Asymmetric Information and Sovereign Debt: Theory Meets Mexican Data*

Harold Cole† Daniel Neuhann‡ Guillermo Ordoñez§

August 2020

Preliminary and Incomplete

Abstract

We combine novel data and theory to show that asymmetric information among investors is an important friction in primary sovereign debt markets. We exploit a unique dataset of Mexican auctions for Cetes bonds. Auctions are pay-your-bid, and our data includes all bids made by all individual bidders from 2001 to 2017. We document that the largest bidders tend to bid at higher prices (that is, their bids are more likely to be accepted), but on average they do not pay more for the bonds they buy (that is, their accepted bids are executed at the average price). We construct a model in which investors can differ in wealth, risk-aversion, market power and/or information. Only heterogeneous information can qualitatively account for our findings. We calibrate the model and find that heterogenous information about rare disasters can quantitatively match key aspects of Cetes yield dynamics and bidding behavior.

*We thank Isaac Vivas Escobedo and Gilberto Montano Calvillo from the Bank of Mexico for help with the data and Collum Freedman, Dohan Kim and Juan Sagredo for outstanding research assistance. The usual waiver of liability applies. Cole and Ordonez received support from the NSF through grant 1851707.

†University of Pennsylvania (e-mail: colehl@sas.upenn.edu)

‡University of Texas at Austin (e-mail: daniel.neuhann@mccombs.utexas.edu)

§University of Pennsylvania (e-mail: ordonez@econ.upenn.edu)

1 Introduction

Governments finance their fiscal needs by selling debt claims in sovereign bond auctions (the primary market for government bonds). Primary market prices experience periods of high volatility and high average yields, often termed crises, and more tranquil periods. This is particularly true for emerging economies, such as Mexico, where sovereign debt pays a substantial premium on the order of the equity risk premium even when accounting for default losses (**?**)). In this paper we introduce a unique dataset on primary markets for Cetes bonds in Mexico for the period June 2001 - September 2017 to study the determination of bond prices at auction. Cetes are domestically-denominated zero-coupon bonds which are sold at small face values and in large lots to a wide variety of investors. They are the most important public debt instrument in Mexico, representing 25% of all government securities in 2001.In the sample period, they are auctioned weekly using a "pay-as-you-bid" protocol.

Figure 1 shows the real *marginal price* in Cetes auctions (computed as the annual yield deflated by the yearly CPI inflation), for the four most common maturities of 28, 91, 182 and 364 days. Our sample a is relatively tranquil period for Mexico: prices are high on average and unconditional volatility is moderate, (at least relative to the prior two decades which each include a debt crisis. Conditional volatility is particularly low: prices change in small increments and are well predicted by the previous week's marginal price.

Figure 1: Marginal Prices

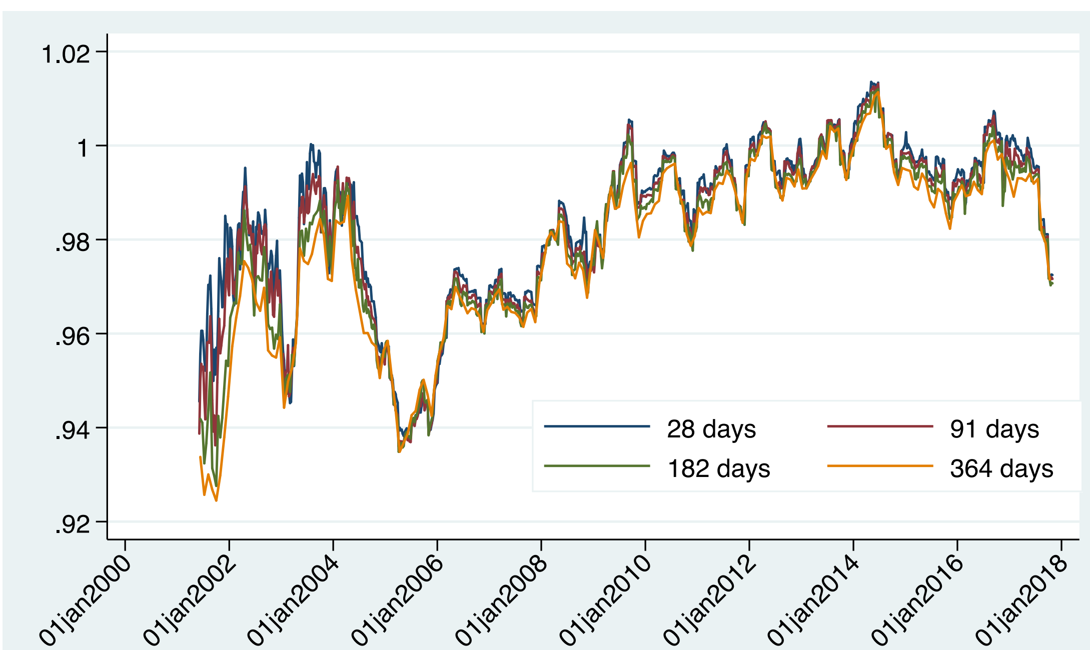

Our goal is to understand the bidding behavior and asymmetries among investors that underlie these prices. We collect bid-level data that includes information on the quantities and prices bid for each investor at each auction. This allows us to document several key facts. First, there were on average 20 bidders in each auction, with each bidder submitting an average of three bids per auction. The large number of bids and bidders suggest that the auctions were fairly competitive. Second, the share of an investor's bids that are accepted, which we label the in-the-money share (ITM share), differs substantially across bidders. The largest bidder at an auction has on average 86% of bids accepted. The remaining bidders have only 33% of their bids accepted on average. Third, the largest bidders do not overpay on average relative to the remaining bidders, where we measure overpayment as the ratio of average price paid to the marginal price.

The combination of the last two facts is both surprising and informative about investor hetereogeneity. It is surprising because the largest bidders can only have a higher ITM share if they are bidding at higher prices. However, since marginal prices are determined by other investors' bids, this bidding aggression would seem to lead to overpaying on average. What can account for this apparent inconsistency? And, what does it suggest about the nature of the shocks driving primary market prices? Answering these questions is relevant to understanding the fundamental determinants of sovereign bond prices in primary markets, critical not only for government funding but also for economic performance (**?**)).

We build a model that can accommodate a wide array of heterogeneity across investors, and asks which type of heterogeneity is most promising in explaining these patterns. The model we propose is a discriminatory-price (DP) auction that captures the protocol used by Mexico to sell bonds in primary markets during this period. We assume investors have CRRA preferences and differ in their wealth, their market power, their risk aversion and their information about the quality of the bond (as measured by its default probability). This model features a particularly tractable framework to study information heterogeneity by considering the *Walrasian limit* where bidders act as price-takers, yet being amenable to discuss the role of heterogenous market power when information is symmetric.

In the model, as in the data, investors submit multiple bids consisting of a price and a commitment to buy a given number of bonds at the bid price. The government runs down the list of bids in descending order of price until it obtains the desired revenue. This protocol thus implies a lowest-accepted, or marginal, price, with all bids at prices above the marginal price also being accepted. We allow for a *demand shock* about which there is common uncertainty and a *quality shock* (default risk) about which there can be heterogenous uncertainty. To decide on an optimal bidding strategy, an investor must forecast the distribution of marginal prices, which requires forecasts of both the demand and the quality shocks. This is an easier task for informed investors, who we assume know the probability of default and hence know which price schedule will ensue. Uninformed investors, on the other hand, do not know the quality shock and thus do not know which of the possible marginal price schedules will be operative. As a result, they face more extreme *price risk*: bids made at high prices will be accepted even if the probability of default is high and the marginal price is low. This risk can be avoided by not bidding at high prices, but this implies not buying bonds with low default probability, not earning a risk premium on infra-marginal bond purchases.

We find that it is not possible to account for the differences in bidding behavior that we document between the largest bidders and the rest relying only on the first three sources of heterogeneity we consider. If largest bidders were wealthier, they would simply scale their bids with wealth, buying more bonds, but not featuring higher ITM relative to the rest. If largest bidders had more market power, they would have an incentive to shade bids at higher prices, and should display lower ITM than the rest, not higher. Finally, differences in risk aversion could in principle accommodate large differences in the ITM share if largest bidders were bidding more aggressively at high prices, but this would mean they are more willing to pay more than the rest for bonds, which is inconsistent with largest bidders not overpaying relative to the rest.

Asymmetric information about the quality of bonds is the only source of heterogeneity with the potential to generate the qualitative patterns we have documented. If largest bidders were better informed about the bond quality, they would have a natural incentive to bid more in response to positive information and less to negative. This makes less informed investors more timid about bidding at high prices. This timidity comes from adverse selection with respect to the bond quality and creates a *winner's curse* effect: the concern of uninformed investors to be only successful in winning their bid (and paying a high price) when the bond is low quality. As a response, less informed investors tend to bid mostly at lower prices, knowing they will buy less frequently. Asymmetric information is then consistent with informed bidders displaying a larger ITM, as they buy both good and bad bonds. It is also consistent with informed bidders not overpaying in average: when bonds are good quality information induces higher marginal prices, and hence buying at higher prices (what increases the ITM) does not imply overpaying. At the same time, when bonds are good quality informed investors do not bid at high prices, then not overpaying either. Since less informed investors largely bid at low prices, for low quality bonds, they do so at prices similar to the informed investors.

Not does only this logic suggests that asymmetric information is an important friction but also that there is a sizable mass of both informed and uninformed bidders. This is because informed investors must be prevalent enough to largely clear the market when the bond quality is high, to raise the marginal price enough to avoid overpayment. At the same time, they cannot be so prevalent that there are small differences between the ITM shares of the largest bidders and the rest. To accommodate our finding it is needed that the share of informed bidders is intermediate so the adverse selection effect for the less informed bidders is strong enough to have a large impact on their bidding.

This evidence that asymmetric information plays an important role in sovereign debt primary markets is in principle surprising, as it is commonly assumed that information about a country's finances is publicly and widely available. To examine the nature of this asymmetry in greater detail, we calibrate a version of the model to Cetes data and explore the set of minimum parameters about the information structure that allows us to match quantitatively both the dynamic behavior of marginal prices and the heterogenous behavior of investors that we document.

As we discussed, a critical dynamic property of Cetes marginal prices is the relatively high unconditional volatility and low conditional volatility. Based on this piece of evidence, we make a distinction in the model between publicly observed information, which determines a public state, and privately observed information, which leads to the information heterogeneity we uncover. For the public information we have in mind standard fundamentals, such as GDP growth or inflation, or the past week's auction prices. For the privately observed information we have in mind information that is difficult or costly to acquire, process and evaluate. Within the Mexican context, a particularly pertinent example is knowledge of the inner workings of the government, such as the financial negotiations that took place between Clinton and Congress over the 1995 bailout.[1]

Our calibration shows that a surprisingly small amount of information heterogeneity of this sort is sufficient to generate a large difference in bidding behavior consistent with a very low conditional uncertainty, especially when the government's debt is close to risk-free. However, when the average default risk is modestly higher, a stronger adverse selection effect is needed to induce the same level of timidity by uninformed investors, but without generating counterfactual levels of unconditional volatility of prices. The

[1] On January 30, 1995, at exactly the moment when the Mexican government was informing the Clinton Administration that without an emergency injection of funds it would have to default, the Speaker of the House, Newt Gingrich, was informing the Clinton Administration that the bailout bill was stalled in the Congress. See Chun, John H. "Post-Modern Sovereign Debt Crisis: Did Mexico Need an International Bankruptcy Forum." Fordham L. Rev. 64 (1995): 2647. The relevance of political uncertainty for sovereign default in emerging markets has been also highlighted by **?**).

version of our model that best achieves this result includes a small probability of an extremely bad (*black swan*) event of the sort seen in Mexico during the 1990s and 1980s.[2] When we assume that this low-probability bad event has not occurred in the sample data (as it did not occur during the timeframe we consider), as in the ”peso-problem” literature on asset pricing anomalies, we generate results that are quite close to the data within a very sparse model.[3]

This combination of a fairly large amount of unconditional uncertainty coming through changes in public information, and the low amount of conditional uncertainty coming through the combination of heterogeneous information about a rare disaster, along with a common modest degree of demand uncertainty, provides a plausible set of sources for both Cetes price movements and bidding behavior. This is relevant because we show in the model that intermediate levels of asymmetric information increase the yield on bonds vis-a-vis situations in which information is more symmetric across investors.

Related Literature: Our paper fills an important gap in the sovereign debt literature, which has typically focused on bond yields in secondary market, but has neglected the specifics of how a government sells its bonds and the role of investors heterogeneity in general but asymmetric information in particular, in determining bonds prices in primary markets.[4] To focus squarely on the determination of primary market prices, we neglect some of the issues studied in the literature, but expand on others. First, most papers study sovereign default as the outcome of governments' strategic choice, but use a parsimonious model of investor optimization (see, for instance, **?**), **?**), **?**)). We take the opposite route, and focus on the auction mechanics and investors choices while entirely neglecting strategic considerations on the part of the government.[5]

Second, most of the literature generates a fixed mapping between the bond quality and its price by assuming that investors are risk neutral and then requiring that the return, adjusted for the probability of default, equals the risk-free rate. While there has been some attention to the impact of the timing of decisions and of debt maturity in sovereign markets (see **?**)), the actual mechanics of how sovereign bonds are sold in reality through auctions and their impact on observed prices has been largely ignored. Our paper not

[2]Disaster risk has been argued to play a large role in both asset pricing and macroeconomic fluctuations. See for example **?**), **?**) and **?**).

[3]A peso-problem refers to a circumstance in which some infrequent or unprecedented event, such as an economic disaster, may have a substantial impact on asset prices. However, the infrequency of the event makes it hard to estimate its empirical probability and it even may not have occurred in the time series being considered despite affecting the assessments of the investors. See for example **?**), **?**), **?**), and **?**).

[4]Exceptions are **?**) and **?**), who study how the liquidity of bonds in secondary markets affect their price in primary markets.

[5]See for example **?**), the review articles by **?**) and **?**), and the recent quantitative literature by **?**), **?**), **?**), **?**).

only focuses on the neglected roles of auction mechanics and information heterogeneity, but argues that the interaction of these factors along with investor risk aversion are driving pricing and bidding in primary markets. The nature of the information shocks, both public and private, heterogeneous and common, is consistent with the referred literature on rare disasters and the "peso-problem".

The paper also contributes theoretically on circumventing some of the challenges that standard auction models face in accommodating asymmetric information.[6] We propose a novel auction model with three key characteristics: (i) the good being auctioned is perfectly divisible, (ii) the number of bidders is large, and (iii) there is both uncertainty about the good quality and about the mass of investors who participate in the auction. Given these three characteristics, the price-quantity strategic aspects of standard auction theory become less relevant, and a price-taking, or Walrasian, analysis emerges as a good approximation.[7] Finally, our paper is related to a recent effort to empirically document the implications of information sharing across dealers on the revenue of governments (see **?**)).

In the next section we describe our novel dataset on Cetes auctions and describe our main findings on bidding behavior and outcomes. In Section 3 we provide a tractable model of discriminatory-price auction with several sources of bidder heterogeneity. In Section 4 we show that only information asymmetry can qualitatively accommodate the patterns of bidding behavior that we uncover. Once we have identified asymmetric information as the main friction in primary markets, we explore its implications for bond prices in Section 5. Finally, in Section 6 we calibrate the model to match relevant moments on price dynamics and bidding behavior in order to understand the quantitative extent and nature of asymmetric information for the Mexican case. Section 7 concludes.

[6]For a discussion see, for instance, **?**) who characterize an optimal mechanism in the context of initial public offering auctions under pure common values in the presence of better informed dealers (investment banks) and retail investors. Another example is **?**), who study a divisible good uniform-price auction with risk neutral bidders with asymmetric dispersed information in linear strategies.

[7]Recent auction literature shows that price-taking arises as the number of bidders get large. A recent example is **?**), who show that the equilibria of large double auctions with correlated private values are essentially fully revealing and approximate price-taking behavior when the number of bidders goes to infinity. Another is **?**) who show a similar result when bidders have affiliated values and prices are on a fine grid.

2 Institutional Background and Data

2.1 Cetes Bonds

We study auctions of Mexican Federal Treasury Bills (Cetes), which are zero-coupon pure discount bonds with typical maturities of 28, 91, 182 and 364 days. They are the leading instrument in Mexican money markets and main source of federal government funding since 1978. Since its inception, the primary market for Cetes has consisted of public auctions, with the auction protocols alternating between uniform-price and discriminatory-price since 1978. Our data comes is from the archives of the Mexican central bank and we are the first to compile it. We focus on the period June 2001 to September 2017 during which Cetes of all maturities were regularly sold in weekly pay-your-bid auctions.[8] We have data on all bids submitted (not just those that were executed) and all bidders. Moreover, we have a numeric identifier that allows us to construct the set of submitted bids for each bidder.

2.2 Basic Facts about Prices and Bids

We observe a total of 2,717 Cetes auctions. On average there are 20 bidders at each auction, with each bidder submitting an average of 3 bids per auction. Table 1 shows summary statistics for each maturity.

Table 1: Summary Data on Cetes Auctions. 2001-2017

Maturity (days)	Auctions	Bidders per auction	Bids per auction
28	857	19.4	59.6
91	857	19.2	64.8
182	789	17.2	60.0
364	214	17.3	66.7

As a first step, we establish a number of basic facts about prices and bids that allow us to discriminate between sources of investor heterogeneity in our model. During the sample period, Mexico experienced relatively stable inflation and macroeconomic conditions. This is reflected in relatively low average yields and mild conditional volatility of marginal auction prices. Table 1 shows the time-series moments of marginal prices (MP)

[8] In October 5, 2017 Mexico switched protocol to uniform-price auctions.

by maturity. We report the autocorrelation between marginal prices at subsequent auctions as a measure of conditional uncertainty. 364-day bonds are auctioned monthly, all other maturities are auctioned weekly.

Table 2: Time series properties of marginal prices

Maturity	Avg. MP	St. Dev. MP	Autocorrelation MP
28	0.984	0.017	0.984
91	0.983	0.018	0.983
182	0.982	0.018	0.992
364	0.978	0.019	0.956

The high degree of autocorrelation suggests that the unconditional uncertainty of the marginal price (its standard deviation) is much higher than the uncertainty conditional on the prior week's auction results. To further characterize this uncertainty, we run a regression of marginal prices 28-day bonds on a constant and one lag,

$$p_t^{28d} = \beta_0 + \beta_1 p_{t-1}^{28d} + \epsilon_t. \tag{1}$$

We estimate $\beta_1 = 0.98$, which implies that lagged prices are very informative.[9] This is not surprising because auctions take place every week. We also estimate $R^2 = 0.97$, which implies that the conditional uncertainty is indeed quite low during this period. A naive interpretation of this result might suggest that publicly observable prices from previous weeks encode all relevant information for pricing bonds in the current auction. However, we will show that even a small amount of conditional uncertainty can have significant effects on bidding strategies and prices.

Next we establish basic facts about bidding patterns. One drawback of our data is that the numeric bidder identifier is auction-specific, so that we cannot track bidders across auctions. To uncover heterogenous bidding behavior, we compare the bidding behavior of the largest bidder at an auction to all other bidders. The largest bidder is the bidder who demands the most bonds, as measured by bid price multplied by bid amounts. We make this distinction anticipating certain findings from our model. In particular, the largest bidder would naturally stand out as (i) the wealthiest, (ii) the one with most market power, (iii) the least risk-averse, and (iv) the one with the strongest incentive to be informed, where each of these factors are potential sources heterogeneity in our model.

[9] As can be guessed from Figure 1 and Table 5, the result for other maturities is very similar.

Figure 2 shows a histogram of the fraction of bids that are accepted for the largest bidder at an auction and all other bidders, aggregated across all auctions and maturities. We call the fraction of accepted bids the *in-the-money share* (the ITM). For the largest bidder, the mode of the ITM is 1 (typically, all of their bids are accepted), but there is much more dispersion for smaller bidders. On average, the largest investors have an average 84% of their bids executed, while only 33% of the rest of the bidders' bids are executed.

Figure 2: In the Money Shares Largest vs. Rest

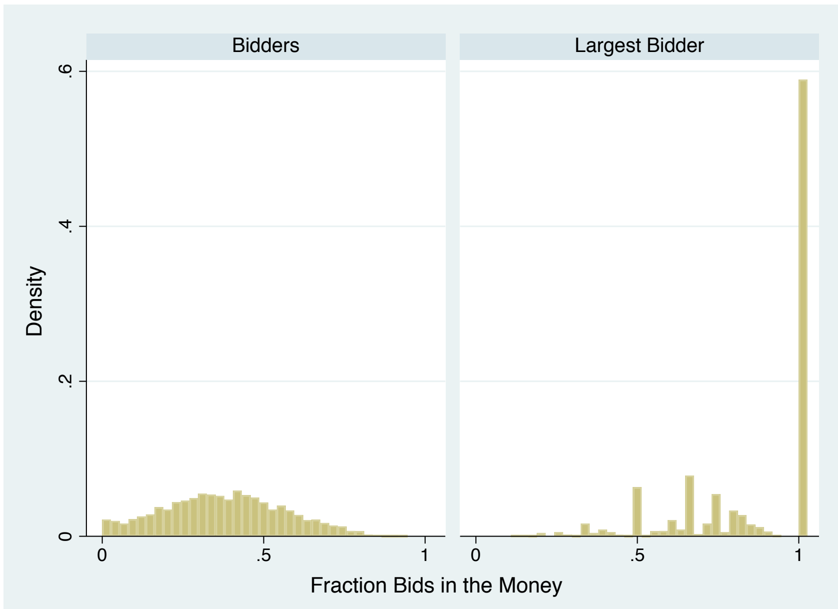

What accounts for differences in in-the-money shares across bidders? One possibility is that large bidders systematically bid higher prices. To investigate this issue, we construct a measure of *overpayment*, defined as the ratio of the average price paid (weighted by bids executed at each price) to the marginal price. Since all bids above the marginal price are accepted, a ratio greater than one indicates that the bidder overpaid for at least some bids. In Figure 3 we show a histogram (for all auctions and all maturities) of the overpayment for the largest bidder and the rest of bidders. We find that the distribution of overpayment is very similar for large bidders and other bidders. Hence it is unlikely that excessively high bids by the largest bidders account for the observed differences in in-the-money shares.

The combination of aggressive bidding by the largest bidder (as measured by the differences in-the-money shares) without concomitant differences in overpaying is puzzling: in a pay-your bid auction, having more bids accepted is typically associated with higher bids. In the next section we construct a model that can replicate these facts by relying on asymmetric information about default risk.

Figure 3: Ratio between Weighted Price Paid and the Marginal Price

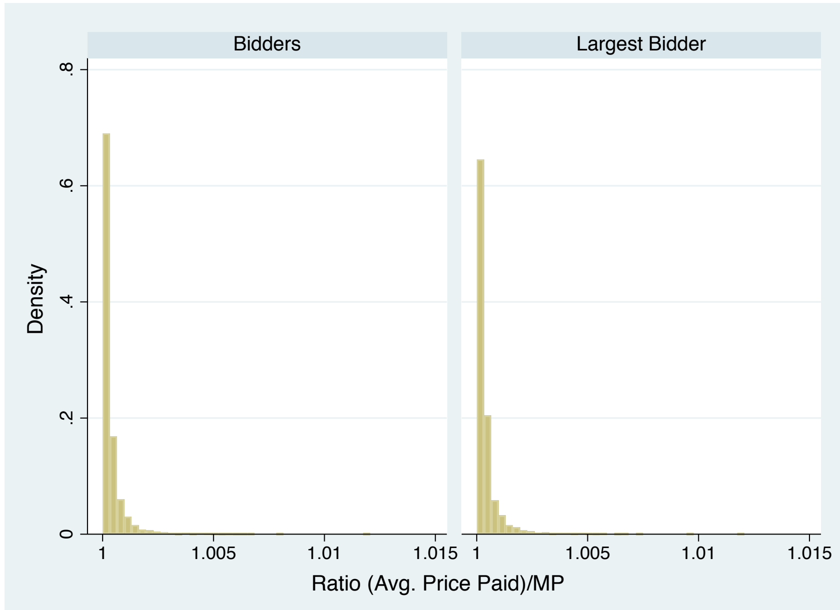

3 A Model of the Primary Debt Market

We now construct a model of sovereign debt auctions that allows for rich heterogeneity across bidders, including differences in wealth, risk aversion, market power, and information. We find that only information asymmetry can rationalize the bidding patterns discussed in the previous section.

3.1 Environment

There is a single period with two dates ($t = 1, 2$), and a single good (the numeraire). The economy is populated by a government and a measure one of risk-averse investors. Investors consume only in period two. Their objective is to maximize their expected utility over second-period consumption given a strictly concave flow utility function that satisfies the Inada conditions. Each investor has wealth W_j in period one and can either invest it in a risk-free bond (storage) or the risky bond being auctioned by the government. There is no borrowing. The government is modeled mechanically: it needs to raise D units of the numeraire good in period one by auctioning multiple units of a bond that promises repayment in period two. Without loss of generality (and like *Cetes*), bonds are zero-coupon and pure-discount: each bond promises a claim to one unit of the numeraire good in the period two. Bonds are risky in that the government may default on its promises. If the government defaults, investors cannot recover any of their investment.

The government's default probability κ_θ is random and determined by the realization of an exogenous state of the world $\theta \in \{g, b\}$. We assume that $\kappa_g < \kappa_b$ and that the ex-ante probability of each state is given by $f(g)$ and $f(b)$ respectively, with $f(g) + f(b) = 1$. Since the default probability determines the expected repayment of the bond, we refer to the realization of κ as a *quality shock*. The bond with default probability κ_g is a good quality bond and the one with default probability κ_b is a bad quality bond. In our simple one-period model, we capture different bond maturities by the length of the period. If we view defaults as random events that occur with some (constant) arrival rate, then longer maturities are associated with higher default probabilities κ_θ.

To allow for the fact that not all bids are always accepted, we introduce *demand shocks*. Specifically, we assume that a random share of investors η cannot participate in the auction and instead invests in the risk-free bond only. The demand shock η is discrete and lives on an arbitrarily fine discrete grid $\mathcal{H} \equiv \{\eta_0, ..., \eta_M\}$ with length M. We index the demand shock by $k \in \{0, M\}$. Without loss of generality, let η_k be strictly increasing in k and denote the probability of η_k by $h(\eta_k)$.

We refer to $s = (\theta, \eta)$ as the state of the world and to the set of states by $\mathcal{S} = \{g, b\} \times \mathcal{H}$. The cumulative distribution function over states is denoted by Γ. Consistent with our mechanical modeling of the government we assume that it observes neither θ nor η before the auction. This precludes signaling by the government.

Remark 1. *We later consider a repeated version of the static model in which the possible realizations of θ are governed by an aggregate shock process that is publicly observed at the end of the previous period. While this process allows us to account for dynamic evolution of bond prices, it does not impact the analysis of the single-period auction. Hence we abstract from it for now.*

Remark 2. *A natural interpretation of η is that it represents the fraction of investors who suffer a liquidity or hedging shock. Another interpretation is that some investors randomly have access to more favorable investment opportunities and this do not invest in bonds. In the context of the auction literature, the demand shock η can be therefore be thought of as a correlated private value shock, while the quality shock θ is a common value shock. The demand shock can also be thought as a supply shock to the government financing needs $D\psi$, where $\psi = 1/(1 - \eta)$.*

3.2 Investor Heterogeneity

Investors can differ in terms of their fundamental type and their information type. The *fundamental type* $j \in \{1, 2\}$ indexes both the investor's utility function U_j over second period consumption and their initial wealth W_j, capturing differences in both risk-aversion and wealth. Further, higher wealth can also capture the extent of market power of the bidder. To hold down notational clutter we will only consider two fundamental types at a time; thus we will assume that either they have the same preferences over consumption but differ in terms of their wealth, or that they have the same wealth level but differ in their preferences over consumption. We will also assume for simplicity that there are an equal mass of the two fundamental types. This does not affect our results.

The *information type* refers to whether investors are informed (I) about θ and thus know its realization, or uninformed (U) about θ and do not know its realization. We denote by $i \in \{I, U\}$ the type of investor and by $n \in [0,1]$ the share of informed investors. The remaining share $1-n$ consists of uninformed investors. The fraction n determines the degree of asymmetric information in the sense that it measures the relative mass of investors with superior information about the quality of the bond. No investor is informed about the demand shock η, which means that all investors face some uncertainty.

To reduce notational clutter, we assume that there is no correlation between fundamental and information types and that both fundamental and information types are symmetrically impacted by η.

3.3 Auction Protocol and Strategies

We now describe the discriminatory-price auction protocol used to sell bonds. A *bid* is a pair $\{\tilde{P}, \tilde{B}\}$ representing a commitment to purchase $\tilde{B}$ units of the bond at a price $\tilde{P}$, should the government decide to accept the bid. Each investor is free to submit as many bids as desired at the beginning of the auction. There is no short-selling, $\tilde{B} \geq 0$. The government treats each bid independently, sorts all received bids from the highest to the lowest bid price, and accepts all bids in descending price-order until it raises D in revenue. We refer to the lowest accepted price in state s as the state-contingent *marginal price* $P(s)$. All bids at prices above the marginal price are accepted, all bids below are rejected. We refer to bids at or above the marginal price as *in the money*, and to bids below the marginal price as *out of the money*.[10]

Since bonds pay one unit of the numeraire after repayment and zero after default, the range of feasible prices is $[0,1]$. A bidding strategy for an investors maps any price in $[0,1]$ into a weakly positive bid quantity. Since all prices above the marginal price are accepted

[10] If there is excess demand at the marginal price, we assume that the government rations bids pro-rata. In the model, this does not occur in equilibrium. In our data, there is some rationing because prices are restricted to a fine grid. However the extent of rationing of the bids at the marginal price is roughly uniformly distributed between 0 and 1, suggesting that it is not playing a key role. As rounding bids does not add anything to our insights, we follow the literature that assumes the set of possible bid prices is in a continuum, as in the seminal work of **?**). For a treatment of bidders restricted to submit discrete bidpoints in a uniform-price auction of a perfectly divisible good see **?**).

and bids are executed at the bid price, it is without loss of generality to restrict attention to bidding strategies that assign zero bids to any price that is not marginal in at least one state of the world. This simplification allows us to define bidding strategies as follows.

Let $P(s)$ denote the marginal price in state $s = (\theta, \eta)$, and define $\mathcal{P}$ to be the set of marginal prices. A bidding strategy for an uninformed investor of fundamental type j is a function $B_j^U(s)$ denoting the number of bids at marginal price $P(s)$. A bidding strategy for an informed investor of fundamental type j is a function $B_j^I(s|\hat{\theta}(s))$ denoting the number of bids at marginal price $P(s)$, where $\hat{\theta}(s)$ is the realized quality shock associated with state s that is observed by the informed investor. Observe that bids in any two states with a common price are perfect substitutes because they are accepted and rejected in the identical set of states. The precise allocation of bids across such states is thus irrelevant. A useful

Since there are two fundamental types with mass $\frac{1}{2}$ each, the primary market clearing condition which ensures that the government raises revenues equal to D in state s is

$$D = (1-\eta)\left[\frac{1}{2}\sum_{j\in\{1,2\}}\sum_{\tilde{s}:P(\tilde{s})\geq P(s)}\left[nB_j^I(\tilde{s}|\hat{\theta}(\tilde{s})) + (1-n)B_j^U(\tilde{s})\right]P(\tilde{s})\right]. \tag{2}$$

The left-hand side is the debt level, the right-hand side is the sum over bids submitted at prices above the state-contingent marginal price $P(s)$, evaluated at the bid price $P(\tilde{s})$.

3.4 In-the-money shares and average prices paid

To map bids into the data moments described above, we define the in-the-money share ITM of informed investors in state $[\theta, \eta]$ as

$$ITM_j^I(\theta,\eta) = \frac{\sum_{P(\theta,\hat{\eta})\geq P(\theta,\eta)} B_j^I(\theta,\hat{\eta})h(\hat{\eta})}{\sum B_j^I(\theta,\hat{\eta})h(\hat{\eta})} \tag{3}$$

The big-weighted average price paid by informed investors in state (θ, η) is

$$AP_j^I(\theta,\eta) = \sum_{P(\theta,\hat{\eta})\geq P(\theta,\eta)} P(\theta,\hat{\eta})\frac{B_j^I(\theta,\hat{\eta})h(\hat{\eta})}{\sum B_j^I(\theta,\hat{\eta})h(\hat{\eta})}. \tag{4}$$

Overpayment is the ratio between the weighted average price and the marginal price,

$$Overpayment_j^I = \frac{AP_j^I(\theta,\eta)}{P(\theta,\eta)} \tag{5}$$

The analogous definitions for uninformed investors are

$$ITM_j^U(\theta,\eta) = \frac{\sum_{P(\hat{\theta},\hat{\eta})\geq P(\theta,\eta)} B_j^U(\hat{\theta},\hat{\eta})f(\hat{\theta})h(\hat{\eta})}{\sum B_j^U(\hat{\theta},\hat{\eta})f(\hat{\theta})h(\hat{\eta})}, \tag{6}$$

$$AP_j^U(\theta,\eta) = \sum_{P(\hat{\theta},\hat{\eta})\geq P(\theta,\eta)} P(\hat{\theta},\hat{\eta})\frac{B_j^U(\hat{\theta},\hat{\eta})f(\hat{\theta})h(\hat{\eta})}{\sum B_j^U(\hat{\theta},\hat{\eta})f(\hat{\theta})h(\hat{\eta})}, \tag{7}$$

and

$$Overpayment_j^U = \frac{AP_j^U(\theta,\eta)}{P(\theta,\eta)} \tag{8}$$

3.5 Uninformed Investor Decision Problem

We start defining the problem of the uninformed investor. If the government ends up defaulting in the second period, the uninformed investor of fundamental type j simply consumes the unit payoff from his risk-free bonds, which we denote by $\mathcal{B}_{RF,j}^U(s)$. If the government does not default, then the investor additionally consumes the unit payoff from his total holdings of the risky bond, which we denote by $\mathcal{B}_{R,j}^U(s)$. Hence the expected payoff of an uninformed investor of fundamental type j is the probability-weighted sum over conditional payoffs in each state (θ,η), i.e.

$$V_j^U = \sum_{\theta\in\{g,b\}}\sum_{\eta\in\mathcal{H}} \left\{U(\mathcal{B}_{RF,j}^U([\theta,\eta]))\kappa_\theta + U\left(\mathcal{B}_{RF,j}^U([\theta,\eta]) + \mathcal{B}_{R,j}^U([\theta,\eta])\right)(1-\kappa_\theta)\right\} f(\theta)h(\eta). \tag{9}$$

The total number of risky bonds purchased by an uninformed bidder in each state, $\mathcal{B}_{R,j}^U(s)$, is the sum of in-the-money bids,

$$\mathcal{B}_{R,j}^U(s) = \sum_{s':P(s')\geq P(s)} B_j^U(s'). \tag{10}$$

The total expenditure on risky bonds determines the investor's holding of the risk-free bond $\mathcal{B}_{RF,j}^U(s)$. Since the price at which a bid is executed depends on the state of the world s, so does the expenditure on risky bonds. Given that the price of risk-free bonds is normalized to one, we have

$$\mathcal{B}_{RF,j}^U(s) = W_j - \left[\sum_{s':P(s')\geq P(s)} B_j^U(s')P(s')\right]. \tag{11}$$

The short-sale and borrowing constraints for an uninformed investor are:

$$B_j^U(s) \geq 0 \quad \text{and} \quad \mathcal{B}_{RF,j}^U(s) \geq 0 \qquad \forall s \in \mathcal{S}. \tag{12}$$

3.6 Informed Investor Decision Problem

Informed investors make bids conditional on the realized quality shock $\theta \in \{g, b\}$ since they can observe it by virtue of being informed. Hence we can restrict attention to bid functions of the form $B_j^I(s, \theta)$. Define $\mathcal{B}_{R,j}^I(s, \theta)$ and $\mathcal{B}_{RF,j}^I(s, \theta)$ to be the total purchases of the risky bond and the risk-free bond in state $s = [\theta, \eta]$. respectively, given that $\theta(s) = \theta$ denotes the quality shock associated with state s. The expected payoff to an informed investor of fundamental type j is

$$V_j^I = \sum_{\eta \in \mathcal{H}} \left\{ U(\mathcal{B}_{RF,j}^I([\theta, \eta], \theta))\kappa_\theta + U\left(\mathcal{B}_{RF,j}^I([\theta, \eta], \theta) + \mathcal{B}_{R,j}^I([\theta, \eta], \theta)\right)(1 - \kappa_\theta) \right\} h(\eta), \quad \forall \theta \in \{g, b\}. \tag{13}$$

The total purchases of risky bonds are

$$\mathcal{B}_{R,j}^I(s, \theta) = \sum_{s' : P(s') \geq P(s)} B_j^I(s', \theta), \qquad \forall \theta \in \{g, b\}, \tag{14}$$

and the holdings of the risk-free bond are

$$\mathcal{B}_{RF,j}^I(s, \theta) = W_j - \left[\sum_{s' : P(s') \geq P(s)} B_j^I(s', \theta) P(s') \right], \qquad \forall \theta \in \{g, b\}, \tag{15}$$

The short-sale and borrowing constraints are

$$B_j^I(s, \theta) \geq 0 \quad \text{and} \quad \mathcal{B}_{RF,j}^I(s, \theta) \geq 0 \qquad \forall s \in \mathcal{S} \ \text{ and } \ \theta \in \{g, b\}. \tag{16}$$

3.7 Equilibrium Definition

The definition of an equilibrium is as follows.

Definition 1 (Auction Equilibrium)**.** *An auction equilibrium is a price schedule $P : \mathcal{S} \to [0, 1]$, and bidding fuctions $B_j^U : \mathcal{S} \to [0, \infty)$ and $B_j^I : \mathcal{S} \times \{g, b\} \to [0, \infty)$ for $j \in \{1, 2\}$, such that*

1. *Uninformed investors choose $B_j^U(s)$ $\forall s \in \mathcal{S}$ to maximize (9) subject to (10), (11) and (12).*
2. *Informed investors choose $B_j^I(s, \theta)$ $\forall s \in \mathcal{S}$ and each realized $\theta \in \{g, b\}$ to maximize (13) subject to (14), (15) and (16).*

3. *The auction-clearing condition* (2) *is satisfied for all* $s \in \mathcal{S}$,

4 Role of Investor Heterogeneity on Auction Equilibrium

We can now characterize some properties of the equilibrium.

Proposition 1. *The price function* $P(\theta, \eta)$ *is strictly decreasing in* η *conditional on* θ.

Proof. Monotonicity in η follows directly from the auction-clearing condition for a realization of θ. $\square$

Corollary 1. *A bid made at price* $P(\theta, \eta)$ *is in-the-money for all* $\hat{\eta} \geq \eta$ *given* θ. *If there exists a* $\bar{\eta}$ *such that* $P(\bar{\theta}, \bar{\eta}) = P(\theta, \eta)$ *for* $\bar{\theta} \neq \theta$, *then it is also in-the-money for all* $\hat{\eta} \geq \bar{\eta}$ *given* $\bar{\theta}$.

4.1 Investors' Optimal Bidding - Information Types

The informed investor's problem is relatively simply because they know the realized quality shock, which we denote here as $\theta = \theta^*$. The first-order condition for bid $B_j^I([\theta^*, \eta^*])$ for each state $s = [\theta^*, \eta^*]$ is

$$\sum_{\eta} \left\{ \begin{array}{c} -U_j'(\mathcal{B}_{RF,j}^I([\theta^*, \eta]))\kappa_{\theta^*} P([\theta^*, \eta^*]) \\ +U_j' \left(\begin{array}{c} \mathcal{B}_{RF,j}^I([\theta^*, \eta]) \\ +\mathcal{B}_{R,j}^I([\theta^*, \eta]) \end{array} \right) (1 - \kappa_{\theta^*})(1 - P([\theta^*, \eta^*])) \end{array} \right\} \mathcal{I}\left\{ P\left([\theta^*, \eta^*]\right) \geq P([\theta^*, \eta]) \right\} h(\eta)$$
$$- \chi_j^I([\theta^*, \eta^*]) = 0, \qquad \forall \theta^* \in \{g, b\}. \tag{17}$$

where $\chi_j^I(s)$ is the multiplier on the nonnegativity constraint for the short-sale constraint, and $\mathcal{I}\{\cdot\}$ is an indicator function.[11] The marginal price $P([\theta^*, \eta])$ determines the set of states the bid is in the money, but not the price at which each accepted bid is executed, as all in-the-money bids are executed at the bid price in the discriminatory-price protocol.

Notice that the sum over the indicator function indicates the "in-the-money" share of bids for informed investors. A bid made for state $[\theta^*, \eta^*]$ is not in the money when η is such that $P[\theta^*, \eta^*] < P[\theta^*, \eta]$, which according to Proposition 1 corresponds to states for which $\eta < \eta^*$. This implies that all bids are in the money only when $\eta = \eta_M$. For all $\eta < \eta_M$ a fraction of bids will not be in the money. For uninformed investors of fundamental type j, the first-order conditions for $B_j^U([\theta^*, \eta^*])$, given that uninformed investors do not know

[11] Notice that the constraint that risk-free bonds have to be non-negative is never binding given the assumption of Inada conditions for preferences.

the realized state $\theta = \theta^*$ is,

$$\sum_{\theta \in \{g,b\}} \sum_{\eta} \left\{ \begin{array}{c} -U_j'(\mathcal{B}^U_{RF,j}([\theta,\eta]))\kappa_{\theta^*} P([\theta^*,\eta^*]) \\ +U_j' \begin{pmatrix} \mathcal{B}^U_{RF,j}([\theta,\eta]) \\ +\mathcal{B}^U_{R,j}([\theta,\eta]) \end{pmatrix} (1-\kappa_{\theta^*})(1-P([\theta^*,\eta^*])) \end{array} \right\} \times$$
$$\mathcal{I}\left\{ P\left([\theta^*,\eta^*]\right) \geq P([\theta,\eta])\right)\right\} f(\theta)h(\eta) \quad - \chi^U_j\left([\theta^*,\eta^*]\right) = 0.$$

Comparing this expression to the informed investors' first-order condition, it is clear that the uninformed face the same basic tradeoffs as the informed, but take expectations over all possible quality shocks, not only the realized one. As we show below, this leads to a form of adverse selection: when bidding at high marginal prices, uninformed investors also expect these bids to be accepted after bad quality shocks, leading to a downward revision of the asset quality they expect to acquire. Uninformed investors thus have weaker marginal incentives to bid at high prices than informed investors.

As we discuss later, when analyzing the auction equilibrium with asymmetric information, it is possible that in-the-money shares are different across information types, and while the average prices that they pay conditional on buying is similar.

4.2 Investors' Optimal Bidding - Fundamental Types

Here we show that fundamental type differences cannot accommodate the bidding pattern heterogeneity we document in Cetes data.

Differences in Wealth: How does differences in wealth affect the bidding behavior of an investor of information type i when there is symmetric information (all investors have the same information type)? As preferences are constant-relative-risk-aversion (CRRA), it turns out that wealth impacts on bidding behavior only by scaling up bids.

Proposition 2. *With CRRA preferences, if $B^i_j([\theta^*,\eta^*])$ is a solution to the first-order condition for a bidder of type (i,j) then the solution to the problem of a (i,j') type bidder is*

$$B^i_{j'}([\theta^*,\eta^*]) = \alpha B^i_j([\theta^*,\eta^*])$$

where $\alpha = W_{j'}/W_j$.

Proof. Define $W_{j'} = \alpha W_j$. By scaling all bids $B^i_j([\theta^*,\eta^*])$ by α then, by the property of CRRA utility functions $U_{j'}'(\mathcal{B}^I_{RF,j'}([\theta^*,\eta^*])) = \alpha U_j'(\mathcal{B}^I_{RF,j}([\theta^*,\eta^*]))$ (from equation 15) and $U_{j'}'(\mathcal{B}^I_{RF,j'}([\theta^*,\eta^*]) + \mathcal{B}^I_{R,j'}([\theta^*,\eta^*])) = \alpha U_j'(\mathcal{B}^I_{RF,j}([\theta^*,\eta^*]) + \mathcal{B}^I_{R,j'}([\theta^*,\eta^*]))$ (from equation 14).

Replacing these results in the first-order condition for informed investors in (17) it is clear that scaling all bids by α is indeed optimal when wealth is scaled by α. $\square$

From applying this result to equation (3) or (6) it is clear that a wealthier bidder would bid and buy more bonds, but not a larger fraction of his total bids, which is the same across wealth levels, state by state.

Differences in Market Power: One may wonder, however, whether the largest bidder, by being potentially wealthier and buying a larger fraction of the auction. If one of the bidders were sufficiently large to impact on the price, she would internalize the fact that by bidding less at a price $P(\theta, \eta)$ she would be lowering the demand in all states for which that bid is in-the-money (for all states $\eta' > \eta$ for a given θ). In equilibrium the marginal price schedule from this point on would decline to increase the demand for bonds by other, smaller, price-taker, bidders. This additional term in the large bidder first-order condition would encourage her to bid less, especially at the very top of the price schedule where it would depress the largest number of prices in the schedule. However, the fact that she had bid less at the top would mean that the large bidder's cumulative purchases would be lower, and this in term would give a stronger incentive to bid more at lower prices through the terms in (17). These two effects (bid less than price takers at low values of η with high prices and more at high values of η with low prices) would imply that bidders large enough to have an impact on prices would display lower in-the-money share than their smaller, price-taker, counterparts. We summarize these insights in the next Corollary

Corollary 2. *With CRRA preferences, wealthier bidders who have market power as a result, do not have a larger in-the-money share.*

Differences in Risk Aversion: Can differences in risk aversion accommodate the bidding patterns we document in Cetes data? To see why this cannot be the case, denote the coefficient of risk aversion by σ and assume the fundamental type 1 is more risk averse than type 2, this is $\sigma_1 > \sigma_2$.

Naturally, more risk averse individuals will demand less risky bonds. To see this, assume otherwise, that the two types bid the same, this is $B_2^i([\theta^*, \eta^*]) = B_1^i([\theta^*, \eta^*])$ in each state $s = [\theta^*, \eta^*]$. The first order condition (17) would be strictly negative for type $j = 1$ for all bid choices at which type $j = 2$ makes a positive bid. This is the case because the marginal utility from low consumption (in case of default), which enters negatively, increases faster than the marginal utility from high consumption (in case of repayment), which enters positively. In words, higher risk aversion lowers the marginal incentive to invest in a risky asset. It is not clear, however, in which section of the price schedule this reduction in bidding is more prevalent, and then it may create in-the-money share differences consistent with the patterns in the data.

However, even if differences in risk aversion could in principle rationalize in-the-money share differences between the largest bidders and the rest, those positively correlate with overpayment. To see this, if at state (θ, η) the fundamental type j has a higher in-the-money share that type j', as $P(\theta, \hat{\eta}) > P(\theta, \eta)$ it should be the case, from equations (4) that $\frac{AP_j^I(\theta,\eta)}{P(\theta,\eta)}$ is also higher for type j than for type j'. The same holds for uninformed investors from equation (7).

Corollary 3. *With CRRA preferences, investors whose preferences induce a higher in-the-money share also overpay more.*

Intuitively, the positive correlation between in-the-money shares and overpayment (conditional on buying) comes from bidding on a single pricing schedule, which is the case when all investors have the same informational type (either they are all informed or uninformed). As we show next, differences in informational type creates different price schedules for different quality shocks θ, which breaks the correlation between the in-the-money share and overpayment, as data suggests.

5 Price Effects of Asymmetric Information

Since heterogeneity in wealth or risk aversion cannot be important factors in accounting for the heterogeneity in bidding behavior that we observe in the data, henceforth we will drop the fundamental type distinction in order to focus on heterogeneity in information. We start our characterization of an auction equilibrium when there is no information heterogeneity (as a benchmark) and then we introduce asymmetric information. For expositional simplicity, among the CRRA preference family, we assume log utility.

5.1 Symmetric Information (or Ignorance)

To start our analysis of symmetric information, we consider first the extreme example in which investors know the demand shock η. In case of symmetric information the investors also know the quality shock θ. In case of symmetric ignorance, investors are equally uninformed about the quality shock. In this case our auction equilibrium is equivalent to a standard competitive equilibrium, as there is a representative investor. In both cases the solution follows the same construction. Under symmetric information it is based on $\kappa = \kappa_\theta$ and under symmetric ignorance on $\kappa = E(\kappa_\theta)$. We can simplify notation to $P(\eta)$ for prices and $B(\eta)$ for bond purchases. The auction-clearing condition in state η is

$$(1-\eta)B(\eta)P(\eta) = D. \tag{18}$$

The investor's f.o.c. is given by

$$\frac{\kappa P(\eta)}{W - B(\eta)P(\eta)} = \frac{(1-\kappa)(1-P(\eta))}{W + B(\eta)(1-P(\eta))}.$$

And we can solve for prices in closed form using the two previous expressions

$$P(\eta) = 1 - \kappa - \kappa \frac{\frac{D}{1-\eta}}{W - \frac{D}{1-\eta}}. \tag{19}$$

Plugging this price back into the market clearing condition allows us to derive an expression for the level of total bond purchases. At this benchmark, no one bids above the marginal price, and all bids are in-the-money in the sense that there is no incentive to bid at anything but the correct marginal price.

Now assume that the demand shock η is not known but all investors still share beliefs about the probability of default, κ. As total expenditures are monotonically increasing in η and bids must be strictly positive, $B(\eta) > 0$ for all $\eta \in \mathcal{H}$, the short-sale constraint for the investors cannot bind for any η. In this simple case we can rewrite the first-order condition (17) for the representative investor at η^*,

$$\sum_{\eta \geq \eta^*} \left\{ \begin{array}{c} -U'(\mathcal{B}_{RF}(\eta))\kappa P(\eta^*) \\ +U'\left(\mathcal{B}_{RF}(\eta) + \sum_{\hat{\eta} \leq \eta} B(\hat{\eta})\right)(1-\kappa)(1-P(\eta^*)) \end{array} \right\} h(\eta) = 0 \qquad \forall \eta^*. \tag{20}$$

Notice that the cumulation of marginal utilities is determined by the holdings of risk-free bonds, which are the residual of expenditures evaluated at the bid price rather than at the marginal price. As a result, the system of first-order conditions is not block-recursive and must be solved simultaneously. This introduces both computational and analytical complexity. For convenience in exposition and computation we can rewrite this system defining the following vectors of prices, returns, and bids

$$\vec{P} = \begin{bmatrix} P(\eta_0) \\ \vdots \\ P(\eta_M) \end{bmatrix}, \left(1 - \vec{P}\right) = \begin{bmatrix} 1 - P(\eta_0) \\ \vdots \\ 1 - P(\eta_M) \end{bmatrix}, \vec{B} = \begin{bmatrix} B(\eta_0) \\ \vdots \\ B(\eta_M) \end{bmatrix}$$

and the following triangular matrices of dimension $M \times M$

$$\mathbf{P} = \begin{cases} P_{rk} = P(\eta_r) \text{ if } r \leq k \\ P_{rk} = 0 \text{ o.w.} \end{cases}, \quad \mathbf{1} - \mathbf{P} = \begin{cases} 1 - P_{rk} = 1 - P(\eta_r) \text{ if } r \leq k \\ 1 - P_{rk} = 0 \text{ o.w.} \end{cases}.$$

Price vector $\mathbf{P}$ must then solve the stacked system of first-order conditions

$$-U'\left(W - \mathbf{P} \times \vec{B}\right) \cdot \vec{P} \cdot \kappa + U'\left(W + [\mathbf{1} - \mathbf{P}] \times \vec{B}\right) \cdot \left[1 - \vec{P}\right] \cdot [1 - \kappa] = 0. \tag{21}$$

5.2 Asymmetric Information

In this case, heterogeneity comes through $n \in (0, 1)$, the fraction of the investors who are informed about θ; i.e. whether the default risk is κ_g or κ_b. Given the discrete grid for demand shock $\eta \in \mathcal{H}$, there is a discrete set of states $\{s_0, \ldots, s_N\}$ indexed by r, with $N = M \times 2$. Without loss of generality, order states in decreasing order of prices, $P(s_r) > P(s_{r+1})$, and let

$$\vec{P} = \begin{bmatrix} P(s_0) \\ \vdots \\ P(s_N) \end{bmatrix}, \vec{B}^U = \begin{bmatrix} B^U(s_0) \\ \vdots \\ B^U(s_N) \end{bmatrix}, \left(1 - \vec{P}\right) = \begin{bmatrix} 1 - P(s_0) \\ \vdots \\ 1 - P(s_N) \end{bmatrix}, \varkappa = \begin{bmatrix} \kappa(\theta_0) \\ \vdots \\ \kappa(\theta_N) \end{bmatrix},$$

$$\mathbf{P} = \begin{cases} P_{rk} = P(s_r) \text{ if } r \leq k \\ P_{rk} = 0 \text{ o.w.} \end{cases}, \quad \mathbf{1} - \mathbf{P} = \begin{cases} 1 - P_{rk} = 1 - P(s_r) \text{ if } r \leq k \\ 1 - P_{rk} = 0 \text{ o.w.} \end{cases}.$$

Given this notation, the system of first-order conditions that pins down the optimal bids of uninformed investors is

$$U'\left(W - \mathbf{P} \times \vec{B}^U\right) \cdot \vec{P} \cdot \varkappa + U'\left(W + [\mathbf{1} - \mathbf{P}] \cdot \vec{B}^U\right) \cdot \left[1 - \vec{P}\right] \cdot [1 - \varkappa] \leq 0,$$
$$\text{and} \quad B^U(s_r) = 0 \text{ if the inequality is slack.}$$

For informed investors, the analogous system of equations is

$$U'\left(W - \mathbf{P} \times \vec{B}^I\right) \cdot \vec{P} \cdot \kappa_{\theta(s)} + U'\left(W + [\mathbf{1} - \mathbf{P}] \times \vec{B}^I\right) \cdot \left[1 - \vec{P}\right] \cdot [1 - \kappa_{\theta(s)}] = 0.$$
$$\text{and} \quad B^I(s_r) = 0 \text{ for all } \theta \neq \theta(s_r).$$

This leads to the following results. Observe first that the gap between the high-quality and low-quality price schedules is increasing in n. When n is large, uninformed investors face a substantial adverse selection problem: if they bid on the high-quality schedule, they will overpay in the bad state since the government is sure to accept their high-state bids and executes them at the high bid price. For n sufficiently large ($n = 0.6$ in our example), the uninformed may therefore refrain from placing any bids on the high-price schedule. This has two effects. First, an uninformed investor who bids only on the low-price schedule knows that, *conditional* on a bid being accepted, the state must be bad. Hence they choose the same portfolio as informed investors when $\theta = b$, and the composition of informed and uninformed investors is irrelevant in determining bad bond prices, this is the low-quality schedule is locally independent of n. Second, precisely because the uninformed do not participate at high prices, informed investors have to buy more bonds per capita in the good state, and are therefore disproportionately exposed to the government's default risk. Since there are fewer participants as n decreases, the high-price schedule must decline. This discussion leads to the next Proposition.

Proposition 3. *When the gap between price schedules in the informed equilibrium is large enough, for all large n such that $B^U(g, \eta) = 0$ for all η, then $B^U(b, \eta) = B^I(b, \eta)$ for all η.*

Proof. Assume $B^U(g, \eta) = 0$ for all η, then the uninformed know that conditional on their bids being on the money the state should be $\theta = b$. As in the bad state informed and uninformed are identical, in terms of their characteristics and knowledge, $B^U(b, \eta) = B^I(b, \eta)$ for all η. □

We use this case in which n is large enough to discourage bidding of uninformed at high-quality bond prices to show that different information types can display differences in-the-money shares and still overpay similarly conditional on acquiring bonds, consistent with the patterns documented for Cetes data.

Proposition 4. *Assume $B^U(g, \eta) = 0$ for all η, uninformed investors bids are in average less in-the-money than informed investors, but their overpayment is similar.*

Proof. Assume $B^U(g, \eta) = 0$ for all η. The average in-the-money shares for uninformed investors is,

$$\mathbb{E}[ITM^U(\theta, \eta)] = f(b)\mathbb{E}_{|\theta=b}[ITM^U(b, \eta)].$$

As informed investors observe θ when making their bids, their average in-the-money shares is,

$$\mathbb{E}[ITM^I(\theta, \eta)] = f(b)\mathbb{E}_{|\theta=b}[ITM^I(b, \eta)] + f(g)\mathbb{E}_{|\theta=g}[ITM^I(g, \eta)],$$

From Proposition 3, $B^U(b, \eta) = B^I(b, \eta)$ for all η and then $\mathbb{E}_{|\theta=b}[ITM^U(b, \eta)] = \mathbb{E}_{|\theta=b}[ITM^I(b, \eta)]$.

Since $ITM^I(g, \eta_M) = 1$, the expectation $\mathbb{E}_{|\theta=g}[ITM^I(g, \eta)]$ is positive and then

$$\mathbb{E}[ITM^U(\theta, \eta)] < \mathbb{E}[ITM^I(\theta, \eta)]$$

Also, since $B^U(g, \eta) = 0$ for all η

$$\mathbb{E}\left[\frac{AP^U(\theta, \eta)}{P(\theta, \eta)}\right] = \mathbb{E}\left[\frac{AP^I(b, \eta)}{P(b, \eta)}\right].$$

From Proposition 3, $B^U(b, \eta) = B^I(b, \eta)$ for all η. Assume that

$$\mathbb{E}_{|\theta=g}\left[\frac{AP^I(g, \eta)}{P(g, \eta)}\right] \approx \mathbb{E}_{|\theta=b}\left[\frac{AP^I(b, \eta)}{P(b, \eta)}\right].$$

which is the case with different price schedules. In such case,

$$\mathbb{E}\left[\frac{AP^U(\theta, \eta)}{P(\theta, \eta)}\right] \approx \mathbb{E}\left[\frac{AP^I(\theta, \eta)}{P(\theta, \eta)}\right].$$

□

Proposition 5. *Fix $n \in (0, 1)$. The quality-contingent price schedules converge to each as other as $n \to 0$. That is, $\lim_{n\to 0} P([g, \eta]; n) = P([b, \eta]; n)$ for all $\eta < \eta_M$.*

Proof. Denote by $X^U([\theta, \eta]; n)$ the uninformed investors' total expenditures on risky bonds. Then $\lim_{n\to 0} X^U([\theta, \eta]; n) \to D/(1-\eta)$ for all θ by auction-clearing. As uninformed bids are made unconditionally on θ, $\lim_{n\to 0} X^U([g, \eta]; n) \to X^U([b, \eta]; n)$ and $\lim_{n\to 0} P([g, \eta]; n) \to P([b, \eta]; n)$. By Proposition 1, $P(\theta, \eta; n)$ is strictly decreasing in η given θ and n. When $n \to 0$, prices must then be sorted by η. That is, there is always a ϵ small enough such that for $\eta' - \eta = \epsilon$, i.e $P([\theta, \eta]; n) < P([\theta', \eta']; n) < P([\theta, \eta']; n)$. This proof does not apply at extreme values of η, and then convergence will not happen at $\eta = 0$ and $\eta = \eta_M$. □

5.3 Numerical Illustration of Asymmetric Information Equilibrium

Solving for equilibrium prices is analytically intractable. To illustrate the main properties of the auction equilibrium with asymmetric information, we use a numerical example based on log-preferences and arbitrary parameters. In the next section we discipline parameters with Mexican Cetes data.

As per capita supply in our model is given by $D/(1 - \eta)$, and it is useful that supply is uniformly distributed, we make a change in variables and use the previously mentioned supply shocks $\psi = (1-\eta)^{-1}$. The demand shock η distributed between 0 and η_M is equivalent to a supply shock ψ distributed between 1 and $\psi_M = (1-\eta_M)^{-1}$. We will assume that the ψ shocks are uniformly distributed on this interval. We use the parameters in Table 3 below.

Table 3: Illustration Parameterization

$\kappa_g = 0.15$	$\kappa_b = 0.35$
$W = 250$	$D = 60$
$f(b) = 0.5$	$\psi \sim U[1, 1.20]$

Symmetric Information (or Ignorance): We start out by computing the symmetric information (and ignorance) price function (19) with information about the supply shock (complete information), as in subsection 5.1, and without such information. The results are plotted in the first panel of Figure 4. The price schedules with and without complete information, given a default probability, are very close. In both cases schedules are fairly flat, with the complete information price starting higher and falling faster: investors bid more aggressively when they do not have to worry about overpaying because of uncertainty about the supply shock. One can also see that the schedules under symmetric information in each quality shock or under symmetric ignorance (the intermediate schedule) are very similar in shape, showing that information extremes are very similar in shape and just differ in levels.

Figure 4: Symmetric Benchmarks

(a) Price Schedules

(b) Bid Schedules

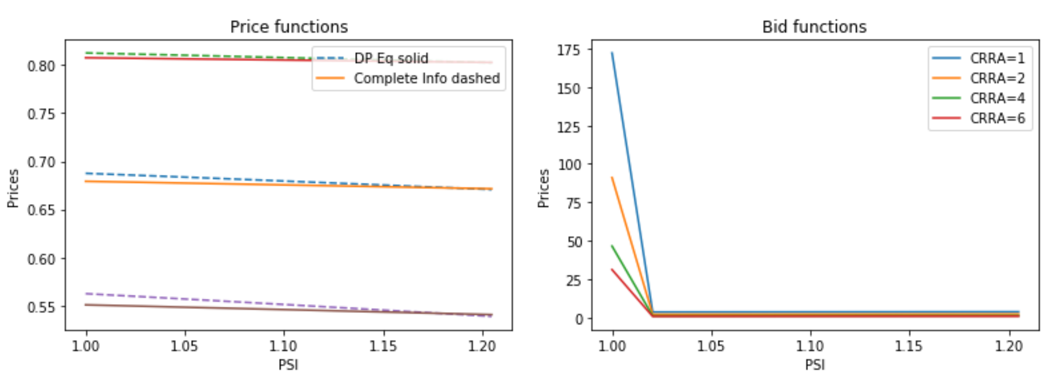

In the second panel of Figure 4 we show optimal bid schedules for different coefficient of risk aversion (CRRA) fixing the price schedule under symmetric ignorance. The first thing to note from the figure is that the bulk of investor's bids are made at the highest price. This is simply because the lion share of the supply happens at that price and only incrementally increases with ψ, and the equilibrium price schedule must induce this pattern of bids in aggregate to clear the market. As we discussed in subsection 4.2, a higher risk aversion shades down the bids, but the effect on in-the-money share is analytically unclear. In this particular illustration the share of the bids that are in-the-money is essentially identical, ranging from 0.83 at $\sigma = 1$, to 0.80 at $\sigma = 6$. While more risk averse investors bid less, they do not bid "less aggressively," since they shade down their bids by the same factor everywhere, with overpayment being virtually identical too.

Asymmetric Information and the Role of n: Now we perform comparative statics with respect to asymmetric information, n. Figure 5 shows equilibrium prices as we shrink n from one to zero. Adverse selection of uninformed bidders and the concentration of default risk among informed investors play a prominent role in understanding the forces behind equilibrium prices under asymmetric information. Each example in Figure 5 is chosen to highlight one of these forces.

Observe first that the gap between the high-quality and low-quality price schedules is increasing in n. When n is large, uninformed investors face a substantial adverse selection problem: if they bid on the high-quality schedule, they will overpay in the bad state since the government is sure to accept their high-state bids and executes them at the high bid price. For n sufficiently large ($n = 0.6$ in our example), the uninformed may therefore refrain from placing any bids on the high-price schedule. This has two effects. First, an uninformed investor who bids only on the low-price schedule knows that, *conditional* on a bid being accepted, the state must be bad. Hence they choose the same portfolio as informed investors when $\theta = b$, and the composition of informed and uninformed investors is irrelevant in determining bad bond prices, this is the low-quality schedule is locally independent of n. Second, precisely because the uninformed do not participate at high prices, informed investors have to buy more bonds per capita in the good state, and are therefore disproportionately exposed to the government's default risk. Since there are fewer participants as n decreases, the high-price schedule must decline. This discussion leads to the next Proposition.

When n is sufficiently small ($n = 0.4$ in our example), prices on the high-quality schedule are low enough such that the uninformed investors are less worried about adverse selection and begin to bid on both price schedules. Since bids on the high-price schedule are also executed in the bad state, there is less residual demand that needs to be met by marginal bids on the low-price schedule. Hence the low-price schedule rises.

The adverse selection effect continues to operate as n decreases further (to $n = 0.1$ in our example). In particular, because the per-capita bids of the uninformed remain below

Figure 5: DP auction equilibrium as n falls.

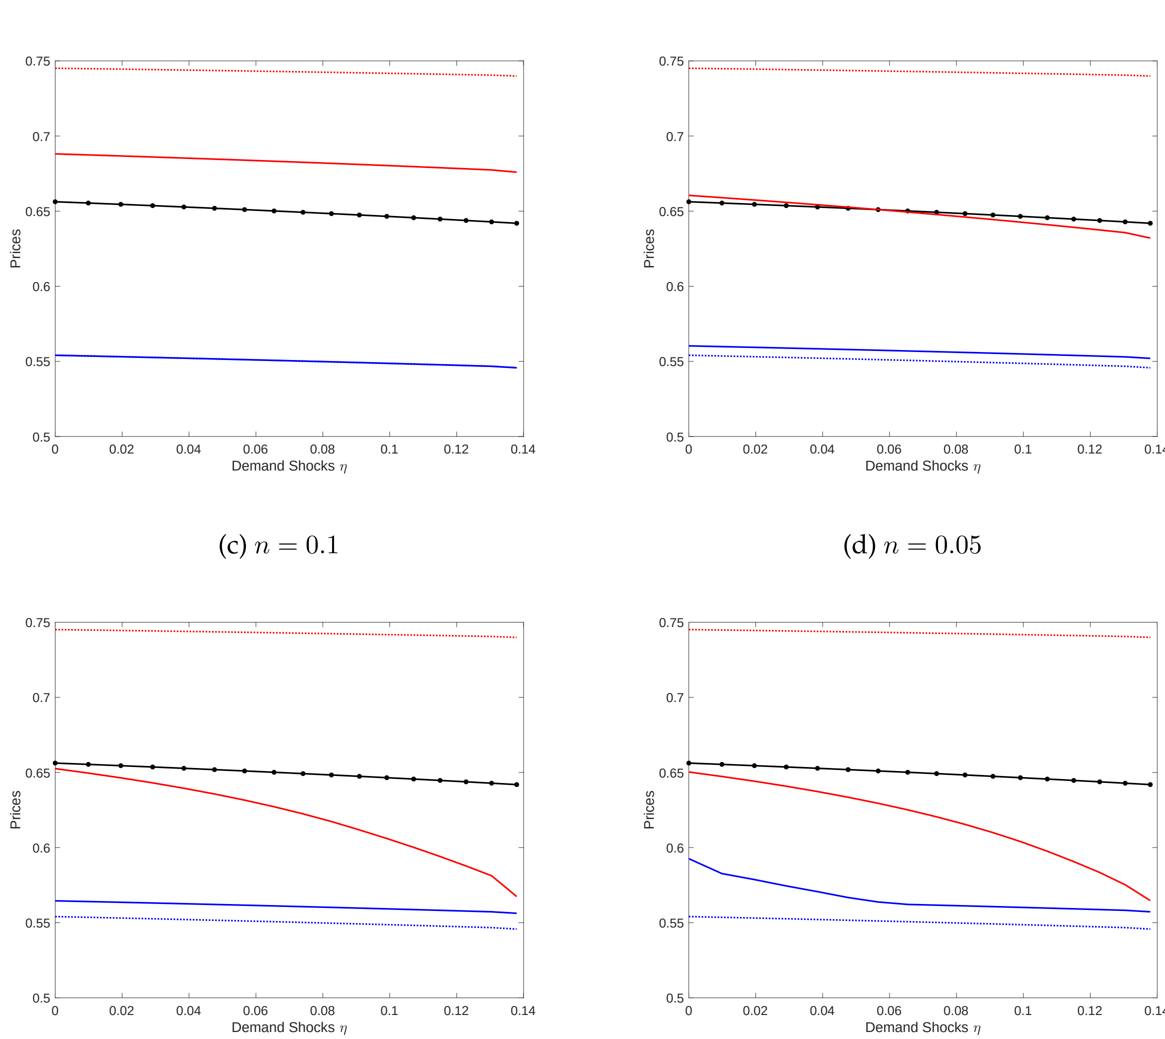

(a) $n = 0.6$ (b) $n = 0.3$ (c) $n = 0.1$ (d) $n = 0.05$

those of the informed on the high-price schedule, reductions in n continue to further concentrate default risk in informed portfolios. This forces a large fraction of the high-quality schedule to drop below the *uninformed* price schedule. That is, the adverse selection effect may be severe enough that prices are lower than in the uninformed equilibrium for both high- *and* low-quality bonds. Finally, when n is very small ($n = 0.02$ in the figure), price schedules start overlapping. Uninformed investors are now willing to participate fully on both schedules and prices converge to the uninformed price schedule as $n \to 0$.

Yields: The yield of a government bond sold at price P is the promised return,

$$Yield = \frac{1-P}{P}.$$

We compute the quantity-weighted average yield using the individual yields on all sold bonds and considering all participating investors (both informed and uninformed).[12] This average yield captures the *risk-neutral component* of the government's payoff and the *risky component*, given by the variation in the average yield conditional on the quality of the bond, capturing the government's exposure to demand shocks.

In Figure 6 we compare certain key features of our equilibria as we increase the fraction of informed investors, n. The first panel shows the average ITM shares for informed and uninformed bidders. As n increases from 0, the participation of the uninformed on the high-quality bond schedule decreases and completely ceases close to $n = 0.4$. At this level of asymmetric information the difference of ITM share between the informed and the uninformed is maximized. In the second panel of Figure 6 we plot the average overpayment by information type, which is the average price paid relative to the marginal price. As one can see from the figure, an informed share of 0.4 or higher generates very similar overpayments by the informed and the uninformed, as the uninformed only participate on the low-quality bond schedule.

Figure 6: Examining Impact of Informed Share

(a) ITM Shares

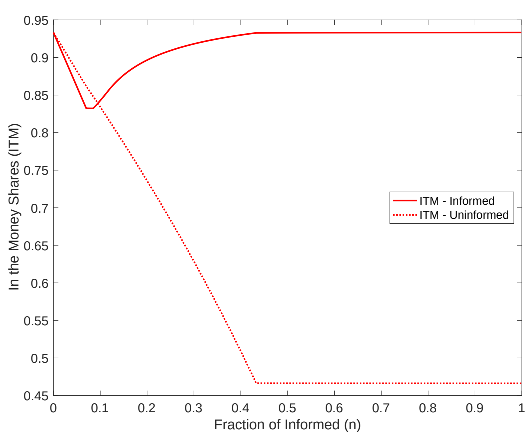

(b) Overpayment and Yield

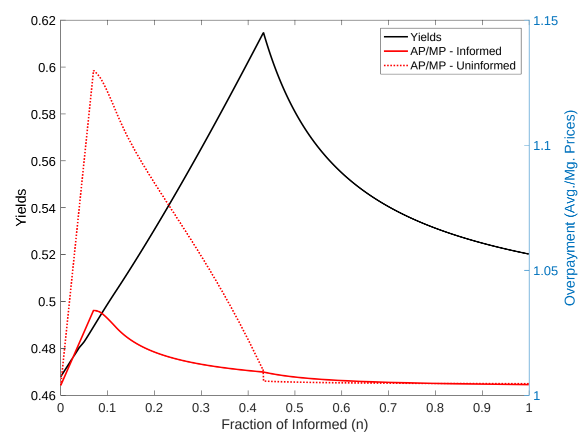

In the second panel of Figure 6 we also plot average yields, which are hump-shaped:

[12]Relatedly, the government's debt burden can be defined as $D/P = D(1 + Yield)$: the government faces a higher average debt burden if bonds trade at a higher average yield.

increasing for low levels of n and declining for high levels of n. If n is low, uninformed investors bid on both schedules, but face adverse selection because bids on the high-quality schedule are also accepted if the bond is bad. This effect depresses prices on the high-quality schedule to an extent that is not fully compensated by an increase of the low-quality schedule. Indeed, the *adverse selection effect* grows in importance as n rises, up to the point at which uninformed investors stop bidding on the high-quality schedule. When only informed investors participate, further increases in n lead to lower per-capita risk exposure. As a result, informed investors' information rents are gradually competed away as n grows. In that case, bad-state prices are unaffected because the uninformed choose the same portfolio as the informed conditional on the bad state. This *cannibalization effect* raises prices in the good state and reduces the yield as $n \to 1$.

An important insight of this analysis is that only intermediate levels of n have both the potential to generate a higher ITM share for informed, and at the same time a similar overpayment than uninformed. However, it is exactly at those levels of asymmetric information when a government is forced to offer the highest average yield for their bonds and then face the largest debt burden.

6 The Extent and Nature of Asymmetric Information

6.1 Calibration Exercise

In this section we calibrate both investors' bidding behavior and the dynamics of Cetes marginal prices. The goal of this calibration is twofold. First, examining the quantitative extent to which the bidding behavior of investors participating in Cetes auctions can be accounted for by our simple model with asymmetric information. Second, matching the dynamics of marginal prices informs us what sort of information investors have asymmetric access to. In particular, if a price in a week's auction can predict quite well the price of the next week's auction, what sort of information asymmetries are relevant? By calibrating the model we show that, even though conditionally there is not much uncertainty for investors (what we denote as the *public regime* being common knowledge) there is enough uncertainty within a regime that can explain the patterns of bidding behavior.

The goal of this section is not to match dynamics and cross-investor bidding behavior moments perfectly, but instead to provide a proof of concept that both set of moments can be captured quite well with plausible parameters within the simplest structure of our model, without imposing extra bells and whistles. In principle we can always include more parameters and richness into the structure to match the moments perfectly, but this is beyond the goal and scope of the paper.

Calibrated Parameters: The two-period model we have discussed above has to be extended to accommodate dynamics. We do so by repeating the structure dynamically, with a long-lived government and short-lived (for two periods) investors, who buy bonds in the first period of life and consume in the second. At each auction period the country is in one of two possible *public regimes*, $z \in \{1, 2\}$. Each public regime has, as in the text, two possible states $\theta \in \{b, g\}$. Public regimes follow a symmetric transition matrix

$$\begin{bmatrix} \rho & 1-\rho \\ 1-\rho & \rho \end{bmatrix}$$

parameterized by the single parameter ρ. Because of our symmetry assumption each public regime has equal unconditional probability to occur. Observe that investors can typically learn the public shock from past prices. This is grounded int he following observations:

In this simple dynamic setting, the minimum set of parameters we need to calibrate are: the default probabilities, $\kappa_{g,z}$ and $\kappa_{b,z}$, and the probability of the good state $f(g_z)$ in each public regime $z \in \{1, 2\}$, the transition probability ρ, the maximum demand shock η_M (we maintain our assumption that η follows a uniform distribution), and most importantly, the mass corresponding to the informed largest bidder (which we call n_{big}) and the fraction of informed investors, n. In what follows we combine the description of the moments we match with the role of these parameters on their determination. We focus on calibrating Cetes bonds of 28 days of maturity, but similar results hold for the other maturities.

6.2 Disciplining Moments

The first two important moments in the data are the average and standard deviation of marginal prices, which are 0.98 and 0.017 respectively. The standard deviation gives us a natural measure of the *unconditional* volatility of prices. These moments have an immediate mapping with the mean and variance of κ across the two regimes and the two states in each regime.

The second set of moments are (i) the coefficient $\beta_1 = 0.98$ from the price regression in (1), and the conditional volatility of prices as measured by the regression's $R^2 = 0.97$. This high predictability of prices is informative about the nature of information, in particular the persistence of information and the degree of *conditional* (on previous period's auctions) volatility of prices. In terms of the calibration this information points towards highly persistent public regimes (high ρ), without much volatility of within regime states.

How much volatility within regimes and how different those regimes are, however, should also be consistent with differences in bidding behavior. The in-the-money share of the largest bidder is in average 0.84 (this is, in average, the largest bidder buys 84% of his submitted bids), while the in-the-money share of the rest of investors is just 0.33 in average. At the same time, all bidders (both the largest and the rest) pay in average 0.1% above marginal price. According to our model, if the largest bidder were informed and there were no demand uncertainty, both his in-the-money share and overpay would be 1. The fact that this is not the case, helps us discipline η_M, this is the range of demand shocks that prevents the largest bidder to always be in the money.

Also according to our model, the rest of investors have a low in-the-money share because part of them are uninformed, who choose not to bid at high prices. As n_{big} is defined as the "size" of the largest bidder, the total fraction of informed investors in the market, n determines the composition of the rest of investors and the strength of adverse selection. Adverse selection comes from the spread between κ_g and κ_b in a given public regime. This creates a tension in the fitting exercise, as both increasing $\eta_M$and the gap in default probabilities raise the conditional uncertainty in the model given the public regime.

In the data, the largest bidder buys in average 38% of the bonds. Formally,

$$\frac{n_{big} \sum B_I}{n \sum B_I + (1-n) \sum B_U} = 0.38$$

Under the assumption that the largest bidder is informed, this value puts a lower bound on the fraction of informed investors. Further, the difference in-the-money shares between the largest bidder and the rest puts an upper bound. To see this notice that the in-the-money share of the rest of bidders is a combination between informed and uninformed investors

$$\frac{1-n}{1-n_{big}} ITM^U + \frac{n-n_{big}}{1-n_{big}} ITM^I,$$

As n is maximized when $ITM^U = 0$, then $\frac{n-n_{big}}{1-n_{big}} 0.84 = 0.33$ and

$$n_{big} < n < n_{big} + \frac{ITM^U}{ITM^I}(1-n_{big}).$$

The fact that n in intermediate is consistent with the range of informed investors that creates enough adverse selection to shy uninformed investors to bid at high prices. From this discussion, η_M and n_{big} are relatively well identified by in-the-money shares and participation of the largest bidder respectively, while the combination of n, ρ, κ and the within-regime state probabilities jointly affects the other moments.

6.3 Capturing Moments

Under the denomination "Baseline Model," we have chosen a set of parameters (listed in Table 4) that minimizes the sum of squared errors between the model generated moments (second column of Table 5) and the data targets (first column of Table 5).[13] We match quite well the mean and standard deviation of marginal prices, as well as the in-the-money share of the largest bidder and the extent of overpayment for both the largest and the rest. Qualitatively, but not quantitatively, we do well with respect to the in-the-money share of the rest of the investors (lower than for the largest bidder but much larger than in the data, 0.62 versus 0.33) and the extent of predictability (positive but much smaller than in the data, 0.7 versus 0.97).

Table 4: Calibrated Parameters

Common Parameters	Values	Model Specific	Baseline Model	Black Swan Model
κ_{g1}	0.001	κ_{g2}	0.019	0.02
κ_{b1}	0.014	κ_{b2}	0.029	0.50
$f(g_1)$	0.65	$f(g_2)$	0.65	0.95
ψ_{max}	1.3	n	0.40	0.39
ρ	0.999	n_{big}	0.22	0.17

The main reason the calibration fails in these dimensions is the tension of matching bidding behavior given the high unconditional volatility and the low conditional volatility we observe in the data. To see why, we look more closely at our outcomes. Some model elements behind the Baseline calibration are plotted in Figure 7. The price schedules for both public regimes, plotted in the first panel, show that we are able to accommodate both a moderately high degree of unconditional volatility and a low degree of conditional volatility by having two high price schedules in public regime 1, and two low price schedules in public regime 2, along with each public regime being highly persistent.

As default rates in average have to be very close to 0, a very small gap between κ_g and κ_b is sufficient to induce enough adverse selection. The second and third panel of Figure 7 show both the informed and uninformed bid schedules for public regimes 1

[13] The level of the persistence parameter reflects the fact that this is a weekly model; $.999^{52} = 0.95$.

Table 5: Calibration Targets: Data vs. Model

Target	Data	Baseline Model	BS Model	PP Model
Mean Price	0.98	0.98	0.96	0.97
Std. Price	0.02	0.02	0.10	0.02
Regression β	0.98	0.84	0.14	0.96
Regression R^2	0.97	0.70	0.02	0.92
LB ITM share	0.84	0.87	0.88	0.88
Rest ITM share	0.33	0.62	0.43	0.41
UI ITM share		0.45	0.18	0.16
Overpay LB	1.001	1.001	1.003	1.002
Overpay Rest	1.001	1.004	1.005	1.004
Share LB	0.38	0.38	0.37	0.38

and 2, respectively. The high unconditional volatility means, however, that one of the public regimes had to include moderately high default rates (lower prices). But this in turn means that the risk-free and the risky bond are close substitutes, and then getting the informed to hold all of the risk when the bond is good pushed the price of bad bonds so low that the uninformed also bid on it. This can be seen in the public regime 2 bid schedules: the uninformed bid more on the high-quality schedule than the low-quality schedule. In particular, they bid their largest amount at the highest possible price, thereby insuring a fairly high ITM for uninformed in public regime 2. In this baseline calibration, uninformed bidders would display a lower in-the money share only if they were much more concerned about adverse selection in public regime 2, which would be the case with a much higher κ_{b2}, which would, however, distort other moments, such as average prices.

Figure 7: Baseline Figures

(a) Price Schedules

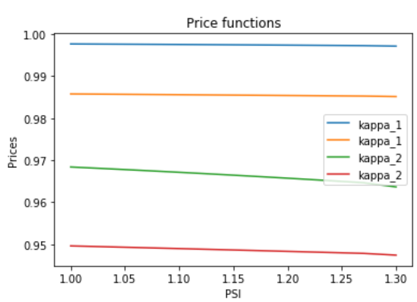

(b) Bid Schedules 1

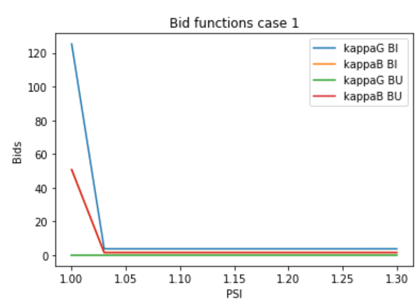

(c) Bid Schedules 2

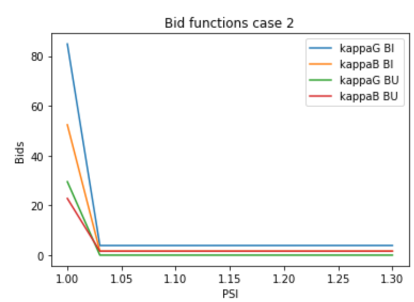

As an alternative calibration, which we label "Black Swan" (BS), we increase κ_{b2} considerably, while making κ_{g2} much more likely. This alternative prevents average prices on becoming too large while increasing adverse selection concerns for uninformed bidders in public regime 2 (we maintain default probabilities in public regime 1, which induce sufficient adverse selection already). To gauge the plausibility of this calibration, note that the likelihood that a BS event does not occur in a decade is 31%, consistent with the fact that over the three decades before our dataset, Mexico experienced two major crises.[14]

Figure 8 plot prices and bid functions for this alternative calibration, while the generated moments are reported in the third column of Table 5. The first panel of the Figure shows that the price schedule for bad quality bonds in public regime 2 is now much lower than in the baseline case, consistent with having a very high default probability. The small possibility of this extreme outcome still has a substantial impact on adverse selection and on discouraging the uninformed to bid at high prices in public regime 2 (see third panel of the Figure). As a result, the in-the money share of the rest of investors declines substantially, though not as low as in the data.

Figure 8: Black Swan Figures

(a) Price Schedules (b) Bid Schedules 1 (c) Bid Schedules 2

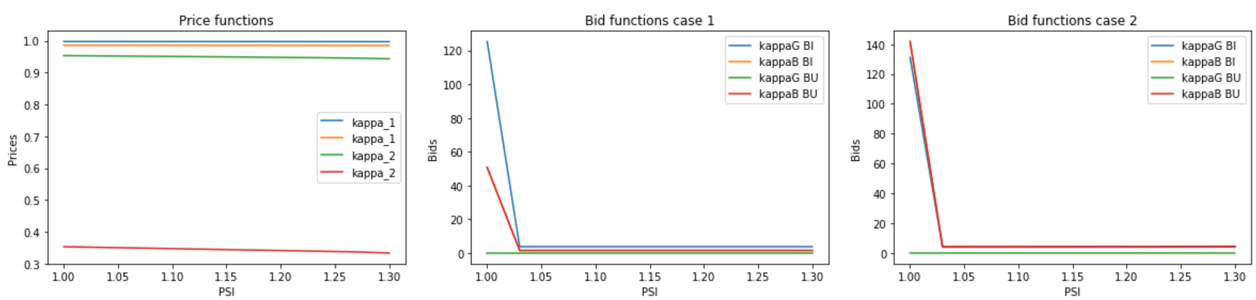

Our Black Swan alternative, however, fails more dramatically in the dynamic components. Even though the probability of the black swan is fairly low, it substantially raises the standard deviation of prices (to 0.10), greatly depresses our measure of persistence (β falls to 0.14) and increases conditional uncertainty (R^2 falls to 0.02).[15]

[14]The probability of not having a black swan in a decade obeys the following backwards recursion: denote π^i_j is the probability that a black swan (a bad state in public regime 2) *does not happen* j weeks from the end of the decade, where i denotes the public regime. Then

$$\pi^1_j = .999 * \pi^1_{j-1} + (1 - .999) * [.95 \times \pi^2_{j-1}]$$
$$\pi^2_j = .999 * [.95 \times \pi^2_{j-1}] + (1 - .999) * \pi^1_{j-1}.$$

The unconditional probability of not having a crisis over a decade is then $(\pi^1_{520} + \pi^2_{520})/2 = 0.31$.

[15]This dynamic failure comes from the nonlinear nature of these metrics and the high weight associated with large pricing errors.

One way to rationalize the data given this alternative calibration is arguing that the data suffers from a *peso problem*, a phrase attributed to Milton Freedman to explain the gap between Mexican and U.S. deposit rates during the 1970s. A peso problem arises in asset pricing models when market participants anticipate the possibility of a discrete change in the probability distribution generating outcomes, and hence their subjective distribution differs from the distribution which has generated the historical data (see **?**) and **?**)). The possibility of such a shift seems particularly salient in Mexico during this period 2001-2017 from which we obtain our targets. In our considered timeframe Mexico did not experience any major crisis, but it did suffered sovereign debt crises in both the 1980s and 1990s, with very high and volatile bond spreads.

To examine the extent to which a peso-problem can account for the data, we kept the parameters at the Black Swan calibration for the purposes of computing the equilibrium outcomes, hence the price and bid functions are as shown in Figure 8. But, in computing the time series implications of the model, we set the probability of a high-quality bond in public regime 2 at 1 rather than 0.95. This is equivalent to assuming that while a black swan event is anticipated as possible by the investors in our model, it did not occur in the sample realization. We label this alternative as a "Peso Problem" (PP) calibration and the results are reported in the last column of Table 5. This version of our model does very well, matching almost all of the data moments closely. In particular, the dynamic moments: the regression β has risen to 0.96, while the R^2 has risen to 0.92, which are quite close to the data.

This quantitative exploration sheds light on the sources of information asymmetries. Investors do not differ much in their information about usual price movements, because they can all access publicly available data on the health of the country's finances, the previous auction results and even the functioning of secondary markets. The joint patterns of bidding behavior and dynamic bond price evolution, however suggests that some investors may be particularly well informed about events leading to large price swings, and then more difficult to access information (such as internal political decisions and other sorts of "insider" information).

7 Conclusion

By using a unique dataset of Mexican primary Cetes (domestically denominated bonds) weekly auctions between 2001 and 2017, we document a set of 3 key facts about prices and bidding behavior patterns. First, the unconditional price volatility is sizable, the volatility conditional on the previous auction price is low, which suggests high predictability of prices. Second the largest bidders tend to buy a larger fraction of submitted bids than the rest, which in a "pay as you bid" auctions (as the one conducted by Mexico in the considered period) suggests that they bid at higher prices. At the same time, and despite the low conditional volatility, the rest of the bidders tend to buy a fairly small fraction of their bids. Third, and prima facie inconsistent with the previous fact, the largest bidders do not pay more conditional on buying.

We then construct a *Walrasian* model of price-discriminating sovereign debt auctions in which participating investors can differ in their wealth, on their market power, risk-aversion and/or information. We use the implications of our model to show that the documented heterogeneity in bidding patterns across investors cannot be explained by differences in wealth, market power or risk aversion. However, it is consistent with the largest investors being more or better informed about the probability the bonds default than the rest of bidders.

Finally, we perform a calibration of the model that is informative both about the extent of asymmetric information and its nature. First, to conform to the data, our model implies an intermediate fraction of investors informed, which according to the model generates a larger debt burden for Mexico compared to situations in which most of investors were informed or most were ignorant. Second, to generate bidding by the less informed investors that is consistent with the data, we need not only a small amount information heterogeneity with respect to information about usual price movements, but also, and critically, information asymmetries about the advent of low probability events that generate large and sudden price swings (such as liquidity crises or currency runs).

The goal of this paper was providing a tractable model of sovereign debt primary markets with asymmetric information that is useful to confront our novel empirical findings on bidders behavior and to identify the main frictions behind these auctions. As we make a case of the existence of asymmetric information and its relevance for sovereign bond yields, a natural extension is understanding its implication more generally. In **?**), for instance, we explore the role of information asymmetry for spillovers across countries. By endogenizing information acquisition and extending the setting to accommodating many countries and secondary markets, we show the possibility of multiplicity on informational regimes and the possibility that a country suffers from a shock in an unrelated country through endogenous asymmetric information. Our model, however, is tractable enough to explore many other dimensions, such as the incentives of the government to disclose information or to affect information asymmetry by using different auction protocols. We leave these extensions for future research.