

Asymmetric Information and Sovereign Debt: Theory Meets Mexican Data

Harold Cole

University of Pennsylvania

Daniel Neuhann

University of Texas at Austin

Guillermo Ordoñez

University of Pennsylvania

Using bid-level data from discriminatory auctions for Mexican government bonds, we demonstrate that asymmetric information about default risk is a key friction in sovereign bond markets. We document that large bidders achieve higher bid-acceptance rates than other bidders despite paying no more for executed bids. We then propose a new model of primary markets in which investors may differ in wealth, risk aversion, market power, and information. Only asymmetric information can qualitatively account for our empirical finding, and asymmetric information about rare disasters can quantitatively match bidding and yield moments. Counterfactuals reveal substantial effects of asymmetric information on yields.

We thank Isaac Vivas Escobedo and Gilberto Montano Calvillo from the Bank of Mexico for help with the data and Collum Freedman, Dohan Kim, Juan Sagredo, and German Sanchez for outstanding research assistance. We also thank Greg Kaplan (the editor), three anonymous referees, Yan Bai (discussant), Carlo Galli (discussant), William Fuchs, Travis Johnson, Aaron Pancost, and seminar participants at the Barcelona Graduate School of Economics Summer Forum, the National Bureau of Economic Research International Finance and Macro,

Electronically published June 2, 2022

Journal of Political Economy, volume 130, number 8, August 2022.

© 2022 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/720139>

I. Introduction

Emerging-market sovereign bonds typically offer a risk premium on the order of the US equity premium even when accounting for default losses, and they experience recurring episodes in which yield spreads are elevated and volatile for extended periods of time (Aguiar et al. 2016). These patterns, which have been widely documented using secondary-market price data, have their roots in primary markets, where the government sells bonds to investors directly by conducting auctions. Since bond sales are a key source of financing for most emerging-market governments, the pricing of these bonds in primary markets has important aggregate consequences. We investigate how demand-side characteristics—such as investors' preferences, wealth, market power, and access to information—affect sovereign yields in primary bond markets and ultimately secondary markets.

We focus in particular on asymmetric information among primary bond market investors. At least since Milton Friedman's proposal that the United States sell bonds through uniform price auctions, it has been thought that discriminatory price auctions (also called multiple price or "pay as you bid" auctions) are particularly sensitive to asymmetric information about the risk that bonds default.¹ This may be a particularly pressing concern in emerging markets where sovereign default risk is more salient than in advanced economies. Nevertheless, discriminatory price auctions are extensively used in the world, providing a unique setting to examine the nature and consequences of information frictions in sovereign debt markets.

In discriminatory auctions, bids are executed at the bid price in descending order of prices. In the absence of frictions, an investor achieves a relatively high likelihood of obtaining bonds only if she offers relatively high prices. Hence, multiunit discriminatory bond auctions typically generate a strong positive correlation between the share of an investor's bids that are accepted (her in-the-money [ITM] share) and the average price paid by the investor relative to the marginal price (her overpayment). We test this simple relationship in an emerging-market economy by constructing a data set covering weekly discriminatory auctions of Mexican domestically denominated zero-coupon CETES (certificados de la Tesorería de la Federación) bonds (the dominant source of bond financing for the Mexican government) from 2001 to 2017. Our data allow us to identify all bids and the

the St. Marten's Conference on the Sand, the Toulouse School of Economics, and the University of Texas at Austin for comments and feedback. Cole and Ordoñez received support from the National Science Foundation through grant 1851707. Replication files are provided as supplementary material online. This paper was edited by Greg Kaplan.

¹ Milton Friedman first proposed this idea in *Hearings before the Joint Economic Committee*, 86th Congress, 1st Session, Washington, DC (October 30, 1959, 3023–26), which was partly reproduced in the *Wall Street Journal* article "How to Sell Government Securities" (August 28, 1991). This idea was discussed in more detail by Goldstein (1962).

bond allocation to every bidder at each auction. We do not find the expected positive relation. Instead, we find that investors with distinctively larger acceptance rates pay no more than other investors on average.

To understand these bidding patterns, we propose a new theoretical model of primary markets with discriminatory pricing in which a large number of heterogeneous bidders can bid for multiple units of a risky bond or invest in a risk-free outside option. Our theory points toward the existence of some investors who can more accurately forecast the fundamental value of bonds than others (i.e., there is asymmetric information about the common value of bonds) to explain our empirical patterns.

In a discriminatory auction, asymmetric information leads to the winner's curse: uninformed bids at high prices are executed above the marginal price when the realized bond value is low. We show that equilibrium bidding strategies are strongly asymmetric when the winner's curse is severe: informed investors always bid, while uninformed investors bid only at low prices. Hence, uninformed investors have lower acceptance rates (their bids are accepted only if the realized bond value is low) but they pay similar prices when their bids are accepted (because they know that low-price bids are accepted only when the realized price is low). Thus, asymmetric information can rationalize the bidding facts if the largest investor is informed while some smaller investors are not. This is a natural assumption as long as there is a fixed cost component of information acquisition.

While asymmetric information can account for the bidding facts, differential bidding behavior could also be due to investor characteristics other than access to information. To evaluate this possibility, we allow investors in our model to differ in wealth, risk preferences, and market power as well as information. While other forms of heterogeneity may lead to differences in the number of bonds purchased or prices offered, asymmetric information is the only source of heterogeneity able to break the positive correlation between ITM shares and overpayment. Hence, the bidding facts provide strong evidence that asymmetric information is an important concern in Mexican primary bond markets.

The theoretical mechanism also generates predictions that inform us about the nature of asymmetric information when taking into account data within and across auctions. For instance, the theory places restrictions on the stochastic process for fundamentals, such as the default risk distribution, because we require a sufficiently severe winner's curse to match the bidding facts. Since fundamentals and asymmetric information jointly determine prices, we can thus evaluate our mechanism by asking whether the observed bidding behavior is consistent with the dynamics of observed prices. We document, for instance, that CETES primary prices have a striking dynamic property in our sample: while there is substantial unconditional volatility, an auction's marginal price is highly predictive of next week's marginal price. This implies that primary-market prices must contain

substantial information about fundamental bond values across auctions, while there must be sufficient residual uncertainty within auctions to induce the required winner's curse.

A calibrated version of our model can simultaneously account for the bidding patterns and the level and dynamic behavior of prices. To capture fluctuations in the macroeconomic environment, we assume that fundamentals are driven by a stochastic process in which default risk varies across two public regimes and two states of the world within each regime. Regimes are publicly observable (since they can typically be inferred from past auction prices or other data sources), while there is asymmetric information within regimes (informed investors know the state of the world, uninformed investors do not). Our model matches the data when default risk is very low in one public regime and there is a small chance of a "rare disaster" with high default risk in the other regime.

The key quantitative challenge is matching the low ITM shares of uninformed investors despite mild conditional uncertainty. Doing so is straightforward in the regime with low default risk. Here even a mild winner's curse is sufficient to deter the uninformed from bidding at high prices, since investing in the risk-free outside option is a close substitute. It is more difficult in the regime with higher average default risk because the higher risk premium makes it attractive to submit bids even given the risk of the winner's curse. Here we can rationalize the observed bidding behavior if the winner's curse is driven by a small risk of an extreme event, such as a debt crisis of the size that has occurred several times in Mexico in previous decades. The fact that such a rare and extreme event is not seen in our sample means that the data indicate a so-called peso problem, which has previously been advanced to explain Mexican exchange rate data. While the bidding patterns we document are qualitatively consistent with the existence of asymmetric information, this quantitative exploration that also accommodates the dynamic patterns is informative about the nature of asymmetric information. Investors fear that an extreme event is possible in some states of the world, and informed investors may have privileged access to such information, perhaps because of political connections.²

The degree of asymmetric information we infer from the data has important consequences for investor welfare and government funding costs. We measure the value information for an investor as the additional wealth that would make an uninformed investor as well off as an informed investor.

² In the context of the 1995 Mexican tequila crisis, a particularly pertinent example was knowledge of the inner workings of financial negotiations between President Clinton and the US Congress over a bailout. On January 30, 1995, at exactly the moment when the Mexican government was informing the Clinton administration that without an emergency injection of funds it would have to default, the Speaker of the House, Newt Gingrich, was informing the Clinton administration that the bailout bill was stalled in the Congress. See Chun (1995).

Unsurprisingly, the value of information is low in the low-risk regime in which a rare disaster is impossible and prices are high. However, it is much higher in the regime in which a rare disaster is a possibility, reaching 1% of investor wealth. This shows that the winner's curse is a substantial impediment to auction participation when default becomes more likely. This has clear implications for government funding costs. We find, for instance, that the average annualized yield on government debt would have declined by 1.4 percentage points ($\approx 45\%$) if all investors had been symmetrically informed and the disaster did not occur on path, while yield volatility would have declined by half under the same conditions.

We confirm additional model predictions in the bidding data. First, while our baseline analysis distinguishes only between the largest bidder and the rest, other relatively large bidders are also likely to be more informed than smaller bidders. We find that the second- and third-largest bidders behave similarly to the largest bidder but that ITM share differences with the smaller bidders are less pronounced. This suggests some gradation in the quality of information. Second, according to our model, differences in bidding outcomes are driven by differences in bidding behavior at high prices. To assess this mechanism *ex post*, we identify good and bad states of the world by running predictive regressions and isolating auctions in which the realized marginal price is 2 standard deviations above or below the predicted price. Consistent with the theory, we find that ITM shares of smaller bidders are distinctly lower (higher) in auctions with unusually high (low) marginal prices.

Finally, we extend our model to include secondary markets. This extension provides two insights. First, and perhaps surprisingly, the fact that secondary-market prices reflect information revealed at auction may strengthen adverse selection in the primary market. The reason is that secondary markets allow informed investors to sell some of their purchases at a risk-free arbitrage profit. This induces them to bid more aggressively at auction, which strengthens the winner's curse by widening the gap between prices conditional on good and bad news. Furthermore, since uninformed investors have the option to trade in the secondary market without fear of adverse selection, the arbitrage spread between primary and secondary markets earned by informed traders provides an observable measure of the winner's curse.

The second insight from incorporating secondary markets is empirical. We use daily CETES secondary-market data for different maturities and provide evidence of the existence, nature, and extent of asymmetric information in primary markets. First, we show that auctions are information events: secondary-market yields are twice as volatile on auction days than on nonauction days, and an unexpected increase of 1% in primary-market prices generates an unexpected increase of 0.74% in same-day secondary-market prices. Second, the information content in auctions is about the fundamental quality of bonds: unexpected increases in primary prices

predict increases in next week's auction price, and these unexpected changes are highly correlated across different maturities on auction days. Third, the extent of asymmetric information (as reflected in bidding behavior) is correlated with arbitrage gains. We find that the average difference in ITM shares between the largest investors and the rest was larger before 2009 (except for 1-year CETES), suggesting a more severe winner's curse. Consistent with the model, the average spreads between primary and secondary yields on auction days were 0.26% before 2009 and only 0.11% after 2009. This difference is consistent in magnitude with our counterfactual exercise. For the 1-year CETES, the average spreads were roughly constant, suggesting no change in the severity of the winner's curse. Consistent with this, the ITM shares of the largest bidder were also roughly constant between these two periods.

Related literature.—The existing literature on sovereign default studies bond pricing by focusing on governments' strategic default decisions while using a parsimonious model of investor optimization (see, e.g., Chatterjee and Eyigungor 2012; Mendoza and Yue 2012; Hatchondo, Martinez, and Sosa-Padilla 2016). In this literature, sovereign debt prices are modeled using competitive pricing rules under which risk-adjusted bond yields are equal to the risk-free rate.³ While 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), 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 auction mechanics and investors' choices while entirely neglecting strategic considerations on the part of the government. Our paper argues that the neglected roles of auction mechanics and information heterogeneity (and their interaction) drive 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 challenges for standard auction models in accommodating asymmetric information among risk-averse investors.⁵ In this context, our framework can be viewed as an auction model

³ See, e.g., 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), and Bocola and Dovis (2019).

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

⁵ For a discussion, see, e.g., 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.

with three key characteristics: (i) the good for sale is perfectly divisible, (ii) the number of bidders is large and/or bidders behave competitively, and (iii) there is uncertainty about both good quality and 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

A major benefit of this approach is that we can characterize the equilibrium mapping from parameters and state variables to outcomes. This is currently not feasible for a strategic discriminatory auction model (see Hortaçsu and McAdams 2010, n. 9) or requires strong functional form assumptions (as in Février, Preget, and Visser 2004). Thus, we do not need to resort to the resampling methods that are common in the literature but require strong independence assumptions across different auctions. We also do not need to assume that valuation differences come from either private value shocks or conditionally independent signals about the common value. In fact, for simplicity we assume that the common value signal regarding default risk is perfectly correlated among the informed and effectively perfectly uninformative for the uninformed. We also consider risk-averse investors, which is important when modeling sovereign bond risk premia.

Besides the theoretical contribution, our paper also complements empirical efforts to measure the implications of information for government revenues in auctions, such as Boyarchenko, Lucca, and Veldkamp (2021), who analyze the impact of information sharing among dealers on US Treasury prices, or Hortaçsu and Kastl (2012), who measure dealers' information advantage from observing order flow in Canadian Treasury auctions.

Finally, our paper contributes to recent work studying the interaction between primary and secondary sovereign bond markets, such as Chaumont (2020) and Passadore and Xu (2020). These papers argue that secondary-market liquidity can affect primary-market prices by directly altering the risk of default and thus the fundamental value of bonds. We instead focus on the extent of the winner's curse induced by asymmetric information in primary markets, holding default risk and the fundamental value of bonds fixed.

Section II describes our data and establishes basic facts about bidding patterns and price dynamics. In section III, we develop a tractable primary-market model with a discriminatory pricing protocol and rich bidder

⁶ Recent auction literature shows that price taking arises as the number of bidders gets 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.

heterogeneity. In section IV, we show that information asymmetry is critical for accommodating the empirical bidding patterns we document. We also conduct a detailed analysis of yields under various levels of asymmetric information, which is useful for conducting counterfactuals. In section V, we calibrate the model to match relevant moments on price dynamics and bidding behavior to quantify the extent of asymmetric information in Mexican bond markets and conduct counterfactuals. In section VI, we validate further testable implications of the model, both in primary markets and in their relation with secondary markets. Section VII concludes.

II. Institutional Background and Data

We study auctions of Mexican federal Treasury bills (CETES), which are domestically denominated 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 debt funding since 1978. The primary market for CETES consists of closed-system (allocations and prices are disclosed only after the auction is closed) public auctions conducted by the Mexican central bank, which acts as a financial broker.⁷ Auctions take place weekly—almost always on Tuesdays from 10 to 11 a.m.—and results are disclosed at 11:30 a.m.⁸

In 2000, the Ministry of Finance (Secretaría de Hacienda y Crédito Público [SHCP]) created the Market Makers program to foster the development of primary and secondary fixed-rate government securities markets. Market makers are credit institutions and brokerage firms (including Banamex, Bank of America, JP Morgan, and others) appointed by the Ministry of Finance to present bids at competitive prices in each primary auction for a minimum amount of bids totaling 20% of the amount offered (or one divided by the number of bidders, whatever is smaller). Market makers must also permanently quote purchase (bid) and sale (offer) prices in the secondary market to provide liquidity.⁹ This program induced regular competition in primary markets and fostered secondary-market stability. More details about these rules and the identities of market makers are provided in appendix A.

⁷ The other instruments auctioned by the central bank are *bondes* (floating-rate government securities with maturities of 3, 5, and 7 years), *bonos* (fixed-rate bonds placed at 3, 5, 10, 20, and 30 years), and *udibonos* (inflation-hedged instruments).

⁸ Occasionally, 364-day maturity bonds were auctioned monthly instead of weekly.

⁹ The benefits of being a market maker consist of access to securities lending through the central bank facility, exclusive participation in syndicated placements, and access to a next-day purchase option to buy additional securities at the auction price on the day after the auction. A market maker not complying on submitting bids according to these rules cannot exercise the next-day purchase option.

TABLE 1
SUMMARY DATA ON CETES AUCTIONS, 2001–17

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

Our data are from the archives of the Mexican central bank. We focus on the period June 2001–September 2017, which has three key advantages. First, CETES were consistently auctioned using a discriminatory-price protocol throughout this period (a switch to a uniform-price protocol occurred in October 2017). Second, Mexico experienced relatively stable inflation and did not suffer any major crisis during the period. Hence, CETES auctions took place on a regular weekly schedule for all maturities. Third, our data include all bids submitted (not only those that were executed) by all bidders for all auctions. 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 three bids per auction. Table 1 shows summary statistics for each maturity.

A. *Bidding Patterns*

We now establish basic facts about bidding patterns. We cannot track bidders across auctions because the numeric bidder identifier is auction specific. To uncover heterogeneous bidding behavior, we therefore compare the bidding behavior of the largest bidder at an auction with all other bidders, where the largest bidder is the bidder who buys the most bonds in an auction.¹⁰ 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 the 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 1 shows a histogram of the share of all bids that are accepted at an auction, per bidder, aggregated across all auctions and maturities. We call this fraction of accepted bids the ITM share and distinguish between the largest bidder and the rest of bidders. For the largest bidder, the mode of the ITM across auctions is one (typically, all of their bids are accepted), while there is much more dispersion for smaller bidders. On average, the

¹⁰ Our results are robust to considering the top two or three bidders and to defining the largest bidder as the one who submits the most bids.

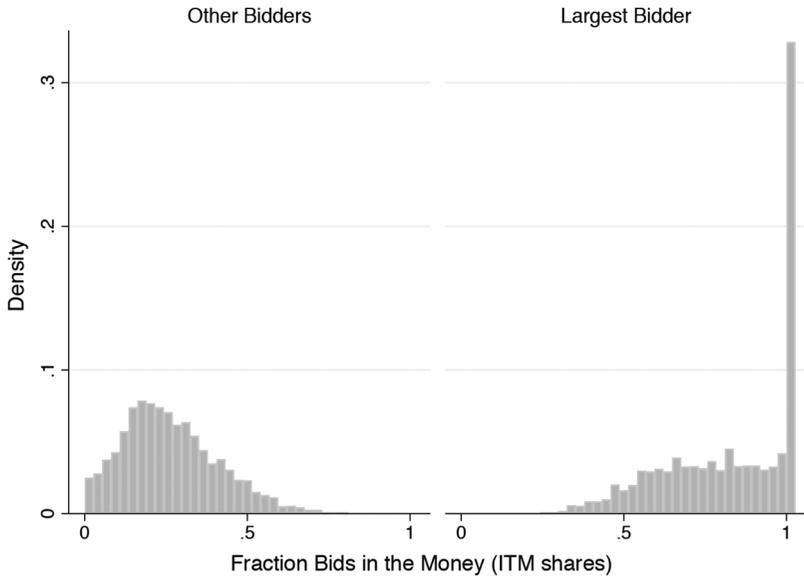


FIG. 1.—ITM shares, largest versus the rest. A bidder's ITM share is computed as the fraction of the bidder's bids (in pesos) that were allocated to the bidder. We compute this share for the largest bidder and for all other bidders in each of the 2,717 auctions in the sample. Source: Bank of Mexico. A color version of this figure is available online.

largest investor at a given auction has 84% of bids executed, while only 33% of remaining bidders' bids are executed.

One possible explanation for this finding 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 2, we show a histogram (for all auctions and all maturities) of the overpayment for the largest bidder and the rest of bidders.

The distribution of overpayment is very similar for large bidders and other bidders. The combination of high ITM 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. Averaging across auctions, one would therefore expect to see a positive correlation between average ITM shares and average overpayment. 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.

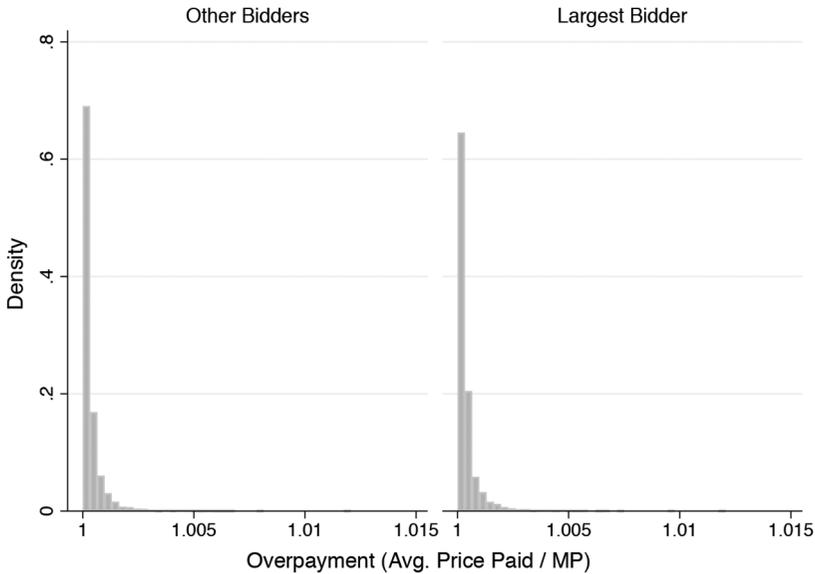


FIG. 2.—Ratio between weighted price paid and the marginal price (MP). A bidder's overpayment is the quantity-weighted average price paid for purchased bonds divided by the marginal price (the minimum price accepted at the auction). We compute this ratio for the largest bidder and for all other bidders in each of the 2,717 auctions in the sample. Source: Bank of Mexico. A color version of this figure is available online.

B. Price Dynamics

We now establish a number of basic facts about price dynamics. During our sample period, Mexico experienced relatively stable macroeconomic conditions and inflation. This is reflected in relatively low average yields and mild conditional volatility of auction prices. Figure 3 shows the real marginal price in CETES auctions, defined by the lowest accepted price in an auction and computed using the annual yield deflated by annual consumer price index inflation, for the four most common maturities of 28, 91, 182, and 364 days.

Table 2 shows the time-series moments of marginal prices by maturity. 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. Our results suggest that the unconditional uncertainty of the marginal price (its standard deviation) is much higher than the uncertainty conditional on the previous week's auction results. To further characterize the conditional uncertainty, we regress marginal prices of the 28-day bond on a constant and one lag,

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

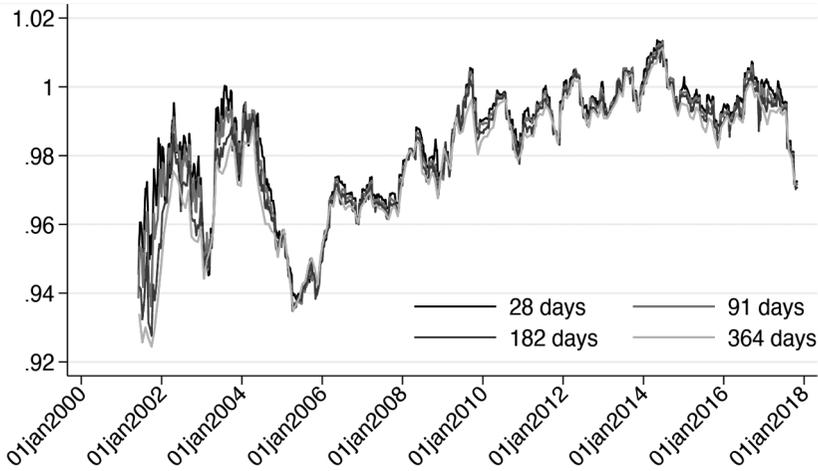


FIG. 3.—Marginal prices of CETES bonds of different maturity. Marginal prices are constructed from the maximum discount d accepted at the corresponding auction. We obtain the corresponding real yield as described by the Bank of Mexico, $r = [(1 - d)/d] - \pi$, where π represents the inflation rate of the corresponding period. The annualized real marginal price of a bond with nominal value one is $P = 1/(1 + r)$. Source: Bank of Mexico. A color version of this figure is available online.

We estimate $\beta_1 = 0.98$ and $R^2 = 0.97$, which implies that lagged prices are very informative and that conditional uncertainty is indeed quite low during this period.¹¹

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 show that even a small amount of conditional uncertainty can have significant effects on bidding strategies and prices.

III. 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 V, 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

¹¹ As can be guessed from fig. 3 and table 4, running the same regression for other maturities leads to very similar results. Note also that our approach is essentially a single-factor term structure model estimated on data from one bond only. While there may be additional information encoded in other contemporaneous variables, our goal is to show that even a very simple approach implies strong predictability. Later, we calibrate to the results from this regression using an indirect inference approach.

TABLE 2
TIME-SERIES PROPERTIES OF MARGINAL PRICES (MPs)

Maturity	Average MP	Standard Deviation MP	Autocorrelation MP
28	.984	.017	.984
91	.983	.018	.983
182	.982	.018	.992
364	.978	.019	.956

heterogeneity that can rationalize the bidding patterns previously discussed. Our baseline model abstracts from secondary markets, but appendix B extends the model to include secondary markets and shows that all key findings are robust to this change.

A. Environment

There is a single period with two dates ($t = 1, 2$) and a single good (the numeraire). Bonds are sold at auction at the first date, and consumption occurs at the second date. 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 1 and cannot borrow. Investors invest their wealth in either a risk-free asset (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.¹² The 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 $f(g) \in (0, 1)$ and $f(b) = 1 - f(g)$, respectively; the unconditional default probability is

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

Since the default probability determines the expected value 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 1-period model, we can capture different bond maturities by the length of the period. If we view

¹² We explore robustness to positive recovery rates in app. C.

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 ITM shares, we need an additional source of uncertainty. Hence, we introduce a supply shock ψ to the government financing needs, where ψD represents the total revenue that the government must raise (e.g., 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 $h(\psi_k)$.¹⁵ The set of possible states of the world is $\mathcal{S} = \{g, b\} \times \mathcal{H}$.

B. 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 also allow us to capture differences in market power. For simplicity, we consider only two fundamental types of equal mass at a time; thus, we assume that investors either have the same preferences but differ in their wealth or 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 a type-specific set of plausible states \mathcal{F}_j^i that collects 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 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.

C. Pricing Protocol and Strategies

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

¹⁵ An alternative interpretation is that there is a demand shock, which precludes a share η of investors from participating in the auction (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 they personally suffered such a shock. In the interest of parsimony, we therefore use supply shocks.

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), and all bids below are rejected (they are out of the money).¹⁴ The set of marginal prices associated with the good-quality shock $P(g, \cdot)$ is the high price schedule, and the set of marginal prices associated with the bad-quality shock $P(b, \cdot)$ is the low price schedule. The state-contingent yield is

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

Since bonds pay off at least zero and at most one unit of the numeraire, the range of 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.¹⁵ Since there is a single marginal price associated with each state, we can equivalently define bidding strategies as functions that map 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 set of feasible states \mathcal{F}_j^i into bids consisting of quantity $B_j^i(s) \geq 0$ and bid price $P(s)$.

Strictly speaking, informed investors need not choose bids for states of the world they know will not occur (i.e., states that are not associated with the realized quality shock). However, to compactly describe decision problems, it is convenient to assume that they first choose bids for all possible states \mathcal{S} and then discard any bids that are not associated with the realized

¹⁴ If 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 zero and one, suggesting that it is not playing a key role. As rounding bids does not add any insights, we follow the literature that assumes that the set of possible bid prices is in a continuum, as in the seminal work of Wilson (1979). For a treatment of a uniform-price auction of a perfectly divisible good in which bidders are restricted to bids at discrete points, see Kastl (2011).

¹⁵ Observe 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.

quality shock. This ensures that bidding strategies are well defined for all investors and possible states.

D. Decision Problems and Equilibrium Definition

An investor's objective function depends on the final portfolio acquired at auction. Hence, it is convenient to define sets of executed bids $\mathcal{E}_j^i(s)$, which collect all bids by an investor of type $\{i, j\}$ that are executed in state $s = \{\theta, \psi\}$. Since each bid is associated with a state-specific marginal price, the elements of these sets are states of the world.

For uninformed investors, the executed bid set includes all states with marginal prices above the realized marginal price. For informed investors, there is the extra requirement that states must correspond to the realized quality shock. Hence, we have

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

The total quantity of bonds purchased in state s by an investor of type $\{i, j\}$ is

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

and the total amount expended on bonds in state s is

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

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 ITM share in state s is the ratio of accepted bids to submitted bids,

$$\text{ITM}_j^i(s) = \frac{\sum_{\tilde{s} \in \mathcal{E}_j^i(s)} B_j^i(\tilde{s})}{\sum_{\tilde{s} \in \mathcal{F}_j^i} B_j^i(\tilde{s})}. \quad (5)$$

The bid-weighted average price, AP, paid in state s is

$$\text{AP}_j^i(s) = \sum_{s' \in \mathcal{E}_j^i(s)} \frac{P(s') B_j^i(s')}{\sum_{\tilde{s} \in \mathcal{E}_j^i(s)} B_j^i(\tilde{s})}. \quad (6)$$

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

$$\Omega_j^i(s) = \frac{\text{AP}_j^i(s)}{P(s)}. \quad (7)$$

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 ,

$$V_j^i(\tilde{B}_j^i) = \mathbb{E}_{\theta,\psi}[\kappa_\theta U_j(w_j^i(s)) + (1 - \kappa_\theta)U_j(w_j^i(s) + B_j^i(s))|\mathcal{F}_j^i], \tag{8}$$

and the associated decision problem is

$$\max_{\tilde{B}_j^i} V_j^i(\tilde{B}_j^i) \quad \text{s.t.} \quad B_j^i(s) \geq 0; w_j^i(s) \geq 0 \text{ for all } s \in \mathcal{S}, \tag{9}$$

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

The market-clearing condition ensuring that the government raises revenue ψD in state $s = (\theta, \psi)$ given share n of informed investors and share $1/2$ of fundamental type j is

$$\psi D = \frac{1}{2} \sum_{j \in \{1,2\}} (nX_j^I(s) + (1 - n)X_j^U). \tag{10}$$

We are now ready to state our equilibrium definition.

DEFINITION 2 (Equilibrium). An 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 (9) for all types and
2. the market clearing condition (10) is satisfied for all $s \in \mathcal{S}$.

IV. Equilibrium Characterization

A. Optimal Bids and Equilibrium Prices

We begin by characterizing optimal bidding strategies and equilibrium prices. Since preferences satisfy Inada conditions, borrowing constraints do not bind. However, the nonnegativity constraint on bids may bind for some investors in some states of the world.

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.¹⁶ 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 that have a lower marginal price. That is,

¹⁶ Acceptance sets are complements of executed bid sets. The former collect all states with marginal prices that are lower than the bid price, and the latter collect all bids that were submitted at higher prices and are thus executed in the current state.

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

First-order conditions are necessary and sufficient for optimality in the investor's decision problem. Let $E_j^i[y] \equiv E[y|\mathcal{F}_j^i]$ denote the type-specific expectations operator, and define investor $\{i, j\}$'s expected marginal rate of substitution for bids at $P(s^*)$

$$M_j^i(s^*) = \frac{\mathbb{E}_j^i \left[\sum_{s \in \mathcal{A}_j^i(s^*)} \kappa(s) U_j'(w_j^i(s)) \right]}{\mathbb{E}_j^i \left[\sum_{s \in \mathcal{A}_j^i(s^*)} (1 - \kappa(s)) U_j'(w_j^i(s) + \mathcal{B}_j^i(s)) \right]},$$

where the numerator is expected marginal utility after default and the denominator is expected marginal utility after repayment. The proof of the next proposition shows that the associated first-order condition for optimal bids is

$$M_j^i(s^*) \geq \frac{1 - P(s^*)}{P(s^*)} \quad \text{and with equality if } B_j^i(s^*) > 0. \quad (11)$$

This condition states that the investor trades off yields against the marginal gain from repayment relative to the marginal loss after default across all states in which the bid will be accepted. Importantly, consumption levels are indexed by s rather than s^* when constructing the marginal rate of substitution. This is because the realized portfolio differs across states in the acceptance set. Even given identical bids at each price, uninformed investors thus face different optimality conditions than informed investors because acceptance sets may contain more states (in particular, low realizations of the quality shock). For some investors, yields may thus be too low to warrant participation, in which case the nonnegativity constraint on bids binds.

The following result characterizes equilibrium prices and shows that bond yields satisfy asset pricing relationships that are sensitive 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:

$$\frac{1 - P(s)}{P(s)} = \min_{i,j} M_j^i(s) \quad \text{for all } s \in \mathcal{S}. \quad (12)$$

- (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). By the investor's decision problem, necessary and sufficient condition for optimality of the bid at marginal price $P(s^*)$ associated with state $s^* = (\theta^*, \psi^*)$ for an informed investor of fundamental type j is

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

where $\chi_j^I(s^*)$ is the Lagrange multiplier on the nonnegativity constraint on bids. Since informed investors know the quality shock, expectations are taken only with respect to the supply shock ψ . Similarly, for an uninformed investor of fundamental type j , the necessary and sufficient condition for bids at marginal price $P(s^*)$ is

$$\mathbb{E}_{\theta, \psi} \left[\sum_{s \in \mathcal{A}_j^U(s^*)} - U_j^U(w_j^U(s)) \kappa_{\theta} P(s^*) + U_j^U(w_j^U(s) + B_j^U(s))(1 - \kappa_{\theta})(1 - P(s^*)) \right] - \chi_j^U(s^*) = 0, \tag{13}$$

where expectations are taken with respect to both the quality shock and the supply shock. Rearranging these equations gives (11). 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. Equation (12) then follows.

ii). The government can raise more revenue only if at least one type of agent consumes less after a default. Since agents have constant relative risk aversion (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 that $P(g, \psi) = P(b, \psi)$, 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 (10), $\lim_{n \rightarrow 0} \sum_{j \in (1,2)} (1/2) X_j^U(\theta, \psi) \rightarrow D\psi$ for all ψ . As uninformed bid unconditionally of θ , $\lim_{n \rightarrow 0} \sum_{j \in (1,2)} (1/2) X_j^U(g, \psi) \rightarrow \sum_{j \in (1,2)} (1/2) X_j^U(b, \psi)$ 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 an ϵ small enough such that $\psi' - \psi = \epsilon$ —that is, $P(\theta, \psi) < P(\theta', \psi') < P(\theta, \psi')$. This proof does not apply at extreme values of ψ , and hence convergence will not happen at $\psi = \psi_0$ and $\psi = \psi_M$. QED

Statement i 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 to standard frameworks 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 across different states of the world and implies that bids at high prices affect willingness to pay in all states with lower marginal prices. Statement ii 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.

Statement iii 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 is realized. This discourages uninformed investors from bidding at high prices and has implications for risk sharing, average yields, and equilibrium bidding patterns. Statement iv shows that the wedge disappears as the share of informed investors approaches zero.

The following simple examples provide insights into the basic mechanics of bidding and motivate our analysis of heterogeneous bidders and asymmetric information.

EXAMPLE 1 (Homogeneous investors with complete information). Suppose that all investors are ex ante symmetric, have log preferences,

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 bids. For any ψ , the market-clearing condition is $B(\psi)P(\psi) = \psi D$ 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 risk premium is equal to the default probability multiplied by the ratio of the per capita debt burden ψD to net wealth $W - \psi D$. Hence, prices are strictly ordered by the demand shock and the quality shock. Since every investor can correctly forecast the marginal price in every state, all bids are ITM in every state, and every investor pays the same price. In contrast to the data, ITM shares and overpayment are thus equal to one for all investors:

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

The next example shows that supply uncertainty can create variation in ITM 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, have log preferences, and have common expectations about the quality shock. In contrast to the previous example, assume that they are uncertain about the supply shock. Further, assume that the supply shock has two possible realizations, $\psi \in \{\psi_1, \psi_2\}$ with $\psi_1 < \psi_2$, and 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 ITM 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

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

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

$$\begin{aligned} \text{ITM}(2) &= 1, & \text{AP}(2) &= \frac{B(1)P(1) + B(2)P(2)}{B(1) + B(2)} > P(2) \\ \Omega(2) &= \frac{B(1)(P(1)/P(2)) + B(2)}{B(1) + B(2)} > 1. \end{aligned}$$

Hence, investors choose symmetric strategies, and—contrary to the data—ITM and overpayment are positively correlated.

The positive correlation between ITM 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.

B. Using Heterogeneity to Account for the Empirical Bidding Patterns

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

1. Fundamental Heterogeneity

We first consider the effects of fundamental heterogeneity in 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 ITM shares or average prices paid, are invariant to wealth heterogeneity. (This result readily extends to constant absolute risk aversion utility, which has no wealth effects.) 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, ITM shares and average prices paid are the same for all types.

Proof. Multiplicative separability of optimal policies is a standard property of CRRA preferences. The remaining statements then follow immediately. QED

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 ITM 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 effects 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 ITM shares and overpayment across investors. Hence, fundamental heterogeneity of any form cannot generate the facts documented in section II.

Proof. Under symmetric information, there is a single 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 ITM shares and overpayment for each type j in each state as monotone functions of this ratio only. For state 1, we obtain

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

where $\text{ITM}_j(1)$ is strictly increasing in ρ_j . In state 2, we have

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

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

$$\mathbb{E}[\text{ITM}_j(\psi)] > \mathbb{E}[\text{ITM}_{j'}(\psi)] \Leftrightarrow \mathbb{E}[\Omega_j(\psi)] > \mathbb{E}[\Omega_{j'}(\psi)].$$

QED

In a pay-your-bid protocol with symmetric information, an investor can attain a relatively high acceptance rate only by submitting a relatively large share of bids at high marginal prices. But this implies that the investor will also overpay when purchasing those bids. Hence, ITM shares and overpayment are both governed by the ratio of bids ρ_j , and any form of fundamental heterogeneity that leads to differences in ITM shares must necessarily lead to a positive comovement with overpayment.¹⁷

This intuition is particularly striking if the largest investor is also the one with the largest market power. As discussed in Hortaçsu (2011), investors with market power tend to shade their bids, which implies placing bids at lower prices than investors without market power. If this is the case, the largest investors would also achieve lower ITM shares than investors without market power (i.e., large investors exercise their market power by offering lower prices while accepting lower acceptance probabilities). Indeed, Hortaçsu, Kastl, and Zhang (2018, 153) study US Treasury auctions and show that “the fact that primary dealers are winning a smaller share of their tenders than direct and indirect bidders suggests that primary dealers bid systematically higher (lower) yields (prices) in these auctions.” This result is counterfactual in our data, as we have shown that the largest bidder is the one who buys the largest fraction of their bids.

The fact that fundamental heterogeneity cannot account for the bidding patterns does not imply that investors are homogeneous in wealth, market power, or risk aversion. Instead, we claim that in spite of these differences there is strong evidence that asymmetric information is an important source of investor heterogeneity.

2. Informational Heterogeneity

We now show that heterogeneous information can break the positive association between ITM shares and overpayment that arises under fundamental heterogeneity. This allows us to match the empirical lack of correlation between ITM shares and overpayment documented in section II. 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 and that there is a group of investors (namely, the uninformed) who submit bids at the price schedule that is not operational given the realized state. Thus, informed investors target bids to the realized quality shock, while uninformed investors face a winner’s curse

¹⁷ When considering richer supply shock specifications, numerical analysis shows that this result is robust to considering differences in risk aversion. This is because more risk-averse agents will submit uniformly lower bids (i.e., bids scaled by an approximately constant factor) at all prices. Measures that depend on ratios of bids, such as ITM shares or overpayment, are thus not differentially affected by such preference heterogeneity.

because they may submit bids on the “wrong” schedule (the one associated with the unrealized quality shock).

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 ITM shares and overpayment. However, the lack of supply shocks leads to counterfactual predictions for informed investors’ ITM 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 investors’ ITM 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 $nB^I(g)P(g) = D$ and $[nB^I(b) + (1 - n)B^U(b)]P(b) = B^I(b)P(b) = D$, where

$$B^I(\theta) = \frac{W(1 - \kappa(\theta) - P(\theta))}{P(\theta)(1 - P(\theta))}. \tag{15}$$

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 consider ITM shares and overpayment. Informed investors submit state-contingent bids at the marginal price in every state, and all of their bids are always accepted (their ITM share is always one). 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 ITM shares than informed investors (ITM shares are either zero or one and thus below one on average). Like informed investors, however, uninformed investors never overpay if their bids are accepted because they submit bids only at $P(b)$.

Last, we must verify the optimality of $B^U(g) = 0$. A sufficient condition 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

$$1 - \bar{\kappa} < P(g) \quad \Leftrightarrow \quad \bar{\kappa} > \kappa(g) \frac{nW}{nW - D}. \quad (16)$$

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. QED

The key insight is that uninformed investors stop bidding on the high price schedule if the winner's curse is sufficiently severe, in which case they can bid on the low price schedule as if they were informed because bids on the low price schedule are accepted if and only if $\theta = b$. Asymmetric participation given $\theta = g$ and symmetric bidding given $\theta = b$ leads to differences in ITM shares without differences in prices paid conditional on bid execution. Hence, the bidding facts strongly suggest the presence of asymmetric information. (Of course, this does not imply the absence of other investor differences.)¹⁸

REMARK. The existence of two distinct state-contingent price schedules is a necessary but not sufficient condition for breaking the correlation between ITM shares and overpayment across auctions. The other necessary condition is that some investors submit bids on an "incorrect" price schedule (i.e., one associated with an unrealized state of the world). The argument works as follows. If some investors want to bid more in some states of the world than in others (e.g., owing to state-contingent preferences), then there will be distinct state-contingent price schedules. But if there is symmetric information about the realized state, there is a single operational price schedule for each state (no bids will be submitted on the other schedule because it is known to be irrelevant). Hence, our analysis of fundamental heterogeneity is valid state by state, and we would observe a positive correlation between ITM shares and overpayment state by state. For there to be investors with a low ITM share (as in the data), these investors must bid on the low price schedule when the high price schedule is operative. This is the key feature of asymmetric information that is difficult to rationalize by other sources of heterogeneity.

The key counterfactual implication of proposition 4 is that ITM shares are always equal to one for some investors and that overpayment is always equal to one for all investors. We use the following numerical example to show that the combination of supply shocks and information heterogeneity can qualitatively account for the empirical bidding patterns.

¹⁸ We could easily accommodate a known distribution of other sources of fundamental heterogeneity (e.g., wealth) in our setting. Given our assumption of a continuum of investors, an individual's realization of wealth works like a private value shock; scaling bids up or down in proportion to the realized wealth. The equilibrium marginal price function and the aggregate number of bids at each potential marginal price would depend only on the mean level of wealth, and hence the analysis would be invariant to this extension.

Parameters are chosen for ease of exposition; we formally calibrate parameters to data in section V.

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 default probabilities satisfy $\kappa_g = 0.15$ and $\kappa_b = 0.35$. Supply shock 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 values of the share of informed investors n . The high price schedule is shown in black open circles, and the low price schedule is shown in gray filled circles. Solid gray horizontal lines show three benchmark prices. In descending order, they are the high price schedule when all investors are informed ($n = 1$), the unconditional price schedule when no investor is informed ($n = 0$), and the low price schedule when all investors are informed ($n = 1$).

Price schedules are very sensitive to the winner’s curse. When $n = 0.6$, informed investors participate at high prices but uninformed investors drop out to avoid the winner’s curse. Since informed investors must therefore bear more risk per capita, the high price schedule falls to provide a sufficient risk premium. Conversely, the low price schedule is locally independent of n because uninformed investors can submit bids at low prices as if they were informed.

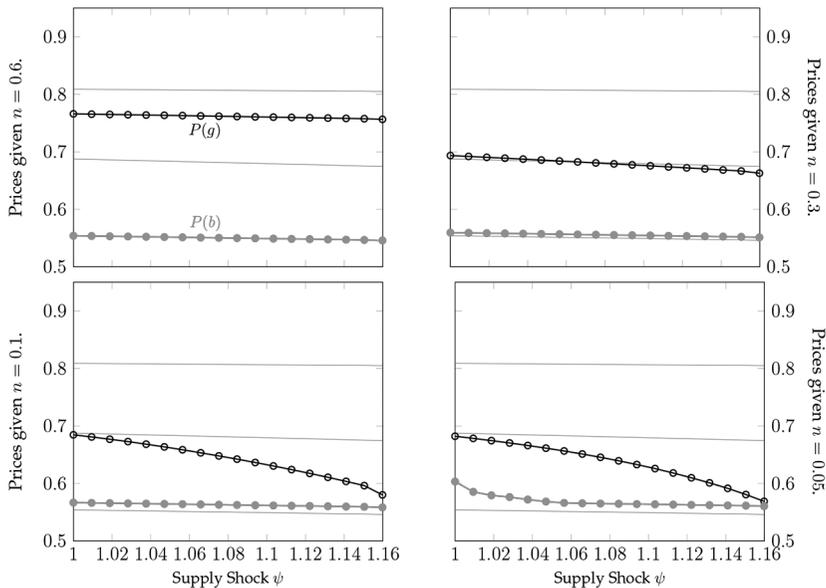


FIG. 4.—Equilibrium prices for various values of n . A color version of this figure is available online.

The risk premium is even higher when $n = 0.3$ because there are now even fewer informed investors who bear risk conditional on a good shock. This leads to a decline in the high price schedule that weakens the winner's curse, which induces uninformed investors to bid on both price schedules. Since bids on the high price schedule are also executed in the bad state, there is less residual demand that needs to be met by marginal bids on the low price schedule, and the low price schedule rises.

The winner's curse effect continues to operate as n decreases to $n = 0.1$. Because high-price bids of the uninformed remain below those of the informed on a per capita basis, a further decrease in n further concentrates default risk among informed investors. This forces a large fraction of the high price schedule to drop below the benchmark uninformed price schedule (the unconditional price schedule that would obtain if no investor were informed). The presence of informed investors may thus lead to lower prices compared with the case with symmetric ignorance even when there is good news. Finally, when n is very small (around $n = 0.02$), price schedules start overlapping in the sense that there are prices that 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 $n \rightarrow 0$. (See proposition 1.)

To facilitate a comparison of model outcomes with data moments, figure 5 plots overpayment and ITM shares for informed (black open circles) and uninformed (gray filled circles) investors as a function of the share of informed investors n . The right panel also shows the quantity-weighted average yield on the right axis.

Foreshadowing our quantitative results, we take as given that large bidders at an auction are informed, while smaller bidders are uninformed. Matching the empirical bidding patterns requires that both information

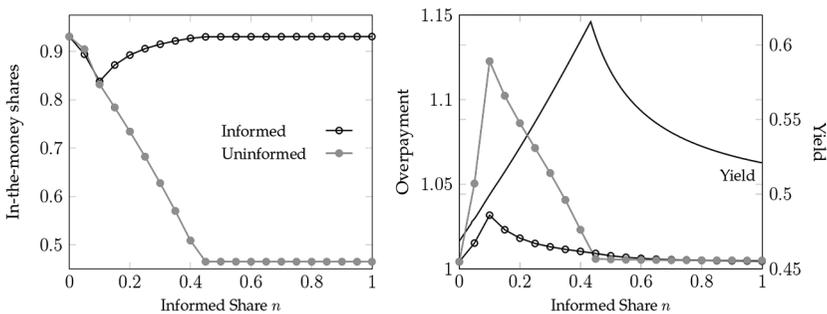


FIG. 5.—Impact of informed share on ITM shares, overpayment, and yields. A color version of this figure is available online.

types have similar overpayment, while informed investors have substantially larger ITM 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 investor types; however, the lack of uninformed participation in the high state implies sharp differences in ITM shares. The presence of supply shocks further allows us to capture the empirical fact that ITM shares are below one and that overpayment is above one on average for all investors.

Notably, average yields are particularly high when n is around 0.5, which 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.

C. The Effects of Secondary Markets

Like other sovereign bonds, CETES can be traded in secondary markets. In appendix B, we extend our model to include secondary markets and show that asymmetric information remains the only source of heterogeneity that can break the positive association between ITM shares and overpayment. Specifically, we extend proposition 4 and show that the winner's curse at auction may be stronger in the presence of secondary markets. This is because secondary-market prices respond to information revealed at auction, and uninformed investors face the risk that bonds bought at high prices at auction trade at low secondary-market prices. This price risk is amplified by the fact that informed dealers bid more aggressively because they are less exposed to default risk if they sell part of their portfolio in secondary markets and may capture arbitrage profits across markets. (We show below that CETES trade at a 3% premium above the primary-market price.)

V. Calibration

We now calibrate parameters to the data to 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. In a second step, we ask whether matching

the dynamics of marginal prices can inform us about the sort of information investors have access to.

To make use of the time-series dimension of our data, we consider an infinitely repeated version of our basic model. The government issues bonds in every period, and there are successive generations of 1-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.

In section II, we showed that CETES prices have high unconditional volatility but low conditional volatility. To account for this and to distinguish public and private news, we introduce common variation in beliefs by including in our information structure public regimes that 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 transition probability 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 $\{\theta, z\}$. Without loss of generality, we assume that default risk is lower on average in public regime 1 than in regime 2. Since investors can typically learn the public shock from past prices, we take as given that the current public regime is known at the start of every auction. The parameters that determine the stochastic process for default risk and bond supply are thus 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 characteristics. While we focus on information heterogeneity as the key driver of portfolio differences, we nevertheless require a model-consistent measure of investor size because there are sharp differences in bidding behavior between the largest bidder and the rest. 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 use investor mass to distinguish size.

It is difficult to obtain a precise measure of investors' average portfolio exposure to CETES bonds. As our benchmark, we 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 app. D). Because this is a rough measure, in appendix C we provide robustness checks for various measures of investor exposure to CETES default risk (nonzero recovery rates, lower values of D/W , higher risk aversion) and show that our results are not sensitive to these choices. Last, we assume that there are no other fundamental differences between investors.

Our theoretical results show that we can match empirical bidding patterns only if the largest investor is informed. We therefore take as given that there is a large bidder with mass n_{big} who is informed. This can be justified by assuming a fixed cost to information acquisition. We then calibrate n_{big} and the total share of informed investors $n \geq n_{\text{big}}$. This implies that some of the smaller investors may also be informed.

We fix the coefficient of relative risk aversion to be equal to one (log utility), which is low relative to the asset pricing literature. This implies that our quantitative results are due to the pricing protocol and asymmetric information rather than to investor preferences.

A. *Disciplining Moments*

We now discuss how parameters are pinned down by data moments. We focus throughout on the 28-day bond. This maturity is particularly suitable for our analysis of asymmetric information because the time horizon is long enough for news to materialize but short enough that investors need not forecast the far-off future.

We calibrate the volatility and level of default risk using the mean and standard deviation of marginal prices. These moments provide measures of the unconditional level and volatility of prices. Their values are 0.98 and 0.017, respectively. In the model, these moments are determined by the mean and variance of default probabilities across and within regimes.

The second set of moments are the coefficient $\beta_1 = 0.98$ from regression (1) and the associated $R^2 = 0.97$. These moments capture the persistence and predictability of prices from one auction to the next, and they 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 ITM share of the largest bidder is 0.84 on average (the largest bidder buys 84% of his submitted bids on average), while the average ITM share of the remaining investors is just 0.33. Additionally, all bidders' average overpayment is 1.001, which implies a quantity-weighted price that is 0.1% above the marginal price on average. Our model demonstrates that an informed bidder achieves ITM shares and overpayment of one 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 that prevent informed bidders from having all bids accepted in every state.

Our model also suggests that the remaining investors have low ITM 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. Hence, the difference $n - n_{\text{big}}$ is the mass of informed investors among the remaining bidders. Bidding data then allow us to bound the total share of informed investors. The largest bidder buys on average 38% of the bonds. Hence, we must have

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

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

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

As n is maximal when $\text{ITM}^U = 0$, it follows that $(n - n_{\text{big}})/(1 - n_{\text{big}}) 0.84 = 0.33$ and

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

An intermediate value of n ensures that the winner's curse is severe enough to discourage uninformed bids at high prices. Taken together, ψ_M and n_{big} are identified by ITM 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.

B. How Well Can the Model Match the Data?

We now assess the model's ability to match the data moments discussed above. We consider two approaches. The first approach is the baseline calibration described in section V.B.1, in which we choose parameters to minimize the sum of squared errors between the data targets and model-generated moments. This calibration is quantitatively successful in most dimensions but can match two important moments only qualitatively: the ITM shares of small bidders and the time-series predictability of prices.

The second approach is the peso problem calibration in section V.B.2, which shows that we can achieve a quantitative match on these dimensions

TABLE 3
CALIBRATED PARAMETERS

Common Parameter	Value (1)	Remaining Parameter (2)	Baseline (3)	Rare Disaster (4)
κ_{g1}	.001	κ_{g2}	.019	.02
κ_{b1}	.011	κ_{b2}	.029	.50
$f(g_1)$.65	$f(g_2)$.65	.95
ψ_M	1.3	n	.40	.40
ρ	.999	n_{big}	.22	.17

if we posit the existence of a “rare disaster” state with high default risk that is considered plausible by investors but did not materialize in our sample period. Importantly, 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. Based on these results, section V.B.3 provides an interpretation of the nature of asymmetric information that may be present in Mexican bond markets.

1. Baseline Calibration

We now describe our baseline calibration approach. To economize on the number of parameters to be calibrated, we impose some symmetry across public regimes by restricting (i) the persistence of regimes, (ii) their supply shock distribution, and (iii) their quality shock probabilities to be the same. Table 3 reports common parameters (fixed across both calibration approaches) and moment-matching parameters for the baseline calibration in columns 2 and 4, respectively.¹⁹ Data and simulated moments are in columns 1 and 2 of table 4, respectively.

We match well the mean and standard deviation of marginal prices, as well as the ITM share of the largest bidder and overpayment for the largest bidder and the rest. Qualitatively (but not quantitatively), we can also account for the ITM share of the rest of the investors (lower than for the largest bidder but much larger than in the data, 0.62 vs. 0.33) and the extent of predictability (positive but much smaller than in the data, 0.7 vs. 0.97).

The quantitative fit is imperfect along these dimensions because of the high unconditional volatility and low conditional volatility in the data. Since prices are also high on average, matching price moments requires a default risk process in which there is one regime with relatively low default risk (regime 1) and another with relatively high default risk (regime 2). Moreover, default risk cannot be too variable within each regime. The price schedules implied by these benchmark parameters, for each regime

¹⁹ Note in the table that a weekly persistence of $\rho = 0.999$ reflects an annual persistence of $0.999^{52} = 0.95$.

TABLE 4
CALIBRATION TARGETS: DATA VERSUS MODEL

Target	Data (1)	Baseline (2)	Peso Problem (3)
Mean price	.98	.98	.97
Standard deviation of price	.02	.02	.02
Regression β	.98	.84	.96
Regression R^2	.97	.70	.92
ITM of largest bidder	.84	.87	.88
ITM of other bidders	.33	.62	.41
ITM of uninformed bidders		.45	.16
Overpayment of largest bidder	1.001	1.001	1.002
Overpayment of other bidders	1.001	1.004	1.004
Bid share of largest bidder	.38	.38	.38

and each quality shock, are depicted in the upper left panel of figure 6. Both price schedules are relatively high in regime 1, and both are relatively low in regime 2.

In regime 1, the risk-free asset is a close substitute for bonds. Thus, uninformed investors do not bid at high prices in this regime despite a mild winner's curse. Yet averaging across regimes generates ITM shares for the uninformed that are too high. This is because the high risk premium in regime 2 makes it attractive for uninformed investors to participate despite the winner's curse. To further improve model fit, we would therefore require uninformed investors to submit fewer high-price bids in regime 2. However, this would force informed investors to absorb all of the (high) default risk in regime 2. Since informed investors are willing to do so only if prices are very low on average, the high implied risk premium weakens the winner's curse, inducing uninformed investors to again participate at high prices. This can be seen in the bottom left panel of figure 6: uninformed investors actually bid more on the high-quality schedule than on the low-quality schedule for low realizations of the supply shock.

Starting from the baseline calibration, uninformed bidders would thus achieve a lower ITM share only if they were much more concerned about the winner's curse in regime 2. As we discuss in the next section, it is possible to induce such behavior without strongly affecting average prices by allowing for an unlikely but severe worst-case scenario in regime 2 (a rare disaster).

2. Rare Disasters and Peso Problems

Our sample is a relatively tranquil period for Mexico: unconditional volatility is moderate, and there was no government debt crisis. This is in marked contrast to the previous 2 decades in which each saw such crises. Since our baseline calibration aims to match in-sample price data, the calibrated default risk process does not permit the possibility of a severely bad

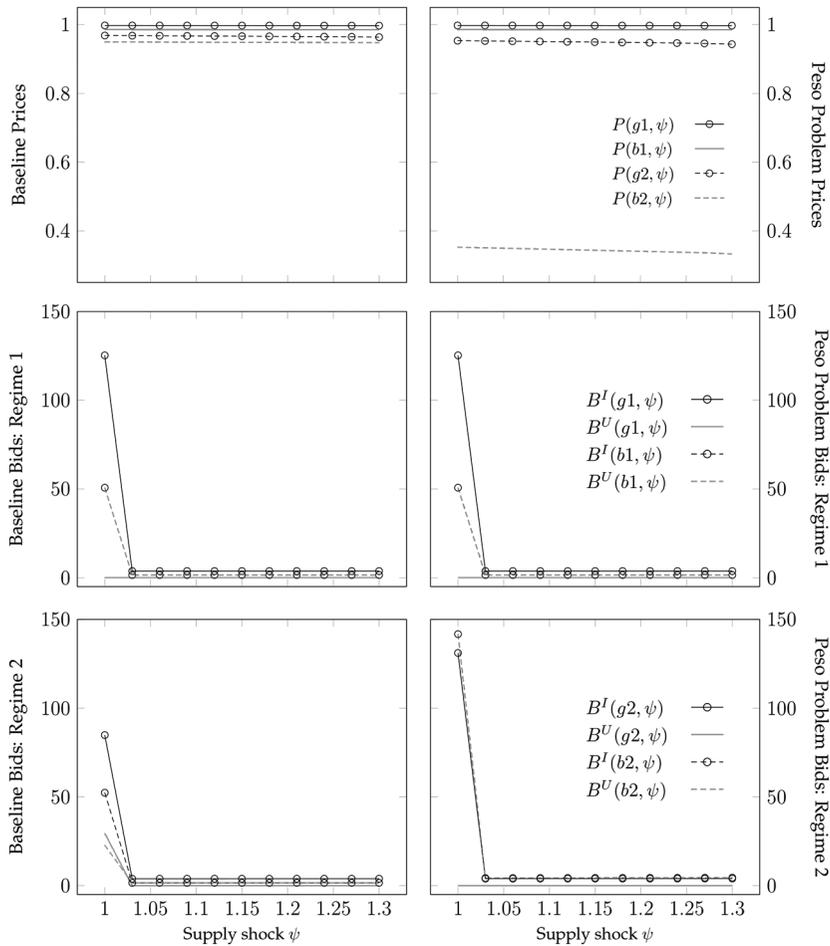


FIG. 6.—Prices and bids in baseline (left) and peso problem (right) calibrations. A color version of this figure is available online.

shock. However, investors may be wary of such an event even if it has not yet materialized. Going back to Milton Friedman’s explanation for the observed gap between US and Mexican deposit rates, this difference between beliefs and sample realizations has been referred to as the peso problem.²⁰

Starting from the baseline calibration, we introduce the risk of a rare disaster by raising the worst-case default probability κ_{12} that occurs in the bad

²⁰ 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 have not yet occurred (e.g., the rare disaster), so that their subjective probability distribution differs from the data-generating process that has generated a particular sample in historical data (see Rogoff 1980; Evans 1996).

state of regime 2 to 50% while lowering the probability of this state to 5% from 45%. Hence, this state now represents a severe crisis event whose probability is consistent with historical record. Specifically, the implied likelihood that the rare disaster does not occur within a decade is 31%, consistent with the fact that there were two crises in the 3 decades preceding our sample.²¹

Default probabilities in regime 1 are unchanged from the baseline calibration. This ensures that average prices are not too high while increasing the severity of the winner's curse for uninformed bidders in regime 2. We maintain the same share of informed investors as in the baseline calibration but reduce the mass of the largest investor to match the share of bids purchased. To account for the peso problem, we compute simulated moments by discarding sample paths in which a crisis occurs (or, equivalently, computing moments assuming that the probability of the good shock in regime 2 is equal to one, even though investors perceive a positive probability of a bad shock, here a disaster).

The right panels of figure 6 show prices and bids for the rare disaster parameters. The low price schedule in 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. The bottom right panel shows that the winner's curse discourages the uninformed from bidding at high prices in regime 2. The associated moments are reported in column 3 of table 4. The ITM share of the rest of investors declines substantially from 0.62 to 0.41 and now better matches the data.²² The rare disaster model also successfully matches dynamic moments and the predictability of prices: regression coefficient β has risen to 0.96, and the regression R^2 has risen to 0.92, both of which are quite close to the data.

3. An Interpretation on the Nature of Asymmetric Information

Our quantitative exploration provides some insights into the nature of asymmetric information. Specifically, our results suggest that investors are

²¹ The probability that no disaster occurs within one decade obeys the following recursion. Denote by π_j^i the probability that the bad state in public regime i does not occur j weeks from the end of the decade. Then

$$\begin{aligned}\pi_j^1 &= 0.999 \times \pi_{j-1}^1 + (1 - 0.999) \times [0.95 \times \pi_{j-1}^2], \\ \pi_j^2 &= 0.999 \times [0.95 \times \pi_{j-1}^2] + (1 - 0.999) \times \pi_{j-1}^1.\end{aligned}$$

The unconditional probability that no disaster occurs within one decade is then $(\pi_{520}^1 + \pi_{520}^2)/2 = 0.31$.

²² ITM shares are still not quite as low as in the data. Here it is important to recognize that we do not drop any bids in our sample. If we treated very low bids as "spoofs" that would not be ITM in any scenario, we would achieve an even closer fit of data moments.

relatively similar in their access to information about typical price movements (as captured by public regimes) because they can all access publicly available data on the country's finances, previous auction results, and 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 a rare disaster. This information is difficult to access for all investors in real time, as it depends on access to decision makers, internal political decisions, and other nonpublic information that some bidders can obtain through political connections.

C. *Information and Government Financing Costs*

We can now quantitatively evaluate the extent to which asymmetric information raises government financing costs. We compare quantity-weighted average yields and yield volatility in the calibrated model (where $n = 0.4$) to two counterfactual benchmarks: symmetric ignorance, where $n = 0$, and symmetric information, where $n = 1$. We continue to presume that no disaster occurs on path. Hence, the comparisons should be understood as measuring counterfactual yields in sample.

The results are in table 5. We find that the government would have benefited from the presence of more informed investors: yields fall by approximately 1.4 percentage points ($\approx 45\%$ of prevailing borrowing costs) when all investors are informed, while volatility falls by approximately half. This is because informed investors bid more aggressively for bonds with low default risk, and there are smaller price differences across public regimes when investors are aware that a crisis did not materialize in regime 2. Indeed, prices are unambiguously higher if $n = 1$ than if $n = 0.4$ because more investors (indeed, all investors) can participate on the low-price schedule without fear of the winner's curse.

Interestingly, the effect on mean yields is reversed in the symmetric ignorance counterfactual. Here yields rise slightly because all investors are now worried about a possible disaster, and positive news (i.e., the fact that the disaster will not materialize in a particular period) is not priced in. Volatility falls because prices do not respond to quality shocks in

TABLE 5
YIELDS IN CALIBRATION AND INFORMATION COUNTERFACTUALS

	Mean Yield (%)	Standard Deviation of Yields (%)
Calibration: $n = .4$	3.03	3.38
Symmetric information: $n = 1$	1.66	1.63
Symmetric ignorance: $n = 0$	3.39	2.79

either regime, but volatility falls by less than under symmetric information because the fear of a disaster drags down average prices in regime 2.

An important qualification to the conclusion that the government prefers more informed investors is that we compute counterfactuals only under the presumption that a disaster does not materialize in sample. However, it stands to reason that the presence of more informed investors might be costly to the government during a disaster because adverse shocks would be more rapidly impounded into prices. Owing to our focus on the peso problem, this trade-off is difficult to capture in our present model because quality shocks are transitory in our theory but crises are typically persistent in practice. We thus leave this alternative counterfactual exercise and a more careful modeling of disasters that occur on path for future work.²³

D. *The Value of Information*

We can measure the value of information in our calibrated model. Since investors have log preferences and differ only in information, the portfolio problem in (8) is homogeneous of degree one in wealth for each information type, and indirect utility can be written as $V^i = \hat{V}^i + \log(W)$, where \hat{V}^i represents indirect utility when wealth is normalized to one. The value of information is the proportional wealth adjustment factor a that makes an investor indifferent between being informed or not. This factor is defined by

$$\hat{V}^I + \log(W) = \hat{V}^U + \log(aW) \Rightarrow a = e^{[\hat{V}^I - \hat{V}^U]}.$$

We separately compute this factor for each public regime. In regime 1, we find that $a_1 = 0.9997$, which implies that the value of information is only 0.03% of wealth. This is expected because the mean and volatility of default risk is very low. In contrast, in regime 2 we find that $a_2 = 0.9903$, indicating that the value of information is considerably higher, roughly 1% of wealth. This is because regime 2 has the possibility of a rare disaster that generates a more severe winner's curse.

To understand these findings, recall that uninformed investors submit bids only on the low-quality schedule. Conditional on a bad realization of the quality shock, they thus obtain the same payoffs as informed investors, but conditional on a good-quality shock, only informed investors capture the surplus risk premium on inframarginal bond purchases. In regime 1, the default risk of the good bond is so low that the risk-free asset is almost a perfect substitute for the risky bond, and the uninformed investor does

²³ To capture persistent crises such as the Mexican tequila crisis, the model would need to allow for an additional persistent public regime that is on the path of play only after a bad-quality shock in regime 2.

not lose much from not buying those bonds. In regime 2, the default risk of the good bond is substantially higher than in regime 1 (0.02 vs. 0.001, from table 3) and the good state is substantially more likely (0.95 vs. 0.65); hence, uninformed investors lose more from not buying those bonds, and the value of information becomes substantially large.

VI. Reexamining the Data

In addition to fitting the data using a calibration exercise and performing counterfactuals with respect to the information environment, we can use our model to develop and test additional empirical predictions. This allows us to provide additional validation for our model in primary-market data and offers some insights into the link between primary and secondary markets.

A. Testable Predictions within Primary Markets

Our results thus far use the largest bidder as the benchmark for an informed investor. However, our calibration suggests that some of the smaller bidders are also informed. This interpretation is consistent with the data if ITM shares and overpayment of some smaller investors looked like they too were informed. Continuing to use number of bonds purchased as a measure of private incentives to become informed, we therefore examine whether the second- and third-largest bidders are relatively closer in their bidding behavior to those of the largest bidders than to those of the rest.

In line with this prediction, table 6 shows that ITM shares of the second- and third-largest bidder by bonds purchased lie between those of the largest and the rest of bidders and that there is no difference in average overpayment. (We show these statistics for 28-day CETES, which is the focus of our calibration; the same results hold for all maturities.)²⁴

An additional prediction of the model is that, because of the winner's curse, informed and uninformed investors bid differently on the high-price schedule but similarly on the low-price schedule. Hence, the model predicts that there should be state-contingent differences in realized portfolios. Specifically, uninformed investors should have particularly low ITM shares in good states, while informed investors' ITM shares are driven primarily by the supply shock and are thus less sensitive to the quality shock.

²⁴ Our model considers only two information types: informed and uninformed. The fact that ITM shares steadily decline in investor size suggests some gradation in the quality of the information among investors in the data; i.e., there may be highly informed, partially informed, and uninformed types.

TABLE 6
 BIDDING BEHAVIOR OF SECOND- AND THIRD-LARGEST BIDDERS

Statistic	Largest	Second-Largest	Third-Largest	Rest
Bidder's share of bonds sold	.38	.23	.13	.36
ITM share	.84	.68	.55	.28
Overpayment	1.001	1.001	1.001	1.001

To examine these predictions in our data, we need to take a stance on when the realized quality shock was high or low. This is not trivial, because prices are jointly affected by quality and supply shocks. In the 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 1 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 that we have a low-quality one. Table 7 documents the predicted pattern. Small investors' ITM shares are significantly higher in the bad state than in the good state, but the largest investors' ITM shares are very similar. This coherence between model and data provides strong additional evidence of the model's basic mechanism and the importance of asymmetric information.

B. Relation with Secondary Markets

The existence of a secondary market for CETES provides additional opportunities to measure the information content of auctions, the nature of this information, and the relevance of adverse selection for government financing. We show three sets of results. First, information revealed at auction affects same day's secondary-market prices, suggesting that auctions indeed contain information. Second, information revealed at auction also affects the primary-market price of bonds auctioned a week later, suggesting that such information is about bonds' fundamental quality. Third, we show that dealer banks earn rents by selling at a markup in the secondary market and that this markup increases when adverse selection at auction is likely to be

TABLE 7
 ITM SHARES FOR DIFFERENT QUALITY SHOCKS

	Positive Pricing Error	Negative Pricing Error
Largest bidder's ITM	.86	.85
Remaining bidders' ITM	.24	.35

more severe (as indicated by greater differences in ITM shares across investors). Since markups are not captured by the government, increases in adverse selection are detrimental for governments' financing costs.

1. Auction News and Secondary-Market Prices

We have argued that some investors are more informed than others about the fundamental value of bonds. This suggests that information revealed through marginal auction prices should also be reflected in subsequent secondary-market prices. We provide evidence for this channel by showing that (i) secondary prices are significantly more volatile on auction days compared with nonauction days and (ii) these changes can be attributed to unforeseen movements in primary prices.

We use secondary yields at each business day closing (2 p.m.) as reported by the central bank of Mexico, which obtains the information from Proveedor Integral de Precios (PIP) and Valuación Operativa y Referencias de Mercado (Valmer). In what follows, we present results for 28-day CETES, but similar results are obtained for the other maturities. Our measure of secondary price volatility is the average squared 1-day change of secondary prices. We find that it is 0.0036 on days without an auction, and it more than doubles (and is statistically different at a 99% confidence level) to 0.0077 on auction days.²⁵

Increased secondary-market price volatility on auction days alone is not sufficient to identify information revelation at auction. To better identify the information channel, we obtain the elasticity of secondary prices with respect to unexpected changes in auction prices by regressing the unexpected log change of secondary prices ($\Delta \log \text{Sec}_i$) on the unexpected log change on primary prices ($\Delta \log \text{Prim}_i$) on auction days t . (As we discuss in detail in app. E, we measure the unexpected change as the realized price minus the predicted price from a predictive regression that includes observable past prices.)

Column 1 of table 8 shows that this elasticity is 0.74 for 28-day maturities and highly statistically significant, implying that a 1% unexpected increase in primary prices is associated with a 0.74% increase in the secondary-market price.

To get a sense for how quickly information contained in auction prices is impounded into secondary-market prices, we also regress the unexpected log changes in secondary prices 1 and 2 days after the auction on the

²⁵ Auction days represent roughly one-quarter of all observations, as auctions are weekly. We drop observations when auctions happen on Mondays to avoid potentially large accumulations of information on weekends. However, results are almost identical if we also include these days.

TABLE 8
ELASTICITY OF SECONDARY-MARKET PRICES TO INFORMATION RE-
LEASED AT AUCTION

Dependent Variable	$\Delta \log \text{Sec}_t$ (1)	$\Delta \log \text{Sec}_{t+1}$ (2)	$\Delta \log \text{Sec}_{t+2}$ (3)
$\Delta \log \text{Prim}_t$.742*** (.021)	.017 (.022)	.015 (.027)
Constant	-.004*** (.0004)	.001** (.0004)	.001* (.0005)
Observations	707	707	707
R^2	.639	.001	.000

* $p < .1$.

** $p < .05$.

*** $p < .01$.

unexpected log change in primary prices on the auction day. As shown in columns 2 and 3, these effects are not statistically significant, implying an immediate pass-through of auction information to secondary markets.

2. Auction News and Subsequent Primary-Market Prices

Even though information at auction is quickly reflected in secondary prices, is the information about fundamentals that affect underlying default risk? If so, information revealed through marginal auction prices should also be reflected in subsequent primary-market prices. To capture the nature of information, we decompose primary prices, P_b , into an expected price at auction, \hat{P}_t (based on previous secondary-market prices) and an orthogonal unexpected change in this price, ε_{P_t} . We then regress primary prices on these two components of the previous week's prices. The results for 28- and 91-day bonds (which both were sold in consistent weekly auctions in our sample) are shown in table 9. The details about the empirical approach are in appendix E.

TABLE 9
PERSISTENT EFFECTS OF INFORMATION REVEALED AT AUCTION

Dependent Variable: P_t	28-Day	91-Day
Expected primary price: \hat{P}_{t-1}	.997*** (.003)	.998*** (.003)
Unexpected component: $\varepsilon_{P_{t-1}}$	1.120*** (.056)	1.313*** (.086)
Constant	.018 (.017)	.015 (.016)
Observations	706	706
R^2	.994	.995

*** $p < .01$.

The highly statistically significant coefficient for the unexpected change in primary prices strongly suggests that information revealed at auction has persistent effects and is related to fundamental bond values. This interpretation is reinforced by noticing that the unexpected changes ε_{p_i} of 28- and 91-day bonds have a statistically significant positive correlation of 0.81.

3. Arbitrage Gains and ITM Shares

Our theoretical analysis of secondary markets in appendix B suggests that informed market makers (or dealer banks) can capture arbitrage gains by selling bonds at a markup in the secondary market. This arbitrage gain is not exploited by uninformed traders who face the winner's curse at auction.

To verify this prediction in the data, we measure arbitrage profits as the spread between the primary-market CETES yield (reported by the central bank at 11:30 a.m. on auction days) and the secondary-market yield of the same bonds at market close (at 2 p.m. on auction days).²⁶ These auction-day arbitrage gains are shown in figure 7. As predicted, arbitrage gains are almost generally positive, but there is a sharp difference across two periods: the pre-2009 period with a high spread for 28-, 91-, and 182-day CETES (and low prices in both the secondary and the primary markets) and the post-2009 period with a low spread (and higher prices in both markets). These differences are large in terms of magnitude: while the spread for 28-day CETES averaged 26 basis points in the first period, it was around 11 basis points in the second period. The only maturity that does not follow this pattern is the longest maturity of 364-day CETES, for which there is little difference across periods.

Our theoretical model offers a simple perspective on the difference in arbitrage gains across these periods. Uninformed investors are willing to pay high markups only if participating in the auction is very costly. This is the case when the winner's curse is severe and uninformed investors do not bid at high prices. This lowers primary-market prices and depresses ITM shares for the uninformed.

We find evidence consistent with this mechanism. Table 10 shows ITM shares for the largest bidder and the remaining bidders for all maturities before and after 2009. While ITM shares for the largest bidder are very similar in both subperiods (suggesting similar supply shocks in both periods), ITM shares for the rest rose significantly for each maturity in the second subperiod, suggesting less adverse selection.

²⁶ This difference between primary- and secondary-market yields at auction days is the empirical counterpart to the object $P(\theta) - \hat{P}(\theta)$ in the model with secondary markets in app. B.

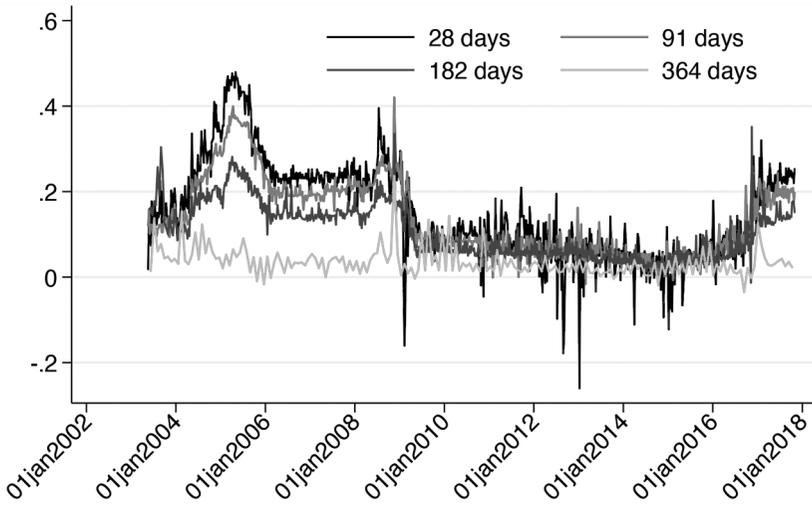


FIG. 7.—Difference between the annualized nominal yield in primary markets (reported the day of the corresponding auction) and the annualized nominal yield in secondary markets (at closing on the day of the auction). Source: Bank of Mexico. A color version of this figure is available online.

Two additional observations are consistent with our theory. First, ITM shares for the rest were higher after 2009 only for the maturities for which the spread declined (and the ITM shares increased more for the maturities for which the spread declined more) but not for the 364-day CETES (in fact, they fell slightly). Second, post-2009 primary yields were not only smaller in levels but also less volatile, consistent with a weaker winner's curse. Second, the magnitude of the effects are consistent with our quantitative results. In our counterfactual in section V.C, we show that eliminating adverse selection (by making all investors informed) would translate to a 1.43 percentage point yield reduction, while the observed reduction (not elimination) of adverse selection between the two subperiods generated a reduction of 0.15%.

TABLE 10
ITM SHARES BEFORE AND AFTER 2009

MATURITY	BEFORE 2009		AFTER 2009	
	Largest	Rest	Largest	Rest
28-day	.84	.27	.85	.38
91-day	.80	.18	.82	.28
182-day	.81	.20	.79	.24
364-day	.82	.28	.80	.25

The relation between the extent of adverse selection (as captured by differences in ITM shares between the largest investors and the rest) and the spread between primary and secondary markets shows that the gap between an investor's willingness to pay and the government's funding costs is determined by the extent of asymmetric information among auction participants.²⁷ Hence, the primary/secondary-market spread can serve as a measure of the winner's curse at auction.

VII. Final Remarks

In this paper, we have provided evidence (for an emerging economy) that asymmetric access to information about the fundamental value of sovereign bonds has a significant impact on primary debt auctions. We have documented the existence, nature, and quantitative implications of asymmetric information using a number of complementary strategies. First, we compiled a unique data set of weekly discriminatory price auctions of Mexican domestically denominated CETES bonds between 2001 and 2017 with bidder identifiers. We then compared the observed bidding patterns with equilibrium outcomes from a model that allows for rich heterogeneity among investors. We showed that asymmetric information is the only source of heterogeneity that can break the positive correlation between the frequency and the average price of purchases, which is the hallmark of discriminatory price auctions but is at odds with our data.

Besides this theory-guided qualitative evidence of asymmetric information, we calibrated our model to show that the bidding behavior within auctions and the time-series properties of prices across auctions are jointly consistent with the theory when investors are symmetrically informed about economic characteristics that are public but asymmetrically informed about specific details that may be relevant for estimating default risk and identifying rare disasters.

We also developed theoretical predictions regarding linkages between primary and secondary markets and provided additional evidence that auctions reveal information (primary-market price surprises are immediately incorporated into secondary-market prices), that information is about fundamental bond values (primary-market price surprises affect subsequent primary-market prices), and that markups in secondary markets are related to the degree of asymmetric information in primary markets.

Having made a case for the existence of asymmetric information and its relevance for sovereign bond yields, a natural next step is understanding its implications more generally. In Cole, Neuhann, and Ordoñez (2020), for instance, we explore the role of information asymmetry for

²⁷ In 2019, Mexico sold US\$80 billion of CETES. A 15-bps yield reduction implies that reducing adverse selection across both subperiods would have generated savings of US\$120 million.

cross-country spillovers. By allowing for endogenous information acquisition and extending the setting to multiple countries, we show that multiple information regimes may coexist (some with asymmetric information, some without) and that a country may suffer from a shock in an unrelated country through endogenous asymmetric information. Our tractable model also allows for many other extensions, such as government incentives to disclose information or to manage prevailing information heterogeneity by using different auction protocols. We leave these extensions for future research.

Appendix A

Institutional Details of the Primary and Secondary Markets for Mexican Government Securities

This section describes the structure and functioning of Mexico's primary market of government securities. In Mexico, the federal government issues and places four different instruments in the local debt market, with the Bank of Mexico acting as a financial broker in their placement:

1. CETES: federal Treasury certificates (Certificados de la Tesorería de la Federación) are the oldest debt instruments (first issued in January 1978) and a fundamental pillar in Mexico's money market. These instruments belong to the family of zero-coupon bonds. Currently, their maximum term is 1 year, although previously they were issued for up to 2 years. In 2020, CETES corresponded to 17% of government securities outstanding.
2. Bondes: federal government development bonds (bonos de desarrollo del gobierno federal) are floating-rate government securities. They pay monthly interest in pesos, which compound on a daily basis. Currently, these bonds are traded at 3-, 5-, and 7-year terms. In 2020, bondes and bondes D corresponded to 19% of government securities outstanding.
3. Bonos: fixed-rate federal government development bonds (bonos de desarrollo del gobierno federal con tasa de interés fija) were first issued in January 2000. Today, they are issued and placed at 3, 5, 10, 20, and 30 years. Bonos pay interest every 6 months, and their interest rate is determined on the issue date. In 2020, bonos corresponded to 40% of government securities outstanding.
4. Udibonos: federal government development bonds denominated in investment units (bonos de desarrollo del gobierno federal denominados en unidades de inversión) were created in 1996 and are inflation-hedged instruments. In 2020, udibonos corresponded to 18% of government securities outstanding.

In the period we consider (2001–18), CETES and bondes D were auctioned using discriminatory pricing and bonos and udibonos were auctioned using uniform pricing. The competitive bidding process can be summarized in four steps: (i) offering (previous Friday at 12 p.m.), (ii) bidding (Tuesday from

10 to 11 a.m.), (iii) result disclosure (Tuesday at 11:30 a.m.), and (iv) settlement (2 days later, on Thursday).

A1. *Market Makers Program*

In 2000, the Ministry of Finance (SHCP) created the Market Makers program to foster the development of the fixed-rate government securities market. Market makers are credit institutions and brokerage firms appointed by the Ministry of Finance that must present bids at competitive prices in each primary auction of securities and must also permanently quote purchase (bid) and sale (offer) prices in the secondary market to provide liquidity.

This program has reinforced the primary placement of securities. Since its implementation, no auction of CETES, bonos, and udibonos has been declared void. During its existence, 13 institutions have been appointed as market makers, including two brokerage firms (Invex and Merrill Lynch). A maximum number of 10 market makers has been achieved (from September 2007 to February 2008) and a minimum number of five (from May 2001 to July 2002). In 2014, for instance, the market makers were Banamex, Bank of America, Barclays, BBVA Bancomer, Deutsche Bank, HSBC, JP Morgan, and Santander.

Criteria to become a market maker.—To be designated as a market maker (of CETES and bonos), the main requirement is to participate and maintain a significant and diversified level of trading in these securities. To this effect, the government constructs an index as follows:²⁸

$$\text{Index of market maker} = \text{Activity index} + \text{Incentives} - \text{Penalties},$$

where

$$\text{Activity index} = 3\% \text{ CETES component} + 97\% \text{ Bonos component}.$$

For CETES, only the participation in auctions matters:

$$\text{CETES component} = 100\% \text{ Auction participation},$$

where the variable Primary is the amount assigned to this particular market maker in the primary auctions as a percentage of the total value of the auction during that period, while for bonds, diversity in the participation in secondary markets also matters. More specifically,

$$\begin{aligned} \text{Bonos component} &= 25\% \text{ Auction participation} + 40\% \text{ Clients} \\ &\quad + 35\% \text{ Interbank through brokers}, \end{aligned}$$

where the variable Clients measures transactions where the counterparty is not a bank or a brokerage house and Interbank through brokers refers to market operations done through authorized trading mechanisms in which the counterparties can be only banks or brokerage houses.

Finally, Incentives and Penalties points can be gained or lost depending on other factors, such as diversification in different CETES and bonos segments and participation in derivative markets. For example, if the market maker operations

²⁸ Numbers are calculated each month using data for the previous 6 months.

are distributed on at least half of the current/still-active CETES or bond emissions, they gain points, and if their operations are concentrated on less than 40% of the current/still-active emissions, they lose points.

To be designated a market maker, it is necessary to obtain an index of at least 7% for at least three consecutive months. To remain a market maker, that grade should be maintained every month. Market makers are designated on March, June, September, and December.

Duties of market makers.—Market makers have obligations in both primary and secondary markets.

1. Duties on primary markets: market makers are requested to present a minimum amount of bids in each primary auction for the type of securities they are market makers of. This minimum amount is the lowest amount resulting from 20% of the amount offered by the federal government or one divided by the number of market makers. Market makers cannot submit more than 60% of the amount offered. Hence, by having five market makers for each instrument, the bids are guaranteed to be at least the security offer at auction.
2. Duties on secondary markets: market makers must create and develop this market by quoting bid and ask prices daily, between 9 a.m. and 1 p.m., for all relevant ranges in those fixed-rate securities they operate and for a minimum nominal amount of MXN\$20 million.

Rights of market makers.—Market makers have rights in terms of transactions that they can carry out with the Ministry of Finance and Bank of Mexico. These include the following:

1. Exclusive participation in syndicated placements.
2. Access to securities lending through the central bank facility. This facility guarantees that market makers can satisfy the demand for a certain security even if they have none in their inventories at the date when they sell it.
3. Additional purchase of securities by exercising an option for the amount originally placed in the weekly primary auction. This option consists of buying the next working day after the primary auction takes place up to an additional 25% of the amount placed at the same allocation price.

A2. *Secondary Market*

A survey conducted by the Emerging Markets Traders Association during 2020 revealed that Mexican instruments were the most frequently traded local markets debt, at US\$625 billion, 17% of overall volumes in the world. The data obtained from the Bank of Mexico show that approximately 50% of daily trading volume consists of bonds and around 35% of CETES. Volume shows not only the depth of secondary markets for Mexican government securities but also spreads between purchase (bid) and sale (offer) prices, which come close to those of developed economies.

To induce a fair valuation of financial instruments, a valuation committee was established in the mid-1990s, which comprised both authorities and members of

the financial sector and established the estimation guidelines and criteria for price valuation. As a result, there are two price vendors in Mexico: PIP and Valmer. These price vendors in markets may facilitate the development of local market reference indexes. For example, PIP has two government bond indexes, while Valmer has six indexes referenced to bonos. Regarding CETES, PIP has seven indexes and Valmer has six.

The Mexican bond market is operated either through telephone services or by brokers. There are five brokers in the Mexican market: Enlace (since 1993), Remate Lince and Eurobrokers (both since 1996), Servicios de Integración Financiera (since 1998), and Tradition (incorporated in 2004).

According to the definition of the International Monetary Fund, the Mexican market may be classified as a periodic market, with working hours from 7 a.m. to 2:30 p.m. Most transactions take place during these hours, as compared with most developed-country markets, which operate for nearly 24 hours.

The Mexican government introduced the CETES Directo program on November 26, 2010, which was a successful effort to extend savings and investment on government securities by small and medium retail investors. This scheme is already used by countries such as the United States (Treasury Direct), Brazil (Tesouro Direto), and Spain (Tesoro Público), among others. The program gives small investors the opportunity to buy government securities online at low cost and without intermediaries. The price at which government securities can be purchased is linked to auction results in the previous week.

Appendix B

Secondary Markets Exacerbate Adverse Selection in Primary Markets

Here we consider a variant of our model that includes secondary markets. Motivated by the Market Makers program in Mexico, we assume that there is a mass one of market makers. These are primary dealers whose duties are to participate in primary auctions and then sell a fraction α of their bond inventory to a mass one of regular investors in secondary markets. We assume that the wealth of market makers and regular investors is W and \tilde{W} , respectively, and that both have the same utility function. Even though in Mexico there are a few large primary dealers, we maintain our assumption that they compete when bidding, then solving the same portfolio optimization problem as atomistic investors in our benchmark. We also assume that secondary markets are centralized. However, we leave open the possibility that they exploit their size to affect the supply of bonds available for sale in secondary markets.

The timing we consider is that secondary markets open after an auction concludes. Since the outcome of an auction is announced immediately after the auction concludes (usually at 11:30 a.m. on Tuesdays in Mexico), regular investors can infer the state θ when submitting orders. We denote by hats the objects in secondary markets: $\hat{B}(\theta)$ denotes regular investors' bond demand, and $\hat{P}(\theta)$ denotes secondary prices in each state. This extension is a slight change to the definition of the investor's problem in (9) and obtains by introducing effective prices $\tilde{P}(s) = P(\tilde{s}) - \alpha\hat{P}(s)$ in equation (3) and multiplying $B_j^i(s)$ by $(1 - \alpha)$ in equation (8).

Even though this simple extension opens up a number of interesting questions, we focus on informational implications in primary markets. Specifically, we show in the next proposition that the existence of secondary markets exacerbates adverse selection in primary markets and weakens the necessary and sufficient conditions for the existence of an equilibrium in which uninformed investors do not bid on the high price schedule. To facilitate the comparison with the benchmark without secondary markets, we follow the strategy of proposition 4, which assumes no supply shock and log preferences. Hence, there are only two possible prices $P(g)$ and $P(b)$ at auction, one for each possible quality shock.

PROPOSITION 5. Secondary markets exacerbate the extent of adverse selection in primary markets that breaks the positive association between investors' ITM shares and overpayment.

Proof. We solve backwards, starting from secondary markets. As regular investors know the state θ when submitting orders, their bond demand is

$$\hat{B}(\theta) = \frac{\hat{W}(1 - \kappa(\theta) - \hat{P}(\theta))}{\hat{P}(\theta)(1 - \hat{P}(\theta))}.$$

Since we focus on an equilibrium in which only informed investors bid at $P(g)$, the supply of bonds in secondary markets is $\alpha n B^g(g)$ in the good state and $\alpha[nB^g(b) + (1 - n)B^g(b)] = \alpha B^g(b)$ in the bad state. From market clearing, secondary-market prices are

$$\hat{P}(g) = 1 - \frac{\kappa(g)\hat{W}}{\hat{W} - n\alpha B^g(g)} \quad \text{and} \quad \hat{P}(b) = 1 - \frac{\kappa(b)\hat{W}}{\hat{W} - \alpha B^g(b)},$$

which are increasing in \hat{W} and decreasing in $\kappa(\theta)$ and in the elements that increase supply: n in the good state, the fraction of bonds sold by dealers, α , and the allocations assigned in each state to informed dealers, $B^g(\theta)$. Notice also that $\hat{P}(g) > \hat{P}(b)$, for two reasons: (i) a lower default probability and (ii) a lower supply of bonds.

Turning to primary markets, informed dealers' bids are

$$B^g(\theta) = \frac{W((1 - \alpha)(1 - \kappa(\theta)) - \tilde{P}(\theta))}{\tilde{P}(\theta)((1 - \alpha) - \tilde{P}(\theta))},$$

where $\tilde{P}(\theta) = P(\theta) - \alpha\hat{P}(\theta)$ and primary prices $P(\theta)$ come from auction clearing in each state, $nP(g)B^g(g) = D$ and $P(b)B^g(b) = D$. Even though primary prices do not have a closed-form solution, even in this simplified setting, we can show two properties: primary-market prices are increasing in (i) the extent of dealers' participation on secondary markets, α , and (ii) dealers' arbitrage gains from selling in secondary markets, $P(\theta) - \hat{P}(\theta)$.

i). Primary-market prices are increasing in α : we focus on the case without arbitrage gains, $P(\theta) = \hat{P}(\theta)$, which implies that $\tilde{P}(\theta) = (1 - \alpha)P(\theta)$. In this case,

$$B^g(\theta)|_{P(\theta)=\hat{P}(\theta)} = \frac{1}{(1 - \alpha)} \frac{W(1 - \kappa(\theta) - P(\theta))}{P(\theta)(1 - P(\theta))} > B^g(\theta)|_{\alpha=0},$$

where $B^i(\theta)_{|\alpha=0}$ denotes the informed investors' demand for bonds in the absence of secondary markets, as in equation (15). Intuitively, dealers are less exposed to the government supply as they unload to secondary markets a fraction of their purchases (at no premium). This effect operates as reducing the supply in primary markets to $(1 - \alpha)D$, which we showed increases $P(\theta)$.

ii). Primary prices increase with $P(\theta) - \hat{P}(\theta)$: let us focus on the extreme case in which regular investors are close to risk neutral or are very wealthy, such that $\hat{P}(\theta) = 1 - \kappa(\theta)$, and then $\tilde{P}(\theta) = P(\theta) - \alpha(1 - \kappa(\theta))$. In this case,

$$B^i(\theta)_{|\hat{P}(\theta)=1-\kappa(\theta)} = \frac{W(1 - \kappa(\theta) - P(\theta))}{[P(\theta) - \alpha(1 - \kappa(\theta))][1 - P(\theta) - \alpha\kappa(\theta)]} > B^i(\theta)_{|\alpha=0}.$$

Intuitively, arbitrage gains give more incentives to dealers to demand bonds, raising the corresponding primary-market prices.

As in proposition 4, we compute the sufficient condition for the optimality of $B^U(g) = 0$, which is that uninformed bidding at $P(g)$ earns negative expected returns if the government is expected to default with probability $\bar{\kappa}$. This is now the relevant measure of both default and trading risk, as any uninformed bids at $P(g)$ would be executed in every state. If $\theta = b$, uninformed would overpay for low-quality bonds, holding a fraction $(1 - \alpha)$ and selling a fraction α at low prices $\hat{P}(b)$ in secondary markets. Expected returns for the uninformed if bidding at $P(g)$ are

$$(1 - \bar{\kappa})(1 - P(g)) - \bar{\kappa}P(g) - \alpha[1 - \bar{\kappa}] - \mathbb{E}(\hat{P}),$$

with $\mathbb{E}(\hat{P}) = f(g)\hat{P}(g) + f(b)\hat{P}(b)$, which is negative when

$$1 - \bar{\kappa} < P(g) + \alpha[1 - \bar{\kappa}] - \mathbb{E}(\hat{P}).$$

This sufficient condition is more likely to hold in the presence of secondary markets (as compared with the sufficient condition in proposition 4) for two reasons. First, as we discussed above, $P(g)$ is higher in the presence of secondary markets and thus adverse selection is stronger. Second, the uninformed sell a fraction of their purchases on secondary markets at a price below the fair value.

Appendix C

Robustness

We undertake several robustness experiments aimed at examining how changes to our calibration would affect our results—in particular, our stylized facts about bidding patterns. All experiments involve reducing the risk of default that investors face and hence have somewhat similar aspects. In a first experiment, we relax the assumption that default is complete, and instead we assume that the recovery rate upon default is 50%. In the second experiment, we reduce the level of funds being raised by the government relative to the wealth of investors, by reducing D/W dramatically from 0.2 to 0.04. In these experiments, we did not change any of the other parameters of the model from our benchmark calibration to clearly isolate the impact of these changes on the results. The results are presented in table C1, which is an extension of table 4.

TABLE C1
ROBUSTNESS EXPERIMENTS

Target	Data	PP Model	Partial Default	Lower D/W	Lower D/W , Higher CRRA
Mean price	.98	.97	.98	.99	.99
Standard price	.02	.02	.01	.01	.01
Regression β	.98	.96	.96	.87	.85
Regression R^2	.97	.92	.91	.75	.73
LB ITM share	.84	.88	.88	.88	.88
Rest ITM share	.33	.41	.48	.44	.44
Overpay LB	1.001	1.002	1.0004	1.0001	1.0002
Overpay rest	1.001	1.004	1.0007	1.0002	1.0004
Share LB	.38	.38	.38	.38	.38

NOTE.—LB = largest bidder; PP = peso problem.

As one would expect, reducing the risk exposure (either by reducing the losses upon default or by reducing the amount of debt purchased) raises the average marginal price and lowers its variance. While ITM and average-overpayment results are largely unchanged, though the ITM share of the rest rose slightly, the regression coefficients of the second experiment are quite different, lowering both the coefficient on the lagged marginal price and the explanatory power of the regression.

The reason that it is more difficult to match the regression coefficients when investors face less risk is that the volatility of prices falls; hence, a single lag of the price is now less informative on the margin.

The reason there is little impact on bidding patterns (this is ITM and overpayment) is that adverse selection remains strong even when risk exposure declines. More specifically, even though average prices increase, the optimal bidding strategy of uninformed bidding is still characterized by not bidding on the good-state schedule. Intuitively, as the price in the good schedule approaches one, the uninformed do not miss much from avoiding bidding for good-quality bonds (close substitutes of risk-free assets) while can still overpay in case of buying bad-quality bonds.

To examine the extent to which increasing risk aversion could offset the impact of lowering the amount of debt coming due, D/W , on the regression coefficient and R^2 , we did a final third experiment in which we also increased the relative risk aversion of our investors from 1 (log) to 2. This had only a modest impact in terms of increasing the pricing of risk, which is consistent with the myriad results in asset pricing that point to very high degrees of risk aversion being consistent with aggregate consumption risk and various risk premia in the data.

Appendix D

Stability of Mexico's Fiscal Situation during the Sample Period

2001–17

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 quarterly revenues raised by auctioning CETES of all maturities in Mexico (in real terms) and plot it in figure D1 as a fraction of quarterly real GDP. As can be seen, every quarter, Mexico raises (or rolls over) around 5% of quarterly 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 precrisis levels in 2010. This stability comes from both figures growing on average 0.5% per quarter in real terms over this sample period.

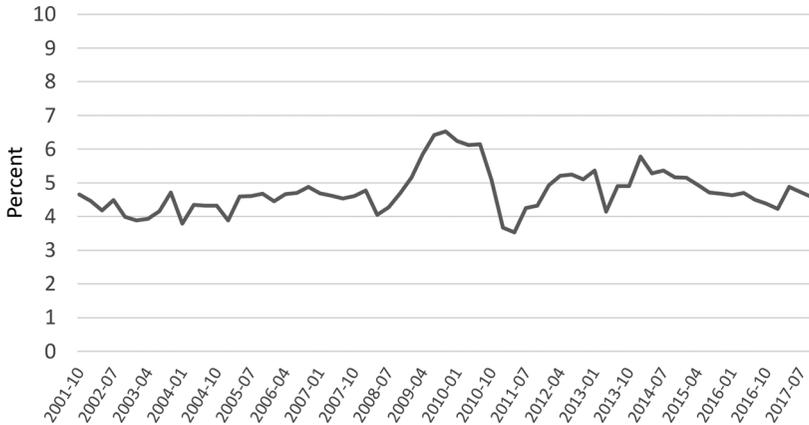


FIG. D1.—Real CETES revenue (in all maturities) as a fraction of real GDP. A color version of this figure is available online.

Appendix E

Empirical Strategy: Relating Primary and Secondary Prices

E1. Empirical Strategy (Sec. VI.B.1)

Here we explain how we construct unexpected changes in secondary and primary prices. To predict primary-market prices, we estimate

$$P_t = \alpha_p + \beta_p P_{t-7} + \sum_{k=1}^4 \gamma_{p,k} S_{t-k} + \varepsilon_{p,t}, \quad (\text{E1})$$

where P_t represents the primary price when the auction is performed at period t , P_{t-7} represents the primary price of the previous week's auction, and S_{t-k} represents the secondary prices in the four business days before the auction. This implies that there are as many observations as auctions.

To predict secondary-market prices, we estimate

$$S_t = \alpha_s + \sum_{k=1}^4 \gamma_{s,k} S_{t-k} + \varepsilon_{s,t}. \quad (\text{E2})$$

We do not include primary-market prices so that we can later compute elasticities of unexpected changes in secondary prices to unexpected changes in primary auctions. Notice also that this specification allows for more observations and statistical power, since secondary-market prices are observed daily (except weekends).

Finally, we compute the unexpected changes as follows:

$$\Delta \log \text{Prim}_t = \log(P_t) - \log(\hat{P}_t) \quad \text{and} \quad \Delta \log \text{Sec}_t = \log(S_t) - \log(\hat{S}_t),$$

where \hat{P}_t and \hat{S}_t represent the predicted primary and secondary prices at time t from equations (E1) and (E2), respectively.

E2. Empirical Strategy (Sec. VI.B.2)

Here we construct expected and unexpected changes in primary prices running the regression (E1) but without previous primary prices as an independent variable. This implies that we can decompose each primary price between the predicted and unexpected prices as $P_t = \hat{P}_t + \varepsilon_{p_t}$, where \hat{P}_t represents the prediction based on the coefficients obtained from the regression. We then take the primary prices and regress them on the two components of the previous primary price for the same maturity—that is,

$$P_t = \gamma_0 + \gamma_1 \hat{P}_{t-1} + \gamma_2 \varepsilon_{p_{t-1}} + \varepsilon_t. \quad (\text{E3})$$

References

- Aguiar, Mark, and Manuel Amador. 2013. "Sovereign Debt: A Review." Working Paper no. 19388, NBER, Cambridge, MA.
- Aguiar, Mark, Manuel Amador, Hugo Hopenhayn, and Iván Werning. 2019. "Take the Short Route: Equilibrium Default and Debt Maturity." *Econometrica* 87:423–62.
- Aguiar, Mark, Satyajit Chatterjee, Harold Cole, and Zachary Stangebye. 2016. "Quantitative Models of Sovereign Debt Crises." In *Handbook of Macroeconomics*, vol. 2, edited by John B. Taylor and Harald Uhlig, 1697–755. Amsterdam: North-Holland.
- Aguiar, Mark, and Gita Gopinath. 2006. "Defaultable Debt, Interest Rates and the Current Account." *J. Internat. Econ.* 69 (1): 64–83.
- Arellano, Cristina. 2008. "Default Risk and Income Fluctuations in Emerging Markets." *A.E.R.* 98 (3): 690–712.
- Barro, Robert J., and José F. Ursúa. 2012. "Rare Macroeconomic Disasters." *Ann. Rev. Econ.* 4 (1): 83–109.
- Biais, Bruno, Peter Bossaerts, and Jean-Charles Rochet. 2002. "An Optimal IPO Mechanism." *Rev. Econ. Studies* 69:117–46.
- Bocola, Luigi, and Alessandro Dovis. 2019. "Self-Fulfilling Debt Crises: A Quantitative Analysis." *A.E.R.* 109:4343–77.
- Boyarchenko, Nina, David Lucca, and Laura Veldkamp. 2021. "Taking Orders and Taking Notes: Dealer Information Sharing in Treasury Markets." *J.P.E.* 129 (2): 607–45.
- Chatterjee, Satyajit, and Dean Corbae. 2007. "On the Aggregate Welfare Cost of Great Depression Unemployment." *J. Monetary Econ.* 54 (6): 1529–44.

- Chatterjee, Satyajit, and Burcu Eyigungor. 2012. "Maturity, Indebtedness, and Default Risk." *A.E.R.* 102 (6): 2674–99.
- Chaumont, Gaston. 2020. "Sovereign Debt, Default Risk, and the Liquidity of Government Bonds." Working paper, Univ. Rochester, Rochester, NY.
- Chun, John H. 1995. "Post-modern Sovereign Debt Crisis: Did Mexico Need an International Bankruptcy Forum?" *Fordham Law Rev.* 64 (6): 2647–700.
- Cole, Harold, Daniel Neuhann, and Guillermo Ordoñez. 2020. "Information Spillovers in Sovereign Debt Markets." Working paper, Univ. Pennsylvania, Philadelphia.
- Eaton, Jonathan, and Mark Gersovitz. 1981. "Debt with Potential Repudiation: Theoretical and Empirical Analysis." *Rev. Econ. Studies* 48 (2): 289–309.
- Evans, Martin. 1996. "Peso Problems: Their Theoretical and Empirical Implications." In *Handbook of Statistics*, vol. 14, edited by G. S. Maddala and C. R. Rao, 613–46. Amsterdam: North-Holland.
- Fevrier, Philippe, Raphaelae Preget, and Michael Visser. 2004. "Econometrics of Share Auctions." Manuscript, Center Res. Econ. and Statis., Paris.
- Fudenberg, Drew, Markus Mobius, and Adam Szeidl. 2007. "Existence of Equilibrium in Large Double Auctions." *J. Econ. Theory* 133 (1): 550–67.
- Goldstein, Henry. 1962. "The Friedman Proposal for Auctioning Treasury Bills." *J.P.E.* 70 (4): 386–92.
- Gourio, François. 2012. "Disaster Risk and Business Cycles." *A.E.R.* 102 (6): 2734–66.
- Hatchondo, Juan Carlos, Leonardo Martinez, and Cesar Sosa-Padilla. 2016. "Debt Dilution and Sovereign Default Risk." *J.P.E.* 124 (5): 1383–422.
- Hortaçsu, Ali. 2011. "Recent Progress in the Empirical Analysis of Multi-unit Auctions." *Internat. J. Indus. Org.* 29 (3): 345–49.
- Hortaçsu, Ali, and Jakob Kastl. 2012. "Valuing Dealers' Informational Advantage: A Study of Canadian Treasury Auctions." *Econometrica* 80:2511–42.
- Hortaçsu, Ali, Jakob Kastl, and Allen Zhang. 2018. "Bid Shading and Bidder Surplus in the U.S. Treasury Auction System." *A.E.R.* 108 (1): 147–69.
- Hortaçsu, Ali, and David McAdams. 2010. "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market." *J.P.E.* 118 (5): 833–65.
- Kastl, Jakob. 2011. "Discrete Bids and Empirical Inference in Divisible Good Auctions." *Rev. Econ. Studies* 78 (3): 974–1014.
- Manzano, Carolina, and Xavier Vives. 2020. "Market Power and Welfare in Asymmetric Divisible Good Auctions." Working paper, IESE Bus. School, Barcelona.
- Mendoza, Enrique, and Vivian Yue. 2012. "A General Equilibrium Model of Sovereign Default and Business Cycles." *Q.J.E.* 127:889–946.
- Passadore, Juan, and Yu Xu. 2020. "Illiquidity in Sovereign Debt Markets." Working paper, Einaudi Inst. Econ. and Finance, Rome.
- Reny, Philip J., and Motty Perry. 2006. "Toward a Strategic Foundation for Rational Expectations Equilibrium." *Econometrica* 74 (5): 1231–69.
- Rogoff, Kenneth. 1980. "Essays on Expectations and Exchange Rate Volatility." PhD diss., Massachusetts Inst. Tech.
- Wilson, Robert. 1979. "Auctions of Shares." *Q.J.E.* 93:675–89.