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

Ricardo Alves Monteiro*

This version: May 2022†

Abstract

Does demand for sovereign debt change during high default risk events? Using a dataset containing individual bids on Portuguese debt auctions, I investigate whether investors' demand for sovereign debt changes during the debt crisis. I find that aggregate bid functions are, on average, five times more elastic leading up and during the crisis. That is, on average, in order to increase the amount raised by 1%, the price would need to decrease, in percentage terms, by five times more than it had before the crisis. I then decompose the changes in demand into two components: a fundamental component, due to changes in valuation, and a strategic component, that arises from investors' market power. Although the role of market power is negligible in normal times, it gets more pronounced leading up and during the crisis. The auction mechanism loses efficiency during that period as the government is not able to extract the full surplus from strategic investors. At their peak, inefficiency costs jump to 0.6% of the issued amount. Finally, I discuss a possible mitigation strategy. Everything else constant, shorter maturities should be used to avoid higher inefficiency costs.

***Keywords**: Treasury auctions, default risk, demand, market power, inefficiency costs.*

***JEL**: D44, E62, F34, G12, H63.*

---

*University of Minnesota, e-mail: alves030@umn.edu. 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 Basseto, Daniel Belchior, José Cardoso da Costa, Cristina Casalinho, Doireann Fitzgerald, Stelios Fourakis and participants at the Minnesota-Wisconsin International Macro Workshop for excellent comments. All remaining errors are my own.

†First version: September 2021. Preliminary and incomplete, comments are welcome.

1 Introduction

In the sovereign debt literature, the government is typically seen as the only strategic player. Investors play a passive role in the process of issuing debt, as they are willing to lend as long as they break even. While this competitive behavior from investors is deemed reasonable when thinking about the secondary market, it is usually not the case in the primary market. In fact, given the typical mechanism chosen by governments to issue debt - through auctions with a limited number of participants - it is not obvious that investors are not strategic players in the debt issuing game being played. This apparent disconnect between the usual assumptions and the actual mechanism of issuing debt motivates this paper. The goal is to have a better understanding of how investors' demand for sovereign debt with different maturities evolves around high default risk events, while taking into account the non competitive nature of the market.

Does demand change during high default risk events? If so, how much of this change is due to a change in valuation and how much is due to investors' market power? Here I argue that investors' strategic considerations play a relevant role in determining the demand faced by the government in the debt issuing game being played.

To answer these questions, I use a dataset containing all primary market bids on Portuguese sovereign debt auctions from 2003 to 2020 and present a quantitative exercise for Portugal around the sovereign debt crisis of 2010-2012. Portuguese debt is mainly issued through auctions with a limited number of potential participants. The non-competitive nature of the auctions provides investors with market power: their bid functions influence the market clearing price and they internalize that effect.

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. That can be attributed to the decrease in the marginal price for larger amounts of debt. That is, the marginal price of the auction is significantly lower than the average auction price. Conversely, marginal yields are significantly higher than average yields. The government is not willing to issue a larger amount, as there is a big drop in price (jump in the yield) for those extra units. Figure 1 shows the emergence of this trend for 12 month treasury bills during the crisis period.

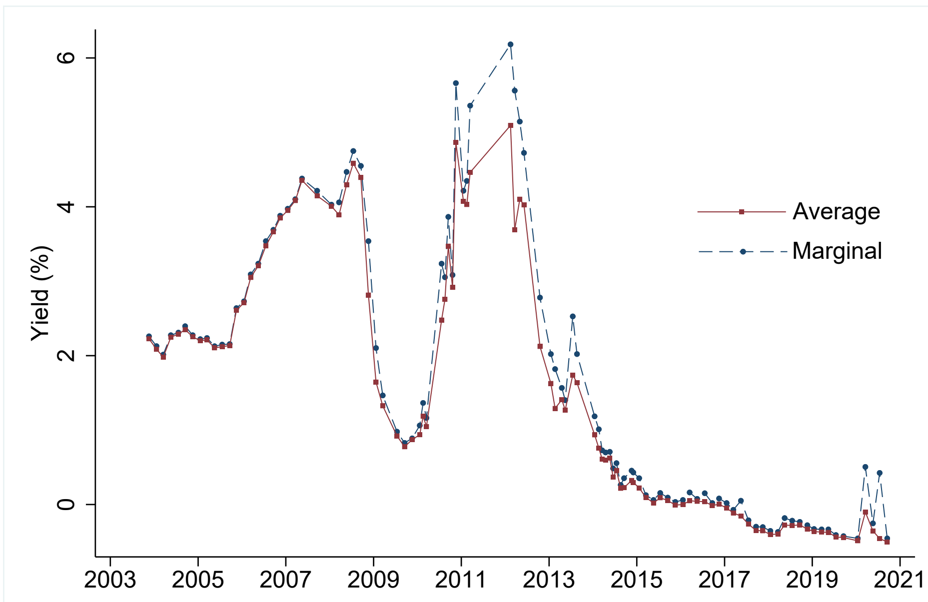

Figure 1: Marginal and Average Yields of 12 month Treasury Bills

More important than the spread between the marginal and average prices is the evolution of the price elasticity, a key statistic for the government, a monopolist in the debt market, as it pins down the optimal choice. I find that the bid schedules for both short and long maturities get significantly more elastic leading up and during the sovereign debt crisis. That is, in order to increase the amount raised by 1%, the price needs to decrease, in percentage terms, by more than it had before the crisis. As a result, the government issues less than it originally expected.

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?

To filter the data, I will introduce an environment based on Wilson (1979) framework, and more specifically on Hortaçsu and McAdams (2010) and Kastl (2011). The auction model uses a multiple price protocol - pay as you 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 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 condition I 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 the fact that each can influence the price.

I find that market power plays a limited role during normal times. However, leading up and during the crisis the wedge between bids and valuation gets more pronounced. The mechanism follows. Bids are more dispersed leading up and during the crisis. The increased dispersion of bids implies that the subjective distributions of the aggregate state are less precise than before the crisis. Particularly, the likelihood ratio of step $k+1$ in a bid function being a winning step, relative to, step k being the last winning step, increases. This leads to a larger wedge between bids and valuations as investors bid 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 biding below their willingness to pay. Let the inefficiency be measured as the ratio of the aggregate wedge over the amount raised in a given auction. At their peak, inefficiency costs go up to 0.6% of the issued amount, during the crisis.

Finally, a more normative analysis should follow. 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. A more thorough analysis of optimal maturity choice accounting for the inefficiency costs of the mechanism is left as future research.

Related Literature

The quantitative work that started with Aguiar and Gopinath (2006) and Arellano (2008), 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[1]. Since those initial quantitative models, there have been substantial developments in the literature with the study of maturity choice, self fulfilling crisis and the role of timing on the existence of such multiple equilibria. Examples of such are Arellano and Ramanarayanan (2012) and quantitative models based on Cole and Kehoe (2000), such as Bocola and Dovis (2019). More recent work includes Ayres et al. (2018) and Aguiar et al. (2019). Importantly, until recently the focus was on the government's strategic choices under assumptions that are consistent with the secondary market for sovereign debt. This paper aims to provide a broader understanding of investors' strategic considerations in the context of the primary market for sovereign debt.

Cole et al. (2021) motivation 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

[1] See Aguiar and Amador (2014) for a survey.

in demand during a high default risk event and the role that market power plays on the evolution of bids.

Bigio et al. (2021) use 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.

The environment I will use in this paper to filter the data is based on the seminal work of Wilson (1979) and more precisely on Hortaçsu and McAdams (2010), and Kastl (2011). Kastl (2020) provides a review of the literature and methods applied to financial auctions and particularly to treasury bond auctions.

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

2 Data: Background and Evidence

The data was provided by the *Agência de Gestão da Tesouraria e da Dívida Pública - IGCP, E.P.E.*, the institute responsible for issuing and managing public debt in Portugal.

Issuance of treasury bills in the primary market is done through auctions, while treasury bonds are launched in syndicated operations and then issued through auctions as well. Both treasury bills and bonds were auctioned using a multiple price protocol up to 2011.

From 2011 onward, treasury bonds were auctioned using a single price protocol.

Only institutions that have been granted the status to do so by the agency may participate in treasury bill and bond auctions. From here on out these are denominated ”dealers”. Dealers are permitted to submit multiple bids[2] as long as the total value does not exceed the upper limit of the overall amount announced for the competitive phase of the auction.

The data comprises all bill and bond auctions held from 2003 and 2004, respectively, and up to 2020. As such, the time series includes the sovereign debt crisis of 2010-2012, which enables us to analyze changes in demand during that period. Most importantly, the data comprises all individual bids (price and amount) that were placed in each auction, even if they were not executed.

Maturity	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

Table 1: Summary Data on Treasury Bond and Bill auctions

One can observe 400 Treasury bill auctions and 161 Treasury bond auctions. The most common maturities are 12 and 3 months for the treasury bills and 10 and 5 years for the treasury bonds. In bill auctions the number of bids averages 39 and in bond auctions it averages 60. Table 1 presents some summary data for the most common bill and bond auctions.

[2]For treasury bill auctions each dealer may submit up to five bids per auction, for treasury bond auctions a limit is not specified.

2.1 Motivating Evidence

Below I show some evidence that motivates the shift in investors' demand for Portuguese sovereign debt while approaching the sovereign debt crisis.

2.1.1 Short maturities

Figures 2 and 3 respectively present the aggregate bid functions and the amounts raised by the Portuguese Government in 3 month and 12 month treasury bill auctions over time. The analysis focus on the crisis event so the figures show the evolution of demand leading up to the sovereign crisis and the recovery period afterwards. Each panel represents an individual auction at a given date. Within each panel, one can see the aggregate bid function, the amount that the government is able to raise at each given price. The aggregate bid function for 3 month treasury bills is obtained by aggregating individual bids (amount that the investor is willing to buy) at each price[3]. The prices are normalized by the marginal price of the auction, i.e. the minimum price accepted in the auction. Finally, the red line identifies the amount raised in each auction.

Before the crisis, up to 2010, the prices of individual bids present almost no dispersion. Leading up to the crisis this is no longer the case. Particularly, there are several bids with prices significantly below the auction price. Starting in 2013, there is a recovery of the schedule to its previous shape, with almost no dispersion in bids.

With this data, one can see how the price elasticity, i.e. the necessary percentage change in price so that the amount raised by the government increases 1%, is changing during the same period.

Interestingly, for both maturities the average price elasticity increases during the sovereign debt crisis and returns to previous levels afterwards.

[3] In the appendix I present the schedules using yields instead of prices.

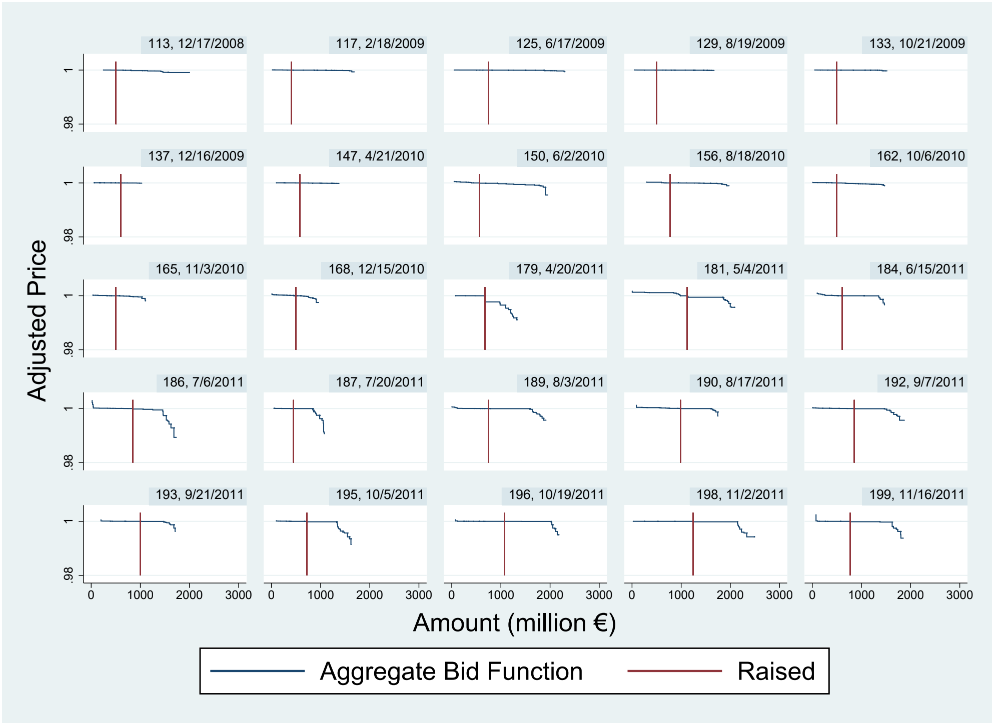

(a) Leading to the sovereign crisis

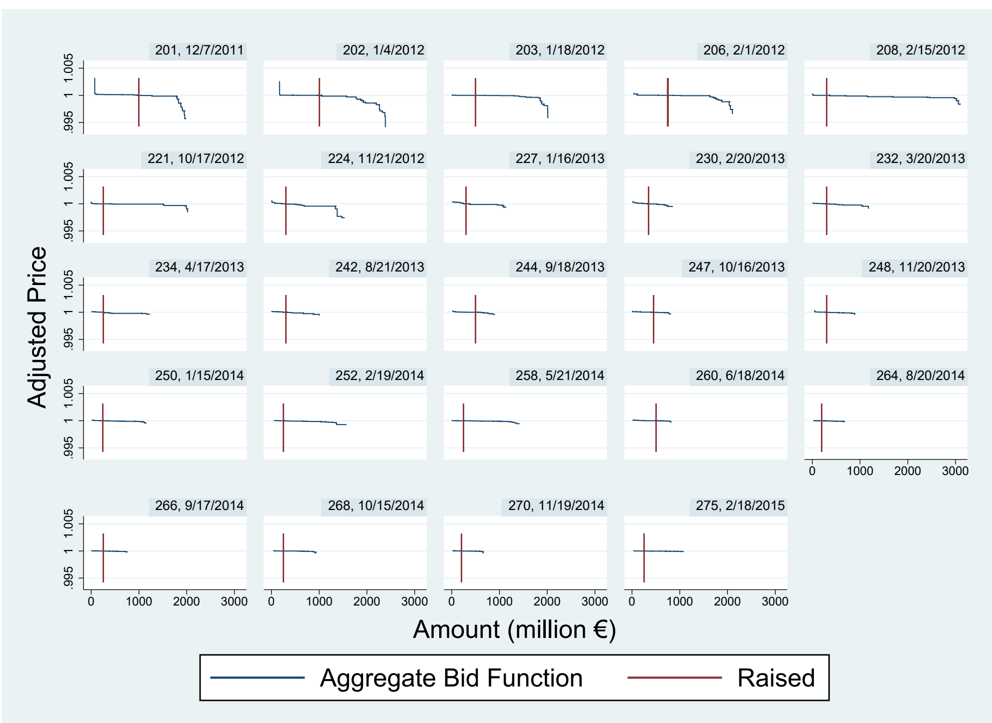

(b) Recovery

Figure 2: Bids for 3 month treasury bills in the primary market

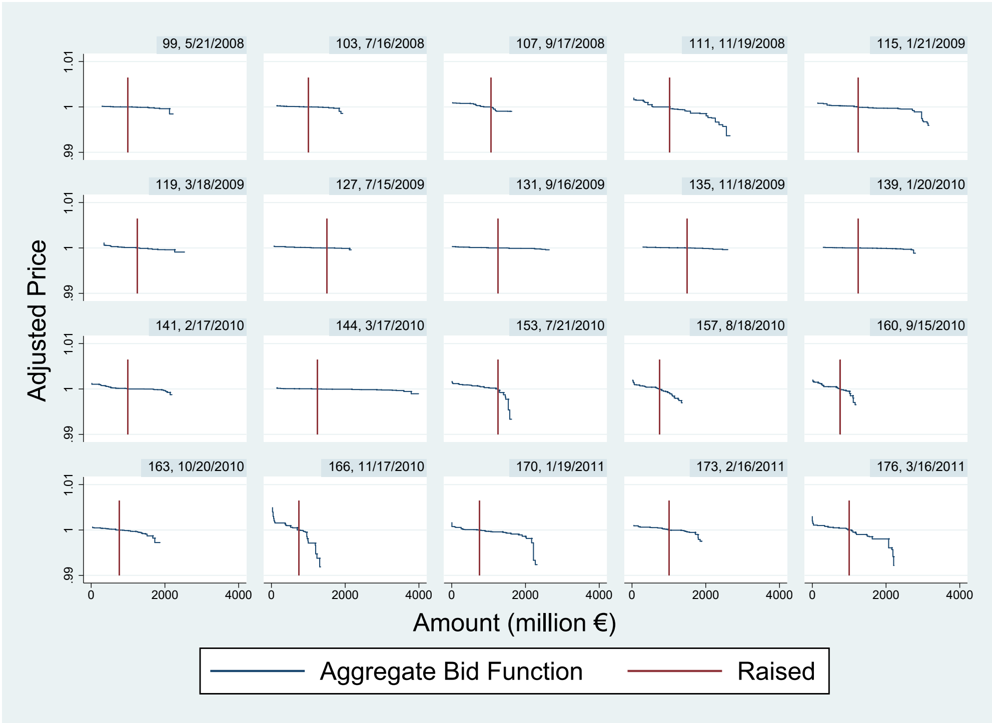

(a) Leading to the sovereign crisis

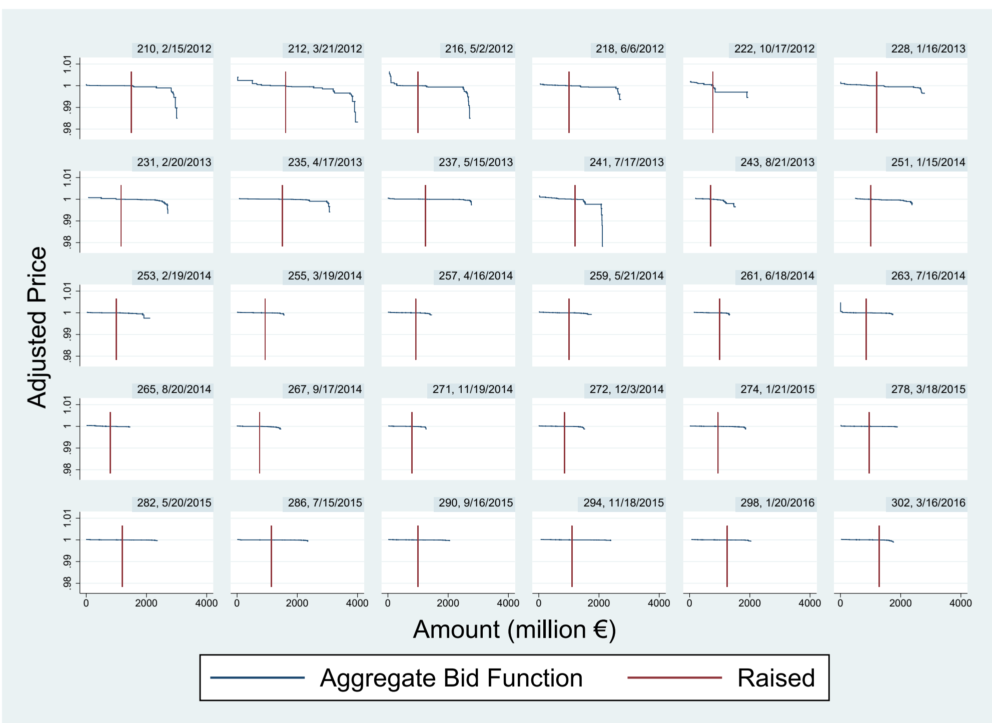

(b) Recovery

Figure 3: Bids for 12 month treasury bills in the primary market

2.1.2 Long maturities

Figures 4 and 5 present the schedules and the amounts raised 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 schedules leading up to the sovereign crisis and the recovery period afterwards[4].

It is worth noting that, in contrast to what happens with treasury bills, it seems that the schedules for treasury bonds do not recover to the pre-crisis shape. One potential explanation for this behavior is the fact that the auction protocol for treasury bonds switched from a multiple price protocol to a single price protocol. It is a well known fact that the winner's curse is a potential outcome of a multiple price auction. By moving from a multiple price to a single price protocol, bidders are less likely to shade their bids as they will end up paying the marginal price regardless.

2.1.3 Discussion

Investors' bid functions are indeed changing leading up and during high default risk events. Particularly, as the sovereign debt crisis approaches, the schedule tends to become more elastic in the sense that, at a given pair (price, amount raised), in order to increase the amount raised by 1% on average the price needs to decrease, in percentage terms, by more than it was needed before the crisis.

This increased 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 single price protocol, whereas all auctions before the crisis were executed under

[4]The elements within the figures are as in Figures 1 and 2 above.

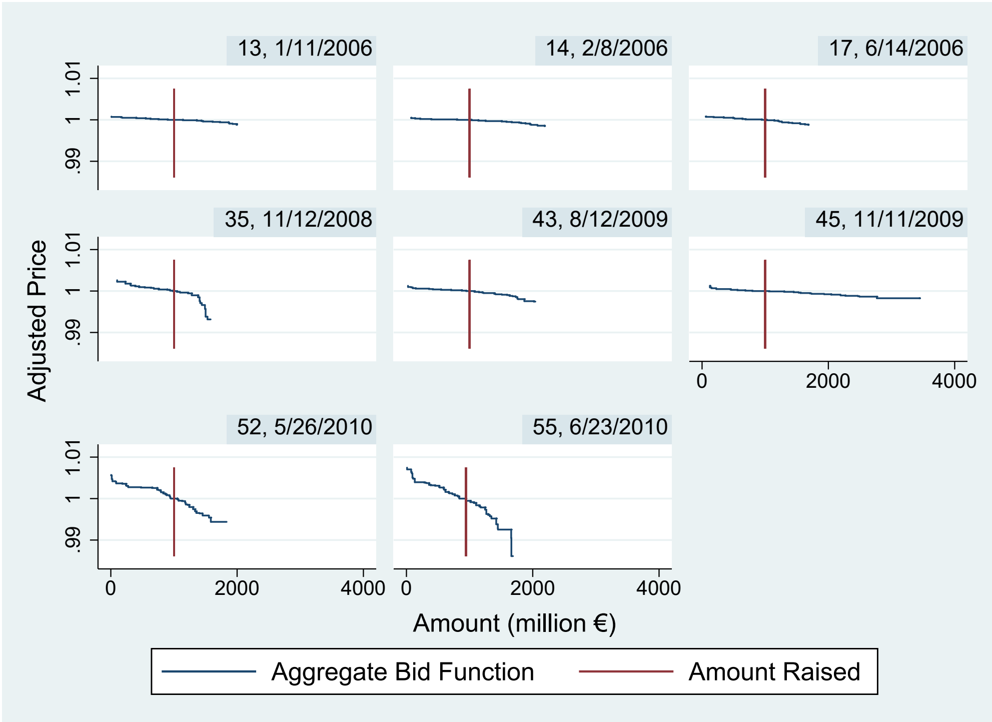

(a) Leading to the sovereign crisis

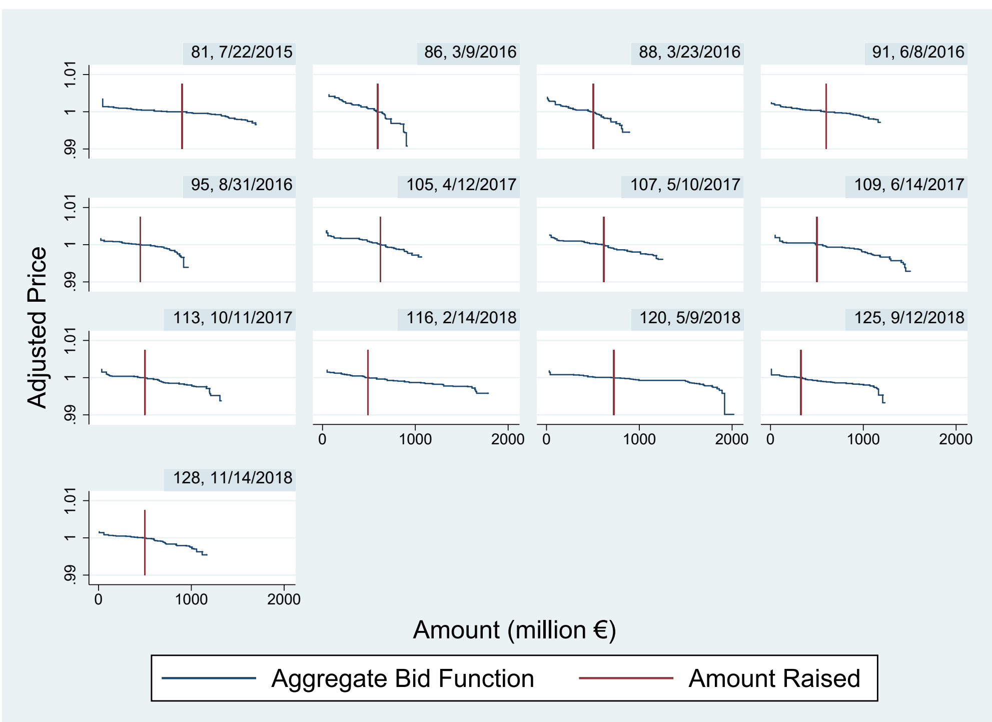

(b) Recovery

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

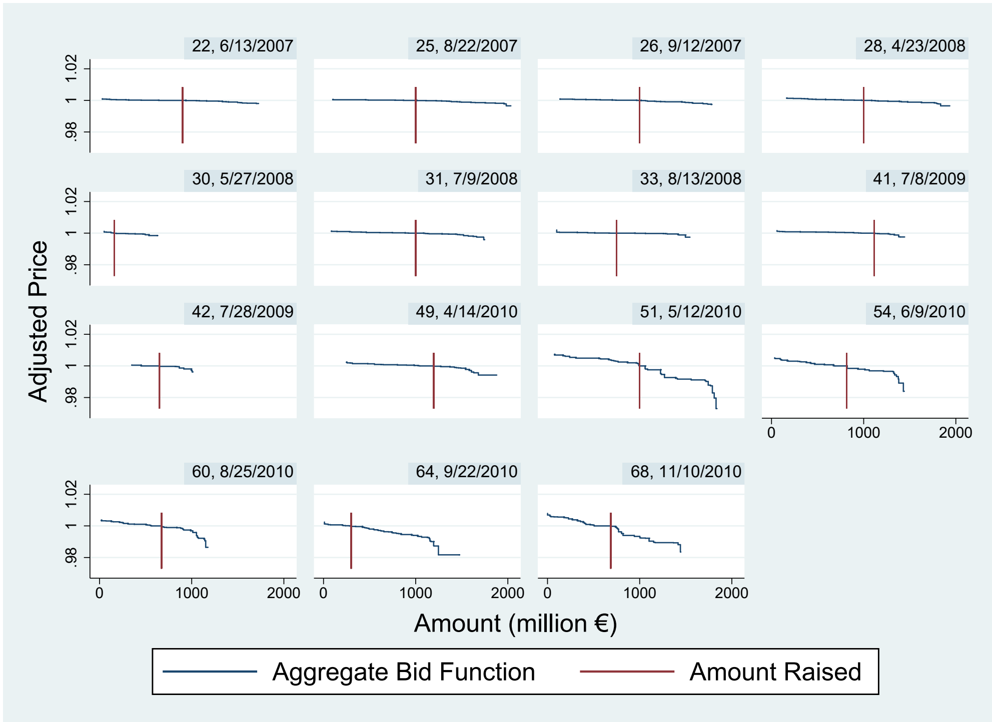

(a) Leading to the sovereign crisis

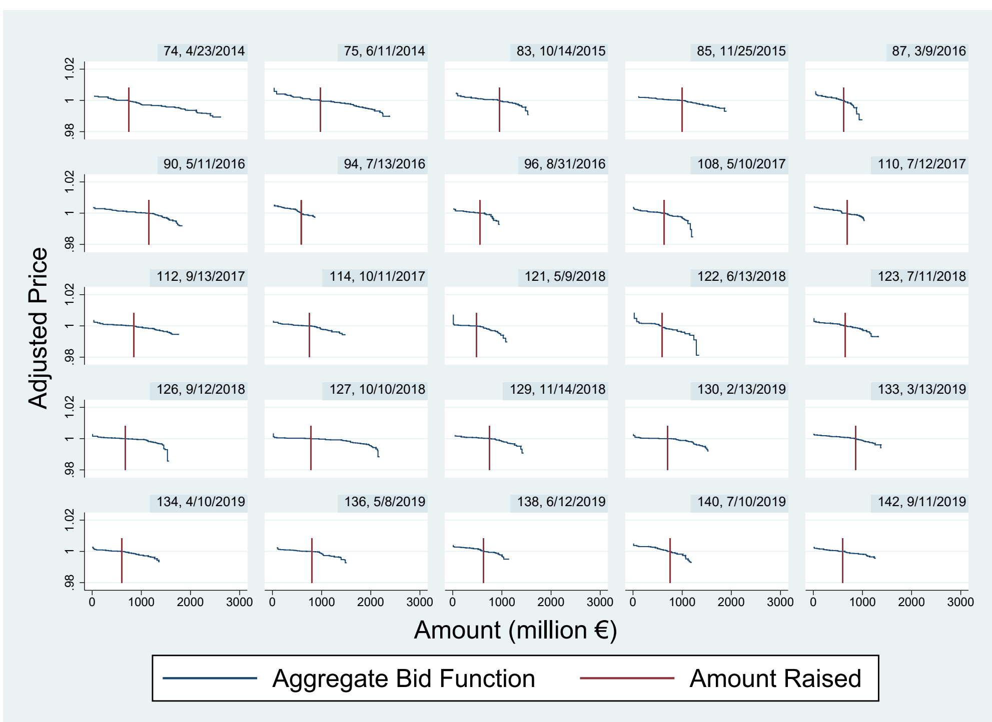

(b) Recovery

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

a multiple price protocol. As a result, one cannot disentangle the recovery from the crisis from the change in auction protocol[5]. One can argue that the described change in the auction mechanism would lead to a smaller wedge between the bids and valuation (less shading) as there is no winner's curse as in the multiple price protocol. This in turn leads to potentially steeper bid functions, i.e. with less shading on the first steps. Noticeably, this argument can help us understand the apparent lack of recovery of the schedule to its previous shape.

2.2 Elasticity Measures

Let q be the marginal price and B the amount of debt raised in an auction. Then, the price elasticity is defined as $\mathcal{E} = \frac{\partial q}{\partial B}\frac{B}{q}$, the necessary change in price such that the amount raised increases by 1%. A small (absolute) value of this elasticity means that large increase in the amount raised are associated with small decreases in price.

In order to compute this elasticity we need to estimate the slope coefficient $\frac{\partial q}{\partial B}$. I follow Albuquerque et al. (2022) in computing the marginal elasticity of an auction, the main measure of elasticity (ME) used throughout this paper. This measure uses bids from untapped liquidity, next to the marginal price of the auction. 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 bid functions, I use the quantity price pairs to estimate a linear regression model of the price on the amount raised and a constante. The slope coefficient in the model is an estimate of $\frac{\partial q}{\partial B}$. To get the elasticity one just multiplies the slope estimate by the ratio of the amount raised to the marginal price of the auction.

I compute three other measures of elasticity. First, I compute total elasticity (TE), that

[5] 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 single price protocol.

differs from ME in that it uses all the aggregate bid function points to estimate the slope coefficient from a simple linear regression model. Secondly, I compute the point elasticity (PE) that looks at the change in the amount raised going from the auction price to the price immediately below. This is the least reliable measure as it relies on a single variation and for that reason includes a significant amount of noise. Finally, I compute the average point elasticity across the aggregate bid function for each auction.

The focus on marginal elasticity is arguably the most reasonable for two reasons. First, it gives us the most conservative changes for the elasticity over time; secondly, and perhaps more important, the marginal elasticity gives us an estimate of the elasticity around the marginal price of the auction taking into account the approximate slope in that region. This last argument turns out to be relevant, particularly leading up and during the crisis, where the price schedules tend to present a quasi-kink. This distorts the slope estimate around the auction price and it is usually the case that TE is greater than ME. Figure 18 shows the comparison of the elasticity measures over time. Unsurprisingly, PE is the noisiest measure considered. ME and TE are relatively close with TE greater than ME during the crisis period[6].

Focusing on shorter maturities (treasury bills), Figure 7 shows the marginal elasticity time series for 3 and 12 month treasury bills, as well as the average increase during the crisis period considering all treasury bill maturities. Note that the longer maturity bills have a more pronounced increase in elasticity during the crisis period when compared to the shortest maturity. Taking into account all treasury bill issuances, we have a new fact to document: *on average, the price elasticity increased by a factor of 5 leading up and during the crisis*.

[6]Figure 18 in the appendix illustrates this difference in a given auction during the crisis period

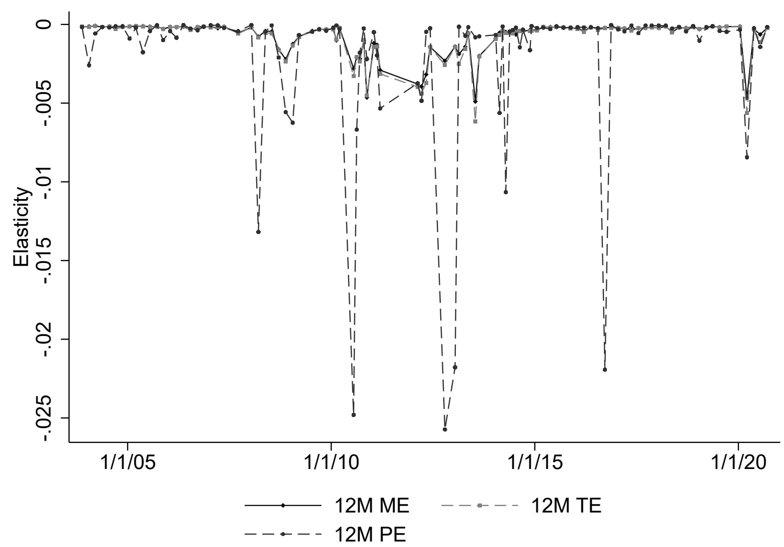

Figure 6: Comparison of ME, TE and PE over time

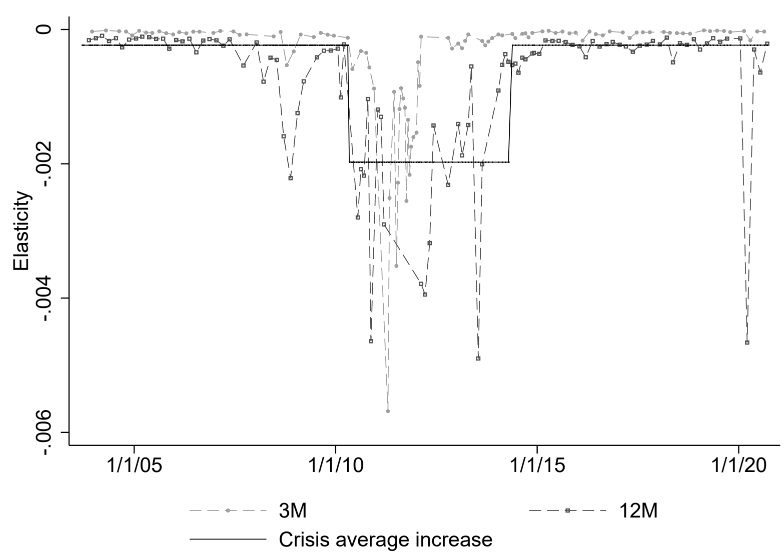

Figure 7: Average increase in ME for treasury bills during the crisis

3 Environment

In this section, I present the model that I am going to use 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 (2011). Importantly, it enables the understanding of what is driving the shifts in demand.

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. This 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 his bid function as "shading".

The shading term can be thought of as the investor's market power. Note that, as there is a limited number of investors in any given auction, the strategy of a single investor may change the equilibrium price of the security being auctioned[7]. In this sense, investors might be pivotal. The limited number of potential dealers and consequential lack of perfect competition among buyers is crucial for the existence of market power.

In a first price auction with multiple price protocol (pay as you bid) it is intuitive that the shading term is likely to be positive. This is so to avoid the winner's curse[8].

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

[7]Suppose the targeted amount for the auction is 400 million euros. Consider two scenarios: (i) four dealers bid for 150 million euros, two at 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 strategy affects the market clearing price that is €99 instead of the €98.9 in (i).

[8]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 she could have paid only €98.

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*[9] *distributed on $[\underline{B}, \overline{B}]$ 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 follows I'll assume that $v(b, s_i, s_{-i}) = v(b, s_i)$, that is values are assumed to be private. Furthermore, it follows from assumption 1 that I will work in the special case of independent private values (IPV).

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 a 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 $i = 1, \ldots, N$ has an action set:*

$$A_i = \left\{ \begin{array}{c} \left(\vec{p}_i, \vec{b}_i, K_i\right) : \dim(\vec{p}) = \dim(\vec{b}) = K_i \in \{1, \ldots, K\} \\ p_{ik} \in P \equiv \{l\} \cup [0, \bar{p}],\ b_{ik} \in [0, 1],\ p_{ik} > p_{ik+1},\ b_{ik} < b_{ik+1} \end{array} \right\}$$

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

Let $V(b, s_i) = \int_0^b v(x, s_i)dx$. Then, the expected utility of bidder i of type s_i employing strategy $a_i(\cdot \mid S_i \equiv \bigtimes_{i=1}^N s_i) \in A_i$ can be written as in the equation that follows.

[9] 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. Finally data suggests that often the target is not met. As such, it seems reasonable to treat the issued amount as a random variable.

$$EU(s_i) = \sum_{k=1}^{K_i} \Big[\overbrace{\Pr(p_{ik} > P^c > p_{ik+1} \mid s_i)}^{k \text{ is the last winning step}} V(b_{ik}, s_i) - \overbrace{\Pr(p_{ik} > P^c \mid s_i)}^{k \text{ is winning step}} \overbrace{p_{ik}(b_{ik} - b_{ik-1})}^{\text{added cost of step } k} \Big]$$

$$+ \sum_{k=1}^{K_i} \Pr(p_{ik} = P^c \mid s_i) \, E_{B, s_{-i} | s_i} \Big[V(B_i^c(B, \mathbf{S}, \mathbf{a}(\cdot \mid S)), s_i)$$

$$- p_{ik}(B_i^c(B, \mathbf{S}, \mathbf{a}(\cdot \mid S)) - b_{ik-1}) \mid p_{ik} = P^c \Big]$$

Where the random variable B_i^c is the market clearing quantity bidder i obtains if the state is (B, S) and bidders submit bids specified in the vector $\mathbf{a}(\cdot \mid S) = [a_1(\cdot \mid s_1), \ldots, a_N(\cdot \mid s_N)]$. 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.

3.1 Equilibrium

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

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

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

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 covered 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:

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

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.

4 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 others' bidding strategies affect bider 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:

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

Define an indicator of excess supply:

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

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

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

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

The estimator ξ 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 $p \in [0, \overline{p}]$.

Note that it is not feasible to compute ξ by summing over all permutations of bids. Instead, I will use the resampling procedure first proposed for the multi-unit auction environment in Hortaçsu (2002)[10]. 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[11]. 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[12].

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.

[10] The author shows the asymptotic properties of the estimator being used.

[11] This relies on the modeling assumption that all bidders are identical ex-ante apart from the realization of the private signal. Figure 19 in the Appendix illustrates potential different states of the world taking into account the number of potential participants in a given auction.

[12] 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.

Example 1 *In Figure 8 one can see, in panel (a) an illustration of the resampling procedure for Dealer 1 in a given auction, where each upward sloping residual supply curve corresponds to a different sample; and in panel (b) the corresponding market clearing price distribution.*

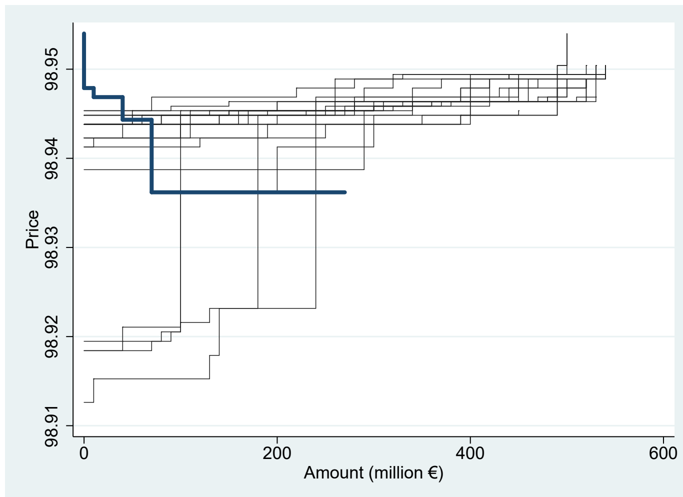

(a) Different realizations of residual supply for dealer 1

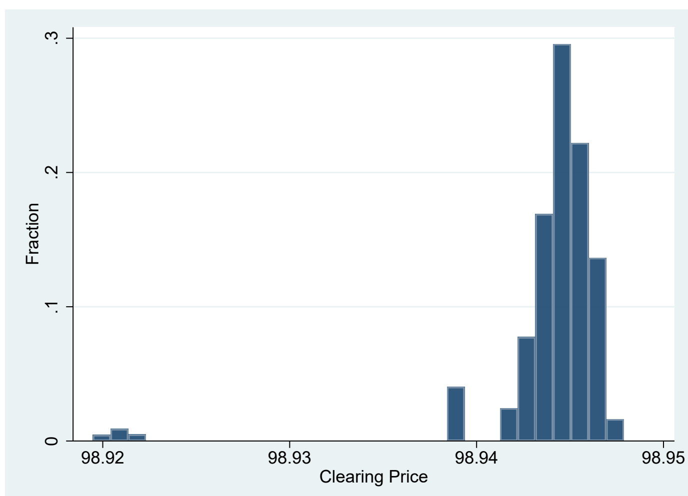

(b) Market clearing price distribution

Figure 8: Illustration of the resampling procedure for Dealer 1

5 Market Power and Inefficiency Costs

Having 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 9 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.*

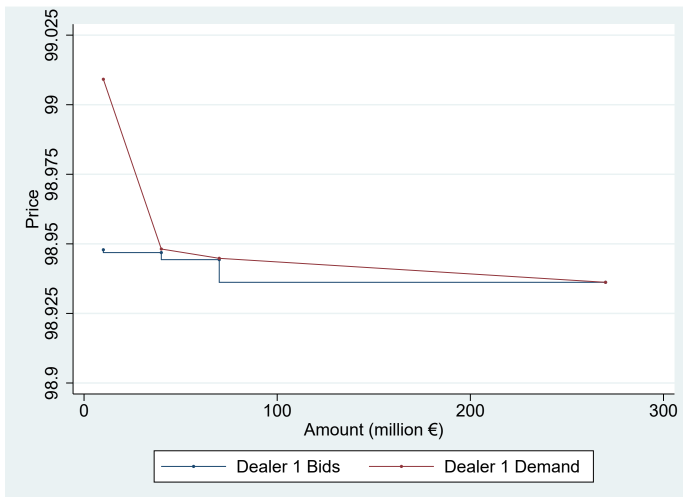

(a) An auction before the crisis

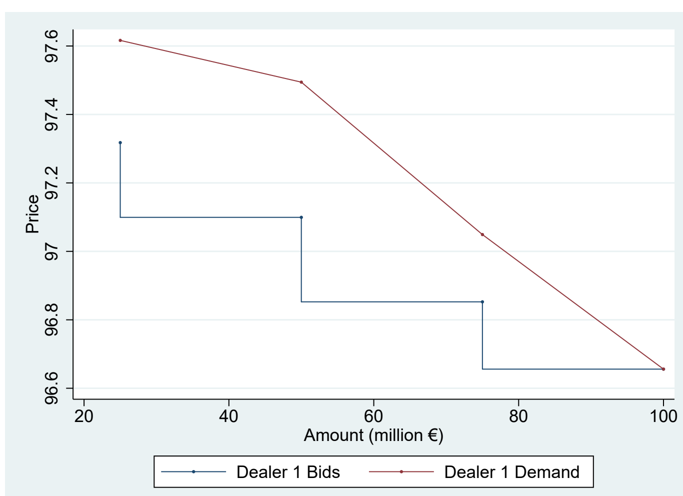

(b) An auction during the crisis

Figure 9: Valuation and bid function for Dealer 1

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[13], i.e. exclusively with ITM bids (and valuations), one can have a better sense of the effective average wedge.

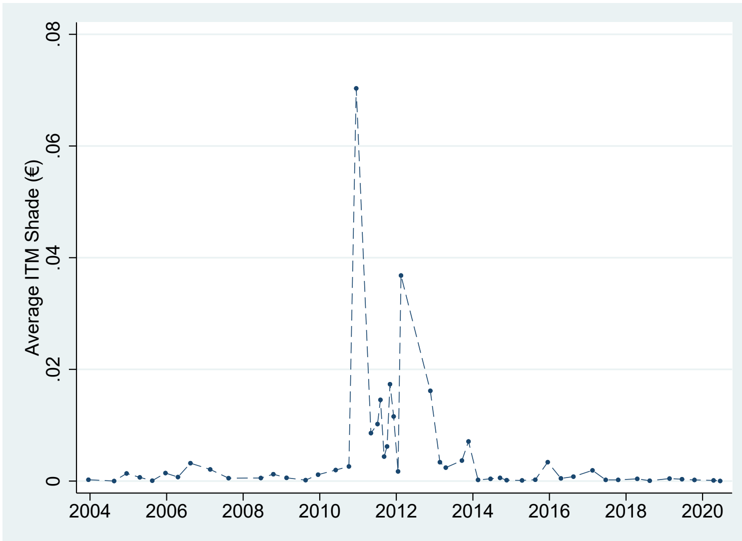

(a) 3 month treasury bill auctions

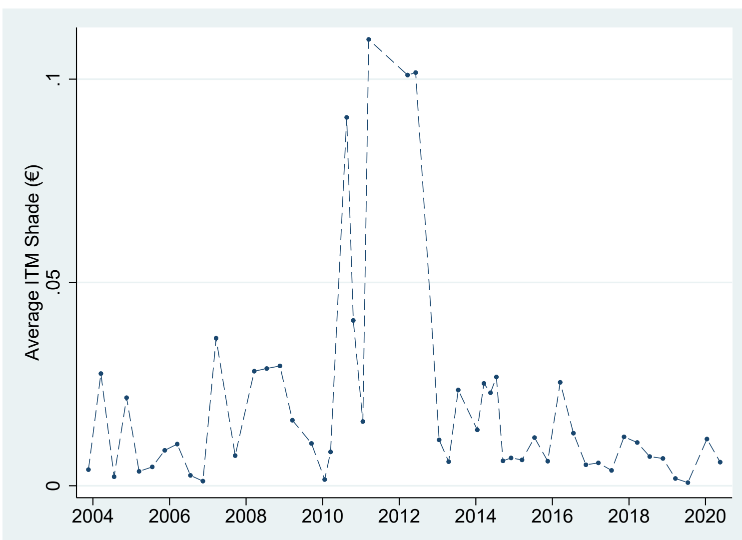

(b) 12 month treasury bill auctions

Figure 10: Average In the Money shade in treasury bill auctions

[13]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.

Figure 10 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.

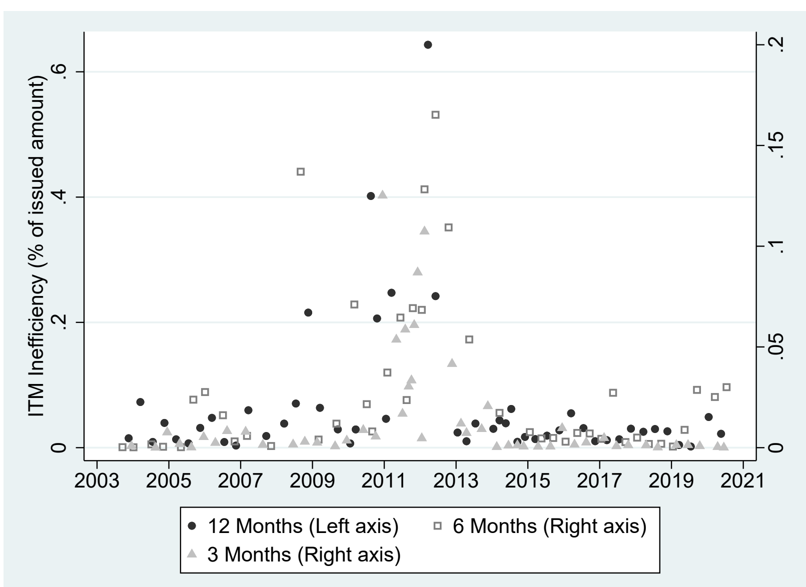

Figure 11: Inefficiency costs as a % of raised amount in treasury bill auctions

Figure 11 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.

5.1 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[14]. 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 $P[P^c < p|s_i],\ \forall p \in [0, \overline{p}]$, 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[15].

[14] 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.

[15] 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.

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

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

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

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

The increased dispersion in market clearing prices during the crisis likely leads to a higher $\theta(k, k+1 \mid s_i)$, 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 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 12 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 get a sense of the importance of the role of market power and consequent inefficiency by analyzing the dispersion of bids in an auction.

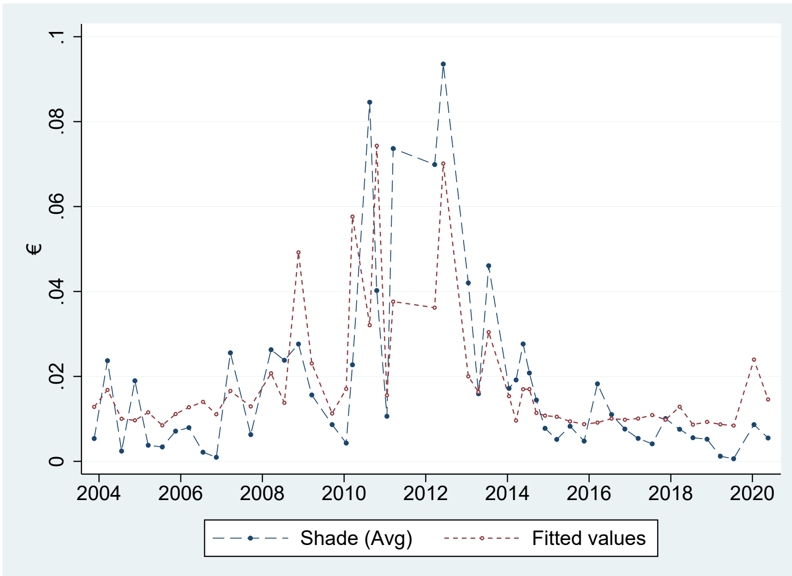

Figure 12: Prediction of Inefficiency with observables

6 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). The problem is simply:

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

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

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[16].

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 (α)	Crisis increment (β)
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 matu-

[16]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

rities. 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 0.16%. 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%.

As such, only taking into account 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 13 shows the evolution of inefficiency costs for 10 year treasury bonds. The figure shows a spike in the inefficiency costs as the crisis approaches. Granted that the main reason behind the decision to stop issuances of long bonds had most likely to do with the price drop that they faced, this trend also suggests that avoiding such long bonds is, everything else constant, a good strategy to mitigate inefficiency costs.

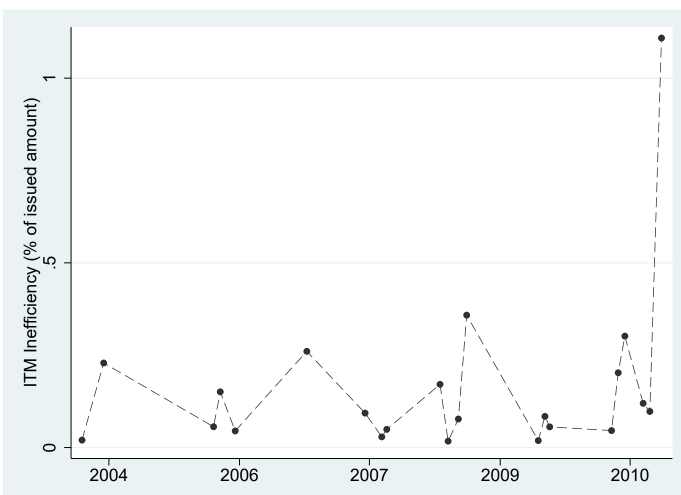

Figure 13: Inefficiency costs as a % of raised amount in 10 year treasury bond auctions

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

7 Concluding Remarks

Using micro data for Portuguese sovereign debt auctions, containing individual bids from 2003 to 2020, I document a key pattern in investors' demand during a high default risk event - the Portuguese sovereign debt crisis: leading up and during the crisis, bids get more disperse and the aggregate bid function faced by the government gets (5 times) more elastic. This is true across short and long maturities for the duration of the crisis. Afterwards, these schedules tend to recover to their pre crisis shape.

This fact fits the description put forward by the the Portuguese government: shifts in demand were responsible for lower than expected amounts issued during the crisis. Particularly, this is more evident as the distance between the marginal price and the average price of the auction increases leading up and during the crisis.

I then present 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 a given auction. Crucially, market power allows bids to differ from valuations. With the model, it is possible to filter the data and separate between 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 and during the crisis period. In fact, the shading terms for 3 and 12 month treasury bills, at their respective peeks, 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*". As the wedge itself, the ratio of such inefficiency costs as a percentage of the amount raised tend to be negligible in normal times, but go 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. A more thorough analysis of optimal maturity choice accounting for the inefficiency costs of the mechanism is left as future research.

Appendix A - Schedules

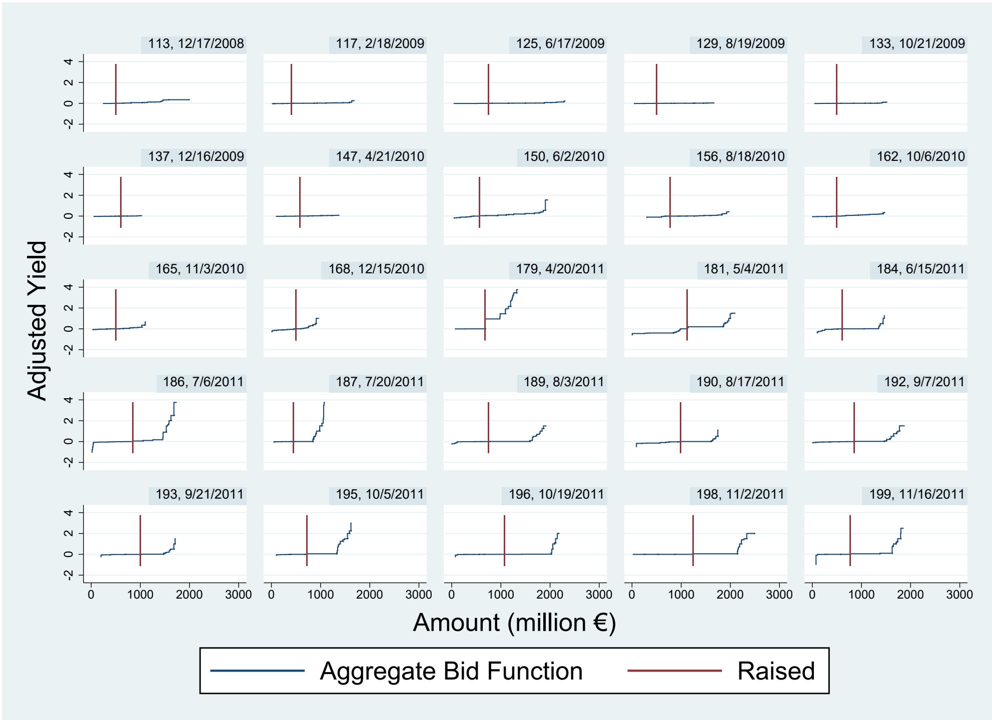

(a) Leading to the sovereign crisis

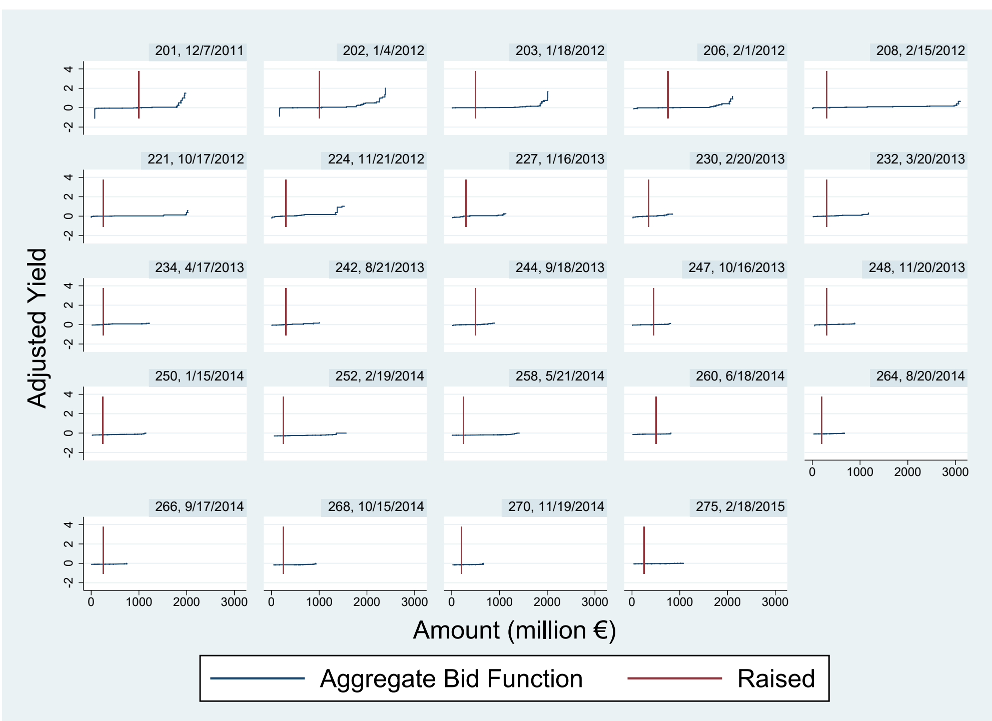

(b) Recovery

Figure 14: Demand schedule for 3 month treasury bills in the primary market

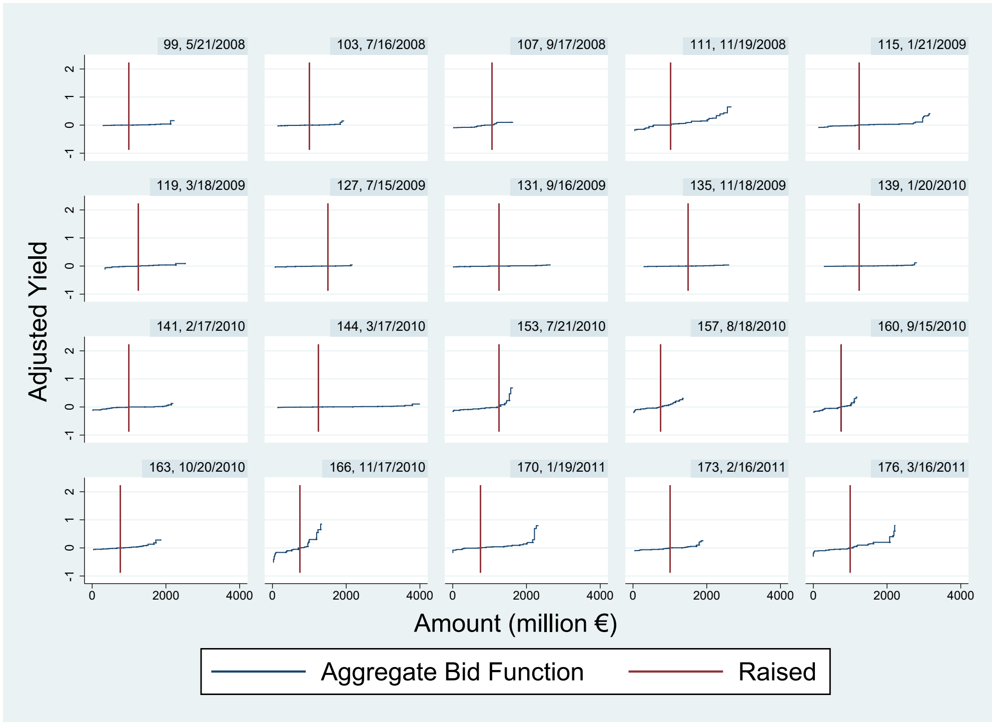

(a) Leading to the sovereign crisis

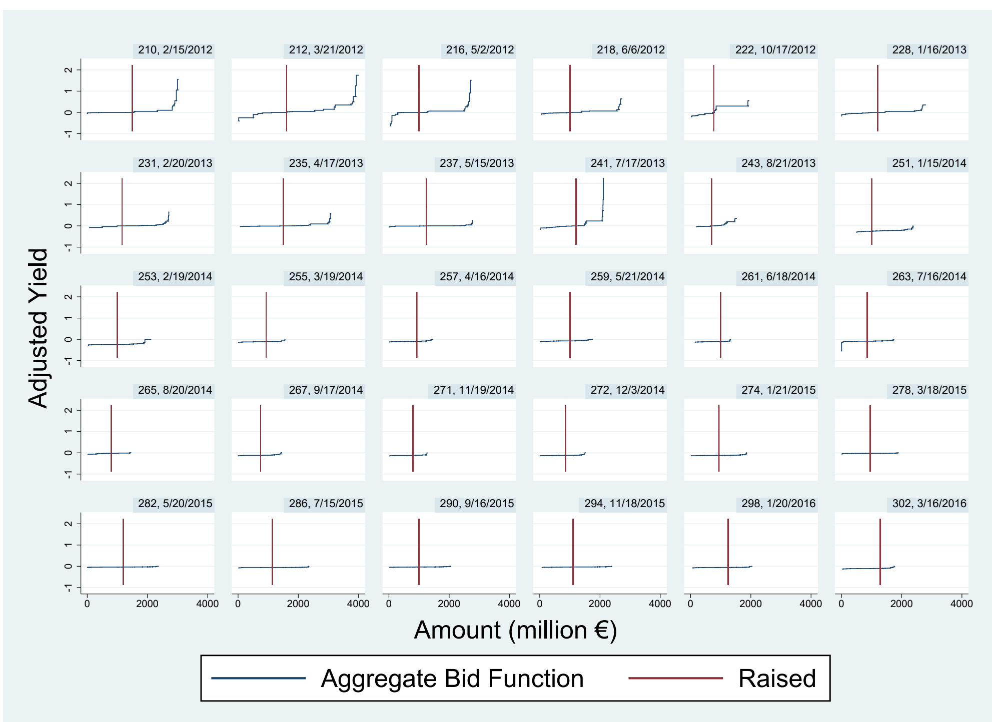

(b) Recovery

Figure 15: Demand schedule for 12 month treasury bills in the primary market

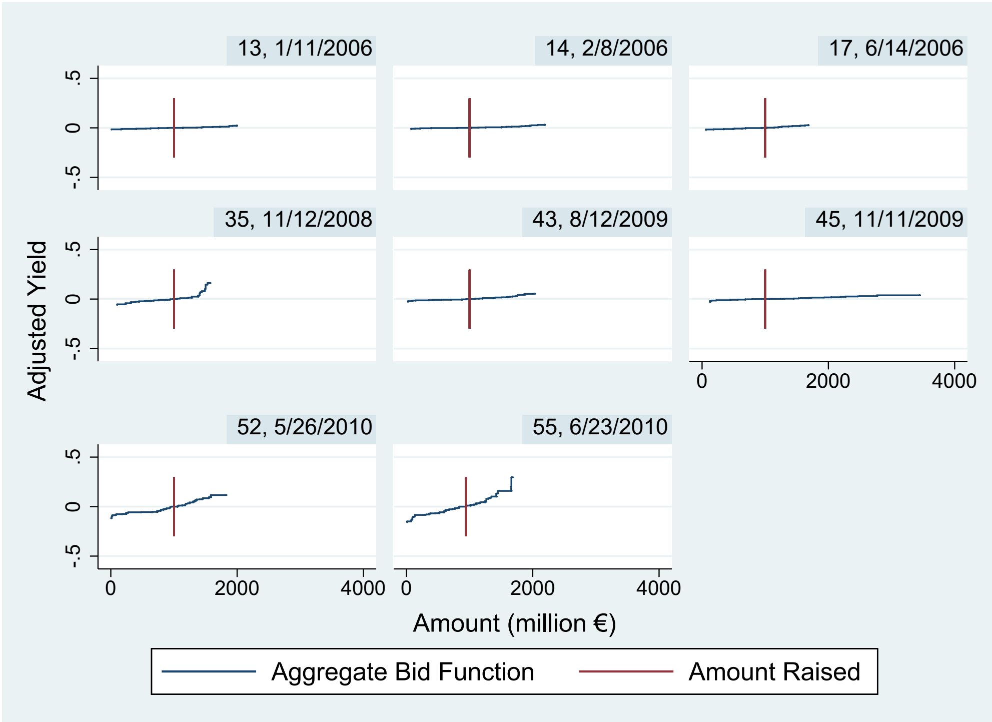

(a) Leading to the sovereign crisis

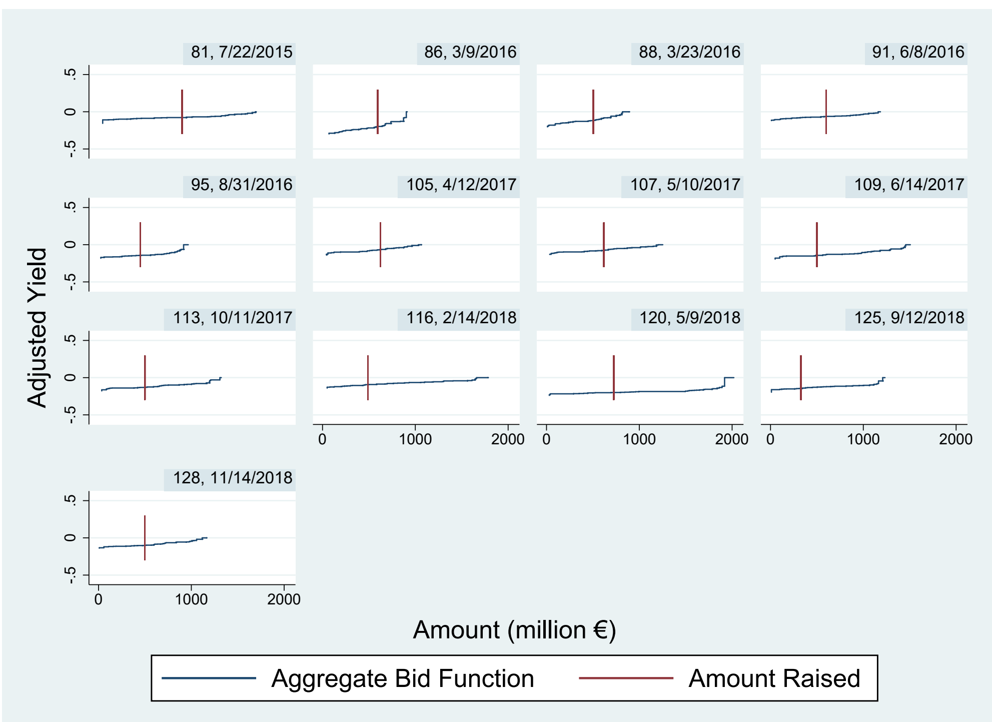

(b) Recovery

Figure 16: Demand schedule for 5 year treasury bonds in the primary market

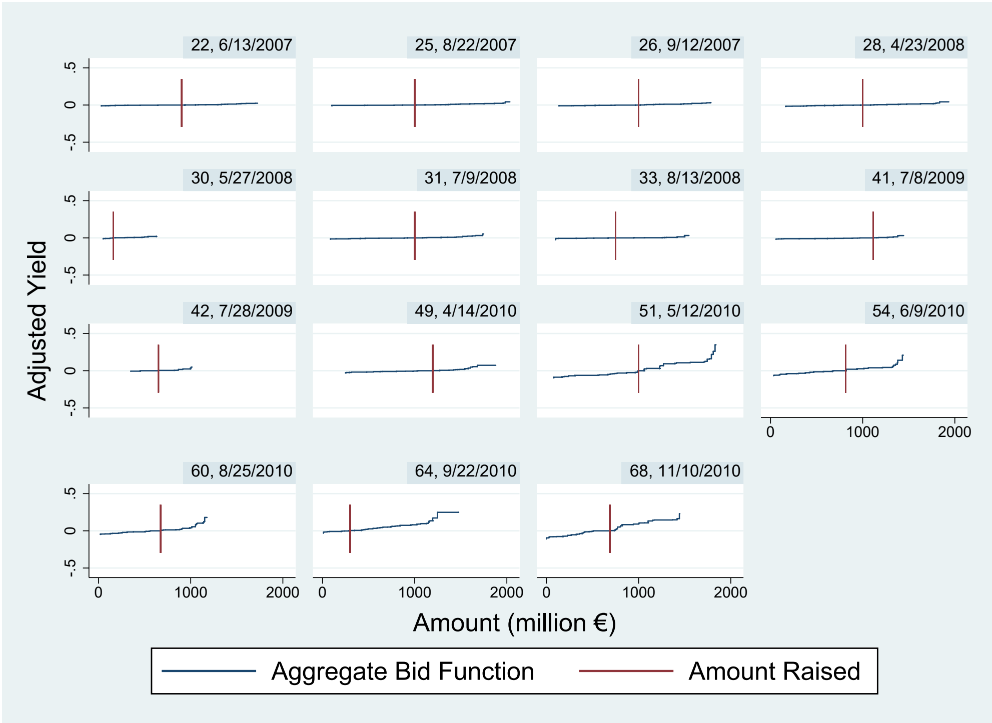

(a) Leading to the sovereign crisis

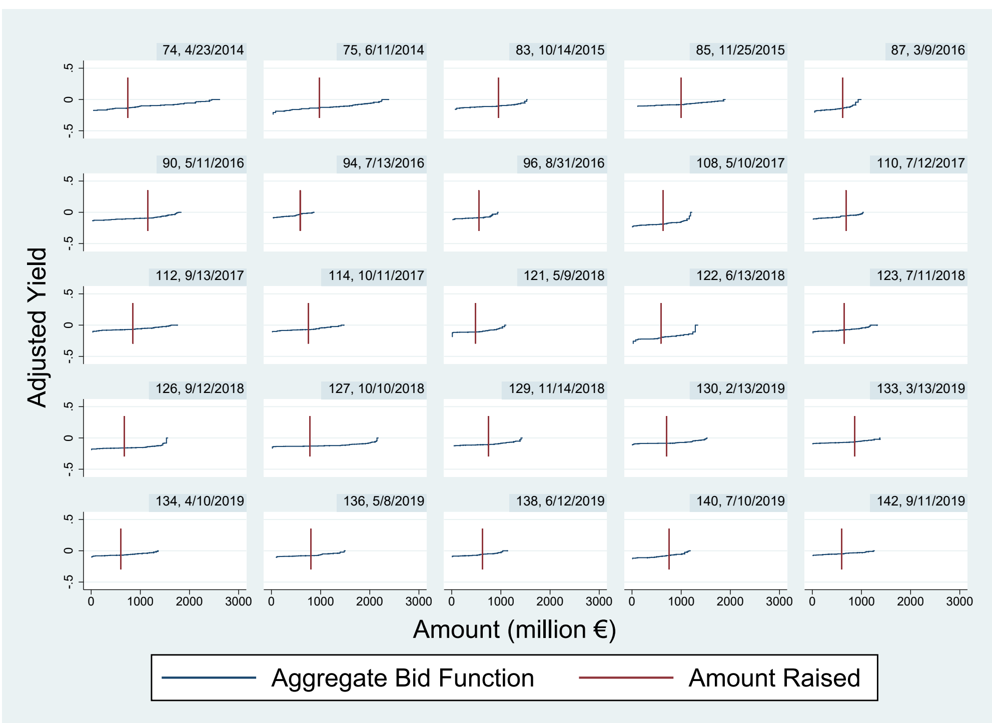

(b) Recovery

Figure 17: Demand schedule for 10 year treasury bonds in the primary market

Appendix B - Elasticities Comparison

The figure below illustrates 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 optimal behavior – the auction price is such that the government avoids the cliff after the quasi-kink.

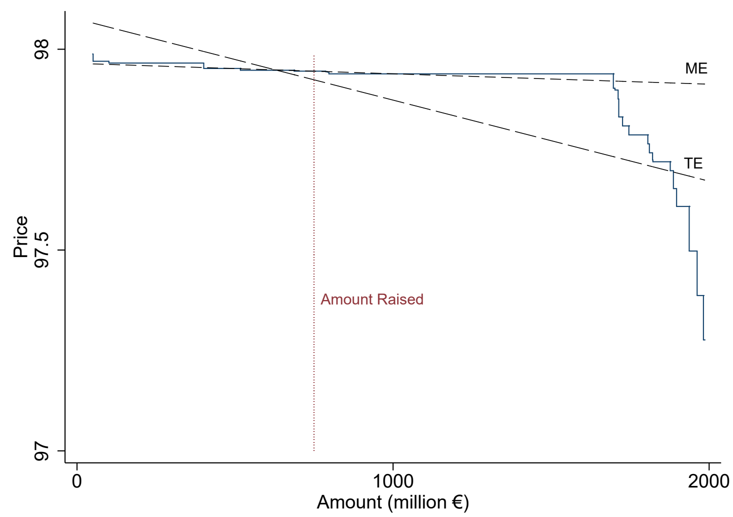

Figure 18: Comparison of ME and TE in a given auction during the crisis period

Appendix C - Resampling Procedure

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

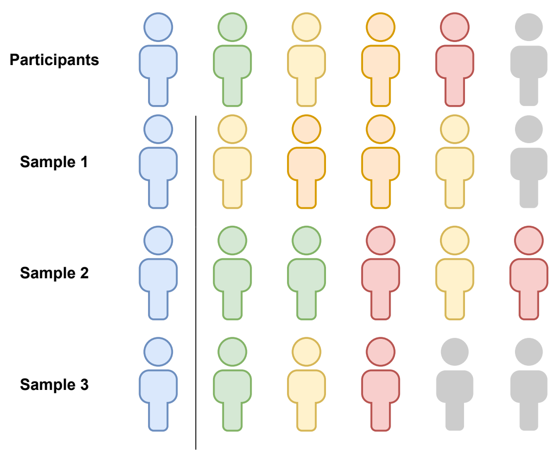

Figure 19: Illustration of the resampling procedure and different aggregate states

References

Mark Aguiar and Manuel Amador. *Sovereign Debt*, pages 647–87. North-Holland, 2014.

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.

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.

Joāo Ayres, Gaston Navarro, Juan Pablo Nicolini, and Pedro Teles. Sovereign default: The role of expectations. *Journal of Economic Theory*, 175(C):803–812, 2018.

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.

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.

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. *Review of Economic Studies*, 78:978–1014, 2011.

Jakub Kastl. Auctions in financial markets. *International Journal of Industrial Organization*, 70, 2020.

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