The Demand for Safe Assets[1]

Filippo Cavaleri Angelo Ranaldo Enzo Rossi

December 19, 2024

Abstract

This paper examines how heterogeneity in investment horizons determines the demand for safe assets, bidding strategies in auctions, and post-auction price dynamics. We model a uniform-price double auction with resale where long-term investors hold assets to maturity, while dealer banks distribute the asset in secondary markets. Pure private (common) values emerge when only long-term investors (dealers) participate. Using unique data on Swiss Treasury bond auctions revealing bidders' identities, our empirical findings support key predictions: (1) substantial heterogeneity in demand schedules, with steeper demand curves for dealer banks; (2) Dealer banks' demand becomes steeper with increased demand risk and bid dispersion; and (3) demand elasticity positively predicts post-auction returns.

Keywords: auction, asset demand, safe asset, private and common values, government bonds
JEL Classification: D44, G12, D82, G14.

[1]First draft: December 2024. Filippo Cavaleri is at the University of Chicago, Booth School of Business and Department of Economics; Angelo Ranaldo is at the University of Basel - Swiss Finance Institute; Enzo Rossi is at the Swiss National Bank. We thank Ben Brooks, Zhiguo He, Ali Hortaçsu, Adrian Martinez, Stefan Nagel, Eric Richert, Quentin Vandeweyer, and Daniel Wittwer for comments. We thank Marco Gortan for excellent research assistance. Results are preliminary, and all remaining errors are own.

1 Introduction

Sovereign debt is predominantly issued through auctions, including by countries whose debt is regarded as a safe asset.[2] Safe assets, when held to maturity, carry no fundamental risk and feature minimal information asymmetries regarding their terminal payoff.[3] These characteristics make safe government debt appealing to a diverse set of investors, each with unique investment horizons and expected holding periods (OECD, 2019). This heterogeneity in valuations is evident in bidding behavior at auctions, where bidding strategies often display significant dispersion and variation over time. Motivated by this, we address two fundamental questions, both theoretically and empirically. How does demand heterogeneity influence the pricing of safe assets in the presence of a resale market? In what ways do differences in expected holding periods and investment horizons shape bidding strategies and post-auction price dynamics?

Theoretically, we analyze heterogeneity in the demand for safe assets using a multi-unit uniform-price double auction model within the linear-quadratic framework (Vives, 2011). The key innovation lies in our analysis of safe government debt, which is free from the fundamental uncertainty or information asymmetries about liquidation value commonly assumed in the literature (Kyle, 1989; Vayanos, 1999). Our main methodological contribution is modeling demand heterogeneity for safe assets through differences in investment horizons. We consider two types of investors. First, *long-term investors* buy the asset at the auction and hold it to maturity. Second, short-term investors, referred to as *dealer banks*, purchase the asset with the intention of distributing (at least) a portion of it in the secondary market. In line with the preferred-habitat view (Vayanos & Vila, 2021), we interpret long-term investors as regional banks, insurance companies, and pension funds, which are typically less active in secondary markets. Conversely, we identify short-term investors as dealer banks, implicitly acting as primary dealers and market makers, whose option to resell exposes them to aggregate demand risk or inventory risk (Boyarchenko, Lucca, & Veldkamp, 2021; Fleming, Nguyen, & Rosenberg, 2024). The buy-and-hold behavior of long-term investors shields them from secondary market price fluctuations, allowing them to operate with pure private values. In contrast, dealer banks' intent to resell implies they also operate with common values driven by future market conditions.

Our model has three periods. First, dealer banks and long-term investors bid in an auction for safe government bonds. Second, dealer banks trade in the secondary market with the investors that did not participate to the auction, while long-term investors abstain from trading. Lastly, the bond pays off a deterministic amount, and the game ends. All agents incur quadratic holding costs each period in which they hold the asset. They draw marginal cost intercepts from a common distribution with an unknown mean, referred to as the average cost, which each agent observes privately.Secondary market prices reflect the average cost, rendering private cost intercepts noisy signals of post-auction capital gains. Uncertainty in post-auction gains drives risk-averse dealer banks to demand a risk premium, a key distinction from Du and Zhu (2017), where traders are risk-neutral. In our setup, the distribution of post-auction returns conditional

[2] More than 80% of OECD countries primarily issue government debt through auctions (OECD, 2016).

[3] See Dang, Gorton, Holmström, and Ordoñez (2017), Gorton (2017), and Gorton and Ordoñez (2022).

on prices and private information is *endogenously* linked to demand risk, cost dispersion, and the ratio of dealer banks to total participants. Distinct from the models of Rostek and Yoon (2021) and Allen and Wittwer (2023), where the volatility of asset returns is given, our results show that the subjective volatility of post-auction returns, and thus subjective risk premia, vary across investor types and depend on how many dealer banks relative to long-term investors bid at the auction.

In our model, dealer banks participate in the secondary market while long-term investors do not, so that we associate heterogeneity in *investment horizons* to heterogeneity in expected holding periods. We characterize a Bayes-Nash equilibrium in demand schedules to demonstrate how such heterogeneity influences bidding behavior and post-auction return dynamics. A key novelty of our approach is the analysis of an *asymmetric* equilibrium, where demand schedules are symmetric within types but differ across types. We demonstrate that an equilibrium with downward-sloping demand curves always exists and is never symmetric. Dealer banks, because of anticipated future resales, benefit from learning about average costs from auction prices, while long-term investors, who are not exposed to fluctuations in the secondary market price, lack such incentives. Our model nests pure private values and pure common value as special cases. Pure private values arise when only long-term investors participate. By contrast, a pure common value arises when only dealer banks participate. Under private values, bidding strategies are unaffected by demand uncertainty, cost dispersion, or risk premia. In intermediate cases, two predictions emerge: (i) demand schedules steepen as either demand uncertainty or cost dispersion increases, with a more pronounced effect for dealer banks; (ii) equilibrium prices include compensation for aggregate demand risk, tied explicitly to demand uncertainty and cost dispersion. Lower demand elasticity during auctions therefore positively predicts post-auction returns. The horizon of return predictability depends on which investors are more inelastic. When only dealer banks' demand is less elastic, post-auction return predictability is short-lived until dealer banks offload their inventories in the secondary market. However, when long-term investors also exhibit inelastic demand, post-auction returns remain predictable for longer periods, up to one month. This prediction aligns with preferred-habitat models, where mean-reverting supply shocks impact short-term return predictability, and the effect is more pronounced when dealer banks' demand elasticity is lower (Greenwood & Vayanos, 2014). In contrast, shifts in long-term investors' elasticity influence how quickly and at what price dealer banks can offload their holdings in the secondary market, thus leading to return predictability at longer horizons.

Our second major contribution to the literature is empirical. We validate our model using a unique and novel dataset of hand-collected Swiss government bond auction data. The Swiss setting is particularly well-suited to our research questions for at least two reasons. First, Swiss Treasuries are widely recognized as safe assets.[4] Second, Switzerland does not formally have a primary dealer system, meaning that auctions are open to a wide range of participants. This set includes large banks *de facto* acting as primary dealers as well as long-term oriented in-

[4]Swiss sovereign debt has never faced downgrades or negative outlooks. S&P rated U.S. debt AA+ in August 2011 and gave Germany a negative outlook in December 2011. Japan and UK are currently rated A+ and AA. Additionally, the combination of a low debt-to-GDP ratio and prudent debt management policies minimizes concerns about rollover risk and auction failures.

vestors such as regional banks, pension funds, and insurance companies. A key innovation of our dataset is that we directly observe the *identities* of individual bidders. Thus, we can distinguish dealer banks performing market-making activities from investors more long-term oriented such as pension funds, insurance companies, and smaller regional banks. Accordingly, we bring the model to the data by interpreting large, systemic banks as dealer banks and classifying all other bidders as long-term investors.[5] Our ability to measure demand elasticities at the bidder level is a significant improvement over prior studies, which typically rely on aggregate measures (Albuquerque, Cardoso-Costa, & Faias, 2024) or use only allotment and price data (Boyarchenko et al., 2021). Furthermore, we can compare auction allotments to the secondary market allocation and validate whether the ownership structure varies over time through subsequent resales. Our dataset spans over forty years of Treasury auctions, covering bonds with maturities from two to fifty years. This extensive time series provides a rare opportunity to study bidding behavior across varying economic conditions and regulatory changes, including the Basel III reform.

Five key findings emerge from our paper. First, we document substantial cross-sectional heterogeneity in the level and the slope of demand schedules. Dealer banks are more cautious and submit steeper demand curves relative to long-term investors. Across all bidders, demand elasticities decline with bond maturity, consistent with a duration exposure channel (Allen, Kastl, & Wittwer, 2024; Greenwood & Vayanos, 2014). Second, as in our theory, dealer banks are more sensitive to demand risk than long-term investors. Dealer banks' demand schedules become steeper than those of long-term investors when bond return volatility rises before the auction. Through the lens of our model, this provides evidence of a common value in Treasury bond valuations driven by future resale in the secondary market. Third, our empirical analysis supports the hypothesis that dealer banks' demand elasticity decreases relative to other long-term investors when cross-sectional bid dispersion is higher. Fourth, we use a difference-in-differences design to analyze bidding behavior before and after the implementation of Basel III capital requirements, which primarily affect dealer banks through their market-making activities (BIS, 2017). Comparing dealer banks (*treatment*) to long-term investors (*control*), we find that Basel III regulatory costs lead to steeper demand schedules for dealer banks relative to the control group. Surprisingly, we also observe that dealer banks bid at a significantly lower discount to the secondary market after the reform. Fifth, we explore the predictive power of changes in demand elasticity among dealer banks versus long-term investors for post-auction returns. Consistent with our theoretical predictions, we find that lower dealer banks' elasticity predicts higher bond returns up to two days after the auction. In contrast, when long-term investors also become less elastic, returns are predictable for up to one month post-auction.

1.1 Related Literature

First, we contribute to the literature on safe assets by studying the determinants and implications of demand heterogeneity for safe government debt across the primary and secondary markets.[6]

[5] We apply the official definition of systemically important banks in Switzerland, which is determined collaboratively by the Swiss National Bank (SNB) and the Swiss Financial Market Supervisory Authority (FINMA).

[6] The shortage of safe assets and the macroeconomic and financial stability implications are studied by Caballero and Krishnamurthy (2009), Caballero, Farhi, and Gourinchas (2017), and Caballero and Farhi (2017).

Much of the literature focuses on the supply of safe assets, their scarcity (see, e.g., Caballero & Farhi, 2017; Caballero, Farhi, & Gourinchas, 2016), substitutability (Krishnamurthy & Vissing-Jorgensen, 2012a), and overall effects of their supply (Benigno & Nisticò, 2017; Infante, 2020). We document significant heterogeneity in safe asset demand. Our main contribution is to theorize and provide consistent evidence that heterogeneity in investment horizons introduces a demand risk premium. As a result, our paper links the concept of safe assets to heterogeneity in demand and the investor composition in the primary market. Another strand of the literature predicts cross-sectional dispersion in (quasi-)safe assets and convenience yields.[7] We shed light on the mechanisms governing the demand for safe assets. Barro, Fernández-Villaverde, Levintal, and Mollerus (2022) theorize heterogeneous risk-averse agents with rare disasters. Our paper makes a novel contribution by theoretically modeling and empirically analyzing private and common components in safe asset valuations, which are revealed in the primary market through strategic interactions among bidders.

Second, we contribute to the extensive literature on bidding behavior in Treasury auctions. The earlier literature on Treasury auctions explores the theoretical predictions of the share auction model of Wilson (1977) and Back and Zender (2001).[8] We depart from this literature by taking the auction mechanism as given and studying the implications of demand heterogeneity on bidding behavior and post-auction returns. The debate on the best way to sell government debt is very old (Friedman, 1991). Counterfactual exercises to assess which auction format would maximize government revenue have been the subject of many studies, often concluding that switching between discriminatory and uniform price auctions entails little efficiency gains (Hortaçsu & McAdams, 2010). Although our model is silent on optimal auction design, we emphasize that governments should also take into account the composition of the investor base, equipping them with an additional tool above and beyond the auction format (discriminatory versus uniform) or information that is disclosed prior and post auction (Dworczak, 2020). Furthermore, our unique sample covers more than four decades of auctions for bonds with maturities ranging from 2 to 50 years, providing plenty of variation in economic conditions and asset characteristics.[9] Our assumption about investment horizons implies that dealer banks operate with a common value, whereas long-term investors operate with pure private values. Structural models of multi-unit auctions generally assume private values for computational tractability (Hortaçsu & McAdams, 2010; Hortaçsu et al., 2018; Richert, 2024). Hortaçsu and Kastl (2012) test for common values in Canadian auctions for 3-month and 12-month Treasury bills by looking at whether primary dealers update their bids after observing customer orders, but find little evidence in support of common value components. We complement these results by inspecting Treasury auctions of

[7] See the discussions on the variations in the perceived safety and liquidity premiums associated with different assets in Krishnamurthy, 2002; Krishnamurthy & Vissing-Jorgensen, 2012b; Stein, 2012; Sunderam, 2015; Caballero, Farhi, & Gourinchas, 2016; Nagel, 2016; Moreira & Savov, 2017; He, Krishnamurthy, & Milbradt, 2019.

[8] See Hamao and Jegadeesh (1998) for Japan, Nyborg, Rydqvist, and Sundaresan (2002) for Sweden, Keloharju, Nyborg, and Rydqvist (2005) for Finland, Armantier and SbaÏ (2006) and Février, Préget, and Visser (2002) for France, Goldreich (2007) and Hortaçsu, Kastl, and Zhang (2018) for the US, Hortaçsu and Kastl (2012) and Allen and Wittwer (2023) for Canada, Hortaçsu and McAdams (2010) for Turkey, Beetsma, Giuliodori, Hanson, and de Jong (2020) for euro countries, Umlauf (1993) for Mexico, and Kastl (2011) for the Czech Republic.

[9] In comparison, the closest data in terms of comprehensiveness is Allen, Hortacsu, Richert, and Wittwer (2024) and Allen and Wittwer (2023) who observe Canadian auctions from 1999 to 2022.

longer-term bonds in a setting where the primary market is open to the general public. We show that when bidders have different investment horizons there is significant heterogeneity in how bidding strategies respond to demand uncertainty and secondary market liquidity, suggesting the presence of a common value for securities with longer maturity.

A more recent asset pricing literature uses Treasury auctions to explore the price impact of anticipated and repeated supply shocks. For the US, Lou, Yan, and Zhang (2013) show that yields drift up in the days before the auction and revert in the days after the auction. Beetsma et al. (2020) and Albuquerque et al. (2024) also document similar auction cycles in Italy and Portugal, respectively. Albuquerque et al. (2024) show that return predictability is driven by auctions where the aggregate elasticity of demand is low. We contribute to this literature by measuring demand elasticities at the bidder level and by showing that the price impact of Treasury auctions is more persistent when both long-term investors and dealer banks submit steeper demand curves. In particular, our findings support the view that demand curves for safe government bonds are downward sloping (Gabaix & Koijen, 2023; Vayanos & Vila, 2021), and we illustrate a specific channel, which is heterogeneity in investment horizons, that is reflected in bidding strategies at the auction. From the perspective of a government, our framework hints at a trade-off between paying a demand risk premium (only dealer banks) or having a less liquid secondary market (only long-term investors). The reduction in demand risk premia when long-term investors participate in the auction is akin to the interpretation of insurance companies and pension funds as asset insulators (Chodorow-Reich, Ghent, & Haddad, 2020). Furthermore, our assumption on buy-and-hold behavior, and the mechanism through which it can potentially impact the final allocation, finds empirical support in Musto, Nini, and Schwarz (2018), where assets held by buy-and-hold agents eventually become more illiquid.

Third, our theoretical model extends the literature on uniform-price double auctions by studying asymmetric equilibria with heterogeneous agents and deterministic payoffs. To characterize equilibria with downward sloping demand schedules, we build on the linear-quadratic setting of Vives (2011). As in Vives (2011), and in contrast to Kyle (1989) and Klemperer and Meyer (1989), equilibrium existence does not rely on aggregate demand shocks or noise trading. Different from Vives (2011) and Rostek and Weretka (2012), however, our model simultaneously admits pure private and common value components so that the average correlation in values across agents participants is not constant. Common values endogenously emerge through future resale in the secondary market, so that the investment horizon determines incentives to learn from prices. Our empirical analysis reveals that bidding strategies are significantly heterogeneous across investor types and respond differently to changes in risk. Existence of an equilibrium with asymmetric strategies combined with an asset free of fundamental risk is distinct from Vayanos (1999), Rostek and Weretka (2012), Du and Zhu (2017), and Rostek and Yoon (2021). Rostek and Weretka (2012) study a similar environment in the linear-quadratic setting, allowing a rich pattern of correlation in the signals in a setting with quadratic payoffs and uncertainty about marginal costs. Du and Zhu (2017) present a model of sequential double auctions with quadratic flow costs where common value shocks arise through signals about a stochastic terminal payoff. In our setting, agents are risk averse, and bidders use private signals to update their estimates of future resale prices. In the framework of Rostek and Yoon (2021), asset returns are exogenously specified and the conditional distribution of asset returns is unrelated to signal dispersion and prior uncertainty. Duffie, Malamud, and Manso (2009) and Chen and Duffie (2021) emphasize welfare implications but take expected returns as given and explore symmetric equilibria. In Kyle, Obizhaeva, and Wang (2017), agents are overconfident and each of them perceive their signal to be more informative than the other agents, but the equilibrium still remains symmetric.

1.2 Organization

The rest of the paper is structured as follows. Section 2 presents institutional setting and describes the data. Section 3 explains the model. Section 4 presents the empirical results. Section 5 concludes and discusses policy implications.

2 Swiss Treasury Auctions

This section describes the institutional setting of Swiss Treasury bond auctions and presents descriptive statistics of our novel hand-collected bids data. We emphasize the institutional background relevant to our analysis, and refer to Ranaldo and Rossi (2016) for further details.

2.1 Institutional Background

Switzerland has been one of the first OECD countries to issue government debt exclusively through auctions for all medium- and long-term maturities. The Swiss National Bank (SNB) has conducted sealed-bid uniform price auctions for Swiss government bonds on behalf of the Treasury since 1980. Auction participants submit competitive demand schedules, which consist of multiple price-quantity pairs that specify the amounts they are willing to buy (*bid quantity*) at each price (*bid price*). Bidders can also submit non-competitive quantity bids, which are filled with certainty.[10] Bidders can simultaneously submit competitive and non-competitive bids, and can always abstain from bidding. The SNB does not impose any restrictions on participation, bid steps, and maximum individual awards on competitive bids. When the auction closes, the Treasury compares aggregate demand and aggregate supply, net of non-competitive bids, to determine the market clearing price, which is the lowest accepted bid. Bids below the market clearing price are rejected, while bids above it are fully allocated at that price. Bids at the market clearing price may be prorated.

A key feature of Swiss bond emissions is that the auctions have always been open to the general public rather than being restricted to a limited group of primary dealers.[11] This feature makes the Swiss setting well-suited to our analysis of heterogeneity in safe asset demand for two reasons. First, while large banks *de facto* act as primary dealers, open participation attracts a broad spectrum of investors with different investment horizons and expected holding periods. The SNB groups all bidders into six categories, namely cantonal banks, big banks, foreign investors, other

[10] The rules of noncompetitive bidding underwent several changes, see Annex I in Ranaldo and Rossi (2016).

[11] The literature on Treasury auctions often studies settings where only primary dealers participate to the auction, e.g. Allen, Kastl, and Wittwer (2024) for Canada, Albuquerque et al. (2024) for Portugal, Kang and Puller (2008) for South Korea, Nyborg et al. (2002) for Sweden and Keloharju et al. (2005) for Finland.

banks (including private banks, trade banks, exchange banks, small credit banks), regional banks (including savings banks and Raiffeisen banks), and a residual category that includes insurance companies, pension funds, and individuals. Since we directly observe the identities of the bidders, we can separate large and systemic banks from all the other participants and measure bond demand at the bidder level. In both the theory and the empirical section, following the definition applied by the SNB and FINMA[12] we will refer to large and systemic banks as *dealer banks*, and to all other investors such as regional banks, insurance companies and pension funds, as *long-term investors*. Second, the absence of an official primary dealers system reduces artificial demand for Treasury debt driven by contractual obligations of a primary dealer system.[13] Without contractual constraints, all bidders are treated equally and retain full flexibility in their bidding strategies, creating a more transparent and competitive auction environment.
A measurement challenge arises because direct access for most non-bank financial institutions has been curtailed after the switch to electronic bidding in 2001. While major non-bank investors can still participate directly, smaller institutions may find it more convenient to route their bids through a direct bidder, typically a dealer bank. In such cases, we cannot distinguish dealer banks' bids made for their accounts from those made on customers' behalf. We impute such indirect bids to the corresponding direct bidder. By doing so, we obtain an upper bound on dealer banks' exposure to demand risk, as the actual portion of bids intended for secondary market distribution is likely smaller. Because of this, the observed behavior of dealer banks may more closely resemble that of long-term investors. Therefore, our approach understates the heterogeneity in the impact of market conditions and risk on bidding strategies, strengthening our conclusions. Furthermore, talks with practitioners indicate that indirect bidders delegate a non-negligible portion of strategy formulation to direct bidders, rather than specifying price-quantity pairs themselves.

The auction process begins with an announcement by the Treasury. The time between announcement and auction has gradually declined from several days to just one day since 1998. The bidding window opens at 9:30 AM and closes at 11:00 AM. Although the settlement date occurs several days after the auction, the securities begin trading immediately after the auction closes. The Treasury established an advance notice period to help participants prepare for the auctions. The auction announcement includes the coupon, the maturity, and, starting in August 1993, any amount it may want to reserve for subsequent sales in the secondary market, referred to as cancel the own tranche. The disclosure of the emission size has become progressively less precise over time. Initially, the Treasury announced an approximate borrowing target until October 1991, after which it began providing a maximum borrowing amount until November 1999. Since January 2000, no information on the emission size has been disclosed prior to an auction. After the auction, the Treasury releases summary statistics, including the total volume

[12] In Switzerland, the designation of systemically important banks is a collaborative process between the SNB and the Swiss Financial Market Supervisory Authority (FINMA). For more information, see https://www.finma.ch/en/enforcement/recovery-and-resolution/too-big-to-fail-and-financial-stability/systemically-important-banks/.

[13] US auctions are open to the general public, but primary dealers face constraints on participation and maximal awards, see Hortaçsu et al. (2018) and Boyarchenko et al. (2021). Payne and Szöke (2024) show how financial regulation creating captive demand for Treasuries can distort yields on government debt.

of received and accepted bids, the market clearing price and yield, the sum of non-competitive bids, payment date, (possible) fungibility with a previous issue, and the own tranche.

2.2 Data Description

The top panel of Table 1 presents descriptive statistics on the uniform-price auctions for Swiss government bonds. Our data cover the period from 1980 to 2023 and contain bidder-level demand schedules for 530 auctions. The total issue size is on average CHF 357.5 million (roughly USD 406 million at the current exchange rate). Though not seemingly enormous, this is significant given Switzerland's low public debt, with a debt-to-GDP ratio ranging from 8 to 24% during our sample period. The average bid volume is CHF 609.9 million, which far exceeds the average issue size. The average cover ratio (bid volume to issue size) is 1.77, indicating excess demand for Swiss Treasury bonds. However, since the Treasury does not announce the issue size in advance and can adjust it based on observed demand, the minimum cover ratio is effectively one. There is significant variation over time in auction participation, ranging from a minimum of five to a maximum of 73 participants. The average number of participants is 16.5, which is comparable to other economies (Hortaçsu & Kastl, 2012; Kastl, 2011).

	N	Mean	SD	Min	Median	Max
Auction variables						
Maturity	530	15.01	9.08	2.00	11.96	50.00
Issue size	530	356'879	237'243	56'700	284'375	1'553'470
Bid volume	530	607'540	461'792	99'100	475'487	4'676'315
Cover ratio	530	1.77	0.69	1.00	1.59	8.63
Bid steps	530	68.03	51.77	12.00	50.00	306.00
Participants	530	16.41	11.25	5.00	12.00	73.00
Auxiliary variables						
Volatility	330	0.42	0.29	0.05	0.34	2.00
Yield spread	359	0.02	0.03	-0.05	0.02	0.19
Inflation	344	1.09	1.47	-1.44	0.73	6.57
KOF Barometer	317	-0.01	1.05	-5.69	0.05	3.55
SARON	250	0.33	1.09	-0.75	0.02	3.39
Slope	327	0.74	0.65	-1.36	0.71	2.09

Table 1: Sample summary statistics. Maturity is the difference in years between settlement date and maturity date. Issue size (thousands CHF) is the total quantity issued in each auction, which is equal to the sum of the allocated quantities less the own tranche. Bid volume is the total volume of bids. Cover ratio is the ratio between bid volume and issue size. Volatility is the standard deviation of bond returns in the month prior to the auction. Yield spread is the difference between the auction market clearing yield and the secondary market yield. Slope is the difference between the 10-year and the 2-year yield. The auction sample is from 1980 to present. The secondary market sample is from 2000 to present.

The average maturity is 15 years, ranging from a minimum of 2 years up to a maximum of 50 years. The left panel of Figure 1 plots the distribution of bond maturities throughout the sample. Almost half of the 530 auctions issued medium-term bonds with a maturity between 10 and 20 years. However, our sample also includes a considerable number of emissions of bonds with maturities exceeding 30 years. Auctions with fractional maturities are typically security reopenings.

The scatter plot in the right panel of Figure 1 plots issue size and maturity for new emissions and security reopenings over time. The market size is proportional to the issue size. Most of the auctions after 2000 are security reopenings of already existing CUSIPs, with identical coupon rate and maturity date. In this period, the Treasury relied on these reopenings to manage liquidity in the secondary market. Over time, new bond emissions have become less frequent, larger in size, and have longer maturities. Until 1990, the Swiss government only issued two bonds with maturity longer than 20 years. In contrast, after 1995, the maturity of newly issued securities regularly exceeds 25 years.

(a) Histogram of bond maturities.

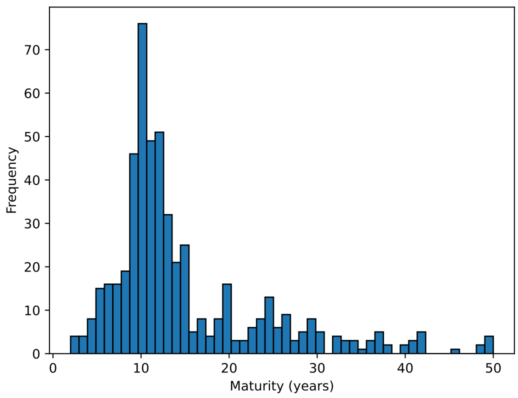

(b) Bond emissions over time.

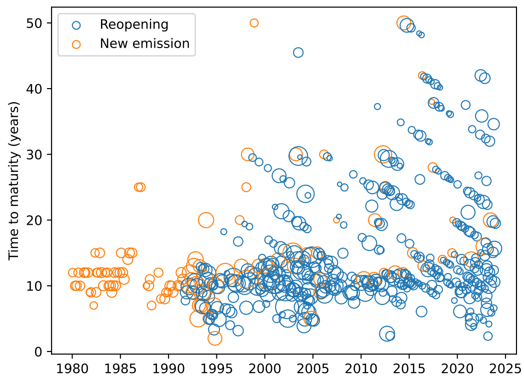

Figure 1: The left panel presents a histogram of bond maturities. The right panel plots bond issuance over time, separating issuance of new securities (orange circles) from reopenings (blue circles). Marker size is proportional to the issue size. The sample covers 530 auctions from 1980 to present.

2.3 Auxiliary Variables

We obtain secondary market prices for all outstanding Swiss government bonds from Bloomberg, which provides comprehensive coverage. The secondary market sample starts in 2000, since prior to that prices are sparse and contain significant gaps. For bond reopenings, we measure return volatility as the standard deviation of daily returns in the month prior to each auction closing day, provided there are at least 15 observations. The secondary market yield spread, henceforth *yield spread*, is the difference between the auction market clearing yield and the prevailing secondary market yield at the auction close. We measure secondary market liquidity as the relative bid-ask spread (RBAS), given by the ratio of the bid-ask spread to the midprice.

Most of the other auxiliary variables are from the SNB data portal. We retrieve monthly CPI inflation data from December 1982 to the present. Daily Treasury yields and interest rate data include overnight (SARON), one-week (SAR1W), two-week (SAR2W), and one-month (SAR1M) rates from June 1999 onward, as well as 10-year and 2-year yields from January 1988 onward. We compute the slope of the term structure as the difference between the 10-year and the 2-year yield. We finally use the KOF Economic Barometer as a proxy of the business cycle.

The bottom panel of Table 1 presents summary statistics of the auxiliary variables that we use in the main analysis in Section 4. A positive yield spread indicates that auctions are underpriced. On average, buying at the auction is cheaper than in the secondary market. The yield spread is 2 basis points, up to a maximum of 19 basis points. These numbers are comparable in magnitude to earlier findings in Cammack (1991) for US three-month T-bill auctions. Throughout our sample, there is significant variation in economic and financial conditions. The average overnight rate (SARON) is 0.33%. While the average overnight rate is slightly positive, our sample includes periods in which interest rates fell in negative territory down to negative 75 basis points. The average slope of the term structure is 74 basis points. The monthly inflation rate ranges from -1.44% to 6.57%.

3 Theoretical Framework

We propose a theoretical model of uniform price auctions in the linear-quadratic setting to study the implications of heterogeneity in investment horizons on bidding strategies. We then use the model to guide the empirical analysis.

3.1 Environment

Assets and Timing There are three periods, $t = 0, 1, 2$. The financial market consists of a two period bond that pays off a unit of the consumption good in period 2 with certainty. There are no other risky assets. The exogenous risk-free rate is set to zero and it is the numeraire. Let p_t denote the price of the two period bond at time t, and Q_t its outstanding supply.

Agents and Preferences The economy is populated by a continuum of agents. Each agent has CARA utility over terminal wealth W_{i2} given by

$$-\exp(-\gamma W_{i2})$$

where γ denotes the coefficient of risk aversion. At time $t = 0$, the government issues Q_a units of the bond through a uniform price auction. At $t = 1$, agents trade the two period bond in the secondary market. The bond matures at $t = 2$ and the game ends.

We assume that only a finite (exogenous) number $N = n + m$ of bidders participates in the auction. Of those N, there are $j = 1, \ldots, n$ dealer banks and $k = 1, \ldots, m$ long-term investors. The difference between the two types is their expected holding period, which we refer to as *investment horizon*. Dealer banks participate in the secondary market, whereas long-term investors buy and hold until maturity. The heterogeneity in investment horizons is sufficient to break symmetry of the linear equilibrium, and the slope of the demand schedule will be different across types. The framework nests the special cases of pure common values ($m = 0$) and pure private values ($n = 0$).

The $t = 1$ budget constraint for agents that participate in the secondary market, that is the n dealer banks plus the general public who does not bid in the auction, is

$$W_{i2} = W_{i1} + (1 - p_1)q_{i1} - \lambda_i q_{i1} - \frac{\kappa}{2}q_{i1}^2$$

where q_{i1} denotes bond quantities and p_1 is the equilibrium price. To simplify the analysis, we assume that the secondary market is competitive, so that no agent has market power. This assumption implies that we do not have to keep track of beliefs about the auction allocation and information leakages in the primary market to determine the equilibrium in the secondary market.

We introduce uncertainty in the demand for safe assets by assuming that agents pay a quadratic cost $\Lambda(q_{it}) = \lambda_i q_{it} + \frac{1}{2}\kappa q_{it}^2$ for each period in which they hold the bond. We specify the cost function such that the intercept of the marginal cost λ_i varies across investors, but the slope of the marginal costs is constant and publicly known (Vives, 2011). A positive value of λ_i can be interpreted as a regulatory requirement or balance sheet constraints. A negative value λ_i can be interpreted as a non-pecuniary benefit from holding the asset in the spirit of Krishnamurthy and Vissing-Jorgensen (2012a). Since the bond payoff is deterministic, the quadratic term ($\kappa \neq 0$) ensures that bond demand is bounded.

Agents that only participate in the secondary market (the competitive fringe) are endowed with exogenous wealth W_{i1} that they obtain from sources outside of the model (e.g. labor income). The n dealer banks who participate in the secondary market are endowed with wealth W_{i0} at $t = 0$ and buy bonds in the auction such that

$$W_{j1} = W_{j0} + (p_1 - p_0)q_{j0} - \lambda_j q_{j0} - \frac{\kappa}{2}q_{j0}^2$$

where q_{j0} is the quantity of bonds purchased at the auction and p_0 is the equilibrium price in the primary market. The m long-term investors who participate in the auction cannot retrade in the secondary market, so that their budget constraint is

$$W_{k2} = W_{k0} + (1 - p_0)q_{k0} - 2\left(\lambda_j q_{k0} + \frac{\kappa}{2}q_{k0}^2\right)$$

We associate long-term investors with regional banks, insurance companies, and pension funds that buy in the auction and hold to maturity with minimal retrading in the secondary market. For simplicity, we interpret the auction as a new emission, so that the initial endowment of the asset is zero. Without loss of generality, we normalize initial wealth to zero, $W_{k0} = W_{j0} = 0$.

Information Structure The intercept of the marginal cost λ_i varies across agents. Agent i obtains a draw of λ_i such that $\lambda_i = \lambda + \varepsilon_i$ prior to bidding in the auction, where λ is an unknown parameter. Agent i privately observes λ_i and never reveals it to the public. We assume that $\lambda \sim \mathcal{N}(\bar{\lambda}, \sigma_\lambda^2)$ and that $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, where $\mathbb{C}\mathrm{ov}(\varepsilon_i, \varepsilon_j) = 0$ for $i \neq j$. The prior variance σ_λ^2 is interpreted as aggregate demand uncertainty, whereas the variance σ_ε^2 is interpreted as dispersion in holding costs across agents. The key distinction from Vives (2011) is that each of the signals enters each agents' payoff directly, so that an equilibrium does not collapse in the pure common value case. This is because agents' strategies will still depend on λ_i even in situations in which the price is a sufficient statistics of the aggregate information in the market. As opposed to Kyle (1989), we do not explicitly rely on noise traders or random supply. Further, the asset payoff is deterministic and there is no uncertainty about asset fundamentals. In contrast, aggregate uncertainty at $t = 0$, arises endogenously through the resale market as a function of how many dealer banks relative to long-term investors bid in the auction. All the agents in the economy agree about the terminal payoff of the asset, and there is no ex-ante information asymmetry. In Kyle (1989) and subsequent literature, informed traders observe a private signal about the a stochastic liquidation value, so that information sets are different across informed and uninformed investors. We use the convention that $\int_0^1 \varepsilon_i di = 0$ so that $\int_0^1 \lambda_i di = \lambda$.

Primary Market We model the auction as a multi-unit uniform price auction in the quadratic-normal setting (Vives, 2011). Auction participants compete in demand schedules. Their strategy is a price-contingent downward sloping demand schedule $\{q_{j0}(p)\}_{j=1}^{n}$, $\{q_{k0}(p)\}_{k=1}^{m}$. The auction rules are as in the canonical uniform price mechanism. The solution concept we adopt is a Bayesian-Nash equilibrium in demand schedules. The assumption that the number of bidders is finite and common knowledge is motivated by the fact that auctions are dominated by dealer banks. Further, although there is no formally established primary dealer system in Switzerland, the number of participants in the auctions tend to be much smaller than the number of investors trading in the secondary market.
The left panel of figure 2 plots the number of participants and the fraction of each issuance absorbed by the top three bidders. The number of bidders peaks in the early 1990s and has since then been declining. The top three bidders absorb around 75% of each issuance on average.

(a) Market concentration **(b)** Non-competitive bids and own tranches.

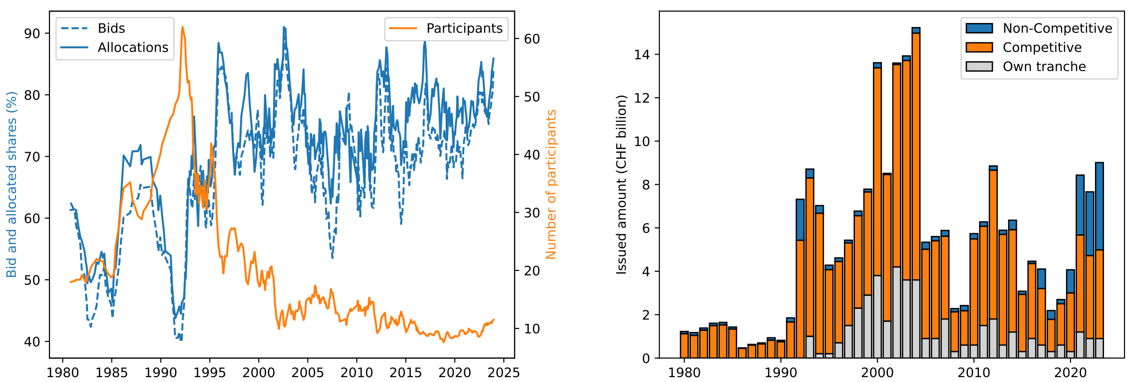

Figure 2: Market concentration and issuance composition. The left panel shows the time series of market concentration (left axis) and the number of bidders (right axis). Market concentration is defined as the share of total bid volume or issuance absorbed by the top three bidders. The right panel breaks down issuances by competitive bids, non-competitive bids, and own tranches. The sample is from 1980 to present.

Market Clearing From the perspective of dealer banks, supply at $t = 1$ is uncertain. Purchases of long-term investors $\sum_{k=1}^{m} q_{k0}$ effectively reduce the residual outstanding supply, so that dealer banks must form expectations about how much will be available for trading in the secondary market. To avoid this additional layer of complications, we assume that the government reopens the market and issues additional $Q_1 = \sum_{k=1}^{m} q_{k0}$ of the same bond at $t = 1$. This is consistent with the idea that the Treasury relies on reopenings to improve secondary market liquidity. Furthermore, the Treasury retains a sizable fraction of each issuance that reaches the secondary market only at a later stage. The right panel of figure 2 decomposes allocations across competitive bids, non-competitive bids, and own tranches. The Treasury retains part of the issue size for liquidity management reasons. Most of the issues are absorbed by competitive bids, but we observe an increasing share obtained by non-competitive bidders in the most recent years. Accordingly, the market clearing conditions in the secondary market is

$$Q_a - \sum_{k=1}^{m} q_{k0} + Q_1 = Q_a = \int_i q_{i1} di$$

Finally, the market clearing condition in the primary market is

$$\sum_{j=1}^{n} q_{j0} + \sum_{k=1}^{m} q_{k0} = Q_a$$

The market clearing price p_0 equalizes aggregate demand and supply. As usual, if there is no such price or if multiple prices exist, then no trade takes place. Strictly downward-sloping demand schedules are sufficient to rule out trivial equilibria (Rostek & Weretka, 2012).

3.2 Equilibrium in the Secondary Market

We solve the model backwards, starting from $t = 1$. The secondary market is large and competitive such that nobody is large enough to influence the price. Proposition 1 characterizes the equilibrium price of the bond and the quantity demanded by each agent. The proof, which is in the Appendix, solves for equilibrium prices and quantities by looking for a symmetric price-taking equilibrium in demand functions. What is important to our theory is that sets of investors in the primary and the secondary market are different.

Proposition 1 (Secondary market equilibrium)**.** *The equilibrium in the secondary market fully reveals the average cost λ. The market clearing price is*

$$p_1^* = 1 - \lambda - \kappa Q_a \tag{1}$$

Further, equilibrium demand is

$$q_{1i}^* = \frac{\lambda}{\kappa} - \frac{\lambda_i}{\kappa} + Q_a \tag{2}$$

The price p_1^* declines with the average intercept λ and the slope κ. Agents with a below-average marginal cost intercept λ_i purchase a quantity in excess of the per-capita supply Q_a, and vice-versa. The quadratic term in the cost function ensures that the solution for asset demand is bounded. The assumption of a continuum of agents is convenient because it implies that the equilibrium bond price p_1^* in Proposition 1 is fully revealing. Indeed, in the aggregate, the idiosyncratic noise ε_i vanishes, and agents learn the average intercept λ from the market clearing price. As opposed to Kyle (1989), agents submit demand schedules that depend on λ_i even when the price eventually reveals λ. The reason is that the private signal λ_i is payoff-relevant, and so will generate heterogeneity in agents' demand even with in the full-information case.

Remark on Auction Outcomes In auctions with resale, information disclosure by the auctioneer has implications on the final allocation of the asset, see e.g. Dworczak (2020). The outcome of the auction game reveals information about agents' types and asset holdings, which impact trading in the resale market through private information or market power. The assumption of a perfectly competitive secondary market that fully reveals λ implies that any post-auction differences in beliefs about λ are irrelevant. Further, although they are the only sellers in the secondary market, dealers do not exert market power. This modelling strategy is common in games where auction participants interact with each other in an aftermarket (Haile, 2001, 2003).

3.3 Equilibrium in the Primary Market

Taking p_1^* and q_{1i}^* as given, we next characterize a linear equilibrium in demand schedules. There are three main differences from Allen and Wittwer (2023), Malamud and Rostek (2017), and Rostek and Yoon (2021). First, the capital gain of the bond between the auction and the secondary market is endogenous. The average intercept λ is unknown ex-ante. Hence, holding the bond from period 0 to period 1 is still risky as the aggregate demand for bonds might be less than expected. We refer to this risk as aggregate demand risk. A higher value of λ implies that holding costs are higher, demand weaker, and prices lower. Second, we introduce two types of bidders that with different investment horizons, so that the equilibrium is no longer symmetric in general. Third, there is no uncertainty about the liquidation value of the asset.

A strategy is a mapping from the signal space to the space of non-increasing continuous functions as in Vives (2011). This assumption is made for tractability, although bidders usually submit discrete step functions rather than continuous demand schedules (Hortaçsu and Kastl (2012); Kastl (2011)). The construction of the equilibrium is standard. First, we conjecture that agents play linear strategies and characterize the random (inverse) residual supply curve that each bidder faces. Second, given the conjectured strategies, we rewrite terminal consumption as a quadratic function of the common value λ. Third, we derive best responses and solve for the unknown coefficients as a function of the model primitives.

In the model, the intercept of marginal costs λ_i plays two important roles. First, it is a noisy signal about future capital gains. Retrading in the secondary market generates a common value component due to the fact that future prices linearly depend on the same unknown parameter λ for all dealer banks. The common value impacts equilibrium strategies and also how bidding behavior responds to changes in uncertainty. In fact, dealer banks will have an incentive to learn about λ from the price p_0. Second, it introduces a linear penalty for holding the bond that is heterogeneous across agents. As a result, the signal enters directly into the payoff function, and an equilibrium exists even in the case of pure common values.

Inverse Residual Supply Heterogeneity in investment horizons prevents the equilibrium from being symmetric. However, within each type, each agent will submit the same demand schedule. We construct a Bayes-Nash equilibrium in demand schedules by conjecturing that dealer banks (indexed by D) and long-term investors (indexed by L) submit demand schedules of the form

$$\begin{aligned} q_j &= b_D - a_D p - c_D \lambda_j \quad &: \quad j \in \mathcal{N} = \{1, \ldots, n\} \\ q_k &= b_L - a_L p - c_L \lambda_k \quad &: \quad k \in \mathcal{M} = \{1, \ldots, m\} \end{aligned}$$

where we omit the time subscript to simplify notation. Market clearing requires

$$\sum_{j=1}^{n} q_j + \sum_{k=1}^{m} q_k = Q_a$$

From the perspective of dealer j, the (inverse) residual supply is

$$p = d_D q_j + I_{j,D}$$

where $d_D \doteq [(n-1)a_D + m a_L]^{-1}$ and $I_{j,D} \doteq d_D \{(n-1)b_D + m b_L - h_{j,D} - Q_a\}$ are the (endogenous) slope and intercept of the inverse (residual) supply curve. Therefore, for dealer j, the price p is informationally equivalent to the total signal $h_{j,D} = c_D \sum_{j' \neq j} \lambda_{j'} + c_L \sum_{k=1}^{m} \lambda_k$. By submitting a price-contingent schedule, each dealer bank essentially conditions on the random intercept of the inverse residual supply function $I_{j,D}$. The endogenous coefficient d_D determines the price impact of each dealer bank.

A similar argument reveals that the residual supply faced by long-term investors is $p = d_L q_k + I_{k,L}$, where the slope is $d_L \doteq [n a_D + (m-1)a_L]^{-1}$ and the total signal is $h_{j,L} = c_D \sum_{j=1}^{n} \lambda_j + c_L \sum_{k' \neq k}^{m} \lambda_{k'}$, so that the intercept is $I_{j,L} \doteq d_M \{n b_D + (m-1)b_L - h_{j,L} - Q_a\}$. From the perspective of a long-term investor, p is informationally equivalent to $h_{j,L} \neq h_{j,D}$. Hence, each type of agents faces a different (inverse) residual supply curve. As usual, the slope of the inverse residual supply that each agent faces depends on the strategies of the other type.

Dealers' Problem We rewrite dealers banks' objective substituting in the equilibrium p_1^* and q_{j1}^* in the secondary market from Proposition 1. Given the strategies of the other $n-1$ dealers and the m long-term investors, the terminal wealth of dealer j is

$$W_{j2} = (p_1^* - p_0)q_j - \lambda_j q_j - \frac{\kappa}{2} q_j^2 + (1 - p_1^*) q_{j1}^* - \lambda_j q_{j1}^* - \frac{\kappa}{2} (q_{j1}^*)^2$$

The structure is reminiscent of Vayanos and Wang (2011), with the difference that trading at time $t = 0$ occurs through a uniform price auction. The terminal payoff is a quadratic function of the unknown average cost λ, that is $W_{j2} = \pi_0(q_j) + \lambda \pi_1(q_j) + \frac{\lambda^2}{2\kappa}$ where the coefficients $\pi_0(q_j)$ and $\pi_1(q_j)$ are given in the Appendix. The best response of agent j therefore solves

$$\max_{q_j} -\mathbb{E}_0^D \left[e^{-\gamma \left\{ \pi_0(q_j) + \lambda \pi_1(q_j) + \frac{\lambda^2}{2\kappa} \pi_2 \right\}} \middle| \lambda_j, p_0 \right] \tag{3}$$

The next Lemma shows that the objective in (3) is equivalent to a quadratic function of q_{ja}.

Lemma 1 (Dealers' objective). *Given strategy profiles $q_{j'} = b - a_D p_0 - c_D \lambda_{j'}$ for $j' \neq j$ and*

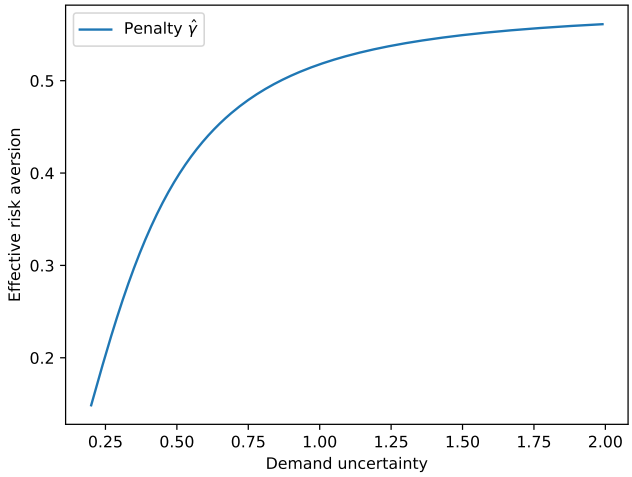

Figure 3: Effective risk aversion $\hat{\gamma}$. The model parameters are $\gamma = 5$, $\kappa = 3$, $\sigma_\varepsilon = 1$, $m = 4$, and $n = 3$. Demand uncertainty is in the range $\sigma_\lambda \in (0.2, 2)$.

$q_k = b - a_L p_0 - c_L \lambda_k$, *dealers j's best response solves*

$$\max_{q_j} \pi_0(q_j) + \mu_\lambda \pi_1(q_j) - \frac{\hat{\gamma}}{2}\left(\pi_1(q_j) + \mu_\lambda \kappa^{-1}\right)^2 - p_0 q_j \tag{4}$$

where $\mu_\lambda \doteq \mathbb{E}_0^D[\lambda|\lambda_j, p_0]$ *is dealer j's posterior mean* λ *conditional on* λ_j *and the price p. Given* $\Sigma_\lambda \doteq \mathbb{V}ar_0^D[\lambda|\lambda_j, p_0]$*, the effective risk aversion coefficient* $\hat{\gamma} > 0$ *is defined as*

$$\hat{\gamma} \doteq \frac{\gamma}{\Sigma_\lambda^{-1} + \gamma\kappa^{-1}}$$

The penalty $\hat{\gamma}$ *increases with the dispersion in marginal costs* σ_ε^2 *and declines with the number of bidders* $n + m$*. As* $n + m \to \infty$*, the penalty approaches zero* $\hat{\gamma} \to 0$.

The objective (4) shows that dealer banks, and dealer banks only, are subject to a penalty from bond holdings for three reasons. First, dealer banks must pay a quadratic holding cost, since $\kappa > 0$. Second, dealer banks refrain from trading due to price impact d_D. The third penalty endogenously arises through the interaction of risk aversion and imperfect information about the average marginal cost λ. When λ is known, the bond price p_1^* is deterministic, and the curvature in the objective comes solely from the quadratic cost κ and price impact. When λ is stochastic, however, the capital gain between the auction and the secondary market is also stochastic, and risk-averse agents demand an additional compensation, which we interpret as inventory risk in the spirit of Fleming et al. (2024). As a result, uncertainty about aggregate demand and signal dispersion impact bond demand through the effective risk aversion $\hat{\gamma}$. As an illustration of the mechanism, Figure 3 plots $\hat{\gamma}$ as a function of demand uncertainty σ_λ. This third penalty arises endogenously through the resale motive, and not because of uncertainty about the asset fundamentals. In fact, the penalty is higher when the dispersion in marginal costs σ_ε is large, when the number of participants is low, and when the prior variance σ_λ^2 is high. A higher number of bidders lowers the posterior variance $\mathbb{V}\text{ar}(\lambda|h_j, \lambda_j) = \Sigma_\lambda$, as the market aggregates more information about λ, and hence about future capital gains.

In quasi-linear settings with quadratic utility, demand uncertainty and bid dispersion do affect equilibrium strategies, but they typically do not affect the post-auction distribution of asset returns. This happens because only the uncertainty about payoffs is penalized, and that is independent of the distribution of signals and, consequently, demand uncertainty. For a safe asset with no fundamental uncertainty, penalizing fundamental uncertainty only is less suitable. This explicit penalty for demand uncertainty is absent in Rostek and Yoon (2021) because the payoff distribution is assumed to be independent of investors' private signals, and private information does not convey any information about the asset risk and return. In contrast, in Allen and Wittwer (2023) the volatility of the asset is independent of the signal precision and the number of participants. In our setting, a larger number of participants makes the price in the primary market more informative and mitigates demand uncertainty. This does not mean that the uncertainty over the quantity allocation is penalized, since bidders always condition their demand on the market clearing price.

Long-term Investors' Problem Taking equilibrium strategies as given, long-term investor k chooses a demand schedule to maximize

$$\max_{q_k} -\mathbb{E}_0^L \left[e^{-\gamma\left\{q_k(1-p)-2\lambda_k q_k-\kappa(q_k)^2\right\}} | \lambda_k, p \right]$$

Given that long-term investors do not access the secondary market, they pay the holding costs for purchases q_k twice, once for each period. However, their consumption is deterministic and they face no uncertainty. In fact, long-term investors' terminal value does not depend on the uncertain parameter λ. The problem simplifies to

$$\max_{q_k} q_k(1-p) - 2\lambda_k q_k - \kappa\,(q_k)^2 \tag{5}$$

As opposed to dealers, long-term investors refrain from trading only because of price impact and the quadratic holding cost κ, but not because of inventory risk associated with the stochastic capital gain. Furthermore, since their terminal wealth does not depend on λ, long-term investors have no incentive to learn from the equilibrium price p.

Best Responses and Equilibrium Strategies The solutions to problem (4) and (5) determine the best response of each agent and, therefore, the equilibrium in the auction game. The next Proposition characterizes necessary conditions on the coefficients to define a Bayes-Nash equilibrium in demand schedules.

Proposition 2 (Equilibrium of the auction game). *In a linear Bayes-Nash equilibrium, in demand schedules, each agent type submits the demand schedules*

$$\begin{aligned} q_{j0} &= b_D - a_D p_0 - c_D \lambda_j \quad &: \quad j = 1, \ldots, n \\ q_{k0} &= b_L - a_L p_0 - c_L \lambda_k \quad &: \quad j = 1, \ldots, m \end{aligned}$$

The slopes $\boldsymbol{a} = (a_D,\, a_L)$ *and the loadings on the private cost* $\boldsymbol{c} = (c_D, c_L)$ *solve the system*

$$a_D = (\kappa + \hat{\gamma}(\boldsymbol{c}) + d_D(\boldsymbol{a}))^{-1} \left\{1 - \mu^p_\lambda(\boldsymbol{a}, \boldsymbol{c})(\hat{\gamma}(\boldsymbol{c})\kappa^{-1} - 1)\right\} \tag{6}$$

$$c_D = (\kappa + \hat{\gamma}(\boldsymbol{c})) + d_D(\boldsymbol{a}))^{-1} \left(\hat{\gamma}(\boldsymbol{c})\kappa^{-1} + 1 - \mu^\lambda_\lambda(\boldsymbol{c})(\hat{\gamma}(\boldsymbol{c})\kappa^{-1} - 1)\right) \tag{7}$$

$$a_L = (2\kappa + d_L(\boldsymbol{a}))^{-1} \tag{8}$$

$$c_L = 2(2\kappa + d_L(\boldsymbol{a}))^{-1} \tag{9}$$

where $d_D(\boldsymbol{a})$ *and* $d_L(\boldsymbol{a})$ *are the slope of dealer banks' and long-term investors' residual (inverse) supply. The penalty* $\hat{\gamma}(\boldsymbol{c})$ *and the posterior loadings* $\mu^p_\lambda(\boldsymbol{c})$*,* $\mu^\lambda_\lambda(\boldsymbol{c})$ *are given in Lemma 1.*

Proposition 2 shows that, if an equilibrium exists, then the coefficients on price p and holding cost must satisfy equations (6) through (9). We refer to a_D and a_L as the slope of the demand curve, and we say that demand curves are steeper when either a_D and a_L are small. The coefficients $\boldsymbol{c} = (c_D, c_L)$ are the demand loadings on the intercept of marginal holding costs, whereas the coefficients $\boldsymbol{b} = (b_D, b_L)$ denote the demand intercepts.

Given equilibrium strategies, the equilibrium price in the primary market p_0 is

$$p_0^* = \frac{1}{na_D + ma_L}\left\{nb_D + mb_L - c_D \sum_j \lambda_j - c_L \sum_k \lambda_k\right\}$$

The revenue of the government is $R_0 \doteq p_0^* Q_a$ and the capital gain from the auction to the secondary market is $p_1^* - p_0^* = 1 - \lambda - \kappa Q_a - p_0^*$.

3.4 Special Cases

We characterize the equilibrium demand schedules for two special cases that have closed-form solutions. The first case is the pure private value (PV) case in which only long-term investors bid in the auction. The second case is the pure interdependent values environment, which we henceforth refer to as common value (CV), in which only dealer banks bid in the auction.

Pure Private Values We characterize the (unique) Bayes-Nash equilibrium in demand functions when only long-term investors participate to the auction.

Proposition 3 (Private values)**.** *When only long-term investors bid* ($n = 0$)*, the equilibrium is symmetric and each bidder submits a linear demand schedule* $q_{k0} = a_L^{PV} - b_L^{PV} p_0 - c_L^{PV} \lambda_k$*, where*

$$a_L^{PV} = \frac{m-2}{m-1}\frac{1}{2\kappa}$$
$$c_L^{PV} = \frac{m-2}{m-1}\frac{1}{\kappa}$$

Furthermore, $\mu^\lambda_\lambda = 0$.

The equilibrium strategies in Proposition 3 do not depend on risk aversion, demand uncertainty, and cost dispersion. The coefficient on price a_L^{PV} only depends on the slope of marginal costs only κ and the total number of participants m. As the number of long-term investors grow large, the equilibrium converges to the perfectly competitive case.

Long-term investors buy at the auction and hold the asset to maturity. As a result, the post-auction capital gain is not payoff relevant. The private holding cost λ_k is a perfect signal of long-term investors' own payoffs. Therefore, long-term investors' consumption is deterministic, since movements in post-auction prices do not impact their consumption. In addition, since their payoff is independent of λ, long-term investors have no incentive to learn from prices.

Pure Common Values Proposition 4 characterizes the equilibrium when only dealer banks participate in the auction. Unlike in Kyle (1989), the fundamental value of the asset is known, and the signal λ_j is directly payoff-relevant.

Proposition 4 (Common values)**.** *When only dealer banks bid ($m = 0$), the equilibrium is symmetric and each bidder submits a linear demand schedule $q_{j0} = a_D^{CV} - b_D^{CV} p_0 - c_D^{CV} \lambda_j$, where*

$$a_D^{CV} = \frac{(n-2)(\hat{\gamma}+\kappa) + n\tau_\lambda(\hat{\gamma}-\kappa)}{(n-1)(\hat{\gamma}+\kappa)\left[(\hat{\gamma}+\kappa) - n\tau_\lambda(\hat{\gamma}-\kappa)\right]}$$
$$c_D^{CV} = \frac{(n-2)(\hat{\gamma}+\kappa) + n\tau_\lambda(\hat{\gamma}-\kappa)}{(n-1)(\hat{\gamma}+\kappa)\kappa}$$

where $\tau_\lambda = \frac{\sigma_\lambda^2}{\sigma_\varepsilon^2 + n\sigma_\lambda^2}$. Furthermore, $\mu_\lambda^\lambda(\boldsymbol{c}) = 0$.

As in the private value case, both a_D^{CV} and c_D^{CV} vary with the number of bidders and with the slope of marginal holding costs. On the other hand, bidding strategies are now sensitive to demand uncertainty σ_λ^2 and cost dispersion σ_ε^2. As it is standard in the literature, an increase in the number of participants leads to flatter demand curves primarily because more bidders share the aggregate risk, but also because price impact declines (Keloharju et al., 2005; Kyle, 1989). First, an increase in the prior uncertainty about aggregate demand makes the post-auction capital gain riskier, and risk-averse bidders incorporate a demand risk-premium through the penalty $\hat{\gamma}$. Second, dealers have an incentive to learn about the common value component λ, since λ is payoff relevant. Dealer banks use prices to learn about the future price of the asset and to reduce demand uncertainty. To see why these are two distinct forces, note that a_D^{CV} and c_D^{CV} are different than the private value case even when dealers are risk-neutral, that is $\gamma = 0$. The posterior mean loading on the private signal $\mu_\lambda^\lambda(\boldsymbol{c})$ is zero as in the pure private value case of Proposition 3. This is because the price p_0 is a sufficient statistics of the aggregate information in the market.

The equilibrium in Proposition 4 is in contrast to Kyle (1989) and Vives (2011). The equilibrium with pure common value does not collapse even without uncertainty or noise traders. The common value component arises endogenously through the future resale. This is also in contrast to Biais, Martimort, and Rochet (2000), where agents are risk-neutral and common values come from information asymmetries about a stochastic liquidation value. In our setting with safe assets, there is no information asymmetry about the asset payoff.

3.5 General Case

We establish existence of a Bayes-Nash equilibrium for the general case in which both long-term investors and dealers participate in the auction. Proposition 5 also proves that there is always an equilibrium in which all bidders submit downward sloping demand schedules, implying existence and uniqueness of the market clearing price p_0.

Proposition 5 (Equilibrium existence)**.** *Let $\boldsymbol{c} = (c_L, c_D)$, $n > 1$, $m > 1$. There exists a Bayes-Nash equilibrium with downward sloping demand schedules $a_D > 0$ and $a_L > 0$ such that*

$$
\begin{aligned}
a_D(\boldsymbol{c}) &= c_L \frac{2 - m + \kappa(m-1)c_L}{2n(1 - c_L\kappa)} \\
a_L(\boldsymbol{c}) &= \frac{1}{2} c_L
\end{aligned}
$$

The coefficients $\boldsymbol{c}^* = (c_L^*, c_D^*)$ solve the system of equations $f(\boldsymbol{c}) = (f_1(\boldsymbol{c}), f_2(\boldsymbol{c})) = \boldsymbol{c}$ defined by

$$
\boldsymbol{c} = \left(\frac{1 - \mu_\lambda^p(\boldsymbol{c})(\hat{\gamma}(\boldsymbol{c})\kappa^{-1} - 1)}{\kappa + \hat{\gamma}(\boldsymbol{c}) + d_D(\boldsymbol{c})} \cdot \frac{2n(1 - c_L\kappa)}{2 - m + \kappa(m-1)c_L} \quad , \quad \frac{\hat{\gamma}(\boldsymbol{c})\kappa^{-1} + 1 - \mu_\lambda^\lambda(\boldsymbol{c})(\hat{\gamma}(\boldsymbol{c})\kappa^{-1} - 1)}{\kappa + \hat{\gamma}(\boldsymbol{c}) + d_D(\boldsymbol{c})} \right)
$$

where $d_D(\boldsymbol{c}) = [(n-1)a_D(\boldsymbol{c}) + m a_L(\boldsymbol{c})]^{-1}$ and

$$
\begin{aligned}
\mu_\lambda^\lambda(\boldsymbol{c}) &= \frac{m\sigma_\lambda^2 c_L(c_L - c_D)}{[(n-1)c_D + mc_L]^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)} \\
\mu_\lambda^p(\boldsymbol{c}) &= -\frac{\sigma_\lambda^2((n-1)c_D + mc_L)\frac{c_L}{2}\frac{2-\kappa c_L}{1-\kappa c_L}}{[(n-1)c_D + mc_L]^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)} \\
\Sigma_\lambda(\boldsymbol{c}) &= \frac{\left[c_D^2(n-1) + c_L^2 m\right]\sigma_\varepsilon^2\sigma_\lambda^2}{[(n-1)c_D + mc_L]^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)}
\end{aligned}
$$

Further, the coefficients $c_L \in \left(\frac{1}{\kappa}\frac{m-2}{m-1}, \frac{1}{\kappa}\right)$ and $c_D > 0$ are strictly positive.

In any equilibrium with downward sloping demand curves, c_L and c_D are strictly positive. Furthermore, c_L lies between the pure private value case of Proposition 4 and the perfectly competitive case. The next proposition shows that there does not exist an equilibrium in which strategies are symmetric across investor types.

Proposition 6 (Asymmetry)**.** *Suppose $m > 0$ and $n > 0$. Then, the equilibrium strategies are not symmetric across types. Thus, $\mu_\lambda^\lambda(\boldsymbol{c}) \neq 0$.*

Intuitively, equations (8) and (9) imply that $c_L = 2a_L$, whereas the ratio between a_D and c_D is typically different from two. On the one hand, dealer banks find it optimal to learn about λ in the auction. As a result, the equilibrium strategies of dealer banks depend directly on the posterior mean μ_λ. On the other hand, long-term investors are indifferent about what they learn from prices, since λ is not payoff-relevant and their terminal wealth is deterministic. Given that long-term investors face no risk, their strategies do not directly respond to changes in demand uncertainty σ_λ^2 and bid dispersion σ_ε^2. The denominator of equations (6) and (7) is $\kappa + \hat{\gamma} + d_D$, which means that dealer banks restrict trading because of holding costs, risk aversion, and price impact. In contrast, the denominator of equations (8) and (9) is $\kappa + d_L$, which implies that long-term investors restrict trading only because of price impact and holding costs, but not because of risk. Importantly, this does not imply that the equilibrium strategies of long-term investors are independent of demand uncertainty and cost dispersion. The reason this is not the case is that changes in σ_λ^2 and σ_ε^2 do affect the equilibrium strategies of dealer banks, impacting the slope of the residual demand that long-term investors face. As a result, even if the bond is completely safe from the perspective of long-term investors, a_L and c_L still vary with demand uncertainty.

We now derive a sufficient condition on bidders' composition such that dealer banks are more sensitive in absolute terms to demand uncertainty relative to long-term investors

Proposition 7 (Sensitivity to demand risk)**.** *If $(m-1)(m-2) > n$, then*

$$\left|\frac{\partial a_D}{\partial \sigma_\lambda^2}\right| > \left|\frac{\partial a_L}{\partial \sigma_\lambda^2}\right|$$

Proposition 7 demonstrates that an increase in demand uncertainty impacts the slope of the demand schedules of both investor types. While we are not able to provide a sharper analytical characterization of the sign of these derivatives, the numerical results below indicate that $\frac{\partial a_D}{\partial \sigma_\lambda^2} < \frac{\partial a_L}{\partial \sigma_\lambda^2} < 0$. An increase in demand uncertainty implies that both types submit steeper demand schedules. However, the effect is stronger for dealer banks relative to long-term investors.

Numerical Results We present numerical solutions for the general case in which both dealer banks and long-term investors participate in the auction, solving for equilibrium strategies in two steps. First, equations (6) through (9) in Proposition 2 form a non-linear system of four equations with four unknowns. This system is solved numerically given parameter values for the slope of marginal holding costs κ, risk aversion γ, demand uncertainty σ_λ, cost dispersion σ_ε, expected demand $\bar{\lambda}$ and supply Q_a. Second, given a solution to equations (6)–(9), we compute $b_L = a_L$. Substituting a_D, a_L, c_D, c_L and b_L into the non-linear equation that characterizes b_D yields another nonlinear equation, which we also solve numerically.

Figure 4 plots the price loadings $\boldsymbol{a} = (a_D, a_L)$ of dealer banks and long-term investors against the ratio of dealer banks to total participants. When a_D increases, we say that bidders submit flatter demand curves. When there are no dealer banks, the equilibrium converges to the pure private value case of Proposition 3. In the pure private value case, demand curves are the flattest. This is because long-term investors are not affected by demand risk σ_λ, and only restrict trading because of the holding cost κ and price impact. At the other extreme, when the fraction of dealer banks approaches one, the equilibrium approaches to the common value case, where demand curves are the steepest. A high proportion of dealer banks intensifies adverse selection, and dealer banks refrain from trading also because of the risky capital gain. For intermediate cases in which the ratio of dealer banks to total participants is between zero and one, long-term investors submit flatter demand curves relative to dealer banks. The gap between the slope of the demand curves shrinks as the ratio of dealer banks to total participants approaches zero.

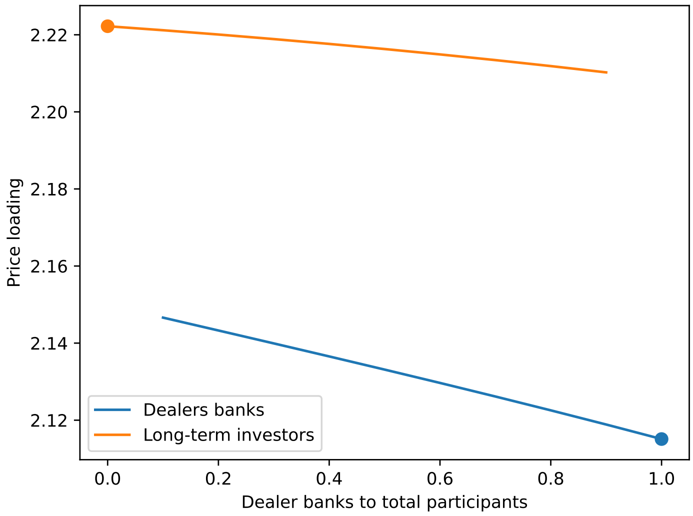

Figure 4: Convergence to special cases. The figure plots the coefficients a_D and a_L as the auction participant composition shifts from zero dealer banks to only dealer banks. The markers indicate the pure private value case (orange) and the pure common value case (blue), with total participants fixed at $N = m + n = 10$. The model parameters are $\gamma = 3$, $\kappa = 0.2$, $\sigma_\varepsilon = 0.6$ and $\sigma_\lambda = 1$.

Figure 5a plots the price loading $\boldsymbol{a} = (a_D, a_L)$ for dealer banks and long-term investors. When uncertainty about average costs is higher, both types submit steeper demand curves. Our calibration suggests that $\frac{\partial a_D}{\partial \sigma_\lambda} < 0$ and $\frac{\partial a_L}{\partial \sigma_\lambda} < 0$. Further, the effect is stronger for dealer banks than for long-term investors, so that $\frac{\partial a_D}{\partial \sigma_\lambda} < \frac{\partial a_L}{\partial \sigma_\lambda} < 0$ as in Proposition 7. The relative magnitudes of the price loadings across bidder types depend on the level of demand uncertainty. When prior uncertainty is large, dealer banks have very noisy estimates of future prices, so that the post-auction capital gain is riskier. Because of this, the risk-averse dealer banks act more cautiously and submit steeper demand curves. In contrast, long-term investors adjust their strategies only to the extent that an increase in σ_λ influences the price impact of their trades. This effect vanishes in the pure private value case, in which strategies are independent of demand uncertainty.

Conversely, Figure 5b reveals that the impact of higher demand uncertainty on the signal loading may be asymmetric across types. In the calibration with $\gamma = 3$, long-term investors trade more aggressively on their own private signal than dealer banks, so that $c_L > c_D$ for all $\sigma_\lambda \in (0, 1)$. Higher demand uncertainty leads dealer banks to trade even more aggressively on their own signal, whereas long-term investors trade less aggressively. This is because the heterogeneity in investment horizons generates different incentives to learn from prices across investor types. Long-term investors refrain from trading too aggressively on their private information because doing so will reduce capital gain uncertainty and increase dealer banks' willingness to pay.

Asymmetric responses of signal loadings $\boldsymbol{c}$ to σ_λ is not a universal feature across all parameter configurations. We find that when the coefficient of risk aversion is small, e.g. $\gamma = 1$, an increase in demand uncertainty impacts the signal loading of long-term investors and dealer banks in the same direction. When risk aversion is low, the reduction in the demand risk premium is very small, and long-term investors trade more aggressively. The asymmetry emerges for a large enough coefficient of risk aversion because learning about the capital gain becomes particularly valuable when dealer banks' risk aversion is moderately large. Long-term investors find it beneficial to trade against uninformed dealer banks and keep demand risk premia high to lower the auction price. This mechanism only emerges when long-term investors and dealer banks simultaneously participate in the auction. As shown Propositions 3 and 4 for the pure private and pure common value cases, the impact of an increase in σ_λ always has the same sign for all bidders. Our main takeaway is that heterogeneity in investment horizons can potentially affect the way in which the primary and the secondary markets aggregate information.

(a) Loading on price. **(b)** Loading on private signal.

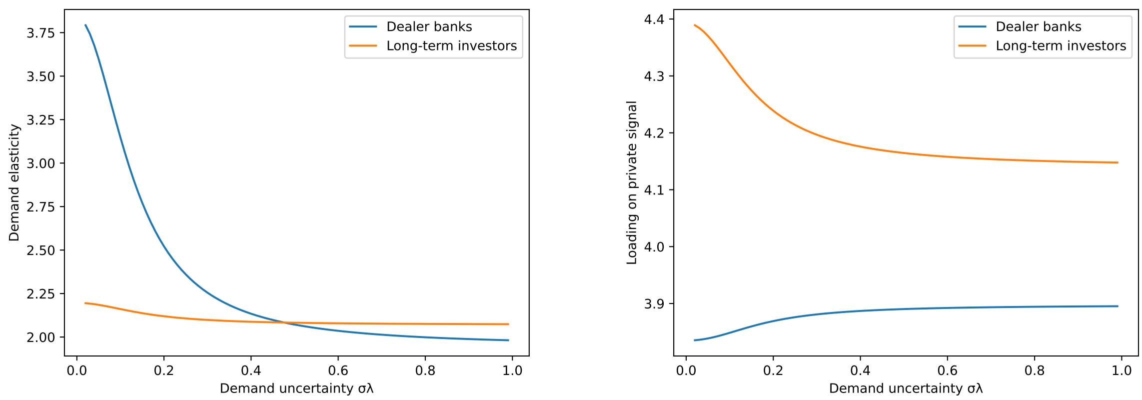

Figure 5: The left panel plots the price loading $\boldsymbol{a} = (a_D, a_L)$ for dealer banks and long-term investors. The right panel plots the loading on the private signal $\boldsymbol{c} = (c_D, c_L)$ also for dealer banks and long-term investors. The model parameters are $\gamma = 3$, $\kappa = 0.2$, $\sigma_\varepsilon = 0.6$, and $\sigma_\lambda \in (0.2, 1)$. The investor composition is $m = 4$ long-term investors and $n = 3$ dealer banks. The average holding cost is $\lambda = 0.25$, and the prior mean is $\bar{\lambda} = 0.05$. Bond supply at the auction is $Q_a = 1$.

Informational Advantages and Risk Aversion A common feature of competitive rational expectations equilibria (Admati, 1985; Diamond & Verrecchia, 1991) is that some traders have an informational advantage due to more precise private signals. Heterogeneity in the precision of private information also leads to asymmetric equilibria in which bidding strategies depend on traders' information sets. Competitive equilibria typically imply that traders with an information advantage submit flatter demand curves relative to uninformed traders (Vives, 2008). The same holds in models of imperfect competition where traders no longer take prices as given (Kyle, 1989), although price impact complicates an explicit characterization of the slope of the demand schedules. Because primary dealers generally have informational advantages (Boyarchenko et al., 2021; Hortaçsu & Kastl, 2012), such source of heterogeneity would imply that dealer banks' demand curves are flatter. Conversely, consistent with the empirical findings in Section 4, our model instead implies that dealer banks' may submit *steeper* demand curves ($a_L > a_D$) relative to long-term investors when demand uncertainty is sufficiently large.

Post-auction Capital Gains The post-auction capital gain is $R_{0|1} \doteq p_1^* - p_0^* = 1 - \lambda - \kappa Q_a - p_0^*$. From the perspective of dealer bank j, the subjective distribution of the post-auction capital gain conditional on the private signal λ_j and p_0 is

$$R_{0|1} \mid p_0, \lambda_j \sim \mathcal{N}(1 - \mu_\lambda - \kappa Q_a - p_0^*, \Sigma_\lambda)$$

where μ_λ and Σ_λ are given in Proposition 5.

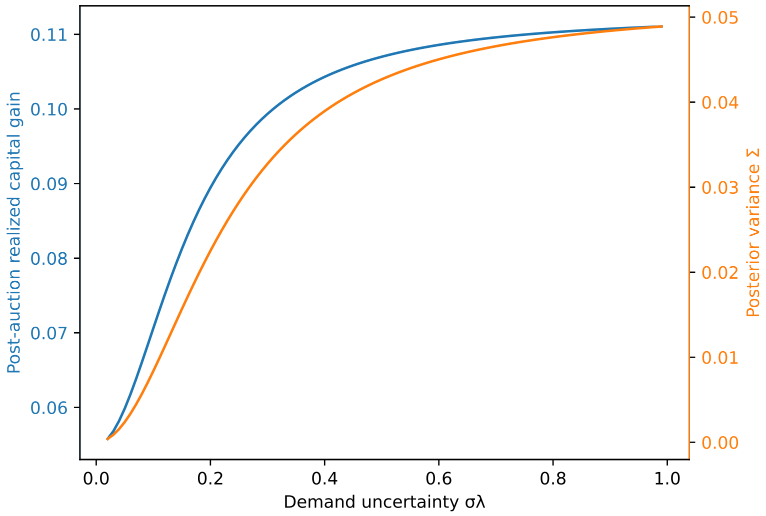

Figure 6: Realized post-auction capital gain and posterior variance Σ_λ. The blue line plots the *realized* post-auction capital gain $p_1^* - p_0^*$ against demand uncertainty. The orange line plots the posterior variance for dealer banks Σ_λ against demand uncertainty. The model parameters are $\gamma = 3$, $\kappa = 0.2$, $\sigma_\varepsilon = 0.6$, and $\sigma_\lambda \in (0.2, 1)$. The investor composition is $m = 4$ long-term investors and $n = 3$ dealer banks. The average holding cost is $\lambda = 0.25$, and the prior mean is $\bar{\lambda} = 0.05$. Bond supply at the auction is $Q_a = 1$.

Figure 6 plots the *realized* capital gain against demand uncertainty σ_λ. Risk-averse dealer banks require an additional compensation for their exposure to aggregate risk between the auction and the secondary market. The aggregate risk premium is a function of the total supply allocated to dealer banks $Q_{a,D} = \sum_{j=1}^{n} q_{j0}$ and of aggregate demand uncertainty σ_λ^2. Accordingly, the blue line shows that, in our calibration, the realized capital gain monotonically increases in demand uncertainty σ_λ. On the one hand, an increase in σ_λ leads to steeper demand curves. On the other hand, risk-averse dealer banks charge an additional risk premium to hold the bond that is positively related to the conditional variance of the capital gain Σ_λ. The orange line in Figure 6 shows that higher aggregate demand risk translates into higher posterior variance.

3.6 Predictions

Guided by our theory, we now formulate four empirical predictions on bidding behavior and post-auction price dynamics that we test in Section 4. The first two predictions follow directly from Propositions 5 and 7 and relate to the behavior of the slope of demand schedules across investor types in response to an increase in uncertainty.

Prediction 1 (Proposition 5)**.** *Both long-term investors and dealer banks submit downward sloping demand schedules. Dealer banks' demand schedules are steeper.*

Prediction 2 (Proposition 7)**.** *Both long-term investors and dealer banks submit steeper demand curves as demand uncertainty increases. Dealer banks exhibit greater sensitivity to demand uncertainty and holding cost dispersion relative to long-term investors.*

Informed by our numerical solution, we turn to post-auction return predictability. Dealer banks are more cautious when demand uncertainty is high and submit steeper demand schedules. This demand rigidity captures a risk premium that dealer banks require until they can resell their inventories in the secondary market. The more rigid dealer banks' demand is, the higher the associated demand risk premium. Empirically, this suggests that inelastic demand from dealer banks at the auction predicts positively predicts post-auction excess bond returns. On the other hand, the elasticity of long-term investors can also influence post-auction price dynamics. If dealer banks exhibit inelastic demand while other investors remain elastic, market makers can quickly sell their positions, resulting in a short-lived impact on bond returns. Conversely, if long-term investors also display inelastic demand, it will take longer for market makers to offload their inventories. Hence, when both long-term investors and dealer banks exhibit inelastic demand, we expect returns to remain predictable over longer horizons. This prediction finds support in Albuquerque et al. (2024). Yet, our model precisely describes the mechanism and our data on bidders' identities allows us to distinguish between dealers and long-term investors, so that we expand on their work by testing additional implications related to investor composition and heterogeneity in expected holding periods.

Prediction 3 (Return Predictability)**.** *Inelastic demand from dealer banks at the auction positively predicts post-auction returns. Additionally, when both dealer banks and long-term investors exhibit inelastic demand, returns become predictable over longer horizons.*

Our theory assumes a competitive secondary market, so that there are neither trading frictions nor transaction costs at $t = 1$. Therefore, the framework cannot provide explicit predictions about how secondary market liquidity impacts bidding behavior and, vice versa, how the auction allocation affects secondary market liquidity. However, we propose extending the empirical analysis with an additional prediction based on heuristic reasoning. We conjecture that agents' expectations of greater secondary market liquidity should alleviate demand rigidity, particularly for dealer banks with shorter expected holding periods and market-making obligations (Amihud & Mendelson, 1986). Our hypothesis is that a less liquid secondary market likely impairs market-making activities, adversely affecting dealer banks' profits from resale. On the other hand, we expect long-term investors with longer expected holding periods to be less affected by secondary market liquidity, and even sort into cheaper securities that are highly illiquid (Musto et al., 2018). Although this final prediction is partly outside the model, it expands the scope of our analysis and provides an additional test of our assumption on heterogeneity in investment horizons.

Prediction 4 (Secondary Market Liquidity)**.** *Dealer banks submit steeper demand curves when they anticipate the secondary market to be more illiquid, whereas the bidding strategies of long-term investors remain essentially unaffected.*

4 Demand Heterogeneity and Investment Horizons

We explore demand heterogeneity for Swiss Treasury bonds and empirically validate our theoretical model by testing Predictions 1 through 4. First, we document heterogeneity in bidding behavior and explore its determinants over time. We then demonstrate that dealer banks' demand curves become steeper relative to all the other bidders when demand uncertainty is higher, measured in terms of bond return volatility or bid dispersion, and when the secondary market is more illiquid.

4.1 Bidder Demand Schedules

We document heterogeneity in demand for safe government debt by inspecting cross-sectional differences in bidding strategies. Henceforth, j indexes auctions, i indexes bidders, and k indexes bid steps. We construct two measures of bidding heterogeneity: one for the level and one for the slope of demand schedules. First, we calculate the quantity-weighted yield *discount* for each bidder as

$$\text{Discount}_{ij} = \sum_{k=1}^{K} w_{ijk} B_{ijk} - y_j$$

where $w_{ijk} = \frac{Q_{ijk}}{\sum_{k=1}^{K} Q_{ijk}}$. For security reopenings, discount$_{ij}$ is the difference between the average bid yield of bidder i and the secondary market yield y_j of the same bond at the end of the auction closing day. A positive value means that bidder i is bidding, on average, less than the secondary market price. Using the discount as a normalized measure of bid level is standard in the literature, as it captures bidders' strategic bid shading (Keloharju et al., 2005; Nyborg et al., 2002). The discount is different from the yield spread in Table 1, which is the difference between the auction market clearing price b_j and the secondary market yield y_j.

Second, we collect price-quantity pairs $(B_{ijk}, \mathcal{Q}_{ijk})$ such that $\mathcal{Q}_{ijk}$ equals the sum of the quantity bid at bid price B_{ijk} or above. We include non-competitive bids by adding the non-competitive bid quantity $\mathcal{Q}_{ijk}^{nc}$ to all quantity bids $\mathcal{Q}_{ijk}$. Non-competitive bids are always fully allocated regardless of the price. As a result, these bids do not have an associated price, but they shift demand schedules upwards at each bid step. We then estimate the elasticity of demand β_{ij} (TE, total elasticity) by regressing the log bid quantity $q_{ijk} = \log \mathcal{Q}_{ijk}$ onto a constant and the log bid price $b_{ijk} = \log b_{ijk}$. We interpret the coefficient β_{ij} as the percentage reduction in bid quantity in response to a one percent increase in bid price. We use the elasticity of demand rather than the absolute slope for better cross-sectional comparisons, as elasticities are scale-invariant and offer a more consistent measure across bidders with significantly different bidding volumes (Albuquerque et al., 2024). We only consider demand schedules with five or more bid steps. A higher value of β_{ij} means that demand is more elastic and that demand curves are flatter. As alternative measures, we estimate the elasticity of demand by dropping the highest and the lowest bids ($\hat{\beta}_{ij}$, intermediate elasticity IE) and by considering winning bids only ($\tilde{\beta}_{ij}$, winning elasticity WE). For all three measures, we report the negative logarithm of the coefficient estimate

Table 2 reports summary statistic on bidding behavior. Across all the 530 auctions, we observe 8'699 unique demand schedules. On average, bidders submit four bid steps, with a median of three bid steps. The average demand schedule accounts for approximately 6.09% of the total bidding volume. Similarly, each bidder obtains, on average, 6.09% of the total supply. In many cases, however, the allocation is highly concentrated among relatively few bidders, with the top bidders receiving a sizable share of the issue as in Figure 2. The average total elasticity (in logs) is 5.08 (level is 160.77), and IE and WE are roughly the same. The average magnitude of our estimates is comparable to those of Albuquerque et al. (2024) for Portuguese Treasury auctions. The range of demand elasticities is substantial, reflecting significant heterogeneity across bidders. The standard deviation of total elasticity is approximately one, and the elasticity of demand ranges from 0.56 to 7.50 (in logs). As shown in Table 9 in Appendix B.1, the demand elasticity is negatively associated with volatility, bid dispersion, illiquidity, and maturity, whereas it is positively associated with the number of participants and total supply. The positive relationship with supply suggests that the government may increase issuance om the margin when observing flatter demand curves. Auction participants bid higher yields than in the secondary market, so that the yield discount is on average positive. There is, however, substantial cross-sectional dispersion in yield discounts, ranging from a minimum of −1.45% to a maximum of 3.94%. Some bidders bid prices that are regularly higher than in the secondary market.

	N	**Mean**	**SD**	**Min**	**Median**	**Max**
Bid steps	8'699	4.15	4.11	1.00	3.00	38.00
Bid quantity	8'699	37'015	97'165	1.00	10'000	4'040'380
Bid share	8'699	6.09	10.56	0.00	1.59	92.71
Allocated quantity	8'699	21'743	53'244	0.00	4'859	1'026'580
Allocated share	8'699	6.09	11.62	0.00	1.19	98.56
TE β_{ij}	2'279	5.08	0.99	0.56	5.18	7.50
IE $\hat{\beta}_{ij}$	2'279	5.19	1.00	−4.34	5.29	8.01
WE $\tilde{\beta}_{ij}$	1'601	5.26	1.14	0.36	5.35	7.85
Discount$_{ij}$	3'482	0.02	0.12	−1.45	0.02	3.94

Table 2: Descriptive statistics of bidding behavior. Bid steps is to the number of price-quantity pairs submitted by each bidder. Bid quantity is the total bid volume. Bid share is the fraction of total bid volume tendered by each bidder. Allocated quantity denotes the awarded volume per bidder, and allocated share is the fraction of total supply allocated to each bidder. Total elasticity (TE) is the demand elasticity using all bids. Intermediate elasticity (IE) excludes the highest and the lowest bids. Winning elasticity (WE) uses only winning bids. Discount is the difference between the quantity-weighted bid yield and the secondary market yield. The sample period is from 1980 to present.

Figure 7 demonstrates the heterogeneity in bidding behavior by plotting time-series averages of total elasticity and bidder yield discounts. The left panel shows that there is significant dispersion in the level of demand. Approximately two-thirds of the bidders offer bid yields that exceed those in the secondary market. In contrast, the remaining bidders are more aggressive, often bidding higher than the price in the secondary market. The right panel displays the average time-series demand elasticity. The elasticity of the demand schedules (in logs) exhibits significant cross-sectional variation. The two most elastic bidders have log demand elasticities of 6.74 and 6.64, whereas the least elastic bidder has a log elasticity of 4.48. In levels, this disparity translates into a remarkable difference of 757.32. As in our theoretical framework for large enough σ_λ, dealer banks are typically less elastic, on average, than long-term investors. All dealer banks, except one, have a demand elasticity of no more than 5.22. Conversely, most long-term investors submit significantly flatter demand curves.

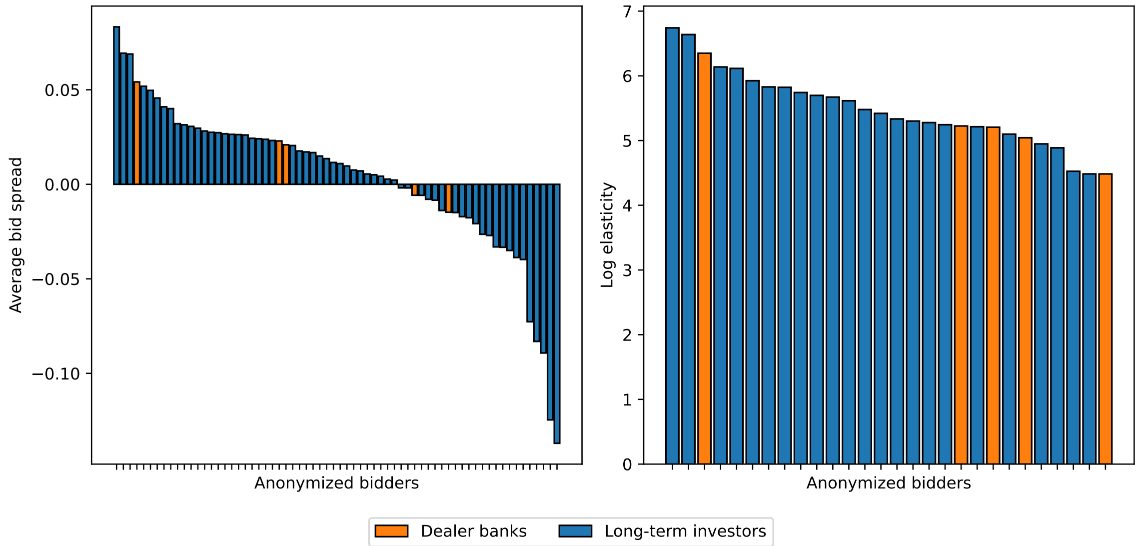

Figure 7: Time series averages of yield discount and demand elasticity. The left panel displays average yield discounts for all bidders. The right panel shows the total demand elasticity (TE). The elasticity is calculated using all bid steps submitted by a bidder, provided there are at least five. We report the log of the negative estimated elasticity. The sample period is from 1980 to present.

Given that long-duration bonds are more sensitive to interest rate movements, we expect yield discounts and demand elasticities to vary with bond maturity. While the relation between yield discounts and bond maturity is ex-ante ambiguous, demand elasticities typically decline as bond maturity increases. Table 3 presents the average yield discounts and demand elasticities separately for dealer banks and long-term investors for different maturity intervals.

	Yield discount				Log demand elasticity			
Maturity	$[2, 10)$	$[10, 15)$	$[15, 20)$	$[20, 50]$	$[2, 10)$	$[10, 15)$	$[15, 20)$	$[20, 50]$
Dealer banks	0.009	0.002	0.029	0.008	5.293	5.055	4.757	4.272
Long-term investors	0.025	0.013	0.034	0.025	5.627	5.495	4.860	4.511
Cantonal banks	0.017	0.019	0.053	0.035	5.533	5.505	4.671	4.565
Regional banks	0.072	0.042	0.069	0.048	5.865	5.454	4.603	4.177
Foreign-controlled	0.021	0.017	0.035	0.027	5.767	5.624	4.808	4.442
Non-finance companies	0.035	0.011	0.047	0.064	–	4.883	4.851	4.440
Others	0.025	-0.003	-0.016	-0.011	5.621	5.340	5.735	4.651

Table 3: Time series average yield discount and demand elasticity by maturity interval. The elasticity is calculated using all bid steps submitted by a bidder, provided there are at least five. We report the log of the negative elasticity. We group long-term investors into cantonal banks, regional banks, foreign-controlled entities, non-finance companies, and others. The sample period is from 1980 to present.

Two clear patterns emerge across all maturity intervals. First, dealer banks bid lower secondary market discounts as compared to long-term investors, so that dealer banks' underpricing is less severe. Decomposing longer-term investors across the five bidder categories, we observe that long-term investors' underpricing is a robust pattern across all categories with the exception of other bidders. Auction underpricing relative to dealer banks is particularly severe for regional banks and non-finance companies. Second, as predicted by the theory, dealer banks submit steeper demand schedules compared to long-term investors. We interpret this pattern through the lens of aggregate demand risk: both bidder types shade their bids to account for price impact and holding costs, but dealer banks also shade their bids due to risk aversion, reducing their elasticity even further. Another consistent finding is that demand elasticity decreases with bond maturity across all investor groups, aligning with Greenwood and Vayanos (2014). In addition, there is substantial heterogeneity across different groups of long-term investors.

4.2 Sensitivity to Aggregate Demand Risk

We formally test Predictions 1 and 2 by comparing the responses of dealer banks and long-term investors to an increase in demand uncertainty. We measure aggregate demand uncertainty as the volatility of daily bond returns $\sigma_{j-21,j}$ in the month prior to the auction, provided that there are at least fifteen observations. Second, we proxy holding cost dispersion as the cross-sectional standard deviation $\sigma_{b,j}$ in quantity-weighted bid yields within each auction.

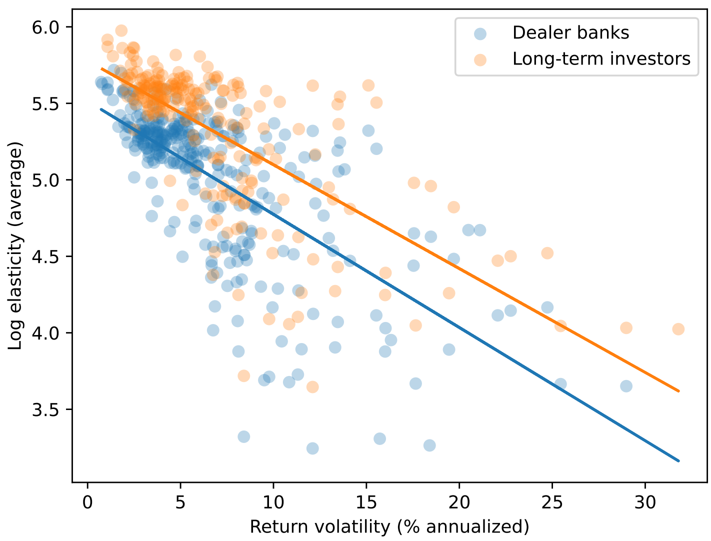

Figure 8: Bond return volatility and adjusted demand elasticity. The blue dots indicate dealer banks. The orange dots indicate long-term investors. Return volatility is the standard deviation of daily bond returns in the month prior to the auction, provided that there are at least 15 observations. Log elasticity is the average total elasticity β_{ij}. We adjust the demand elasticity for maturity by projecting β_{ij} onto a constant and bond maturity, and computing the fitted values. The solid lines represent linear fits of adjusted elasticity onto return volatility. The sample covers reopenings from 2000 to present.

Figure 8 plots the annualized return volatility in the month prior to the auction date against the average elasticity of demand separately for dealer banks and long-term investors, adjusted for bond maturity. Dealer banks typically submit steeper demand curves as compared to long-term investors. Further, there is a negative relation between return volatility and demand elasticity for both investor types. In periods of high volatility, the elasticity of demand declines. However, as shown by the solid lines, the effect is stronger for dealer banks. An increase in volatility leads to a steepening of bank dealers' demand curves relative to long-term investors.

Return Volatility We estimate the linear regression model

$$\beta_{ij} = b_0 + b_1 \cdot \sigma_{x,j} + b_2 \cdot \sigma_{x,j} \times \mathbb{1}\{\text{DB}\}_i + b_3 \cdot x_j + \varepsilon_{ij} \tag{10}$$

where β_{ij} is total demand elasticity, and $\sigma_{x,j}$ is one of return volatility $\sigma_{j-21,j}$ or bid dispersion $\sigma_{b,j}$. The indicator $\mathbb{1}\{\text{DB}\}_i$ equals one if bidder i is a dealer bank. The vector of exogenous controls x_j includes the number of participants, the log issue size, maturity, and the relative bid-ask spread. We also include inflation, the short-term rate, a business cycle indicator, and the slope of the yield curve. We do not control for bidder yield discounts, which is an endogenous outcome of the auction. The coefficients of interest are b_1 and b_2, where b_2 captures the differential response of dealer banks to higher volatility or bid dispersion. Predictions 1 and 2 imply that b_1 and b_2 should both be negative. We report estimates for $\sigma_{j-21,j}$ in Table 4.

	Log demand elasticity				
	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}
$\sigma_{j-21,j}$	-1.58^{***}	-1.23^{***}	-0.48^{**}	-0.44^{**}	-0.30
	(0.15)	(0.15)	(0.19)	(0.19)	(0.19)
$\sigma_{j-21,j} \times \mathbb{1}\{\text{DB}\}_i$		-0.58^{***}	-0.56^{***}	-0.57^{***}	-0.57^{***}
		(0.13)	(0.12)	(0.12)	(0.12)
Maturity			-0.03^{***}	-0.03^{***}	-0.02^{**}
			(0.01)	(0.01)	(0.01)
Participants			0.02^{**}	0.02^{***}	0.02^{***}
			(0.01)	(0.01)	(0.01)
Log issue size				-0.08	-0.07
				(0.06)	(0.06)
RBAS_j					-0.23^{***}
					(0.08)
Constant	5.59^{***}	5.62^{***}	5.61^{***}	6.59^{***}	6.48^{***}
	(0.08)	(0.08)	(0.14)	(0.67)	(0.67)
Macro	✓	✓	✓	✓	✓
Adj. R^2	0.25	0.27	0.30	0.30	0.31
N	993	993	993	993	993

Table 4: Coefficient estimates of regression (10). The dependent variable is the total elasticity (TE) of demand (in logs) at the bidder level. $\sigma_{j-21,j}$ denotes the volatility of daily bond returns in the month prior to the auction, whereas $\mathbb{1}\{\text{DB}\}$ is an indicator equal to one if bidder j is a dealer bank. Maturity refers to time to maturity at the auction date. Participants is the number of bidders. Log issue size is the logarithm of total supply. RBAS_j is the relative bid-ask spread in the secondary market at the auction date. Macro controls include inflation, the short-term rate (SARON), the slope of the yield curve, and the KOF economic barometer. The sample is from 2000 to present and only considers security reopenings. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.

The first column shows that the elasticity of demand declines in response to an increase in volatility. Return volatility explains roughly one quarter of the variation in total demand elasticity. The second column demonstrates that the effect is stronger for dealer banks, and the coefficient estimate is economically and statistically significant. A one-percentage-point increase in daily return volatility is associated to a demand elasticity (in logs) that is 0.64 lower for dealer banks compared to long-term investors. As in Table 3, the elasticity of demand declines with maturity. Consistent with our theory, a higher number of participants leads to flatter demand curves as the auction becomes more competitive. Controlling for secondary market spreads and issue size does not affect the results. Furthermore, there is no apparent relation between the emission size and bidding behavior. Bidders do not seem to submit steeper demand schedules when supply is larger. In Table 10 in Appendix B.1, we obtain quantitatively and qualitatively similar results using the intermediate elasticity (IE) and the winning elasticity (WE). Robustness checks in Appendix B.1 confirm that our results are not driven by outlier auctions around the sudden removal of the EUR/CHF floor (Auer, Burstein, & Lein, 2021) or during the Covid period.

Cost Dispersion We repeat the same exercise with the standard deviation of bid yields $\sigma_{b,j}$, which we interpret as a proxy of holding costs dispersion σ_ε. Table 5 reports the results with bid dispersion and illustrates a similar pattern as in Table 4.

The elasticity of demand declines with bid dispersion, but the overall effect is not statistically significant. However, the effect is negative and statistically significant for dealer banks, as shown in the second column. Coefficient estimates are qualitatively similar after controlling for maturity, the number of participants, issue size, relative bid-ask spread, and return volatility. As a first robustness check, we consider three alternative measures of bid dispersion, namely the cross-sectional standard deviation of equally-weighted bid yields, the interquartile range of quantity-weighted bid yields, and the interquartile range of equally-weighted bid yields. We report coefficient estimates using these three measures in columns two to four of Table 13 in Appendix B.1. For comparison, the first column on Table 13 is identical to the fourth column of Table 5. Across all different measures and specifications, the coefficient on the interaction between bid dispersion and the dealer bank indicator is negative and statistically significant, as predicted by our theory. We also separately test whether the total effect $b_1 + b_2$ is negative and statistically significant, and we reject the null that $b_1 + b_2 \geq 0$. Overall, an increase in bid dispersion is associated with dealer banks submitting steeper demand schedules relative to all other investors. The magnitude of the coefficient is larger when we consider the interquartile range because of the different units.

	Log demand elasticity			
	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}
$\sigma_{b,j}$	-0.23	0.62	0.44	0.56
	(0.50)	(0.41)	(0.36)	(0.38)
$\sigma_{b,j} \times \mathbb{1}\{\text{DB}\}_i$		-1.96^{**}	-1.91^{**}	-1.90^{**}
		(0.93)	(0.81)	(0.87)
Maturity			-0.04^{***}	-0.02^{**}
			(0.00)	(0.01)
Participants			0.02^{**}	0.03^{***}
			(0.01)	(0.01)
Log issue size				-0.08
				(0.06)
$\sigma_{j-21,j}$				-0.65^{***}
				(0.19)
RBAS$_j$				-0.23^{***}
				(0.08)
Constant	4.94^{***}	4.96^{***}	5.61^{***}	6.53^{***}
	(0.06)	(0.07)	(0.13)	(0.68)
Macro	✓	✓	✓	✓
Adj. R^2	0.12	0.13	0.27	0.30
N	1'087	1'087	1'087	993
p-sum		0.08	0.04	0.07

Table 5: Coefficient estimates of regression (10). The dependent variable is the total elasticity (TE) of demand (in logs) at the bidder level. $\sigma_{b,j}$ is the cross-sectional standard deviation of quantity-weighted bid yields, whereas $\mathbb{1}\{\text{DB}\}_i$ is an indicator equal to one if bidder i is a dealer bank, and zero otherwise. Maturity refers to time to maturity at the auction date. Participants is the number of bidders. $\sigma_{j-21,j}$ denotes the volatility of the bond in the month prior to the auction. Log issue size is the logarithm of supply. RBAS$_j$ is the relative bid-ask spread at the auction date. Macro controls include inflation, the short-term rate (SARON), the slope of the yield curve, and the KOF economic barometer. The sample is from 2000 to present and only considers security reopenings. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively. p-sum is the p-value from testing the null hypothesis that $b_1 + b_2 \geq 0$.

A major challenge in constructing proxies for holding cost dispersion is that bidders' holding costs are unobservable. While the cross-sectional volatility in bid yields $\sigma_{b,j}$ is related to σ_ε, the dispersion in bid yields is an endogenous outcome that directly reflects bidders' strategies. In particular, the slope of their demand schedules and the number of bid steps influence this dispersion measure. As a result, the connection between the demand slope and the observed cross-sectional dispersion in bid yields is partly mechanical, and unrelated to how heterogeneity in investment horizons shapes bidding behavior. Because the alternative proxies discussed in Appendix B.1 also face the same limitation, we construct a measure of cross-sectional dispersion that does not rely on the standard deviation of bid yields. For each auction j, given price-quantity pairs $(B_{ijk}, \mathcal{Q}_{ijk})$ for bidder i, we estimate the regression

$$q_{ijk} = \delta_0 + \mathbb{1}\{\text{DB}\}_i + \delta_1 b_{ijk} + \delta_2 \mathbb{1}\{\text{DB}\}_i \times b_{ijk} + \gamma_i + \varepsilon_{ijk} \tag{11}$$

where γ_i is a set of bidder fixed effects and $\mathbb{1}\{\text{DB}\}_i$ is an indicator equal to one if bidder i is a dealer bank. Again, we add non-competitive bid quantities to $\mathcal{Q}^{nc}_{ijk}$ to each bid step. The proxy for σ_ε is the standard deviation of the regression residuals $\hat{\sigma}_{\varepsilon,j} = \sqrt{\mathbb{V}\text{ar}(\varepsilon_{ijk})}$. Table 14 in Appendix B.2 reports estimates of specification (10), but measuring holding cost dispersion using $\hat{\sigma}_{\varepsilon,j}$. Table 14 corroborates our findings that dealer banks steepen their demand curves relative to long-term investors when dispersion in holding costs is higher. The coefficient estimate on the interaction term is negative and statistically significant in all four columns. The total effect b_1+b_2 of higher bid dispersion on dealer banks is also negative and statistically significant across most specifications. In summary, Table 4, Table 5, Table 13, and Table 14 strongly suggest that bidding strategies respond to changes in aggregate uncertainty that we proxy by either bond return volatility or bid dispersion. The effect is stronger for dealer banks, which are significantly more sensitive to uncertainty relative to long-term investors.

4.3 Holding Costs and Bidding Behavior

We next investigate how changes in the regulatory environment influence bidding behavior in the primary market. Basel III regulations impose higher capital requirements and introduce additional buffers for large systemic banks, especially those engaged in market-making activities (Duffie, 2016). These regulatory measures directly impact the dealer banks in our sample. As a result, we hypothesize that dealer banks will submit steeper demand curves than other investors, such as pension funds, insurance companies, and regional banks, who are not subject to comparable regulatory constraints. To this purpose, we implement a difference in differences (DiD) design around the gradual implementation of the Basel III regulatory framework, which we interpret as an increase in the slope of marginal holding costs κ of dealers only. Although the regulation was first endorsed in 2013, the capital requirements have been gradually phased in starting in January 2015. Therefore, we implement a difference in differences (DiD) design around January 1, 2015. To increase statistical power we consider a symmetric window of five years around January 1, 2015. In robustness checks reported in Appendix B.2, we perform the same analysis using a three year and a four year window, and estimates are qualitatively and quantitatively similar. The pre-implementation period is from January 2010 to January 2015. The post-implementation is from January 2015 to January 2020. The DiD design is

$$z_{ij} = b_0 + b_1 \cdot \mathbb{1}\{\text{Basel III}\}_j + b_2 \cdot \mathbb{1}\{\text{DB}\}_i + b_3 \cdot \mathbb{1}\{\text{Basel III}\}_j \cdot \mathbb{1}\{\text{DB}\}_i + b_4 \cdot x_j + \varepsilon_{ij} \quad (12)$$

The indicator Basel III_j equals one if auction j occurred post 2015 when Basel III requirements were being phased in. Similarly, $\mathbb{1}\{\text{DB}\}_i$ is an indicator variable equal to one if bidder i is a dealer bank. We expect a negative coefficient, that is liquidity and capital requirements make systemic banks less elastic relative to other bidders who are not subject to Basel III. We consider two outcome variables z_{ij}, that is the total elasticity of demand (TE β_{ij}) and the yield discount (discount_{ij}). The vector of controls x_j includes bond maturity, the number of participants, and the issue size. Through the lens of Proposition 2, we expect an increase in dealer banks' holding cost to cause a steepening of their demand curves relative to long-term investors that are not subject to the regulation. Accordingly, when the dependent variable is the elasticity of demand, we expect the coefficient b_3 to be negative.

Table 6 reports coefficient estimates using the log total demand elasticity and the yield discount as the dependent variable. As in Figure 7, the first and the second columns show that dealer banks submit, on average, steeper demand schedules relative to other auction participants. The first column shows that, in the post-Basel III period, the gap between dealer banks' and long-term investors' demand elasticity has grown. The coefficient estimate of -0.57 is highly statistically and economically significant. The introduction of the Basel III framework is associated to a reduction of 0.57 (in log units) in dealers' elasticity of demand relative to other investors. The magnitude and the statistical significance remain stable after controlling for maturity, number of participants, and issue size. Robustness checks in Appendix B.2 exclude auctions around the sudden EUR/CHF floor removal on January 15, 2015 and recover similar estimates.

	Log elasticity		Yield discount	
	TE β_{ij}	TE β_{ij}	Discount$_{ij}$	Discount$_{ij}$
$\mathbb{1}\{\text{Basel III}\}_j$	-0.11	-0.10	0.02	0.02^{**}
	(0.14)	(0.13)	(0.01)	(0.01)
$\mathbb{1}\{\text{DB}\}_i$	-0.21^{**}	-0.27^{***}	0.01	0.01
	(0.10)	(0.09)	(0.01)	(0.01)
$\mathbb{1}\{\text{Basel III}\}_j \times \mathbb{1}\{\text{DB}\}_i$	-0.57^{***}	-0.40^{**}	-0.03^{**}	-0.03^{**}
	(0.19)	(0.16)	(0.01)	(0.02)
Maturity		-0.05^{***}		-0.00
		(0.00)		(0.00)
Participants		-0.02		-0.00
		(0.02)		(0.00)
Log issue size		-0.00^{**}		0.00
		(0.00)		(0.00)
Constant	5.35^{***}	6.59^{***}	0.01	0.02
	(0.08)	(0.26)	(0.01)	(0.04)
Adj. R^2	0.10	0.31	0.00	0.00
N	562	562	1'375	1'375

Table 6: Coefficient estimates of the difference in differences specification (12). In the first and second column, the dependent variable is the total elasticity of demand (in logs). In the third and in the fourth column, the dependent variable is the quantity-weighted yield discount. $\mathbb{1}\{\text{Basel III}\}_j$ is an indicator equal to one if the auction occurs after January 2015. $\mathbb{1}\{\text{DB}\}_i$ is an indicator equal to one if bidder i is a dealer bank. The sample period is from January 2010 to December 2019 and spans a five-year window around the introduction of the Basel III regulations. Robust standard errors are reported in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.

The third and the fourth columns reveal that large banks have been offering lower yields (higher prices) than in the secondary market after the introduction of capital requirements. The fourth column implies that dealer banks bid at significantly lower discounts of 3 basis points relative to the secondary market. However, the DiD specification explains very little of the variation in bidder level yield discounts. Taken together, our results indicate that regulatory requirement cause dealer banks' demand schedules to become steeper.

4.4 Relation to Secondary Market Outcomes

In this final section, we further examine how bidding strategies influence outcomes in the secondary market and, vice versa, how secondary market conditions influence bidding behavior. To this purpose, we first study how demand elasticities respond to secondary market liquidity. Second, we explore the connection between bidding behavior and post-auction return predictability.

Market Liquidity Our assumptions of investment horizons and resale imply that dealer banks with a shorter expected holding period behave more cautiously when the secondary market is expected to be illiquid. In contrast, long-term investors with longer holding periods are less affected by liquidity conditions (Amihud & Mendelson, 1986). To test our conjecture that dealer banks' demand schedules become less elastic in response to higher illiquidity, we estimate the linear regression model

$$\beta_{ij} = b_0 + b_1 \cdot \mathrm{RBAS}_j + b_2 \cdot \mathrm{RBAS}_j \times \mathbb{1}\{\mathrm{DB}\}_i + b_3 \cdot x_j + \varepsilon_{ij} \tag{13}$$

We measure liquidity in the secondary market using the relative bid-ask spread RBAS_j at the auction close date. Due to the strong autocorrelation in bid-ask spreads, the current bid-ask spread well captures expectations about future liquidity conditions in the secondary market.[14]

	Log demand elasticity				
	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}
RBAS_j	-0.63^{***}	-0.43^{***}	-0.10	-0.08	-0.00
	(0.05)	(0.06)	(0.09)	(0.09)	(0.09)
$\mathrm{RBAS}_j \times \mathbb{1}\{\mathrm{DB}\}_i$		-0.31^{***}	-0.31^{***}	-0.32^{***}	-0.33^{***}
		(0.06)	(0.06)	(0.06)	(0.06)
Maturity			-0.03^{***}	-0.03^{***}	-0.02^{***}
			(0.01)	(0.01)	(0.01)
Participants			0.02^{***}	0.03^{***}	0.03^{***}
			(0.01)	(0.01)	(0.01)
Log issue size				-0.11^{**}	-0.08
				(0.05)	(0.06)
$\sigma_{j-21,j}$					-0.69^{***}
					(0.19)
Constant	5.63^{***}	5.65^{***}	5.56^{***}	6.80^{***}	6.50^{***}
	(0.07)	(0.08)	(0.13)	(0.63)	(0.67)
Macro	✓	✓	✓	✓	✓
Adj. R^2	0.25	0.27	0.29	0.30	0.31
N	1'087	1'087	1'087	1'087	993

Table 7: Coefficient estimates of regression (13). The dependent variable is the total elasticity (TE) of demand (in logs) at the bidder level. RBAS_j is the relative bid-ask spread at the auction date, whereas $\mathbb{1}\{\mathrm{DB}\}_i$ is an indicator equal to one if bidder i is a dealer bank, and zero otherwise. Maturity is the time to maturity at the auction date. Participants is the number of bidders. $\sigma_{j-21,j}$ denotes the volatility of the bond in the month prior to the auction. Log issue size is the logarithm of supply. Macro controls include inflation, the short-term rate (SARON), the slope of the yield curve, and the KOF economic barometer. The sample is from 2000 to present and only considers security reopenings. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.

Consistent with Prediction 4, the last three columns in Table 7 reveal that a higher relative bid-ask spread leads to a steepening of dealer banks' demand curves only, whereas there is no

[14] Using moving averages or fitted values produces similar results.

effect on the demand elasticity of long-term investors. The estimates on the interaction term are economically and statistically significant across all specifications. The significance of the unconditional effect b_1 vanishes once we control for bond maturity and the number of participants. Both the sign and the significance of these estimates are consistent with our interpretation on heterogeneity in investment horizons. Furthermore, the effect remains highly significant after controlling for return volatility, suggesting that market liquidity impacts bidding behavior over and above aggregate risk. Long-term investors with a longer holding period are not affected by secondary market liquidity as they do not plan to resell the security right after the auction. On the other hand, dealer banks are more exposed to changes in liquidity conditions around the auction through their market making activities. We obtain the same results when we consider the average bid-ask spread in the month prior to the auction.

Return Predictability Lastly, we validate our mechanism by studying the relation between bidding behavior and post-auction return predictability. We test Prediction 3 by estimating predictive regressions of post-auction excess return on demand elasticities. For each auction j, we construct one-day, two-day, one-week and one-month post-auction excess bond returns $rx_{j,j+h}$, for $h \in \{1, 2, 5, 21\}$. We compute returns relative to market clearing price at the auction given that this is the price at which auction participants purchased the securities. We convert returns into excess returns by subtracting the horizon-matched risk free rate. For $h \in \{1, 2\}$ we use the SARON, whereas for $h = 5$ and $h = 21$ we use the SAR1W and SAR1M rates, respectively. Accordingly, we project bond excess return on the average demand elasticity $\bar{\beta}_j = \frac{1}{I}\sum_{i=1}^{I} \beta_{ij}$

$$rx_{j,j+h} = b_0 + b_1 \cdot \bar{\beta}_j + b_2 \cdot x_j + \varepsilon_j \tag{14}$$

where x_j includes bond maturity, the number of participants, the issue size, and the relative bid-ask spread. The intuition is that dealer banks require a risk premium to hold the securities between the auction and the secondary market because the post-auction capital gain is stochastic. In our theory, this risk premium increases with demand uncertainty and the aggregate quantity purchased by these agents at the auction. Because demand elasticities decline with aggregate uncertainty, a less elastic average demand positively predicts post-auction returns. We further separate dealer banks and long-term investors to assess whether the horizon of return predictability depends on which investor bids less elastic demand schedules. Intuitively, the demand risk premium vanishes as soon as dealer banks are able to distribute their inventories in the secondary market. This becomes less likely if long-term investors also have inelastic demand, and the resale price can potentially be lower. Consequently, a more rigid demand solely from dealer banks should predict returns only at a short horizon. In contrast, when long-term investors are also less elastic, it will likely take longer for dealer banks to unwind their inventories. Because of this, dealer banks will bear inventory risk for a longer period, leading to return predictability at longer horizons. To this purpose, we estimate

$$rx_{j,j+h} = b_0 + b_1 \cdot \bar{\beta}_j^{\mathrm{LT}} + b_2 \cdot \bar{\beta}_j^{\mathrm{DB}} + b_3 \cdot x_j + \varepsilon_j \tag{15}$$

The top panel of Table 8 reports coefficient estimates of the predictive regression (14). We find short-term return predictability until one week after the auction. For very short horizons, the regression R^2 statistics are relatively sizable. The first column documents that the average demand elasticity explains up to 12% of the one-day post-auction excess bond returns. However, although coefficients estimates have the expected sign and are statistically significant, their magnitudes do not follow any clear pattern.

The bottom panel of Table 8 reveals a different picture once we distinguish between long-term investors and dealer banks. A decline in the demand elasticity of dealer banks positively predicts post-auction excess bond returns only up to two days after the auction. However, estimates for horizons longer than one week are statistically indistinguishable from zero. In contrast, the elasticity of demand of long-term investors predicts post-auction bond returns up to one month after the auction. Further, the coefficient on $\bar{\beta}_j^{\mathrm{LT}}$ increases monotonically with the horizon.

	$rx_{j,j+1}$	$rx_{j,j+2}$	$rx_{j,j+5}$	$rx_{j,j+21}$
Aggregate elasticity				
$\bar{\beta}_j$	-0.20^{***}	-0.17^{**}	-0.28^{**}	-0.28
	(0.07)	(0.08)	(0.14)	(0.28)
Constant	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Adj. R^2	0.12	0.03	0.00	0.04
N	342	342	338	314
Dealer banks vs. long-term investors				
$\bar{\beta}_j^{\mathrm{LT}}$	-0.19^{**}	-0.27^{***}	-0.41^{***}	-0.70^{***}
	(0.09)	(0.09)	(0.11)	(0.27)
$\bar{\beta}_j^{\mathrm{DB}}$	-0.22^{**}	-0.14	-0.09	0.35
	(0.09)	(0.11)	(0.14)	(0.28)
Constant	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Adj. R^2	0.14	0.06	0.03	0.05
N	226	226	223	205

Table 8: Coefficient estimates of regressions (14) and (15). The top panel regresses h-day ahead excess returns onto the average total elasticity (TE) across all participants in the auction, which is computed as $\bar{\beta}_j = \frac{1}{I}\sum_{i=1}^{I}\beta_{ij}$. The bottom panel regresses h-day ahead excess returns onto the average demand elasticity of systemic banks $\bar{\beta}_j^{\mathrm{DB}}$ and of all the other bidders $\bar{\beta}_j^{\mathrm{LT}}$ separately. Excess returns are computed based on the auction price. Controls include maturity, number of bidders, the relative bid-ask spread, and the log issue size. The sample is from 2000 to present. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively

5 Conclusion

This paper demonstrates, both theoretically and empirically, that heterogeneity in investment horizons plays a critical role in determining bidding behavior and post-auction price dynamics of safe assets. Using a novel dataset with detailed bidder identities, we show that short-term oriented dealers are more sensitive to demand uncertainty, resulting in steeper demand curves and asset return predictability lasting only a few days. In contrast, long-term investors' inelastic demand extends return predictability to longer horizons, up to one month. These findings emphasize how changes in investor base can impact the risk-return profile of safe assets and their post-auction dynamics.

Our paper provides at least three important implications for both policy makers and academics. First, our theoretical findings offer new insights into how investors price safe assets and how differences in investment horizons influence bidding behavior in the primary market. Second, we propose a novel approach to Treasury auction design, emphasizing the critical role of bidder composition. Our findings reveal that both *how* an asset is sold and *to whom* it is sold significantly impact its return distribution. Unlike much of the existing literature, which focuses on auction rules and post-auction disclosure of results, we highlight the heterogeneity among auction participants and their investment horizons (dealers versus long-term investors). Our results hint at a trade-off between secondary market liquidity (supported by dealers) and post-auction price volatility (driven by their short-term trading horizons and inventory offloading). This trade-off allows us to derive the costs and benefits of a primary dealer system. On the cost side, primary dealership induces a risk premium that the government must bear. On the benefit side, a well-established primary dealer system ensures a more liquid secondary market. Our policy recommendation to Debt Management Offices is to integrate secondary market dynamics into the auction design process, rather than treating auctions in isolation. Lastly, our paper revisits the concept of a safe asset by linking it to demand risk and investment horizons, extending the traditional view that focuses on fundamental risk or information asymmetry (Dang et al., 2017). This broader perspective highlights the role of investor heterogeneity as a key determinant of asset prices, showing that market structure can influence prices and, ultimately, the final allocation, even in settings without fundamental risk or information asymmetry.

References

Admati, A. R. (1985). A noisy rational expectations equilibrium for multi-asset securities markets. *Econometrica*, *53*(3), 629–657.

Albuquerque, R., Cardoso-Costa, J. M., & Faias, J. A. (2024). Price elasticity of demand and risk-bearing capacity in sovereign bond auctions. *The Review of Financial Studies*, *37*(10), 3149–3187.

Allen, J., Hortacsu, A., Richert, E., & Wittwer, M. (2024). Entry and exit in treasury auctions. *Working paper*, 1–77.

Allen, J., Kastl, J., & Wittwer, M. (2024). Estimating demand systems with bidding data. *Working paper*, 1–61.

Allen, J., & Wittwer, M. (2023). Intermediary market power and capital constraints. *Working paper*, 1–62.

Amihud, Y., & Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics*, *17*(2), 223–249.

Armantier, O., & SbaÏ, E. (2006). Estimation and comparison of treasury auction formats when bidders are asymmetric. *Journal of Applied Econometrics*, *21*(6), 745-779.

Auer, R., Burstein, A., & Lein, S. M. (2021). Exchange rates and prices: Evidence from the 2015 swiss franc appreciation. *American Economic Review*, *111*(2), 652–686.

Back, K., & Zender, J. F. (2001). Auctions of divisible goods with endogenous supply. *Economics Letters*, *73*(1), 29-34.

Bank for International Settlements. (2017). Structural changes in banking after the crisis. *Committee on the Global Financial System (CGFS) Papers*, *No. 60*. (Available at https://www.bis.org/publ/cgfs60.pdf)

Barro, R. J., Fernández-Villaverde, J., Levintal, O., & Mollerus, A. (2022). Safe Assets. *The Economic Journal*, *132*(646), 2075-2100.

Beetsma, R., Giuliodori, M., Hanson, J., & de Jong, F. (2020). Determinants of the bid-to-cover ratio in eurozone sovereign debt auctions. *Journal of Empirical Finance*, *58*, 96-120.

Benigno, P., & Nisticò, S. (2017). Safe assets, liquidity, and monetary policy. *American Economic Journal: Macroeconomics*, *9*(2), 182–227.

Biais, B., Martimort, D., & Rochet, J.-C. (2000). Competing mechanisms in a common value environment. *Econometrica*, *68*(4), 799–837.

Boyarchenko, N., Lucca, D. O., & Veldkamp, L. (2021). Taking orders and taking notes: Dealer information sharing in treasury auctions. *Journal of Political Economy*, *129*(2), 607–645.

Caballero, R. J., & Farhi, E. (2017). The safety trap. *Review of Economic Studies*, *85*(1), 223–274.

Caballero, R. J., Farhi, E., & Gourinchas, P.-O. (2016). Safe asset scarcity and aggregate demand. *American Economic Review*, *106*(5), 513–518.

Caballero, R. J., Farhi, E., & Gourinchas, P.-O. (2017). The safe assets shortage conundrum. *Journal of Economic Perspectives*, *31*(3), 29–46.

Caballero, R. J., & Krishnamurthy, A. (2009). Global imbalances and financial fragility. *American Economic Review Papers and Proceedings*, *99*(2), 584–588.

Cammack, E. B. (1991). Evidence on bidding strategies and the information in treasury bill auctions. *Journal of Political Economy*, *99*(1), 100–130.

Chen, D., & Duffie, D. (2021). Market fragmentation. *American Economic Review*, *111*(7), 2247-2274.

Chodorow-Reich, G., Ghent, A., & Haddad, V. (2020). Asset insulators. *The Review of Financial Studies*, *34*(3), 1509–1539.

Dang, T. V., Gorton, G., Holmström, B., & Ordoñez, G. (2017). Banks as secret keepers. *American Economic Review*, *107*(4), 1005–1029.

Danthine, J.-P., & Moresi, S. (1993). Volatility, information and noise trading. *European Economic Review*, *37*(5), 961–982.

Diamond, D. W., & Verrecchia, R. E. (1991). Disclosure, liquidity, and the cost of capital. *The Journal of Finance*, *46*(4), 1325–1359.

Du, S., & Zhu, H. (2017). What is the Optimal Trading Frequency in Financial Markets? *The Review of Economic Studies*, *84*(4), 1606–1651.

Duffie, D. (2016). Financial regulatory reform after the crisis: An assessment. *ECB Forum on Central Banking*.

Duffie, D., Malamud, S., & Manso, G. (2009). Information percolation with equilibrium search dynamics. *Econometrica*, *77*(5), 1513-1574.

Dworczak, P. (2020). Mechanism design with aftermarkets: Cutoff mechanisms. *Econometrica*, *88*(6), 2629–2661.

Fleming, M., Nguyen, G., & Rosenberg, J. (2024). How do treasury dealers manage their positions? *Journal of Financial Economics*, *158*, 103885.

Friedman, M. (1991). How to sell government securities. *Wall Street Journal*.

Février, P., Préget, R., & Visser, M. (2002). Econometrics of share auctions. *Working paper*, 1–37.

Gabaix, X., & Koijen, R. (2023). In search of the origins of financial fluctuations: The inelastic markets hypothesis. *Working paper*, 1–109.

Goldreich, D. (2007). Underpricing in discriminatory and uniform-price treasury auctions. *Journal of Financial and Quantitative Analysis*, *42*(2), 443–466.

Gorton, G. (2017). The history and economics of safe assets. *Annual Review of Economics*, *9*, 547–586.

Gorton, G., & Ordoñez, G. (2022). The supply and demand for safe assets. *Journal of Monetary Economics*, *125*, 132–147.

Greenwood, R., & Vayanos, D. (2014). Bond supply and excess bond returns. *Review of Financial Studies*, *27*(3), 663–713.

Haile, P. A. (2001). Auctions with resale markets: An application to u.s. forest service timber sales. *American Economic Review*, *91*(3), 399–427.

Haile, P. A. (2003). Auctions with private uncertainty and resale opportunities. *Journal of Economic Theory*, *108*(1), 72–110.

Hamao, Y., & Jegadeesh, N. (1998). An analysis of bidding in the japanese government bond auctions. *The Journal of Finance*, *53*(2), 755–772.

He, Z., Krishnamurthy, A., & Milbradt, K. (2019). A model of safe asset determination. *American Economic Review*, *109*(4), 1230–1262.

Hortaçsu, A., & 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.

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

Hortaçsu, A., Kastl, J., & Zhang, A. (2018). Bid shading and bidder surplus in the us treasury auction system. *American Economic Review*, *108*(1), 147–169.

Infante, S. (2020). Private money creation with safe assets and term premia. *Journal of Financial Economics*, *136*(3), 828–856.

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

Kastl, J. (2011). Discrete Bids and Empirical Inference in Divisible Good Auctions. *The Review of Economic Studies*, *78*(3), 974–1014.

Keloharju, M., Nyborg, K. G., & Rydqvist, K. (2005). Strategic behavior and underpricing in uniform price auctions: Evidence from finnish treasury auctions. *The Journal of Finance*, *60*(4), 1865–1902.

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

Krishnamurthy, A. (2002). The bond/old-bond spread. *Journal of Financial Economics*, *66*(2-3), 463–506.

Krishnamurthy, A., & Vissing-Jorgensen, A. (2012a). The aggregate demand for treasury debt. *Journal of Political Economy*, *120*(2), 233–267.

Krishnamurthy, A., & Vissing-Jorgensen, A. (2012b). The aggregate demand for treasury debt. *Journal of Political Economy*, *120*(2), 233–267.

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

Kyle, A. S., Obizhaeva, A. A., & Wang, Y. (2017). Smooth Trading with Overconfidence and Market Power. *The Review of Economic Studies*, *85*(1), 611–662.

Lou, D., Yan, H., & Zhang, J. (2013). Anticipated and Repeated Shocks in Liquid Markets. *The Review of Financial Studies*, *26*(8), 1891–1912.

Malamud, S., & Rostek, M. (2017). Decentralized exchange. *American Economic Review*, *107*(11), 3320–3362.

Moreira, A., & Savov, A. (2017). The macroeconomics of shadow banking. *Journal of Finance*, *72*(6), 2381–2432.

Musto, D., Nini, G., & Schwarz, K. (2018). Notes on bonds: Illiquidity feedback during the financial crisis. *The Review of Financial Studies*, *31*, 2983–3018.

Nagel, S. (2016). The Liquidity Premium of Near-Money Assets. *The Quarterly Journal of Economics*, *131*(4), 1927–1971.

Nyborg, K. G., Rydqvist, K., & Sundaresan, S. M. (2002). Bidder behavior in multiunit auctions: Evidence from swedish treasury auctions. *Journal of Political Economy*, *110*(2), 394–424.

OECD. (2016). *OECD Sovereign Borrowing Out,look 2016.* OECD Publishing, Paris. (Available at https:/https://doi.org/10.1787/sov_b_outlk-2016-en)

OECD. (2019). *Oecd sovereign borrowing outlook 2019.*

Payne, J. P., & Szöke, B. (2024). Convenience yields and financial repression. *Working Paper*.
Ranaldo, A., & Rossi, E. (2016). Uniform-price auctions for swiss government bonds: Origin and evolution. *SNB Economic Studies*, *10*, 1–46.
Richert, E. (2024). Indirect inference techniques as an equilibrium computation method. *Working paper*, 1–55.
Rostek, M., & Weretka, M. (2012). Price inference in small markets. *Econometrica*, *80*(2), 687–711.
Rostek, M., & Yoon, J. H. (2021). Exchange design and efficiency. *Econometrica*, *89*(6), 2887–2928.
Stein, J. C. (2012). Monetary policy as financial stability regulation. *Quarterly Journal of Economics*, *127*(1), 57–95.
Sunderam, A. (2015). Money creation and the shadow banking system. *Review of Financial Studies*, *28*(4), 939–977.
Umlauf, S. R. (1993). An empirical study of the mexican treasury bill auction. *Journal of Financial Economics*, *33*(3), 313–340.
Vayanos, D. (1999). Strategic Trading and Welfare in a Dynamic Market. *The Review of Economic Studies*, *66*(2), 219–254.
Vayanos, D., & Vila, J.-L. (2021). A preferred-habitat model of the term structure of interest rates. *Econometrica*, *89*(1), 77–112.
Vayanos, D., & Wang, J. (2011). Liquidity and Asset Returns Under Asymmetric Information and Imperfect Competition. *The Review of Financial Studies*, *25*(5), 1339–1365.
Vives, X. (2008). *Information and learning in markets: The impact of market microstructure.* Princeton, NJ: Princeton Univ. Press.
Vives, X. (2011). Strategic supply function competition with private information. *Econometrica*, *79*(6), 1919–1966.
Wilson, R. (1977). A bidding model of perfect competition. *The Review of Economic Studies*, *44*(3), 511–518.

Appendices

A Mathematical Results

A.1 Proofs

A.1.1 Proof of Proposition 1

Proof. We construct a linear symmetric equilibrium by conjecturing that demand is

$$q_{i1} = a - bp_1 - c\lambda_i$$

Market clearing requires

$$\int_i (a - bp_1 - c\lambda_i)di = a - bp_1 - c\lambda = Q_a$$

so that the price is informationally equivalent to λ. As a result, $\mathbb{E}[\lambda|\lambda_i, p] = \lambda$, and the secondary market equilibrium is fully revealing. Agent i's objective is

$$-\max_{q_{i1}} \mathbb{E}_1\left[\, \exp(-\gamma W_{i2})\right]$$

subject to the budget constraint

$$W_{i2} = W_{i1} + (1-p_1)q_{i1} - \lambda_i q_{i1} - \frac{\kappa}{2}q_{i1}^2$$

Given $\mathbb{E}[\lambda|\lambda_i, p] = \lambda$, all the uncertainty is resolved, and the problem is rewritten as

$$\max_{q_{i1}} W_{i1} + (1-p_1)q_{i1} - \lambda_i q_{i1} - \frac{\kappa}{2}q_{i1}^2$$

The first-order condition is

$$(1-p_1) - \lambda_i - \kappa q_{i1} = 0 \Longrightarrow q_{i1} = a - bp_1 - c\lambda_i = \frac{1-p_1-\lambda_i}{\kappa}$$

Matching coefficients gives

$$c = \frac{1}{\kappa} \quad : \quad a = \frac{1}{\kappa} \quad : \quad b = \frac{1}{\kappa}$$

which verifies the initial conjecture. Market clearing requires

$$\frac{1-p_1-\int_i \lambda_i di}{\kappa} = \frac{1-p_1-\lambda}{\kappa} = Q_a \Longrightarrow p_1 = 1 - \lambda - \kappa Q_a$$

Plugging p_1^* into the demand function gives the desired result and completes the proof.

■

A.1.2 Proof of Lemma 1

Proof. From Proposition 1, the equilibrium in the secondary market is $p_1^* = 1 - \lambda - \kappa Q_a$ and $q_{1i}^* = \frac{\lambda}{\kappa} - \frac{\lambda_i}{\kappa} + Q_a$. Plugging these expressions into the dealers' budget constraint gives

$$\begin{aligned}
W_{j2} &= (p_1^* - p)q_j - \lambda_j q_j - \frac{\kappa}{2}q_j^2 + (1-p_1^*)q_{j1}^* - \lambda_j q_{j1}^* - \frac{\kappa}{2}(q_{j1}^*)^2 \\
&= (1 - \lambda - \kappa Q_a - p)q_j - \lambda_j q_j - \frac{\kappa}{2}q_j^2 + \frac{\lambda^2}{2\kappa} + \frac{\lambda_j^2}{2\kappa} - \frac{\lambda\lambda_j}{\kappa} + Q_a(\lambda - \lambda_j) + \frac{\kappa}{2}Q_a^2 \\
&= (1 - \kappa Q_a - p)q_j - \lambda_j q_j - \frac{\kappa}{2}q_j^2 + \frac{\lambda_j^2}{2\kappa} - Q_a\lambda_j + \lambda\left(Q_a - q_j - \frac{\lambda_j}{\kappa}\right) + \frac{\lambda^2}{2\kappa} + \frac{\kappa}{2}Q_a^2
\end{aligned}$$

Let

$$\begin{aligned}
\pi_0(q_j) &\doteq (1 - \kappa Q_a - \lambda_j)q_j - \frac{\kappa}{2}q_j^2 \\
\pi_1(q_j) &\doteq \left(Q_a - q_j - \lambda_j \kappa^{-1}\right)
\end{aligned}$$

so that

$$W_{j2} = \frac{\lambda_j^2}{2\kappa} - Q_a\lambda_j + \frac{\kappa}{2}Q_a^2 + \pi_0(q_j) + \pi_1(q_j)\lambda + \frac{\lambda^2}{2\kappa} - pq_j$$

is a quadratic function of λ. As a result, the objective is rewritten as

$$-\mathbb{E}_0^D\left[-e^{-\gamma W_{i2}}|p,\lambda_j\right] = -\mathbb{E}_0^D\left[-e^{-\gamma\left(\pi_0(q_j)+\pi_1(q_j)\lambda+\frac{\lambda^2}{2\kappa}-pq_j\right)}\middle|p,\lambda_j\right]e^{-\gamma\left(\frac{\lambda_j^2}{2\kappa}-Q_a\lambda_j+\frac{\kappa}{2}Q_a^2\right)}$$

Further, let $\mu_\lambda = \mathbb{E}_0^D[\lambda|\lambda_j,p]$ and $\Sigma_\lambda = \mathbb{V}\mathrm{ar}_0^D[\lambda|\lambda_j,p]$ denote the posterior mean and variance of λ from the perspective of a dealer. It follows that the random variable λ is rewritten as $\lambda = \mu_\lambda + \eta$, where $\eta \sim \mathcal{N}(0,\Sigma_\lambda)$. Then, terminal wealth is written as

$$W_{j2} = \frac{\lambda_j^2}{2\kappa} - Q_a\lambda_j + \frac{\kappa}{2}Q_a^2 + \frac{\mu_\lambda^2}{2\kappa} + \pi_0(q_j) + \mu_\lambda\pi_1(q_j) + \eta\left(\pi_1(q_j)+\mu_\lambda\kappa^{-1}\right) + \eta^2\frac{1}{2\kappa} - pq_j$$

where the elements $\frac{\lambda_j^2}{2\kappa} - Q_a\lambda_j + \frac{\kappa}{2}Q_a^2 + \frac{\mu_\lambda^2}{2\kappa}$ do not depend on q_j and only enter as multiplicative constants. The objective is an expectation of a quadratic form of normal variables, so that

$$\mathbb{E}_0^D\left[-e^{-\gamma W_{i2}}|p,\lambda_j\right] = \frac{1}{\sqrt{\det\Sigma_\lambda\det\left(\Sigma_\lambda^{-1}+\gamma\kappa^{-1}\right)}}e^{-\gamma\pi_0(q_j)-\gamma pq_j+\frac{\gamma^2}{2}\frac{\left(\pi_1(q_j)+\mu_\lambda\kappa^{-1}\right)^2}{\Sigma_\lambda^{-1}+\gamma\kappa^{-1}}}e^{-\gamma\left(\frac{\lambda_j^2}{2\kappa}-Q_a\lambda_j+\frac{\kappa}{2}Q_a^2+\frac{\mu_\lambda^2}{2\kappa}\right)}$$

It follows that the dealers' problem is equivalent to

$$\max_{q_j}\pi_0(q_j) + \mu_\lambda\pi_1(q_j) - \frac{\hat{\gamma}}{2}\left(\pi_1(q_j)+\mu_\lambda\kappa^{-1}\right)^2 - pq_j$$

where we define dealers' effective risk aversion as

$$\hat{\gamma} \doteq \frac{\gamma}{\Sigma_\lambda^{-1}+\gamma\kappa^{-1}}$$

In the final step of the proof, we characterize the posterior distribution of λ conditional on p and λ_j. The price p is informationally equivalent to the total signal $h_{j,D}$. Hence, the joint distribution of the signals and the parameter λ is normal and given by

$$\theta_j \doteq \begin{pmatrix}\lambda\\ \lambda_j\\ h_{j,D}\end{pmatrix} \sim \mathcal{N}\left(\mathbb{E}_0^D[\theta_j], \mathbb{V}\mathrm{ar}_0^D(\theta_j)\right)$$

The mean and the variance-covariance matrix are given by

$$\mathbb{E}_0^D[\theta_j] = \begin{pmatrix}\bar{\lambda}\\ \bar{\lambda}\\ \bar{\lambda}\bar{c}\end{pmatrix} \quad : \quad \mathbb{V}\mathrm{ar}_0^D(\theta_j) = \begin{pmatrix}\sigma_\lambda^2 & \sigma_\lambda^2 & \bar{c}\sigma_\lambda^2\\ \sigma_\lambda^2 & \sigma_\lambda^2+\sigma_\varepsilon^2 & \bar{c}\sigma_\lambda^2\\ \bar{c}\sigma_\lambda^2 & \bar{c}\sigma_\lambda^2 & \bar{c}^2\sigma_\lambda^2+\sigma_\varepsilon^2[c_D^2(n-1)+c_M^2 m]\end{pmatrix}$$

where we define

$$\bar{c} \doteq [(n-1)c_D + mc_L]$$

The projection theorem for normal distributions implies

$$\mu_\lambda \doteq \mathbb{E}_0^D[\lambda | h_{j,D}, \lambda_j] = \mu_\lambda + \Sigma_{\lambda s}\Sigma_{ss}^{-1}(s - \mu_s)$$

where s is short-hand notation for the vector of signals λ_j and $h_{j,D}$. Some algebra yields

$$\mu_\lambda = \bar{\lambda} + \frac{1}{\bar{c}^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)} \begin{pmatrix} \sigma_\lambda^2 \left[c_D^2(n-1) + c_L^2 m\right] & \bar{c}\sigma_\lambda^2 \end{pmatrix} \begin{pmatrix} \lambda_j - \bar{\lambda} \\ h_{j,D} - \bar{\lambda}\bar{c} \end{pmatrix}$$

Using the definition of the total signal $h_{j,D} = nb_D - na_D p + mb_L - ma_L p - c_D\lambda_j - Q_a$ gives

$$\begin{aligned}
\mu_\lambda = {} & \bar{\lambda}\frac{\sigma_\varepsilon^2[c_D^2(n-1) + c_L^2 m]}{\bar{c}^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)} + \lambda_j \frac{m\sigma_\lambda^2 c_L(c_L - c_D)}{\bar{c}^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)} \\
& - \frac{\bar{c}\sigma_\lambda^2[na_D + ma_L]}{\bar{c}^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)} p + \frac{\bar{c}\sigma_\lambda^2(nb_D + mb_L - Q_a)}{\bar{c}^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)} \\
= {} & \bar{\lambda}\bar{\mu}_\lambda + \lambda_j \mu_\lambda^\lambda + \mu_\lambda^p p
\end{aligned}$$

where

$$\bar{\mu}_\lambda \doteq \frac{\bar{\lambda}\sigma_\varepsilon^2[c_D^2(n-1) + c_L^2 m]}{\bar{c}^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)} + \frac{\bar{c}\sigma_\lambda^2(nb_D + mb_L - Q_a)}{\bar{c}^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)}$$

and μ_λ^λ and μ_λ^p are the posterior loadings on the private signal and the price.
Finally, the posterior variance is

$$\Sigma_\lambda \doteq \mathbb{V}\mathrm{ar}_0^D\left(\lambda | \lambda_j, p\right) = \Sigma_{\lambda,\lambda} - \Sigma_{\lambda,s}\Sigma_{s,s}^{-1}\Sigma_{s\lambda} = \frac{\left[c_D^2(n-1) + c_L^2 m\right]\sigma_\varepsilon^2\sigma_\lambda^2}{\bar{c}^2\sigma_\lambda^2 + \left[c_D^2(n-1) + c_L^2 m\right](\sigma_\varepsilon^2 + \sigma_\lambda^2)}$$

It immediately follows that, as the number of bidders grows large,

$$\lim_{n\to\infty} \Sigma_\lambda = 0 \quad : \quad \lim_{m\to\infty} \Sigma_\lambda = 0$$

which is the desired result and completes the proof.

In the special case that $m = 0$, that is the pure common value, μ_λ and Σ_λ simplify to

$$\mu_\lambda = \bar{\lambda}\frac{\sigma_\varepsilon^2}{n\sigma_\lambda^2 + \sigma_\varepsilon^2} - \frac{\sigma_\lambda^2 na_D}{c_D[n\sigma_\lambda^2 + \sigma_\varepsilon^2]}p + \frac{\sigma_\lambda^2(nb_D - Q_a)}{c_D[n\sigma_\lambda^2 + \sigma_\varepsilon^2]} = \bar{\lambda}\tau_\varepsilon - \frac{na_D}{c_D}\tau_\lambda p + \tau_\lambda\frac{(nb_D - Q_A)}{c_D}$$

and

$$\Sigma_\lambda \doteq \mathbb{V}\mathrm{ar}_0^D\left(\lambda | \lambda_j, p\right) = \Sigma_{\lambda,\lambda} - \Sigma_{\lambda,s}\Sigma_{s,s}^{-1}\Sigma_{s\lambda} = \frac{\sigma_\varepsilon^2\sigma_\lambda^2}{\sigma_\varepsilon^2 + n\sigma_\lambda^2}$$

$\blacksquare$

A.1.3 Proof of Proposition 2

Proof. Taking other bidders' strategies as given, the dealers maximize the objective

$$\max_{q_j} \pi_0(q_j) + \mu_\lambda \pi_1(q_j) - \frac{\hat{\gamma}}{2}\left(\pi_1(q_j) + \mu_\lambda \kappa^{-1}\right)^2 - pq_j$$

where $p = I_{j,D} + d_D q_j$ The first-order condition for q_j is

$$\pi_0'(q_j) + \mu_\lambda \pi_1'(q_j) - I_{j,D} - 2q_j d_D = \hat{\gamma}\left(\pi_1(q_j) + \mu_\lambda \kappa^{-1}\right)\pi_1'(q_j)$$

From Lemma 1, it follows that

$$\begin{aligned}\pi_0'(q_j) &= 1 - \kappa Q_a - \lambda_j - \kappa q_j \\ \pi_1'(q_j) &= -1\end{aligned}$$

As a result, the first-order condition is rewritten as

$$1 - \kappa Q_a - \lambda_j - \kappa q_j - \mu_\lambda - p - q_j d_D = \hat{\gamma}\left(q_j + \lambda_j \kappa^{-1} - \mu_\lambda \kappa^{-1} - Q_a\right)$$

where we substitute $I_{j,D} = p - d_D q_j$. Solving for q_j gives

$$b_D - a_D p - c_D \lambda_j = q_j = \frac{1 - p + (\hat{\gamma} - \kappa)Q_a - \lambda_j(1 + \hat{\gamma}\kappa^{-1}) + \mu_\lambda(\hat{\gamma}\kappa^{-1} - 1)}{\kappa + d_D + \hat{\gamma}}$$

Matching coefficients on p and λ_j gives

$$\begin{aligned}c_D &= (\kappa + \hat{\gamma} + d_D)^{-1}\left\{(1 + \hat{\gamma}\kappa^{-1}) - \mu_\lambda^\lambda(\hat{\gamma}\kappa^{-1} - 1)\right\} \\ a_D &= (\kappa + \hat{\gamma} + d_D)^{-1}\left\{1 - \mu_\lambda^p(\hat{\gamma}\kappa^{-1} - 1)\right\}\end{aligned}$$

which gives equations (6) and (7). The demand intercept b_D solves

$$b_D = (\kappa + \hat{\gamma} + d_D)^{-1}\left\{1 + (\hat{\gamma} - \kappa)Q_a + \bar{\mu}_\lambda(\hat{\gamma}\kappa^{-1} - 1)\right\}$$

Taking other bidders' strategies as given, the long-term investors maximize

$$\max_{q_k} q_k(1 - p) - 2\lambda_k q_k - \kappa\left(q_k\right)^2$$

where $p = I_{k,L} + d_L q_k$. The first-order condition is

$$1 - I_{k,L} - 2q_k d_D - 2\lambda_k - 2\kappa q_k = 0$$

Using $I_{k,L} = p - d_L q_k$ and solving for q_k gives

$$b_L - a_L p - c_L \lambda_k = q_k = \frac{1 - p - 2\lambda_k}{d_L + 2\kappa}$$

Matching coefficients on p and λ_k gives

$$a_L = (d_L + 2\kappa)^{-1}$$
$$c_L = 2(d_L + 2\kappa)^{-1}$$

The demand intercepts b_L is

$$b_L = (d_L + 2\kappa)^{-1} = a_L$$

This is the desired result and it completes the proof. ∎

A.1.4 Proof of Proposition 3

Proof. If $n = 0$, the slope of the inverse residual supply is $d_L = a_L^{-1}(m-1)^{-1}$. It follows that

$$a_L = (a_L^{-1}(n-1)^{-1} + 2\kappa)^{-1}$$

Solving for a_L gives

$$a_L^{\text{IPV}} = \frac{m-2}{m-1}\frac{1}{2\kappa}$$

Since $c_L = 2a_L$, it immediately follows that $c_L^{\text{IPV}} = \frac{m-2}{m-1}\frac{1}{\kappa}$, concluding the proof. As $m \to \infty$, the IPV case approaches the price-taking benchmark. ∎

A.1.5 Proof of Proposition 4

Proof. Setting $m = 0$, it immediately follows that $d_D = a_D^{-1}(n-1)^{-1}$ and that

$$\mu_\lambda^\lambda = 0$$

Therefore, the system of equations simplifies to

$$c_D = \frac{1 + \hat{\gamma}\kappa^{-1}}{\kappa + \hat{\gamma} + d_D}$$
$$a_D = \frac{1 + \frac{na_D}{c_D}\tau_\lambda(\hat{\gamma}\kappa^{-1} - 1)}{\kappa + \hat{\gamma} + d_D}$$

where we substitute $\mu_\lambda^p = -\frac{na_D}{c_D}\tau_\lambda$. Hence

$$a_D = (\kappa + \hat{\gamma} + d_D)^{-1}\left\{1 + na_D\tau_\lambda(\kappa + \hat{\gamma} + d_D)\frac{\hat{\gamma} - \kappa}{\hat{\gamma} + \kappa}\right\}$$

Rearranging and using the definition of d_D gives

$$\frac{a_D(n-1)(\kappa + \hat{\gamma}) + 1}{n-1} = \left[1 - n\tau_\lambda\frac{\hat{\gamma} - \kappa}{\hat{\gamma} + \kappa}\right]^{-1}$$

Solving for a gives

$$a_D = \left\{(n-1)\left[1 - n\tau_\lambda \frac{\hat{\gamma}-\kappa}{\hat{\gamma}+\kappa}\right]^{-1} - 1\right\} \frac{1}{(n-1)(\kappa+\hat{\gamma})} = \frac{(n-2)(\hat{\gamma}+\kappa)+n\tau_\lambda(\hat{\gamma}-\kappa)}{(n-1)(\hat{\gamma}+\kappa)\left[(\hat{\gamma}+\kappa)-n\tau_\lambda(\hat{\gamma}-\kappa)\right]}$$

Finally, c_D is given by

$$c_D = \frac{\hat{\gamma}\kappa^{-1}+1}{\kappa+\hat{\gamma}+(n-1)^{-1}a_D^{-1}} \frac{\hat{\gamma}\kappa^{-1}+1}{\kappa+\hat{\gamma}+\frac{(\hat{\gamma}+\kappa)[(\hat{\gamma}+\kappa)-n\tau_\lambda(\hat{\gamma}-\kappa)]}{(n-2)(\hat{\gamma}+\kappa)+n\tau_\lambda(\hat{\gamma}-\kappa)}} = \frac{(n-2)(\hat{\gamma}+\kappa)+n\tau_\lambda(\hat{\gamma}-\kappa)}{(n-1)(\hat{\gamma}+\kappa)\kappa}$$

In summary

$$\begin{aligned} a_D^{\mathrm{CV}} &= \frac{(n-2)(\hat{\gamma}+\kappa)+n\tau_\lambda(\hat{\gamma}-\kappa)}{(n-1)(\hat{\gamma}+\kappa)\left[(\hat{\gamma}+\kappa)-n\tau_\lambda(\hat{\gamma}-\kappa)\right]} \\ c_D^{\mathrm{CV}} &= \frac{(n-2)(\hat{\gamma}+\kappa)+n\tau_\lambda(\hat{\gamma}-\kappa)}{(n-1)(\hat{\gamma}+\kappa)\kappa} \end{aligned}$$

which is the desired result and completes the proof. ■

A.1.6 Proof of Proposition 5

Proof. Downward sloping demand schedules, that is $a_D > 0$ and $a_L > 0$, require that $c_L \in \mathcal{C}_L \doteq \left(\frac{1}{\kappa}\frac{m-2}{m-1}, \frac{1}{\kappa}\right)$. Further, $c_D > 0$. To see why, suppose instead that $c_D \leqslant 0$. Given $a_D > 0$ and $a_L > 0$, it follows that $\kappa + \hat{\gamma}(\boldsymbol{c}) + d_D(\boldsymbol{c}) > 0$. Hence, $c_D = f_2(\boldsymbol{c}) \leqslant 0$ is negative only if

$$\mu_\lambda^\lambda(\boldsymbol{c})(\hat{\gamma}(\boldsymbol{c})\kappa^{-1}-1) \geqslant \hat{\gamma}(\boldsymbol{c})\kappa^{-1}+1 \iff \mu_\lambda^\lambda(\boldsymbol{c}) \leqslant \frac{\hat{\gamma}(\boldsymbol{c})\kappa^{-1}+1}{\hat{\gamma}(\boldsymbol{c})\kappa^{-1}-1} = -\left(\frac{2\gamma\Sigma_\lambda(\boldsymbol{c})}{\kappa}+1\right) \leqslant -1$$

Substituting the definition of $\mu_\lambda^\lambda(\boldsymbol{c})$ and rearranging, $\mu_\lambda^\lambda(\boldsymbol{c}) \leqslant -1$ implies that

$$m\sigma_\lambda^2 c_L c_D \geqslant [(n-1)c_D + mc_L]^2\sigma_\lambda^2 + \left[c_D^2(n-1)+c_L^2 m\right](\sigma_\varepsilon^2+\sigma_\lambda^2) + m\sigma_\lambda^2 c_L^2$$

All terms on the right-hand side are strictly positive, so the inequality cannot hold for $c_D \leqslant 0$. Further, since $n > 1$, there is no $\boldsymbol{c} \in \mathbb{R}^2$ such that $\mu_\lambda^\lambda(\boldsymbol{c}) \leqslant -1$. Hence $f_2(\boldsymbol{c}) > 0$ for all $\boldsymbol{c} \in \mathbb{R}^2$.

I next derive bounds for $c_D = f_2(\boldsymbol{c})$. First, $f_2(\boldsymbol{c}) > 0$ for all $\boldsymbol{c} \in \mathbb{R}^2$, so that $m_D = 0$ is a lower bound. Given that $f_2(c_L, c_D)$ is continuous in c_D, $f_2(c_L, 0) > 0$ and since $\lim_{c_D\to\infty} f_2(c_L, c_D) < \infty$ is finite for all $c_L \in \mathcal{C}_L$, the intermediate value theorem implies that there is $\bar{c}_D(c_L) < \infty$ such that $\bar{c}_D(c_L) = f_2(c_L, \bar{c}_D(c_L))$ for all $c_L \in \mathcal{C}_L$. Hence, there is an upper bound $M_D < \infty$ such that $f(c_L, c_D) \leqslant M_D$. Thus, the set $\mathcal{F}_D = [0, M_D]$ satisfies $0 < f_2(\boldsymbol{c}) \leqslant M_D$ for all $\boldsymbol{c} \in \mathcal{C}_L \times \mathcal{F}_D$.

Second, consider $f_1(\boldsymbol{c})$, $c_D \in \mathcal{F}_D$ and $c_L \in \mathcal{C}_L = \left(\frac{1}{\kappa}\frac{m-2}{m-1}, \frac{1}{\kappa}\right)$. The function $f_1(\boldsymbol{c})$ is continuous over $\mathcal{C}_L \times \mathcal{F}_D$. As c_L approaches the left endpoint $\frac{1}{\kappa}\frac{m-2}{m-1}$, $f_1(\boldsymbol{c}) \to \infty$ for all $c_D \in \mathcal{F}_D$. The limit of $f_1(\boldsymbol{c})$ as c_L approaches the right endpoint $\frac{1}{\kappa}$ is

$$\lim_{c_L\to\frac{1}{\kappa}} f_1(c_L, c_D) = \frac{n\cdot\sigma_\lambda^2((n-1)c_D + m\kappa^{-1})\cdot(\hat{\gamma}(\kappa^{-1}, c_D)\kappa^{-1}-1)}{(\kappa+\hat{\gamma}(\kappa^{-1}, c_D))\left[((n-1)c_D+m\kappa^{-1})^2\sigma_\lambda^2 + \left(c_D^2(n-1)+m\kappa^{-2}\right)(\sigma_\varepsilon^2+\sigma_\lambda^2)\right]} \cdot \frac{1}{\kappa}$$

which is finite. The denominator is strictly positive, whereas the numerator is strictly negative given that $(\hat{\gamma}(\kappa^{-1}, c_D)\kappa^{-1} - 1) < 0$. Hence $\lim_{c_L \to \frac{1}{\kappa}} f_1(c_L, c_D) < 0$. By the intermediate value theorem, there exists c_L^* for each $c_D \in \mathcal{F}_D$ with the property that $c_L^* = f_1(c_L^*, c_D)$ such that $c_L^* \in \mathcal{C}_L$. Define the function $g(c_L^*, c_D) \doteq f_1(c_L^*, c_D) - c_L^*$, and note that $g(c_L^*, c_D) = 0$. By the implicit function theorem, there exists a continuously differentiable function $c_L^* = h(c_D)$ in an open set $U \subset \mathbb{R}$, $c_D \in U$. Because the conditions for the implicit function theorem are satisfied for all $c_D \in \mathcal{F}_D$, h is continuous over $\mathcal{F}_D$.

Substituting $c_L^* = h(c_D)$ into $f_2(\boldsymbol{c})$ gives

$$c_D = f_2(h(c_D), c_D) = \frac{\hat{\gamma}(h(c_D), c_D)\kappa^{-1} + 1 - \mu_\lambda^\lambda(h(c_D), c_D)(\hat{\gamma}(h(c_D), c_D)\kappa^{-1} - 1)}{\kappa + \hat{\gamma}(h(c_D), c_D) + d_D(h(c_D), c_D)}$$

The range of $h(c_D)$ is $\mathcal{C}_L$, so that $f_2(h(c_D), c_D)$ is continuous and maps the compact and convex set $\mathcal{F}_D$ into itself. Brouwer's fixed point theorem applies, and there exists a fixed point $\boldsymbol{c}^* = (c_L^*, c_D^*)$, $f(\boldsymbol{c}^*) = \boldsymbol{c}^*$, such that $a_D > 0$ and $a_L > 0$. Both (c_L^*, c_D^*) are strictly positive. ■

A.1.7 Proof of Proposition 6

Proof. By way of contradiction, suppose that there is a symmetric equilibrium with $a = a_D = a_L$ and $c = c_D = c_L$. The number of bidders is $N = m + n$. It follows that $\mu_\lambda(\lambda_j) = 0$ and that

$$\mu_\lambda(p) = -\frac{Na\sigma_\lambda^2}{c\left[N\sigma_\lambda^2 + \sigma_\varepsilon^2\right]} = -\frac{Na}{c}\tau_\lambda$$

Since the slope is $d_D = d_A = d = (N-1)^{-1}a^{-1}$, it immediately follows that

$$\begin{aligned}
a &= \left(\kappa + \hat{\gamma} + (N-1)^{-1}a^{-1}\right)^{-1}\left\{1 + \frac{Na}{c}\tau_\lambda(\hat{\gamma}\kappa^{-1} - 1)\right\} \\
c &= \left(\kappa + \hat{\gamma} + (N-1)^{-1}a^{-1}\right)^{-1}(\hat{\gamma}\kappa^{-1} + 1) \\
a &= ((N-1)^{-1}a^{-1} + 2\kappa)^{-1} \\
c &= 2((N-1)^{-1}a^{-1} + 2\kappa)^{-1}
\end{aligned}$$

Since all these equations must hold at the same time, the third and the fourth line imply that $a = 2c$. However, it must then be the case that

$$1 + 2N\tau_\lambda(\hat{\gamma}\kappa^{-1} - 1) = 2\left(\hat{\gamma}\kappa^{-1} + 1\right)$$

Rearranging, we see that

$$(N\tau_\lambda - 1)2\frac{\hat{\gamma}}{\kappa} - 1 - 2N\tau_\lambda = 0$$

However, since $N\tau_\lambda - 1 = -\tau_\varepsilon < 0$, $\hat{\gamma} > 0$, $\kappa > 0$ and $N > 0$, all terms on the right hand side are negative. It follows that the equation cannot hold, which is a contradiction. ■

A.1.8 Proof of Proposition 7

Consider the interval $c_L \in \left(\frac{1}{\kappa}\frac{m-2}{m-1}, \frac{1}{\kappa}\right)$ in which demand schedules are downward sloping for both types. Differentiating a_D with respect to c_L gives

$$\frac{\partial a_D}{\partial c_L} = \frac{1}{2n}\frac{1-(m-1)(1-\kappa c_L)^2}{(1-\kappa c_L)^2}\frac{\partial a_D}{\partial c_L} = \frac{1}{2}\cdot\frac{1}{n}\left\{\frac{1}{(1-\kappa c_L)^2}-(m-1)\right\}$$

Since $\frac{\partial a_L}{\partial c_L} = \frac{1}{2}$, it is sufficient to show that

$$\frac{1}{n}\left\{\frac{1}{(1-\kappa c_L)^2}-(m-1)\right\} > 1$$

Define the function $f(x)$

$$f(x) = \frac{1}{n}\left\{\frac{1}{(1-\kappa x)^2}-(m-1)\right\}$$

in the interval $x \in \mathcal{X} \doteq \left[\frac{1}{\kappa}\frac{m-2}{m-1}, \frac{1}{\kappa}\right)$. First, note that $f'(x) > 0$ for $x \in \mathcal{X}$. In fact, since $x > 0$, $n > 0$.

$$f'(x) = \frac{1}{n}\frac{2\kappa}{(1-\kappa x)^3} \quad\longrightarrow\quad f'(x) > 0 \quad \text{if} \quad x < \frac{1}{\kappa}$$

Hence, $f(x)$ is minimal for $x_0 = \frac{1}{\kappa}\frac{m-2}{m-1}$. Dealers are more sensitive if $f(x_0) > 1$, or

$$f(x_0) = \frac{1}{n}(m-1)(m-2) > 1 \iff (m-1)(m-2) > n$$

which completes the proof.

A.2 Normal Random Variables and CARA Utility

Let

$$\mu = \begin{bmatrix}\mu_\theta \\ \mu_s\end{bmatrix} \quad : \quad \Sigma = \begin{bmatrix}\Sigma_{\theta,\theta} & \Sigma_{\theta,s} \\ \Sigma_{s,\theta} & \Sigma_{s,s}\end{bmatrix}$$

The conditional density of θ given s is normal with conditional mean

$$\mathbb{E}[\theta|s] = \mu_\theta + \Sigma_{\theta,s}\Sigma_{s,s}^{-1}(s-\mu_s)$$

and variance-covariance matrix

$$\mathbb{V}\mathrm{ar}(\theta|s) = \Sigma_{\theta,\theta} - \Sigma_{\theta,s}\Sigma_{s,s}^{-1}\Sigma_{s,\theta}$$

In a two-dimensional case, since $\mu_s = \bar{\theta}$, the conditional expectation can be written as

$$\mathbb{E}[\theta|s] = \frac{\mathbb{C}\mathrm{ov}(\theta,s)}{\mathbb{V}\mathrm{ar}(s)}s + \left(1-\frac{\mathbb{C}\mathrm{ov}(\theta,s)}{\mathbb{V}\mathrm{ar}(s)}\right)\bar{\theta} = \xi s + (1-\xi)\bar{\theta}$$

which is the standard Bayesian updating rule with normal priors and normal likelihood. The expression can be generalized to accommodate for common and private values.

Another useful result is the expectation of a quadratic form of normal random variables. Let $\omega = c + b'z + z'Az$ and $z \sim \mathcal{N}(0, \Sigma)$. Then

$$-\mathbb{E}[e^{-\rho\omega}] = (\det \Sigma)^{-\frac{1}{2}} \left(\det\left(\Sigma^{-1} + 2\rho A\right)\right)^{-\frac{1}{2}} e^{-\rho\left[c - \frac{1}{2}\rho b'\left(\Sigma^{-1} + 2\rho A\right)^{-1} b\right]}$$

A proof can be found in Danthine and Moresi (1993).

B Additional Results

B.1 Determinants of demand heterogenenity

Correlation between demand elasticity and auxiliary variables

	Log demand elasticity						
	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}
$\sigma_{j-21,j}$	-1.78^{***}						-0.91^{***}
	(0.13)						(0.16)
$\sigma_{b,j}$		-0.31					-0.08
		(0.52)					(0.42)
RBAS$_j$			-0.68^{***}				-0.30^{***}
			(0.04)				(0.07)
Maturity				-0.05^{***}			-0.01
				(0.00)			(0.01)
Participants					0.07^{***}		0.03^{***}
					(0.01)		(0.01)
Log issue size						0.31^{***}	-0.02
						(0.05)	(0.06)
Constant	5.88^{***}	5.15^{***}	5.59^{***}	5.93^{***}	4.19^{***}	1.15^{**}	5.73^{***}
	(0.06)	(0.04)	(0.04)	(0.06)	(0.09)	(0.57)	(0.66)
Adj. R^2	0.22	-0.00	0.22	0.20	0.09	0.04	0.28
N	993	1087	1087	1188	1188	1188	993

Table 9: Determinants of demand elasticity. TE is the total elasticity of demand using all bids. $\sigma_{j-21,j}$ denotes the volatility of the bond in the month prior to the auction. RBAS$_j$ is the relative bid-ask spread at the auction date. Maturity refers to time to maturity at the auction date. Participants is the number of bidders. Log issue size is the logarithm of supply. he sample is from 2000 to present. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.

Alternative measures of demand elasticity

	Log demand elasticity		
	TE β_{ij}	IE β_{ij}	WE β_{ij}
$\sigma_{j-21,j}$	-0.44^{**}	-0.73^{***}	-0.45
	(0.19)	(0.19)	(0.30)
$\sigma_{j-21,j} \times \mathbb{1}\{\text{DB}\}_i$	-0.57^{***}	-0.35^{***}	-0.54^{***}
	(0.12)	(0.12)	(0.19)
Maturity	-0.03^{***}	-0.03^{***}	-0.03^{***}
	(0.01)	(0.01)	(0.01)
Participants	0.02^{***}	0.01	0.03^{**}
	(0.01)	(0.01)	(0.02)
Log issue size	-0.08	-0.09^{*}	-0.08
	(0.06)	(0.05)	(0.10)
Constant	6.59^{***}	7.02^{***}	6.64^{***}
	(0.67)	(0.65)	(1.16)
Macro	✓	✓	✓
Adj. R^2	0.30	0.38	0.21
N	993	993	660

Table 10: Coefficient estimates of regression (10). $\sigma_{j-21,j}$ denotes the volatility of the bond in the month prior to the auction, whereas $\mathbb{1}\{\text{DB}\}$ is a binary variable equal to one if bidder i is a dealer bank. Maturity refers to time to maturity at the auction date. Participants is the number of bidders. Log issue size is the logarithm of supply. Macro controls include inflation, the short term rate (SARON), the slope of the yield curve, and the KOF economic barometer. The sample is from 2000 to present and only considers security reopenings. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively. We include alternative measures of the steepness of demand curves. TE is the total elasticity of demand using all bids. IE is the demand elasticity obtained by dropping the highest and the lower bids. WE is the demand elasticity computed using winning bids.

Robustness checks: dropping auctions around EURCHF floor removal

	Log demand elasticity			
	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}
$\sigma_{j-21,j}$	-1.58^{***}	-1.24^{***}	-0.55^{***}	-0.51^{***}
	(0.15)	(0.16)	(0.20)	(0.20)
$\sigma_{j-21,j} \times \mathbb{1}\{\mathrm{DB}\}_i$		-0.60^{***}	-0.58^{***}	-0.59^{***}
		(0.13)	(0.13)	(0.12)
Maturity			-0.02^{***}	-0.03^{***}
			(0.01)	(0.01)
Participants			0.02^{**}	0.02^{**}
			(0.01)	(0.01)
Log issue size				-0.07
				(0.06)
Constant	5.59^{***}	5.63^{***}	5.59^{***}	6.42^{***}
	(0.09)	(0.08)	(0.14)	(0.67)
Macro	✓	✓	✓	✓
Adj. R^2	0.25	0.27	0.29	0.29
N	975	975	975	975

Table 11: Coefficient estimates of regression (10). The dependent variable is the total elasticity (TE) of demand (in logs) at the bidder level. $\sigma_{j-21,j}$ denotes the volatility of the bond in the month prior to the auction, whereas $\mathbb{1}\{\mathrm{DB}\}$ is an indicator equal to one if bidder i is a dealer bank. Maturity refers to time to maturity at the auction date. Participants is the number of bidders. Log issue size is the logarithm of supply. Macro controls include inflation, the short term rate (SARON), the slope of the yield curve, and the KOF economic barometer. The sample is from 2000 to present and only considers security reopenings. We drop auctions around the EURCHF cap removal on January 15, 2015. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.

Robustness checks: dropping auctions around the Covid-19 shock.

	Log demand elasticity			
	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}
$\sigma_{j-21,j}$	-1.61^{***}	-1.26^{***}	-0.53^{***}	-0.48^{**}
	(0.16)	(0.16)	(0.20)	(0.20)
$\sigma_{j-21,j} \times \mathbb{1}\{\text{DB}\}_i$		-0.59^{***}	-0.59^{***}	-0.60^{***}
		(0.13)	(0.13)	(0.12)
Maturity			-0.03^{***}	-0.03^{***}
			(0.01)	(0.01)
Participants			0.02^{**}	0.02^{**}
			(0.01)	(0.01)
Log issue size				-0.08
				(0.06)
Constant	5.58^{***}	5.61^{***}	5.61^{***}	6.49^{***}
	(0.10)	(0.09)	(0.15)	(0.69)
Macro	✓	✓	✓	✓
Adj. R^2	0.26	0.28	0.30	0.31
N	959	959	959	959

Table 12: Coefficient estimates of regression (10). The dependent variable is the total elasticity (TE) of demand (in logs) at the bidder level. $\sigma_{j-21,j}$ denotes the volatility of the bond in the month prior to the auction, whereas $\mathbb{1}\{\text{DB}\}$ is an indicator equal to one if bidder i is a dealer bank. Maturity refers to time to maturity at the auction date. Participants is the number of bidders. Log issue size is the logarithm of supply. Macro controls include inflation, the short term rate (SARON), the slope of the yield curve, and the KOF economic barometer. The sample is from 2000 to present and only considers security reopenings. We drop auctions between March 2020 and September 2020. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.

Alternative measures of bid dispersion

	Log demand elasticity			
	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}
$\sigma_{b,j}$	0.35			
	(0.41)			
$\sigma_{b,j} \times \mathbb{1}\{\text{DB}\}_i$	-1.14^{*}			
	(0.65)			
$\hat{\sigma}_{b,j}$		0.42		
		(0.37)		
$\hat{\sigma}_{b,j} \times \mathbb{1}\{\text{DB}\}_i$		-1.10^{*}		
		(0.62)		
$\bar{\sigma}_{b,j}$			-3.14^{*}	
			(1.88)	
$\bar{\sigma}_{b,j} \times \mathbb{1}\{\text{DB}\}_i$			-5.44^{***}	
			(1.54)	
$\tilde{\sigma}_{b,j}$				-2.00
				(1.97)
$\tilde{\sigma}_{b,j} \times \mathbb{1}\{\text{DB}\}_i$				-5.57^{***}
				(1.63)
Constant	6.42^{***}	6.40^{***}	7.01^{***}	6.82^{***}
	(0.68)	(0.67)	(0.69)	(0.68)
Controls	✓	✓	✓	✓
Macro	✓	✓	✓	✓
Adj. R^2	0.33	0.33	0.35	0.35
N	993	993	993	993
p-sum	0.07	0.07	0.00	0.00

Table 13: Coefficient estimates of regression (10) using alternative proxies of cross-sectional cost dispersion. First, $\sigma_{b,j}$ is the standard deviation of quantity-weighted bid yields. Second, $\hat{\sigma}_{b,j}$ is the standard deviation of equally-weighted bid yields. Third, $\bar{\sigma}_{b,j}$ is the interquartile range of quantity-weighted bid yields. Fourth, $\tilde{\sigma}_{b,j}$ is the interquartile range of quantity-weighted bid yields. Controls include maturity, number of bidders, quantity-weighted yield spread, return volatility in the previous month, the relative bid-ask spread, and the log issue size. Macro controls include inflation, the short-term rate (SARON), the slope of the yield curve, and the KOF economic barometer. The sample is from 2000 to present and only considers security reopenings. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.p-sum is the p-value from testing the null hypothesis that $b_1 + b_2 \geq 0$.

Cost dispersion as volatility of residuals

	Log demand elasticity			
	TE β_{ij}	TE β_{ij}	TE β_{ij}	TE β_{ij}
$\hat{\sigma}_{\varepsilon,j}$	-0.02	0.03	0.04^{*}	0.04
	(0.02)	(0.03)	(0.03)	(0.03)
$\hat{\sigma}_{\varepsilon,j} \times \mathbb{1}\{\text{DB}\}_i$		-0.08^{***}	-0.08^{***}	-0.07^{**}
		(0.03)	(0.03)	(0.03)
Maturity			-0.05^{***}	-0.02^{**}
			(0.00)	(0.01)
Participants			0.01	0.02^{**}
			(0.01)	(0.01)
Log issue size				-0.07
				(0.06)
$\sigma_{j-21,j}$				-0.67^{***}
				(0.20)
RBAS_j				-0.23^{**}
				(0.09)
Constant	4.84^{***}	4.84^{***}	5.65^{***}	6.38^{***}
	(0.06)	(0.06)	(0.14)	(0.76)
Macro	✓	✓	✓	✓
Adj. R^2	0.12	0.13	0.27	0.30
N	1'087	1'087	1'087	993
p-sum		0.04	0.07	0.13

Table 14: Coefficient estimates of regression (10). The dependent variable is the total elasticity (TE) of demand (in logs) at the bidder level. $\hat{\sigma}_{\varepsilon,j}$ is the standard deviation of regression (11) whereas $\mathbb{1}\{\text{DB}\}_i$ is a binary variable equal to one if bidder i is a dealer bank, and zero otherwise. Maturity refers to time to maturity at the auction date. Participants is the number of bidders. $\sigma_{j-21,j}$ denotes the volatility of the bond in the month prior to the auction Log issue size is the logarithm of supply. RBAS_j is the relative bid-ask spread at the auction date. Macro controls include inflation, the short-term rate (SARON), the slope of the yield curve, and the KOF economic barometer. The sample is from 2000 to present and only considers security reopenings. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively. p-sum is the p-value from testing the null hypothesis that $b_1 + b_2 \geq 0$.

B.2 Difference in Differences (DiD) Design

Robustness checks: three-year window

	Log elasticity		Yield discount	
	TE β_{ij}	TE β_{ij}	Discount$_{ij}$	Discount$_{ij}$
$\mathbb{1}\{\text{Basel III}\}_j$	-0.11	-0.04	0.01	0.02*
	(0.16)	(0.14)	(0.02)	(0.01)
$\mathbb{1}\{\text{DB}\}_i$	-0.22	-0.29***	0.00	-0.00
	(0.14)	(0.11)	(0.02)	(0.02)
$\mathbb{1}\{\text{Basel III}\}_j \times \mathbb{1}\{\text{DB}\}_i$	-0.71***	-0.49**	-0.04*	-0.04*
	(0.24)	(0.20)	(0.02)	(0.02)
Maturity		-0.05***		-0.00
		(0.01)		(0.00)
Participants		-0.01		-0.00
		(0.03)		(0.00)
Log issue size		-0.00		0.00
		(0.00)		(0.00)
Constant	5.37***	6.58***	0.02	0.04
	(0.10)	(0.36)	(0.02)	(0.03)
Adj. R^2	0.12	0.36	0.00	0.01
N	350	350	881	881

Table 15: Coefficient estimates of the difference in differences specification (12). In the first and second column, the dependent variable is the total elasticity of demand (in logs). In the third and in the fourth column, the dependent variable is the quantity-weighted yield spread. $\mathbb{1}\{\text{Basel III}\}_j$ is a binary variable equal to if the auction occurs after January 2015. $\mathbb{1}\{\text{DB}\}_i$ is a binary variable equal to one if bidder i is a dealer bank. The sample period is from January 2012 to December 2017 and spans a three-year window around the introduction of the Basel III regulations. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.

Robustness checks: four-year window

	Log elasticity		Yield discount	
	TE β_{ij}	TE β_{ij}	Discount$_{ij}$	Discount$_{ij}$
$\mathbb{1}\{\text{Basel III}\}_j$	-0.14	-0.06	0.02	0.03***
	(0.15)	(0.14)	(0.01)	(0.01)
$\mathbb{1}\{\text{DB}\}_i$	-0.22*	-0.28***	0.01	0.01
	(0.12)	(0.11)	(0.01)	(0.01)
$\mathbb{1}\{\text{Basel III}\}_j \times \mathbb{1}\{\text{DB}\}_i$	-0.66***	-0.46**	-0.04**	-0.04**
	(0.21)	(0.18)	(0.02)	(0.02)
Maturity		-0.05***		-0.00
		(0.00)		(0.00)
Participants		-0.00		-0.00
		(0.02)		(0.00)
Log issue size		-0.00		0.00
		(0.00)		(0.00)
Constant	5.36***	6.42***	0.01	0.03
	(0.09)	(0.30)	(0.01)	(0.03)
Adj. R^2	0.12	0.34	0.00	0.01
N	465	465	1130	1130

Table 16: Coefficient estimates of the difference in differences specification (12). In the first and second column, the dependent variable is the total elasticity of demand (in logs). In the third and in the fourth column, the dependent variable is the quantity-weighted yield spread. $\mathbb{1}\{\text{Basel III}\}_j$ is a binary variable equal to if the auction occurs after January 2015. $\mathbb{1}\{\text{DB}\}_i$ is a binary variable equal to one if bidder i is a dealer bank. The sample period is from January 2011 to December 2018 and spans a four-year window around the introduction of the Basel III regulations. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.

Robustness checks: dropping auctions around EURCHF floor removal

	Log elasticity		Yield discount	
	TE β_{ij}	TE β_{ij}	Discount$_{ij}$	Discount$_{ij}$
$\mathbb{1}\{\text{Basel III}\}_j$	-0.13	-0.11	0.02	0.03**
	(0.14)	(0.13)	(0.01)	(0.01)
$\mathbb{1}\{\text{DB}\}_i$	-0.21^{**}	-0.27^{***}	0.01	0.01
	(0.10)	(0.09)	(0.01)	(0.01)
$\mathbb{1}\{\text{Basel III}\}_j \times \mathbb{1}\{\text{DB}\}_i$	-0.56^{***}	-0.42^{**}	-0.04^{***}	-0.04^{**}
	(0.19)	(0.16)	(0.02)	(0.02)
Maturity		-0.05^{***}		-0.00
		(0.00)		(0.00)
Participants		-0.02		-0.00
		(0.02)		(0.00)
Log issue size		-0.00^{*}		0.00
		(0.00)		(0.00)
Constant	5.35***	6.56***	0.01	0.02
	(0.08)	(0.26)	(0.01)	(0.04)
Adj. R^2	0.10	0.30	0.00	0.00
N	541	541	1302	1302

Table 17: Coefficient estimates of the difference in differences specification (12). In the first and second column, the dependent variable is the total elasticity of demand (in logs). In the third and in the fourth column, the dependent variable is the quantity-weighted yield spread. $\mathbb{1}\{\text{Basel III}\}_j$ is a binary variable equal to if the auction occurs after January 2015. $\mathbb{1}\{\text{DB}\}_i$ is a binary variable equal to one if bidder i is a dealer bank. The sample period is from January 2010 to December 2019 and spans a five-year window around the introduction of the Basel III regulations, but we drop auctions around the EURCHF cap removal on January 15, 2015. Robust standard errors are in parentheses. *, **, *** correspond to significance levels of 10%, 5% and 1%, respectively.