

# A Debt Crisis with Strategic Investors: Changes in Demand and the Role of Market Power

Ricardo Alves Monteiro

WP/25/19

**IMF Working Papers** describe research in progress by the author(s) and are published to elicit comments and to encourage debate.

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2025  
JAN**



**IMF Working Paper**

Institute for Capacity Development

**A Debt Crisis with Strategic Investors: Changes in Demand and the Role of Market Power**  
**Prepared by Ricardo Alves Monteiro\***Authorized for distribution by Ali Alich  
January 2025

**IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**ABSTRACT:** This paper documents changes in investors' demand for sovereign debt during a debt crisis. Using a dataset containing individual bids on Portuguese debt auctions, I document that bid functions become more inelastic during the crisis. That is, investors require bigger drops in price to buy additional units of debt, increasing the government's marginal cost of issuing debt. I then decompose these changes in demand into two components: a fundamental component, due to changes in the valuation of the security, and a strategic component, that arises from investors' market power. I find that, although the role of market power is negligible in normal times, it gets more pronounced leading up to and during the crisis. The government is not able to extract the full surplus from strategic investors, and, as a result, the auction mechanism loses efficiency during that period. Finally, I discuss a possible mitigation strategy. Everything else constant, the use of shorter maturities could avoid higher inefficiency costs.

JEL Classification Numbers:	D44, E43, F34, G12, H63
Keywords:	Debt auctions; default risk; demand elasticity; market power
Author's E-Mail Address:	<a href="mailto:ralvesmonteiro@imf.org">ralvesmonteiro@imf.org</a>

\* I thank Agência de Gestão da Tesouraria e da Dívida Pública - IGCP, E.P.E. for access to the dataset. I am grateful to my advisors Manuel Amador and Tim Kehoe for their support and insightful discussions. I thank Cristina Arellano, Marco Bassetto, Tamon Asonuma, Daniel Belchior, José Cardoso da Costa, Cristina Casalinho, Doireann Fitzgerald, Stelios Fourakis, William Jungerman, Leonardo Martinez, Francisco Roldan, and participants at the Minnesota-Wisconsin International Macro Workshop, and at the International Monetary Fund Institute for Capacity Development for excellent comments. All remaining errors are my own.

WORKING PAPERS

# **A Debt Crisis with Strategic Investors: Changes in Demand and the Role of Market Power**

Prepared by Ricardo Alves Monteiro<sup>1</sup>

---

<sup>1</sup> I thank Agência de Gestão da Tesouraria e da Dívida Pública - IGCP, E.P.E. for access to the dataset. I am grateful to my advisors Manuel Amador and Tim Kehoe for their support and insightful discussions. . I thank Cristina Arellano, Marco Bassetto, Tamon Asonuma, Daniel Belchior, José Cardoso da Costa, Cristina Casalinho, Doireann Fitzgerald, Stelios Fourakis, William Jungerman, Leonardo Martinez, Francisco Roldan, and participants at the Minnesota-Wisconsin International Macro Workshop, and at the International Monetary Fund Institute for Capacity Development for excellent comments. All remaining errors are my own.

# 1 Introduction

Governments around the world maintain enormous stocks of sovereign debt. Most of this debt is issued in auctions and only then traded in the secondary market. In sovereign debt auctions, investors submit bids consisting of the highest price they are willing to pay to purchase a unit of debt, and how much they are willing to buy. Then, the government chooses which bids to accept. It follows that the government effectively takes investor's aggregate bid function as given and chooses its preferred price quantity pair. The government is a monopolist in sovereign debt auctions.

The elasticity of demand is a key statistic for a monopolist as it pins down its optimal decision. Demand, however, is rarely observed by the researcher, let alone its elasticity. In this paper, I use a proprietary dataset containing all individual bids submitted on Portuguese sovereign debt auctions from 2003 to 2020 to estimate the elasticity of demand. The time-series includes the European sovereign debt crisis of 2010-2014, providing a unique opportunity to assess how demand and the price elasticity of demand change during high default risk events. The crisis period starts in 2010, after the newly elected Greek government disclosed that its budget deficits were far higher than previously thought, this led to the downgrade of Greek debt and to a sharp increase in spreads. The focus soon followed to Portugal that by April 2011, after a bailout from the International Monetary Fund and European institutions, lost part of its market access. This period lasted through 2013, and by 2014 Portugal had a full return to normal, pre-crisis, procedures.

The Portuguese agency that issues sovereign debt mentions changes in demand around the debt crisis as the explanation for lower amounts being issued during the period. Analyzing the difference, at the auction level, between the average price and the marginal price, the lowest price accepted, is informative. Figure 1 depicts this difference as the spread between the marginal yield<sup>1</sup>, the maximum yield accepted by the government, and the average of the yields bid in the auction, weighted by the amount bid, for 12 month Treasury Bills. In normal times, the average and marginal prices are very close, and spreads are essentially zero. During the crisis, however, the marginal price of the auction is lower than the average price. Conversely, marginal yields are higher than average yields. This difference between prices highlights why it is misleading to use average yields as a tool to assess the cost of issuing an additional unit of debt. As pointed by [Aguilar and Amador \(2021\)](#), the yield curve does not reflect the marginal cost of borrowing, it is the elasticity of the bond price with respect to the government policy – the inverse of the elasticity of demand – that determines the marginal cost of issuing a given security. It tells us how big of a drop in price (jump in yield) investors require for the government to issue a larger amount of debt.

I find that the bid functions for both short and long maturity securities get significantly more inelastic leading up to and during the sovereign debt crisis. That is, to increase the amount of debt issued by 1%, the price needs to decrease, in percentage terms, by more than it had before the crisis. Particularly, I find that the inverse price elasticity for Treasury Bills is, on average, thirteen times higher leading up to and during the crisis, from 0.012 to

---

<sup>1</sup> A yield is the interest rate required by investors, such that the present value of the claim – one unit promises to pay 1 euro at maturity – is consistent with the price submitted in the investor's bid.

0.15 basis points. As for Treasury Bonds, I find that the same measure of elasticity is, on average, 26% higher in the lead up to the crisis period<sup>2</sup>, from 0.29 to 0.36 basis points.

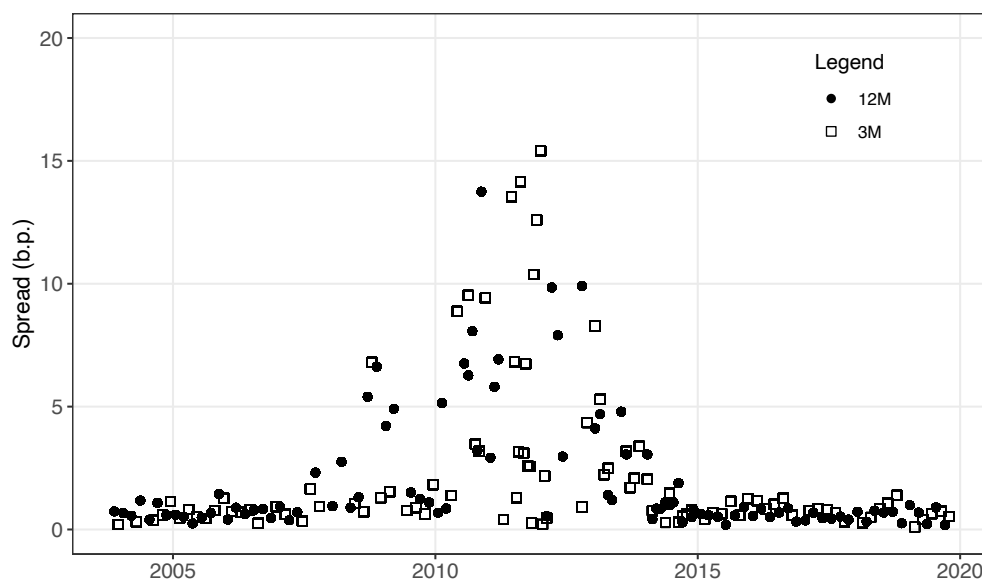


Figure 1: Spread between Marginal and Average Yields of Treasury Bills

Primary dealership models are widely used by debt management offices. In these, only a limited number of authorized dealers – the primary dealers – participates in the auctions, and then act as market makers by selling those securities in the secondary market. As an example, in Europe, according to [AFME \(2020\)](#), at least 20 countries use a primary dealership model.

A small number of participants suggests the existence of strategic behavior from investors. In particular, a single bid influences the price that clears the auction and investors internalize that effect. As a result, bids submitted might differ from the dealers' willingness to pay. This wedge, between bids and valuations, is a direct consequence of dealer's market power. This is in contrast with the competitive nature of secondary markets for debt, open to a much larger number of investors.

The non-competitive nature of the market and investors' strategic bidding motivates the decomposition exercise that follows. How much of the observed shifts in bid functions are due to shifts in the valuation of the asset and how much of these shifts are due to the market power of investors and their strategic decisions? Is this decomposition of bids constant across maturities?

The goal of the exercise is to have a better understanding of how investors' demand for sovereign debt with different maturities evolves around high default risk events, while considering the non-competitive nature of the market. With this decomposition, it is possible to estimate what is the elasticity of the actual investors' willingness to pay.

<sup>2</sup>Note that Treasury Bonds stopped being issued in mid 2011 around the bailout. This explains why the increase is not as noticeable for Treasury Bonds.

To filter the data, I introduce an environment based on [Wilson \(1979\)](#) framework, and more specifically on [Hortaçsu and McAdams \(2010\)](#) and [Kastl \(2011b\)](#). The auction model uses a discriminatory price protocol – pay-as-bid – and treats all investors as identical *ex-ante*. Investors differ *ex-post* on the realization of the idiosyncratic private signals regarding the security being auctioned. Given the private realization of their signal, as well as their subjective expectation of the aggregate state, each investor then submits a discrete bid function that maximizes their expected utility.

In the dataset, I observe the equilibrium object, the discrete bid functions. Through a necessary equilibrium condition, I then back out the primitive, investors' true valuations. With both investor's valuations and bids, I assess how the wedge between the two evolves around the crisis. This wedge represents investor's market power: bidding below the valuation is possible as investors internalize that a single bid can influence the clearing price.

I find that market power plays a limited role during normal times. However, leading up to and during the crisis the wedge between bids and valuation gets more pronounced. The mechanism follows. In normal times, the bid schedule is mostly flat and as such, investors have no room to exercise market power and are price takers for all intents and purposes. Leading up to and during the crisis, bids get more dispersed and bid schedules become steeper. The increased dispersion of bids implies that the subjective distributions of the price that clears the auction are less precise than before the crisis. Particularly, the likelihood ratio of bid  $k + 1$  in a bid function being a winning bid, relative to, bid  $k$  being the last winning bid, increases. This leads to a larger wedge between bids and valuations as investors bid below value to avoid the winner's curse.

A consequence of the larger wedge between bids and valuations is that the auction mechanism becomes less efficient during the crisis. That is, the government is not able to extract the full surplus from investors as they are bidding below their willingness to pay. Let the inefficiency be measured as the ratio of the aggregate wedge over the amount issued in each auction. I find that, at their peak, inefficiency costs go up to 0.6% of the issued amount, during the crisis.

Finally, a more normative analysis should follow. Particularly, what can the government do to mitigate these inefficiency costs when issuing debt during a crisis? I briefly look at maturity choice as a mitigation device, as different securities face different wedges but also need to be rolled over at different frequencies. A more thorough analysis of optimal maturity choice accounting for the inefficiency costs of the mechanism is left as future research.

The remainder of the paper is organized as follows: section 2 presents a literature review; section 3 introduces the data while providing relevant institutional background and evidence for changes in demand leading up to and during the crisis; section 4 presents the model used to filter the data and back out investors' valuations of the securities being auctioned; section 5 discusses the estimation procedure; section 6 discusses the role of market power and presents the rise of inefficiency costs leading up to and during the crisis; section 7 discusses a possible mitigation strategy from the government through maturity choice; section 8 concludes.

## 2 Literature Review

This paper builds on the sovereign debt and default literature, with an emphasis on the auction framework used to issue debt and investors' market power. Related papers here include [Cole et al. \(2021\)](#) and [Bigio et al. \(2021\)](#). Each uses data for sovereign debt auctions for other countries. The motivation of [Cole et al. \(2021\)](#) is similar, the authors present a model that focus on investors' choices in an auction setting with information heterogeneity. However, in their sample there are no meaningful high default risk episodes, and while the authors focus on matching moments and patterns in micro data for Mexico, they do not tackle the changes in demand during a high default risk event and the role that market power plays on the evolution of bids. [Bigio et al. \(2021\)](#) uses micro data on sovereign debt auctions from Spain to assess liquidity costs. Their focus is on optimal debt-maturity management in the presence of such costs. I also discuss maturity choice but as a mitigation strategy to the inefficiency costs created by investors' market power. Importantly, with this data encompassing the high default risk event, I can to present estimates for the price elasticity of the aggregate bid functions and evaluate how this elasticity evolves around the crisis.

This paper also relates to those estimating elasticities of demand for sovereign debt. Related papers include [Albuquerque et al. \(2022\)](#) and [Moretti et al. \(2024\)](#). [Albuquerque et al. \(2022\)](#) uses bid level data for Portuguese sovereign debt auctions to estimate the elasticity of demand and assess its predictive power for same-bond post-auction returns in the secondary market. To do so they focus on uniform price auctions after the sovereign debt crisis. This paper estimates elasticity and analysis its evolution around the sovereign debt crisis. It then uses this information to better understand the implications of default risk for strategic bidding. [Moretti et al. \(2024\)](#) estimates the elasticity of demand for sovereign debt in the secondary market. It then incorporates the inelastic demand into a sovereign debt model to assess its impact on government's supply of bonds and default risk. In contrast, this paper estimates the elasticity of demand for sovereign debt in the primary markets. As noted before, the secondary market price denotes the average price of debt, which can differ substantially from the marginal price of debt at the auction. The latter is what determines how much debt the government issues at the auction.

In terms of methodology, the auction model used to filter the data is based on [Hortaçsu \(2002\)](#), and [Kastl \(2011b\)](#). [Hortaçsu \(2002\)](#) presents a model based on [Wilson \(1979\)](#) of a multi-unit discriminatory price auction with a finite number of symmetric risk-neutral bidders with independent private values. In this model, a single bid affects the bid functions through changes in the distribution of the price that clears the auction. They construct a non-parametric estimator of the distribution exploiting a re-sampling technique. [Kastl \(2011a\)](#) builds on this framework by allowing for discrete-step bid functions. [Kastl \(2020\)](#) provides a review of the literature and methods applied to financial auctions and particularly to treasury bond auctions. This paper uses the tools developed by the authors mentioned and focus instead on the impact that a high default risk event has on dealers' bidding patterns and, particularly, to bid shading over time.

This paper also relates to the quantitative sovereign debt literature. The work that started with [Aguilar and Gopinath \(2006\)](#) and [Arellano \(2008\)](#), based on the classic setting of [Eaton and Gersovitz \(1981\)](#), focus on

sovereign default as the outcome of the government's financing problem provided there are competitive investors that are willing to lend as long as they break even<sup>3</sup>. Since those initial quantitative models, there has been substantial developments in the literature with the study of maturity choice and self-fulfilling crisis, to name a few. Examples of such are [Arellano and Ramanarayanan \(2012\)](#) and quantitative models based on [Cole and Kehoe \(2000\)](#), such as [Bocola and Dovis \(2019\)](#). Finally, [Aguiar et al. \(2019\)](#) points out that the price elasticity of demand is the crucial element that determines how much the government borrows and whether it prefers to borrow long-term debt or short-term debt. This paper provides estimates for the price elasticity of demand and aims to offer a broader understanding of investors' strategic considerations in the context of the primary market for sovereign debt and its interaction with default risk.

This paper also relates to the strand of literature on sovereign debt management and maturity choice. Previous work on the topic include [Arellano and Ramanarayanan \(2012\)](#), [Broner et al. \(2013\)](#), [Sánchez et al. \(2018\)](#) and [Aguiar et al. \(2019\)](#). In this paper, I present the inefficiency costs induced by the auction mechanism for different maturities. The fact that inefficiency costs vary across maturities introduces a new margin to be considered for maturity choice.

Finally, the paper is related to [Alves Monteiro and Fourakis \(2023\)](#) in that it shares the same dataset and studies the crisis period in Portugal. This other paper however, has a different goal. Instead of focusing on the game between investors in an auction when they exercise market power given a random bond supply, it focus on the strategic interaction between the set of investors and a government that also optimizes.

### 3 Data: Background and Evidence

Auction data was provided by the *Portuguese Treasury and Debt Management Agency* (IGCP, Portuguese acronym). The data comprises all auctions of Treasury Bills (short maturities) and Treasury Bonds (long maturities) held from 2003 and 2004, respectively, and up to 2020. As such, the time series includes the sovereign debt crisis of 2010-2014, which enables the analysis of changes in demand during that period. Importantly, the data comprises all individual bids (price and amount) that were placed in each auction, even if they were not executed.

Issuance of Treasury Bills in the primary market is done through auctions. Treasury Bonds are launched for the first time in syndicated operations<sup>4</sup>. New issuances of a line that has already been launched are done through auctions. Both types of securities were auctioned using a discriminatory price protocol, where investors *pay-as-bid*, up to 2011. From 2014 onward, Treasury Bonds were auctioned using a uniform price protocol, where bids are executed at the marginal price of the auction. For a thorough analysis of the impact of the auction protocol on the outcomes of the auctions refer to [Alves Monteiro and Fourakis \(2023\)](#).

The IGCP uses a primary dealership model to issue bills and bonds. Only primary dealers, a group of financial

---

<sup>3</sup>See [Aguiar and Amador \(2014\)](#) for a survey.

<sup>4</sup>A syndicate is a group of banks that is given the mandate to place a specific amount of government bonds. It follows a book building process that allows for permanently monitoring of orders and intervention in the allocation of such orders by the IGCP.

intermediaries, participate in the auctions. Dealers are permitted to submit multiple bids<sup>5</sup> provided that the total value does not exceed the upper limit of the overall amount announced for the auction.

An auction is as follows: i) the government announces an auction and the characteristics of the security being auctioned, as well as a target for the size of the issuance; ii) the auction takes place and investors submit bids that consist of a price and amount pair; iii) the auction closes and the government orders bids in descending order of price; iv) the government chooses the minimum price it is willing to accept, determining the size of the issuance; v) bids above the minimum price are executed and investors pay either the minimum price (in a uniform price auction) or the price they bid (in a discriminatory price auction).

Table 1 presents some summary data for the most common bill and bond auctions. One can observe 400 Treasury Bill auctions and 161 Treasury Bond auctions. The most common maturities are 12 and 3 months for the bills and 10 and 5 years for bonds. In bill auctions the number of bids averages 39 and in bond auctions it averages 56. Dealers (mean) refer to the average number of dealers present in the auctions of each security. Steps (mean) refer to the average number of bids submitted by a single dealer. Issued (mean, M€) refer to the average amount issued by the IGCP in auctions of each type of security.

Table 1: Summary Data on Treasury Bond and Bill auctions

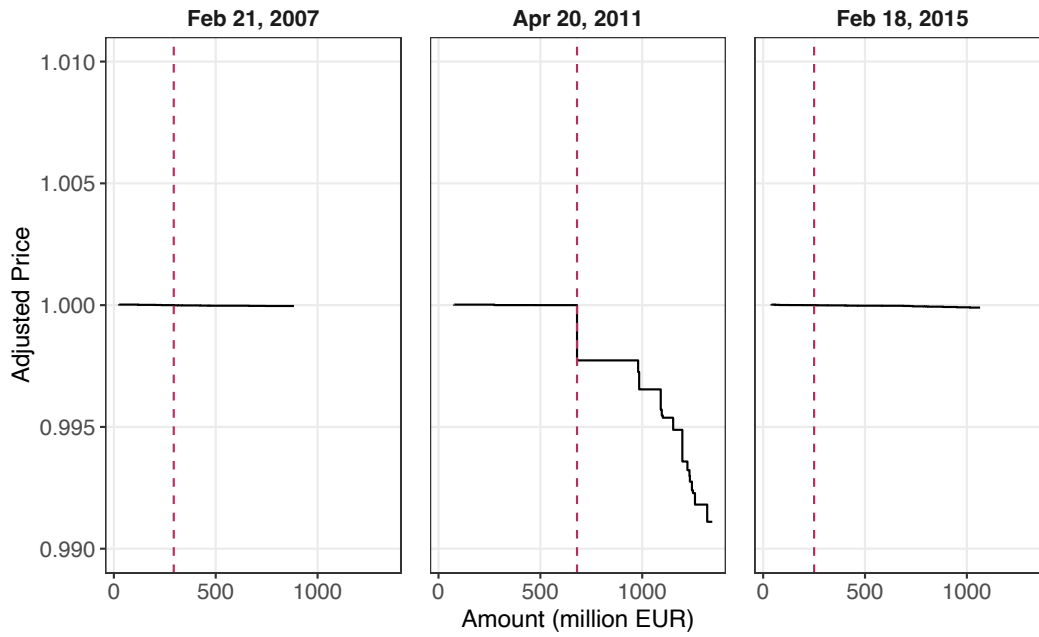
Security	Auctions	Bids (mean)	Dealers (mean)	Steps (mean)	Issued (mean, M€)
3 Months	101	35.2	14.5	2.4	471.0
6 Months	88	36.4	14.7	2.4	505.6
12 Months	101	44.0	15.4	2.8	1,037.5
All Bills	400	38.7	14.8	2.5	703.1
5 Years	21	55.9	18.9	2.8	732.3
6 Years	14	56.5	18.2	3.0	754.1
10 Years	52	59.1	17.9	3.2	805.8
All Bonds	161	56.4	17.9	3.0	756.0

### 3.1 Changes in Demand

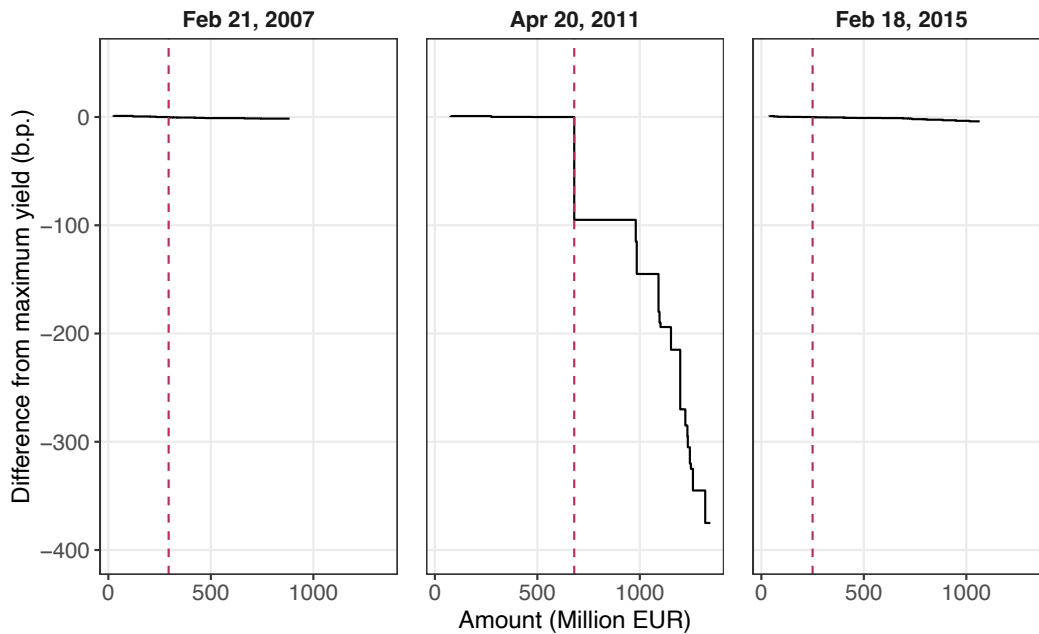
Below I show evidence that motivates the shift in investors' demand for Portuguese sovereign debt while approaching and during the sovereign debt crisis. Figures 2 and 3, respectively, present the aggregate bid functions (downward step functions) and the amounts issued (dashed line) by the Portuguese Government in 3-month and 12-month treasury bill auctions over time. In panel (a), the aggregate bid functions are presented as price and amount pairs: the amount that the government is able to raise at each given price. Panel (b) presents an alternative representation of the aggregate bid function: the yield required by investors for borrowing each amount to the government. The aggregate bid function is obtained by aggregating individual bids.

<sup>5</sup>For Treasury Bill auctions each dealer may submit up to five bids per auction, for Treasury Bond auctions a limit is not specified.

The analysis focus on the crisis event. The figures present a representative auction before, during and after the sovereign debt crisis<sup>6</sup>.



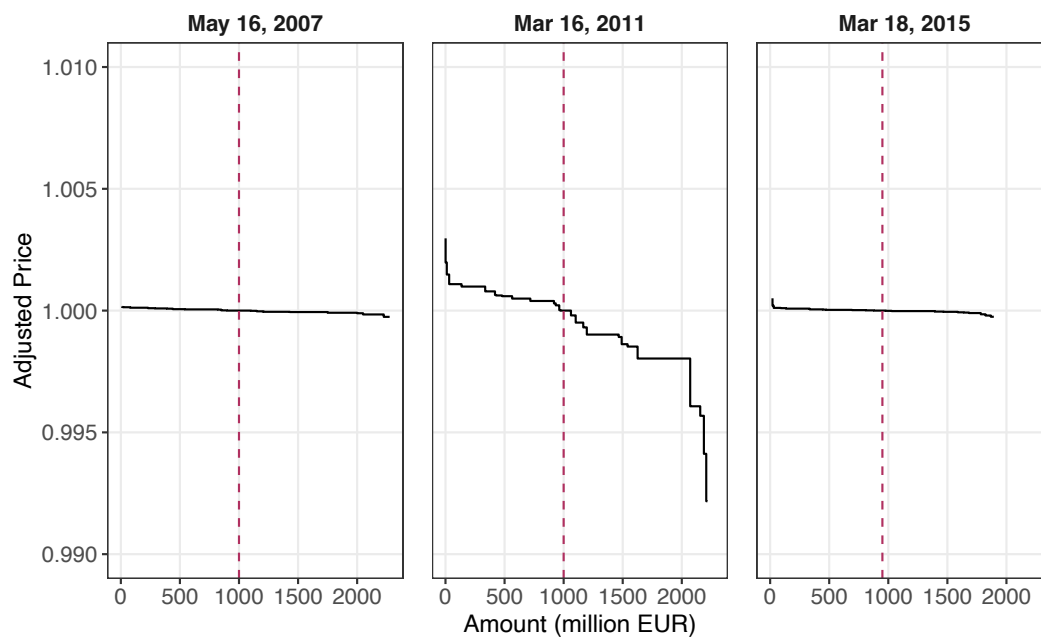
(a) Changes in Demand (Prices)



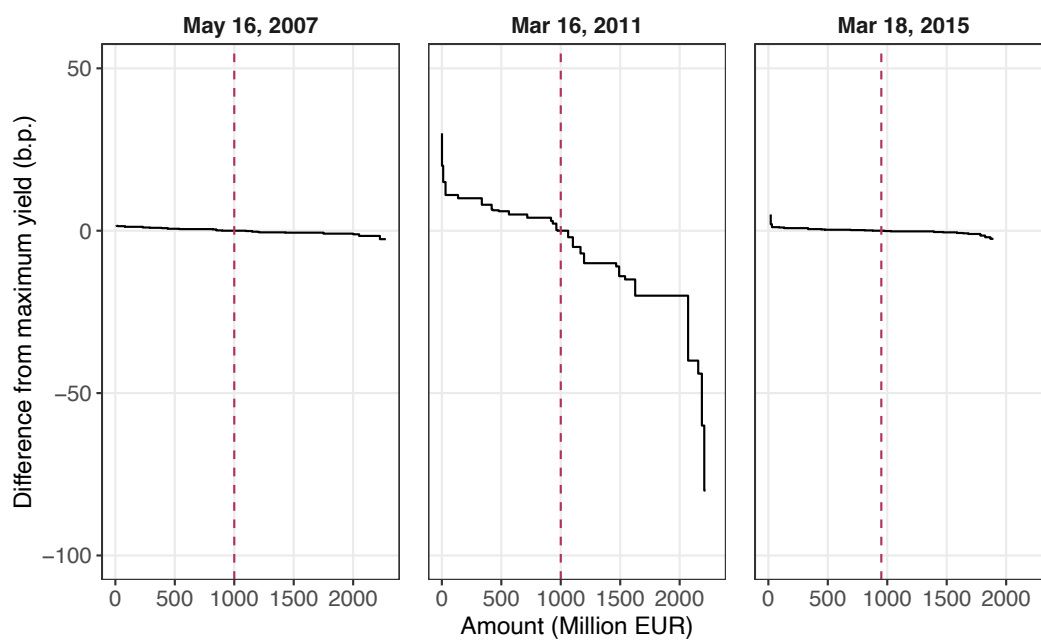
(b) Changes in Demand (Yields)

Figure 2: Aggregate bid functions for 3 month treasury bills

<sup>6</sup>In the appendix I present a sequence of all auctions from 2007 to 2016. Exact dates differ depending on the maturity being issued.



(a) Changes in Demand (Prices)



(b) Changes in Demand (Yields)

Figure 3: Aggregate bid functions for 12 month treasury bills

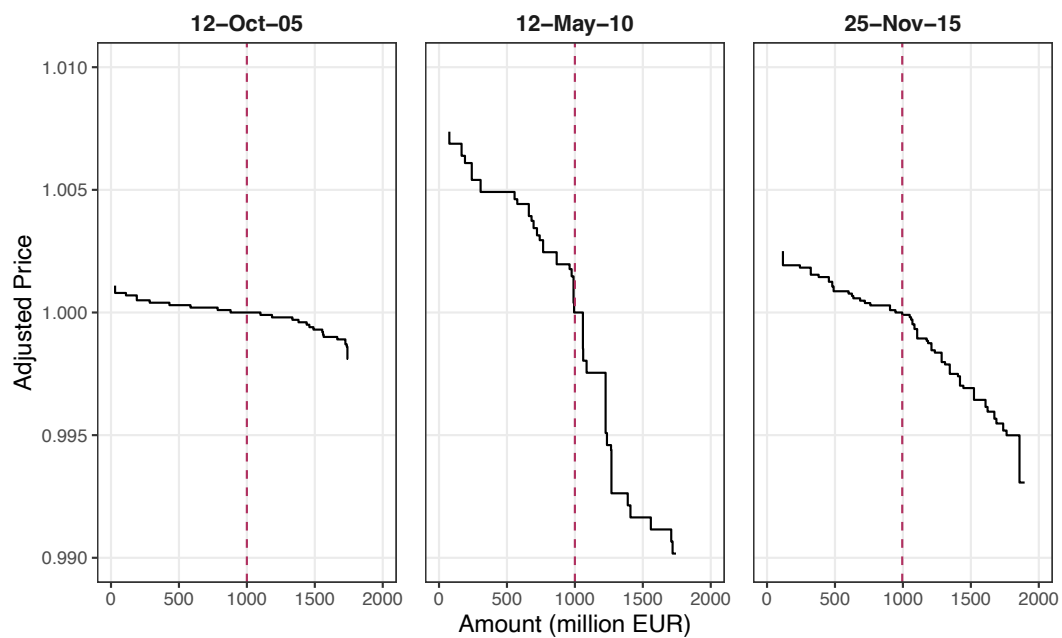
In panel (a), prices are normalized by the marginal price of the auction, i.e. the minimum price accepted by the government. That is, bids with an adjusted price above 1 are executed and those with an adjusted price below 1 are not executed. Panel (b) depicts the difference, in basis points, between the maximum annualized yield accepted by the government in the auction and the annualized yields associated with each submitted bid. It follows that bids with a positive difference in yields are executed and those with yields above the maximum yield accepted (and so, a negative difference) are not executed. Finally, the maroon dashed line identifies the amount of debt issued in each auction.

During the period before the crisis, represented by the auctions in 2007, there is almost no dispersion across the prices of individual bids. That is, investors submit similarly flat bid functions. During the crisis period, represented by the auctions in 2011, this is no longer the case. Particularly, there are several bids with prices significantly below the auction price, and, consequently, requiring yields far above the maximum accepted yield. Note that a bid with a price 1% below the marginal price in the auction, requires an annual yield 400 basis points or 100 basis points above the maximum accepted yield, for 3 month and 12 month bills, respectively. These are meaningful differences. After the crisis period, represented by the auctions in 2015, there is a recovery of the aggregate bid function to its previous shape, with almost no dispersion in the prices bid.

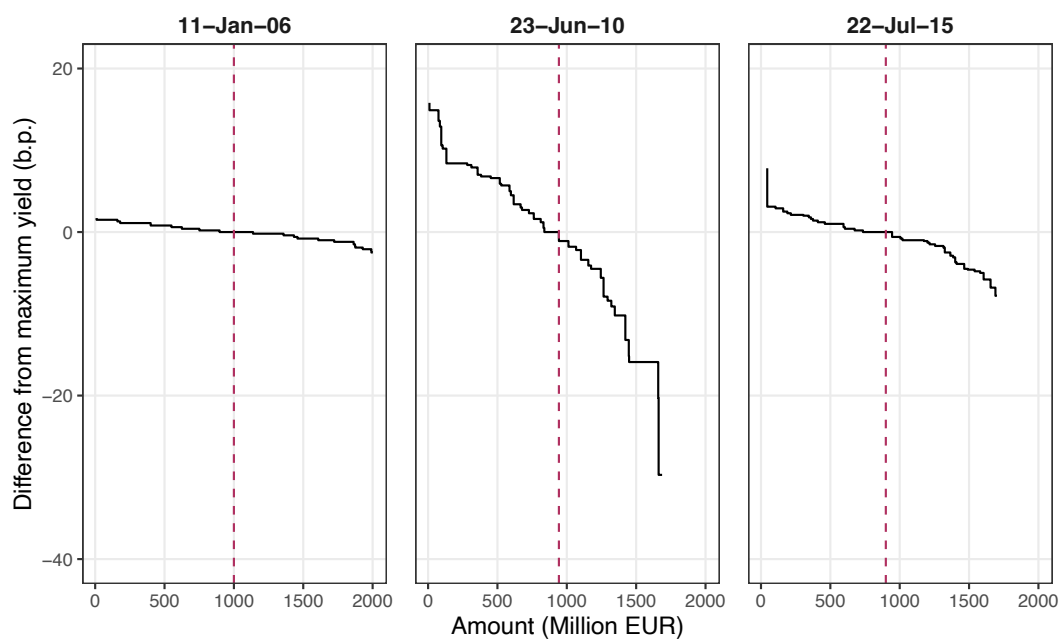
Figures 4 and 5 present the aggregate bid functions and the amounts issued by the Portuguese Government in 5 year and 10 year treasury bond auctions over time, respectively. As before, the analysis focus on the crisis event so the figures show us the evolution of the bid functions leading up to the sovereign crisis and the recovery period afterwards. The elements within the figures are as in Figures 2 and 3 above.

As described for the short maturities, the aggregate bid functions for long maturity debt become steeper leading up to and during the crisis period. It is worth noting, however, that, in contrast to what happens with treasury bills, the bid functions for treasury bonds do not recover to their pre-crisis shape. A potential explanation for this behavior is the fact that the auction protocol for treasury bonds switched from a discriminatory price protocol to a uniform price protocol after the crisis.

It is a well known that the winner's curse is a potential outcome of a discriminatory price auction. In particular, if dealers pay-as-bid then there is an incentive to bid lower prices. By moving from a discriminatory price to a uniform price protocol, bidders are less likely to shade their bids as they will end up paying the marginal price of the auction regardless. This is consistent with steeper bid functions under a uniform price protocol. The effect of the auction protocol on investors' bids is studied in detail in [Alves Monteiro and Fourakis \(2023\)](#).

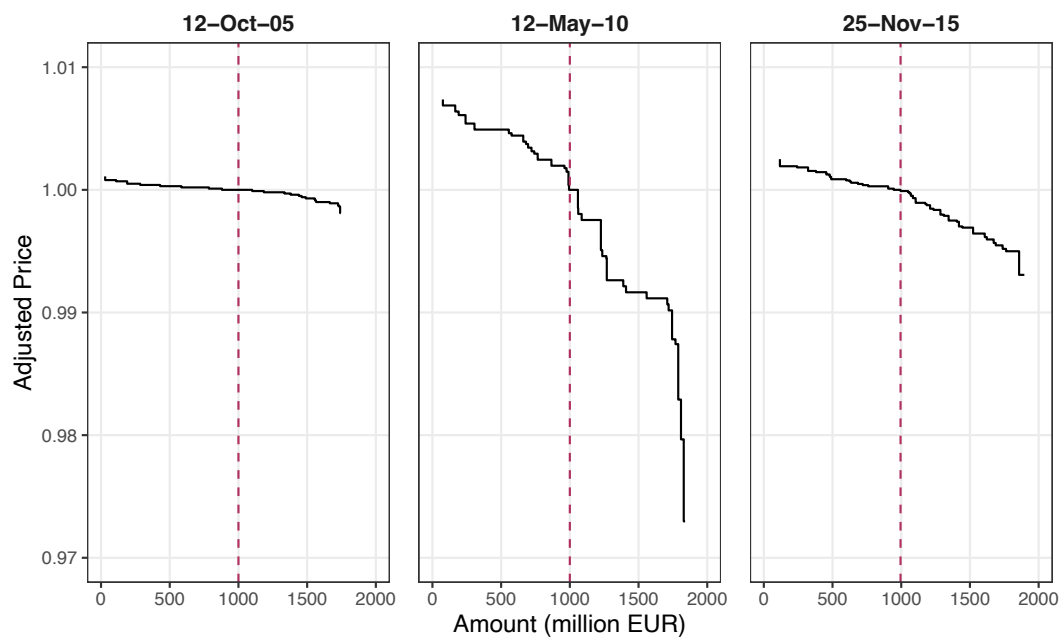


(a) Change in Demand (Prices)

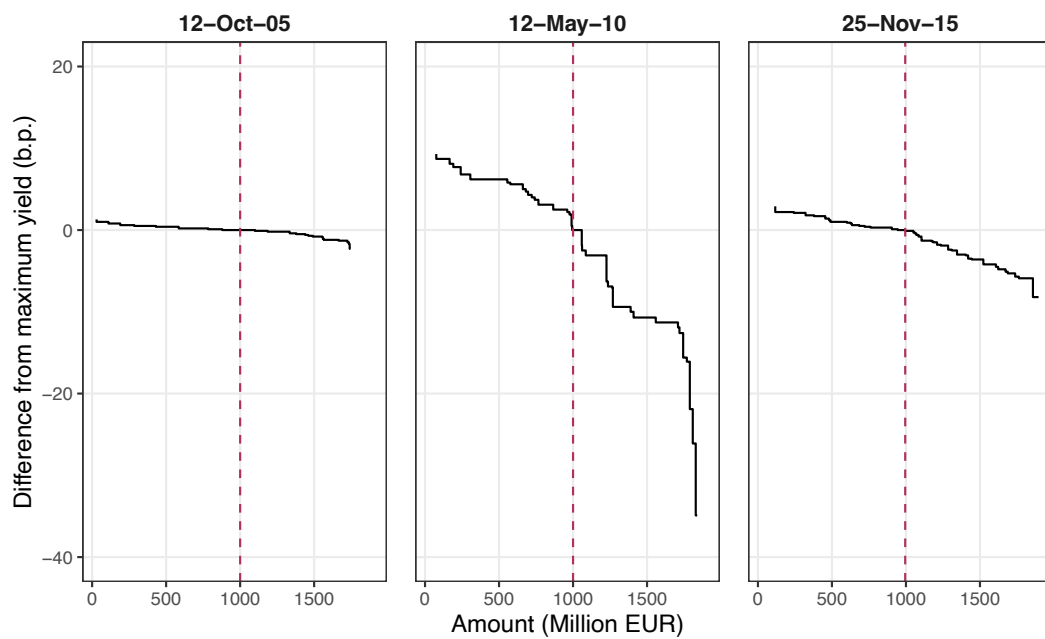


(b) Change in Demand (Yields)

Figure 4: Bids schedule for 5 year treasury bonds in the primary market



(a) Change in Demand (Prices)



(b) Change in Demand (Yields)

Figure 5: Bids schedule for 10 year treasury bonds in the primary market

**Discussion.** As mentioned before, with this data, one can estimate the inverse price elasticity, i.e. the percentage change in price such that the amount issued by the government increases by 1%, and assess how it changes during the same period. Investors' bid functions are indeed changing leading up to and during high default risk events. As the sovereign debt crisis approaches, the aggregate bid function tends to become more inelastic in the sense that, at a given pair (price, amount issued), in order to increase the amount issued by 1% on average, the price needs to decrease, in percentage terms, by more than it had before the crisis.

This increase in the inverse elasticity happens for both short and long maturities. However, the change itself is not homogeneous across all maturities. Moreover, after the crisis period the elasticity tends to correct to its previous levels, particularly for shorter maturities. For longer maturities, it is important to mention that the first auction after the crisis was executed under a uniform price protocol, whereas all auctions before the crisis were executed under a discriminatory price protocol. As a result, one cannot disentangle the recovery from the crisis from the change in auction protocol<sup>7</sup>. One can argue that the described change in the auction mechanism would lead to a smaller wedge between the bids and valuations as there is no winner's curse as in the discriminatory price protocol. This in turn leads to potentially steeper bid functions, i.e. with less shading on the first steps. This argument can help us understand the apparent lack of recovery of the bid functions to their previous shape.

## 3.2 Elasticity Measures

Let  $p$  be the marginal price and  $B$  the amount of debt issued in an auction. Then, the demand elasticity is defined as  $\mathcal{E} = -\frac{\partial B}{\partial p} \frac{p}{B}$ . Throughout, I will report the inverse of the demand elasticity,  $\eta = -\frac{\partial p}{\partial B} \frac{B}{p}$ , the necessary change in price such that the amount issued increases by 1%. A small (absolute) value of this inverse elasticity means that large increases in the amount issued are associated with small decreases in price.

In order to compute the inverse elasticity we need to estimate the slope coefficient  $\frac{\partial p}{\partial B}$ . I first compute the inverse marginal elasticity (ME) of an auction, the main measure of elasticity used throughout this paper. I follow [Albuquerque et al. \(2022\)](#) and use bids from untapped liquidity, next to the marginal price of the auction. More precisely, I use the four price points from unsubscribed bids next to the cut-off price, together with the cut-off price point itself. After constructing the aggregate bid function, adding up all individual bids, I use the quantity price pairs to estimate a linear regression model of the price on the amount issued and a constant. The slope coefficient in the model is an estimate of  $\frac{\partial p}{\partial B}$ . To get the elasticity one just multiplies the slope estimate by the ratio of the marginal price of the auction to the amount issued. I also compute the inverse total elasticity (TE), that differs from ME in that it uses all the bids to estimate the slope coefficient from a simple linear regression model.

These different elasticity measures provide different information. The marginal elasticity provides an estimate of the elasticity around the marginal price of the auction considering the approximate slope in that region. As such, it is more informative than the total elasticity to assess the increase in cost needed for a larger issuance

<sup>7</sup>During the crisis there were no issuances of treasury bonds (only treasury bills). As such, the first auction after the crisis, in 2014, was also the first with the uniform price protocol.

around the observed issuance amount. The total elasticity, however, provides a more comprehensive picture of whether there was a shift in demand as it is uniquely driven by bidder behaviour. This turns out to be relevant, particularly leading up and during the crisis, where the aggregate bid functions tend to present a quasi-kink – they are relatively flat at first and then abruptly get steeper –, that distorts the estimate around it. Whether or not the marginal price is close to this drop in the aggregate bid function is captured by the difference between ME and TE, as the latter uses a slope coefficient estimated with all bids, diluting the quasi-kink. If the government chooses a marginal price close to the quasi-kink, then untapped liquidity will have much lower prices and ME is likely higher, in absolute value, than TE. If, however, the government avoids the quasi-kink and there is untapped liquidity at a relatively flat price then ME is likely lower, in absolute value, than TE.

As described earlier, the government chooses the size of the issuance and the respective marginal price of the auction, given the aggregate bid function. The difference between ME and TE informs whether the government avoided the steep part of the bid function or ventured close to it, shedding light on the need for funds at a given auction. While TE is pinned down by bidder behavior, ME is in part determined by the government optimal and discretionary decision, after observing the bid function.

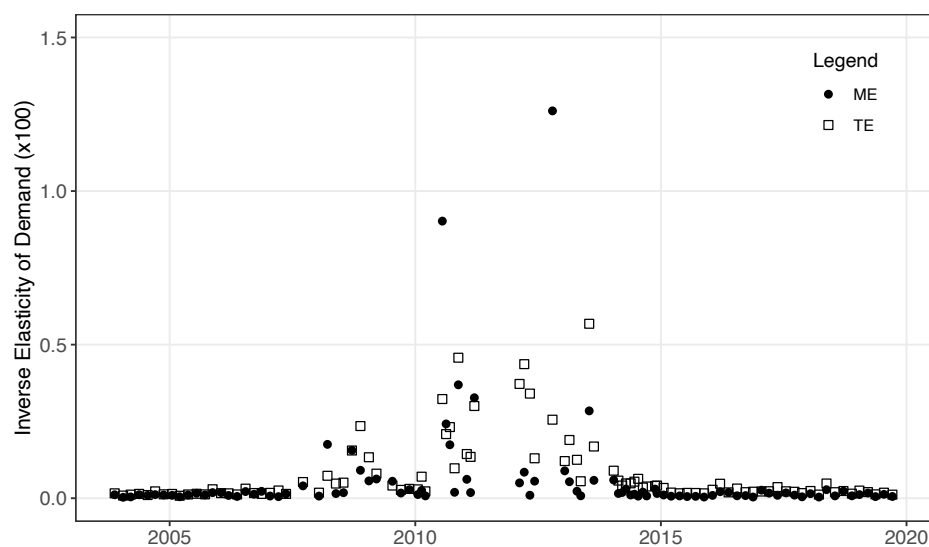


Figure 6: Comparison of ME and TE for 12 month treasury bills over time

Figure 6 shows the comparison of the two elasticity measures over time. The two measures are relatively close, particularly before and after the crisis. During the crisis, there is a marked increase in the inverse elasticity, regardless of the measure used. As discussed before, whether ME or TE is larger for a given auction depends on whether the government, respectively, is close or was able to avoid (by issuing less debt) the quasi-kink in the aggregate bid function<sup>8</sup>. Note that, the highest elasticity values come from ME and during the crisis, leading to the interpretation that the government was more willing to (or needed to) approach the steeper part of the bid function.

<sup>8</sup>Figures 21 and 22 in the appendix illustrate these differences in two auctions during the crisis.

Focusing on all securities with maturities shorter or equal to one year (Treasury Bills), Figure 7 shows the average increase in ME during the crisis period. To do that, I divide the time-series into a crisis period, ranging from 2010 to the end of 2014, and to normal times, before and after the crisis, up to the end of 2019. The maroon line represents the average inverse elasticity in each of those periods. The figure also presents the time series for ME for 3 and 12 month treasury bills as an illustration.

Considering all Treasury Bill issuances, the estimates above suggest that, *on average, the inverse elasticity of demand increased by a factor of 13 leading up and during the crisis*. This means that leading up to and during the crisis, to increase the amount issued in an auction by 1%, the decrease in price would need to be, on average, thirteen times larger than in normal times, from 0.012 to 0.15 basis points. Note also that the longer maturity bills (12 months) have a more pronounced increase in elasticity during the crisis period when compared to the shorter 3 month maturity.

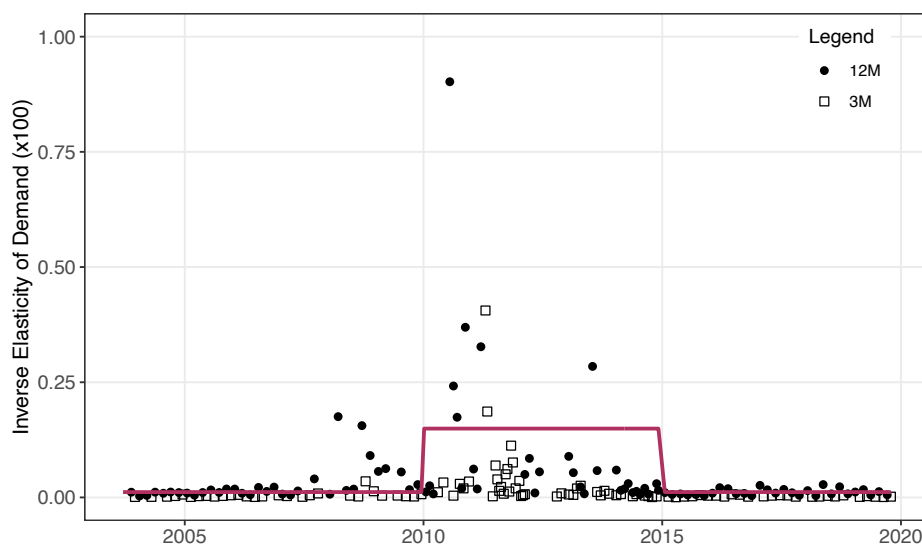


Figure 7: Average increase in inverse ME for treasury bills during crisis

Until here I have focused on Treasury Bills, securities with maturity shorter than a year. Treasury Bonds, securities with longer maturities, have two particular characteristics during this sample that make the analysis less straight forward. First, they were not issued between mid 2011 and 2014, during the crisis. Secondly, the first auction of treasury bonds after the crisis used a different protocol, a uniform price protocol. For these reasons, in order to evaluate the evolution of the inverse elasticity over time I divide the sample into three distinct periods: before the crisis (discriminatory price protocol), crisis and after the crisis (uniform price protocol).

Figure 8 shows the comparison of the two elasticity measures over time for 10 year Treasury Bonds. The two measures are relatively close, particularly before the crisis. Leading up to the crisis, before the bailout in 2011, there is a marked increase in the inverse elasticity, regardless of the measure used. This is similar to the evolution of the inverse elasticity observed for Treasury Bills, albeit with larger values for the inverse elasticity

measures. The marked difference comes in the period after the crisis. For treasury bonds, the elasticity does not recover to the pre-crisis level. I have argued before that this lack of recovery is caused by the change in auction protocol. When investors “pay-as-bid” they bid the expected value of the asset conditional on their bid being accepted. This contrasts with the uniform price protocol, where investors bid the value of the asset, as they will only be charged the minimum price accepted. It follows that bids under a discriminatory price protocol tend to be flatter than those under a uniform price protocol as investors bid below value to protect against dilution from a government that does not commit to an issuance amount ex-ante, and marginal dilution decreases along the bid function.

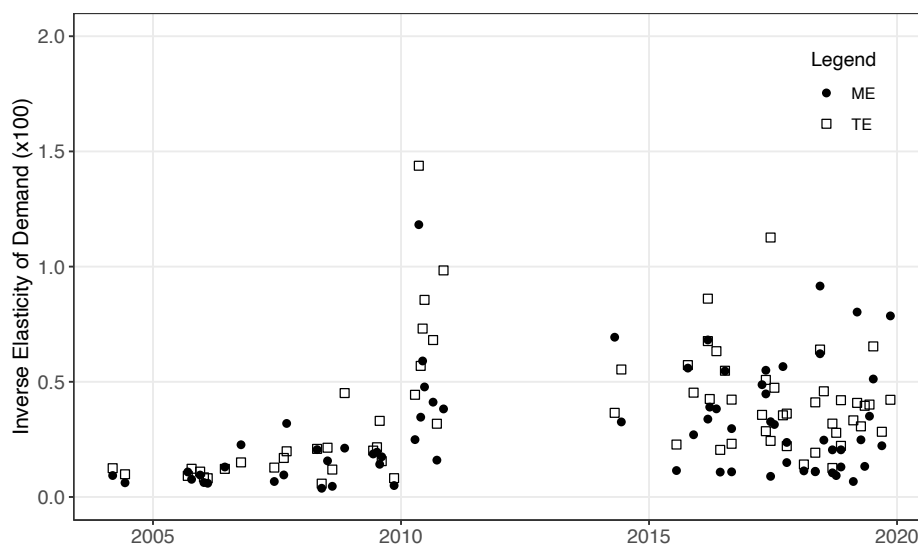


Figure 8: Average increase in inverse ME for treasury bills during crisis

Focusing on all securities with maturities longer than one year (treasury bonds), Figure 9 shows the evolution of ME over the three periods: before the crisis, before 2010; during the crisis, from 2010 to the end of 2014; and after the crisis up to the end of 2019. The maroon line represents the average ME in each of those periods. The figure also presents the time series for ME for 5, 6 and 10 year Treasury Bonds as an illustration.

Considering all Treasury Bond issuances, the estimates above suggest that, *on average, the inverse elasticity of demand for Treasury Bonds increased by 26% leading up to the crisis, from 0.29 to 0.36 basis points.*

## 4 An Auction Model

In this section, I present the model used to filter the data in order to isolate the role played by investor’s market power. The environment is based on [Wilson \(1979\)](#) framework and more specifically on [Hortaçsu and McAdams \(2010\)](#) and [Kastl \(2011b\)](#). Importantly, it enables the understanding of what is driving the shifts in demand.

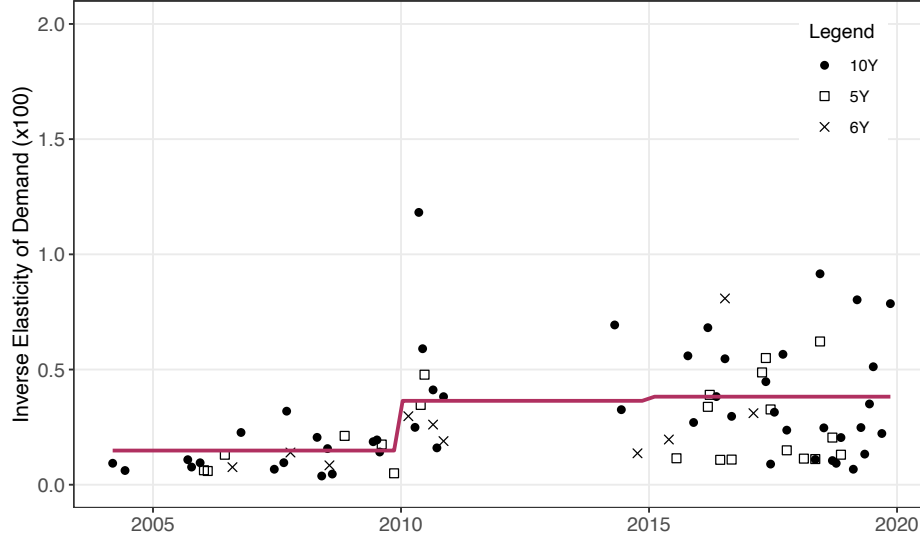


Figure 9: Average increase in inverse ME for treasury bills during crisis

A bid function may not be a good approximation for the investor's willingness to pay – the demand function. Everything else constant, different auction mechanisms will induce different bid functions. These bid functions may be closer or further apart from the agent's actual valuation of the good being auctioned. As in the related literature, I will often refer to the wedge between the agent's valuation of the asset and their bids as “shading”.

The shading term can be thought of as investor's market power. Note that, as there is a limited number of dealers in any given auction, the strategy of a single dealer may change the equilibrium price of the security being auctioned<sup>9</sup>. In this sense, dealers might be pivotal and act as such. The limited number of potential dealers and consequential lack of perfect competition among them is crucial for the existence of market power.

In an action using a discriminatory price protocol (pay-as-bid), it is intuitive that the shading term, difference between valuation and bid, is likely to be positive. This is to avoid the winner's curse<sup>10</sup>.

Let  $T$  be the number of auctions and  $N_t$  be the number of potential bidders in an auction  $t \in \{1, \dots, T\}$ . Let  $s_i$  be a private signal that  $i$  observes. This signal affects the underlying value for the auctioned good.

**Assumption 1** Bidder's signals are independent and identically distributed according to a distribution function  $F$  with density  $f$ .

**Assumption 2** Supply  $B$  is a random variable<sup>11</sup> distributed on  $[\underline{B}, \bar{B}]$  with strict positive density conditional on  $s_i \forall i$ .

Obtaining a share  $b$  of the supply  $B$  is valued according to a marginal valuation function  $v(b, s_i, s_{-i})$ . In what

<sup>9</sup>Suppose the targeted amount for the auction is 400 million euros. Consider two scenarios: (i) four dealers bid for 150 million euros, two at €98.9 and the other two at €99, and one of the dealers does not participate; (ii) the same four dealers bid as in (i) and investor A decides to participate and bids for 100 million euros at €99.1. In (ii) investor A's strategy affects the market clearing price that is €99 instead of the €98.9 in (i).

<sup>10</sup>Suppose investor A values the asset being auctioned at €99 and as such bids for it at €99; suppose further that the market clearing price of the auction is €98; it follows that investor A is going to pay €99 for the asset while they could have paid only €98.

<sup>11</sup>Although the agency announces a targeted amount to be issued, ex-post there is no commitment to the target. Furthermore, the target is sometimes presented as a range. Data suggests that often the target is not met. As such, it seems reasonable to treat the issued amount as a random variable.

follows I assume that  $v(b, s_i, s_{-i}) = v(b, s_i)$ , that is values are assumed to be private<sup>12</sup>. Furthermore, it follows from assumption 1 that I will work in the special case of independent private values (IPV). This is a standard assumption in the literature that allows for the estimation of investors' valuations using the resampling procedure detailed in section 4. The justification for the use of independent private values relies on private information being driven by idiosyncratic factors, such as the structure of the balance sheet, investment opportunities, liquidity constraints or risk-bearing capacity of each dealer. This simplifying assumption is arguably more reasonable before and after the crisis, and less so when default risk is a first order concern. Note however that the distribution of these private shocks may change overtime. That is, the distribution function  $F$  needs not be the same across auctions. Changes in the moments of the distribution can arguably rationalize changes in the common factor in investors' valuations, induced by changes in default risk. At the same time, differences across investors are justified by the idiosyncratic independent shocks. In [Alves Monteiro and Fourakis \(2023\)](#) we abstract from the independent private values assumption, and to do so we need to consider competitive investors showcasing the trade-off between these two assumptions.

**Assumption 3**  $v(b, s_i)$  is non-negative, bounded, strictly increasing in each component of  $s_i \forall b$ , and weakly decreasing and continuous in  $b \forall s_i$ .

The legislation that regulates Portuguese debt auctions establishes a maximum number of bids to be submitted by each dealer. The data corroborates this fact. Hence I restrict the strategy set available to each bidder to step functions with a finite number,  $K$ , of steps.

**Assumption 4** Each bidder  $i = 1, \dots, N$  has an action set:

$$A_i = \left\{ \begin{array}{l} (p_i, b_i, K_i) = (\{p_{ik}, b_{ik}\}_{k \in \{1, \dots, K_i\}}), K_i \in \{1, \dots, K\} \\ p_{ik} \in P \equiv [0, \bar{p}], b_{ik} \in [0, 1], p_{ik} > p_{ik+1}, b_{ik} < b_{ik+1} \end{array} \right\}$$

where  $p$  and  $b$  are, respectively, the vectors of prices and shares of total supply that constitute a bid function with  $K_i$  steps.

Let  $V(b, s_i) = \int_0^b v(x, s_i) dx$  be the utility derived from holding a share  $b$  of the debt being auctioned. Then, the expected utility of bidder  $i$  with type  $s_i$  employing strategy  $a_i(\cdot | s_i) \in A_i$  can be written as in the equation that follows.

$$\begin{aligned} EU_i(s_i) = & \sum_{k=1}^{K_i} \left[ \overbrace{\Pr(p_{ik} > P^c > p_{ik+1} | s_i)}^{\text{probability of obtaining } b_{ik}} V(b_{ik}, s_i) - \overbrace{\Pr(p_{ik} > P^c | s_i)}^{k \text{ is winning bid}} \overbrace{p_{ik} (b_{ik} - b_{ik-1})}^{\text{added cost at bid } k} \right] \\ & + \sum_{k=1}^{K_i} \Pr(p_{ik} = P^c | s_i) E_{B, s_{-i} | s_i} \left[ V(B_i^c(B, S, \mathbf{a}(\cdot | S)), s_i) \right. \\ & \left. - p_{ik} (B_i^c(B, S, \mathbf{a}(\cdot | S)) - b_{ik-1}) | p_{ik} = P^c \right] \end{aligned}$$

Where the random variable  $B_i^c$  is the market clearing quantity obtained by bidder  $i$  when the state is  $(B, S \equiv$

<sup>12</sup>Investor's valuation depends only on the realization of its individual signal.

$\times_{i=1}^N s_i$ ) and bidders submit bids specified in the vector  $\mathbf{a}(\cdot | S) = [a_1(\cdot | s_1), \dots, a_N(\cdot | s_N)]$ . The market clearing price is a random variable denoted by  $P^c$ . The last two lines of the expression above describe the expected utility from bids that may be rationed as the submitted price equals the market clearing price of the auction.

## 4.1 Equilibrium

A **Bayesian Nash Equilibrium** is a set of strategies that maximizes the expected utility for each agent and signal realization:  $a_i(\cdot | s_i) \in \arg\max_{a_i \in A_i} EU_i(s_i) \forall i$  and  $s_i$ .

In an equilibrium, every step  $k < K_i$  in the bid function  $a_i(\cdot | s_i)$  must satisfy:

$$\Pr(p_{ik} > P^c > p_{ik+1} | s_i) [v(b_{ik}, s_i) - p_{ik}] = \Pr(p_{ik+1} \geq P^c | s_i) (p_{ik} - p_{ik+1}) \quad (1)$$

The necessary equilibrium condition above clarifies the trade-off at each step  $k$ . Suppose that a bidder has the following bid function  $\{(b_1, p_1), (b_2, p_2)\}$ . Further, let  $b'_1 > b_1$  and consider moving the first bid from  $(b_1, p_1)$  to  $(b'_1, p_1)$ . If the clearing price is such that only the first bid is executed then the surplus increases by  $\int_{b_1}^{b'_1} [v(x, s_i) - p_1] dx$ . On the other hand, if the clearing price is below  $p_2$  then the loss in surplus is  $(p_1 - p_2) \times (b'_1 - b_1)$ . In equilibrium, strategies must be such that there is no incentive to change the bids chosen: at the margin, the expected gain from deviating equals the expected loss from doing so.

At the last step,  $k = K_i$ , the bid function  $a_i(\cdot | s_i)$  must satisfy:

$$v(\bar{b}, s_i) = p_{iK_i}, \text{ where } \bar{b} = \sup_{\{b, s_{-i}\}} B_i^c(B, \mathbf{S}, \mathbf{a}(\cdot | S))$$

For a given dealer, at the last step there is no trade-off in the sense that there are no bids placed at a lower price. As such, there are no incentives to shade the last step.

## 5 Estimation

With the  $K_i$  equations, for every dealer  $i$ , one can retrieve the true valuations  $v_i(\cdot)$ , provided one can estimate the distribution of market clearing prices  $P^c$  conditional on  $s_i$ . This is a result of the private values assumption: the only way other bidders' bidding strategies affect bidder  $i$  is through the distribution of  $P^c$  conditional on  $s_i$ . This in turn hinges on the distribution of residual supply bidder  $i$  faces.

Formally, we want to estimate the probability of a winning bid conditional on the individual signal  $s_i$ :

$$G(p; B) \equiv \Pr[P^c \leq p | s_i] = \mathbb{E}_{\{B, s_{-i}\}} \mathbb{1} \left( B - \sum_{j \neq i} a(p | s_j) \geq a(p | s_i) \right), \quad \forall p \in [0, \bar{p}]$$

Define an indicator of excess supply:

$$\Phi \left( \{a(p | s_j)\}_{j \neq i}; p, B \right) = \mathbb{1} \left( B - \sum_{j \neq i} a(p | s_j) \geq a(p | s_i) \right)$$

One estimator for  $G(p)$  can be derived as a V-statistic:

$$\xi \left( \hat{F}; p, B \right) = \frac{1}{(NT)^{(N-1)}} \sum_{\alpha_1=(1,1)}^{(T,N)} \dots \sum_{\alpha_{N-1}=(1,1)}^{(T,N)} \Phi(a_{\alpha_1}, \dots, a_{\alpha_{N-1}}, p)$$

where  $\hat{F}$  is the empirical distribution of bids.

The estimator  $\xi$  is simply the proportion of aggregate states ( $S$ ) – over all the permutations of  $N - 1$  individual bids – in which there is excess supply at a price  $p \in [0, \bar{p}]$ .

Note that it is not feasible to compute  $\xi$  by summing over all permutations of bids. Instead, I use the resampling procedure first proposed for the multi-unit auction environment in [Hortaçsu \(2002\)](#)<sup>13</sup>. Essentially for one fixed bidder at a time I draw a random sample of  $N_t - 1$  individual bid functions with replacement. I then construct the residual supply function and find the market clearing price by intersecting it with the fixed individual bid function. This is performed a larger number of times per bidder to obtain the subjective distribution of the market clearing price. Refer to the appendix for a thorough description of the resampling procedure.

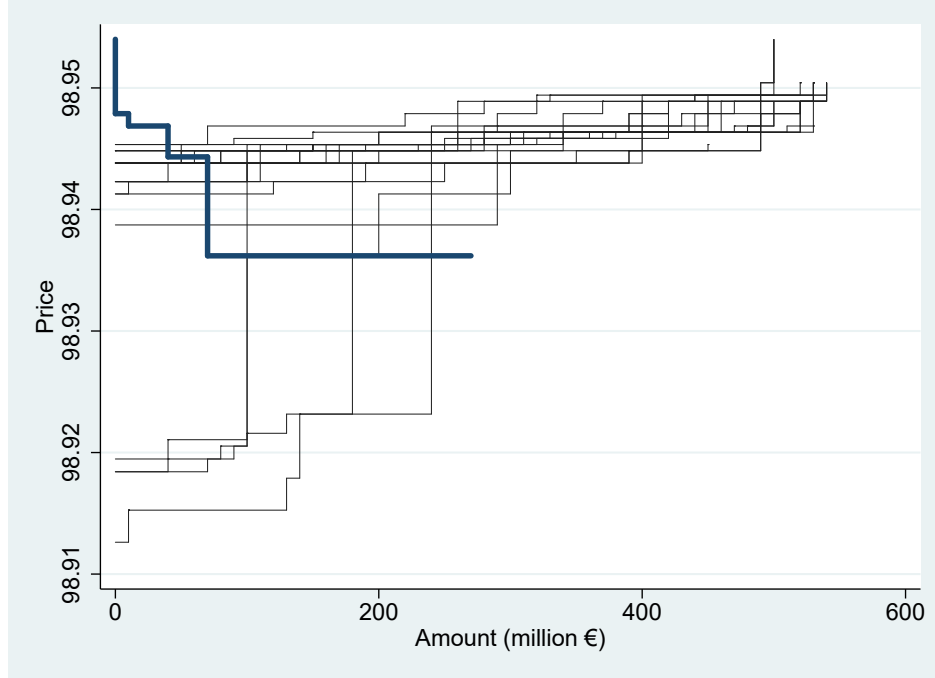


Figure 10: Different realizations of residual supply for dealer 1

<sup>13</sup>The author shows the asymptotic properties of the estimator being used.

**Example 1** In Figures 10 one can see, an illustration of the resampling procedure for Dealer 1 in a given auction: the downward step function is Dealer's 1 bid function and each upward residual supply corresponds to a different sample. In Figure 11 one can see the corresponding market clearing price distribution, computed from the intersection between the bid function and the residual supply functions.

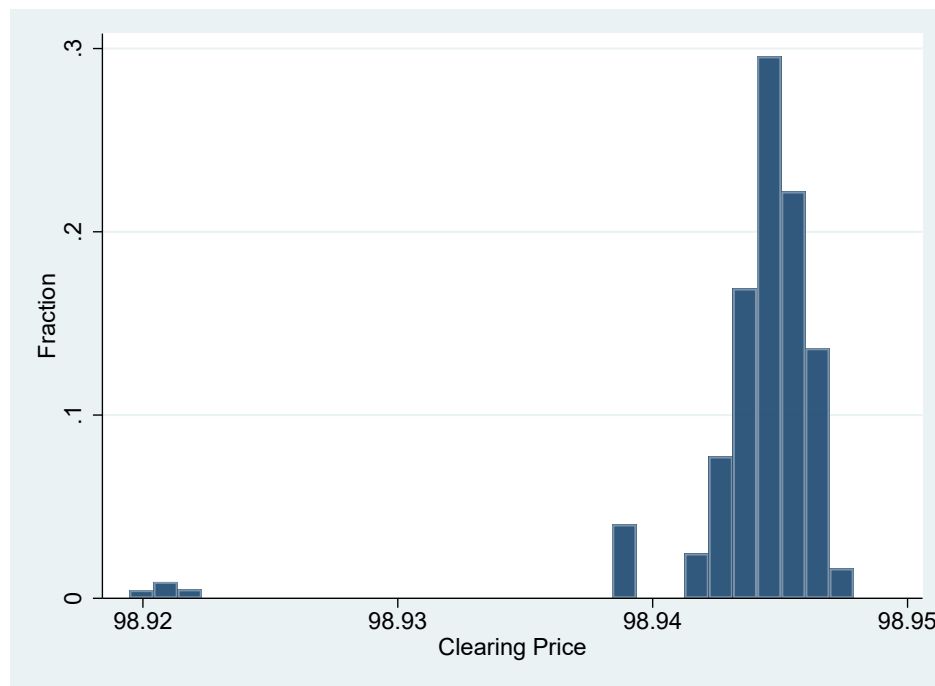


Figure 11: Market clearing price distribution

## 6 Market Power and Inefficiency Costs

With both the valuations and the bid functions for each investor, one can disentangle the effect of decreasing valuations from the role of market power on said bid functions.

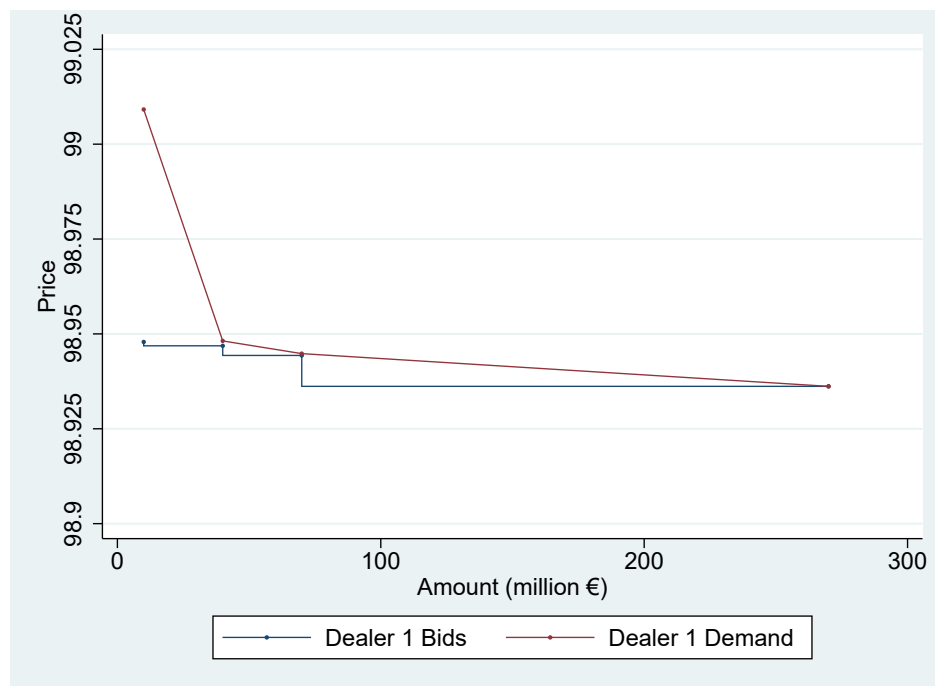
**Example 2** Figure 12 shows the difference in the wedge between bids and valuations for the same dealer. Panel (a) shows the wedge before the crises and panel (b) shows the wedge during the crisis. It is clear that the wedge is much more pronounced during the crisis.

From the example, it follows that the individual willingness to pay of a dealer tends to be more inelastic during the crisis. This is indicative that not only are bid functions more inelastic during the crisis but so are the actual valuations. As a result, the increase in the elasticity when considering the bid functions, instead of the actual willingness to pay of dealers, may be understated.

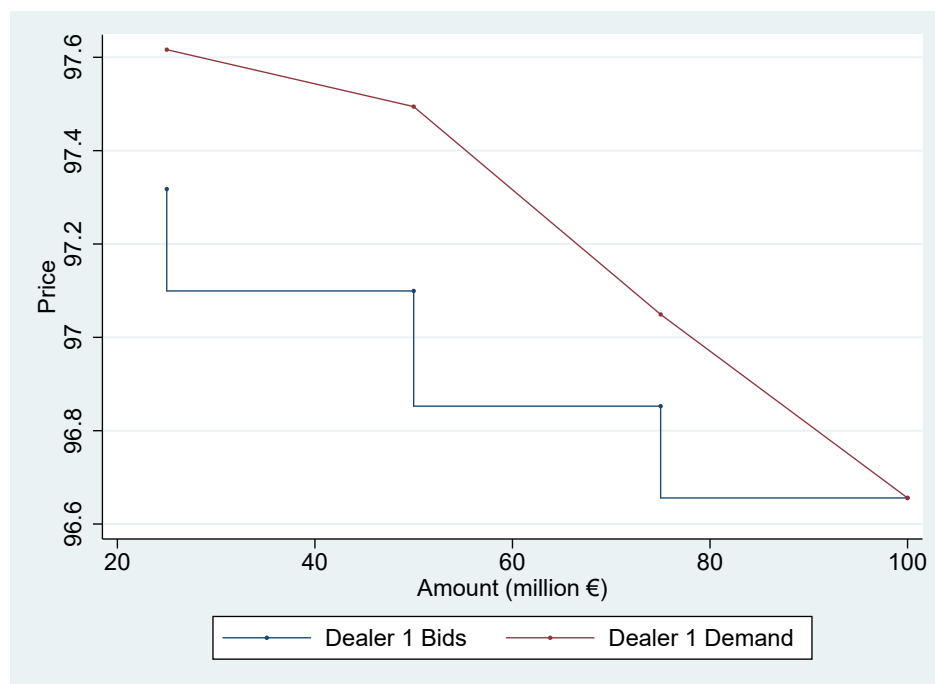
In order to present an aggregate measure of this wedge I look at the "In the Money" (ITM) bids: the winning bids that are executed, not necessarily in full. Computing the average ITM shade<sup>14</sup>, i.e. exclusively with ITM bids

<sup>14</sup>The average ITM shade is computed as follows: (i) for each bidder compute the average shade across bids weighted on the amount of each bid over the total amount bid by the dealer, (ii) average across bidders.

(and valuations), one can have a better sense of the effective average wedge.

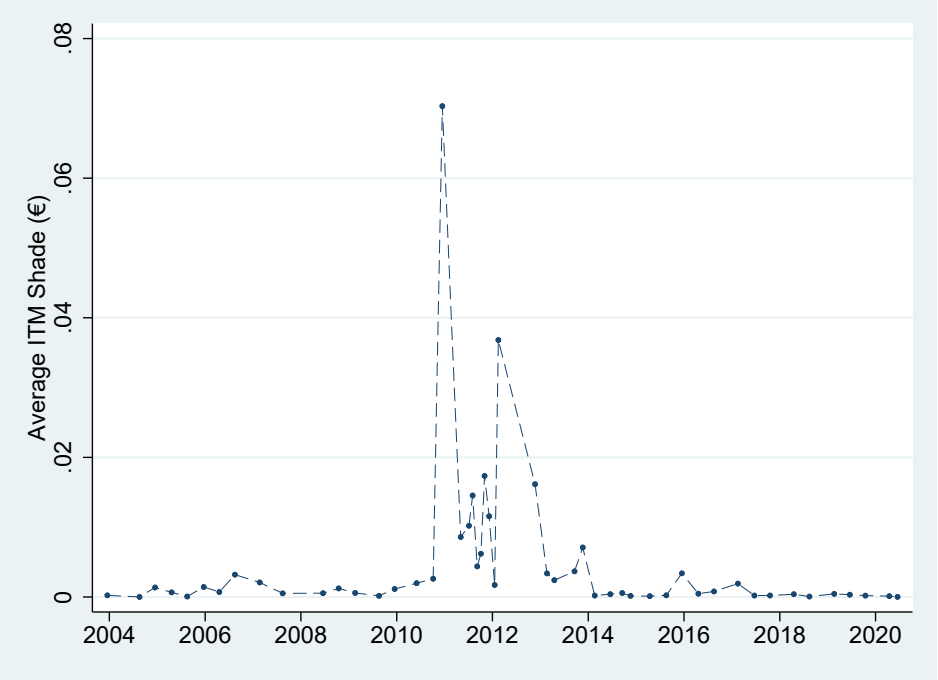


(a) An auction before the crisis

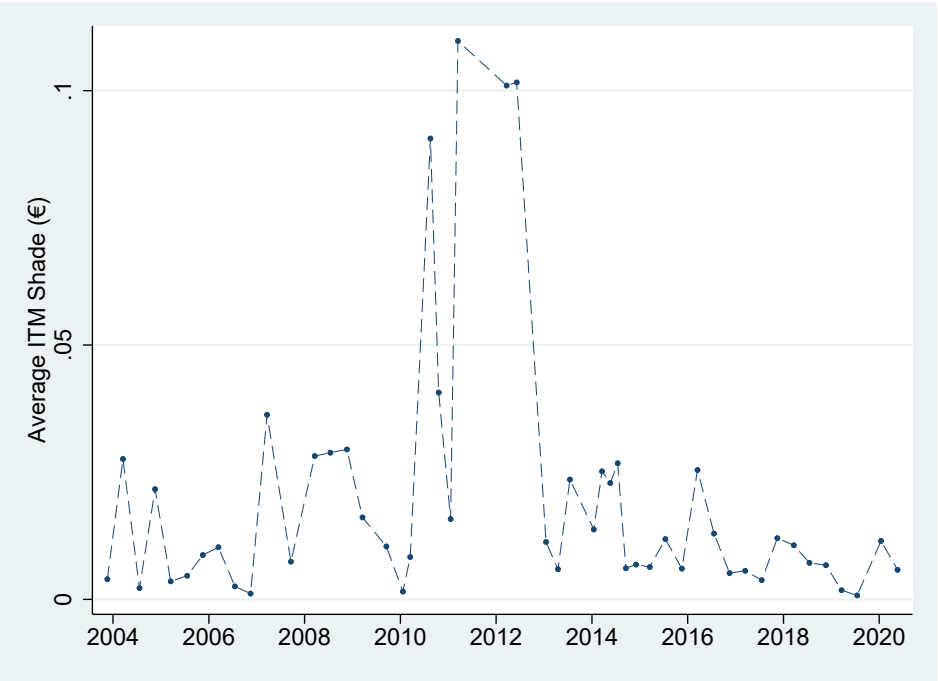


(b) An auction during the crisis

Figure 12: Valuation and bid function for Dealer 1



(a) 3 month treasury bill auctions



(b) 12 month treasury bill auctions

Figure 13: Average In the Money shade in treasury bill auctions

Figure 13 shows the average wedge between ITM bids and the corresponding valuation for each auction of 3 month and 12 month bills in the sample period. Note that the role of market power is not very significant during normal times, as bids tend to be closer to the dealers' valuations. Also, even during the crisis period, the effect of decreasing valuations still dominates over the role of market power on investors' actions. However, the strategic component gets more significant leading up and during the crisis. In fact, the shading terms for 3 and 12 month treasury bills, at their respective peaks, account for approximately a 20 and 10 basis point increase in the average yield of the auction.

The next logical step would be to evaluate how this wedge is linked to the inefficiency costs of the mechanism. Note that the wedge between bids and valuations implies that the government is not extracting all of the dealer's surplus from buying the auctioned securities. Thus, one can think of this wedge as a unitary inefficiency cost of the auction mechanism: the money "*left on the table*" per unit of debt issued.

Figure 14 presents the inefficiency costs as a percentage of the amount raised in a given auction, over time and for three different maturities. Once again, the inefficiency is computed as the sum of the individual wedges for the ITM bids, i.e. the effective inefficiency cost of the auction. Across the three short term maturities depicted, the inefficiency costs increase leading up and during the crisis.

This increase is not surprising since the wedge per unit increases. It is worth mentioning nonetheless that the inefficiency cost differs, in levels, across maturities: it is smaller for shorter maturities.

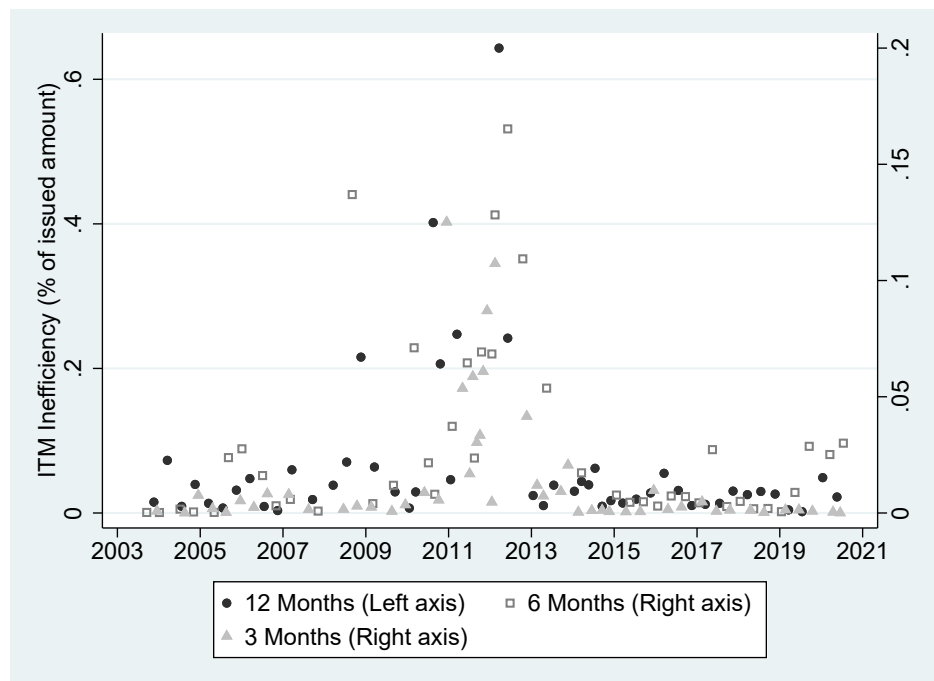


Figure 14: Inefficiency costs as a % of raised amount in treasury bill auctions

## 6.1 What is driving the increase in inefficiency costs?

Investors chose their bid function given their private signal and the corresponding expectation of the aggregate state. Before the crisis, there is little dispersion across bid functions and, as a result, the role of market power is limited. That is, if investors bid below their valuation, they will likely leave the auction empty handed<sup>15</sup>. As the crisis period gets closer, the dispersion in bid functions increases, allowing investors to exploit their market power. Mechanically, the subjective price distribution for a given dealer, in the above notation  $P[P^c < p | s_i]$ ,  $\forall p \in [0, \bar{p}]$ , has more variance during the crisis. For each set of dealers participating in an auction, there is a different market clearing price; moreover, changing the set of participating dealers leads to potentially more disperse market clearing prices than before<sup>16</sup>.

Recall equation 1 and note that it can be written as follows:

$$v(b_{ik}, s_i) - p_{ik} = \theta(k, k+1 | s_i) (p_{ik} - p_{ik+1})$$

where  $\theta(k, k+1 | s_i)$  is the likelihood ratio:

$$\theta(k, k+1 | s_i) = \frac{\Pr(p_{ik+1} \geq P^c | s_i)}{\Pr(p_{ik} > P^c > p_{ik+1} | s_i)}$$

The increased dispersion in market clearing prices during the crisis likely leads to a higher  $\theta(k, k+1 | s_i)$ , both because (i) the numerator is bigger – winning bids at lower prices is more likely than before –, and (ii) the denominator is smaller – the probability of the market clearing price being in a given interval is smaller. To rationalize the observed equilibrium bids, an increase in the likelihood ratio needs to be accompanied by an increase in the wedge. That is, valuations need to be larger than the submitted bids, and by more than they were before the crisis.

Estimating the following equation is indicative of the above explanation:

$$Shade_t = \alpha + \beta Bid\ sd_t + \epsilon_t$$

where the estimated coefficient of interest is  $\hat{\beta} = 0.069$  with an  $R^2 = 0.5$ . That is, an increase in the dispersion of bids of 1 is associated with an increase in the wedge of 0.07, everything else constant. Figure 15 depicts the fitted values against the wedge from the data generating process.

Summing up, leading up and during the sovereign debt crisis: (i) valuations decrease; (ii) bids decrease more than valuations do, due to an increased importance of the market power mechanism; (iii) this wedge between valuations and bid functions generates inefficiency costs when the government issues debt; (iv) these costs tend to be negligible in normal times but go up to 0.6% of the raised amount, during the crisis; and finally, (v) one can

<sup>15</sup>With no dispersion in bids investors will act as if they had no market power. Essentially, if there is no dispersion across bids the market clearing price is pinned down and investors take it as given.

<sup>16</sup>An intuitive way to visualize the increased variance is to think of the resampling procedure and the fact that resampling from a set of bid functions that are more disperse will lead to a more disperse distribution of residual supply and, consequently, market clearing price.

get a sense of the importance of the role of market power and consequent inefficiency by analyzing the dispersion of bids in an auction.

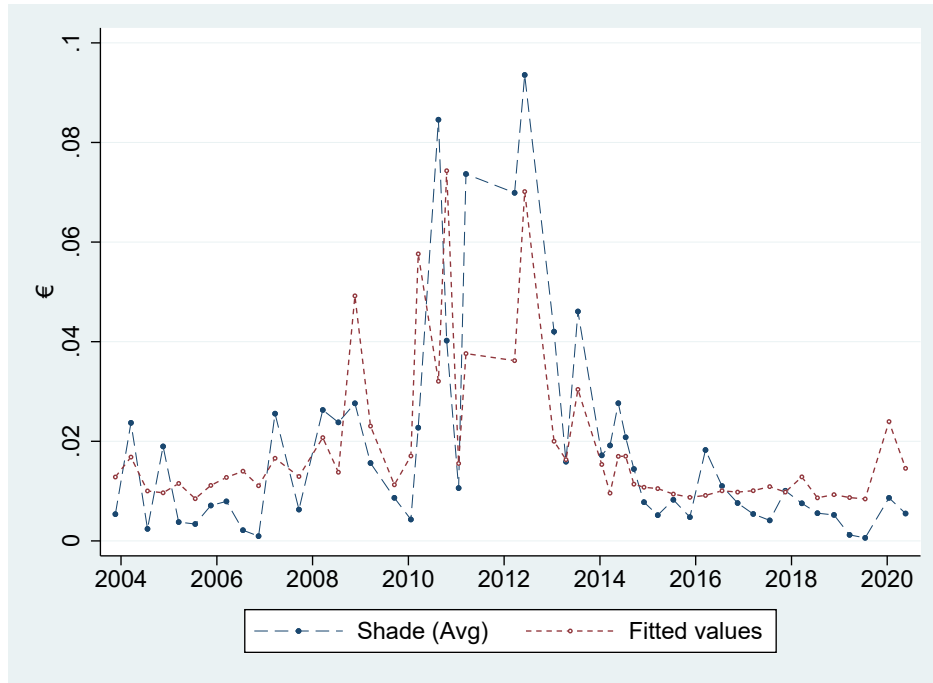


Figure 15: Prediction of Inefficiency with observables

## 7 Mitigation Strategies: Maturity Choice

Inefficiency costs increase during a debt crisis. However, this increase is not homogeneous across maturities. Moreover, shorter maturities need to be rolled over more frequently than longer ones. These observations raise the question of whether the government can employ a mitigation strategy to minimize the annualized inefficiency costs.

In this section, I briefly discuss a simple mitigation strategy: maturity choice. Let's suppose that the government needs to roll over a certain amount of debt,  $B$ , and given the high marginal costs it is facing, is not willing to issue new debt in excess of the amount needed to roll over.

I will abstract from default decisions; this particular problem aims to solve the maturity choice for a government that is indeed repaying its debt. In that spirit, I will also abstract from self fulfilling rollover risk as in [Cole and Kehoe \(2000\)](#). That is, more short-term debt leads to an increased share of debt to be refinanced in any given period. This in turn could make the government more vulnerable to changes in investor sentiment. Abstracting from such considerations, the problem is simply:

$$\min_{\{\{b_{m,t}^j\}_{m=1,\dots,12}\}_{j \in \mathcal{J}}} C_t$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{J}} q_{m+1,t}^j b_{m+1,t}^j = B_{mt}, \quad \forall m = 0, \dots, 11$$

where the inefficiency cost  $C_t$  is the cost associated with the government's strategy for debt issuances over a given year  $t$  and  $\mathcal{J}$  is the set of available maturities that can be issued by the agency<sup>17</sup>.

### A back of the envelope analysis

The following regression gives us a measure of the average increase in inefficiency costs of each maturity  $j$  during the crisis.

$$inef\ ratio_{j,t} = \alpha_j + \beta_j crisis_{j,t} + \epsilon_{j,t}$$

where  $crisis_t$  is a dummy variable that equals to one leading up and during the crisis.

Maturity	Normal times ( $\alpha$ )	Crisis increment ( $\beta$ )
3 Months	0.002%	0.038%
6 Months	0.015%	0.065%
12 Months	0.032%	0.265%

Table 2: Average increase in the inefficiency cost per raised amount

Table 2 shows that the average inefficiency cost per raised amount increases for all maturities. However, the increase in cost is more pronounced for 12 month treasury bills.

A government that decides to issue only 3 month bills has to issue them 4 times during a year. As such, the annualized cost of issuing 3 month treasury bills is actually  $(0.002 + 0.038) \times 4 = 0.16\%$ . Using the same reasoning, the annualized cost of issuing 6 month treasury bills is also 0.16%. Finally, the cost for 12 month treasury bills is simply 0.297%.

Everything else constant, considering these inefficiency costs, it seems that a reasonable mitigation strategy would be to issue shorter maturity bills that, even though would imply more issuances, result in smaller inefficiencies.

From 2010 to 2014, the government did not auction treasury bonds. Hence, part of the government's decision was in fact to restrict issuances to those of shorter maturities. Figure 16 shows the evolution of inefficiency costs for 10 year treasury bonds. The figure shows a spike in the inefficiency costs as the crisis approaches. This trend also suggests that avoiding such long bonds is, everything else constant, a good strategy to mitigate inefficiency costs.

A more thorough analysis of optimal maturity choice accounting for the inefficiency costs of the mechanism is left as future research.

<sup>17</sup>That is, if the government opts to issue only 3 month bills, then  $C$  is the inefficiency cost associated with 4 issues of such bills. If, on the other hand, the government opts to issue 3 month bills twice and then 6 month bills, then  $C$  is the inefficiency cost associated with those 3 issues of treasury bills

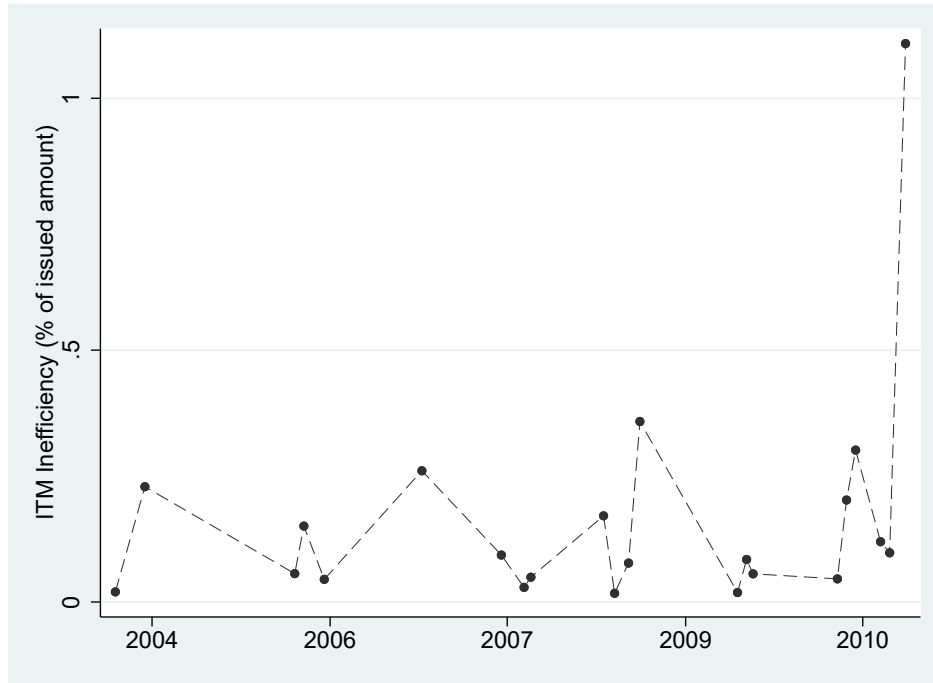


Figure 16: Inefficiency costs as a % of raised amount in 10 year treasury bond auctions

## 8 Conclusion

Using bid level data for Portuguese sovereign debt auctions, from 2003 to 2020, I documented a key pattern in investors' demand during a high default risk event – the Portuguese sovereign debt crisis: leading up to and during the crisis, bids get more disperse and the aggregate bid function faced by the government becomes more inelastic. This is true across short and long maturities for the duration of the crisis. Particularly, for Treasury Bills, the inverse elasticity of demand increases by a factor of 13. After the crisis, bid functions tend to recover to their pre-crisis shape, absent changes in auction protocol. This fact fits the description put forward by the Portuguese government: shifts in demand were responsible for lower than expected amounts issued during the crisis. The government tends to avoid the steeper part of the schedule by issuing lower amounts.

I then presented a model of the discriminatory auctions in which investors' have market power. This market power arises from the non-competitive nature of the auctions: only a small number of investors is able to bid in the auctions. Crucially, market power allows bids to differ from valuations. With the model, I filter the data and separate bids and actual valuations and, consequently, assess the role of investors' market power on the shifts in bid functions. I find that the role of market power is negligible during normal times. However, it gets more significant leading up to and during the crisis period. In fact, the shading terms for 3 and 12 month treasury bills, at their respective peaks, account for approximately a 20 and 10 basis point increase in the average yield of the auction.

I argue that this wedge, between valuations and bids, can be seen as a unitary inefficiency cost of the auction

mechanism, the “*money left on the table*” by the government. The ratio of such inefficiency costs as a percentage of the amount raised tends to be negligible in normal times, but goes up to 0.6% during the crisis.

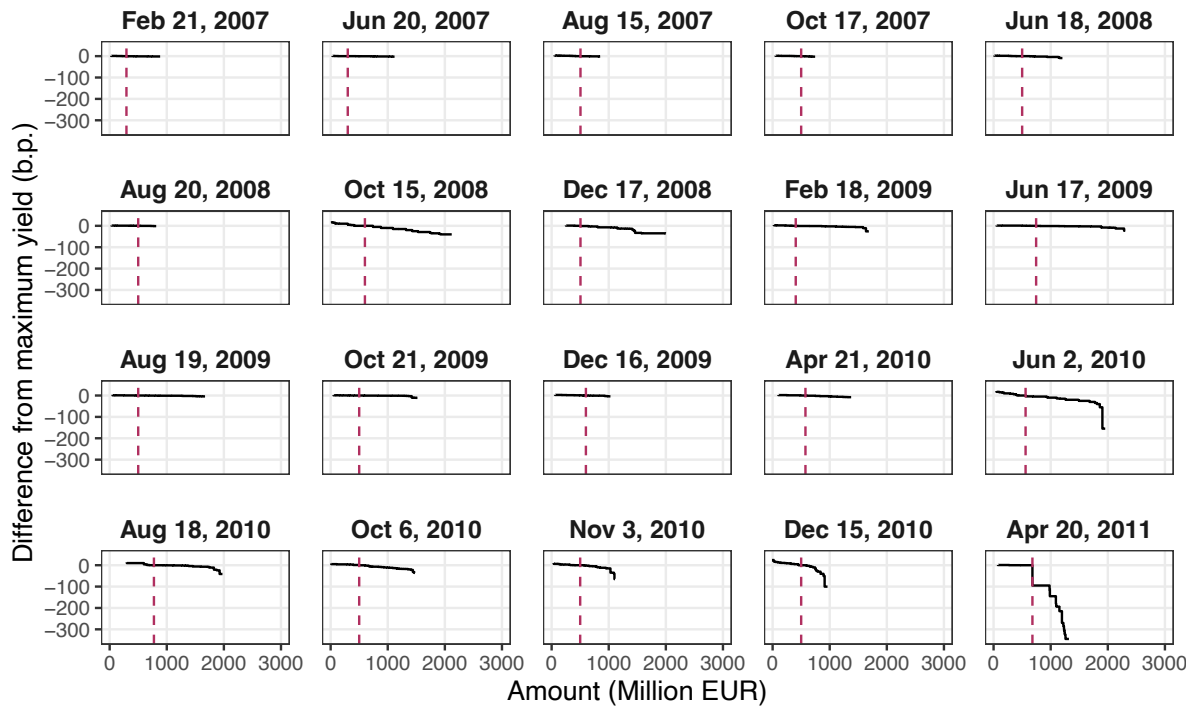
A logical next step would be a more normative analysis: what can the government do to mitigate these inefficiency costs when issuing debt during a crisis? I briefly look at maturity choice as a mitigation device. Short maturities tend to have lower costs but need to be rolled over more frequently. A back of the envelope computation suggests that issuing shorter maturities reduces the inefficiency costs compared to issuing longer maturities.

## References

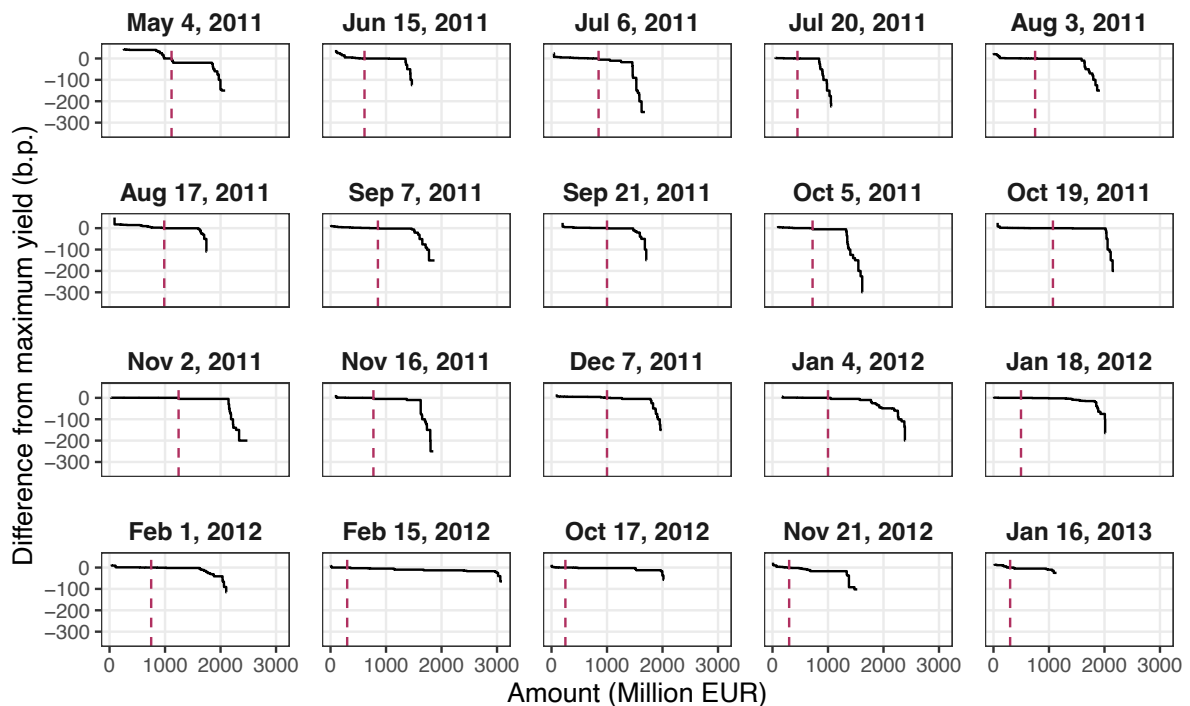
- Association for Financial Markets in Europe AFME. European primary dealers handbook, 2020.
- Mark Aguiar and Manuel Amador. *Sovereign Debt*, pages 647–87. North-Holland, 2014.
- Mark Aguiar and Manuel Amador. *The Economics of Sovereign Debt and Default*. Princeton University Press, 2021.
- Mark Aguiar and Gita Gopinath. Defaultable debt, interest rate and the current account. *Journal of International Economics*, 69(1):64–83, 2006.
- Mark Aguiar, Manuel Amador, Hugo Hopenhayn, and Ivan Werning. Take the short route: Equilibrium default and debt maturity. *Econometrica*, 87(2):423–462, 2019.
- Rui Albuquerque, José Miguel Cardoso-Costa, and José Afonso Faias. Price elasticity of demand and risk-bearing capacity in sovereign bond auctions. March 2022.
- Ricardo Alves Monteiro and Stelios Fourakis. Sovereign debt auctions with strategic interactions. Working paper, 2023.
- Cristina Arellano. Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712, 2008.
- Cristina Arellano and Ananth Ramanarayanan. Default and the maturity structure in sovereign bonds. *Journal of Political Economy*, 120:187–232, 2012.
- Saki Bigio, Galo Nuño, and Juan Passadore. Debt-maturity management with liquidity costs. August 2021.
- Luigi Bocola and Alessandro Dovis. Self-fulfilling debt crises: A quantitative analysis. *American Economic Review*, 109(12):4343–4377, 2019.
- Fernando A. Broner, Guido Lorenzoni, and Sergio L. Schmukler. Why do emerging economies borrow short term? *Journal of the European Economic Association*, 11:67–100, 2013.
- Harold Cole and Timothy Kehoe. Self-fulfilling debt crises. *The Review of Economic Studies*, 67:91–116, 2000.
- Harold Cole, Daniel Neuhann, and Guillermo Ordoñez. Asymmetric information and sovereign debt: Theory meets mexican data. Working Paper 28459, National Bureau of Economic Research, February 2021.
- Jonathan Eaton and Mark Gersovitz. Debt with potential repudiation: Theoretical and empirical analysis. *The Review of Economic Studies*, 48(2):280–309, 1981.
- Ali Hortaçsu. Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. working paper, 2002.

- Ali Hortaçsu and David McAdams. 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, 2010.
- Jakub Kastl. Discrete bids and empirical inference in divisible good auctions. *The Review of Economic Studies*, 78(3):974–1014, 2011a.
- Jakub Kastl. Discrete bids and empirical inference in divisible good auctions. *Review of Economic Studies*, 78: 978–1014, 2011b.
- Jakub Kastl. Auctions in financial markets. *International Journal of Industrial Organization*, 70, 2020.
- Matías Moretti, Lorenzo Pandolfi, Sergio Schmukler, and Germán Villegas Bauer. Inelastic demand meets optimal supply of risky sovereign bonds. 2024.
- Juan M. Sánchez, Horacio Sapriza, and Emircan Yurdagul. Sovereign default and maturity choice. *Journal of Monetary Economics*, 95:72–85, 2018.
- Robert Wilson. Auctions of shares. *The Quarterly Journal of Economics*, 93(4):675–689, 1979.

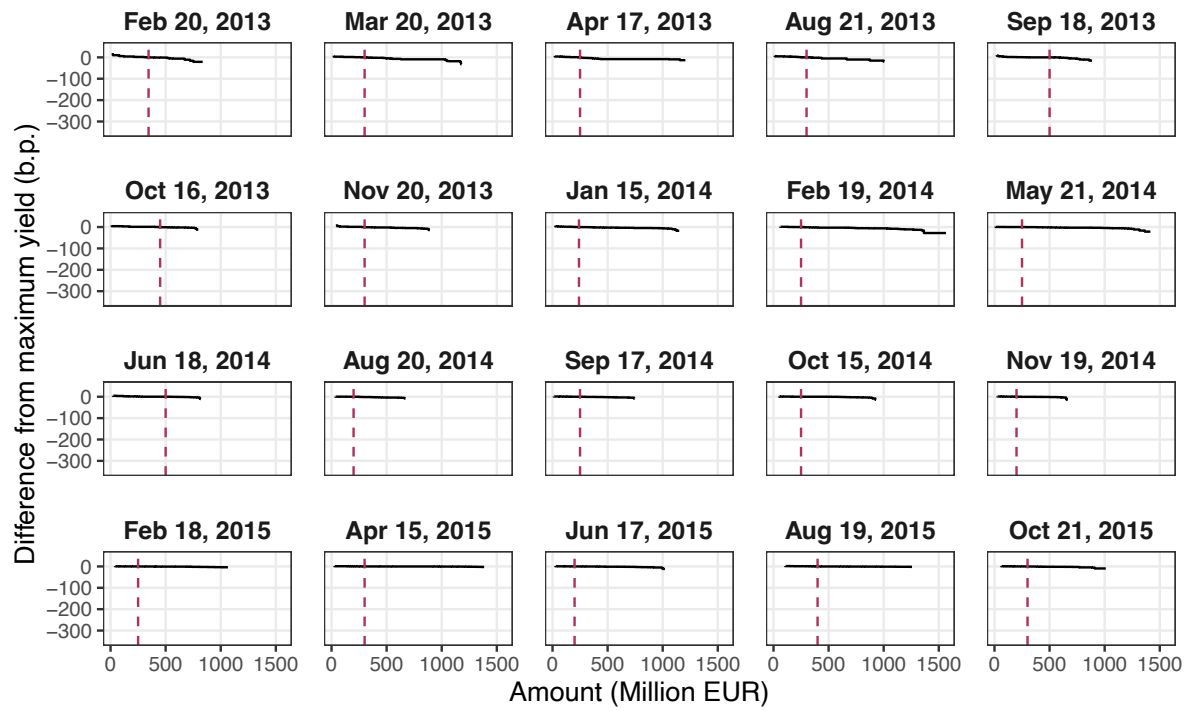
## Appendix A - Bid Functions



(a) Before and leading up to the crisis

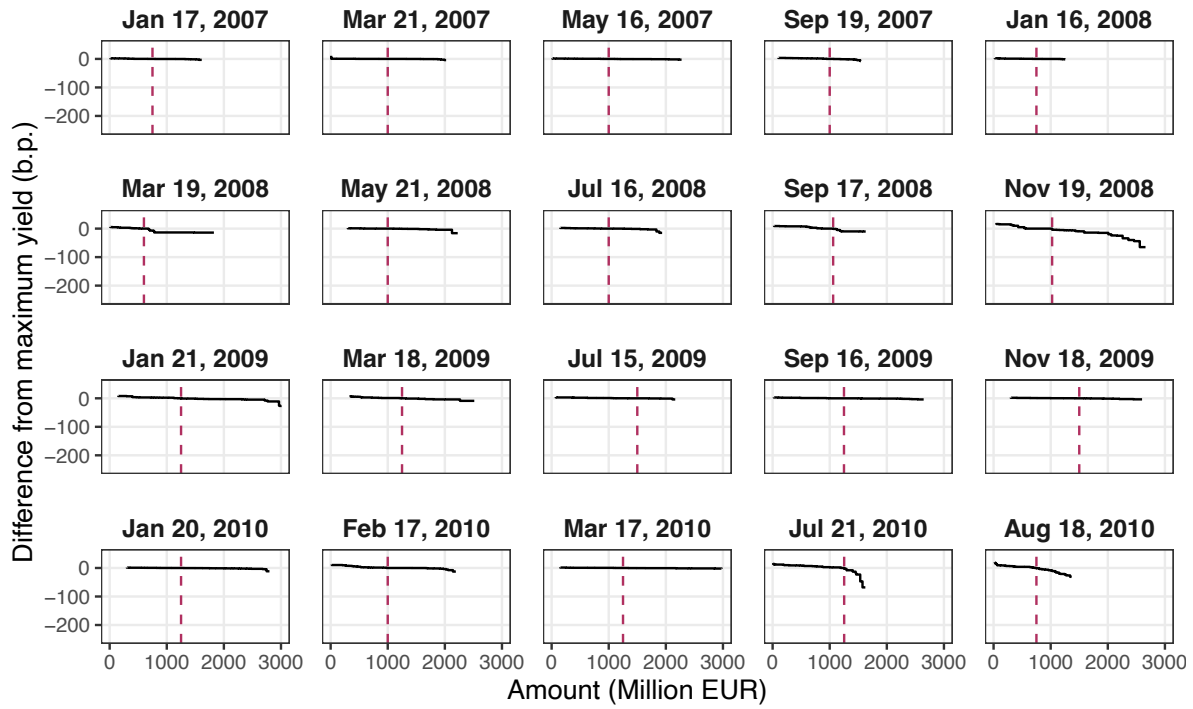


(b) Crisis Period

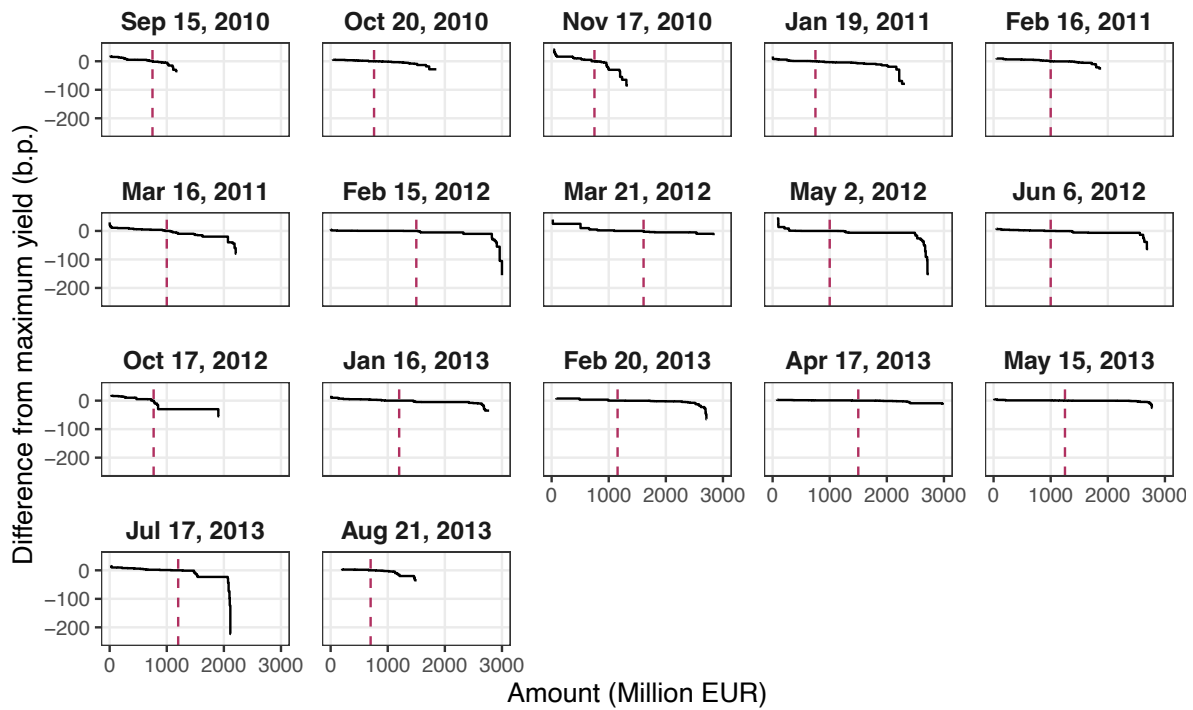


(c) Recovery

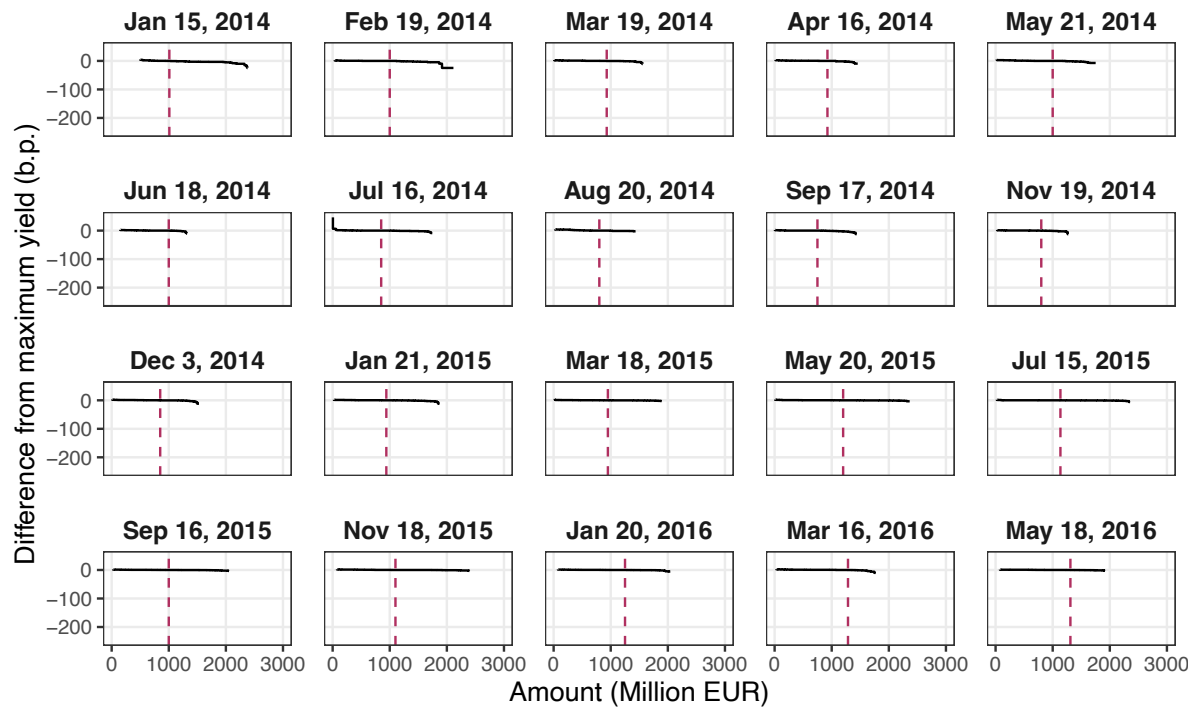
Figure 17: Aggregate Bid Functions for 3 month treasury bills in the primary market



(a) Before and leading up to the crisis



(b) Crisis Period



(c) Recovery

Figure 18: Aggregate Bid Functions for 12 month treasury bills in the primary market

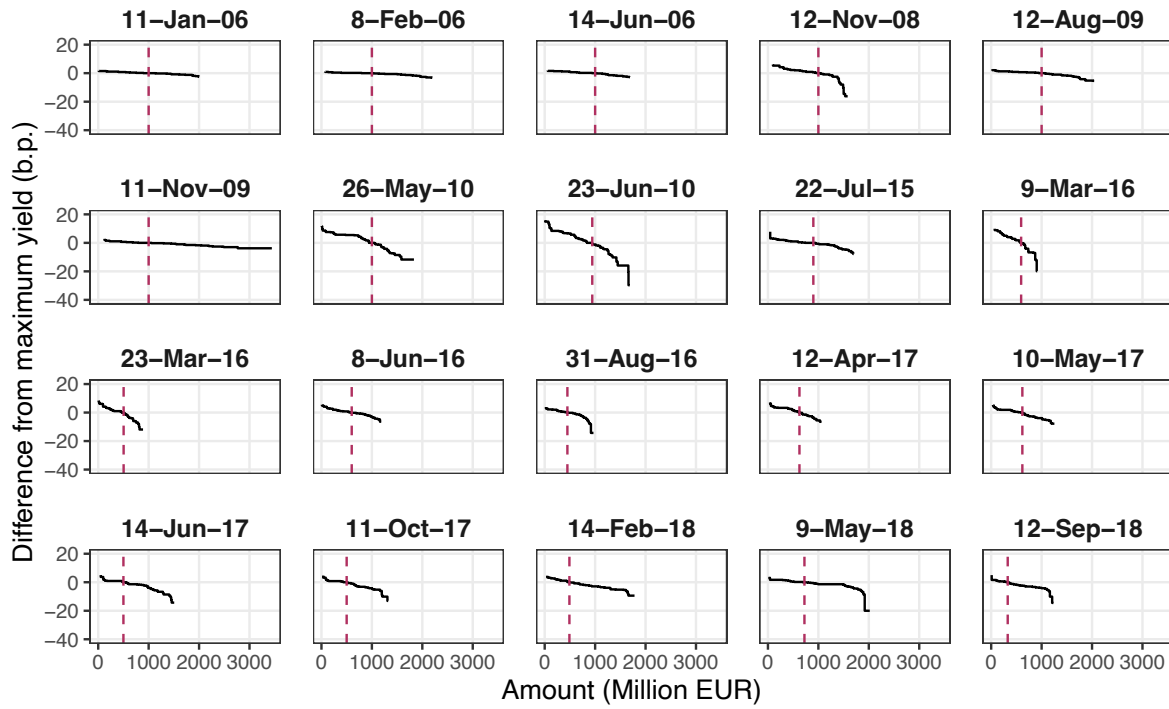
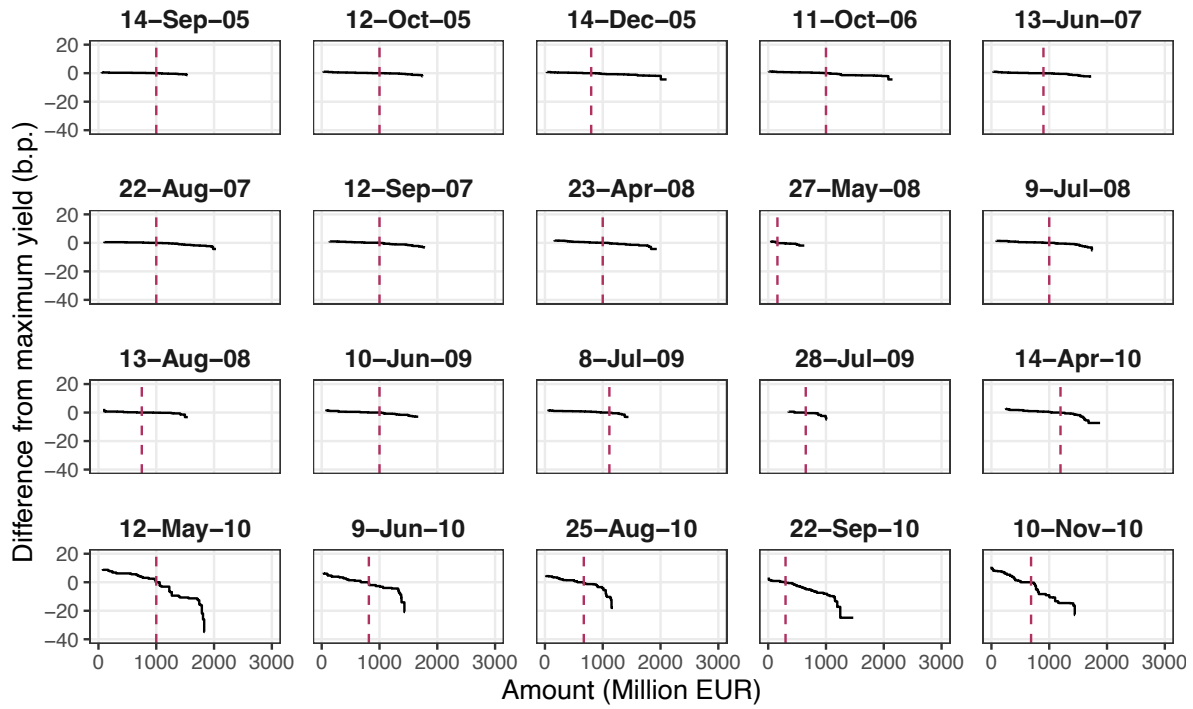
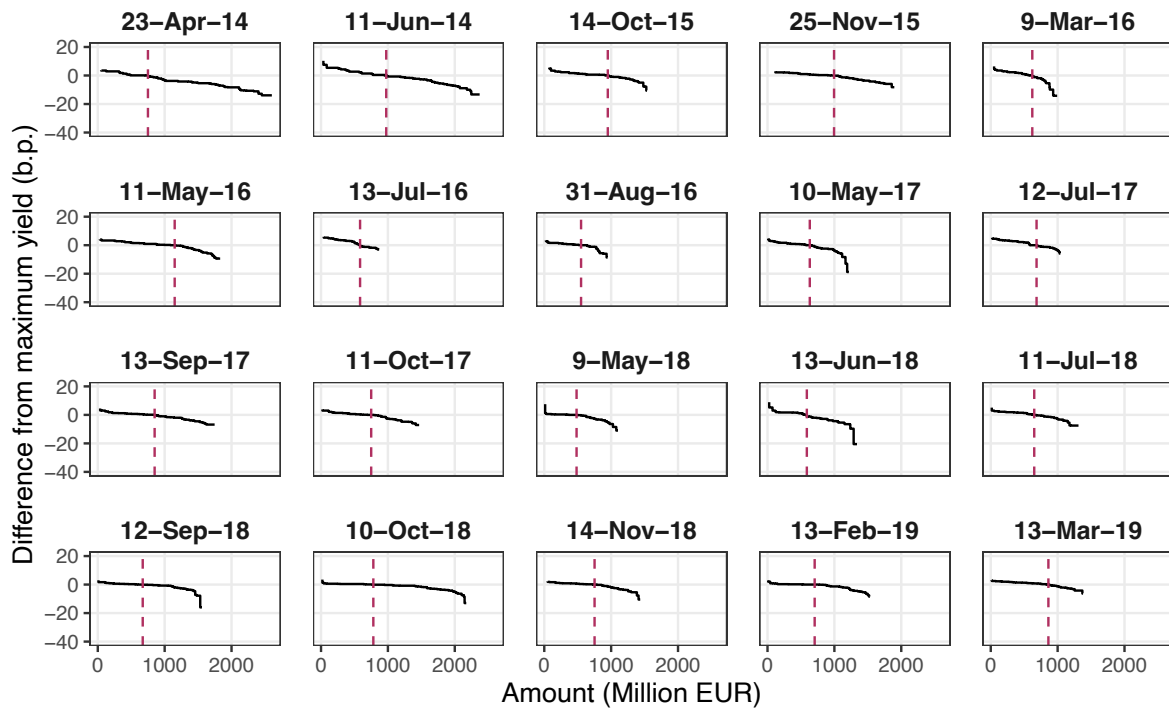


Figure 19: Aggregate Bid Functions for 5 year treasury bonds in the primary market



(a) Leading to the sovereign crisis



(b) Post Crisis

Figure 20: Aggregate Bid Functions for 10 year treasury bonds in the primary market

## Appendix B - Elasticities Comparison

The figures below illustrate the difference between the estimated slopes used to compute ME and TE. This serves to highlight the importance of this difference leading up and during the crisis. The difference in the two measures induced by the quasi-kink in the aggregate bid function, is also indicative of the government's behavior – whether the government avoids the cliff or not, makes ME lower or greater than TE, respectively.

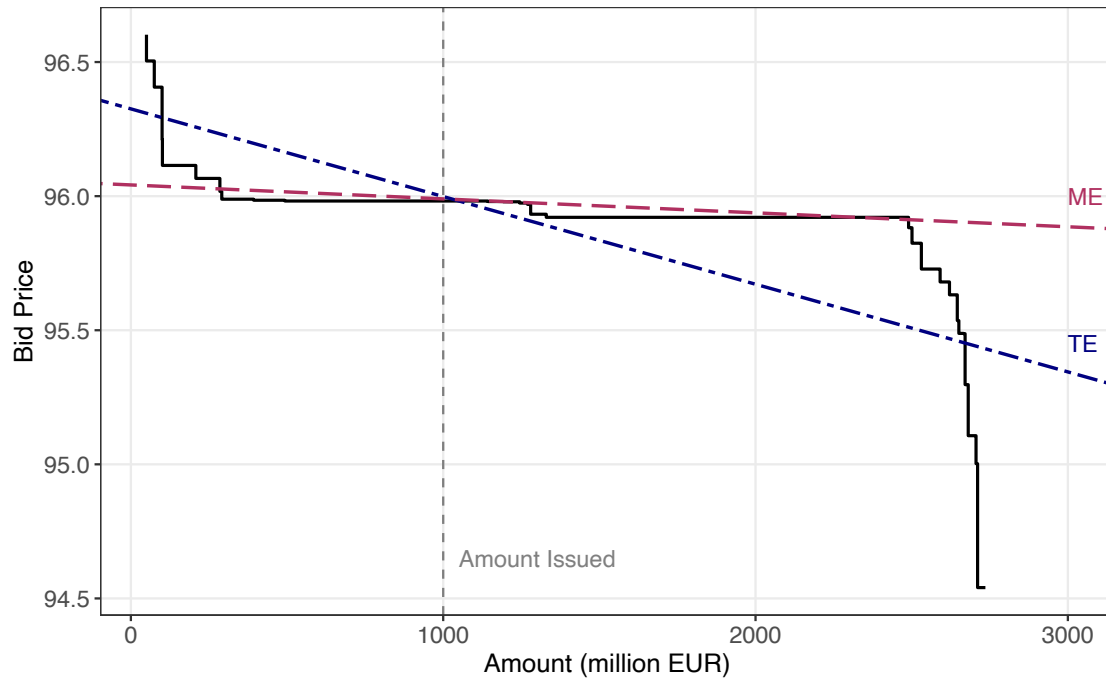


Figure 21: Comparison of ME and TE in a given auction during the crisis period

Figure 21 depicts all bids for the auction of May 2, 2012 of a 12 month Treasury Bill. The figure also depicts the slopes used to compute ME and TE. In this auction, the value of ME ( $\times 10^2$ ) and TE ( $\times 10^2$ ), are respectively 0.01 and 0.34. In the figure, we observe that the government avoids the cliff in the bid function. As such ME is not affected by it. The slope used to estimate TE, however, also uses the bids in the steep part of the bid function.

Figure 22 depicts all bids for the auction of July 21, 2010 of a 12 month Treasury Bill. The figure also depicts the slopes used to compute ME and TE. In this auction, the value of ME ( $\times 10^2$ ) and TE ( $\times 10^2$ ), are respectively 0.90 and 0.32. In this figure, we observe that the government accepts bids in the steeper part of the bid function. As such, the local inverse elasticity estimate, ME, is larger than TE.

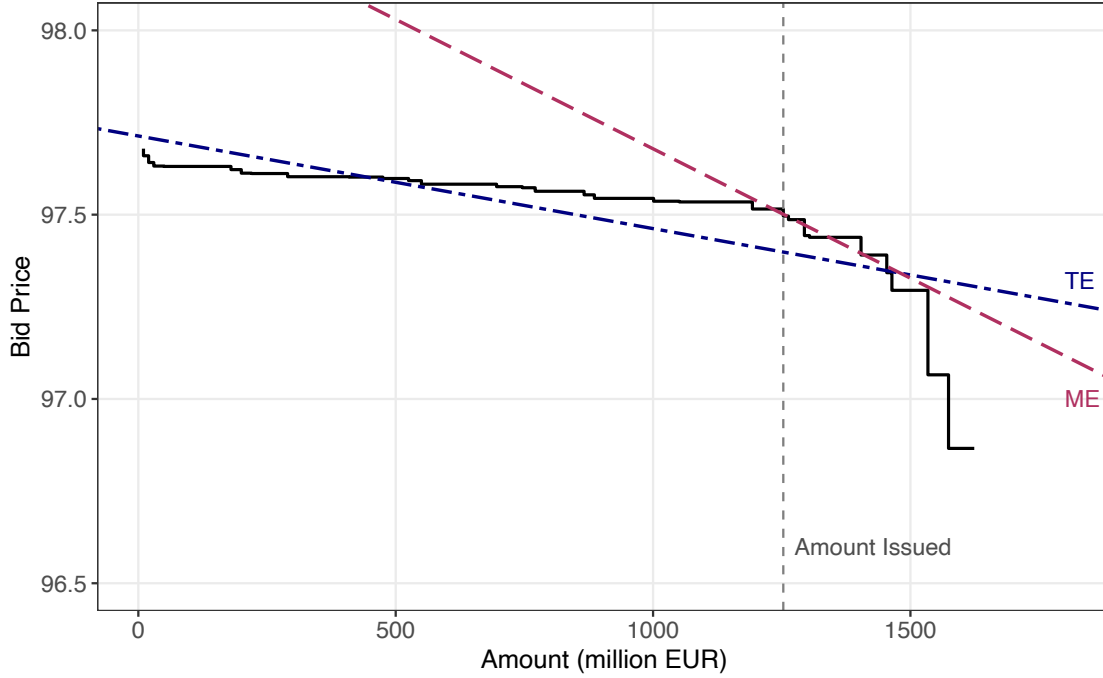


Figure 22: Comparison of ME and TE in a given auction during the crisis period

## Appendix C - Resampling Procedure

The procedure for a given investor  $i$  and auction  $t$  is summarized as follows:

1. Fix a bidder  $i$  among the potential  $N_t$  bidders in auction  $t$ .
2. From the sample of  $N_t$  bid vectors in the data set, draw a random sample of  $N_t - 1$  with replacement, giving the same probability  $1/N_t$  to each bid vector in the original sample.
3. Construct the residual supply function generated by these resampled bid vectors.
4. Find the market clearing price.
5. Repeat steps 1 – 4 a large number of times.

Note that each time step 3 is reached one has a state of the world from the perspective of the fixed bidder: a possible vector of private information<sup>18</sup>. By repeating steps 1 through 4 one hopes to get different possible states of the world in order to properly estimate the distribution of the market clearing price from the perspective of the fixed bidder<sup>19</sup>.

When resampling it is important to note that not all potential dealers submit bids in a given auction. As such, one also needs to resample empty bid vectors.

The figure below illustrates the resampling procedure and different aggregate states. Suppose that each

<sup>18</sup>This relies on the modeling assumption that all bidders are identical ex-ante apart from the realization of the private signal. Figure 23 in the Appendix illustrates potential different states of the world taking into account the number of potential participants in a given auction.

<sup>19</sup>The procedure also relies on having a large number of bidders (and private signals) to resample from. In that sense, to have more observations to resample from I bundle two consecutive auctions.

color is a type (private signal), further let the borderless grey agent be a dealer that decides not to bid after the realization of their signal. Each sample depicted in the figure, together with a realization of the supply for the security being auctioned, represents an aggregate state. Importantly, the non participating agent, an empty vector, is also included in the resampling pool. That happens as the agent that decides not to participate does so due to the realization of the private signal, the only differentiating factor across dealers. As such, non participating dealers are part of the aggregate state, as they constitute an element of the vector of private information.

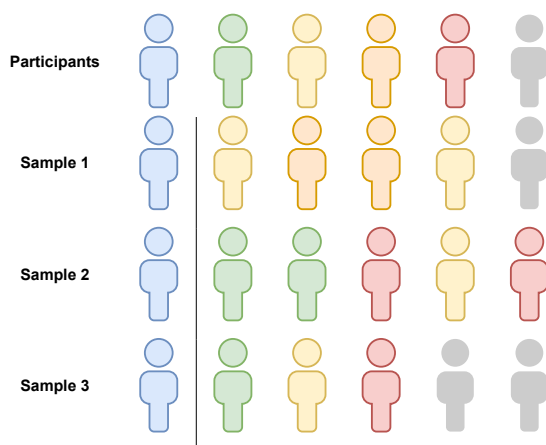


Figure 23: Illustration of the resampling procedure and different aggregate states