

Asymmetric Information and Sovereign Debt: Theory Meets Mexican Data*

Harold Cole[†] Daniel Neuhann[‡] Guillermo Ordoñez[§]

February 2021

Abstract

Using a novel data set containing all bids by all bidders for Mexican government bonds from 2001 to 2017, we demonstrate that asymmetric information about default risk is a key determinant of primary market bond yields. Empirically, large bidders do not pay more for bonds than the average bidder but their bids are accepted more frequently. We construct a model where investors may differ in wealth, risk aversion, market power and information, and find that only heterogeneous information can qualitatively account for these patterns. Moreover, asymmetric information about rare disasters can quantitatively match key moments of bids and yields, both within and across periods.

*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. We also thank seminar participants in Toulouse and St. Marten's Conference on the Sand for comments. Cole and Ordoñez 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@mcombs.utexas.edu)

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

1 Introduction


Governments frequently finance themselves by issuing sovereign bonds. In emerging markets, these bonds usually offer a risk premium on the order of the U.S. equity premium even when accounting for default losses (Aguar et al. (2016)). Sovereign bond markets also experience recurring episodes during which yield spreads are elevated and volatile for extended periods of time. The sovereign debt literature typically analyzes these price movements by assuming competitive price determination in *secondary markets*. However, government finances are determined in *primary markets*, and primary markets do not employ competitive pricing protocols. In this paper, we analyze the determinants of sovereign bond prices by developing and empirically validating a micro-founded model of primary markets with rich investor heterogeneity and realistic pricing protocols. Using a new data set from Mexico, we find that asymmetric information about default risk can *qualitatively and quantitatively* account for many salient features of bond prices and bidding behavior.

The data set we construct contains primary market bids and prices for *Cetes* bonds in Mexico from 2001 to 2017. *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 our sample period, *Cetes* of different maturities were auctioned weekly using a *discriminatory* (pay-your-bid) pricing protocol. Figure 1 shows the real *marginal* price in *Cetes* auctions, defined by the lowest accepted price in an auction and computed off the annual yield deflated by the yearly CPI inflation, for the four most common maturities of 28, 91, 182 and 364 days. We are interested in understanding how investors' characteristics (such as their risk aversion, wealth or information) and their bidding behavior determine the level and volatility of these marginal prices.

To this end, 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. Second, the share of an investor's total bids that are accepted, which we label the *in-the-money share* (ITM), 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 an investor's average price paid to the marginal price.

The combination of the last two facts is perhaps surprising. In order to have high

Figure 1: Marginal Prices of Cetes Bonds of Different Maturity



in-the-money shares, bidders must submit relatively many bids at relatively high prices. Since bids are executed at the bid price, high in-the-money shares would therefore seem to suggest that large investors should overpay relative to the marginal price 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 critical for understanding the fundamental determinants of government financing costs and economic performance more broadly (Bocola (2016) documents a link between the two).

Our approach is twofold. First, we develop a theoretical framework with a realistic discriminatory pricing protocol, rich multidimensional investor heterogeneity, and risk aversion of the constant relative risk aversion type. Then, we use this framework to assess which type of heterogeneity can qualitatively and quantitatively explain the bidding patterns we document for Mexican primary sovereign bond markets.

Investors may submit multiple bids consisting of a bid price and a commitment to buy a positive quantity of bonds at that price. Bid quantities must be weakly positive but are otherwise unrestricted. The government accepts bids in descending order of price until it obtains its required revenue (i.e., its revenue *target*). We introduce uncertainty along two dimensions: a *supply shock* to the government's revenue target, which can affect marginal prices because the government must sell more bonds to raise more revenue, and a *quality shock* that determines the bond's default risk (high or low) and alters investors' willingness to pay.

Given this protocol, we study the *Walrasian limit* with a large number of bidders who act as price-takers. This leads to the existence of a lowest-accepted price in each state of the world, with all bids at prices above this *marginal* price also executed. This allows for a clean comparison with the literature that assumes price determination in competitive markets under uniform pricing. We refer to marginal prices associated with the good quality shock as the *high price schedule*, and to those associated with the bad quality shock as the *low price schedule*.

Investors may differ in terms of their fundamental characteristics, such as wealth, risk aversion or market power, as well as in terms of their information. While all investors are uninformed about the supply shock, there are some investors who are *informed* about the realization of the quality shock, while the remainder are *uninformed*. Heterogeneous information gives rise to heterogeneous bidding strategies that can be mapped to the data.

Our first main result is that it is not possible to account for the observed differences in bidding behavior between the largest bidders and the rest by relying on fundamental heterogeneity in the absence of asymmetric information. This is because a discriminatory-price protocol induces strong positive co-movement between bid prices and in-the-money shares for any source of fundamental heterogeneity. More specifically, an investor can obtain relatively high in-the-money shares only if she offers relatively high prices, but this directly implies that she would *also* overpay more on average. Since all investors face the same price schedule when there is symmetric information, any form of heterogeneity that induces some investors to bid more at high prices then leads to counterfactual positive co-movement between overpayment and in-the-money shares.

In contrast, asymmetric information *can* account for the empirical bidding patterns. This is because the presence of informed investors induces the winner's curse for uninformed investors, whose bids at high marginal prices associated with the good quality shock will be accepted even if the quality shock is bad. If the winner's curse is severe, uninformed investors optimally choose to no longer submit bids at high marginal prices, only at low marginal prices. But given that they do not participate at high prices, they can bid at low prices *as if* they were informed because all of their bids are now accepted if and only if the realized quality shock is bad. This leads to bidding patterns where informed bidders achieve higher average in-the-money shares because they buy both good and bad bonds, but do not overpay relative to the marginal price because they can target bids to the state of the world. Moreover, uninformed investors do not tend to overpay either, but only because they opt not to participate at high marginal prices in the first place.

The winner's curse is sufficiently strong to produce these bidding patterns only if there are sufficiently many informed investors. If share of informed investors were small,

price differences across quality shocks would also be small, and uninformed investors would continue to bid at every marginal price. Hence the equilibrium bidding patterns we document *reveal* that there is substantial asymmetric information in primary government bond markets. This is perhaps surprising because it is commonly assumed that information about a country's finances is publicly known. What is the information that some investors have and some others do not? Can such information explain both bidding behavior and the aforementioned dynamic properties of marginal prices? We answer these questions by calibrating a version of the model to Cetes data exploring the nature of information that allow us to quantitatively match the observed bidding behavior as well as the dynamic behavior of prices.

A critical dynamic property of Cetes marginal prices is the relatively high unconditional volatility and low conditional volatility. To account for this, we distinguish in the model between publicly observed information, which determines a public state, and privately observed information, which leads to information heterogeneity. Public information includes standard fundamentals, such as GDP growth or inflation, or the past week's auction prices. Private information pertains to 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

To be consistent with within and across properties of prices, the calibration has to generate a large difference in bidding behavior despite low conditional price volatility. The version of our model that best achieves this result feature a small probability of an extremely bad (*black swan*) event of the sort seen in Mexico during the 1990s and 1980s. The winner's curse allows this small probability to have a relatively large impact on uninformed investors' portfolio. If, in line with the "Peso problem" literature in asset pricing, we assume that this low-probability bad event has not occurred in the sample period data (which is the case in our sample), our sparse model generates results that are very close to the data.²

¹On 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 Cuadra and Sapriza (2008).

²A 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 Rogoff (1980), Evans (1996), Geert Bekaert and Marshall (2001), and Burnside et al. (2011).

Taken together, our results show that investors' bidding behavior is qualitatively consistent with the data only if there is asymmetric information, and quantitatively consistent with the data only if there is asymmetric information about infrequent and large crises. This result has important implications for government finances, as the level of asymmetric information that is consistent with bidding patterns in our data also leads to substantially higher bond yields relative to a counterfactual with identical fundamentals but symmetric information. Specifically, we show that *symmetric information* about black swan events would lower sovereign bond yields because more investors would be willing to submit bids at high prices; while *symmetric ignorance* would also reduce yields because investors do not face the winner's curse and enjoy better risk sharing.

Related Literature: Our paper fills an important gap in the sovereign debt literature, as most papers study sovereign default as the outcome of governments' strategic choice, but use a parsimonious model of investor optimization (see, for instance, Mendoza and Yue (2012), Chatterjee and Eyigungor (2012), Hatchondo, Martinez, and Sosa-Padilla (2016)). Consistent with this, the macroeconomic sovereign debt literature focuses on competitive pricing rules under which risk-adjusted bond yields are equal to the risk-free rate. In this literature there has been some attention to the impact of the timing of decisions and of debt maturity in sovereign markets (see Aguiar et al. (2019)). However, the actual mechanics of how sovereign bonds are sold in practice through auctions and their impact on observed prices has been largely ignored. 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.³ Our paper argues that the neglected roles of auction mechanics and information heterogeneity, and their interaction, drives primary market prices when investors are risk averse. The nature of the information shocks we consider (public and private, heterogeneous and common) is also consistent with the rich literature on rare disasters and the "peso problem".⁴

Methodologically, we circumvent some of the challenges that standard auction models face in accommodating asymmetric information.⁵ In this context, our framework can be

³See for example Eaton and Gersovitz (1981), the review articles by Aguiar and Amador (2013) and Aguiar et al. (2016), and the recent quantitative literature by Aguiar and Gopinath (2006), Arellano (2008), Chatterjee and Eyigungor (2012), Bocola and Dovis (2019).

⁴Disaster risk has been argued to play a large role in both asset pricing and macroeconomic fluctuations. See for example Chatterjee and Corbae (2007), Barro and Ursúa (2012) and Gourio (2012).

⁵For a discussion see, for instance, Biais, Bossaerts, and Rochet (2002) 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 Manzano and Vives (2020), who study a divisible good uniform-price auction with risk neutral bidders with asymmetric dispersed information in linear strategies.

viewed as an 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 its supply (or, equivalently, demand). 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.⁶ 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 Boyarchenko, Lucca, and Veldkamp (2017)).

Section 2 describes our novel data set and establishes basic facts about bidding and prices. In Section 3 we develop a tractable primary market model with a discriminatory pricing protocol and rich bidder heterogeneity. In Section 4 we show that only information asymmetry can qualitatively accommodate the empirical bidding patterns we document. We also conduct a detailed analysis of yields under various levels of asymmetric information, which will become useful for conducting counterfactuals. In Section 5 we calibrate the model to match relevant moments on price dynamics and bidding behavior in order to understand the quantitative nature and extent of asymmetric information for the Mexican case. In this section we also empirically validate further testable implications of the model and study counterfactual information environments. Section 6 concludes.

2 Institutional Background and Data

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 the main source of federal government funding since 1978. Since its inception, the primary market for Cetes has consisted of public auctions, with the pricing protocols alternating between uniform and discriminatory.

Our data comes 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, which has three key advantages. First, Cetes were regularly auctioned using a discriminatory-price protocol throughout this period (a switch to a uniform-price protocol occurred on October 5, 2017). Second, Mexico experienced relatively stable inflation and did not suffer any major crisis during the period. Hence Cetes auctions took place on a consistent weekly schedule for

⁶Recent auction literature shows that price-taking arises as the number of bidders get large. An example is Fudenberg, Mobius, and Szeidl (2007), 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 Reny and Perry (2006) who show a similar result when bidders have affiliated values and prices are on a fine grid.

all maturities.⁷ Finally, we were able to gather data on all bids submitted (not just those that were executed) and all bidders.

We observe a total of 2,717 Cetes auctions. Across maturities, we observe an average of 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

2.1 Price Dynamics

As a first step, we establish a number of basic facts about Cetes price dynamics. 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. (We return to this point when we discuss Peso problems).

Table 4 shows the time-series moments of marginal prices MP by maturity, defined by the minimum bid accepted in each auction. The autocorrelation between marginal prices at subsequent auctions is a measure of conditional uncertainty, since it partly determines the predictability of future prices from public information.

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

⁷All maturities except 364 days are auctioned weekly, 364 day bonds are auctioned monthly.

Our results suggest 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 the conditional uncertainty, we run a regression of marginal prices 28-day bonds on a constant and one lag,

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

We estimate $\beta_1 = 0.98$, which implies that lagged prices are very informative.⁸ This is not surprising because auctions take place every week. We also estimate $R^2 = 0.97$, which implies that 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.

2.2 Bidding Patterns

Next we establish basic facts about bidding patterns. One drawback of our data is that the numeric bidder identifier is auction-specific. This means that we cannot track bidders across auctions. To uncover heterogeneous bidding behavior, we compare the bidding behavior of the largest bidder at an auction to all other bidders, where the largest bidder is the bidder who buys the most bonds in an auction.⁹

We make this distinction anticipating potential sources of heterogeneity across bidders. 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/or (iv) the one with the strongest incentives (or lowest costs) to become informed. In our model, we allow for all of these factors as potential sources of heterogeneity.


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

What accounts for differences in in-the-money shares across bidders? One possibility

⁸As can be guessed from Figure 1 and Table 4, the result for other maturities is very similar.

⁹Our results are robust to considering the Top 2 or Top 3 bidders

Figure 2: In the Money Shares Largest vs. Rest




is that large bidders systematically offer higher prices than other investors. 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. It is clear that the distribution of overpayment is very similar for large bidders and other bidders.

This combination of high in-the-money shares for the highest bidder without concomitant differences in overpayment is surprising: in a pay-your bid protocol, bids are accepted relatively frequently only if they are submitted at relatively high prices. In the next section we construct a model that can replicate these facts by relying on asymmetric information about default risk. The basic mechanism is that informed investors can more effectively tailor bids to the prevailing marginal price than uninformed investors

3 A Model of the Primary Debt Market

We now construct a model of primary sovereign debt markets with bidder heterogeneity in wealth, risk aversion, market power, and information. Our baseline model is static and has two dates; in Section 5 we consider a repeated version to incorporate time series information on prices. Our key theoretical result is that asymmetric information is the only form of heterogeneity that can rationalize the bidding patterns previously discussed.

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



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' objective is to maximize expected utility over consumption at the end of the period given a strictly concave flow utility function that satisfies the Inada conditions. Each investor has wealth W in period one and cannot borrow. Investors invest their wealth in either a risk-free bond (storage) or a risky bond offered by the government.

The government is modeled mechanically: it needs to raise a certain number of units of the numeraire at date 1 by selling multiple units of a bond that promises repayment at date 2. Without loss of generality and like *Cetes*, bonds are zero-coupon and pure discount, offering a claim to one unit of the numeraire at date 2.

Bonds are risky in that the government may default on its promises. If the government defaults, investors cannot recover any of their investment. Its default probability κ_θ is random and determined by an exogenous state of the world $\theta \in \{g, b\}$, with $\kappa_g < \kappa_b$. The ex-ante probability of each state is $f(g)$ and $f(b)$ respectively, with $f(g) + f(b) = 1$; the unconditional default probability is

$$\bar{\kappa} = f(g)\kappa_g + f(b)\kappa_b.$$

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 can 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 values of κ_θ .

To generate variation in in-the-money shares, we need an additional source of uncertainty. Hence we introduce a *supply shock* ψ to the government financing needs, where ψD is the total revenue that the government must raise (for example, because it needs to roll over existing obligations) and D is a constant. We assume that the supply shock ψ is discrete and lives on an arbitrarily fine discrete grid $\mathcal{H} \equiv \{\psi_0, \dots, \psi_M\}$ with length M . We index the supply shock by $k \in \{0, M\}$. Without loss of generality, let ψ_k be strictly increasing in k and denote the probability of ψ_k by $h(\psi_k)$.¹⁰

We refer to $s = (\theta, \psi)$ 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 ψ , which precludes signaling by the government.

3.2 Investors

Investors may differ in their *fundamental type* and their *information type*. The fundamental type $j \in \{1, 2\}$ indexes the investor's utility function U_j and initial wealth W_j , where wealth differences will also allow us to capture differences in market power. For simplicity, we only consider two fundamental types of equal mass at a time; thus we will either assume investors have the same preferences but differ in their wealth, or that they have the same wealth but differ in their preferences.

The information type $i \in \{I, U\}$ determines whether an investor is *informed* about the quality shock and knows its realization (denoted by θ^*), or *uninformed* about the quality shock and unaware of its realization. Both information types face some residual uncertainty because of the supply shock ψ . We summarize this uncertainty by defining type-specific set of plausible states \mathcal{F}_j^i which collect all states that type i believes may feasibly occur. Given the two sources of shocks, we have $\mathcal{F}_j^I = \theta^* \times \mathcal{H}$ (the supply shock is the only source of uncertainty for informed investors) and $\mathcal{F}_j^U = \mathcal{S}$ (all states are plausible for

¹⁰An alternative interpretation is that there is a *demand shock* which leads to a random share η of investors not participating in buying government bonds (because of liquidity shocks or access to more favorable investment opportunities). These demand shocks could be thought of as a *correlated private value shock*, while the quality shock θ is a *common value shock*. Supply and demand shocks are largely isomorphic to each other when taking $\psi = 1/(1-\eta)$. However demand shocks somewhat complicate the analysis because investors might update their beliefs with respect to the size of the demand shock once they know whether that they were not personally affected. In the interest of parsimony, we therefore use supply shocks.

uninformed investors). The share of informed investors is $n \in [0, 1]$, and we assume that there is no correlation between fundamental and information types.

3.3 Pricing Protocol and Strategies

As in practice, the government sells bonds via a pay-your-bid auction protocol. 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 (they are *in the money*), all bids below are rejected (they are *out of the money*).¹¹ The set of marginal prices associated with the good quality shock $P(g, \cdot)$ is the *high price schedule*, the set marginal prices associated with the bad quality shock $P(b, \cdot)$ is the *low price schedule*. The state-contingent yield is

$$y(s) = \frac{1 - P(s)}{P(s)}.$$

Bonds pay one unit of the numeraire after repayment and zero after default, hence the range of possible prices is $[0, 1]$. A *bidding strategy* maps any price in $[0, 1]$ into a weakly positive bid quantity. Since investors have rational expectations with respect to the set of possible marginal prices, 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.¹² Since there is a single marginal price associated with each state, we can then equivalently define bidding strategies as functions that maps sets of states into weakly positive bid quantities at each state-contingent marginal price.

Definition 1. Let $P(s)$ denote the marginal price in state s . A *bidding strategy* for an investor of information type i and fundamental type j is a function \tilde{B}_j^i that maps every state in the investor's

¹¹If there is excess demand at the marginal price, the government is assumed to ration pro-rata. While this does not occur in the equilibrium of our model, there is some rationing in our data 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 any insights, we follow the literature that assumes the set of possible bid prices is in a continuum, as in the seminal work of Wilson (1979). For a treatment of bidders restricted bids at discrete points in a uniform-price auction of a perfectly divisible good see Kastl (2011).

¹²Observe that if two states have the same marginal price, bids associated with either state 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.

set of feasible states \mathcal{F}_j^i into bids consisting of quantity $B_j^i(s) \geq 0$ and bid price $P(s)$.

Note that, while informed investors need not choose bids for any states of the world they know will not occur (i.e, states that are not associated with the realized quality shock), it is convenient to assume that they follow a two step procedure: they first choose bids for all possible states \mathcal{S} , and then discard any bids that are not associated with the realized quality shock. This procedure allows state-contingency in bids while ensuring that bidding strategies are well-defined for all possible states. We follow this convention going forward.

3.4 Decision problems and equilibrium definition

To find an optimal bidding strategy, investors must forecast which bids will be accepted in which states of the world. It is useful to conceptualize the trading protocol as consisting of two steps: First, investors choose a bidding strategy. Second, the state of the world is realized and the government chooses which bids to accept.

The outcome of this procedure is conveniently summarized using *sets of executed bids* $\mathcal{E}_j^i(s)$ which collects all bids by an investor of type $\{i, j\}$ that are executed in state s . Since each bid is associated with a state-specific marginal price, the elements of executed bid sets are states. For uninformed investors, the executed bid set includes all states with marginal prices above the realized marginal price. For informed investors, states must correspond to the realized quality shock *and* the state-specific marginal price must be higher than the realized marginal price. This implies

$$\mathcal{E}_j^U(s) \equiv \{\tilde{s} : P(\tilde{s}) \geq P(s)\} \quad \text{and} \quad \mathcal{E}_j^I(s) \equiv \{\tilde{s} : P(\tilde{s}) \geq P(s) \text{ and } \tilde{\theta} = \theta\}.$$

The total quantity of bonds purchased in state s by an investor of information type i and fundamental type j is then equal to

$$\mathcal{B}_j^i(s) = \sum_{\tilde{s} \in \mathcal{E}_j^i(s)} B_j^i(\tilde{s}), \tag{2}$$

and the investor's *expenditures* on bonds is

$$X_j^i(s) = \sum_{\tilde{s} \in \mathcal{E}_j^i(s)} P(\tilde{s}) B_j^i(\tilde{s}). \tag{3}$$

Investment in the risk-free asset in state s is then determined as the residual,

$$w_j^i(s) = W_j - X_j^i(s). \quad (4)$$

The objective function of an investor is the conditional expectation of utility after default and repayment, given the investor's information set \mathcal{F}_j^i ,

$$V_j^i(\tilde{B}_j^i) = \mathbb{E}_{\theta, \psi} \left[\kappa_\theta U_j(w_j^i(s)) + (1 - \kappa_\theta) U_j(w_j^i(s) + \mathcal{B}_j^i(s)) \middle| \mathcal{F}_j^i \right], \quad (5)$$

and the associated *decision problem* is

$$\max_{\tilde{B}_j^i} V_j^i(\tilde{B}_j^i) \quad \text{s.t.} \quad \mathcal{B}_j^i(s) \geq 0; w_j^i(s) \geq 0 \text{ for all } s \in \mathcal{S}. \quad (6)$$

where the constraints are the short-sale constraint on bids and the borrowing constraint.

The *market clearing condition* such that the government raises revenue ψD in state $s = (\theta, \psi)$ given share n of informed investors and share $\frac{1}{2}$ of each fundamental type j is

$$\psi D = \frac{1}{2} \sum_{j \in \{1, 2\}} \left(n X_j^I(s) + (1 - n) X_j^U(s) \right). \quad (7)$$

We are now ready to state our equilibrium definition.

Definition 2 (Primary Market Equilibrium). *A primary market equilibrium consists of a price schedule $P : \mathcal{S} \rightarrow [0, 1]$, and bidding strategies $\tilde{B}_j^i : \mathcal{F}_j^i \rightarrow \mathbb{R}_+^{\mathcal{F}_j^i}$ for all i and j such that*

1. *Bidding strategies solve decision problem (6) for all types.*
2. *The market clearing condition (7) is satisfied for all $s \in \mathcal{S}$.*

3.5 Mapping Bids to the Data

Given an equilibrium, we can map investors' bidding strategies into model counterparts of the data moments discussed in Section 2. This facilitates our empirical analysis.

The *in-the-money share* in state s is the ratio of accepted bids to submitted bids,

$$ITM_j^i(s) = \frac{\sum_{\tilde{s} \in \mathcal{E}_j^i(s)} B_j^i(\tilde{s})}{\sum_{\tilde{s} \in \mathcal{F}_j^i} B_j^i(\tilde{s})}. \quad (8)$$

The bid-weighted *average price* paid in state s is

$$AP_j^i(s) = \sum_{s' \in \mathcal{E}_j^i(s)} \frac{P(s')B_j^i(s')}{\sum_{\tilde{s} \in \mathcal{E}_j^i} B_j^i(\tilde{s})}. \quad (9)$$

Overpayment $\Omega_j^i(s)$ is the ratio of the average price to marginal price

$$\Omega_j^i(s) = \frac{AP_j^i(s)}{P(s)}. \quad (10)$$

These quantities are observable since our data set contains submitted bids, executed bids, and all associated prices for each investor at a given auction.

4 Equilibrium Characterization

4.1 Optimal bids and equilibrium prices

We begin by characterizing optimal bidding strategies and mapping these into equilibrium prices. Two preliminary observations are useful. Since preferences satisfy Inada conditions, every investor must always invest some amount of money in the risk-free asset; hence, the borrowing constraint never binds. Second, the non-negativity constraint on bids may bind for some investors in some states of the world; we use $\chi_j^i(s)$ to denote the associated Lagrange multiplier.

Formulating a bidding strategy requires forming expectations about the states of the world in which a given bid will be accepted. Hence we define *acceptance sets* $\mathcal{A}_j^i(s)$ that collect all states in which a bid at a given marginal price $P(s)$ is accepted. For uninformed investors, the pay-your-bid protocol implies that a particular bid is accepted in all states with lower marginal prices; for informed investors a bid is accepted in all states associated with the realized quality shock which have a lower marginal price. That is,

$$\mathcal{A}_j^U(s) = \{\tilde{s} : P(\tilde{s}) \leq P(s)\} \quad \text{and} \quad \mathcal{A}_j^I(s) = \{\tilde{s} : P(\tilde{s}) \leq P(s) \text{ and } \tilde{\theta} = \theta\}.$$

Acceptance sets are complements of executed bid sets. The former collect all states with marginal prices that are *lower* than the bid price, the latter cumulate all bids that were submitted at *higher* prices and are thus executed in the current state.

Now consider the decision of an informed investor of some fundamental type j . From the investor's decision problem, a necessary and sufficient condition for optimality of the

bid at marginal price $P(s^*)$ associated with state $s^* = (\theta^*, \psi^*)$ is

$$\mathbb{E}_\psi \left[\sum_{s \in \mathcal{A}_j^I(s^*)} -U'_j(w_j^I(s))\kappa_{\theta^*}P(s^*) + U'_j(w_j^I(s) + \mathcal{B}_j^I(s))(1 - \kappa_{\theta^*})(1 - P(s^*)) \right] - \chi_j^I(s^*) = 0. \quad (11)$$

Since informed investors know the quality shock, expectations are taken only with respect to the supply shock ψ . The summation in square brackets is the difference between marginal utility after default and after repayment, weighted by the marginal price and default probability in all states in which a bid at $P(s^*)$ is accepted. While a bid at $P(s^*)$ is always executed at the bid price, consumption levels are indexed by s rather than s^* because the *realized* portfolio is state-contingent.

Uninformed investors face the same basic trade-off as informed investors. The key difference is that they face uncertainty about default probabilities and the sets of states in which a given bid is accepted. This leads to the concern that bids on the high price schedule are also accepted when the realized quality shock is bad. This is captured by differences in acceptance sets $\mathcal{A}_j^U(s)$ given a set of bids. Hence for an uninformed investor of fundamental type j , the necessary and sufficient condition for bids at marginal price $P(s^*)$ can be similarly stated as

$$\mathbb{E}_{\theta, \psi} \left[\sum_{s \in \mathcal{A}_j^U(s^*)} -U'_j(w_j^U(s))\kappa_\theta P(s^*) + U'_j(w_j^U(s) + \mathcal{B}_j^I(s))(1 - \kappa_\theta)(1 - P(s^*)) \right] - \chi_j^U(s^*) = 0. \quad (12)$$

where expectations are taken with respect to both the quality shock and the supply shock.

We can summarize these optimality conditions by defining investor $\{i, j\}$'s expected marginal rate of substitution across default and repayment when bidding at marginal price $P(s^*)$ associated with state s^* given type-specific expectation operator \mathbb{E}^i as

$$M_j^i(s^*) = \frac{\mathbb{E}_{\theta, \psi}^i \sum_{s \in \mathcal{A}_j^i(s^*)} \kappa(s) U'_j(w_j^i(s))}{\mathbb{E}_{\theta, \psi}^i \sum_{s \in \mathcal{A}_j^i(s^*)} (1 - \kappa(s)) U'_j(w_j^i(s) + \mathcal{B}_j^i(s))} \quad \text{for all } s^* \in \mathcal{S}.$$

The numerator is the expected marginal utility of consumption after a default; the denominator is the expected marginal utility of consumption in case of repayment.

This leads to the following characterization of equilibrium prices and portfolios, which shows that bond yields satisfy asset pricing relationships that are sensitivity to the presence of informed investors.

Proposition 1. *Equilibrium prices satisfy the following conditions:*

- (i) *The marginal investor in any state is the investor with the lowest expected marginal rate of*

substitution, and yields are equal to the marginal investor's marginal rate of substitution:

$$\frac{1 - P(s)}{P(s)} = \min_{i,j} M_j^i(s) \quad \text{for all } s \in \mathcal{S}. \quad (13)$$

- (ii) Prices are strictly ordered by the supply shock: $P(\theta, \psi)$ is strictly decreasing in ψ given θ .
- (iii) If there are no informed investors, there is a single price schedule: $P(g, \psi) = P(b, \psi)$ for all ψ . If there are informed investors, then $P(g, \psi) \geq P(b, \psi)$ and strictly for at least one ψ .
- (iv) The high and low price schedules converge for all interior demand shocks as $n \rightarrow 0$. That is, $\lim_{n \rightarrow 0} P(g, \psi) = P(b, \psi)$ for all $\psi_0 < \psi < \psi_M$. Hence the winner's curse disappears if the share of informed investors is sufficiently small and $h(\psi_0)$ and $h(\psi_M)$ are sufficiently low.

Proof. We prove each statement in turn:

- (i) The marginal investor is the investor with the highest marginal willingness to pay. This is equivalent to offering the lowest marginal yield. Market clearing requires that the marginal investor's short-sale constraint does not bind. Pricing equation (13) then follows directly from first-order conditions (11) and (12).
- (ii) The government can raise more revenue only if at least one type of agent consumes less after a default. Since agents have CRRA preferences, marginal utility is convex. Hence the required risk premium must increase in ψ , and prices must fall.
- (iii) The first part follows from the fact that uninformed bids are not state contingent. Since $\kappa(g) < \kappa(b)$ and informed investors can submit state-contingent bids, it is clear that there does not exist an equilibrium where $P(g, \psi) < P(b, \psi)$ for some ψ if $n > 0$. We now show that we must have $P(g, \psi) > P(b, \psi)$ for at least one ψ . Since there exist an equal share of informed investors of every fundamental type, at least one type of informed investor must be marginal in every (g, ψ) . Since the statement follows trivially if informed investors do not bid in some state (b, ψ) , assume that optimal bids are interior in every state. Statements (i) and (ii) then imply that optimal bids in state $s = (\theta, \psi)$ satisfy optimality condition

$$\frac{\kappa(\theta)}{1 - \kappa(\theta)} \frac{u'(W - \sum_{\psi' \leq \psi} P(\theta, \psi') B^I(\theta, \psi'))}{u'(W - \sum_{\psi' \leq \psi} P(\theta, \psi') B^I(\theta, \psi)) + \sum_{\psi' \leq \psi} B^I(\theta, \psi')} = \frac{1 - P(\theta, \psi)}{P(\theta, \psi)}.$$

Now suppose for a contradiction that $P(g, \psi) = P(b, \psi)$ for all ψ . Since uninformed bids are not contingent on the state, market clearing then implies that $X^I(g, \psi) =$

$X^U(b, \psi)$ for all ψ . Given $P(g, \psi) = P(b, \psi)$, then $B^I(g, \psi) = B^U(b, \psi)$ for all ψ . Since $\kappa(g) < \kappa(b)$, the optimality condition must therefore be violated for at least one state.

(iv) By market clearing (7), $\lim_{n \rightarrow 0} \sum_{j \in (1,2)} \frac{1}{2} X_j^U(\theta, \psi) \rightarrow D\psi$ for all ψ . As uninformed bid unconditionally of θ , $\lim_{n \rightarrow 0} \sum_{j \in (1,2)} \frac{1}{2} X_j^U(g, \psi) \rightarrow \sum_{j \in (1,2)} \frac{1}{2} X_j^U(b, \psi)$ and then $\lim_{n \rightarrow 0} P(g, \psi) \rightarrow P(b, \psi)$. As $P(g, \psi)$ is strictly decreasing in ψ given θ and n , when $n \rightarrow 0$, prices must be sorted by ψ . That is, there is always a ϵ small enough such that for $\psi' - \psi = \epsilon$, i.e $P(\theta, \psi) < P(\theta', \psi') < P(\theta, \psi')$. This proof does not apply at extreme values of ψ , and then convergence will not happen at $\psi = \psi_0$ and $\psi = \psi_M$.

□

The first statement shows that price determination works similarly to a canonical asset pricing framework. In every state, bonds are priced by the covariance of payoffs with marginal utility of the investor with the highest marginal willingness to pay. The key difference is the pricing protocol, whereby bids are executed at the bid prices whenever they exceed the marginal price. This leads to rich interactions in optimal bids in different states of the world, and implies that bids at high prices affect willingness to pay in all states with lower marginal prices. The second statement shows that we can nevertheless partially order prices by the supply shock. This is because investors must carry more exposure to default risk when bond supply is high, which raises the required risk premium. The third statement shows that the presence of informed investors drives a wedge between the high and low price schedules because informed investors bid more aggressively when the default probability is low. This immediately implies that uninformed investors face the winner's curse whenever there are informed investors because uninformed bids on the high quality schedule are also accepted when the bad quality shock realizes. This discourages uninformed investors from bidding at high prices, and has implications for risk sharing, average yields and equilibrium bidding patterns. The fourth statement shows that such wedge disappears as informed investors leave the market.

While the general model is complex, simple examples provide insights into the basic mechanics of bidding behavior, and motivates our joint analysis of heterogeneous bidders and asymmetric information.

Example 1 (Homogeneous investors with complete information). *Suppose all investors are ex-ante symmetric, know the realized supply shock ψ and have common expectations about the quality shock, with κ denoting the expected default probability (if all investors are informed, then $\kappa = \kappa_\theta$; if they are uninformed then $\kappa = \bar{\kappa}$). Given these assumptions, simplify notation to $P(\psi)$ for prices and $B(\psi)$ for bond purchases. For any ψ , the market-clearing condition is $B(\psi)P(\psi) =$*

ψD and the first-order condition for optimal bids is

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

The bond price in state ψ is

$$P(\psi) = 1 - \kappa - \kappa \frac{\psi D}{W - \psi D}. \quad (14)$$

the in-the-money share is one, the average price is the marginal price, and no investor overpays

$$ITM(\psi) = 1, \quad AP(\psi) = P(\psi) \quad \text{and} \quad \Omega(\psi) = 1.$$

In the example, investors charge a risk premium that depends on the ratio between the per-capita debt level ψD and wealth W and is proportional to the default probability κ . Since every investor can correctly forecast the marginal price in every state, all bids are in-the-money in every state, and every investor pays the same price.

The next example shows that supply uncertainty can create variation in in-the-money shares and overpayment in some states of the world, but that these two outcomes will be positively correlated. This is at odds with the data.

Example 2 (Homogeneous investors with supply uncertainty). *Continue to assume that all investors are ex-ante symmetric and have common expectations about the quality shock, but are uncertain about the supply shock, which has two possible realizations, $\psi \in \{\psi_1, \psi_2\}$ with $\psi_1 < \psi_2$. We index states by 1 and 2 in parentheses, respectively.*

Since investors are risk averse, marginal prices must fall when fewer investors participate in the auction. Hence $P(1) > P(2)$, and the market-clearing conditions in the two states are

$$P(1)B(1) = \psi_1 D \quad \text{and} \quad P(1)B(1) + P(2)B(2) = \psi_2 D.$$

Combining both expressions implies that expenditures at marginal price $P(2)$ must exactly offset the incremental supply of bonds that is not already purchased at marginal price $P(1)$,

$$P(2)B(2) = (\psi_2 - \psi_1)D.$$

Now consider in-the-money shares and prices paid. In state 1, the government rejects bids at $P(2)$ and all accepted bids are executed at marginal price $P(1)$. In state 1 we then have

$$ITM(1) = \frac{B(1)}{B(1) + B(2)} < 1 \quad AP(1) = P(1) \quad \Omega(1) = \frac{P(1)}{P(1)} = 1.$$

In state 2, all bids are accepted but some are executed at above marginal price $P(2)$. We have

$$ITM(2) = 1, \quad AP(2) = \frac{B(1)P(1) + B(2)P(2)}{B(1) + B(2)} > P(2) \quad \Omega(2) = \frac{B(1)\frac{P(1)}{P(2)} + B(2)}{B(1) + B(2)} > 1.$$

Investors choose symmetric strategies, and ITM and overpayment are positively correlated.

The positive correlation between in-the-money shares and overpayment in the example is a direct implication of the pay-your-bid protocol, but is inconsistent with the empirical facts. This motivates our analysis of investor heterogeneity.¹³

4.2 Using heterogeneity to account for the empirical bidding patterns

In Section 2 we documented two key empirical facts: on average, the largest bidder at an auction has a higher in-the-money share (submits higher bids in average), but does not overpay (does not pay more in average). We have shown with the two previous examples that these facts cannot be accounted for without investor heterogeneity. We now investigate which forms of investor heterogeneity can do so.

4.2.1 Fundamental heterogeneity

We now consider the effects of fundamental heterogeneity due to, for example, wealth, risk aversion, or market power, while assuming that all agents have symmetric information. We first show a useful property of optimal bids under CRRA utility, which is that statistics based on *ratios* of bids, such as in-the-money shares or average prices paid, are invariant to wealth heterogeneity. (The result readily extends to CARA utility, but we do not consider such preferences here.) An immediate upshot is that wealth heterogeneity cannot account for the empirical facts we document. The result also informs our calibration strategy below.

Proposition 2 (Wealth neutrality). *Suppose that all agents are symmetrically informed, have common CRRA preferences, and differ only in their type-specific wealth W_j . Then optimal bidding strategies satisfy the decomposition $B_j(s) = F(W_j)\beta(s)$ for all j , where $\beta(s)$ and $F(\cdot)$ are independent of j . Hence in-the-money shares and average prices paid are the same for all types.*

Proof. The multiplicative separability of the optimal policy function is a standard property of CRRA utility functions. The remaining statements then follow immediately. \square

¹³In Appendix A, we also provide a numerical example along the lines of Examples 1 and 2 (with parameters in Definition 3 below), and show the limited role of investors' risk aversion on bidding behavior.

We now turn to other potential forms of fundamental heterogeneity. Our key finding is that the pay-your-bid protocol leads to a counterfactual positive link between in-the-money shares, average prices paid and overpayment for any form of fundamental heterogeneity. We restrict attention to the analytically tractable case with two possible supply shocks, as in Example 2. This allows for a transparent discussion of the key economic forces. We verify the robustness of these forces in our quantitative analysis.

Proposition 3 (Counterfactual implications of fundamental heterogeneity). *Let $\psi \in \{\psi_1, \psi_2\}$. If all investors are symmetrically informed, any form of fundamental heterogeneity generates positive correlation between in-the-money shares and overpayment across investors. Hence fundamental heterogeneity of any form cannot generate the facts documented in Section 2.*

Proof. Under symmetric information, there is a unique price schedule with two marginal prices $P(1)$ and $P(2) < P(1)$. Define the *ratio* of type j 's bids at the marginal prices as

$$\rho_j = B_j(1)/B_j(2).$$

Following Example 2, we can then derive in-the-money shares and overpayment for each type j in each state as monotone functions of this ratio only. For state 1, we obtain

$$ITM_j(1) = \frac{\rho_j}{1 + \rho_j} \quad \text{and} \quad \Omega_j(1) = 1,$$

where $ITM_j(1)$ is strictly increasing in ρ_j . In state 2, we have

$$ITM_j(2) = 1 \quad \text{and} \quad \Omega_j(2) = \frac{1 + \rho_j(P(1)/P(2))}{1 + \rho_j},$$

where $\Omega_j(2)$ is strictly increasing in ρ_j since $P(1) > P(2)$. Hence, differences in-the-money shares and overpayment across investors must be positively correlated on average. Formally, for any types j and j' distinguished by any fundamental heterogeneity,

$$\mathbb{E}[ITM_j(\psi)] > \mathbb{E}[ITM_{j'}(\psi)] \Leftrightarrow \mathbb{E}[\Omega_j(\psi)] > \mathbb{E}[\Omega_{j'}(\psi)].$$

□

In a pay-your-bid protocol with symmetric information, an investor can have a relatively large share of bids her accepted only if she submits a relatively large share of her bids at high marginal prices. But if this implies that she will also overpay more when the marginal price turns out to be low. Hence in-the-money shares and overpayment are both governed by the ratio of bids ρ_j , and *any* form of fundamental heterogeneity that

leads to differences in in-the-money shares must *necessarily* lead to co-movement with overpayment. As we will show in the next section, this link breaks down when investors have heterogeneous information.

4.2.2 Informational heterogeneity

We now show that heterogeneous information can break the positive association between in-the-money shares and overpayment that arises under fundamental heterogeneity. This allows us to match the empirical lack of correlation between in-the-money shares and overpayment we have documented in Section 2. We establish this result by assuming that there is no fundamental heterogeneity.

The key insight is that asymmetric information leads to distinct quality-contingent price schedules. This allows informed investors to *target* bids to the realized quality shock while uninformed investors face the winner's curse.

The following proposition formalizes this possibility result by constructing in closed form an example without supply shocks. The construction clarifies that heterogeneous information is *sufficient* to break the positive association between in-the-money shares and overpayment. However, the lack of supply shocks leads to counterfactual predictions for informed investors' in-the-money shares. Hence we require supply shocks *and* heterogeneous information to jointly account for all empirical facts.

Proposition 4. *Even absent supply shocks, heterogeneous information can break the positive association between the differences in investor's in-the-money shares and overpayment.*

Proof. We construct an example that satisfies the stated conditions. Let investors have log preferences. Assume no supply shocks, $\psi_M \rightarrow \psi_1 = 1$, and index bids and prices by quality shock θ only.

Now construct an equilibrium where informed investors bid at both prices while uninformed investors submit bids only at $P(b)$. Since the winner's curse applies only to bids at $P(g)$ and $B^U(g) = 0$, uninformed investors submit the same bids on the low price schedule as informed investors, $B^U(b) = B^I(b)$. Hence market-clearing in each state is $nB^I(g) = D$ and $nB^I(b) + (1 - n)B^U(b) = B^I(b) = D$, where

$$B^I(\theta) = \frac{(1 - \kappa(\theta) - P_1(\theta))}{P(\theta)(1 - P(\theta))}.$$

Now observe that the winner's curse is increasing in n because equilibrium prices are

$$P(g) = 1 - \frac{\kappa(g)nW}{nW - D} \quad \text{and} \quad P(b) = 1 - \frac{\kappa(b)W}{W - D}.$$

Now we compute in-the-money shares and overpayment. Informed investors submit state-contingent bids at the marginal price in every state, and all of their bids are always accepted. Uninformed investors submit bids at the low price in every state, but these bids are executed only if $\theta = b$. Hence they have strictly lower average in-the-money shares than informed investors. Like informed investors, however, uninformed investors never overpay if their bids are accepted.

A sufficient condition for the optimality of $B^U(g) = 0$ is that uninformed bids at $P(g)$ earn negative expected returns if the government is expected to default with probability $\bar{\kappa}$. This is the relevant measure of default risk because any uninformed bids at $P(g)$ would be executed in every state. This is the case if

$$1 - \bar{\kappa} < P(g) \quad \Leftrightarrow \quad \bar{\kappa} > \kappa(g) \frac{nW}{nW - D}.$$

This condition is satisfied if $\kappa(b)$ is large relative to $\kappa(g)$, n and $f(b)$ are sufficiently large, and D/W is small. These restrictions ensure that the winner's curse $P(g) - P(b)$ is severe. \square

The key insight is that uninformed investors may stop bidding on the high quality schedule if the winner's curse is sufficiently severe, in which case they can submit bids on the low price schedule *as if* they are informed because they can forecast with certainty that all bid on the low price schedule are accepted if and only if $\theta = b$. Asymmetric participation given $\theta = g$ and symmetric bidding strategies given $\theta = b$ lead to differences in in-the-money shares without differences in overpayment conditional on bid execution.


The key counterfactual implications of the constructed example that in-the-money shares are always equal to one for some investors, and that overpayment is always equal to one for all investors. We now use the following numerical example to show that the combination of supply shocks and information heterogeneity can qualitatively account for the empirical bidding patterns. Parameters are chosen for ease of exposition; we calibrate parameters in order to quantitatively match the data in Section 5.

Definition 3 (Numerical Example). *Preferences satisfy log utility, investor wealth is $W = 250$, the debt level is $D = 50$, the probability of the good state is $f(g) = 0.5$, and state-contingent default probabilities satisfy $\kappa_g = 0.15$ and $\kappa_b = 0.35$. Assume that ψ is uniformly distributed on a grid between 1 and 1.16.*

Figure 4 shows equilibrium prices as a function of supply shock ψ for various shares of informed investors n . The high price schedule is shown in red, and the low price schedule in blue. Dashed lines provide the benchmark where *all* investors are informed ($n = 1$)

while black lines with circle markers show the counterfactual uninformed benchmark with symmetric ignorance ($n = 0$).

Figure 4: DP auction equilibrium as n falls.



The winner's curse also generates rich comparative statics of price schedules as we vary n . When $n = 0.6$, only informed investors participate at high prices. Since these investors must bear substantial risk per-capita, the high price schedule is relatively low in order to provide a sufficient risk premium. At the same time, the low price schedule is locally independent of n because uninformed investors can submit bids at low prices *as if* they were informed.

If n is equal to 0.4, the required risk premium is even higher because there are now even fewer informed investors who share the default risk in the high state. This leads to

a decline in high price schedule that weakens the winner's curse. This in turn induces uninformed investors to begin bidding 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, and the low price schedule *rises*.

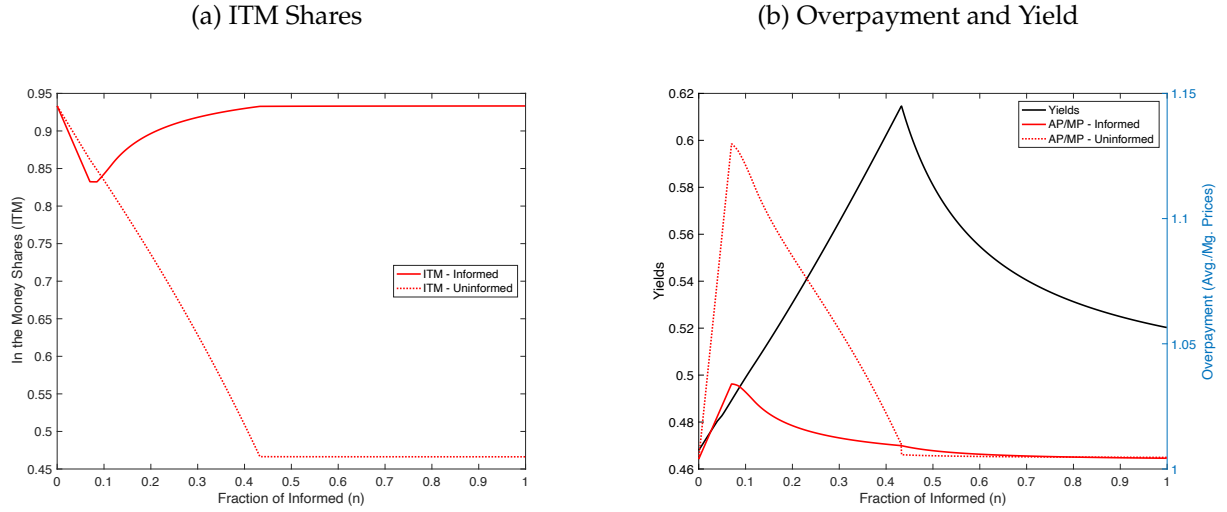
The winner's curse effect continues to operate as n decreases to $n = 0.1$. Because the per-capita bids of the uninformed remain below those of the informed on the high price schedule, reductions in n continue to further concentrate default risk in informed portfolios. This forces a large fraction of the high price schedule to drop below the *uninformed* price schedule (the price schedule that would obtain if no investor were informed). Hence presence of informed investors may lead to lower prices *even where there is good news* compared the case of symmetric ignorance. Note also that overpayment now diverges sharply because uninformed bids submit bids on both schedules, and bids on the high price schedule are executed even if the marginal price is low. Finally, when n is very small (around $n = 0.02$), price schedules start overlapping in the sense that there are prices which are marginal for either a good quality shock and a high realization of the supply shock or a bad quality shock and a low realization of the supply shock. Uninformed investors are now willing to participate fully on both schedules and prices converge to the uninformed price schedule as $n \rightarrow 0$.

Figure 5 plots in red overpayment and in-the-money shares for informed (solid) and uninformed (dotted) investors as a function of the share of informed investors n ; the right panel also shows the equilibrium expected yield in black. Foreshadowing our quantitative results, we will assume that large bidders are an auction are informed, while smaller bidders are uninformed. Matching the empirical bidding patterns then requires that both information types have similar overpayment, while informed investors have substantially larger in-the-money shares.

This is the case when the share of informed investors is large. The underlying mechanism is consistent with Proposition 4 and appears in the price schedules in Figure 4. When the share of informed investors is large, the high price schedule lies substantially above the low price schedule. To avoid the winner's curse, uninformed investors refrain from bidding at the high-price schedule. Given this choice, it is now optimal to submit the same bids on the low price schedule as informed investors. Hence there is now no difference in overpayment between investors types, but the lack of uninformed participation in the high state implies sharp differences in in-the-money shares, however. The presence of supply shocks further allows us to capture the empirical fact that in-the-money shares are below one and that overpayment is above one on average for all investors.

These results suggest that the share of informed investors is a key determinant of price

Figure 5: Examining Impact of Informed Share.



levels. Average yields are particularly high when n is approximately equal to 0.5. Interestingly, this is *also* the level of n that allows the model to match the empirical bidding patterns. In this sense, the bidding patterns we document suggest that heterogeneous information can have substantial implications for government financing costs. In our quantitative analysis, we use counterfactuals with different degrees of information heterogeneity to provide a quantitative assessment of this cost.

5 Calibration

We now calibrate parameters to the data with two goals in mind. First, we explore whether our parsimonious model can quantitatively account for the bidding behavior of investors in Cetes auctions when asymmetric information is the only source of heterogeneity. Second, we ask whether matching the dynamics of marginal prices can inform us about the sort of information investors have access to (public vs. private news.)

Using the time series dimension of our data requires a dynamic framework. We consider an infinitely-repeated version of our basic model. The government issues bonds in every period, and there are successive generations of one-period investors. Investors can observe all past prices but participate in only one auction. Hence price dynamics are determined by news and the evolution of beliefs but not by dynamic portfolio choice.

To distinguish potential sources of news, we introduce common variation in beliefs by including in our information structure *public regimes* which capture all publicly available information relevant to assessing default risk. There are two regimes indexed by $z \in$

$\{1, 2\}$ with symmetric transition matrix

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

parameterized by the single parameter ρ . *Within* each public regime, we replicate the information structure of the basic model with two possible states $\theta \in \{b, g\}$ which may be known by some investors but unknown by others. Default risk then varies with both z and θ and we index default probabilities by subscript $\{\theta, z\}$. Since investors can typically learn the public shock from past prices, we will simply assume that the current public regime is known at the start of every auction.

Taken together, the parameters which determine the stochastic process for default risk and bond supply are default probabilities $\kappa_{g,z}$ and $\kappa_{b,z}$, probability of the good state $f(g_z)$ in each public regime $z \in \{1, 2\}$, transition probability ρ and the maximum supply shock ψ_M (we maintain our assumption that ψ follows a uniform distribution).

The next set of parameters concerns bidder heterogeneity. Given our theoretical analysis, we focus on information heterogeneity as the key driver of portfolio differences. Nevertheless, we require a model-consistent measure of investor size because there are sharp differences in bidding behavior between the largest bidder and the rest.¹⁴ Proposition 2 implies that we cannot separately identify an investor's wealth from her mass because bids are multiplicatively separable in wealth. We therefore assume that all investors have the same wealth, and calibrate the aggregate share of wealth invested in government bonds net of supply shocks to 20%, i.e. $D/W = 0.2$. This is because Cetes are 25% of all debt instruments auctioned by the Mexican government during our sample period, and the ratio of quarterly Cetes issuance to quarterly GDP in Mexico is stable at 5% during our sample (see Appendix B). We then use investor mass to distinguish size, and assume that there are no other fundamental differences between investors.

Since we can match empirical bidding patterns only if the largest investor is informed, we take as given that there is a large bidder with mass n_{big} , and that this investor is informed.¹⁵ This can be justified by assuming a fixed cost to information acquisition. We then calibrate n_{big} and the total share of informed investors $n \geq n_{big}$.

We fix the coefficient of relative risk aversion to be equal to one (log utility), which is

¹⁴While we focus on the single largest bidder, our results are robust to considering the Top 2 or Top 3 largest bidders as well.

¹⁵Notice that we interpret n_{big} as a mass of informed investors, even though in the data this corresponds to a single large investor. Since wealth differences across bidders do not affect their bidding behavior, the bidding strategy of a mass n_{big} of atomistic bidders is the same as the bidding of a large single bidder with the same total level that does not internalize its size.

low relative to the asset pricing literature. This implies that our quantitative results are due to the pricing protocol and asymmetric information rather than investor preferences.

The next section shows how parameters are pinned down by data moments. Throughout, we focus on bonds with 30-day maturities. This frequency is suitable for our investigation of information heterogeneity because there is a sufficiently long interval in which news might materialize, but the time horizon is short enough such that investors need not forecast the far-off future.

5.1 Disciplining moments

We calibrate the volatility and level of default risk using the mean and standard deviation of marginal prices, which are 0.98 and 0.017 respectively. While the first gives a measure of *unconditional* level of prices, the second gives us a measure of the *unconditional* volatility of prices. These two moments then 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 the coefficient $\beta_1 = 0.98$ from the price regression in (1) and the conditional volatility of prices as measured by the regression's $R^2 = 0.97$. While the first measures the persistence of the marginal price from one auction to the next (the *conditional* volatility of prices), the second captures the *predictability* of prices. These two moments are informative about the relative importance of public news because they suggest highly persistent public regimes (high ρ) with relatively little volatility across states within a regime.

Volatility differences within and across regimes must also be consistent with bidding behavior. In the data, the in-the-money share of the largest bidder is 0.84 on average (on average, the largest bidder buys 84% of his submitted bids), while the average in-the-money share of the remaining investors is just 0.33. At the same time, all bidders pay a quantity-weighted price that is 0.1% above the marginal price on average, leading to average overpayment of 1.01. Our model demonstrates that an informed bidder achieves in-the-money shares and overpayment of 1 if there are no supply shocks. The fact that this is not the case helps discipline the maximum supply shock ψ_M , because it is supply shocks which prevent informed bidders from having all bids accepted in every state.

Our model also suggests that the remaining investors have low in-the-money shares because some of them are uninformed and thus choose not to bid at high prices. Since the largest bidder is informed, n_{big} puts a lower bound on the fraction of informed investors and the difference between n and n_{big} represents the share of informed investors among the remaining bidders. Bidding data then allows us to bound the total share of informed

investors. The largest bidder buys in average 38% of the bonds,

$$\mathbb{E} \left[\frac{n_{big} \mathcal{B}^I(s)}{n \mathcal{B}^I(s) + (1 - n) \mathcal{B}^U(s)} \right] = 0.38$$

The difference in in-the-money shares between the largest bidder and the rest puts an upper bound on n , as the in-the-money share of the remaining bidders is a combination of informed and uninformed investors

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

As n is maximal 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 is intermediate is consistent with a degree of adverse selection that discourages uninformed investors from bidding at high prices. Taken together, ψ_M and n_{big} are relatively well identified by in-the-money shares and the share of bids of the largest bidder, respectively, while the combination of n , ρ , κ and the within-regime state probabilities jointly affect the other moments.

5.2 How well can the model match the data?

We now calibrate the model and assess its ability to match the key data moments discussed above. We consider two approaches. In the “Baseline Calibration” in Section 5.2.1, we calibrate parameters to minimize the sum of squared errors between the data targets and model-generated moments (first and second columns of Table 4, respectively).

This calibration is quantitatively successful in most dimensions, but can only qualitatively match two important pricing moments: the time-series predictability of prices, and the in-the-money shares of small bidders. Section 5.2.2 shows that we can achieve a quantitative match on these dimensions if we posit the existence of a *peso problem* whereby a rare disaster state with high default is considered plausible by investors but did not materialize in our sample period. The magnitude of the disaster required to match these moments is in line with previous Mexican default episodes; hence it may be reasonable for investors to consider such an event.

5.2.1 Baseline Calibration

Table 3 reports parameters for the baseline calibration. 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).¹⁶

Table 3: 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_M	1.3	n	0.40	0.40
ρ	0.999	n_{big}	0.22	0.17

The quantitative fit is imperfect along these dimensions because of the high *unconditional* volatility and low *conditional* volatility that is present in the data. To see why this presents a tension in the model, we look more closely at our price and bids outcomes, which are plotted in Figure 6. The price schedules for both public regimes, plotted in the first panel, show that we are able to generate a moderately high degree of unconditional volatility and a low degree of conditional volatility by having two relatively *high* price schedules in public regime 1, and two relatively *low* price schedules in public regime 2, along with each public regime being highly persistent.

The second and third panel of Figure 6 show both the informed and uninformed bid schedules for public regimes 1 and 2, respectively. The high unconditional volatility in the data implies that one of the public regimes has to include moderately high default rates so that prices are low. Hence we assume that default risk is low on average in regime 1, but relatively higher in regime 2.


To be consistent with relatively high average prices, the average default risk in regime 1 must be low. But this implies the risk-free investment is a very close substitute for the

¹⁶Note that the level of the persistence parameter $\rho = 0.999$, may seem large but reflects the fact that this is a weekly model; $.999^{52} = 0.95$.

Table 4: 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

Figure 6: Baseline Figures



risky bonds in this regime because the risk premium is relatively low. Hence small differences in default risk across quality shocks are sufficient to induce uninformed investors to withdraw from bidding at high prices. That is, the risk premium is so small that participating at high prices is not optimal even if the winner’s curse is mild.

To contend with the low average in-the-money shares of uninformed investors in the data, we then require that uninformed investors do not participate on the high price schedule in public regime 2. This would force informed investors to absorb all default risk. But since default probabilities are relatively high in this regime, informed investors are willing to do so only if prices are very low on average. But since this raises the risk premium and weakens the winner’s curse, high average default probabilities are not consistent with asymmetric participation by uninformed investors. This can be seen in bid schedules for public regime 2: uninformed investors actually bid more on the high-quality schedule than the low-quality schedule for low realizations of the supply shock,

generating a fairly high uninformed ITM public regime 2.

Starting from the the baseline calibration, uninformed bidders would thus achieve a lower in-the money share only if they were much more concerned about the winner’s curse in public regime 2. As we discuss in the next section, it is possible to induce such behavior without strongly affecting average prices by allowing for a higher worst-case default risk κ_{b2} (a *rare disaster*). This approach can remedy the shortcomings of the baseline approach if the rare disaster does not materialize in the sample period.

5.2.2 Rare Disasters and Peso Problems

Our sample is a relatively tranquil period for Mexico: unconditional volatility is moderate, and no government debt crisis occurred. This is in marked contrast to the previous two decades which each saw such crises. We now evaluate our model under the plausible assumption that investors form expectations about default risk by assuming that a similar crisis might occur again. Specifically, we consider an alternative “Black Swan” calibration (BS) in which the worst-case default probability κ_{b2} increases to 0.5 from 0.029, but the probability of this state decreases to 0.05 from 0.45 – a black swan event. Default probabilities in public regime 1 are unchanged. These parameters ensure that average prices are not too high while increasing the severity of the winner’s curse for uninformed bidders in public regime 2.

The likelihood that a black swan event (a debt crisis) does *not* occur in a decade is 31%, consistent with the fact that over the three decades before our data set, Mexico experienced two major crises.¹⁷ Hence our calibration is plausible in light of the data.


Figure 7 plots prices and bid functions for this alternative calibration. Generated moments are reported in the third column of Table 4. 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 the bad state having a very high default probability. The small possibility of this extreme outcome substantially strengthens the winner’s curse and discourages the uninformed from bidding 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 (from 0.62 to 0.43), but it is not as low as in the data.

¹⁷The probability of not having a black swan in a decade obeys the following backwards recursion: denote π_j^i 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

$$\begin{aligned}\pi_j^1 &= .999 * \pi_{j-1}^1 + (1 - .999) * [.95 * \pi_{j-1}^2] \\ \pi_j^2 &= .999 * [.95 * \pi_{j-1}^2] + (1 - .999) * \pi_{j-1}^1.\end{aligned}$$

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

Figure 7: Black Swan Figures



If we simulate data given using these parameters, the Black Swan alternative fails to account for the dynamic properties of prices, however. Even though the probability of the disaster state is fairly low, it substantially raises the standard deviation of prices (to 0.10), depresses our measure of persistence (β falls to 0.14) and increases conditional uncertainty (R^2 falls to 0.02). This is due to the nonlinear nature of these metrics and the weight placed on large pricing errors.

One way to rationalize the data using our rare disasters framework is contending that the data suffers from a *peso problem*, a phrase attributed to Milton Friedman when explaining 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 that has not yet occurred, so that their subjective probability distribution differs from the distribution which has generated a particular sample in historical data (see Rogoff (1980) and Evans (1996)). As explained above, this is a salient concern for Mexico given its history of debt crises.

To examine the extent to which a peso problem can account for the dynamic properties of the data, we maintain parameters at the Black Swan calibration when computing the equilibrium outcomes but, when assessing 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 does not occur in the sample realization. We label this alternative as a "Peso Problem" (PP) calibration. The results are reported in the last column of Table 4. This version of our model does very well, matching almost all of the data moments closely. In particular, it is successful in matching dynamic moments: regression coefficient β has risen to 0.96, while the R^2 has risen to 0.92, both of which are quite close to the data.

The quantitative exploration of our model sheds light on the relevant sources of asym-

metric information. Specifically, our findings suggest that investors do not differ much in their information about typical price movements because they can all access publicly available data on the country’s finances, previous auction results and even the functioning of secondary markets. However, the joint patterns of bidding behavior and dynamic bond price evolution suggest that some investors may be particularly well-informed about events leading to large price swings (such as the disaster state). This information is usually more difficult to access without contacts with decision makers because it may depend on internal political decisions and other sorts of non-public information.

5.3 Re-examining the data

In addition to fitting the data using a calibration exercise, we can use our model to develop and test additional empirical prediction.

The first prediction is that, using the largest bidder as a benchmark for an informed bidder, the remaining bidders are composed of both informed and uninformed bidders. This interpretation is consistent with data if the ITM shares and overpayment of some of the other investors looked like they too were informed. Continuing to use number of bond purchased as measure of private incentives to become informed, we examine whether the second and third largest bidders are relatively closer in their bidding behavior to those of the largest bidders than of the rest. This is indeed the case. Table 5 identifies the bidders that bought the second and third largest shares of bonds per auction, and show that their ITM shares lie between those of the largest and the rest of bidders. Moreover, they do not overpay even though their bids are accepted relatively frequently. We show these statistics for 30-day Cetes, which is the focus of our calibration, but the same results hold for all maturities.¹⁸

Table 5: Second and Third Largest Players

Statistic	Largest	Second	Third	Rest
Share Bidder	0.38	0.23	0.13	0.36
ITM Shares	0.84	0.68	0.55	0.28
Overpayment	1.001	1.001	1.001	1.001

¹⁸We only have two information types: informed and uninformed. The fact that the ITM shares decline as we go from the largest to the second to the third suggests that there was some gradation in the quality of the information among investors in the data; i.e. there were highly informed, somewhat informed, and uninformed types.

The second set of predictions concern the contrasting behavior of the largest bidder (the informed) vs. the rest of bidders *conditional* on a good or bad realization of the quality shock. This set of predictions can be summarized as,

1. The largest bidder buys more when the realized quality shock is good and prices are high, but does not overpay more on average.
2. The largest bidder has similar ITM shares for all realizations of the quality shock, but the rest of bidders have lower ITM shares when prices are high.

To examine these predictions in our data, we need to take a stance on when the quality shock is high vs. low, since both the quality shock and the supply shock affect prices jointly. In our model, the high quality shock leads to above-average prices conditional on the public state., and vice versa. Consistent with this logic, we can recover the quality shock ex-post by measuring the pricing error from our regression model in equation (1). To identify the discrete change in default risk induced by a quality shock (as opposed to a supply shock), we consider pricing errors greater than one standard deviation in absolute terms. We interpret a large positive error as indicating that we have a high quality shock, and a large negative error as indicating we have a low quality one. We find that

1. The largest bidder's ITM share was 0.40 conditional on large positive pricing errors and 0.37 conditional on large negative regression errors. Overpayment is of the same magnitude in both cases.
2. While the largest bidder's ITM share was essentially identical in both cases (0.40 vs. 0.37), the remaining bidders had an ITM share of 0.24 for positive errors, significantly lower than the ITM share of 0.35 during negative errors.

This coherence between our model and the data provides strong additional evidence of the basic mechanism in the model and the importance of asymmetric information.

5.4 Counterfactuals

The numerical illustration in the second panel of Figure 5 suggests that the level of asymmetric information that is consistent with the data (this is, n around 0.4) is likely to raise the government's risky debt yield. In light of this, a natural question is big the effect of asymmetric information on yields is in our preferred calibration. There are two ways to think about this question: (i) how does the future look like with and without information asymmetry, and (ii) what would the past data have looked liked. To answer (i) we compute the outcomes at the two information extremes of symmetric information with $n = 1$

and symmetric ignorance with $n = 0$ and compare them to the predictions at $n = 0.4$, keeping the rest of the Black Swan calibration fixed. To answer (ii), we compute the outcomes that we would expect on a sample path in which the Black Swan event did not occur. We compare the average yield computed using the marginal price at an auction and report the results for (i) and (ii) in table 6 respectively.

Table 6: Counterfactuals: comparing avg. yields

	Black Swan	Black Swan + PP
$n = 0.4$	0.07	0.03
$n = 1.0$	0.06	0.02
$n = 0.0$	0.03	0.03

In the Black Swan scenario, in which a bad default event does occur, if all investors were symmetrically informed, expected yields would decline from 0.07 to 0.06. If all were symmetrically ignorant, yields would fall even further to 0.03. Thus the impact of the information friction coming from having roughly half of the investors being informed substantially raises the cost of borrowing for the government. When we compute our results assuming that the Black Swan event does not occur in sample, this lowers the average yield for both the $n = 0.4$ and the $n = 1$ cases, but the yield gap is still one percent. In the case where $n = 0$, there is no impact on the implied yields. Overall, our results suggest that "ignorance may be bliss," particularly when a "disaster" state materializes in sample.

The reduction in yields obtained in these *symmetric counterfactuals* obtains for different reasons. When $n = 0$ all investors are symmetrically ignorant, which implies better risk sharing. When $n = 1$ investors, all investors participate on the high price schedule, lowering per capita debt burdens and raising prices given a good shock. (There is no effect on the low price schedule because uninformed investors always participate at low prices). Hence the government unambiguously receives higher prices if $n = 1$ than if $n = 0.4$.

6 Final Remarks

By compiling a unique dataset of weekly auctions of Mexican Cetes (domestically denominated) bonds between 2001 and 2017, we document three key facts about prices and bidding behavior patterns. First, the unconditional price volatility is sizable, while the volatility conditional on the previous auction price is low, which suggests high pre-

dictability 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, 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 informed about the probability the bonds default than the rest of bidders.

We perform a calibration of the model that is informative both about the extent of asymmetric information and about its nature. First, to conform to the data, our model implies an intermediate share of informed investors. According to the model, this 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).

When we re-examine the data in light of the model’s predictions, more nuanced implications are also supported by the data. We find, for instance, that the second and third largest bidders behavior in terms of their ITM and overpayment are consistent with the notion that these bidders have the next largest incentives to acquire information. We also find that surprisingly high or low marginal prices, as determined by our price regression, have nearly identical ITM shares for the largest bidder, but much lower ITM shares for the rest, exactly as predicted by our model. Having thus validated our model we then use it to conduct counterfactual exercises as to the impact of heterogeneous information. We find that the impact on yields is sizable, suggesting that asymmetric information is an important friction for a government’s ability to raise funds in primary bond markets.

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 in these markets, and their consequences. 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 Cole, Neuhann, and Ordonez (2020), for instance, we explore the role of information asymmetry for spillovers across countries. By endogenizing information acquisition and extending the setting to 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. Furthermore, our model 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.

References

- Aguiar, Mark, and Manuel Amador. 2013. "Sovereign debt: a review." Technical Report, National Bureau of Economic Research.
- Aguiar, Mark, Manuel Amador, Hugo Hopenhayn, and Iván Werning. 2019. "Take the short route: Equilibrium Default and Debt Maturity." *Econometrica* 87:423–462.
- Aguiar, Mark, Satyajit Chatterjee, Harold Cole, and Zachary Stangebye. 2016. "Quantitative Models of Sovereign Debt Crises." *Handbook of Macroeconomics* 2:1697–1755.
- Aguiar, Mark, and Gita Gopinath. 2006. "Defaultable debt, interest rates and the current account." *Journal of international Economics* 69 (1): 64–83.
- Arellano, Cristina. 2008. "Default Risk and Income Fluctuations in Emerging Markets." *American Economic Review* 98 (3): 690–712.
- Barro, Robert J., and José F. Ursúa. 2012. "Rare Macroeconomic Disasters." *Annual Review of Economics* 4.1:83–109.
- Biais, Bruno, Peter Bossaerts, and Jean-Charles Rochet. 2002. "An Optimal IPO Mechanism." *Review of Economic Studies* 69:117–146.
- Bocola, Luigi. 2016. "The Pass-through of Sovereign Risk." *Journal of Political Economy* 124:879–926.
- Bocola, Luigi, and Alessandro Dovis. 2019. "Self-Fulfilling Debt Crises: A Quantitative Analysis." *American Economic Review* 109:4343–4377.
- Boyarchenko, Nina, David Lucca, and Laura Veldkamp. 2017. "Taking Orders and Taking Notes: Dealer Information Sharing in Treasury Markets." Working Paper, New York University.

- Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo. 2011. "Do peso problems explain the returns to the carry trade?" *The Review of Financial Studies* 24(3):853–891.
- Chatterjee, Satyajit, and Dean Corbae. 2007. "On the aggregate welfare cost of Great Depression unemployment." *Journal of monetary Economics* 54.6:1529–1544.
- Chatterjee, Satyajit, and Burcu Eyigungor. 2012. "Maturity, indebtedness, and default risk." *The American Economic Review* 102 (6): 2674–2699.
- Chaumont, Gaston. 2019. "Sovereign Debt, Default Risk, and the Liquidity of Government Bonds." Working Paper, Rochester.
- Cole, Harold, Daniel Neuhann, and Guillermo Ordonez. 2020. "Information Spillovers in Sovereign Debt Markets." Working Paper. University of Pennsylvania.
- Cuadra, Gabriel, and Horacio Saprizza. 2008. "Sovereign Default, Interest Rates and Political Uncertainty in Emerging Markets." *Journal of International Economics* 76:78–88.
- Eaton, Jonathan, and Mark Gersovitz. 1981. "Debt with Potential Repudiation: Theoretical and Empirical Analysis." *The Review of Economic Studies* 48 (2): 289–309.
- Evans, Martin. 1996. "Peso Problems: Their Theoretical and Empirical Implications." *Handbook of Statistics* 14:Chapter 21.
- Fudenberg, Drew, Markus Mobius, and Adam Szeidl. 2007. "Existence of equilibrium in large double auctions." *Journal of Economic theory* 133 (1): 550–567.
- Geert Bekaert, Robert Hodrick, and David Marshall. 2001. "Peso problem explanations for Term structure anomalies." *Journal of Monetary Economics* 48.2:241–270.
- Gourio, François. 2012. "Disaster Risk and Business Cycles." *American Economic Review* 102 (6): 2734–2766.
- Hatchondo, Juan Carlos, Leonardo Martinez, and Cesar Sosa-Padilla. 2016. "Debt Dilution and Sovereign Default Risk." *Journal of Political Economy*, vol. 124 (1383-1422).
- Kastl, Jakub. 2011. "Discrete bids and empirical inference in divisible good auctions." *The Review of Economic Studies* 78 (3): 974–1014.
- Manzano, Carolina, and Xavier Vives. 2020. "Market Power and Welfare in Asymmetric Divisible Good Auctions." Working Paper, IESE.
- Mendoza, Enrique, and Vivian Yue. 2012. "A General Equilibrium Model of Sovereign Default and Business Cycles." *Quarterly Journal of Economics* 127:889–946.

Passadore, Juan, and Yu Xu. 2019. “Illiquidity in Sovereign Debt Markets.” Working Paper, EIEF.

Reny, Philip J, and Motty Perry. 2006. “Toward a strategic foundation for rational expectations equilibrium.” *Econometrica* 74 (5): 1231–1269.

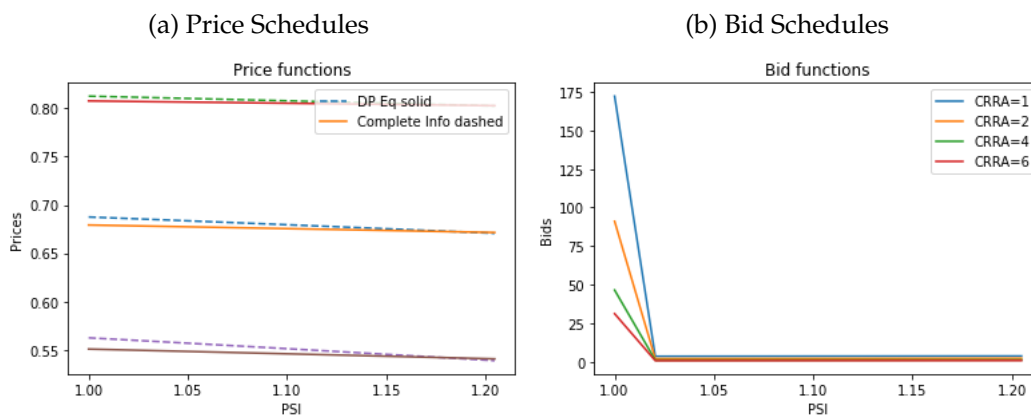
Rogoff, Kenneth. 1980. “Essays on Expectations and Exchange Rate Volatility.” Unpublished Ph.D. Dissertation, M.I.T.

Wilson, Robert. 1979. “Auctions of Shares.” *Quarterly Journal of Economics* 93:675–689.

A Appendix: Symmetric information example

The first panel of Figure 8 compares the price schedules between Examples 1 and 2 (fixing default probability κ_g , κ_b and $\bar{\kappa}$, with and without complete information about the supply shock respectively), for the calibration. As in the main calibration, both schedules are also fairly flat, with the complete information price starting higher and falling faster. This is because investors bid more aggressively when they do not have to worry about overpaying if demand turns out to be low. Differences in default probabilities primarily translate into differences in the *level* of the price schedule, rather than in slope.

Figure 8: Symmetric Benchmarks




The second panel of Figure 8 shows optimal bid schedules for different coefficients of risk aversion under CRRA utility, assuming all investors take as given that the default probability is $\bar{\kappa}$. The bulk of investor’s bids are made at the highest price. This is because the lion share of the supply is sold at this price, with supply only incrementally increases with ψ . More risk aversion leads to more bid shading. However, 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.

B Appendix: Stability in the sample period: 2001-2017.

Here we show that the fiscal revenue obtained by the Mexican government by selling Cetes in primary markets has been stable with respect to GDP during our sample period. This fact maps into the stationarity of D/W imposed in our calibration. To show this stability we compute the quarter revenues raised by auctioning Cetes of all maturities in Mexico (in real terms), and plot it in the next figure as a fraction of quarterly real GDP. As can be seen, every quarter, Mexico raises (or rolls over) around 5% of quarterly GDP

Figure 9: Real Cetes revenue (in all maturities) as a fraction of real GDP



by auctioning Cetes of all maturities. This fraction has been quite stable over our sample period, with a short-lived increase to almost 6% during the global financial crisis of 2009, and returning to pre-crisis levels in 2010. This stability comes from both figures growing in average 0.5% per quarter in real terms over this sample period.