The effect of auction format on efficiency and revenue in divisible goods auctions: A test using Korean Treasury auctions. *Journal of Industrial Economics*, 56(2):290–332.

Kastl, J. (2011). Discrete bids and empirical inference in divisible good auctions. *Review of Economic Studies*, 78:974–1014.

Kastl, J. (2012). On the properties of equilibria in private value divisible good auctions with constrained bidding. *Journal of Mathematical Economics*, 48:339–352.

Klemperer, P. D. and Meyer, M. A. (1989). Supply function equilibria in oligopoly under uncertainty. *Econometrica*, 57(6):1243–1277.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.

Kyle, A. S. (1989). Informed speculation with imperfect competition. *Review of Economic Studies*, 56(3):317–355.

Kyle, A. S. and Xiong, W. (2001). Contagion as a wealth effect. *Journal of Finance*, 56(4):1401–1440.

Lu, L. and Wallen, J. (2024). What do bank trading desks do? Working paper.

Luo, Y. and Takahashi, H. (2023). Bidding for contracts under uncertain demand: Skewed bidding and risk sharing. *RAND Journal of Economics*, forthcoming.

OSFI (2023). Leverage requirements guideline. `http://www.osfi-bsif.gc.ca/Eng/fi-if/rg-ro/gdn-ort/gl-ld/Pages/LR22.aspx`, accessed on 01/13/2023.

Pedersen, L. H. and Gârleanu, N. (2011). Margin-based asset pricing and deviations from the law of one price. *Review of Financial Studies*, 24(6):1980–2022.

Pinter, G. and Üslü, S. (2023). Comparing search and intermediation frictions across markets. Working paper.

Popowicz, J. (2021). End of SLR relief weighs on JP Morgan. `https://www.risk.net/risk-quantum/7856711/end-of-slr-relief-weighs-on-jp-morgan`, accessed 2024-04-08.

Pycia, M. and Woodward, K. (2023). A case for pay-as-bid auctions. Working paper.

Richert, E. (2023). Indirect inference techniques as an alternative to direct computation of equilibrium bid distributions. Working paper.

Rostek, M. and Weretka, M. (2012). Price inferences in small markets. *Econometrica*, 80(2):687–711.

Rostek, M. J. and Yoon, J. H. (2020). Equilibrium theory of financial markets: Recent developments. Working paper.

Rostek, M. J. and Yoon, J. H. (2021). Exchange design and efficiency. *Econometrica*, 89(6):2887–2928.

Siriwardane, E., Sunderam, A., and Wallen, J. (2025). Segmented arbitrage. *Journal of Finance*, forthcoming.

Trebbi, F. and Xiao, K. (2017). Regulation and market liquidity. *Management Science*, 65(5):1949–1968.

Vayanos, D. and Wang, J. (2012). Theories of liquidity. *Foundations and Trends in Finance*, 6(4):221–317.

Vayanos, D. and Wang, J. (2013). Market liquidity: Theory and empirical evidence. In: The Handbook of the Economics of Finance, G. M. Constantinides, M. Harris, and R. Stulz (eds.), North-Holland Press, Volume 2b, Chapter 19.

Vives, X. (2011). Strategic supply function competition with private information. *Econometrica*, 76(6):1919–1966.

Wallen, J. (2023). Markups to financial intermediation in foreign exchange markets. Working paper.

Wang, J. and Zender, J. (2002). Auctioning divisible goods. *Economic Theory*, 19(4):673–705.

Wang, Y., Whited, T. M., Whu, Y., and Xiao, K. (2022). Bank market power and monetary policy transmission: Evidence from a structural estimation. *Journal of Finance*, 77(4):2093–2141.

Wilson, R. (1979). Auctions of shares. *Quarterly Journal of Economics*, 93(4):675–689.

Wittwer, M. (2018). *Pay-as-bid auctions in theory and practice*. PhD thesis, European University Institute.

Wittwer, M. (2021). Connecting disconnected financial markets? *American Economic Journal: Microeconomics*, 13(1):252–282.

Xiong, W. (2001). Convergence trading with wealth effects: an amplification mechanism in financial markets. *Journal of Financial Economics*, 62:247–292.

ONLINE APPENDIX

Market Power and Capital Constraints

by Milena Wittwer and Jason Allen

Appendix A generalizes our benchmark model in various ways. Appendix B provides additional information about our empirics, including statistical tests, and robustness analysis. Proofs are in Appendix C. Appendix tables and figures are at the end.

A Model extensions

We provide a series of model extensions. In Appendix A.1, we introduce inventory positions. In Appendix A.2, we explain how to adjust our framework to analyze trade settings. In Appendix A.3, we illustrate how to model and analyze different types of constraints. In Appendix A.4, we show that the main prediction of our model, which is about how price and price impact depend on capital constraints, can generalize to a discriminatory price auction.

A.1 Inventory positions

Here we explain how to extend the model to account for privately observed inventory positions. To reduce the number of parameters, we shut off the trader's signal about the return of the asset, ϵ_i, by setting $\underline{\epsilon} = \overline{\epsilon} = 0$. It is straightforward to include it.

Each trader i now holds portfolio, z_i, in inventory, which was acquired at price $p_z \in \mathbb{R}^+$.[27] The inventory is drawn iid across traders from some continuous distribution on bounded support and strictly positive density. It is private information of the trader, and may include a variety of security types. In this case, price p_z represents the average per-unit price of a security in inventory; it may be a function of other model primitives: for example, the distribution of asset supply, or the number of traders.

The inventory position affects the trader's payoff, as well as the capital constraint. Trader

[27] If the trader's inventory position only consists of the asset that is for sale in the upcoming auction, we could replace p_z with the auction clearing price, $\boldsymbol{P^c}$.

i's ex-post gross utility is:

$$V(a_i^c, z_i) = \mu a_i^c + \mu_z z_i - \frac{\rho\sigma^2}{2}[a_i^c]^2 - \frac{\rho\sigma_z^2}{2}[z_i]^2 - \rho\iota\sigma\sigma_z a_i^c z_i, \tag{14}$$

with $\mu, \mu_z, \sigma, \sigma_z > 0$, and $\iota \in [-1, 1]$. When $\mu = \mu_z, \sigma = \sigma_z, \iota = 1$ the asset in inventory is identical to the asset at auction. The capital constraint is:

$$\mathbb{E}[\boldsymbol{\theta_i} - \kappa \boldsymbol{P^c a_i^c} - \kappa p_z \boldsymbol{z_i}] \geq 0. \tag{15}$$

Proposition 1 generalizes to this setting. The only difference is that $v(a, \epsilon_i) = \mu + \epsilon_i$ is replaced by $v(a, z_i) = \mu - \rho\iota\sigma_z z_i$. That is, trader i submits the same function as in the benchmark model with $\mu - \rho\iota\sigma_z z_i$ replacing $\mu + \epsilon_i$, and an adjusted Lagrange multiplier, λ.

A.2 Trade setting

In line with the empirical application our benchmark model features a one-sided auction. Here we illustrate how to adjust the setting to accommodate two-sided markets.

The easiest adjustment is to simply let demand curves in the benchmark model, $p_i(\cdot, \epsilon_i)$, map from $\mathbb{R}$ to $\mathbb{R}$ and represent net-demand. Since all traders are ex-ante identical in our benchmark model, all of them either expect to buy from noise traders, or sell to noise traders in this model adjustment. When traders expect to be net-buyers, all model predictions generalize. When they expect to be net-sellers, the capital constraint of the benchmark model needs to be adjusted to reflect the fact that taking short positions doesn't help fulfill capital requirements.

Next, we consider two additional environments, one with symmetric market conditions for buyers and sellers, and one with asymmetric conditions. For other environments, equilibria and predictions about how capital constraints affect prices and price impact, can be derived analogously.

Our trade framework is similar to the one presented in Appendix A.1, but with two groups of traders: N buyers, indexed by B, and N sellers, indexed by S. They seek to trade units of an asset of zero aggregate supply, $A = 0$. Each trader draws an iid inventory position of the asset, $\boldsymbol{z_i^G}$, from a group-specific distribution with $\mathbb{E}[\boldsymbol{z_i^B}] \leq 0$, $\mathbb{E}[\boldsymbol{z_i^S}] > 0$, and

$\mathbb{E}[\boldsymbol{z_i^B}] + \mathbb{E}[\boldsymbol{z_i^S}] > 0$. In addition, we could assume that the support of the seller's distribution is sufficiently negative to make them want to sell with certainty, and vice versa for buyers.

Ex-post, trader i of group G's gross utility is given by (14) with $\mu_z = \mu, \sigma_z = \sigma, \iota = 1$. Each trader face the same capital constraint (15), i.e.,

$$\mathbb{E}[\boldsymbol{\theta_i^G} - \kappa \boldsymbol{P^c a_i^G} - \kappa p_z \boldsymbol{z_i^G}] \geq 0 \text{ for } G \in \{B, S\}, \tag{16}$$

with equity capital, $\boldsymbol{\theta_i^G}$, which is drawn iid from a group-specific distribution.

First, let us consider a symmetric environment, in which the capital constraint is binding in the same way for buyers and sellers in the symmetric equilibrium. For this to be the case, the distribution of capital positions and inventory positions must be such that: $\mathbb{E}[\boldsymbol{\theta_i^B} + \boldsymbol{\theta_i^S}] = \kappa p_z \mathbb{E}[\boldsymbol{z_i^B} + \boldsymbol{z_i^S}] > 0$.

Proposition 4. *Consider a symmetric setting where* $\mathbb{E}[\boldsymbol{\theta_i^B} + \boldsymbol{\theta_i^S}] = \kappa p_z \mathbb{E}[\boldsymbol{z_i^B} + \boldsymbol{z_i^S}] > 0$. *There is a symmetric equilibrium in which traders submit the following net-demand curve:*

$$p(a, z_i^G) = \frac{1}{1 + \lambda\kappa}\left(\mu - \rho\sigma^2 z_i^G - \rho\sigma^2\left(\frac{2N-1}{2N-2}\right)a\right) \quad \textit{and} \tag{17}$$

$$\lambda = \begin{cases} 0 & \textit{if } \mathbb{E}[\boldsymbol{\theta_i^B} - \kappa p_z \boldsymbol{z_i^B}] \geq \tilde{\theta} \\ \frac{1}{\kappa}\left(\frac{\tilde{\theta}}{\mathbb{E}[\boldsymbol{\theta_i^B} - \kappa p_z \boldsymbol{z_i^B}]} - 1\right) & \textit{otherwise.} \end{cases} \tag{18}$$

with $\tilde{\theta} = \frac{\kappa(N-1)}{2N(2N-1)}\Big[2N\mu(\mathbb{E}[\boldsymbol{z_i^S}] - \mathbb{E}[\boldsymbol{z_i^B}]) - \rho\sigma^2\Big(\mathbb{E}[(\boldsymbol{z_i^S})^2] - \mathbb{E}[(\boldsymbol{z_i^B})^2] + (N-1)[\mathbb{E}[\boldsymbol{z_i^S}]^2 - \mathbb{E}[\boldsymbol{z_i^B}]^2]\Big)\Big]$. In this equilibrium, buyer i wins $a_i^{B*} = \frac{N-1}{N(2N-1)}\left(\sum_i z_i^B + \sum_i z_i^S - 2N z_i^B\right)$ and seller i wins $a_i^{S*} = \frac{N-1}{N(2N-1)}\left(\sum_i z_i^B + \sum_i z_i^S - 2N z_i^S\right)$.

The market clears at $P^* = \frac{1}{1+\lambda\kappa}\left(2N\mu - \rho\sigma^2\left(\sum_i z_i^B + \sum_i z_i^S\right)\right)$, and each trader's price impact is $\Lambda = \frac{\rho\sigma^2}{1+\lambda\kappa}\frac{1}{2N-2}$. Therefore, it is still the case that the price and price impact increases when constraints are relaxed—Corollary 1.

Proposition 5. *Consider an asymmetric setting where* $\mathbb{E}[\boldsymbol{\theta_i^B} + \boldsymbol{\theta_i^S}] \neq \kappa p_z \mathbb{E}[\boldsymbol{z_i^B} + \boldsymbol{z_i^S}]$.
(i) In a group-symmetric equilibrium, in which buyers and sellers submit the same net-demand schedule, respectively, a trader i of group G chooses:

$$p^G(a, \epsilon_i^G) = \frac{1}{\beta^G}\left(\alpha^G + \gamma^G z_i^G - a\right) \textit{ for } G \in \{B, S\}. \tag{19}$$

Parameter β^B is the root of the following polynomial:

$$
\begin{aligned}
0 = &\big(2(1+\lambda^B\kappa)^2(1+\lambda^S\kappa)2N(N-1)\big)\\
&+\big((1+\lambda^B\kappa)((1+\lambda^B\kappa)(N-2)^2+2(1-2\lambda^B\kappa+3\lambda^S\kappa+(\lambda^B\kappa-2\lambda^S\kappa-1)N)N-3(1+\lambda^S\kappa)N^2)\big)(\rho\sigma^2\beta^B)^1\\
&+\big(-2(1+\lambda^B\kappa)(N-2)(N-1)+(2+\lambda^B\kappa(4-3N)+2\lambda^S\kappa(N-1)-N)N+(1+\lambda^S\kappa)N^2)\big)(\rho\sigma^2\beta^B)^2\\
&+\big((N-1)(2N-1)\big)(\rho\sigma^2\beta^B)^3. \qquad (20)
\end{aligned}
$$

The other parameters can be expressed as functions of β^B:

$$\beta^S = \frac{\beta^B}{N}\left((N-1)+\frac{1+\lambda^B\kappa}{1+\lambda^B\kappa-\beta^B\rho\sigma^2}\right), \qquad (21)$$

$$\alpha^G = \frac{\beta^G(N-1)+\beta^{-G}N}{1+\lambda^G\kappa+\beta^G(N-1)\rho\sigma^2+\beta^{-G}N\rho\sigma^2}\mu, \qquad (22)$$

$$\gamma^G = -\frac{\beta^G(N-1)+\beta^{-G}N}{1+\lambda^G\kappa+\beta^G(N-1)\rho\sigma^2+\beta^{-G}N\rho\sigma^2}\rho\sigma^2 < 0 \text{ for } G\in\{B,S\} \text{ and } G\neq G. \qquad (23)$$

(ii) This equilibrium exists if the exogenous model parameters are such that $\beta^B>0, \beta^S>0$, and the Lagrange multipliers that solve the binding capital constraints at market clearing are non-negative.

In this equilibrium, the market clears at $P^* = \frac{1}{\beta^B+\beta^S}\left(\sum_G \alpha^G + \frac{1}{N}\sum_G \gamma^G \sum_i z_i^G\right)$, and a trader's price impact is $\Lambda^G = (N\beta^{-G}+(N-1)\beta^G)^{-1}$ for $G\in\{B,S\}$ and $G\neq G$. While it is challenging to provide general statements about how constraints affect the price and price impact, it is straightforward to solve for an equilibrium numerically, and verify the price and price impact effect.

For example, let $N=3, \sigma=\rho=\kappa=\mu=p_z=\mathbb{E}[\boldsymbol{z^S}]=\mathbb{E}[(\boldsymbol{z_i^B})^2]=\mathbb{E}[(\boldsymbol{z_i^S})^2]=1$, $\mathbb{E}[\boldsymbol{z_i^B}]=0$, $\mathbb{E}[\boldsymbol{\theta_i^B}]\approx 0.239$ and $\mathbb{E}[\boldsymbol{\theta_i^S}]$ sufficiently high that $\lambda^S=0$. Then, there exists an equilibrium with $\alpha^B=0.791, \alpha^S=0.808, \beta^B\approx 0.870, \beta^S\approx 0.808, \gamma^B\approx -0.791, \gamma^S\approx -0.808$ and $\lambda^B\approx 0.1$, where we write $\approx$ instead of $=$ to highlight the fact that we are rounding numbers. Appendix Figure A3 shows how the price and price impact of buyers and sellers increase when the constraint for buyers is relaxed in this example.

A.3 Other constraints

Although our focus lies on a particular type of constraint, motivated by Basel III, our framework extends to a wider set of constraints.

Alternative ex-ante constraints. Following the same approach as in the proof of Proposition 1, we can characterize necessary conditions in the presence of the following constraint: $\mathbb{E}[h(\boldsymbol{P^c}, \boldsymbol{a_i^c})] \geq 0$, for any $h(\cdot, \cdot)$ that is differentiable in both input arguments.[28] To solve for an equilibrium in explicit form and prove equilibrium existence, as in Proposition 2, function $h(\cdot, \cdot)$ must take a linear form: $h(P, a) = \Gamma + \Xi a + \Omega p + \Upsilon P a$, with $\Gamma, \Xi, \Omega, \Upsilon \in \mathbb{R}$. This class of constraints includes wealth or budget constraints ($\Gamma > 0, \Xi = \Omega = 0, \Upsilon > 0$), quantity constraints that might come from bidding limits ($\Gamma > 0, \Xi \neq 0, \Omega = \Upsilon = 0$), or constraints on the clearing price which might arise from arbitrage in an outside market ($\Gamma > 0, \Xi = 0, \Omega \neq 0, \Upsilon = 0$) .

Imposing a linear functional form on $h(\cdot, \cdot)$ ensures that there are linear equilibria. In absence of linearity, it becomes challenging to prove that a demand function that fulfills necessary equilibrium conditions is a global optimum. The reason for this is that many sufficient conditions one can derive when maximizing over functions only hold locally (see, for instance, Elsgolts (1977), Chapter 8). To overcome this challenge, the proof of Proposition 2 relies on the property that the objective functional is for any demand function (not just the equilibrium candidate) strictly concave. This must not be the case when the equilibrium candidate is non-linear.

Ex-post capital constraints. In some real-world applications constraints must hold with certainty, i.e., in all states of the world. Here, we characterize equilibrium conditions and illustrate how changes in the constraint can affect the price and price impact when the constraint must hold ex-post. For tractability, we let all traders have the same position, $\theta_i = \theta$ for all i. This allows us to solve for symmetric equilibria.

The trader's maximization problem—the analogue to problem (4)— reads as follows:

$$\max_{p_i(\cdot,\cdot)\in\mathcal{B}} \mathbb{E}\big[V(\boldsymbol{a_i^c}, \boldsymbol{\epsilon_i}) - \boldsymbol{P^c a_i^c}\big] \text{ subject to}$$

$$\text{the capital constraint:} \qquad \theta - \kappa P^c a_i^c \geq 0,$$

$$\text{and market clearing:} \qquad P^c = p_i(a_i^c, \boldsymbol{\epsilon_i}) \text{ for all } a_i^c. \tag{24}$$

[28] It suffices to replace $H(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon)$ in expression (37) in the proof of Proposition 1 with $\left(\frac{\partial h(p(a,\epsilon),a)}{\partial p(a,\epsilon)} p_a(a,\epsilon) + \frac{\partial h(p(a,\epsilon),a)}{\partial a}\right)[1 - G(a, p(a,\epsilon)|\epsilon)]\psi(\epsilon)$, where $p_a(a,\epsilon) = \frac{\partial p(a,\epsilon)}{\partial a}$.

Proposition 6. *Let gross utility $V(a, \epsilon_i)$ be any measurable and bounded function that is continuously differentiable and strictly concave in a for all ϵ_i, and bounded in ϵ_i for all a.*

In any symmetric equilibrium, trader i submits demand function, $p^*(\cdot, \cdot)$, that for all p, a, and ϵ_i satisfies

$$p = v(a, \epsilon_i) + \lambda a \left[\left(\frac{\partial G(a, p|\epsilon_i)}{\partial p} \right) \psi(\epsilon_i) \right]^{-1} - \underbrace{a \left(\frac{\partial G(a, p|\epsilon_i)}{\partial a} \Big/ \frac{\partial G(a, p|\epsilon_i)}{\partial p} \right) (-1)}_{shading}, \tag{25}$$

where $\lambda \geq 0$ is the Lagrange multiplier of the capital constraint, and $G(a, p|\epsilon_i) = \Pr(\boldsymbol{a_i^c} \leq a|\epsilon_i)$ with $\boldsymbol{a_i^c} = \boldsymbol{A} - \sum_{j \neq i} a^*(p, \boldsymbol{\epsilon_j})$ is the probability that trader i, who bids price $p = p^*(a, \epsilon_i)$, wins less than a at market clearance, given that the other traders play the equilibrium strategy.

With ex-post capital constraints it is no longer true that traders bid as if they were bidding in an auction without constraints but with a discounted value (as in Proposition 1).

To build an intuition for the equilibrium condition, we consider a simplified auction environment in which all bidders observe the same signal, and thus do not have private information. Without loss of generality, we normalize this signal to zero, i.e., $\underline{\epsilon} = \overline{\epsilon} = 0$.

Corollary 3. *Let $\underline{\epsilon} = \overline{\epsilon} = 0$, so that $\epsilon_i = 0$ for all traders i, and omit the dependence on ϵ_i in all functions. In any symmetric equilibrium, each trader submits a demand function that satisfies the following condition:*

$$[v(a) - p^*(a)]\phi(aN) = \Lambda(a) \left[\phi(aN) + \lambda(a)\kappa \right] a \text{ for all } a \in \left[0, \frac{\overline{A}}{N} \right], \tag{26}$$

$$\text{where } \Lambda(a) = -\frac{1}{N-1} \frac{p^*(a)}{\partial a}, \tag{27}$$

and $v(a) = \mu - \rho\sigma^2 a$. The unique solution to this differential equation is:

$$p^*(a) = e^{\mathcal{M}(a)} \left(c_1 + \int_1^a \frac{-e^{-\mathcal{M}(y)}(N-1)(\mu - \rho\sigma^2 x)\phi(Nx)}{y(\kappa\lambda(y) + \phi(Ny))} dy \right), \tag{28}$$

with $\mathcal{M}(z) = \int_1^z \frac{(N-1)\phi(Nx)}{x(\kappa\lambda(x)+\phi(Nx))} dx$ for $z = a, y$, $c_1 \in \mathbb{R}$, and $\lambda(a) = \lambda(A/N)$ given by the capital constraint for all $A \in (0, \overline{A}]$. Note that with $c_1 = \mu - \rho\sigma^2 \left(\frac{N-1}{N-2} \right)$ the demand function simplifies to the known solution in an unconstraint auction for $\kappa = 0$.

Intuitively, condition (26) says that the bidder trades off the expected marginal surplus when winning at price p with the expected marginal cost. The expected marginal surplus is the difference between the bidder's value and the price they pay, weighted by the probability density that the market clears at price p and the bidder wins a. The expected marginal cost has two parts. The first part is the increased payments in all states with higher realizations of supply, which is equal to the price impact times the amount won. This is weighted by the probability density that the market clears at (p, a). The second part, $\Lambda(a)\lambda(a)\kappa a$, reflects the marginal cost of the constraint.[29]

Corollary 4. *(i) In the symmetric equilibrium, price impact at amount a is:*

$$\Lambda(a) = \frac{\phi(Na)}{\phi(aN) + \kappa\lambda(a)} \frac{v(a) - p^*(a)}{a} \geq 0. \tag{29}$$

(ii) Fix an amount a. When $\lambda(a)$ increases, price impact at amount a increases as long as $\frac{\partial p^*(a)}{\partial \lambda(a)} < \frac{\phi(Na)}{(\phi(aN)+\kappa\lambda(a))^2}(v(a) - p^*(a))$.

Both the shadow cost of the constraint, $\lambda(a)$, and price impact, as shown in equation (29), are functions of quantity, since the constraint must hold for all states of the world, i.e., for all quantities that can be obtained. When $\lambda(a)$ at a fixed amount a increases, price impact at that point changes due to two opposing effects. A direct effect pushes price impact upwards, while an indirect effect , coming from a change in $p^*(a)$ in response to the changing shadow cost, pushes price downwards. The total change in price impact is positive if the direct effect dominates the indirect effect. This is the case when $\frac{\partial p^*(a)}{\partial \lambda(a)} < \frac{\phi(Na)}{(\phi(aN)+\kappa\lambda(a))^2}(v(a) - p^*(a))$—summarized in Corollary 4 (ii).

A.4 Discriminatory price auction

Here we adjust our benchmark model to the case of discriminatory price auctions, in which bidders pay the prices they offered to pay for all units won, rather than the market clearing

[29]To see why $\Lambda(a)\lambda(a)\kappa a$ is the marginal cost of the constraint, note that, at any fixed $\lambda \geq 0$, the total cost from the constraint is $-\lambda(\theta - \kappa p a)$ at all (a, p) at which the market clears. By market clearing, $p_i^{RS}(a) = p$, where $p_i^{RS}(\cdot)$ is the inverse of trader i's residual supply curve $RS_i(\cdot)$. Taking the derivative of $-\lambda(\theta - \kappa p_i^{RS}(a)a)$ with respect to a, and using the fact that $\frac{\partial p_i^{RS}(a)}{\partial a} = (\frac{\partial RS_i(p)}{\partial p})^{-1} = \Lambda(a)$, we obtain the marginal cost of the constraint.

price. The differences in the payment does not only affect strategic incentives when bidding, but also the capital constraint. Since bidders pay the prices they offer, and winning bids are weakly higher than the auction clearing price, the nominal value of the amount bidders win at auction is the total amount bidders pay, i.e., $\int_0^{\boldsymbol{a_i^c}} p_i(a, \epsilon_i)da$. This is higher than the auction clearing price times the amount won, $\boldsymbol{P^c a_i^c}$. With this, the capital constraint, which is analogue to the benchmark constraint in (4), is:

$$\mathbb{E}\left[\boldsymbol{\theta_i} - \kappa \int_0^{\boldsymbol{a_i^c}} p_i(a, \boldsymbol{\epsilon_i})da\right] \geq 0. \tag{30}$$

In practice, this constraint approximates two scenarios. First, traders buy bonds at auction to keep them on their own balance sheet. In this case, bonds are evaluated according to their book-value, that is, according to the prices paid in the acquisition. Second, traders buy bonds to later sell them to clients. These bonds must be kept on the trading book, and evaluated according to their market-value. The market-value could deviate from the prices traders paid at auction. The deviation is small if traders charge clients prices that are close to their own transaction costs. To account for potential markups, we could multiply the auction payment by a scaler.

We consider two settings: one with continuous demand functions, and one with step functions.

(i) Continuous demand. In the first setting, traders submit continuous demand functions, as in our benchmark model.

Proposition 7. *Let gross utility $V(a, \epsilon_i)$ be any measurable and bounded function that is continuously differentiable and strictly concave in a for all ϵ_i, and bounded in ϵ_i for all a.*

In any symmetric equilibrium, traders behave as if they were bidding in an auction without capital constraints, in which their willingness to pay is $\tilde{v}(a, \epsilon_i)$—given in expression (5)—instead of $v(a, \epsilon_i)$, where $\lambda \geq 0$ is the Lagrange multiplier of the capital constraint. Trader i submits demand function, $p^*(\cdot, \cdot)$, that for all p, a, and ϵ_i satisfies

$$p = \tilde{v}(a, \epsilon_i) - \underbrace{(1 - G(a, p|\epsilon_i))/\frac{\partial G(a, p|\epsilon_i)}{\partial p}(-1)}_{shading}, \tag{6'}$$

where $G(a, p|\epsilon_i) = \Pr(\boldsymbol{a_i^c} \leq a|\epsilon_i)$ *with* $\boldsymbol{a_i^c} = \boldsymbol{A} - \sum_{j \neq i} a^*(p, \boldsymbol{\epsilon_j})$ *is the probability that trader* i*, who bids price* $p = p^*(a, \epsilon_i)$*, wins less than* a *at market clearance, given that the other traders play the equilibrium strategy.*

Proposition 7 is the analogue to Proposition 1; only the shading factor is computed differently due to differences in the payment rules of the auction formats. To show the similarities across auction formats, we use the same notation for $p^*(\cdot, \epsilon_i)$, $G(a, p|\epsilon_i)$, and the Lagrange multiplier, λ, even though all of these are auction-format specific.

Solving for equilibrium demand functions in discriminatory price auctions, and proving equilibrium existence, is much more challenging than in uniform price auctions, even without constraints. This might be one of the reasons for which the literature has focussed on uniform price auctions. Only recently, Pycia and Woodward (2023) have proved pure-strategy equilibrium existence and uniqueness in an auction environment with identical bidders without private information. Earlier contributions that make similar informational restrictions, including Wang and Zender (2002), Holmberg (2009), Ewerhart et al. (2010), and Ausubel et al. (2014), impose additional distributional assumptions. For example, Ausubel et al. (2014) show that there exists a linear equilibrium if and only if supply follows a Generalized Pareto Distribution (GPD). Wittwer (2018) extends their ideas to the case with private information.

Providing novel equilibrium existence results for discriminatory price auctions without capital constraints is beyond the scope of our paper, which focusses on capital constraints. To nevertheless make progress, we follow Wittwer (2018) and impose conditions on equilibrium allocations, which is not ideal. Characterizing conditions on the underlying exogenous model primitives—i.e., the distribution of supply and signals—is challenging when traders have private information. Without private information (when $\overline{\epsilon} = \underline{\epsilon} = \epsilon_i = 0$ for all traders i) we can derive a sharper result: Equilibrium (7') of Proposition 8 exists when supply, $\boldsymbol{A}$, follows a GPD with CDF (31) and $\xi \leq -1$.

Proposition 8. *Let gross utility* $V(a, \epsilon_i)$ *be given by expression (1). Suppose there is a symmetric linear equilibrium in which each trader submits demand function*

$$p^*(a, \epsilon_i) = \frac{1}{1 + \lambda\kappa} \left(\mu + \epsilon_i - \rho\sigma^2 \left(\frac{N-1}{N(1-\xi) - 1} \right) [a + N\nu(\epsilon_i)] \right), \text{ with} \tag{7'}$$

$$\lambda = \begin{cases} 0 & \text{if } \mathbb{E}[\boldsymbol{\theta_i}] \geq \tilde{\theta} \\ \frac{1}{\kappa}\left(\frac{\tilde{\theta}}{\mathbb{E}[\boldsymbol{\theta_i}]} - 1\right) > 0 & \text{otherwise,} \end{cases}$$

where

$$\tilde{\theta} = \frac{\kappa}{N}\Big[\mathbb{E}[\boldsymbol{A}]\left(\mu + \frac{\mathbb{E}[\boldsymbol{\epsilon_i}](1-(1-\sigma^2)\xi)}{1-\xi}\right) + \mathbb{E}[\boldsymbol{A}]\left(\frac{\sigma^2(\underline{\epsilon} + \underline{\epsilon}N(\xi-1) - \overline{A}\rho(\xi-1))\xi}{(1+N(\xi-1)(\xi-1)}\right) + \mathbb{E}[\boldsymbol{A^2}]\left(\frac{(N-1)\rho\sigma^2}{N(1+N(\xi-1))}\right)\Big].$$

This equilibrium exists if the joint distribution of supply and private signals is such that, given all traders submit the demand function (7'), each trader's allocation at the market-clearing price follows a GPD with CDF

$$\Psi_i(a) = 1 - \left(\frac{\nu(\epsilon_i) + \xi a}{\nu(\epsilon_i)}\right)^{-\frac{1}{\xi}}, \tag{31}$$

where $\nu(\epsilon_i) = -\xi\frac{N(1-\xi)-1}{N(1-\xi)\rho}(\epsilon_i - \underline{\epsilon}) - \xi\frac{\overline{A}}{N}$, *and* $\xi \leq -1$.

Propositions 7 and 8 are analogues to Propositions 1 and 2. Therefore, Corollaries 1 and 2 generalize to discriminatory price auctions when equilibria are linear. Price impact, defined as the inverse slope of a bidder's residual supply curve, which is equivalent to $\Lambda^{DPA} = \frac{1}{N-1}\frac{\partial p^*(a,\epsilon_i)}{\partial a}$, is now given by expression (10').

(ii) Step function demand. In the second setting, bidders submit step functions. This is the case in many real-world applications, including Canadian Treasury auctions. It is therefore useful to characterize equilibrium conditions for this environment. For this, we adopt Kastl (2012)'s framework—i.e., Assumptions 1-6 and the equilibrium definition of his paper—with three adjustments. First, we introduce capital constraint, (30). Second, we follow the timing of events of our benchmark model, which implies that each bidder chooses a bidding function for all possible signals prior to observing the signal. Third, for simplicity, we assume that the value function, $v(\cdot, \epsilon_i)$, is strictly decreasing for all ϵ_i, while Assumption 2 in Kastl (2012) allows it to be weakly decreasing. Different to Kastl (2012), we refer to the private signal by ϵ_i, rather than s_i, and let a price bid at step k be $p_k(\epsilon_i)$ instead of b_{ik}, in line with the rest of our paper.

Consider a bidder who observes signal ϵ_i and let all other bidders play the equilibrium strategy. Bidder i obtains the following expected payoff when submitting step function

$\{p_k(\epsilon_i), a_k(\epsilon_i)\}_{k=1}^{K(\epsilon_i)}$ conditional on observing ϵ_i:

$$EU(\epsilon_i) = EV(\epsilon_i) - EP(\epsilon_i), \tag{32}$$

with

$$\begin{aligned} EV(\epsilon_i) = &\sum_{k=1}^{K(\epsilon_i)} \Pr(p_k(\epsilon_i) > \boldsymbol{P^c} > p_{k+1}(\epsilon_i)|\epsilon_i) V(a_k(\epsilon_i), \epsilon_i) \\ &+ \sum_{k=1}^{K(\epsilon_i)} \Pr(p_k(\epsilon_i) = \boldsymbol{P^c}|\epsilon_i) \mathbb{E}[V(\boldsymbol{a_i^c}, \epsilon_i)|p_k(\epsilon_i) = \boldsymbol{P^c}, \epsilon_i], \text{ and} \end{aligned} \tag{33}$$

$$\begin{aligned} EP(\epsilon_i) = &\sum_{k=1}^{K(\epsilon_i)} \Pr(p_k(\epsilon_i) > \boldsymbol{P^c}|\epsilon_i) p_k(\epsilon_i)(a_k(\epsilon_i) - a_{k-1}(\epsilon_i)) \\ &+ \sum_{k=1}^{K(\epsilon_i)} \Pr(p_k(\epsilon_i) = \boldsymbol{P^c}|\epsilon_i) \mathbb{E}[p_k(\epsilon_i)(\boldsymbol{a_i^c} - a_{ik-1})|p_k(\epsilon_i) = \boldsymbol{P^c}, \epsilon_i]. \end{aligned} \tag{34}$$

As before $\boldsymbol{P^c}$ is the market clearing price, and $\boldsymbol{a_i^c}$ is the amount bidder i wins at market clearing. The bidder faces capital constraint (30). With step functions, it looks as follows:

$$\begin{aligned} \mathbb{E}\Big[\boldsymbol{\theta_i} - \kappa\Big[&\sum_{k=1}^{K(\boldsymbol{\epsilon_i})} \Pr(p_k(\boldsymbol{\epsilon_i}) > \boldsymbol{P^c}|\boldsymbol{\epsilon_i}) p_k(\boldsymbol{\epsilon_i})(a_k(\boldsymbol{\epsilon_i}) - a_{k-1}(\boldsymbol{\epsilon_i})) \\ &+ \sum_{k=1}^{K(\boldsymbol{\epsilon_i})} \Pr(p_k(\boldsymbol{\epsilon_i}) = \boldsymbol{P^c}|\boldsymbol{\epsilon_i}) \mathbb{E}\big[p_k(\boldsymbol{\epsilon_i})(\boldsymbol{a_i^c} - a_{ik-1})|p_k(\boldsymbol{\epsilon_i}) = \boldsymbol{P^c}, \boldsymbol{\epsilon_i}\big]\Big]\Big] \geq 0. \end{aligned} \tag{35}$$

The bidder chooses their bidding function to maximize the expectation of (32) subject to (35); market clearing is already guaranteed. Simplifying the objective functional, we obtain the following maximization problem:

$$\max_{\{p_k(\cdot), a_k(\cdot)\}_{k=1}^{K(\cdot)}} \mathbb{E}\Big[EV(\boldsymbol{\epsilon_i}) - (1+\lambda) EP(\boldsymbol{\epsilon_i})\Big], \text{ with } \lambda \geq 0. \tag{36}$$

Similar to the benchmark model, choosing $\{p_k(\epsilon_i), a_k(\epsilon_i)\}_{k=1}^{K(\epsilon_i)}$ to maximize $EV(\epsilon_i) - (1+\lambda)EP(\epsilon_i)$ for all ϵ_i point-wise is equivalent to choosing $\{p_k(\cdot), a_k(\cdot)\}_{k=1}^{K(\cdot)}$ to maximize $\mathbb{E}[EV(\boldsymbol{\epsilon_i}) - (1+\lambda)EP(\boldsymbol{\epsilon_i})]$ for any fixed $\lambda \geq 0$. Therefore, the only substantial difference to Kastl (2012)'s framework without capital constraints is that the bidder's marginal cost from winning a step is no longer just the price they need to pay, $p_k(\epsilon_i)$, but it is the price plus the shadow cost

of the constraint, which is $\lambda\kappa p_k(\epsilon_i)$.

B Additional empirical tests and findings

In Appendix B.1, we explain how we test that the slope coefficients of Figure 4 are higher during the exemption period than during regular time. In Appendix B.2, we summarize our robustness analysis.

B.1 Testing for steepness in the willingness to pay

We conduct a t-test to compare the median slope, β_t, of regression (13) during versus outside the exemption period, by relying on our 100 bootstrapped value estimates, and the assumption that the error in regression (13) comes entirely from measurement noise in the value estimates. For each bootstrap draw d, we estimate regression (13) and collect all β_t^d estimates. We compute the median of these slope parameters across auctions during versus outside the exemption period and take the difference, $\Delta\beta_d$. We repeat this exercise for all bootstrapped values to obtain 100 $\Delta\beta_d$ differences—Appendix Figure A4 shows the distribution of $\Delta\beta_d$'s across draws. With these differences, we compute the t-statistic as the ratio of $\Delta\beta_d$'s mean and their standard deviation, which is divided by the square root of the number of bootstrap draws. The t-statistic is 81.65, so that we can conclude that the difference in slope estimates in Figure 4 is statistically significant.

B.2 Robustness analysis

We conduct a series of robustness checks to validate our risk aversion and shadow cost estimates.[30] We report the median and mean of the risk aversion estimates, in addition to the median and mean of shadow cost estimates, including the respective mean and median of standard errors of the estimates for all regression specifications we estimate in Appendix Table A2. Generally, the magnitude of the risk aversion parameters is affected when changing the units of σ_t^2 or a_{tik} in regression (13), while the magnitude of the shadow cost estimates changes depending on the cross-auction distribution of σ_t^2 during regular times. The higher

[30] A robustness analysis of the value estimates is provided in Allen et al. (2024).

the variance of σ_t^2 across auctions, the more disperse the shadow cost estimates, and the larger the mean due to a few outlying observations.

We start by analyzing the sensitivity of our parameter estimates to the number of steps included in the values functions. In our benchmark specification, we include all functions with at least two steps (which are essentially all functions) to avoid a potential bias coming from omitting functions. Given that we linearly interpolate between steps using our model, we might be concerned about doing this when there are few steps. Our results, however, are qualitatively robust to using value functions with more steps—three to six, where we do not include robustness for seven steps since not all dealers use the maximum allowable number of steps in all auctions.

Next, we estimate equation (13) with quantities expressed in million C$. In our benchmark specification, we normalize quantities by the auction supply to avoid our estimates being affected by the fact that the Bank of Canada issued larger amounts of debt during the exemption period than in regular times. Given that dealers have an obligation to actively participate in the auctions, the increased supply implies that dealers demanded larger amounts during the exemption period (see Appendix Figure A5). Further, since dealers are supposed to bid competitively, and are given a price range when bidding, increasing the total demand decreases the slope in the dealer's bidding function and willingness to pay during the exemption period (relative to the case in which we normalize demand by the supply). The model rationalizes smaller slopes by smaller risk aversion and shadow cost parameters.

Third, we verify robustness with respect to our measure of volatility. We start by smoothing outliers of the return volatility by winsorizing the distribution of volatilities by 2.5% and 5%, respectively. Decreasing the cross-auction variance of the volatility measure results in lower average shadow costs. Similarly, there are fewer large shadow cost outliers, leading to a lower average shadow cost than in our benchmark, when we use the inter-quantile range of yields at which dealers sell prior to the auction instead of the variance of these yields.

Next, we vary the number of trading days we include to measure volatility. In our benchmark specification, we construct volatility using trades where we observe dealers selling the to-be-auctioned security in a five trading day window prior to auction. This is natural given that most pre-auction trading occurs in the one week between the tender open call and

the auction close. The more days we include, the larger the volatility. Since this effect is stronger during the exemption period than during regular times, the average shadow cost is higher than in our benchmark. When we include fewer days to construct the volatility this effect goes in the opposite direction, pushing the average shadow cost downward. However, there is an additional effect. The fewer days we include, the more likely it becomes that the volatility index is missing for some auctions. To avoid dropping these auctions entirely, we use the average volatility of same maturity-type auctions within the year—in our benchmark specification there is no need to do this. This interpolation pushes the average shadow cost estimates upwards.

In addition, we can estimate our model using different volatility indices. One alternative is to use the Implied Volatility Index for Canadian Treasuries, which measures the expected volatility in the Treasury market over the next 30 days (Chang and Feunou (2014)). Given that this volatility drops more strongly during the exemption period than our volatility index, shadow cost estimates are higher when relying on the implied volatility.

Another alternative is to construct return volatility using post-auction trades. We refrain from doing so, because dealers do not know what happens after the auction at the time they bid. Further, post-auction prices likely depend on the realization of the dealer's private information, and with that their willingness to pay, in the auction. This implies that the post-auction volatility—an independent variable in equation (13)—is a function of the dependent variable and would lead to a simultaneous equation bias.

C Proofs

We first present the proofs of all propositions, and then of all corollaries.

Proof of Proposition 1. Let $V(a, \epsilon_i)$ be any measurable and bounded function that is continuously differentiable and strictly concave in a for all ϵ_i, and bounded in ϵ_i for all a.

Consider trader i, and fix all other demand schedules at the equilibrium. To determine the best response, trader i solves maximization problem (4). To simplify this problem, we drop the i-subscript, let $v(a, \epsilon) = \frac{\partial V(a,\epsilon)}{\partial a}$, denote $p_a(a, \epsilon) = \frac{\partial p(a,\epsilon)}{\partial a}$, and abbreviate all functions, for instance, $p(\cdot, \cdot)$ by p when useful. Further, we let $\overline{a}^*(\epsilon)$ be the largest amount the bidder wins when playing the equilibrium strategy, and recall that $\psi(\cdot)$ is the density function of

signals on support $[\underline{\epsilon}, \overline{\epsilon}]$. With this, and auxiliary distribution $G(a, p|\epsilon)$, which is defined in Proposition 1, the trader's maximization problem becomes:

$$\max_{p \in \mathcal{B}} I(p) \text{ subject to } L(p) \geq 0, \text{ with} \tag{37}$$

$$I(p) = \int_{\underline{\epsilon}}^{\overline{\epsilon}} \int_0^{\overline{A}} F(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon) da d\epsilon \tag{38}$$

$$\text{with } F(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon) = [v(a,\epsilon) - p(a,\epsilon) - a p_a(a,\epsilon)][1 - G(a, p(a,\epsilon)|\epsilon)]\psi(\epsilon),$$

$$L(p) = \int_{\underline{\epsilon}}^{\overline{\epsilon}} \int_0^{\overline{A}} H(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon) da d\epsilon \tag{39}$$

$$\text{with } H(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon) = [\mathbb{E}[\boldsymbol{\theta}] - \kappa[p(a,\epsilon) + p_a(a,\epsilon)a][1 - G(a, p(a,\epsilon)|\epsilon)]]\psi(\epsilon).$$

Here we have integrated the inner integral by parts to obtain $I(p)$ and $L(p)$.[31]

A function p^* is optimal if the following conditions are satisfied:

$$\frac{\partial(F + \lambda H)}{\partial p}(p^*(a,\epsilon), p_a^*(a,\epsilon), a, \epsilon) - \frac{d}{da}\left(\frac{\partial(F + \lambda H)}{\partial p_a}(p^*(a,\epsilon), p_a^*(a,\epsilon), a, \epsilon)\right) = 0 \tag{40}$$

for all $a \in [0, \overline{a}^*(\epsilon)]$, and all ϵ. In addition, we need:

$$L(p^*) \geq 0 \text{ and } \lambda \geq 0, \tag{41}$$

$$\frac{\partial(F + \lambda H)}{\partial p_a}(p^*(0,\epsilon), p_a^*(0,\epsilon), 0, \epsilon) = \frac{\partial(F + \lambda H)}{\partial p_a}(p^*(\overline{a}^*(\epsilon),\epsilon), p_a^*(\overline{a}^*(\epsilon),\epsilon), \overline{a}^*, \epsilon) = 0 \text{ for all } \epsilon. \tag{42}$$

Conditions (42) are the natural boundary conditions. They hold automatically given that $\frac{\partial(F+\lambda H)}{\partial p_a}(p^*(a,\epsilon), p_a^*(a,\epsilon), a, \epsilon) = -(1 + \lambda\kappa)a[1 - G(a, p^*(a,\epsilon)|\epsilon)]\psi(\epsilon)$, and $G(0, p^*(0,\epsilon)|\epsilon) = 0$, and $G(\overline{a}^*(\epsilon), p^*(\overline{a}^*(\epsilon),\epsilon)|\epsilon) = 1$ for all ϵ by definition of G.

[31] To simplify the objective function $I(p) = \int_{\underline{\epsilon}}^{\overline{\epsilon}} \int_0^{\overline{A}} [V(a,\epsilon) - p(a,\epsilon)a] g(a, p(a,\epsilon)|\epsilon) da \psi(\epsilon) d\epsilon$ we integrate the inner integral by parts, as follows: $\int_0^{\overline{A}} [V(a,\epsilon) - p(a,\epsilon)a] g(a, p(a,\epsilon)|\epsilon) da = [v(a,\epsilon) - p(a,\epsilon) - a p_a(a,\epsilon)] G(a, p(a,\epsilon)|\epsilon)\big|_0^{\overline{A}} - \int_0^{\overline{A}} [v(a,\epsilon) - p(a,\epsilon) - a p_a(a,\epsilon)] G(a, p(a,\epsilon)|\epsilon) da$. Since $G(\overline{A}, p(\overline{A},\epsilon)|\epsilon) = 1$ and $G(0, p(0,\epsilon)|\epsilon)) = 0$ for any function p by definition of G, this simplifies to $\int_0^{\overline{A}} [v(a,\epsilon) - p(a,\epsilon) - a p_a(a,\epsilon)][1 - G(a, p(a,\epsilon)|\epsilon)] da$. Similarly, for the constraint $L(p)$ we use $\int_0^{\overline{A}} [\mathbb{E}[\boldsymbol{\theta}] - \kappa p(a,\epsilon)]\, da = [\mathbb{E}[\boldsymbol{\theta}] - \kappa[p(a,\epsilon) + a p_a(a,\epsilon)]] G(a, p(a,\epsilon)|\epsilon)\big|_0^{\overline{A}} - \int_0^{\overline{A}} [\mathbb{E}[\boldsymbol{\theta}] - \kappa[p(a,\epsilon) + a p_a(a,\epsilon)]] G(a, p(a,\epsilon)|\epsilon) da = \int_0^{\overline{A}} [\mathbb{E}[\boldsymbol{\theta}] - \kappa[p(a,\epsilon) + a p_a(a,\epsilon)]][1 - G(a, p(a,\epsilon)|\epsilon)] da$.

Dividing (40) by $\psi(\epsilon)$ since $\psi(\epsilon) > 0$ for any ϵ, and simplifying gives:

$$-(1+\lambda\kappa)[1-G(a,p^*(a,\epsilon)|\epsilon)] - [v(a,\epsilon) - (1+\lambda\kappa)(p_a^*(a,\epsilon)a + p^*(a,\epsilon))]\frac{\partial G(a,p^*(a,\epsilon)|\epsilon)}{\partial p} + \frac{d}{da}\big((1+\lambda\kappa)a[1-G(a,p^*(a,\epsilon)|\epsilon)]\big) = 0,$$

where $\frac{d}{da}\big(a[1-G(a,p^*(a,\epsilon)|\epsilon)]\big) = [1-G(a,p^*(a,\epsilon)|\epsilon)] - a\big[\frac{\partial G(a,p^*(a,\epsilon)|\epsilon))}{\partial a} + \frac{\partial G(a,p(a,\epsilon)|\epsilon)}{\partial p}p_a^*(a,\epsilon)\big]$. This rearranges to condition (6), when setting $p = p^*(a,\epsilon)$. □

Proof of Proposition 2. Guess that there is a linear equilibrium of the following form: $a_i(p,\epsilon_i) = \alpha - \beta p + \gamma\epsilon_i$, and let all traders but trader i play this equilibrium. For convenience we drop the trader i supscript. The residual supply trader i faces is $RS(p,\boldsymbol{Z}) = \boldsymbol{Z} - (N-1)\alpha + (N-1)\beta p$, where $\boldsymbol{Z} = \boldsymbol{A} - \gamma\sum_{j\neq i}\boldsymbol{\epsilon_j}$. Following Wittwer (2021), we will maximize over a function $b(Z,\epsilon)$ that maps from Z into prices rather than from quantities into prices directly. This allows us to show that a function that fulfills the necessary conditions of the maximization problem also fulfills sufficient conditions, and is therefore indeed optimal. Maximization problem (4), becomes:

$$\max_{b(\cdot,\cdot)\in\mathcal{B}} \mathbb{E}[V(RS(b(\boldsymbol{Z},\boldsymbol{\epsilon}),\boldsymbol{Z}),\boldsymbol{\epsilon}) - b(\boldsymbol{Z},\boldsymbol{\epsilon})RS(b(\boldsymbol{Z},\boldsymbol{\epsilon}),\boldsymbol{Z})] \tag{43}$$
$$\text{subject to: } \mathbb{E}[\boldsymbol{\theta} - \kappa b(\boldsymbol{Z},\boldsymbol{\epsilon})RS(b(\boldsymbol{Z},\boldsymbol{\epsilon}),\boldsymbol{Z})] \geq 0.$$

Abbreviating $b(\cdot,\cdot)$ by b, this problem is equivalent to

$$\max_{b\in\mathcal{B}} I(b) \text{ subject to } L(b) \geq 0, \text{ with} \tag{44}$$
$$I(b) = \int_{\underline{Z}}^{\overline{Z}}\int_{\underline{\epsilon}}^{\overline{\epsilon}} F(b(\epsilon,Z),Z,\epsilon)\psi(\epsilon)\phi_Z(Z)d\epsilon dZ,$$
$$\text{where } F(b(Z,\epsilon),Z,\epsilon) = V(RS(b(Z,\epsilon),Z),\epsilon) - b(Z,\epsilon)RS(b(\epsilon,Z),Z)$$
$$L(b) = \int_{\underline{Z}}^{\overline{Z}}\int_{\underline{\epsilon}}^{\overline{\epsilon}} H(b(Z,\epsilon),Z,\epsilon)\psi(\epsilon)\phi_Z(Z)d\epsilon dZ,$$
$$\text{where } H(b(Z,\epsilon),Z,\epsilon) = \mathbb{E}[\boldsymbol{\theta}] - \kappa b(Z,\epsilon)RS(b(Z,\epsilon),Z).$$

Here, $\phi_Z(\cdot)$ is the density function of $\boldsymbol{Z}$ which has support $[\underline{Z},\overline{Z}]$, and $\psi(\cdot)$ is the density function of $\boldsymbol{\epsilon}$ on support $[\underline{\epsilon},\overline{\epsilon}]$.

Function b^* is optimal if $L(b^*) \geq 0, \lambda L(b^*) = 0, \lambda \geq 0$, $\frac{\partial(F+\lambda H)}{\partial b}$ evaluated at the optimum is 0 for all Z, ϵ:

$$(\mu + \epsilon)RS' - \rho\sigma^2 RSRS' - (1 + \lambda\kappa)(RS + bRS') = 0, \tag{45}$$

where we abbreviate $RS = RS(b^*, Z)$ and $RS' = \frac{\partial RS(b^*,Z)}{\partial b}$. Rearranging we obtain:

$$\mu + \epsilon - \rho\sigma^2 RS(b^*, Z) = (1 + \lambda\kappa)\Big[b^* + RS(b^*, Z)\Big(\frac{\partial RS(b^*, Z)}{\partial b}\Big)^{-1}\Big]. \tag{46}$$

In addition, natural boundary conditions,

$$\frac{\partial(F + \lambda H)}{\partial b_Z}(b(\underline{Z}, \epsilon), \underline{Z}, \epsilon) = \frac{\partial(F + \lambda H)}{\partial b_Z}(b(\overline{Z}, \epsilon), \overline{Z}, \epsilon) = 0, \tag{47}$$

where b_Z denotes the partial derivative of $b(Z, \epsilon)$ with respect to Z, must be satisfied. This is the case because $F + \lambda H$ is independent of $\frac{\partial b(Z,\epsilon)}{\partial Z}$ and $\frac{\partial b(Z,\epsilon)}{\partial \epsilon}$.

To show that a function b^* that fulfills the necessary conditions is indeed optimal, we rely on the fact that $F(b, Z)$ and $K(b, Z) = F(b, Z) + \lambda H(b, Z)$ are for any Z, ϵ, and $\lambda \geq 0$, strictly concave as functions of b, since

$$-\rho\sigma^2(RS'(b, Z))^2 - 2(1 + \lambda\kappa)RS'(b, Z) < 0 \text{ and } RS''(b, Z) = 0.$$

Therefore, $K(b(Z, \epsilon), Z) - K(b^*(Z, \epsilon), Z) < \frac{\partial K(b(Z,\epsilon),Z)}{\partial b}(b(Z, \epsilon) - b^*(Z, \epsilon)) \leq 0$ for any $b(Z, \epsilon)$ and any Z, ϵ. Multiplying both sides with $\phi_Z(Z)\psi(\epsilon)$ and integrating over Z, ϵ, we see that $\int\int K(b(Z, \epsilon), Z)\phi_Z(Z)\psi(\epsilon)dZd\epsilon < \int\int K(b^*(Z, \epsilon), Z)\phi_Z(Z)\psi(\epsilon)dZd\epsilon$, and similarly for $F(b(Z, \epsilon), Z)$. Thus, $b^*(\cdot, \cdot)$ is indeed optimal.

From here it is straightforward to solve for an equilibrium and show that it is unique within the class of symmetric linear equilibria. For this we rely on the property that function $b^*(\cdot, Z)$ implies a unique demand function $p^*(\cdot, \epsilon)$ for all ϵ. Then we match coefficients of the trader's best reply in (46) with the equilibrium guess and show that these coefficients are unique.

When everyone plays the equilibrium demand function (7), the market clears at

$$P^* = \frac{1}{1 + \lambda\kappa}\left(\mu + \frac{1}{N}\sum_i \epsilon_i - \left(\frac{N - 1}{N - 2}\right)\rho\sigma^2\frac{A}{N}\right), \tag{48}$$

and trader i wins $a_i^* = \frac{A}{N} - \left(\frac{N-2}{N-1}\right)\frac{1}{\rho\sigma^2}\left(N\epsilon_i - \sum_i \epsilon_i\right)$. The binding Lagrange multiplier is pinned down by the binding constraint: $\mathbb{E}[\boldsymbol{\theta_i} - \kappa \boldsymbol{P^* a_i^*}] = 0$. This completes the proof that equilibrium (7) is the unique linear symmetric equilibrium in the general case in which bidders have private information. □

Proof of Proposition 3. Consider bidder i, and assume that all other bidders play the equilibrium strategy. Fix some $\lambda \geq 0$, and note that for any given λ, choosing $\{p_k(\epsilon_i), a_k(\epsilon_i)\}_{k=1}^{K(\epsilon_i)}$ to maximize $EV(\epsilon_i) - (1+\lambda)EP(\epsilon_i)$ for all ϵ_i point-wise is equivalent to choosing $\{p_k(\cdot), a_k(\cdot)\}_{k=1}^{K(\cdot)}$ to maximize $\mathbb{E}[EV(\boldsymbol{\epsilon_i}) - (1+\lambda)EP(\boldsymbol{\epsilon_i})]$. Therefore, we can determine the optimal bids for the event that the bidder observes ϵ_i point-wise for all ϵ_i.

For this, we can follow the proof of Proposition 1 in Kastl (2012) one-by-one with one difference. The marginal expected payment (A.3) in Kastl (2012)'s proof must be multiplied by $(1+\lambda\kappa)$ to account for the shadow cost of the capital constraint. Similarly, the marginal cost of winning a bid at step k is $(1+\lambda\kappa)p_k(\epsilon_i)$ instead of $p_k(\epsilon_i)$. This adjustment also needs to be made in the two auxiliary lemmas that the proof of the proposition relies on. For example, with capital constraints, a bidder places a bid such that $(1+\lambda\kappa)p_k(\epsilon_i) \leq v_i(a_k(\epsilon_i), \epsilon_i)$ when there is a positive probability of a tie at any step.

In the end, we determine the correct λ. This is the λ for which $\lambda\mathbb{E}\Big[\boldsymbol{\theta_i} - \kappa\Big[\sum_{k=1}^{K(\boldsymbol{\epsilon_i})} \Pr(p_k(\boldsymbol{\epsilon_i})) > \boldsymbol{P^c}|\boldsymbol{\epsilon_i})p_k(\boldsymbol{\epsilon_i})(a_k(\boldsymbol{\epsilon_i}) - a_{k-1}(\boldsymbol{\epsilon_i})) + \sum_{k=1}^{K(\boldsymbol{\epsilon_i})} \Pr(p_k = \boldsymbol{P^c}|\boldsymbol{\epsilon_i})\mathbb{E}\big[p_k(\boldsymbol{\epsilon_i})(\boldsymbol{a_i^c} - a_{k-1}(\boldsymbol{\epsilon_i}))|p_k(\boldsymbol{\epsilon_i}) = \boldsymbol{P^c}, \boldsymbol{\epsilon_i}\big]\Big]\Big] = 0$, when all bidders submit step function that satisfy the necessary equilibrium condition for all ϵ_i. □

Proof of Proposition 4. We derive a symmetric equilibrium by simplifying the equilibrium of Proposition 5 under the assumption that $\mathbb{E}[\boldsymbol{\theta_i^B} + \boldsymbol{\theta_i^S}] = \kappa p_z \mathbb{E}[\boldsymbol{z_i^B} + \boldsymbol{z_i^S}] > 0$.

We start by guessing that there is an equilibrium in which the constraint is binding with equal strength for both buyers and sellers, i.e., $\lambda^B = \lambda^S = \lambda \geq 0$. Under this conjecture, we simplify the demand coefficients provided in Proposition 5 to obtain candidate equilibrium demand function (17). This candidate is indeed an equilibrium if the Lagrange multipliers that ensure that the capital constraints hold for both buyers and sellers are non-negative and common across trader groups, i.e., $\lambda^B = \lambda^S = \lambda \geq 0$.

To show that this is the case, and derive the functional form for λ, we determine the λ at which the constraint binds, simultaneously for both groups of traders. The solution is

given in equation (18). Since $\mathbb{E}[\boldsymbol{\theta_i^B} + \boldsymbol{\theta_i^S}] = \kappa p_z \mathbb{E}[\boldsymbol{z_i^B} + \boldsymbol{z_i^S}] > 0$ by assumption, the Lagrange multiplier is identical for buyers and sellers. □

Proof of Propositions 5. We guess that there is a group-symmetric BNE, in which a trader of group G submits the following net-demand curve:

$$a^G(p, z_i^G) = \alpha^G - \beta^G p + \gamma^G z_i^G \text{ with } \beta^G > 0, \text{ for } G \in \{B, S\}.$$

When everyone submits the equilibrium guess, the market clears when total demand equals total supply: $\sum_{i=1}^N a^B(p, z_i^B) = -\sum_{i=1}^N a^S(p, z_i^S)$. Thus, ex-post, buyers and sellers face different types of residual supply curves, and price impact:

$$RS^B(p, z_i^B) = -\sum_{i=1}^N a^S(p, z_i^S) - \sum_{j\neq i}^N a^B(p, z_j^B) \quad , RS^S(p, z_i^B) = -\sum_{j\neq i}^N a^S(p, z_j^S) - \sum_{i=1}^N a^B(p, z_i^B)$$
$$\frac{\partial RS^B(p, z_i^B)}{\partial p} = N\beta^S + (N-1)\beta^B = \frac{1}{\Lambda^B} \qquad , \frac{\partial RS^S(p, z_i^S)}{\partial p} = N\beta^B + (N-1)\beta^S = \frac{1}{\Lambda^S}.$$

Following the same steps as the proof of Proposition 2, we obtain the following necessary and sufficient equilibrium conditions for the buyers and sellers best response, respectively:

$$\mu - \rho\sigma^2(a^B + z_i^B) = (1 + \lambda^B\kappa)\Big[p + a^B\Big(N\beta^S + (N-1)\beta^B\Big)^{-1}\Big], \tag{49}$$
$$\mu - \rho\sigma^2(a^S + z_i^S) = (1 + \lambda^S\kappa)\Big[p + a^S\Big(N\beta^B + (N-1)\beta^S\Big)^{-1}\Big]. \tag{50}$$

In the group-symmetric equilibrium, all buyers and sellers must choose the equilibrium guess. Matching the β coefficients from the best replies to the equilibrium, we can express β^S as a function of β^B, shown in equation (21), and characterize β^B as the root of polynomial (20). For these coefficients to be valid in equilibrium they must be strictly positive. Similarly, we back out the α^G and γ^G coefficients as functions of β^B and β^S from the best responses (49) and (50).

The equilibrium Lagrange multiplier is pinned down by the binding constraint (16) for buyers and sellers, whenever it is binding, similar to the benchmark model. However, now, we can no longer solve for $\lambda^S, \lambda^B \geq 0$ in explicit form, since β^B and β^S are complicated functions of λ^S, λ^B. □

Proof of Proposition 6. The proof is analogous to the proof of Proposition 1, but the constraint is different. Using the same notation as before, the bidder's maximization problem is:

$$\max_{p\in\mathcal{B}} I(p) \text{ subject to } L(p(a,\epsilon),a) = \theta - \kappa p(a,\epsilon)a \geq 0 \tag{51}$$

for all a, ϵ, with $I(p)$ given by (38).

A function p^* is optimal if the following conditions are satisfied:

$$\frac{\partial F}{\partial p}(p^*(a,\epsilon), p_a^*(a,\epsilon), a, \epsilon) - \frac{d}{da}\left(\frac{\partial F}{\partial p_a}(p^*(a,\epsilon), p_a^*(a,\epsilon), a, \epsilon)\right) - \lambda\kappa a = 0, \tag{52}$$

and $L(p^*(a,\epsilon),a) \geq 0, \lambda \geq 0$ for all $a \in [0, \overline{a}^*(\epsilon)]$, and all ϵ. As before, natural boundary conditions,

$$\frac{\partial F}{\partial p_a}(p^*(0,\epsilon), p_a^*(0,\epsilon), 0, \epsilon) = \frac{\partial F}{\partial p_a}(p^*(\overline{a}^*(\epsilon),\epsilon), p_a^*(\overline{a}^*(\epsilon),\epsilon), \overline{a}^*(\epsilon), \epsilon) = 0, \tag{53}$$

are satisfied. Inserting the expressions for the partial derivatives of F like in the proof for Proposition 1, and simplifying gives:

$$(p^*(a,\epsilon) - v(a,\epsilon))\frac{\partial G(a,p|\epsilon)}{\partial p}\psi(\epsilon) = \lambda a + a\frac{\partial G(a,p|\epsilon)}{\partial a}\psi(\epsilon), \tag{54}$$

which rearranges to condition (25). □

Proof of Proposition 7. The proof is analogous to the proof of Proposition 1. There is only one difference, which comes from the fact that bidders pay the prices they bid for all units that they win instead of the market clearing price. This implies that $F(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon)$, and $H(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon)$ in maximization problem (37) are

$$F(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon) = [v(a,\epsilon) - p(a,\epsilon)][1 - G(a, p(a,\epsilon)|\epsilon)]\psi(\epsilon), \tag{55}$$

and

$$H(p(a,\epsilon), p_a(a,\epsilon), a, \epsilon) = [\mathbb{E}[\boldsymbol{\theta}] - \kappa p(a,\epsilon)[1 - G(a, p(a,\epsilon)|\epsilon)]]\psi(\epsilon), \tag{56}$$

respectively. With slight abuse of notation we are using the same notation as in the uniform-price auction. □

Proof of Proposition 8. When equilibrium quantities, $\boldsymbol{a_i^*}$, follow a Generalized Pareto distribution, we can solve for a function that fulfills condition (6') of Proposition 7. For this, we combine the insight that a trader bids as if their true willingness to pay was $\frac{v(a,\epsilon_i)}{1+\lambda\kappa}$ for any given $\lambda \geq 0$, with a known result from the literature on equilibrium existence (e.g., Proposition 7 of Ausubel et al. (2014), Theorem 2 of Wittwer (2018))). In equilibrium, $\lambda > 0$ is pinned down by the capital constraint if the constraint binds, and is zero otherwise. □

Proof of Corollary 1. To show that price impact increases when the constraint is relaxed, which decreases λ, we take the derivative of price impact (8) with respect to λ. It is negative. Similarly, the derivative of the clearing price (48) with respect to λ is negative. □

Proof of Corollary 2. To prove the corollary, we compute the following elasticity:

$$\eta_P = \frac{\partial P^*}{\partial(\lambda\kappa)}\frac{\lambda\kappa}{P^*} \text{ with } \Lambda \text{ defined in (8) and } P^* \text{ defined in (48)},$$
$$\eta_\Lambda = \frac{\partial \Lambda}{\partial(\lambda\kappa)}\frac{\lambda\kappa}{\Lambda} \text{ with } \Lambda \text{ defined in (8)}.$$

Simplifying terms, we obtain that $\eta_P = \eta_\Lambda = \eta = \frac{1}{1+\lambda\kappa} - 1$. □

Proof of Corollary 3. Let $\underline{\epsilon} = \bar{\epsilon} = 0$, so that $\epsilon_i = 0$ for all traders i, and omit the dependence on ϵ_i in all functions. In this case, we can simplify the condition for a symmetric equilibrium (in which each trader wins $\frac{A}{N}$) of Proposition 6. For this, we replace $\psi(\epsilon_i) = 1$ and insert:

$$\frac{\partial G(a,p|\theta_i)}{\partial a} = \phi(Na) \tag{57}$$
$$\frac{\partial G(a,p|\theta_i)}{\partial p} = (N-1)\phi(Na)\left(\frac{\partial a^*(p)}{\partial p}\right) = (N-1)\phi(Na)\left(\frac{\partial p^*(a)}{\partial a}\right)^{-1}. \tag{58}$$

We obtain differential equation (26). This differential equation has unique solution (28). □

Proof of Corollary 4. Statement (i) follows from condition (26) of Corollary 3. Price impact is non-negative since any valid equilibrium demand function must be decreasing, which implies $p^*(a) \leq v(a)$. To prove statement (ii), we fix some point a, and take the partial derivative of price impact (29) with respect to $\lambda(a)$ for a fixed point a. This derivative is positive if $\frac{\partial p^*(a)}{\partial \lambda(a)} < \frac{\phi(Na)}{(\phi(aN)+\kappa\lambda(a))^2}(v(a) - p^*(a))$. □

Appendix Table A1: Bid functions are approximately linear

	mean	median	sd
β_t	0.21	0.16	0.13
R_t^2	0.72	0.74	0.11
Adj. R_t^2	0.64	0.67	0.15
Within R_t^2	0.54	0.56	0.15

Appendix Table A1 shows the point estimate and R^2 from regressing bids on quantities and an auction-dealer fixed effect in each auction, $b_{tik} = \zeta_{ti} + \beta_t a_{tik} + \epsilon_{tik}$, using bidding functions with at least two steps. Bids are in bps of yields and quantities in percentage of supply.

Appendix Table A2: Estimates for all regression specifications

	Median ρ		Mean ρ		Median $\lambda\kappa$		Mean $\lambda\kappa$	
Benchmark	0.0053	(4.438e-05)	0.0048	(5.272e-03)	1.518e-11	(6.193e-03)	0.5292	(1.347e-02)
3 steps	0.0053	(3.379e-05)	0.0049	(4.534e-05)	1.518e-11	(6.835e-03)	0.5571	(1.416e-02)
4 steps	0.0052	(3.378e-05)	0.0048	(4.516e-05)	1.634e-11	(9.320e-03)	0.5098	(1.529e-02)
5 steps	0.0052	(3.340e-05)	0.0047	(3.340e-05)	5.540e-12	(1.539e-02)	0.5494	(2.097e-02)
6 steps	0.0049	(3.494e-05)	0.0049	(4.869e-05)	1.682e-12	(3.403e-11)	0.3389	(1.627e-02)
In mil C$	0.0001	(8.749e-07)	0.0001	(9.970e-05)	3.638e-12	(1.378e-03)	0.3138	(8.806e-03)
2.5%-winsorized	0.0050	(3.214e-05)	0.0048	(4.307e-05)	3.014e-10	(6.193e-03)	0.4115	(1.191e-02)
5.0%-winsorized	0.0050	(3.153e-05)	0.0048	(4.270e-05)	2.405e-10	(6.194e-03)	0.3530	(1.133e-02)
IQR	0.0043	(3.615e-05)	0.0044	(3.857e-05)	1.968e-09	(7.657e-03)	0.4658	(1.293e-02)
2 days	0.0050	(3.881e-05)	0.0045	(4.003e-05)	2.249e-14	(1.421e-04)	0.6188	(1.127e-02)
3 days	0.0047	(3.847e-05)	0.0048	(4.196e-05)	3.513e-13	(4.719e-04)	0.6285	(1.202e-02)
4 days	0.0048	(3.766e-05)	0.0048	(4.360e-05)	2.217e-11	(4.802e-03)	0.5559	(1.200e-02)
6 days	0.0054	(4.386e-05)	0.0049	(3.414e-05)	1.489e-11	(4.898e-0.3)	0.6159	(1.504e-02)
7 days	0.0054	(3.666e-05)	0.0047	(4.067e-05)	1.153e-10	(1.175e-02)	0.6597	(1.607e-02)

Appendix Table A2 shows the median and mean of the risk aversion estimates and the shadow costs, respectively for different regression specifications. The mean and median standard errors are in parentheses. "Benchmark" reports the estimates of our benchmark specification. In columns "N steps" for $N \in \{2, 3, 4, 5, 6\}$ we use only values that correspond to bidding step functions with weakly more than N steps. In column "In mil C$" we use quantities in absolute values, expressed in million C$. In columns "$x$%-winsorized" for $x \in \{2.5, 5\}$ we winsorize the distribution of return volatilities, σ_t^2, by x% to exclude outliers which lead to large shadow costs, and a higher average of these costs. In column IQR we use the inter-quantile range of yields at which dealers trade up to five days prior to the auction as our volatility measure, which also reduces across-auction heterogeneity in volatility. In columns "T days" for $T \in \{2, 3, 4, 6, 7\}$ we construct our volatility measure using trades on T days prior to the auction. We exclude $T = 5$ since this is used in our benchmark specification, and $T = 1$, because there aren't sufficiently many trades for each auction.

Appendix Figure A1: Holders of Canadian government bonds

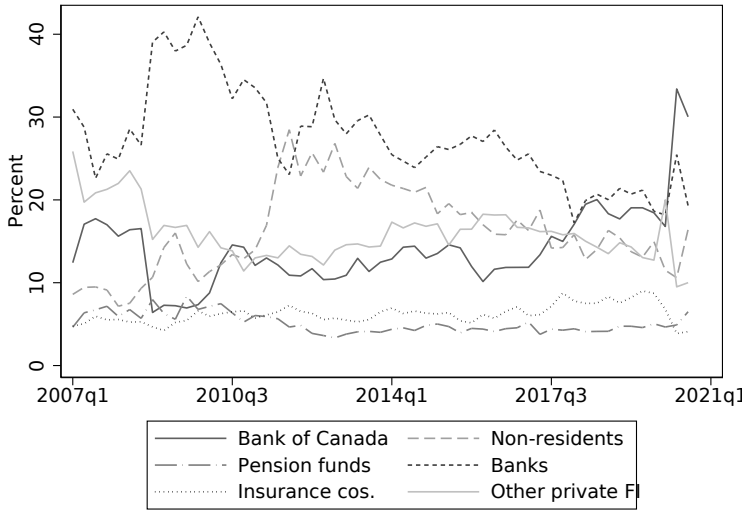

Appendix Figure A1 shows who holds Canadian government bonds and bills from 2007 until 2021 in percentage of par value outstanding: Bank of Canada, Non-residents, Canadian pension funds, Canadian banks, Canadian insurance companies, and other private firms. The bank category holdings are mostly driven by the eight banks we focus on, as they hold over 80% of the assets of all banks.

Appendix Figure A2: How price impact varies in the number of bidders and shadow costs

(a) Direct price impact

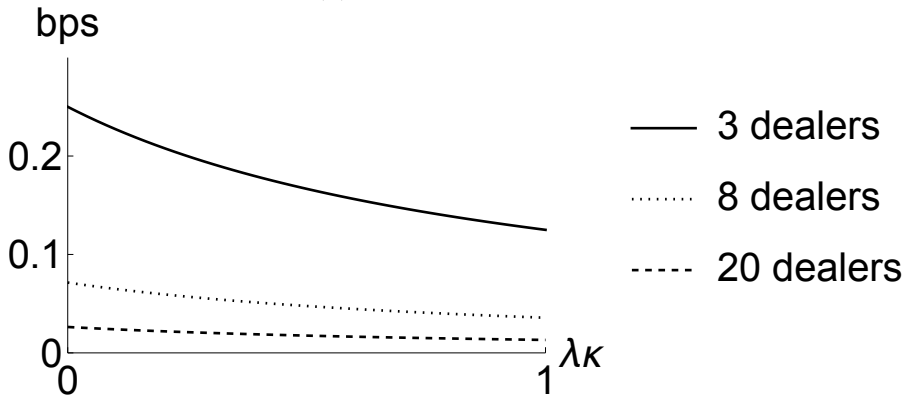

(b) Strategic effect in DPA

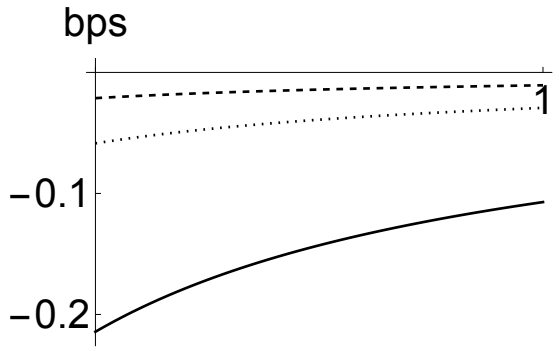

(c) Strategic effect in UPA

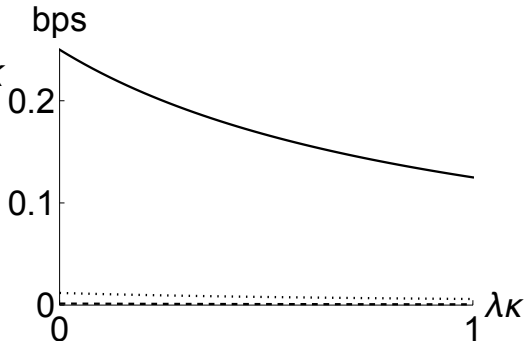

Appendix Figure A2a hows how the direct price impact, $\frac{1}{N-1}\frac{\rho\sigma^2}{1+\lambda\kappa}$, varies in the shadow cost for 3, 8, and 20 bidders, presented by the black, the dotted, and dashed line, respectively, using the average risk aversion estimate, and the average volatility. Figure A2b shows the price impact that arises due to demand reduction, according to equation (10'), in a typical discriminatory price auction (DPA). For illustration, we use the median ξ estimate, which is negative, so that the strategic effect is negative. Figure A2c shows the strategic price impact effect for uniform price auctions (UPA) using equation (10).

Appendix Figure A3: Equilibrium example in an asymmetric trade setting

(a) Effect on slopes: β^B, β^S

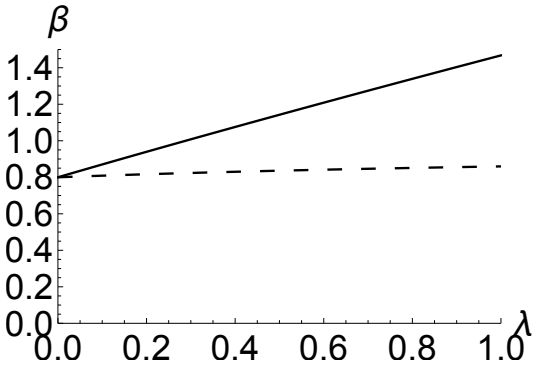

(b) Effect on price: P^*

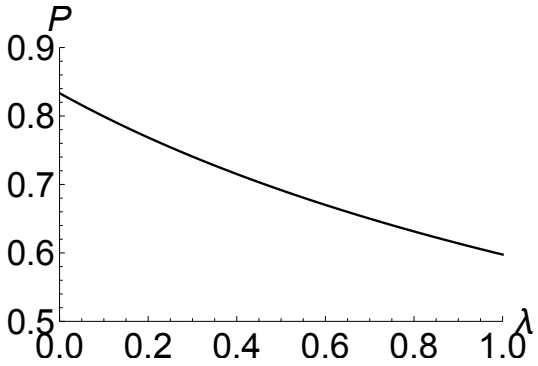

(c) Effect on price impact: Λ^B, Λ^S

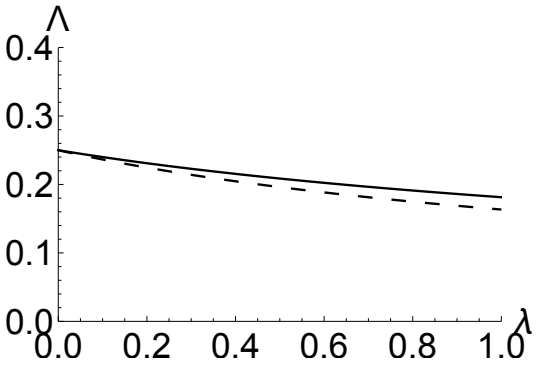

Appendix Figure A3 visualizes the effect of tightening the capital constraint for buyers, i.e., increasing their λ, on the slope coefficients in (a), on the market clearing price in (c), and the price impact of buyers and sellers in (c) for buyers (the solid line) and for sellers (the dashed line). To compute the price, we assume $\sum_i z_i^B = 2$ and $\sum_i = z^S = -1$.

Appendix Figure A4: Difference of slope estimates during and outside the exemption period

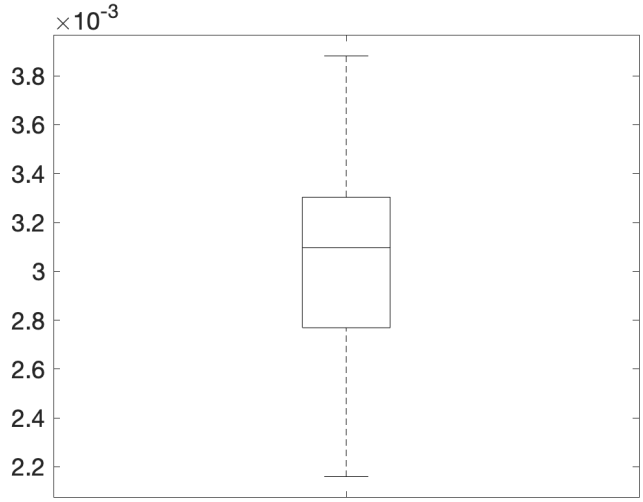

Appendix Figure A4 shows the distribution of the difference in the median slope in all auctions that took place during the exemption period versus the median slope in all auctions that took place during regular times, $\Delta\beta_d$, across bootstraps d.

Appendix Figure A5: Variation in quantities

(a) Supply and total demand

(b) Total demand in percentage of supply

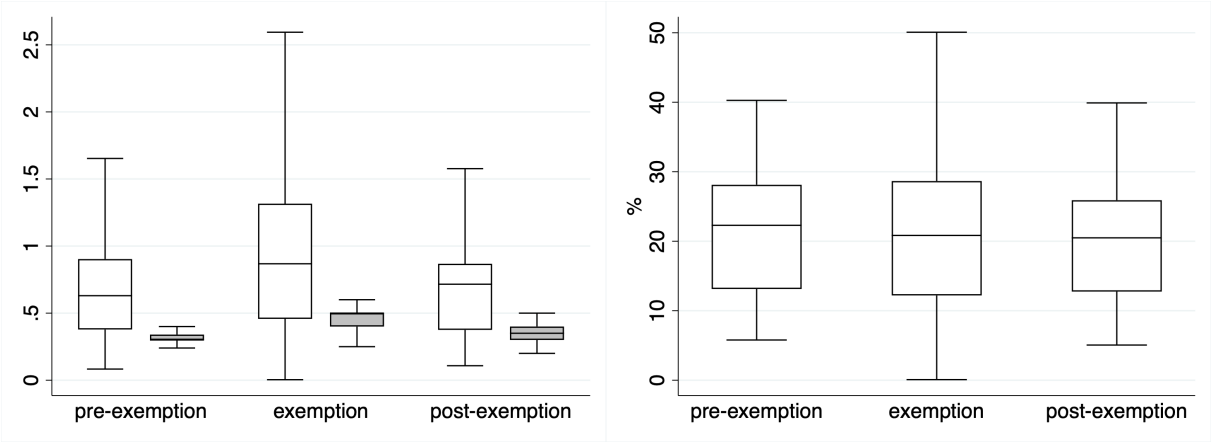

Appendix Figure A5a shows the distribution of the total amount a dealer demands in an auction before, during, and after the exemption period (in white) and the distribution of the supply (in gray). Demand is expressed in million C$, and supply is in 10 million C$ to make the two comparable. Appendix Figure A5b shows the distribution of the total amount demanded as percentage of supply across periods.

Appendix Figure A6: Excess capital holdings

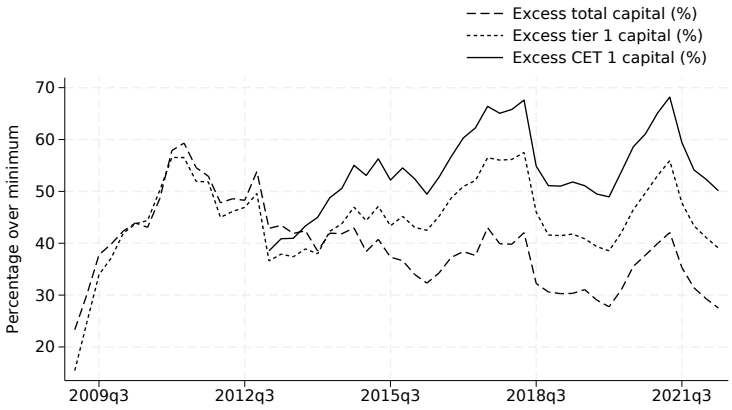

Appendix Figure A6 shows the average excess capital holdings for three different regulatory capital constraints: (i) Total risk-weighted capital, (ii) Tier 1 capital, and (iii) Common equity Tier 1 (CET1) capital. The main take-away is that excess capital holdings increased during the COVID crisis, moving banks away from regulatory thresholds (other than the LR constraint).