Price-elasticity of demand and risk bearing capacity in sovereign bond auctions*

Rui Albuquerque, José-Miguel Cardoso-Costa, and José Afonso Faias

January 15, 2022

Abstract

The paper uses data on bids submitted by primary dealer banks at auctions of sovereign debt to construct the price-elasticity of demand. The price-elasticity of demand correlates strongly with the volatility of returns of identical bonds traded in the secondary market, and correlates weakly with the bid-ask spread of the same bonds, a standard measure of market liquidity. The price-elasticity of demand predicts in sample the return of the same bonds in the secondary market at various horizons. The evidence is consistent with the price-elasticity of demand being a proxy for the risk bearing capacity of primary dealers and of price pressure in the market.

Keywords: Demand elasticity, risk-bearing capacity, primary dealers, price pressure, market liquidity, sovereign bond auctions, supply shock
JEL classification: G12, G20, G24

*We would like to thank Cristina Casalinho, Milena Wittwer, and Adi Sunderam for comments. We would like to thank the Portuguese Treasury and Debt Management Agency for providing the data and helping with the institutional details. Albuquerque is with Boston College, ECGI, and CEPR (rui.albuquerque@bc.edu). Cardoso-Costa is with Banco de Portugal and Nova School of Business and Economics (jmcosta@bportugal.pt). Faias is with Universidade Católica Portugesa – Católica Lisbon School of Business & Economics (jfaias@ucp.pt). Albuquerque and Faias gratefully acknowledge financial support from the Foundation for Science and Technology-FCT under grant PTDC/EGE-OGE/30314/2017. The views expressed in this paper are those of the authors and do not reflect the views of the Banco de Portugal, the Eurosystem, or the Portuguese Treasury and Debt Management Agency. The usual disclaimer applies.

1 Introduction

In a perfectly competitive market, the price-elasticity of demand is infinite meaning that investors will absorb any shock to supply at the equilibrium price. While perfectly competitive markets do not exist, there are markets for which the abundance of liquidity might suggest that they are close to that benchmark. One such market is the U.S. government bond market (e.g. Fleming 2003). However, even in this market, puzzling patterns exist that question the infinite-price-elasticity of demand hypothesis. Fleming and Rosenberg (2008) and Lou, Yan, and Zhang (2013) show that, around auction days, the secondary-market prices of bonds of similar maturity to those being auctioned (and sometimes the same bond) display a V-shaped pattern, which these authors argue is due to primary-dealer banks' limited risk bearing capacity.[1]

This paper studies the price-elasticity of demand in auctions of sovereign bonds. In uniform-price auctions, the cut-off price determines the price paid by every investor whose bids are at the cut-off price or higher. All other bids are unfilled. While bids are unobserved by the marketplace, they can be used to answer a question that is critical to evaluating the risk premium required by primary dealers: how much does the quantity demanded by investors decrease for a marginal increase in the price of the bond?

We analyze a proprietary bid-level data set from the Portuguese Treasury and Debt Management Agency (Portuguese acronym IGCP). These data contain all the bids of all the 73 bond auctions conducted by the Portuguese state from 2014 to 2019, a period characterized by uniform-price auctions.[2] During this period, T-bond auctions in Portugal can occur on predetermined Wednesdays of each month. On the Friday prior to the Wednesday of an auction, the IGCP announced if an auction would occur, the line(s) to be auctioned, and the suggested issuance interval. The IGCP normally kept the issue size close to the top of the indicative

[1]Fleming and Rosenberg (2008) also show that primary-dealer banks tend to sell futures before the auction to hedge some of their inventory risk and Lou et al. (2013) show that primary dealers use the secondary market to short sell similar securities to those being auctioned.

[2]Before 2014, T-bond auctions used a discriminatory price system. Lou et al. (2013) also focus on uniform price auctions.

issuance interval, which in practice means that T-bond auctions were (uninformative) supply shocks anticipated by three business days.

We find that the Portuguese government T-bond market displays the V-shaped pattern first identified by Fleming and Rosenberg (2008) and Lou et al. (2013). On average, in the ten day window around an auction, the secondary-market price of the bond being issued first drops by 15 basis points and subsequently increases by 12 basis points relative to a benchmark representing the performance of all Portuguese government bonds.[3] A large price drop occurs on the 3rd business day prior to the auction (i.e., the Friday the IGCP announces the securities to be auctioned). These numbers are of similar magnitude to those estimated by Lou et al. (2013) for the U.S. T-bond market. We conduct a counterfactual exercise by looking at unused auction dates by the IGCP. We do not observe a V-shaped curve in the secondary-market price at any bond maturity for these 'placebo' dates.

We compute a measure of the marginal elasticity of demand using unsubscribed bids near the cut-off price. There are several noteworthy properties of the absolute value of the estimated marginal elasticity of demand (for brevity, and from now on, we shall discuss the absolute value of the elasticity without explicitly mentioning it). First, the average value of the logarithm of the estimated marginal elasticity is 5.5. This is a high value that suggests that the Portuguese sovereign bond market is fairly liquid.

Second, in the cross-section of all the auctions, the marginal elasticity has a strong negative correlation with the pre-auction volatility of the secondary-market returns of the bond being auctioned. The strong negative correlation suggests that the two measures may be proxies for dealer risk-bearing capacity, since the price volatility is an important metric to assess inventory risk at the dealer level and a low elasticity indicates a low risk-bearing capacity. Second, the correlation between the marginal elasticity and the average relative bid-ask spread in the secondary market of the bond being auctioned is negative but not significant when controlling

[3] Differently from the U.S., we can use the same bond in the secondary market because all auctions in our data are re-openings of an existing bond.

for other factors that may affect the marginal elasticity such as the return volatility. As predicted by Kyle (1989), these facts suggest that the marginal elasticity has additional information content relative to the bid-ask spread. The next findings support the view that the additional information content associated with the marginal elasticity is about dealer risk-bearing capacity.

Third, we find that the V-shaped pattern first documented in Fleming and Rosenberg (2008) and Lou et al. (2013) and observed also in the Portuguese T-bond market, exists only in the sub-sample of auctions where the measured elasticity is low. This evidence is consistent with the hypothesis that primary-dealer risk premium is high in this sub-sample. For the low-elasticity sub-sample, the price drops by about 18 basis points in the five business days prior to the auction date, and then increases by about 60 basis points in the subsequent five trading days.

Fourth, we run in-sample predictive regressions of *post-auction* secondary-market bond returns at various horizons, on the marginal elasticity, the relative bid-ask spread, the price drift prior to the auction (i.e., the price change in the five trading days prior to the auction), the bid-to-cover ratio (Beetsma, Giuliodori, Hanson, and de Jong 2018), and other controls. The marginal elasticity has a negative coefficient at every horizon and is statistically significant for most holding horizons greater than four days. The negative coefficient indicates that a low value of the elasticity (suggestive of a low risk-bearing capacity and high risk premium) is associated with future price increases. Using as a benchmark the 5-day holding return after the auction, the inclusion of the elasticity leads to an increase in adjusted R-squares between 20% and 60%, depending on the regression specification. The volatility of prices has a positive sign as predicted but is statistically insignificant, though it contributes to some loss of significance associated with the marginal elasticity. The evidence suggests that the return volatility and the marginal elasticity are noisy proxies of primary dealers risk-bearing capacity, but that the marginal elasticity appears to be the less noisy of the two.[4] Price drift before the auction is not significant and even carries a positive sign suggesting that it is a poor proxy for risk-bearing

[4]Wittwer (2021) shows that the steepness of dealer demand in Canadian Treasury auctions proxies for dealer capitalization.

capacity.

Our results can be motivated by Kyle (1989) who models financial markets where informed investors' demand functions are downward sloping. The price elasticity of demand for each informed investor is linked to the risk premium and to the slope of the residual supply function that the investor faces, reflecting both investor's risk-bearing capacity and market liquidity. With a low elasticity being associated with a high risk premium, Kyle's model is thus consistent with the observed V-shaped price pattern appearing only around auctions with low price-elasticity of demand, and with the price-elasticity of demand predicting changes in the secondary market price following the auction, even after controlling for other market liquidity proxies. Likewise, Wittwer (2021) introduces a model where the demand curves of capital-constrained dealers are steeper than those of unconstrained banks, and the steepness of the aggregate demand curve reflects the average risk premium linked to primary-dealer constraints.

Our paper extends the literature that studies the predictability of secondary-market yield changes around Treasury auctions. Fleming and Rosenberg (2008) and Lou et al. (2013) interpret their findings as reflecting hedging by primary dealers with limited risk-bearing capacity that is not sufficiently accompanied by other investors. Lou et al. (2013) document that the V-shaped pattern is more pronounced for larger auction sizes, when dealers appear more capital constrained, or when the volatility of interest rates is higher (see also Beetsma, Giuliodori, de Jong, and Widijanto 2016). We complement their work providing further evidence of limited primary-dealer risk bearing capacity by showing that the price elasticity of demand subsumes some of the effects they describe. In subsequent work, Beetsma et al. (2018) show that euro-area sovereign auctions with high bid-to-cover ratios see a more pronounced increase in the secondary-market price on the day of the auction. Forrester (2018) observes a similar finding in the U.S. T-bond market, but only for the 30-year bond. For T-bond auctions in Portugal, the bid-to-cover ratio does not help predict *post-auction* returns in the secondary market and also does not remove the explanatory power of the price-elasticity of demand.

Hendershott and Menkveld (2014) provide evidence consistent with intermediaries in the New York Stock Exchange causing price pressure to mean-revert their inventory. In a dealers market, Hansch, Naik, and Viswanathan (1998) show that an increase (decrease) in inventory may cause the market maker to go from offering a best bid (ask) to offering a best ask (bid). The asset pricing literature also has identified price pressure effects from shocks to both supply and demand (e.g., Shleifer 1986, Kaul, Mehrotra, and Morck 2000, Wurgler and Zhuravskaya 2002, Coval and Stafford 2007, Lou 2012 and Wardlaw 2020). Recent asset pricing literature points to primary dealer risk-bearing as an important driver of financial asset prices. Adrian, Etula, and Muir (2014) show that shocks to the leverage of securities broker-dealers are useful to explain cross sectional variation in expected returns in stocks and bonds, Etula (2013) shows that financial assets and liabilities for U.S. security broker-dealers contain information that can explain commodity returns, and He, Kelly, and Manela (2017) show that equity capital ratio of primary dealer counterparties of the New York Federal Reserve can explain the returns on a broad set of assets. Gabaix and Koijen (2021) show evidence of significant price pressure in the stock market which they hypothesize derives from the behavior of constrained financial intermediaries. Gromb and Vayanos (2010) survey the literature that discusses the implications of institutional constraints for the ability of arbitrageurs to exploit apparent market mispricing.

Bagwell (1992) and Kandel et al. (1999) are the first to estimate demand elasticities in financial markets. Like us, Kandel et al. (1999) are able to describe the whole demand schedule for a financial asset. They have data on 27 IPOs at the Tel Aviv Stock Exchange conducted via uniform-price auctions and find a positive correlation between the abnormal return in the first day after the IPO and the price-elasticity of demand. Our exercise differs from theirs in that they do not have an asset that is a close substitute and trades on a secondary market that primary dealers can use to hedge the risk of the asset being auctioned. In sovereign bond markets, Keloharju, Nyborg, and Rydqvist (2005) also have the bids that form the demand schedule and show that the cut-off price is chosen by the Finnish Treasury to maximize the marginal revenue, which suggests that the marginal elasticity is an important determinant of the issue size, but also that the price elasticity of demand may change considerably around the cut-off price. Their paper motivates our use of a marginal elasticity using subscribed bids as a control variable in the predictive regressions. Neither Keloharju et al. (2005) nor Nyborg, Rydqvist, and Sundaresan (2002), who have data on Swedish sovereign bond auctions, relate the source of risk in the primary market to that in the secondary market as we do.

The paper proceeds as follows. The next section presents a brief description of the institutional setting of the Portuguese T-bond auctions and provides data on these auctions. In Section 3, we provide estimates of the price-elasticity of demand and introduce the remaining variables for the analysis. The main results are presented in Section 4. Section 5 discusses the robustness of the results, and Section 6 concludes.

2 Institutional background

The *Portuguese Treasury and Debt Management Agency* (IGCP) has conducted T-bond auctions using a uniform-price system since April 2014. The adoption of a uniform-price auction method in the aftermath of the euro-area sovereign debt crisis was perceived as more adequate in markets with higher volatility, as was still the case in Portugal at the time. By minimizing the winner's curse (e.g., Keloharju et al. 2005), this method was expected to promote higher participation, including by less informed investors.[5] For the purposes of this study, using data on uniform-price auctions allows for comparability with results using U.S. Treasury auctions (e.g., Lou, Yan, and Zhang 2013) where uniform-price auctions are used since the late 1990s.[6]

[5] Following Greece and Ireland, the Portuguese Republic requested international financial assistance in April 2011 and agreed on a 3-year economic adjustment program that allowed access to funding up to €78 billion (about 50% of the public debt outstanding at the time) from the IMF and EU institutions. For the ensuing 18 months, IGCP did not raise any medium- and long-term funding in the capital markets, though it continued issuing T-bills on a regular basis. In late 2012 and early 2013, IGCP conducted a number of medium-long term operations (through exchange offers and syndicated deals) that served as preparation to a full return to regular capital markets issuance, which occurred in mid-2014 with the end of the EU-IMF Adjustment Program. In April 2014, IGCP announced the first T-bond auction in exactly 3 years.

[6] All Portuguese T-bond auctions conducted until 2011 used a discriminatory-price method. Portuguese T-bills are issued at a discount and have maturities up to 12 months. All T-bills are issued using a discriminatory-price

The IGCP auctions T-bonds using a primary dealership model. A small group of financial intermediaries (the primary dealers) participate in the auctions and are responsible for marketing Portuguese debt securities to final investors in the primary market and for ensuring liquidity in the secondary market. In return, they benefit from intermediation gains by having exclusive direct access to the primary market and from issuance fees when selected as lead managers in syndicated deals.[7]

During our sample period, T-bond auctions can occur on pre-determined dates, the 2nd, 4th, or 5th Wednesdays of each month (though in the whole sample only two auctions occur in August and none in December). The week prior to each of these days, the IGCP contacts primary dealers to have their views on market conditions: it collects their opinion on whether an auction should be conducted, the lines to be auctioned, and the issuance amounts. From 2017 on, this exchange of views took the form of a detailed questionnaire sent out to all primary dealers on Thursday afternoon and returned to the IGCP by Friday morning in the week before the auction takes place. On the Friday prior to a possible auction date, the IGCP announces whether or not an auction will take place, and in the affirmative case, it discloses the specific security (or securities) to be issued and the indicative issuance amount (typically announced as an interval for all the lines being offered).[8]

The auctions comprise a competitive phase whereby bids are submitted electronically through the Bloomberg Auction System between 10:00am and 10:30am Lisbon time. Results are released until 10:45am. Only primary dealers have direct access to the bidding system. Each participant may submit up to 5 bids, in multiples of €1 million, without exceeding the upper limit of the auction indicative amount. Each bid must indicate a price expressed as a percentage of the nominal price, rounded to 2 decimal places. In uniform-price auctions, the IGCP defines a

method.

[7]Most re-openings of T-bonds are conducted via auctions. New T-bond issues are done via syndicated deals, an issuance method that allows for the participation of final investors via a book-building process managed by a subset of primary dealers (the lead managers) specifically selected for the deal.

[8]In practice, the IGCP tends to use only one of the available windows of each month, so the likelihood of conducting an auction when the previous window has not been used is relatively high.

cut-off level and all orders with a price above this threshold are accepted and settled at the cut-off price. Bids at a price equal to the cut-off level may be subject to pro-rata, if the total amount of bids up to that level is higher than the final issuance amount decided by the IGCP.[9] T-bonds were settled on t+3, up to October 2014, and on t+2 since then.

We obtain proprietary bid-level data from the IGCP. The data include all bids submitted by each primary dealer in all T-bond auctions conducted between 2014 and 2019. Table 1 reports summary statistics of T-bond auctions per year. In the beginning of the sample, which coincides with the aftermath of the sovereign debt crisis, the number of auctions was low, with only 4 auctions conducted in 2014. Portugal was slowly returning to issuing T-bonds and syndicated deals played a more important role. Since 2016, the number of T-bond auctions has oscillated between 15 and 16. T-bonds are issued with maturities between 2 and 30 years with an average maturity between 8 and 13 years. The auctions in 2014 and 2015 had an average size of about €1 billion; auction size declined to about €600 million from 2016 onward, a consequence of the IGCP starting to conduct regular simultaneous auctions on two lines.[10]

[Table 1 here]

Over the sample period, there were 24 primary dealers although in any single year the mode of the number of registered dealers was 21. The average number of participants in the auctions was 19, and only 13-15 primary dealers were allocated on average.[11] The average number of bids was quite large, about 60 in the full sample, representing an average of 3 bids per bidder.

[9] In addition to the competitive auctions, a post-auction non-competitive bidding phase is also available during which the IGCP offers an additional amount up to 20% of the auctioned amount for sale to primary dealers, which is allocated according to the dealers' participation in the competitive phase of the last three T-bond auctions. Non-competitive bids are allocated at the cut-off price from the auction. Bids in this phase are to be submitted in the 30-minute period preceding 10:30 a.m. of the following business day after the auction day (Thursday).

[10] While the IGCP typically chooses an issuance amount close to the top of the indicative interval announced, there were several instances in which it defined a cut-off price that resulted in an issuance amount either below or above this threshold. In our sample this occurred in roughly 30% of the auctions, with events below and above the threshold being evenly split. In the most extreme cases, the distance to the threshold was as high as -50% and +30% of the maximum indicative amount.

[11] In Finland, from 1992 to 1999, the sovereign bond market consisted of only 10 primary dealers (see Keloharju, Nyborg, and Rydqvist 2005).

Of these, only 25 bids on average were allocated. The bid-to-cover ratio is the ratio between the total amount bid and the allocated amount. The average bid-to-cover ratio was about 2 in the full sample, reaching an average of 2.54 in 2018.

3 Data description and elasticity measures

Our main data source, described briefly above, is the proprietary bid-level data obtained from the IGCP covering 2014 to 2019. From Bloomberg, we obtain the secondary-market bid, ask, and mid prices at the daily frequency for all T-Bonds for the same time-span, and as a benchmark for market performance, we collect the total return index of Portuguese government bonds. We obtain the spread between 10-year Portuguese T-bonds and German Bunds of the same maturity also from Bloomberg.[12]

3.1 Price elasticity of demand

If P and Q describe price and quantity, respectively, then the absolute value of the price elasticity of demand is $|(\partial Q/\partial P)(P/Q)|$. A large value of the elasticity means that small price decreases are associated with large increases in demanded quantities. This means that a shock to supply, such as a pre-announced re-opening through an auction, is absorbed by demand without much of a price decrease when the elasticity is high.

To construct our main measure of the elasticity of demand, we take the four price points (possibly the same number of bids or more) from unsubscribed bids next to the cut-off price, together with the cut-off price point itself.[13] The pairs (Q_i, P_i) are constructed such that Q_i equals the sum of the quantities bid at each price point P_i. Using these five observations, we estimate a linear regression model of Q_i on P_i and a constant. The slope in the model is an

[12]We use Bloomberg Generic Prices (BGN) for all bonds. These are computed using a Bloomberg proprietary methodology that aims at providing "consensus" prices and is based on different price contributions and other relevant information. The total return index is also computed by Bloomberg and considers the performance of all Portuguese government bonds, weighted by total amount outstanding.

[13]Across all auctions, the mean price interval needed to obtain the desired four price points is €0.08.

estimate of $\partial Q/\partial P$, from which we take the absolute value and multiply by the ratio of the cut-off price to the cut-off quantity to get the elasticity. We label this measure of the marginal elasticity of demand, ME. This measure is close in spirit to the second measure calculated in Kandel et al. (1999) that uses all unfilled orders. ME describes how much the price would have to decline if the IGCP were to increase the quantity sold beyond the subscribed amount.

We construct alternative measures of the elasticity of demand. Figure 1 describes two such measures, in addition to the one used in the paper and described above.[14] The figure plots the demand curve for a 10-year T-bond auctioned on May 11, 2016. Total elasticity (TE) differs from ME in that it uses all the demand price points to estimate the slope $\partial Q/\partial P$ from a linear regression of Q on P and a constant. Gross elasticity (GE) is obtained from the slope of the demand curve estimated using only the maximum price and the cut-off price points and corresponding quantities (see the points identified with the black diamonds). For the auction depicted in the figure, the values of ME, TE, and GE (expressed in logarithms) are 2.22, 2.12, and 2.43, respectively. A feature of this and many other auctions in our sample is that ME and TE are significantly lower than GE, reflecting the demand quasi-kink that at the cut-off price of the auction.[15] Lastly, we construct a measure similar to ME, denoted SE, that uses the four price points (possibly the same number of bids or more) from *subscribed* bids next to the cut-off price, together with the cut-off price point.[16] SE is similar to Kandel et al. (1999)'s first measure that uses the last filled orders in the auction.

[Figure 1 here]

As Figure 1 illustrates, ME is likely to contain different information from that in TE and GE. It is also information that is not available to market participants. The IGCP provides to primary dealers via Bloomberg information about the maximum price and the respective

[14] Another measure, a variant of ME, uses all the price points of the unsubscribed bids next to the cut-off point that fall within a fixed price interval. We take the price interval to be the minimum interval across all auctions that guarantees at least one point besides the cut-off point. The results are qualitative the same.

[15] We shall return to the significance of such demand kinks in subsection 3.4.

[16] Across, all auctions, the mean price interval needed to obtain the desired four price points is €0.07.

demanded quantity, which is the information needed to compute GE and in most cases GE is a close proxy to SE. The IGCP also provides the first unsubscribed price but not the associated quantity, so the market cannot replicate our estimate of ME. We conclude that the market has information about GE and SE, but not about ME or TE. Table 2 contains descriptive statistics of the various elasticity measures. All of the elasticity measures display relatively large figures, but typically ME is lower than SE, and TE is lower than GE, as in the example above, reflecting a potential kink in the demand curve.

[Table 2 here]

3.2 Auxiliary variables

Table 2 contains also descriptive statistics of the other variables used in the analysis (Table A.1 in the Appendix contains all the variable definitions). $SIZE$ and $COVER$ were defined and described before. $RBAS$ is the average of the previous 5-day period (excluding the auction day) of the daily relative bid-ask spread, calculated as the difference between the ask and the bid prices divided by the mid price (in percent). The average bid-ask spread is 0.23%. $DRIFT$ is the log return computed over the 5 days prior to the auction date (in percent) adjusted for the return of the index of Portuguese government bonds. There is a negative drift of 15 basis points with a standard deviation that is about four times as large as the absolute value of its mean. $SPREAD$ is the difference between the 10-year Portuguese T-bond and the 10-year Bund (in basis points). The average spread in the sample is about 200 basis points over the German Bund. The spread declined dramatically over the sample as the Portuguese economy improved, so the higher values are from the earlier part of the sample and the lower values from the later part of the sample. VOL is the standard deviation of log returns in the secondary market of the bond being auctioned (re-opening) over the 20 days prior to the auction date. The average daily volatility of log returns is 0.43%.

Table 3 presents the linear correlations between variables. Larger auctions ($SIZE$) are associated with higher pre-auction returns ($DRIFT$), and with lower bid ask spreads, suggesting that the IGCP sees market appetite for larger auctions in recent price increases and low bid-ask spreads. Surprisingly, with the exception of TE, the correlation between the elasticity measures and $DRIFT$ is not statistically significant. This suggests that while prices may be declining prior to the auction, the price drift may be a noisier measure of the risk-bearing capacity than the marginal elasticity we estimate. Higher $SPREAD$ and higher VOL are associated with higher $RBAS$, and $DRIFT$ is negatively associated with $SPREAD$. All four elasticity measures are strongly positively correlated with each other. We now discuss how the price elasticity is related to other liquidity measures.

[Table 3 here]

3.3 Price elasticity and other liquidity measures

The price elasticity describes the ability of demand to absorb a supply shock, and as such it is a measure of liquidity. In Kyle (1989), the price elasticity carries two components, one linked to risk aversion and the risk-bearing capacity of informed investors and another linked to market liquidity associated with the presence or lack thereof of noise traders.

Table 3 shows that, not surprisingly, all elasticity measures correlate negatively with the relative bid-ask spread, especially GE and TE (only the correlation between SE and $RBAS$ is not statistically significant). All the elasticity measures also correlate significantly with VOL, displaying correlations between -0.66 and -0.33. These high correlations suggest that there is common information between the elasticity measures and both $RBAS$ and VOL. The elasticity measures also correlate positively with $COVER$, an indication that deals that are highly subscribed relative to the allocated amount are also deals where the risk bearing capacity, and hence the demand elasticity, are highest. However, and somewhat unexpectedly, the elasticity does not display a statistically significant correlation with auction size.

Figure 2 plots the time series of quarterly averages of ME, $RBAS$, and VOL and also indicates in bold font (normal font) events by the European Central Bank of programs that increase (decrease) the liquidity of the bond market. ME and $RBAS$ appear to co-move negatively especially after 2016 as is also the case for ME and VOL. The announcement of the introduction of the ECB's Asset Purchase Program in January of 2015 is associated with a spike in the elasticity, while the end of the program is followed with a large decline in the elasticity.

[Figure 2 here]

In order to assess the relative importance of these variables in explaining the cross sectional variation of the marginal elasticity, we regress ME on $RBAS$, $SIZE$, $SPREAD$, VOL, and $DRIFT$. The results are displayed in Table 4.[17] The regression specifications allow for year and/or quarter fixed effects. VOL is the only statistically significant variable across all specifications, with significance levels at 1% or better. The results suggest that the marginal elasticity brings additional content to the traditional liquidity measures such as $RBAS$, after controlling for VOL, but that VOL and ME may capture similar aspects of liquidity that arise from inventory risk.

[Table 4 here]

3.4 IGCP behavior and the cut-off price

The negative difference observed on average between ME and SE (or GE) shows that the demand curve often depicts a kink around the cut-off price (see Table 2). While the elasticity at any given point of the demand curve is completely determined by bidder behavior, the elasticity observed around the cut-off price also depends on the particular choice of cut-off by the seller because the supply is not fixed but is chosen at the auction by the seller.

[17] We do not include $COVER$ because it uses information contemporaneous to that used to construct the ME. Including $COVER$ in the regressions produces the following results: i) it does not change the qualitative nature of the results discussed in the text regarding Table 4; ii) it has a statistically strong positive relation to ME, after controlling for the other variables; and iii) increases significantly the regressions' adjusted R-square. These results are available in the Online Appendix.

Keloharju et al. (2005) discuss the strategic behavior of the Finnish Treasury when conducting uniform-price bond auctions and show that the seller typically chooses the cut-off price to maximize marginal revenue (marginal revenue defined with respect to quantity). Although this does not maximize the total revenue in any particular auction, it may be justified by the fact that the Treasury repeatedly engages the market using auctions. If it were to maximize revenue at any one auction, the Treasury would choose the minimum bid price at that auction, which would likely have a significant negative impact in the secondary market and compromise future auctions. Keloharju et al. (2005) argue that, by maximizing the marginal revenue, around the most preferred price among bidders, the Treasury promotes more competition in subsequent auctions.

In Portuguese Treasury auctions the same pattern is also evident. Figure 3 displays the average bid amount for small intervals of €0.01 around the auction cut-off price (central bar). The figure shows that demand at the cut-off price is generally significantly higher, with about twice the bid amount relative to any other price in the vicinity. Note that in the sample, the average difference between cut-off price and the secondary-market price is €0.06 (i.e., there is overpricing), which means that the IGCP is not choosing the cut-off price to equal the price in the secondary market on average.[18]

[Figures 3 and 4 here]

Figure 4 presents additional evidence that the IGCP maximizes marginal revenue. The figure plots, for the same €0.01 minimum price variation (with the central bar representing the cut-off price), the proportion of auctions where the marginal revenue is maximized. Compared to prices

[18]The literature typically finds that Treasury auctions are underpriced relative to the secondary market, in line with theories that emphasize the winner's curse. Underpricing tends to be higher when auction size is smaller and volatility higher (see Nyborg et al. 2002 and Keloharju et al. 2005) and tends to be lower for uniform-price auctions (see Keloharju et al. 2005 and Goldreich 2007). In contrast, Cardoso-Costa, Faias, Herb, and Wu (2022) show that auctions tend to be overpriced in some Euro area countries and relate this with specific institutional features of the primary dealership model used in these countries. The authors explore the particular case of Portugal, showing suggestive evidence that overpricing may be related with aggressive bidding behavior driven by competition for syndication fees.

in the vicinity of the cut-off price, there is a significantly higher fraction of auctions where the marginal revenue is maximized at exactly the cut-off price.

By choosing a cut-off point that maximizes marginal revenue, the IGCP creates a large difference between the elasticity of demand estimated using price points to the left of the cut-off price (SE) and that estimated using price points to the right of the cut-off price (ME). Because ME uses unsubscribed bids, we hypothesize that ME captures the primary dealers' remaining risk bearing capacity in the aftermath of a bond auction. Controlling for SE in the regressions, or GE, allows us to control for the contemporaneous level of liquidity in the market. Below, we control for SE – and in a robustness test for GE – as a way to capture the kink in demand.

4 Secondary-market price dynamics around Treasury auctions

In this section, we analyze the secondary market price of the bond being auctioned around the auction day. We first analyze bond prices in a window of 11 days centered at the auction day as in Lou et al. (2013). Second, we conduct in-sample predictive return regressions on ME and other variables. These regressions allow us to control for a variety of factors besides understanding the ability of the elasticity to proxy for the risk premium associated with limited risk-bearing capacity.

4.1 Event-study analysis

We plot the cumulative log-return (i.e., the price) in the secondary market of the bond being auctioned starting 5 days prior to the auction date and ending 5 days after the auction date. Because Portuguese government bond prices showed a significant upward trend through most of the sample period, we subtract the log return of the whole portfolio of Portuguese government bonds from the log return of the bond being auctioned and designate this as the abnormal log return. The top panel of Figure 5 displays the average cumulative abnormal log return (solid line) and corresponding 90% confidence bands (grey bands). Note that day 0 represents the close price at the end of the auction day. The close price at the end of the auction day is normalized to 0.

[Figure 5 here]

The top panel of the figure shows that there is a price decline from end-of-day -5 to end-of-day -1 of about 15 basis points, with a large drop occurring on day -3, the Friday before the auction when the line to be auctioned is announced. This price decline starts to reverse on the day of the auction (recall that the day 0 price is the closing price on the day of the auction) and generally continues to increase in the days following the auction. The mean price decline prior to the auction is the mean of $DRIFT$ (of 15 bps). Skipping the day of the auction, the abnormal log return in the 5 days following the auction is about 13 basis points. This V-shaped pattern is consistent with the evidence in Lou et al. (2013) for U.S. Treasury auctions of an inverted V-shaped pattern in yields around auction dates. Lou et al. (2013) as Fleming and Rosenberg (2008) both point to primary dealers' limited risk bearing capacity, and thus price pressure, as an explanation to this phenomenon.

In the bottom panel of Figure 5, we split the auction sample by the median value of ME. Under the hypothesis that the V-shaped pattern in prices is due to primary dealers' limited risk-bearing capacity, we expect to find such pattern only when the elasticity of demand is low. For periods of high elasticity (black line), there is no noticeable price change through the event window. However, for periods of low elasticity (gray line), the V-shaped price pattern is more pronounced than the unconditional pattern in the top panel of the figure. In the subsample of low elasticity, the average price drift prior to the auction is close to -20 bps, and the log-return in the 5 days after the auction is over 50 basis points. This evidence suggests that the liquidity left untouched by the IGCP, captured by the ME metric that uses unsubscribed bids, captures the risk bearing capacity of primary dealers and thus the risk premium they require to accommodate the supply shock.[19]

4.2 Predictive regressions

We use a regression setting to test whether the elasticity of demand predicts the evolution of bond prices after auctions, controlling for other factors. If the elasticity of demand proxies for the risk premium linked to the primary dealers' risk-bearing capacity, then a low elasticity of demand predicts higher returns going forward on average. We run a series of cross-sectional regressions that differ in the left-hand side variable, the log abnormal return for auction i measured from the close on the auction day to day h after the auction, $R_{i,h}$. For each h, we estimate the model

$$R_{i,h} = \beta X_i + \epsilon_{i,h}, \tag{1}$$

where X_i includes the marginal elasticity, ME_i, and control variables for auction i. The control variables are the relative bid-ask spread ($RBAS$), the price drift prior to the auction ($DRIFT$), the auction size ($SIZE$), the bid-to-cover ratio ($COVER$), the spread between Portuguese and German 10-year T-Bonds ($SPREAD$), the return volatility of the bond being auctioned (VOL), and another measure of the elasticity to capture the kink in demand (SE). All of the variables are pre-determined. The marginal elasticity (not disclosed to the market) and the bid-to-cover ratio (disclosed to the market) are measured with information from the auction that may or may not already be incorporated by the market in the closing price at the end of the day of the auction. This may explain some of the difference in results for $COVER$ viz-a-viz Beetsma et al. (2018). Given the evidence in Lou et al. (2013), $DRIFT$ is expected to load negatively. VOL is expected to be positively associated with the holding period return given past results in Lou et al. (2013) and Beetsma et al. (2016).

Our regression model is different from the regression models in Beetsma et al. (2016) and Beetsma et al. (2018) for two reasons. Their models use the time series of the daily yield or daily

[19] We also run the same analysis using TE and GE and do not find any significant explanatory power from these measures of the elasticity of demand. We present these results in the Robustness section.

yield change over the whole sample (the dependent variable of interest). They regress these variables on a dummy for days when auctions occur (or nearby auctions) possibly interacted with other variables. Using time series data brings in a problem of overlapping observations when the holding period horizon is longer than one day, such as in our exercise. We avoid this concern by running cross-sectional regressions with non-overlapping events. In addition, our regression model is predictive in the sense that our dependent variable is measured after the auction. We do this to test whether the price change that follows the auction is positively associated with an existing risk premium. Ideally, we would measure returns starting shortly after the auction results are announced, but because we only have daily data, we measure returns from the close on the auction day. Our regression model is closer in spirit to Lou et al. (2013) because they also use cross-sectional, predictive in sample regressions. However, they measure returns from a trading strategy that spans a window that is centered in the auction day and thus cannot use information that is available only at the auction day such as ME or $COVER$.

Table 5 summarizes the results for a specific horizon of 5 days after the auction date (i.e., $h = 5$). The first two columns show regressions without ME, but where we include variables previously proposed in the literature. They serve as a benchmark for our results. The other three columns include ME as an independent variable. The results show that $RBAS$ is significantly positively associated with the 5-day return, consistent with $RBAS$ capturing a liquidity premium associated with transaction costs (see, for example, Albuquerque, Song, and Yao 2020). $DRIFT$ is weakly positively associated with returns. This means that a price decline prior to the auction is followed by an additional price decline after the auction, controlling for $RBAS$ and other factors. This evidence suggests that $DRIFT$ is not a good proxy for limited-risk bearing capacity. The bid-to-cover ratio does not have any predictive ability. This result contrasts with evidence in Beetsma et al. (2018) and may be due to the timing assumptions in our setting as explained above. $SIZE$ is positively associated with future returns, consistent with Lou et al. (2013), but the coefficients are not statistically significant. Other control variables do not have a statistically significant predictive power except for SE that measures the marginal elasticity of demand allocated in the auction.

Table 5 also shows that ME is a significant predictor of returns. ME is statistically significant at the 1% level or better and carries a negative coefficient estimate in all specifications. In addition, adding ME to the regression contributes to an increase in the adjusted R^2 between 8 and 11 percentage points, depending on the specification. Since the table displays standardized coefficients, an estimate of -0.376 in the last column implies that a one standard deviation decrease in ME (i.e., an increase in the risk premium) translates into a price increase following the auction that is 1/3 of a standard deviation of returns. The evidence suggests that ME is a proxy for the risk premium associated with dealers' risk bearing capacity consistent with Kyle (1989).

[Table 5 about here]

Next, we present results that use the specifications in the last two columns of Table 5, but with returns measured at different horizons. The results are in Tables 6 and 7. Table 6 shows that ME is statistically significant for all horizons from four days after the auction onward. The magnitude of the effect declines somewhat from R_5 to R_{10}, but it is still quite high. Table 7 adds all other control variables. The other controls are mostly not significant statistically, but they remove some of the explanatory power associated with ME, mostly by decreasing the estimate of the coefficient associated with ME, and especially for horizons $h = 7$ and higher. In table 8, we present results where we exclude VOL from the regression because VOL and ME are strongly correlated (see Table 4) and they may be proxies for the same effect. VOL could proxy for the risk-bearing capacity as it is an input to inventory risk or because it could impact other portfolio constraints (i.e., an *input* proxy), whereas ME could proxy for the risk premium priced in investors bids (i.e., an *output* proxy). Adding all controls but VOL still results in lower estimated coefficients associated with ME, but the coefficients remain significant at the 5% level or better. It is possible that the small number of observations in our sample limits our ability to separate the roles of ME and VOL at longer horizons. In any case, the results indicate that ME is a less noisy proxy for dealers' risk-bearing capacity than VOL.

We conclude that ME provides a proxy for risk bearing and for the risk premium demanded by primary dealers for absorbing the supply shock.

[Tables 6, 7 and 8 here]

5 Robustness

In this section we confirm the robustness of our findings to a number of assumptions taken in the baseline scenario.

5.1 Counterfactual analysis

We conduct a counterfactual exercise by looking at unused auction dates by the IGCP – 2nd, 4th and 5th Wednesdays of each month. For these dates, we computed the excess return of 5-year and 10-year bonds (the benchmark maturities) above the performance of the aggregate bond index around these dates. In total, we gathered 82 potential, but unused, auction dates between April 2014 and the end of 2019.[20] We report these results in Figure 6 for three maturities, 2, 5 and 10 years. At each auction date, We find the bonds that are closest to these maturities and check the dynamics of secondary adjusted prices in the window from 5 days before to 5 days after the auction taking place. In these windows, secondary market prices do not present the V-shaped movement observed around executed auctions. Interestingly, the excess return of the 10-year bond seems to increase throughout the whole window, while the excess return of the 2-year and 5-year bonds decrease, but in either case the evolution is gradual and steady, without any discernible movement in the announcement day or throughout the window. The trends observed over the 11-day window are similar to those observed in a wider window comprising

[20]We have excluded Wednesdays falling in late December and early August, as IGCP tends to reduce issuance activity in these periods, as well as those immediately after a syndicated deal, as the IGCP tends to avoid an excessive issuance concentration.

20 days before and after the potential date, suggesting that the 10-year bond tends to outperform the aggregate index (while the 5-year bond underperforms), which may be justified by the fact that this is usually the most liquid on-the-run bond (typically a new benchmark bond issued in the beginning of the year).

[Figure 6 here]

5.2 Alternative measures of elasticity

The results that we found may be just driven by any part of the demand curve. Is there a V-inverted curve in secondary prices of Treasury bond adjusted returns when we split according to the other measures with semi-public information such as GE or private information for the primary dealers such as TE or SE. Thus, we have repeated the event-study analysis, but now splitting the sample using the other proposed elasticity measures. We report these results in Figure 7. Unlike what happens when using ME, there is no statistically significant difference between the average patterns observed in auctions with high/low TE, GE or SE. As already noted inwhen discussing the predictive regressions, in general these measures do not have a strong predictive power of post-auction excess returns. This shows that the predictive power is mainly coming from the marginal unsubscribed bids in the demand curve.

[Figure 7 here]

5.3 Yields' impact

We run the the event-study analysis using yields instead of prices, to confirm the statistical and economic significance of the impact of the marginal elasticity on bonds' excess returns. Yields allows to address in a sole measure the heterogeneity of features between different Treasury bonds, namely maturity and coupon rate. As a good proxy for the average yield of the aggregate bond index is not readily available, we simply plot the evolution of the raw yields of the bonds being auctioned. We report these results in Figure 8. As expected, it follows a typical inverted v-shape, with an average reduction of 2 bps in the 5 day post-auction period. The order of magnitude of these movements is in line with the results obtained by Lou et al. (2013) or Beetsma et al. (2016). The advantage of using prices in the baseline analysis is the fact that we can compute abnormal returns by adjusting Treasury bond raw returns for the Portuguese Treasury bond index returns.

[Figure 8 here]

5.4 Including the year 2020

We do not include in the baseline setup the year of 2020 to show that results are not driven by a potential atypical year. Nevertheless, we repeat all the analysis with an extended sample that includes all the auctions conducted in 2020, except those that have taken place in the peak of the first wave of the pandemic crisis, on March 11, as bond prices were extremely volatile in that period. This increases the sample size to 89 observations. We report one of the tables with the main results for the horizon of 5 days and leave the rest for the online appendix. In Table 9 it is noticeable that the results are qualitatively the same. In fact, auctions with low marginal elasticity hold predictive power of returns for horizons of 5 or more days after the auction day.

[Table 9 here]

5.5 Bid shading

The empirical literature typically finds that Treasury auctions are on average underpriced relative to the secondary market. This is usually associated with bid shading behavior, which tends to be stronger when volatility in the secondary market is higher.[21] In the Portuguese case, auctions are actually overpriced on average. In our sample about 6 bps. Cardoso-Costa et al. (2022) show that in the period between 2004 and 2018, the average overpricing was about 11 bps. This

[21] See Nyborg et al. (2002), Keloharju et al. (2005), or Goldreich (2007).

is a common feature in a number of Euro area countries, in particular when issuance through syndicated deals is a significant fraction of issuance activity.[22]

In any case, primary dealers strategic behavior in Treasury auctions may bias the demand schedule. It is not clear whether bid shading would bias the demand level or slope (our focus here). In order to control for this possibility, we add two additional variables to our predictive regressions: underpricing (UP), measured as the difference between the secondary market price at the end of the auction day and the cut-off auction price, and bid discount ($DISC$), measured as the cross-dealer average of the difference between the secondary market price at the end of the auction day and the average bid price of each dealer. Results are presented in the Online Appendix and are qualitatively and quantitatively identical do the baseline specification. The coefficients associated with these two variables are not statistically significant, except for the 1-day abnormal return, where the coefficient is positive, suggesting that auctions with higher underpricing (bid shading) have a higher excess return in the very short term. Crucially, auctions with low marginal elasticity continue to hold predictive power of returns for horizons of 5 or more days after the auction day. This suggests that the demand schedule revealed in the auction has important information content, even if this demand schedule may be biased due to potential strategic behavior.

6 Conclusion

A common assumption in many financial-asset market models is that supply shocks are absorbed with no price variation, i.e., the price-elasticity of demand is infinite. However, the empirical finance literature has demonstrated many examples where prices move in response to supply shocks, but this literature is yet to demonstrate that in fact the elasticity of demand plays a role in the ability to absorb these shocks. The concern is that some other feature is causing the price

[22]See Cardoso-Costa et al. (2022) who relate this with primary dealership models that promote stronger auction competition, as those primary dealers with higher allocation are typically selected as lead managers in syndicated deals for which they are compensated.

movement.

This paper uses the observed aggregate demand data on auctions of sovereign debt to calculate the price elasticity of demand. It then shows that an apparent price pressure phenomenon in the secondary market around auction days is connected to the price elasticity of demand. This evidence contributes to a burgeoning literature that studies how uninformative flow shocks into asset markets may generate price pressure in those markets.

From a policy perspective, issuers may benefit from knowing the value of the price elasticity of demand when determining the cut-off price of the auction since the elasticity correlates with the price in secondary market in the days after the auction.

In future work, we would like to better understand the incentives of primary dealers in the selection of securities to be auctioned. Are primary dealers interested in securities whose demand is expected to be high post auction, or securities whose price has been going up prior to the auction?

References

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

Albuquerque, R., S. Song, and C. Yao (2020). The price effects of liquidity shocks: A study of the sec's tick size experiment. *Journal of Financial Economics 138*(3), 700–724.

Bagwell, L. S. (1992). Dutch auction repurchases: An analysis of shareholder heterogeneity. *The Journal of Finance XLVII*, 71–106.

Beetsma, R., M. Giuliodori, F. de Jong, and D. Widijanto (2016). Price effects of sovereign debt auctions in the euro-zone: The role of the crisis. *Journal of Financial Intermediation 25*, 30–53.

Beetsma, R., M. Giuliodori, J. Hanson, and F. de Jong (2018). Bid-to-cover and yield changes around public debt auctions in the euro area. *Journal of Banking and Finance 87*, 118–134.

Cardoso-Costa, J.-M., J. A. Faias, P. Herb, and M. Wu (2022). Seeking the winners' curse: Overpricing in treasury auctions.

Coval, J. and E. Stafford (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics 86*, 479–512.

Etula, E. (2013). Broker-dealer risk appetite and commodity returns. *Journal of Financial Econometrics 11*, 486–521.

Fleming, M. J. (2003). Measuring treasury market liquidity. *Federal Reserve Bank of New York Economic Policy Review*, 83–108.

Fleming, M. J. and J. V. Rosenberg (2008). How do treasury dealers manage their positions? *Federal Reserve Bank of New York*.

Forrester, I. J. J. (2018). The effect of treasury auctions results on interest rates: The 1990s experience. *University of Massachusetts, Amherst*.

Gabaix, X. and R. S. J. Koijen (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. Working Paper 20-91, Swiss Finance Institute Research Paper.

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

Gromb, D. and D. Vayanos (2010). Limits of arbitrage. *Annual Review of Financial Economics 2*, 251–275.

Hansch, O., N. Y. Naik, and S. Viswanathan (1998). Do inventories matter in dealership markets? evidence from the london stock exchange. *The Journal of Finance LIII*, 1623–1656.

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

Hendershott, T. and A. J. Menkveld (2014). Price pressures. *Journal of Financial Economics 114*, 405–423.

Kandel, S., O. Sarig, and A. Wohl (1999). The demand for stocks: An analysis of IPO auctions. *Review of Financial Studies 12*, 227–247.

Kaul, A., V. Mehrotra, and R. Morck (2000). Demand curves for stocks do slope down: New evidence from an index weights adjustment. *The Journal of Finance LV*, 893–912.

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

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

Lou, D. (2012). A flow-based explanation for return predictability. *Review of Financial Studies 25*, 3457–3489.

Lou, D., H. Yan, and J. Zhang (2013). Anticipated and repeated shocks in liquid markets. *Review of Financial Studies 26*, 1891–1912.

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

Shleifer, A. (1986). Do demand curves for stocks slope down. *The Journal of Finance XLI*, 579–590.

Wardlaw, M. (2020). Measuring mutual fund flow pressure as shock to stock returns. *The Journal of Finance LXXV*, 3221–3243.

Wittwer, M. (2021). Intermediary capitalization and asset demand. Working paper, Boston College.

Wurgler, J. and E. Zhuravskaya (2002). Does arbitrage flatten demand curves for stocks? *The Journal of Business 75*, 583–608.

Table 1. T-bond auction characteristics

The table reports average statistics per year for the 73 auctions. Maturity is in years. Bid-to-cover is the ratio of all bid amount by the actual issuance amount. Participant bidders (part.) are the bidders that participate with at least one bid at an auction among the registered bidders (regist.). Allocated bidders (alloc.) are the ones that had at least one bid satisfied on that auction.

Year	Nr auctions	Maturity (years)	Size (€million)	Bid-to-Cover	Nr regist. bidders	Nr part. bidders	Nr alloc. bidders	Nr bids	Nr alloc. bids
2014	4	9	981	2.44	22	21	14	91	29
2015	7	12	899	1.79	21	20	15	66	30
2016	15	8	598	1.82	21	19	15	52	27
2017	16	9	650	2.02	21	19	13	56	24
2018	15	10	579	2.54	21	19	13	57	22
2019	16	13	593	1.98	20	18	13	58	26
Overall	73	10	654	2.08	24	19	14	59	25

Table 2. Summary statistics of main variables

The table reports summary statistics for all the variables across all 73 auctions. The variable definitions are in Table A.1. p25, p50, and p75 indicate the observation in percentile 25, 50, or 75, respectively.

	Mean	SD	Min	p25	p50	p75	Max
RBASx100	0.23	0.14	0.05	0.13	0.20	0.28	0.85
DRIFTx100	-0.15	0.69	-4.02	-0.30	-0.07	0.09	1.03
SIZE	654.38	248.64	250.00	500.00	621.00	753.00	1499.00
COVER	2.08	0.50	1.46	1.70	1.93	2.29	3.76
SPREAD	207.84	92.45	61.30	137.40	185.00	299.30	386.60
VOL	0.43	0.27	0.08	0.24	0.35	0.51	1.43
ME	5.51	0.70	4.09	5.02	5.29	6.03	6.90
GE	5.57	0.42	4.46	5.33	5.54	5.88	6.39
SE	5.89	0.79	4.40	5.28	5.99	6.44	7.38
TE	5.35	0.42	4.51	5.08	5.30	5.59	6.52

Table 3. Correlations

The table reports linear correlations among the main variables across all 73 auctions. The variable definitions are in Table A.1. *, **, *** correspond to significance levels of 10%, 5%, and 1%, respectively.

Variables	RBAS	DRIFT	SIZE	COVER	SPREAD	VOL	ME	GE	SE	TE
RBAS	1.00									
DRIFT	-0.28 ***	1.00								
SIZE	-0.23 **	0.34***	1.00							
COVER	-0.04	0.10	-0.34***	1.00						
SPREAD	0.38 ***	-0.20*	-0.03	-0.23**	1.00					
VOL	0.49 ***	-0.44***	-0.11	-0.28**	0.03	1.00				
ME	-0.20 *	0.02	-0.01	0.41***	0.11	-0.42***	1.00			
GE	-0.51 ***	0.14	0.12	0.21*	-0.06	-0.51***	0.38***	1.00		
SE	-0.12	0.09	-0.04	0.48***	-0.00	-0.33***	0.49***	0.45***	1.00	
TE	-0.46 ***	0.26**	-0.05	0.51***	-0.03	-0.66***	0.50***	0.70***	0.48***	1.00

Table 4. Determinants of the marginal elasticity

The table reports standardized coefficients of regressions of ME on RBAS, DRIFT, SIZE, SPREAD, and VOL. The variable definitions are in Table A.1. Some specifications include year and/or quarter fixed effects. *t-stats* calculated using robust standard errors are reported in parenthesis below the coefficients. *, **, *** correspond to significance levels of 10%, 5%, and 1%, respectively.

	ME	ME	ME	ME
RBAS	-0.077	-0.150	-0.046	-0.117
	(-0.67)	(-1.13)	(-0.38)	(-0.85)
DRIFT	-0.170*	-0.200**	-0.109	-0.159
	(-1.68)	(-2.23)	(-1.02)	(-1.48)
SIZE	-0.017	-0.047	-0.043	-0.049
	(-0.14)	(-0.35)	(-0.33)	(-0.37)
SPREAD	0.124	0.042	0.136	0.102
	(1.07)	(0.15)	(1.14)	(0.33)
VOL	-0.457***	-0.451***	-0.480***	-0.466***
	(-4.07)	(-3.71)	(-4.01)	(-3.74)
Year FE	No	Yes	No	Yes
Quarter FE	No	No	Yes	Yes
Observations	73	73	73	73
Adj R^2	0.158	0.206	0.147	0.178

Table 5. Predictive regressions of the 5-day ahead abnormal return

The table reports standardized coefficients of predictive regressions of the 5-day ahead abnormal return on different sets of independent variables. The variable definitions are in Table A.1. *t-stats* calculated using robust standard errors are reported in parenthesis below the coefficients. *, **, *** correspond to significance levels of 10%, 5%, and 1%, respectively.

	R_5	R_5	R_5	R_5	R_5
ME			-0.427***	-0.378***	-0.376***
			(-3.49)	(-3.26)	(-3.06)
RBAS	0.452***	0.424***		0.376***	0.368***
	(3.83)	(3.38)		(3.89)	(3.64)
DRIFT	0.213	0.292**		0.199*	0.227
	(1.56)	(2.09)		(1.68)	(1.63)
SIZE		0.020			0.061
		(0.14)			(0.45)
COVER		-0.220			-0.101
		(-1.48)			(-0.77)
SPREAD		-0.187			-0.105
		(-1.64)			(-1.03)
VOL		0.291*			0.201
		(1.72)			(1.26)
SE		0.226*			0.323**
		(1.72)			(2.37)
Constant	-0.401**	-1.196	3.658***	2.529***	0.493
	(-2.31)	(-1.06)	(3.58)	(2.83)	(0.39)
N	73	73	73	73	73
Adj R^2	0.172	0.258	0.171	0.286	0.342

Table 6. Predictive regressions at various holding horizons

The table reports standardized coefficients of predictive regressions of the h-day ahead abnormal return on several independent variables (h varies from 1 to 10 days). The variable definitions are in Table A.1. *t-stats* calculated using robust standard errors are reported in parenthesis below the coefficients. *, **, *** correspond to significance levels of 10%, 5%, and 1%, respectively.

	R_1	R_2	R_3	R_4	R_5	R_6	R_7	R_8	R_9	R_{10}
ME	-0.080	-0.066	-0.129	-0.223*	-0.355***	-0.363***	-0.259**	-0.294**	-0.290**	-0.279**
	(-1.02)	(-0.58)	(-1.22)	(-1.85)	(-3.26)	(-3.37)	(-2.23)	(-2.46)	(-2.27)	(-2.32)
RBAS	-0.019	0.296**	0.374***	0.454***	0.376***	0.348***	0.398***	0.352***	0.232*	0.254*
	(-0.22)	(2.48)	(2.75)	(3.63)	(3.89)	(3.54)	(3.22)	(2.85)	(1.83)	(1.78)
DRIFT	0.736***	0.605***	0.307*	0.115	0.199*	0.074	0.0477	0.105	0.196*	0.185
	(5.92)	(2.91)	(1.89)	(0.88)	(1.68)	(0.70)	(0.43)	(0.93)	(1.82)	(1.39)
Constant	0.713	0.220	0.701	1.411	2.529***	2.709***	1.748*	2.128**	2.267**	2.285**
	(1.12)	(0.22)	(0.76)	(1.36)	(2.83)	(3.12)	(1.94)	(2.30)	(2.19)	(2.23)
N	73	73	73	73	73	73	73	73	73	73
Adj R^2	0.533	0.334	0.169	0.248	0.286	0.263	0.226	0.208	0.140	0.140

Table 7. Predictive regressions at various holding horizons with all controls

The table reports standardized coefficients of predictive regressions of the h-day ahead abnormal return on all the independent variables (h varies from 1 to 10 days). The variable definitions are in Table A.1. *t-stats* calculated using robust standard errors are reported in parenthesis below the coefficients. *, **, *** correspond to significance levels of 10%, 5%, and 1%, respectively.

	R_1	R_2	R_3	R_4	R_5	R_6	R_7	R_8	R_9	R_{10}
ME	-0.108	-0.104	-0.157	-0.236	-0.378***	-0.293**	-0.233**	-0.284**	-0.288**	-0.231*
	(-0.98)	(-0.89)	(-1.13)	(-1.65)	(-3.06)	(-2.44)	(-2.05)	(-2.50)	(-2.24)	(-1.94)
RBAS	-0.017	0.204	0.360**	0.481***	0.368***	0.335***	0.442***	0.389***	0.280**	0.352**
	(-0.17)	(1.59)	(2.01)	(3.12)	(3.64)	(3.37)	(4.22)	(3.89)	(2.18)	(2.13)
DRIFT	0.743***	0.661***	0.311*	0.116	0.227	0.158	0.079	0.137	0.212	0.174
	(6.41)	(3.42)	(1.93)	(0.79)	(1.63)	(1.43)	(0.57)	(1.01)	(1.52)	(1.19)
SIZE	0.038	0.060	0.136	0.078	0.061	-0.027	0.021	0.017	0.018	0.040
	(0.35)	(0.50)	(0.96)	(0.57)	(0.45)	(-0.20)	(0.16)	(0.13)	(0.13)	(0.25)
COVER	-0.113	-0.065	-0.058	-0.078	-0.101	-0.157	-0.168	-0.165	-0.124	-0.088
	(-1.10)	(-0.57)	(-0.42)	(-0.59)	(-0.77)	(-1.39)	(-1.33)	(-1.24)	(-1.01)	(-0.80)
SPREAD	-0.030	0.072	-0.051	-0.139	-0.105	-0.179*	-0.223**	-0.210*	-0.205	-0.289**
	(-0.29)	(0.58)	(-0.41)	(-1.18)	(-1.03)	(-1.86)	(-2.12)	(-1.76)	(-1.61)	(-2.11)
VOL	0.072	0.239	0.183	0.138	0.201	0.275*	0.182	0.185	0.137	0.084
	(0.58)	(1.29)	(1.08)	(0.93)	(1.26)	(1.68)	(0.89)	(0.93)	(0.68)	(0.49)
SE	0.221**	0.279**	0.271**	0.254*	0.323**	0.255*	0.311*	0.340*	0.283*	0.157
	(2.25)	(2.25)	(2.05)	(1.96)	(2.37)	(1.80)	(1.81)	(1.88)	(1.75)	(1.30)
Constant	-0.394	-1.826	-1.358	-0.215	0.493	0.987	0.018	0.292	0.783	1.349
	(-0.41)	(-1.44)	(-1.11)	(-0.16)	(0.39)	(0.79)	(0.01)	(0.23)	(0.59)	(1.04)
N	73	73	73	73	73	73	73	73	73	73
Adj R^2	0.540	0.375	0.198	0.271	0.342	0.325	0.297	0.285	0.175	0.168

Table 8. Predictive regressions excluding VOL

The table reports standardized coefficients of predictive regressions of the h-day ahead abnormal return on all independent variables but VOL and using SE (h varies from 1 to 10 days). The variable definitions are in Table A.1. *t-stats* calculated using robust standard errors are reported in parenthesis below the coefficients. *, **, *** correspond to significance levels of 10%, 5%, and 1%, respectively.

	R_1	R_2	R_3	R_4	R_5	R_6	R_7	R_8	R_9	R_{10}
ME	-0.122	-0.152	-0.193	-0.263*	-0.418***	-0.348***	-0.269**	-0.321**	-0.315**	-0.248**
	(-1.17)	(-1.26)	(-1.48)	(-1.94)	(-3.58)	(-2.91)	(-2.11)	(-2.52)	(-2.27)	(-2.08)
RBAS	0.014	0.306**	0.438***	0.540***	0.454***	0.452***	0.519***	0.467***	0.338**	0.388**
	(0.15)	(2.53)	(2.70)	(3.73)	(4.44)	(4.05)	(4.04)	(3.58)	(2.37)	(2.55)
DRIFT	0.717***	0.578***	0.247	0.068	0.157	0.062	0.015	0.073	0.164	0.144
	(6.20)	(3.03)	(1.60)	(0.49)	(1.22)	(0.60)	(0.13)	(0.61)	(1.40)	(1.15)
SIZE	0.041	0.072	0.145	0.085	0.071	-0.014	0.030	0.026	0.024	0.045
	(0.38)	(0.60)	(1.05)	(0.64)	(0.52)	(-0.10)	(0.22)	(0.20)	(0.17)	(0.28)
COVER	-0.123	-0.098	-0.084	-0.097	-0.128	-0.195*	-0.193	-0.191	-0.143	-0.099
	(-1.23)	(-0.89)	(-0.64)	(-0.77)	(-1.02)	(-1.68)	(-1.42)	(-1.35)	(-1.11)	(-0.94)
SPREAD	-0.045	0.023	-0.088	-0.167	-0.146	-0.235**	-0.260**	-0.248**	-0.233*	-0.306**
	(-0.40)	(0.18)	(-0.71)	(-1.45)	(-1.46)	(-2.27)	(-2.38)	(-2.02)	(-1.88)	(-2.46)
SE	0.215**	0.259*	0.255*	0.242*	0.306**	0.231	0.295	0.324*	0.271	0.150
	(2.15)	(1.92)	(1.81)	(1.78)	(2.10)	(1.49)	(1.67)	(1.74)	(1.64)	(1.18)
Constant	-0.110	-0.885	-0.640	0.329	1.285	2.070*	0.733	1.021	1.321	1.681
	(-0.14)	(-0.71)	(-0.51)	(0.27)	(1.22)	(1.82)	(0.65)	(0.94)	(1.04)	(1.23)
N	73	73	73	73	73	73	73	73	73	73
Adj R^2	0.544	0.353	0.192	0.271	0.330	0.294	0.290	0.277	0.178	0.177

Table 9. Robustness: predictive regressions of the 5-day ahead abnormal return including the year 2020

The table reports standardized coefficients of predictive regressions of the 5-day ahead abnormal return on different sets of independent variables. The variable definitions are in Table A.1. *t-stats* calculated using robust standard errors are reported in parenthesis below the coefficients. *, **, *** correspond to significance levels of 10%, 5%, and 1%, respectively.

	R_5	R_5	R_5	R_5	R_5
ME			-0.392***	-0.310***	-0.340***
			(-3.22)	(-2.77)	(-2.67)
RBAS	0.413***	0.418***		0.322***	0.349***
	(3.72)	(3.33)		(3.47)	(3.40)
DRIFT	0.160	0.235		0.151	0.194
	(1.11)	(1.58)		(1.17)	(1.30)
SIZE		0.072			0.077
		(0.54)			(0.60)
COVER		-0.145			-0.034
		(-1.00)			(-0.26)
SPREAD		-0.197*			-0.145
		(-1.93)			(-1.54)
VOL		0.327**			0.256*
		(2.29)			(1.70)
SE		0.208*			0.307**
		(1.84)			(2.56)
Constant	-0.294*	-1.571	3.276***	2.209**	-0.0922
	(-1.96)	(-1.50)	(3.39)	(2.52)	(-0.07)
N	89	89	89	89	89
Adj R^2	0.137	0.245	0.144	0.218	0.307

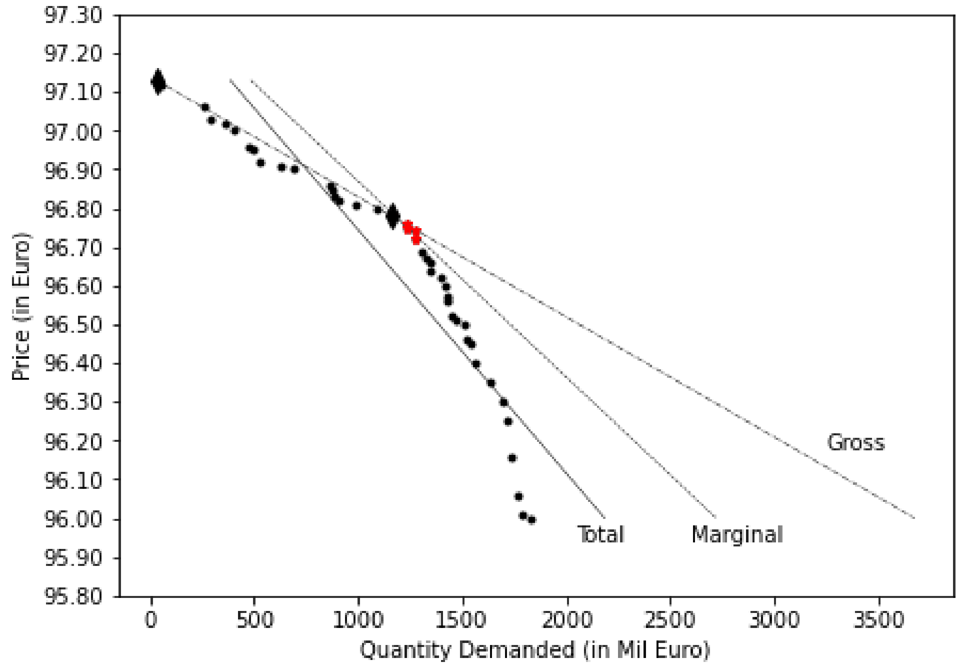

Figure 1. Example of an auction demand curve.
The figure presents all the bid prices for the auction of May 11, 2016 of a 10-year Treasury Bond. The range of bids was between 96.00 and 97.13 with a cut-off price of 96.78. The secondary market price at the end of the day on the same bond was 96.993. The bid amount was €1.83 billion and the allocated amount was €1.15 billion with an initial expected amount announced by the IGCP of €1 billion. The figure also presents three slopes used to construct three different elasticities of demand. *Gross* represents the slope of the gross demand curve using just the cut-off price and the maximum bid price. *Total* represents the slope of the total demand curve using all bid points. *Marginal* represents the slope of the marginal demand curve using the cut-off price and 4 unsubscribed points to the right of the cut-off, which had their demand unsatisfied. In this auction, the value of ME, TE and GE are 2.22, 2.12, and 2.43, respectively.

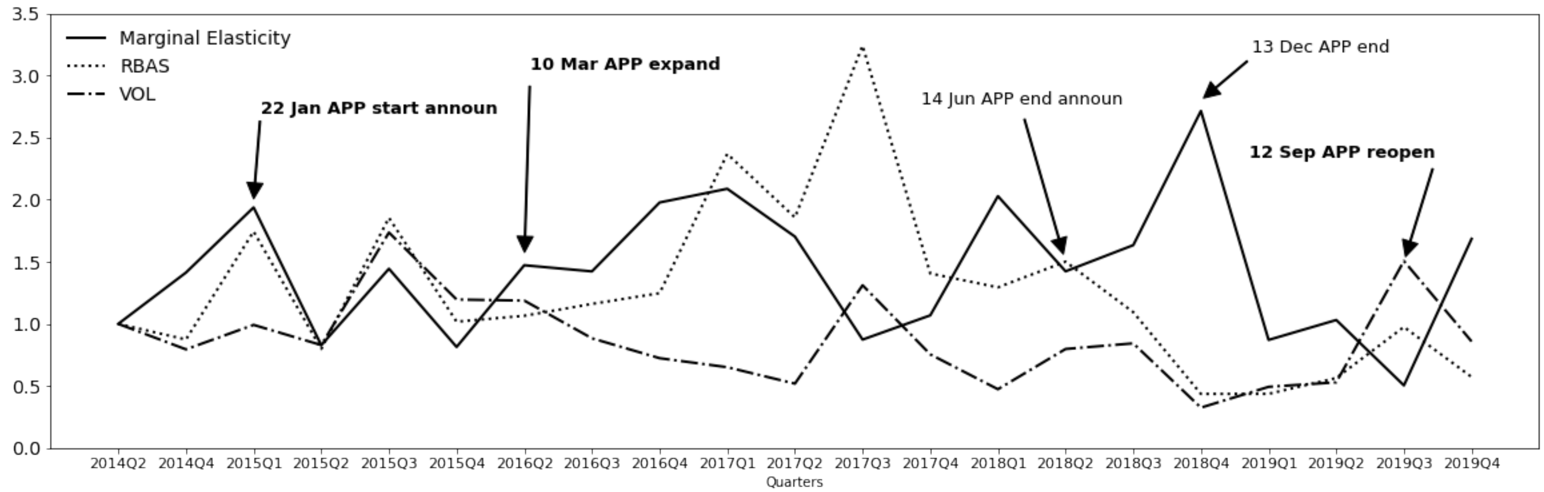

Figure 2. Time series of marginal elasticity, relative bid-ask spread, and volatility.
The figure presents quarterly average values of the marginal elasticity (ME), relative bid-ask spread (RBAS), and volatility (VOL). The three series are normalized to 1 at the initial point for ease of comparison. In boldface, we indicate announcements by the European Central Bank of programs that increase liquidity, and in normal font, we indicate announcements by the ECB of programs that decrease liquidity.

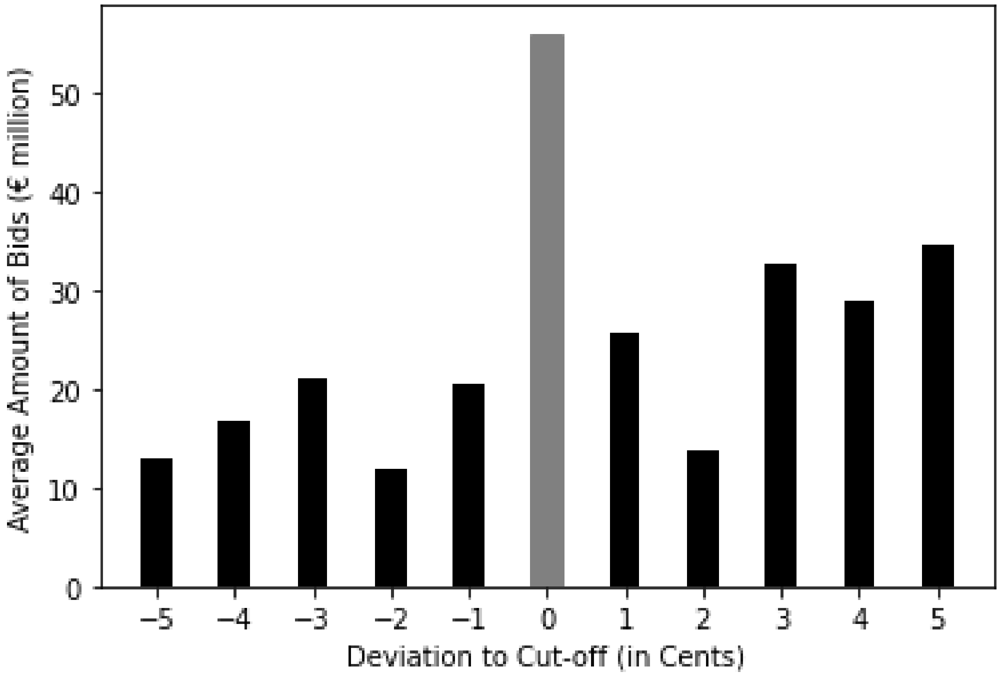

Figure 3. Average bid amount around the cut-off price.
For each auction and conditional on an interval of +/- €0.05 around the cut-off price, the figure plots the average across all auctions of the bid amounts at each each price. Negative values to the left of the cut-off price indicate unsubscribed prices €-0.05 to €-0.01 below the cut-off price. Positive values to the right of the cut-off price indicate subscribed prices €0.01 to €0.05 above the cut-off price.

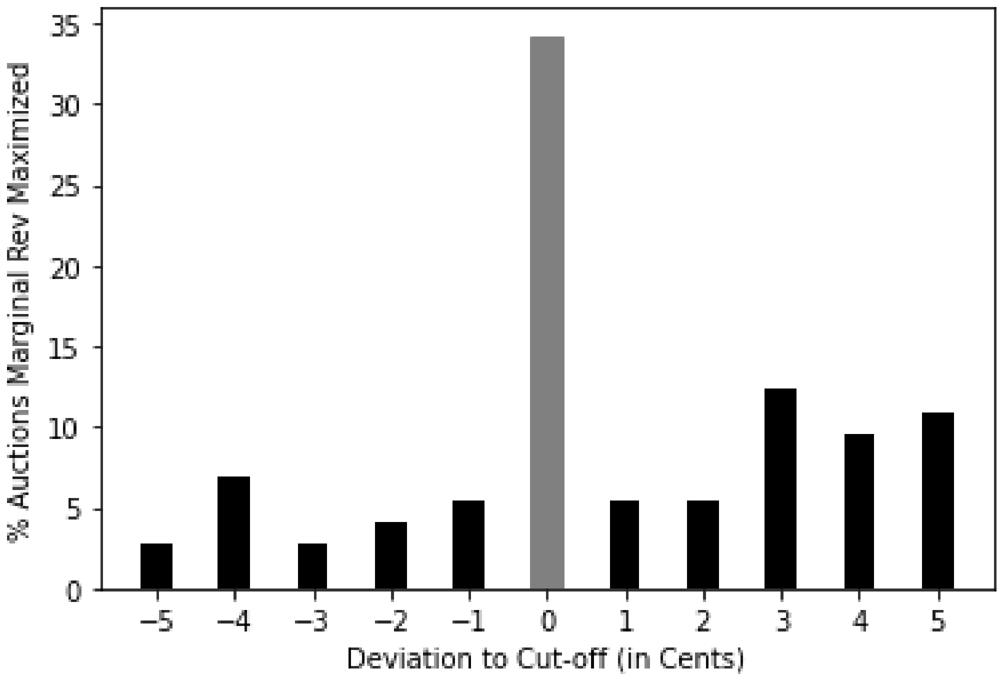

Figure 4. Proportion of bins with maximum marginal revenue around the cut-off price. For each auction and conditional on an interval of +/- €0.05 around the cut-off price, the figure plots the fraction of auctions whose marginal revenue is maximized at each price. Negative values to the left of the cut-off price indicate unsubscribed prices €-0.05 to €-0.01 below the cut-off price. Positive values to the right of the cut-off price indicate subscribed prices €0.01 to €0.05 above the cut-off price.

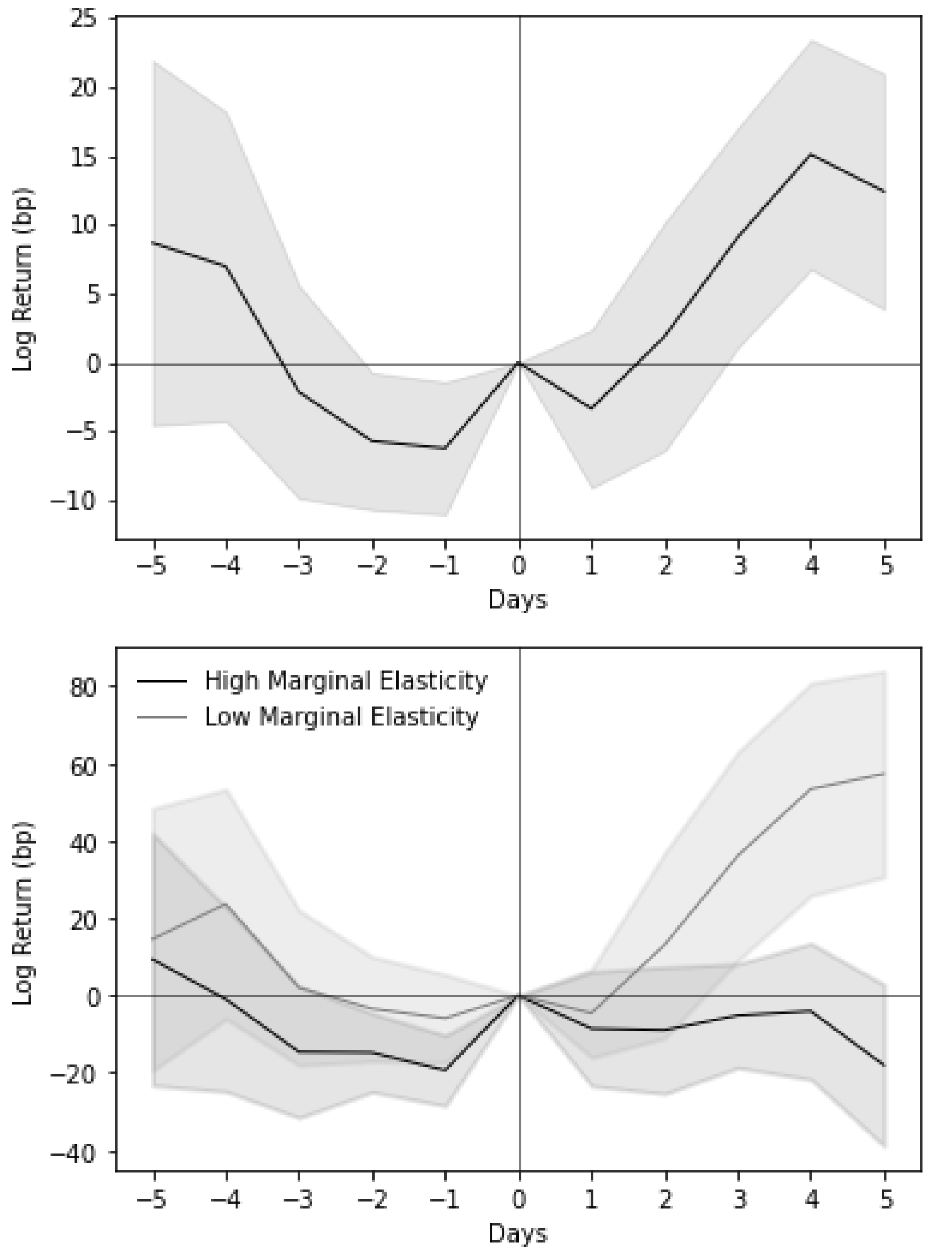

Figure 5. Log-returns around auction day.
The figure displays the average abnormal log-price returns (in bps) between 5 days prior and 5 days after the auction day. The returns are normalized to 0 at the close of auction day (day 0). 90% confidence bands are also reported. The top panel presents the results for all auctions. The bottom panel presents the results partitioning the auctions between low (black line) and high (gray line) marginal elasticity according to the median value of ME.

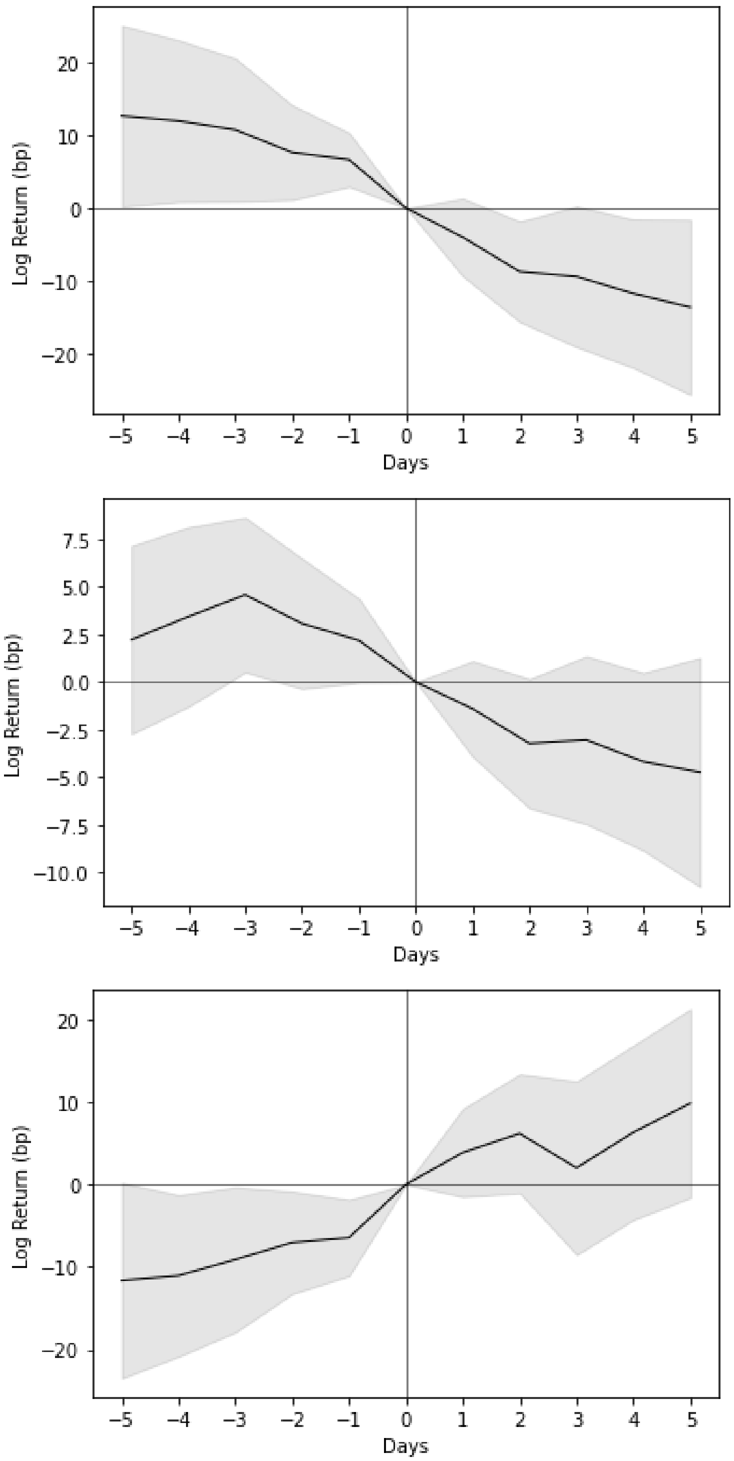

Figure 6. Log-returns around potential auction day.
The figure displays the average abnormal log-price returns (in bps) between 5 days prior and 5 days after the potential auction day. The returns are normalized to 0 at the close of auction day (day 0). 90% confidence bands are also reported. The adjusted returns for the nearest 2-year bond, 5-year bond and 10-year bond are used in the top, middle and bottom graphs, respectively.

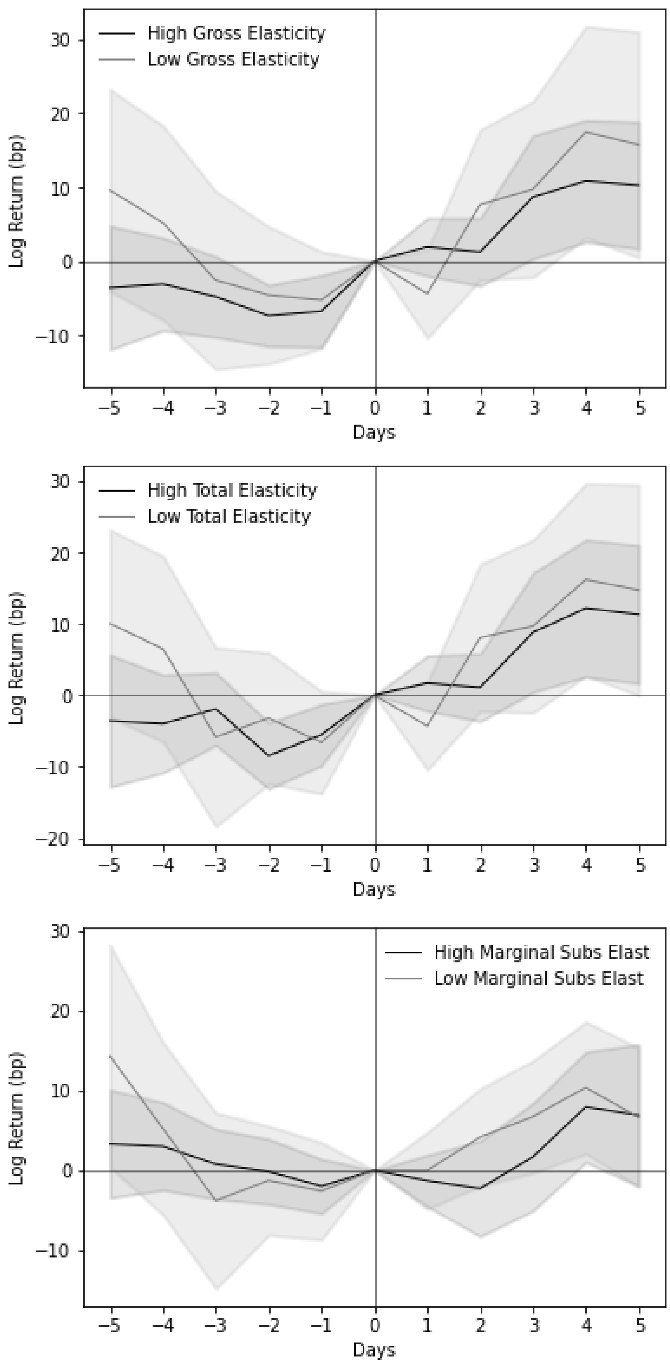

Figure 7. Log-returns around auction day.

The figure displays the average abnormal log-price returns (in bps) between 5 days prior and 5 days after the auction day. The returns are normalized to 0 at the close of auction day (day 0). 90% confidence bands are also reported. The top panel presents the results partitioning the auctions between low (black line) and high (gray line) gross elasticity according to the median value of GE. The middle panel presents the results partitioning the auctions between low (black line) and high (gray line) total elasticity according to the median value of TE. The bottom panel presents the results partitioning the auctions between low (black line) and high (gray line) marginal elasticity subscribed according to the median value of SE.

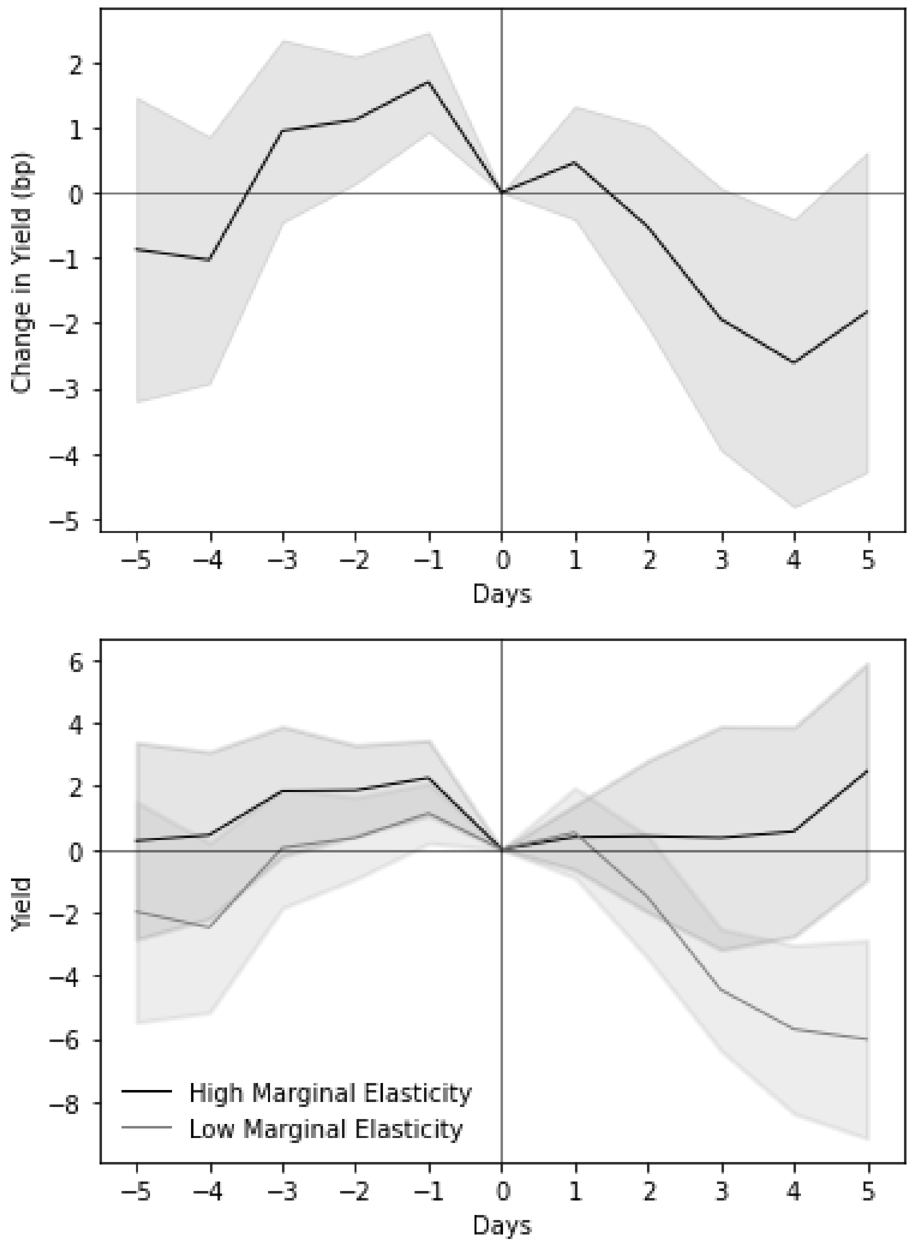

Figure 8. Yield changes around auction day.
The figure displays the average yield change (in bps) between 5 days prior and 5 days after the auction day. The yield changes are normalized to 0 at the close of auction day (day 0). 90% confidence bands are also reported. The top panel presents the results for all auctions. The bottom panel presents the results partitioning the auctions between low (black line) and high (gray line) marginal elasticity according to the median value of ME.

Table A.1. Variable definitions

Variable	Description
COVER	*Bid-to-cover ratio* is the total bid amount divided by the allocated amount. Known at the end of the auction. Source for bids: IGCP.
DRIFT	*Previous 5-day log abnormal return* is the excess log return of a period of 5 trading days until the day before the auction takes place and computed from the T-bond being auctioned over the return of the Portuguese Government Bond index for the same period. This variable is based on Lou et al. (2013). This variable is known before the auction begins. Source for prices: Bloomberg.
GE	*Gross elasticity* is the price-elasticity of demand using two points of the demand curve: the cut-off and the maximum prices and quantities. We compute this elasticity using for the slope the slope of the two points and as price and quantity the ones in the cut-off point. We then present the logarithm of this variable. This variable is known only at the end of the auction and is known by primary dealers at the end of the auction. Source for bids: IGCP.
ME	*Marginal elasticity* is the price-elasticity of demand using the cut-off price/quantity and the first 4 points with aggregated demand that was unsubscribed. We compute this elasticity using for the slope a regression of the five points prices on the quantities and as price and quantity the ones in the cut-off point. We then present the logarithm of this variable. This variable is known only at the end of the auction and is private information by the Treasury department. Source for bids: IGCP.
SE	*Marginal elasticity Subscribed* is the price-elasticity of demand using the cut-off price/quantity and the first 4 points with aggregated demand that was subscribed. We compute this elasticity using for the slope a regression of the five points prices on the quantities and as price and quantity the ones in the cut-off point. We then present the logarithm of this variable. This variable is known only at the end of the auction and is private information by the Treasury department. Source for bids: IGCP.
R_h	*Abnormal log-return from end-of-auction day until day h after the auction*, is the cumulative return of the bond being auctioned in excess of the cumulative return of the Portuguese bond index from the close of auction day until day n. Source for prices: Bloomberg.
RBAS	*Relative bid-ask spread* is the average of the ratios between the difference between ask and bid prices and the mid price over a period of 5 trading days until the day before the auction takes place and computed from the T-bond being auctioned. This variable is known before the auction begins. Source for prices: Bloomberg.
SIZE	*Size* is the allocated amount in the auction (in EUR million). Source: IGCP.
SPREAD	*Spread* is the spread between 10-year Portuguese government bond yield and German government bond yields (in basis points) in the previous day to the auction taking place. This variable is known before the auction begins. Source for time-series: Bloomberg.
TE	*Total elasticity* is the price-elasticity of demand using all the bids. We compute this elasticity by using all points by running a regression of prices on quantities and as price and quantity the ones in the cut-off point. We then present the logarithm of this variable. This variable is known only at the end of the auction and is private information by the Treasury department. Source for bids: IGCP.
VOL	*Volatility* is the standard deviation of log returns of the bond being auctioned over the 20 days until the day before the auction takes place. This variable is known before the auction begins. Source for prices: Bloomberg.