How to Sell Public Debt in Uncertain Times*

Harold Cole† Daniel Neuhann‡ Guillermo Ordoñez§

July 5, 2025

Abstract

How should governments structure primary sovereign bond markets when investors face asymmetric uncertainty about default risk and total demand? Standard protocols either use uniform prices for all investors, or price discriminate based on bid prices ("pay as bid"). Uniform pricing encourages more bidding by uninformed investors but price discrimination captures inframarginal surplus. Based on this trade-off, we analyze a potential protocol that features price discrimination on inframarginal bids and uniform pricing on marginal bids, thereby reducing the mean and variance of public borrowing costs. We also offer new evidence on the information content of primary bond markets which supports our framework.

*We thank Satyajit Chatterjee, William Fuchs, Andrey Ordin, Diego Perez and Milena Wittwer for helpful comments.

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

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

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

1 Introduction

Over the past decade, public debt levels have grown rapidly around the world, reaching $100 trillion in 2024 (International Monetary Fund, 2024). These trends have triggered concerns about the creditworthiness of a number of countries (Arellano et al., 2024), the appropriate valuation of public debt (Jiang et al., 2024; Horn et al., 2024), and the willingness of the private sector to absorb more public debt in the future (Choi et al., 2024). As a result of these concerns, governments around the world face increasingly urgent questions about the appropriate management of public debt issuance going forward.

In this paper, we focus on an important but understudied question in public debt management, which is how to structure the primary market for public debt in periods of uncertainty. Specifically, we study pricing protocols for public debt in a rich framework with risk-averse investors, shocks to default risk, uncertainty about demand or supply, and asymmetric information. The choice between canonical market designs (i.e, discriminatory and uniform price auctions) is shaped by a tradeoff between extracting surplus through price discrimination and fostering more aggressive bidding through uniform pricing. As such, discriminatory auctions outperform uniform auctions when there is little asymmetric information but demand curves are steep. We then analyze a practical and robust pricing protocol with *partial price discrimination* that overcomes this tradeoff and lowers the mean and volatility of public borrowing costs.

There appears to be broad consensus on basic properties of optimal primary bond market design, with the vast majority of debt-issuing countries selling bonds using large multi-unit auctions at regular intervals. However, there is much less agreement on the particular protocols that should govern these auctions. Brenner et al. (2009) show that there is large heterogeneity in the pricing protocol used by governments to sell bonds, with many emerging and advanced economies opting to use either discriminatory or uniform pricing rules. This mirrors a lack of theoretical consensus regarding the relative desirability of standard auction protocols in realistic common-value multi-unit auctions.

Given this background, we take as given the choice to run an auction, and analyze the design of primary bond markets using a model of multi-unit auctions with competitive,

risk-averse bidders who can submit multiple bids for arbitrary quantities.[1] Its key features are (i) that investors have convex marginal utility (e.g. CRRA preferences), which leads to downward sloping demand – essential to capture realistic risk premia in bond markets – and (ii) that there is asymmetric information about two sources of uncertainty: a *quality shock* that determines the likelihood of a sovereign default, and a *quantity shock* that determines the total demand and/or supply of bonds.[2] The quality shock is a common value shock that shifts all investors' demand functions and can induce the winner's curse. The quantity shock instead triggers a move along the demand curve. Given convex marginal utility, the combination of these shocks induces non-linear shifts in marginal valuations that are central to much of asset pricing and macroeconomics, particularly in emerging markets (Morelli et al., 2022; Bai et al., 2025). However, they have received comparatively little attention in the auction literature thus far. We then ask how different protocols affect the mean and variance of new debt that must be issued to achieve a given revenue target (the government's future *debt burden*). We also study how well auctions reveal information to secondary markets.

In periods of little uncertainty, the choice between protocols may be of relatively low importance. If all investors agree on the appropriate price to bid, it is not particularly important how the protocol treats bids at other prices. However, protocols may perform quite differently when substantial uncertainty magnifies the winner's curse that arises under asymmetric information about the common value of bonds, such as default risk.[3] Auction protocols that impose a common price on accepted bids can partially alleviate this concern by insuring uninformed bidders against overpaying. However, such uniform pricing comes at the cost of making the revenue per bond issued more sensitive to the marginal investor's valuation, and it precludes the government from harvesting additional revenue from bids made at high prices. As such, the choice between standard pro-

[1] While we assume a continuum of bidders, Allen et al. (2024) show that even a relatively small number of dealers is sufficient to achieve competitive pricing.

[2] While we model the quality shock as outright default, the basic mechanism is the same for de facto default by way of unexpected inflation. Bigio et al. (2023) provide evidence of demand-based liquidity costs based on Spanish sovereign debt auctions.

[3] Milton Friedman famously argued that the U.S. should switch from largely relying on discriminating-price auctions to uniform-price auctions for this reason (Hearings before the Joint Economic Committee, 86th Congress, 1st Session, Washington, D.C., October 30, 1959, 3023-3026).

tocols, such as the pay-as-bid or uniform price auctions, is driven by a tradeoff between surplus extraction and bidding intensity that is sensitive to country risk, the information environment, financing needs, and the curvature of investor demand.

In particular, we show that the tradeoff hinges on the interplay of two main forces: (i) nonlinear marginal valuations for bonds, which induce infra-marginal rents that can be captured by the government, and (ii) asymmetric information, which may deter bidding by uninformed bidders subject to the winner's curse. The uniform auction is effective at reducing the winner's curse because all investors get to buy at the lowest accepted price. But precisely because all investors pay the lowest price, it also prevents the government from capturing any inframarginal rents. When inframarginal rents are large but the winner's curse is not too severe, the government may therefore prefer to use a discriminatory protocol in order to harvest the revenue from bids at high prices.

We exploit these insights into the determinants of bond prices to propose a novel auction format that combines the most beneficial features of discriminating- and uniform-price auctions. In our proposed format, bids at prices above the lowest accepted price are executed at the bid price *only up to* a *bid quantity tier* τ, after which they are executed at the lowest accepted price. For example, if the marginal price is \$1 and a bidder submits 10 bids at a price of \$2, the first τ bids are executed at \$2 and the remaining $10 - \tau$ at price \$1. This means that the government uses price discrimination on inframarginal bids and uniform pricing on marginal bids. The benefit of this format is that it allows the government to harvest revenue from some high-price bids while removing the winner's curse on marginal bids. Compared to both canonical protocols, the new protocol can reduce the mean and variance of new debt that must be issued to raise a given level of revenue. The new protocol is also simple and has desirable robustness properties: it deviates from canonical protocols only by parameter τ, and small and large values of this parameter recover the standard uniform and discriminatory protocols, respectively.

We first make our points in a setting with exogenous asymmetric information about one shock and perfect information about the other shock. We then extend our model to consider asymmetric information about one shock and symmetric uncertainty about the other, developing a computational methodology to examine bond pricing and the value of

information as the number of informed investors varies. This extension endogenizes the extent of asymmetric information and brings to the fore an additional layer of complexity when comparing protocols. First, for uniform-price auctions, we show that uniformed investors can perfectly replicate the optimal portfolio of informed bidders as long as the fraction of informed bidders is large enough. Hence there are no gains from acquiring information. However, this is not possible when the fraction of informed bidders is low, as marginal prices then no longer reveal the underlying shock. In contrast, in discriminatry-price auctions the value of information is higher because it allows investors to avoid the winners' curse. We interpret these findings as implying that, once we account for endogenous information acquisition, discriminatory auction prices are more likely to reveal the state than uniform auction prices, despite intuitive reasons to believe the opposite.

To examine this auxiliary prediction of our model, we measure how price surprises at an auction (deviations from what one would expect given recent secondary market prices) predict the secondary market price immediately after the close of the auction – an *elasticity of secondary market prices to information released at auction*. We exploit the fact that Mexico switched from discriminatory to uniform auctions for its nominally denominated bonds (called *Cetes*) during a period in which fundamentals seemed largely unchanged, and study how the informativeness of the auctions changed with the change in protocol. We find that discriminatory auctions were more informative (in the sense of the reaction of secondary prices to auctions) than uniform auctions, which is consistent with our results and point toward the presence of asymmetric information in primary bond markets.

1.1 Related literature

Our paper contributes to the literature on the sustainability and optimal issuance of public debt, particularly in periods of uncertainty. The existing literature has analyzed many important aspects of government policy, including debt levels (Eaton and Gersovitz, 1981; Arellano, 2008), maturity (Aguiar et al., 2019; Bocola and Dovis, 2019), indexing (Barro, 2003), reserve management (Alfaro and Kanczuk, 2009; Bianchi et al., 2018; Barbosa-Alves et al., 2024), liquidity management (Bigio et al., 2023) and government market power

(Choi et al., 2024). We complement this literature by analyzing how to design the primary market to best raise a given revenue target in a framework that incorporates many features that are central to assset pricing and macroeconomics. As such, our findings can be applied in any situation in which a government faces quantity and quality shocks and funding needs due to some overarching fiscal policy.

Methodologically, we depart from the sovereign default literature in two main ways. First, we abstract entirely from the government's strategic decision to default or to provide information, and instead focus on investors' decisions. Second, we model investors who are risk averse, asymmetrically informed about a variety of shocks, and bid under realistic pricing protocols. Much of the sovereign debt literature instead relies on the assumption of risk-neutral investors and Walrasian secondary markets to construct a simple mapping between default risk and equilibrium returns.[4] Morelli et al. (2022) and Bai et al. (2025) show that risk premia are critical for understanding global bond prices.

Our focus on the design of the primary market links our work to an important literature on multi-unit auctions. It is well-known that standard revenue equivalence theorems do not apply to multi-unit common value auctions (Maskin and Riley, 1985; Engelbrecht-Wiggans, 1998). As such, rankings over auction protocols are sensitive to particular assumptions regarding the properties of the objects for sale, the slope of demand curves, the source of common valuations, and the information structure (Ausubel et al., 2014; Hortacsu and McAdams, 2018). This motivates us to study a model that incorporates features that are particularly relevant for issuing public debt, including common values, risk aversion with convex marginal utility, post-auction default risk, and asymmetric information about multiple shocks. We also consider outcome variables that are not central to canonical auction theory but are important for government bond markets, such as the volatility of government funding costs and the extent of information that is revealed to secondary markets. This differentiates us from other settings previously studied in the literature.[5]

[4]See for example Eaton and Gersovitz (1981), the review articles by Aguiar and Amador (2014) and Aguiar et al. (2016), and the recent quantitative literature by Aguiar and Gopinath (2006), Arellano (2008), Chatterjee and Eyigungor (2012), Bocola and Dovis (2019). Chaumont (2020) and Passadore and Xu (2022) consider frictional secondary markets.

[5]For instance, Hortacsu and McAdams (2010) compare equilibrium outcomes under uniform and discriminatory auctions, but under private values. Allen et al. (2024) study revenue volatility due to entry and

To accommodate the aforementioned features, we study large auctions with a continuum of bidders. Given this assumption, strategic considerations become less important and a price-taking approach to optimal bids emerges as a useful approximation. Cole et al. (2022) shows that a model of large auctions can account for key bid-level data in Mexican sovereign bond auctions. Other papers show formally that price-taking behavior emerges when the number of bidders get large. A recent example is Fudenberg et al. (2007), who show that the equilibria of large double auctions with correlated private values are essentially fully revealing and approximate price-taking behavior. Reny and Perry (2006) show a similar result when bidders have affiliated values and prices are on a fine grid.[6] In this context, this paper complements our own existing work in Cole et al. (2022) and Cole et al. (2024) in which we took the discriminatory price-protocol as given and showed that information asymmetries played a key role for the level and dynamics of bond prices within and across countries. Here we extend our analysis to allow for different protocols and establish a tradeoff between uniform and discriminatory protocols that is driven by rent extraction through price discrimination and bidder discouragement through the winner's curse. Based on this tradeoff, we also propose a new protocol that can outperform standard protocols in our setting.

Other recent papers study settings more closely related to ours, albeit for different purposes. Gupta and Lamba (2024) study auction performance in an acute crisis in a model with risk averse bidders and an additive common value component. This differs from our setting where the common value component is the default probability. They also suggest that alternative mechanisms could improve auction outcomes, but do not provide fully specified alternative. Alves Monteiro and Fourakis (2024) use a dynamic model to show that the insurance properties of discriminatory auctions first discussed in Cole et al. (2022) interacts with limited commitment on the part of the government. Boyarchenko et al. (2021) study the information advantage of dealers in a model with common values. Wittwer and Allen (2025) use a mean-variance framework to study the

exit of bidders into Treasury auctions, but the mechanism is quite distinct from that in our paper.

[6]Malvey et al. (1995) report that the U.S. Treasury typically receives 75-85 competitive bids or tenders, many of which come from the 37 primary deals. They also receive 850-900 noncompetitive tenders through the book-entry system and another 19,000 through TREASURY DIRECT.

role of dealer market power and capital constraints.

Our framework emphasizes the importance of uncertainty about default risk and demand. While such uncertainty has always been an important consideration for developing market economies, it is also becoming more pressing for a subset of countries that were thought to be able to issue much greater quantities of debt without a concomintant increase in borrowing costs – an "exorbitant prilivege." In particular, recent work casts doubt on the durability of this privilege (Jiang et al., 2024; Choi et al., 2024).

2 Model

2.1 Environment

We study a single-good, two-date ($t = 1, 2$) economy with a government and a measure one of ex-ante identical risk-averse investors with fixed per-capita wealth W.

The government arrives at date 1 needing to raise a fixed amount of revenue, due either to required spending or to roll over existing debt obligations. To sell bonds, the government uses a sealed-bid, multi-unit auction, which takes place at date 1. Each bond promises a unit of consumption at date 2.

Investors can choose to buy government bonds at the auction or invest in risk-free assets whose net return is normalized to zero. Investors have preferences over consumption at date 2 that are represented by a strictly concave, twice continuously differentiable utility function u with convex marginal utility that satisfies the Inada conditions. CRRA preferences, for instance, satisfy these requirements, and we often assume log utility. There is no borrowing: investors cannot bid negative quantities or short the risk-free asset.

Quality uncertainty. Bonds are defaultable zero-coupon debt instruments. They promise a unit payoff at date 2 but pay zero if the government defaults. Default is denoted by a $\delta = 1$ indicator and repayment by $\delta = 0$. The probability of default is given by $\kappa \in (0, 1)$, which is a random variable. We assume that the default probability takes on discrete values $\kappa \in \mathcal{K} \equiv \{\kappa_1, ...\kappa_K\}$ where $\kappa_k < \kappa_{k+1}$. We refer to the realization of the default

probability as a *quality shock* to the bond. Since it is common across all bonds sold at auction, it operates like a common value shock.[7]

Quantity uncertainty. The second source of uncertainty is a *quantity shock* that alters the per-capita revenue the government must raise from each investor participating in the auction. We model this shock by assuming that the revenue the government needs to raise from potential auction participants is ψD, where $D > 0$ is a constant and ψ is a random variable drawn from a set Ψ. For expositional reasons, we focus on discrete shocks, $\psi \in \Psi \equiv \{\psi_1, ...\psi_M\}$ where $\psi_k < \psi_{k+1}$ and $\psi_1 = 1$.

Quantity shocks are an important feature of the data both because governments directly vary the amount sold and because many sovereign auctions feature "noncompetitive bidding" for smaller quantities in which bidders submit the quantity they wish to buy at the yield (price) determined at the auction. Examples include Treasury Direct in the U.S. and Cetes Directo in Mexico. Alternatively, we can think of these shocks as impacting the share of investors who are unable or unwilling to participate in the auction.[8]

While quality shocks affect the value of any given bond, quantity shocks affect marginal valuations only through the number of bonds that must be purchased in equilibrium. Hence, quality shocks induce shifts in the demand curve for bonds, while quantity shocks induce movements along the demand curve. Hence, as will show, they affect equilibrium prices and quantities in different ways.

Public regimes. We allow the distribution of both the quality and quantity shocks to change over time in a way that is known to all investors in the market. In particular, the probability distributions and supports of the quality and quantity shocks are determined

[7]With domestically denominated bonds, inflation is a common form of partial default. Additionally, even when countries go into default, investors generally receive a partial repayment. This feature can be easily incorporated by making the default payment above 0 but less than 1.

[8]In this alternative formulation, the per-capita revenue that must be raised by each investor is $D/(1-\eta)$, where η is the share of investors that do not attend the auction. This leads to the same per-capita revenue requirement if $\psi = 1/(1-\eta)$. Because the supply shock formulation is slightly more tractable, we adopt this variant of our model in the main text. In the demand shock version, investors would need to update their beliefs concerning the size of the demand shock conditional on not experiencing a demand shock themselves.

by the prevailing *public regime* ρ, with the *state of the world* $s = (\kappa, \psi|\rho)$ is drawn conditional on the prevailing regime. Hence, the regime is intended to capture all public information, such as the primary deficit, debt-to-GDP ratios, growth rates, etc.

Conditional on the public regime, the probability distribution over default risk is denoted by $f(\kappa_k|\rho)$, while the cumulative distribution function for the quantity shock is $g(\psi_k|\rho)$. The unconditional *expected* probability of default is

$$\bar{\kappa}(\rho) = \sum_k \kappa_k(\rho) f(\kappa_k|\rho).$$

and the unconditional *expected* quantity shock is

$$\bar{\psi}(\rho) = \sum_k \psi_k(\rho) g(\psi_k|\rho).$$

Allowing for public regime shifts is a useful device for capturing changes in a country's underlying conditions to undertake plausible quantitative analyses.[9] We can also capture the fact that quantity and quality shocks may be correlated unconditionally while maintaining the tractability afforded by conditional independence.

Information environment. All investors know the baseline financing needs D and the prevailing regime ρ, but may differ in their information about the prevailing quantity and quality shocks. We capture information by assigning to investor i a partition of the state space I_i, which we will refer to as the *information set*. Since investors are otherwise symmetric, the information set indexes the investor's type. We use $\mathbb{E}^i$ to denote the expectation operator conditional on type i's information set, and $G^i(\kappa, \psi)$ to denote the associated conditional probability of state (κ, ψ).

Initially, we will study asymmetric information along a single, binary dimension while assuming that there is perfect information along the other. For example, some investors may face uncertainty about the realization of a binary quality shock, while all investors know the realized quantity shock. It is then sufficient to consider two groups of investors,

[9] We explore these possibilities in (Cole et al., 2022, 2024)

informed and *uninformed*, where the informed have perfect information about the realization of the quality shock while the uninformed only knows the common prior. Later we allow this second dimension to be unknown to all, which enriches the analysis in interesting ways and clarifies important differences between uniform and discriminatory auctions. We also extend our model to allow for costly information acquisition.

2.2 Primary Markets

Primary market structure. Governments sell bonds using sealed-bid multi-unit auctions. Investors can submit multiple *bids* consisting of a non-negative quantity and a price. Bids can be made contingent on the investor's information set. The government treats bids independently and executes them in descending order of prices until their sale generates the required revenue ψD. The *marginal price* $P(s)$ is the lowest price accepted by the government in the state $s = (\kappa, \psi|\rho)$. The number of bonds sold is an equilibrium object. We focus on the two predominant protocols used in large multi-unit auctions of common-value goods: the *discriminatory-price* (DP) auction in which all accepted bids are executed at the bid price ("pay as you bid"), and the *uniform-price* (UP) auction in which all accepted bids are executed at the lowest accepted (or marginal) price. Later, we propose a variant of these protocols that combines aspects of both the UP and DP auctions and show that it could improve on the prevailing protocols.

Since bonds pay at least zero and at most one, 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 consider only bidding strategies that assign zero bids to any price that is not marginal in at least one state of the world.[10] Given this restriction, it is convenient to define bidding strategies as a function of the underlying states of the world. That is, we can define another function $B(s) \equiv B'(P(s))$ that maps the state s into a quantity of bids at the associated equilibrium marginal price $P(s)$. Thus, investors must ultimately decide

[10]Excess demand at the marginal price is rationed pro-rata, but rationing does not occur in equilibrium. An investor can avoid rationing by offering an infinitesimally higher price. Moreover, given that marginal prices are distinct, for any equilibrium with rationing, there is an equivalent equilibrium in which bidders scale down their bids by the rationing factor.

how many bonds to bid for at the marginal prices associated with all possible states of the world.

Bid execution and prices. Despite investors having rational expectations over the set of marginal prices, they do not know which one will materialize at the time of bidding. This creates uncertainty about which bids ultimately will be accepted by the government. To capture this concern, we define *executed bid sets* $\mathcal{E}^i(s)$ as the set of state-contingent bids that an investor believes will be executed when the realized state is s.[11]

Since bids are accepted in descending order of bid prices, given investor i's information set, an executed bid set consists of possible states with marginal prices higher than the bid price corresponding to state (κ, ψ),

$$\mathcal{E}^i(\kappa, \psi) = \{(\kappa', \psi') : P(\kappa', \psi') \geq P(\kappa, \psi) \text{ and } (\kappa', \psi') \in I_i\} \tag{1}$$

where we suppress dependence on ρ because all analysis is conducted conditional on the public regime. The total quantity of bonds acquired by an investor i in state (κ, ψ) then is

$$\mathcal{B}^i(\kappa, \psi) = \sum_{\mathcal{E}^i(\kappa,\psi)} B^i(\kappa', \psi'). \tag{2}$$

Observe that, for given bids and marginal prices, the execution set and the total quantity of bonds acquired by the investor does *not* depend on the auction protocol. This is because bids are executed in descending order of bid prices in either protocol.

Of course, investors will, nonetheless, choose different bidding strategies depending on the protocol. This is because the prices at which bids are executed depend on the protocol, and this generates variation in the total *expenditure* on bonds, holding the quantities bid fixed. The *execution price* $\mathcal{P}(\kappa, \psi|\kappa', \psi')$ of a bid at price $P(\kappa, \psi)$ in state (κ', ψ') is

$$\mathcal{P}(\kappa, \psi|\kappa', \psi') = \begin{cases} P(\kappa, \psi) & \text{if the auction protocol is DP} \\ P(\kappa', \psi') & \text{if the auction protocol is UP} \end{cases} \tag{3}$$

[11] Because there is a unique marginal price associated with each state, we can define these sets directly in terms of the underlying states.

Given prices and bids, investor $i's$ *realized expenditure* on bonds $X^i(\kappa, \psi)$ is the product of bid quantities and execution prices for all states in the state-specific execution set. To facilitate aggregation across investors, we say that an investors' expenditure in a state is zero if the execution set is empty. This is the case for states that are ruled out by the investor's information set. Taken together, we define expenditures as

$$X^i(\kappa, \psi) = \sum_{\mathcal{E}^i(\kappa,\psi)} B^i(\kappa', \psi')\mathcal{P}(\kappa, \psi|\kappa', \psi').$$

$$X^i(\kappa, \psi) = \begin{cases} \sum_{\mathcal{E}^i(\kappa,\psi)} B^i(\kappa', \psi')\mathcal{P}(\kappa, \psi|\kappa', \psi') & \text{if } \mathcal{E}^i(\kappa, \psi) \text{ is non-empty} \\ 0 & \text{if } \mathcal{E}^i(\kappa, \psi) \text{ is empty} \end{cases} \tag{4}$$

We illustrate the underlying auction mechanics using a simple example with just three states. We focus on the case of an uninformed investor since (perfectly) informed investors optimally choose to bid only at the realized marginal price.

Example 1 (Auction mechanics with three states)**.** *Consider an uninformed bidder deciding on bid quantities for three states, $s = \in \{s_1, s_2, s_3\}$, with marginal prices $P_1 > P_2 > P_3$. Given the ranking of prices, for any protocol, the execution set for state 1 contains only state 1, the execution set for state 2 contains states 1 and 2, and the execution set for state 3 contains three states.*

The protocol affects bids by determining expenditures on accepted bids. For the **uniform protocol***, uninformed investor's expenditures at auction can be written in matrix form as*

$$\mathbf{X^{UP}} = \begin{bmatrix} X_1 \\ X_2 \\ X_3 \end{bmatrix} = \begin{bmatrix} P_1 & 0 & 0 \\ P_2 & P_2 & 0 \\ P_3 & P_3 & P_3 \end{bmatrix} \begin{bmatrix} B_1 \\ B_2 \\ B_3 \end{bmatrix}$$

where bids B_1 placed at P_1, e.g., are executed at the state-contingent marginal price. Expenditures in state 3 thus depend only on the cumulative sum of bids $B_1 + B_2 + B_3$.

For the **discriminating protocol**, *all bids are executed at the bid price, and expenditures are*

$$\mathbf{X^{DP}} = \begin{bmatrix} X_1 \\ X_2 \\ X_3 \end{bmatrix} = \begin{bmatrix} P_1 & 0 & 0 \\ P_1 & P_2 & 0 \\ P_1 & P_2 & P_3 \end{bmatrix} \begin{bmatrix} B_1 \\ B_2 \\ B_3 \end{bmatrix}.$$

Expenditures depend on the quantities bid at each marginal price, not just on the cumulative sum.

2.3 Decision Problems and Equilibrium Definition

We now describe investors' decision problems and define an equilibrium for our model. Given a set of prices and a bidding strategy, investor is's holdings of the risk-free asset $w^i(\kappa, \psi)$ is given by initial wealth minus expenditures on bonds,

$$w^i(\kappa, \psi) = W - X^i(\kappa, \psi).$$

Hence the consumption process induced by bidding strategy B^i and default outcome δ is

$$c^i(\kappa, \psi, \delta, B^i) = w^i(\kappa, \psi) + (1 - \delta)\mathcal{B}^i(\kappa, \psi) \quad \text{for all } \kappa, \psi \text{ and } \delta.$$

where the first term reflects the risk-free asset and the second bond payments.

Given the investor i's information set, optimal bidding strategies solve the following decision problem, where $\mathbb{E}^i$ denotes expectations under the information set of investor i.

Definition 1 (Bidding problem)**.** *Taking as given the set of marginal prices, investor i's portfolio choice problem given public regime ρ is*

$$\begin{aligned} V^i = &\max_{B^i} \quad \mathbb{E}^i\Big[u(c^i(\kappa, \psi, \rho, \delta, B^i))\Big] \\ &\textit{s.t.} \quad B^i(\kappa, \psi) \geq 0 \quad \textit{and} \quad w^i(\kappa, \psi) \geq 0 \quad \textit{for all } \kappa \textit{ and } \psi \end{aligned} \tag{5}$$

We next turn to our definition of equilibrium. Equilibrium prices and bids must satisfy two conditions in addition to bidder optimality. The first is the *market clearing condition* ensuring that the government raises the required revenue in each state (κ, ψ). This can be

stated directly in terms of expenditures,

$$\sum_i X^i(\kappa, \psi) = \psi D, \qquad \text{for all } (\kappa, \psi). \tag{6}$$

Note that there are four possible types of investors, given by their information sets: informed about both shocks, informed only about one shock, or not informed about either. This implies that there are possibly four levels of expenditures $X^i(\kappa, \psi)$ in this summation, weighted by the share of investors with each information set.

The second equilibrium condition is what we call the *bid overhang constraint*. If execution prices depend on the marginal price, as they do in the uniform auction, the market clearing condition can, in principle, hold at multiple marginal prices. This is because lowering the marginal price allows the auctioneer to execute more bids but also reduces the revenue raised on any given bid. The bid-overhang constraint ensures that the protocol select the equilibrium with the highest marginal price, should such multiplicity occur.

Intuitively, we require that, in state (κ, ψ), the government cannot deviate to accepting bids only up to some higher price $P(\kappa', \psi') > P(\kappa, \psi)$ and still satisfy the revenue requirement in state (κ, ψ). To state this constraint formally, we must define a notion of *counterfactual* expenditures, should the government try to move to a higher marginal price *given* the bids submitted in state (κ, ψ). We define this object as

$$\hat{X}^i(\kappa', \psi' | \kappa, \psi) \equiv X^i(\kappa', \psi') \cdot \mathbb{1}\Big((\kappa, \psi) \in \mathcal{E}^i(\kappa, \psi)\Big)$$

The characteristic function ensures that the government can deviate only to bids actually submitted in (κ, ψ), not bids that would have been submitted had another state been realized. We then have the following definition of the bid-overhang constraint. Note that this constraint never binds in the discriminatory auction, where execution prices are independent of the marginal price.

Definition 2 (Bid-overhang Constraint)**.** *For any* (κ, ψ) *and all* $(\kappa', \psi') \neq (\kappa, \psi)$*,*

$$\sum_i n^i \hat{X}^i(\kappa', \psi') < \psi D \tag{7}$$

We can now formally define an equilibrium. While closely related to the canonical notion of competitive equilibrium, there are two important differences stemming from the fact that we consider price determination by an auction protocol rather than a Walrasian auctioneer. One is the bid overhang constraint. The other is that, depending on the auction protocol, different investors may pay different prices for the same good.[12]

Definition 3 (Walrasian auction equilibrium)**.** *Fixing regime ρ, a Walrasian auction equilibrium consists of pricing functions $P : \{\kappa, \psi\} \to [0, 1]$ and investor strategies $\{B^i\}_i$ such that:*

(i) B^i solves investor i's portfolio choice problem (5)*, given his information set.*

(ii) market clearing condition (6) *and the bid overhang constraint* (7) *are satisfied.*

2.4 Optimal Bidding Strategies

We begin our analysis by establishing fundamental properties of optimal bidding strategies in both uniform and discriminatory auctions. In a large auction, individual bidders take the set of marginal prices as given and choose bidding strategies based on their expectations about the sets of states in which a given bid will be accepted. Since higher-priced bids are accepted first, a bid at a given price is accepted whenever the marginal price is smaller than the bid price. To summarize this fact, it is convenient to define *acceptance sets* $\mathcal{A}^i(\kappa, \psi)$ that collect all states in which a bid at a price $P(\kappa, \psi)$ is expected to be accepted. Formally, these sets are defined as

$$\mathcal{A}^i(\kappa, \psi) = \Big\{(\kappa', \psi') : P(\kappa', \psi') \leq P(\kappa, \psi) \text{ and } (\kappa', \psi') \in I_i)\Big\}.$$

Optimal bidding strategies can be succinctly summarized using the expected marginal rate of substitution across repayment and default, averaged across all states in the acceptance set and adjusted for the execution price $\mathcal{P}(\cdot)$ and the net payoff after repayment $1 - \mathcal{P}(\cdot)$. Denote marginal utility conditional on the default indicator δ by

$$m^i(\kappa, \psi, \delta) = u'(c^i(\kappa, \psi, \rho, \delta, B^i))$$

[12] Appendix C provides a detailed comparison of our model and the canonical competitive equilibrium.

Then the expected marginal rate of substitution is defined as

$$M^i(\kappa,\psi) \equiv \frac{\sum_{(\kappa',\psi')\in\mathcal{A}^i(\kappa,\psi)} \kappa' m^i(\kappa',\psi',1)\mathcal{P}(\kappa,\psi|\kappa',\psi')G^i(\kappa',\psi')}{\sum_{(\kappa',\psi')\in\mathcal{A}^i(\kappa,\psi)} (1-\kappa')\, m^i(\kappa',\psi',0)(1-\mathcal{P}(\kappa,\psi|\kappa',\psi'))G^i(\kappa',\psi')} \tag{8}$$

where the numerator reflects default and the denominator repayment, and G^i is the probability a state is drawn. The system of equations that determines the optimal bidding strategy is therefore given by

$$M^i(\kappa,\psi) \leq 1 \qquad \text{for all } (\psi,\kappa), \tag{9}$$

where each equation holds with equality if and only if the short-sale constraint is slack. Hence the bidding strategy obeys a variant of the canonical risk-return tradeoff familiar from asset pricing, whereby willingness to pay is determined by an expectation over state-contingent marginal utility. The key difference is that prices and quantities in a given state are determined by an auction protocol.

We again provide some intuition using our three-state example.

Example 2 (Bidding problem with three states.)**.** *We now discuss the bidding problem with three states. For the* **uniform protocol***, define the net payoff if the government repays as*

$$\mathbf{Y^{UP}} = \begin{bmatrix} Y_1 \\ Y_2 \\ Y_3 \end{bmatrix} = \begin{bmatrix} 1-P_1 & 0 & 0 \\ 1-P_2 & 1-P_2 & 0 \\ 1-P_3 & 1-P_3 & 1-P_3 \end{bmatrix} \begin{bmatrix} B_1 \\ B_2 \\ B_3 \end{bmatrix}$$

For the **discriminating protocol***, the net payoffs from repayment are*

$$\mathbf{Y^{DP}} = \begin{bmatrix} Y_1 \\ Y_2 \\ Y_3 \end{bmatrix} = \begin{bmatrix} 1-P_1 & 0 & 0 \\ 1-P_1 & 1-P_2 & 0 \\ 1-P_1 & 1-P_2 & 1-P_3 \end{bmatrix} \begin{bmatrix} B_1 \\ B_2 \\ B_3 \end{bmatrix}.$$

Given this structure, we can write the expected utility of an uninformed investor as

$$\sum_s \left\{ u(W - X_s^i)\kappa_s + u(W + Y_s^i)(1-\kappa_s) \right\} Pr\{s\}$$

Differentiation leads to a simultaneous system of first-order conditions

$$\frac{\partial}{\partial B_{s'}} \sum_{s \leq s'} \left\{ u(W - X_s^i)\kappa_s + u(W + Y_s^i)(1 - \kappa_s) \right\} Pr\{s\} = 0 \tag{10}$$

which can be rearranged to give the formulation (8) *in terms of marginal rates of substitution.*

2.5 The Main Tradeoff: Rent Extraction versus Participation

We now describe the key tradeoff underlying our model. Recall that optimal bids are determined by a system of equations requiring that the marginal rate of substitution (MRS) is equal to one, state by state. Given our assumptions on preferences, the MRS (which is inversely related to the marginal willingness to pay) is strictly increasing and convex in the quantity of bonds.

Lemma 1. *$M^i(\kappa, \psi)$ is strictly increasing and convex in $B^i(\kappa, \psi)$.*

This has two consequences which create a tradeoff between auction protocols:

(i) Bidders earn an infra-marginal surplus on bonds purchased at the marginal price.

(ii) The efficient allocation is symmetric, with each investor acquiring a symmetric per-capita share of bonds in every state.

To extract the infra-marginal surplus, the government must execute some bids at prices above the marginal price. The discriminatory auction allows this, but the uniform auction precludes it. Holding bids fixed, the discriminatory protocol thus necessarily raises more revenue for the government. But, precisely because the discriminatory auction leads to price dispersion, it discourages uninformed bidders from bidding, distorting the bond allocation away from the efficient allocation. This reduces investors' average willingness to pay and leads to lower prices. Overall, the revenue ranking between the two protocols is determined by a tradeoff between *surplus extraction* and *bidder participation* that depends on the severity of the winner's curse for the uninformed and the slope of investors' demand. As we will show, this tradeoff is sensitive both to the information environment

and country fundamentals, such as the level of debt and the mean and volatility of default risk. The next section characterizes this tradeoff in a minimalist variant of our model.

3 Simple Information Structure

We first consider an information structure in which there is perfect information about the realization of one shock, but asymmetric information about the other. In particular, we assume that a share $n \in (0, 1)$ of investors is perfectly informed about the second shock, while the remainder has no information beyond the prior. We call investors who are perfectly informed about both shocks *informed*, and the remainder *uninformed*.

Given that we have assumed that the uncertain shock is binary, this information structure implies there are only two relevant states of the world, say s_1 and s_2, with associated marginal prices $P(s_1)$ and $P(s_2)$. For example, if all investors know the realized quantity shock but only some investors know the realized quality shock, the relevant states of the world correspond to the two potential realizations of the quality shock. Hence we refer to this information structure as the *two-state model*. Further assuming log utility, we obtain a tractable characterization of the central tradeoff between protocols. Without loss of generality, we order states such that $P(s_1) > P(s_2)$.

The main reason for the tractability of the two-state model is that it circumvents two constraints that may bind when there is uncertainty about both shocks: the bid-overhang constraint and the short-sale constraint on informed investors. The reason is that, with two states, informed investors face no uncertainty at all and thus are the marginal bidders in every state. As such, the government can never deviate to a higher marginal price and still clear the market. The sort-sale constraint does not bind either for the same reason. This is summarized in the Lemma below and proved in the Appendix. In Section 4, we show that these constraints can bind under more general information structures. While this introduces additional nuance in analysis, it does not alter the central tradeoff.

Lemma 2. *In the two-state model, the bid-overhang constraint and the short-sale constraint on informed investors never bind.*

3.1 Benchmarks without Asymmetric Information

We begin by providing a benchmark for the case in which no investor is informed. This highlights an important difference between the role of quantity and quality shocks, which is that prices must depend on the realized quantity shock even if no investor is informed. The same is not true for quality shocks, which are reflected in prices only to the extent that some investors condition their bids on the realized shocks.

Proposition 1 (Uninformed Benchmark)**.** *Suppose that all investors are uninformed ($n = 0$).*

(i) In the two-state model with unobserved quality shocks, all investors evaluate bonds using the unconditional default probability $\bar{\kappa}$. Hence, there is a single marginal price $\bar{P}$ and a single optimal bid $\bar{B}$ that is invariant in the protocol. For any protocol, prices and quantities are

$$\bar{P} = 1 - \frac{\bar{\kappa}}{1 - \frac{\psi D}{W}} \qquad \bar{B} = \frac{1 - \bar{\kappa} - \bar{P}}{\bar{P}(1 - \bar{P})W}. \tag{11}$$

(ii) In the two-state model with unobserved quantity shocks, marginal prices must respond to the realized shock even if no bids are conditional on the shock. This is because the government must execute different quantities to satisfy the state-contingent revenue requirement. Prices and quantities generically differ across the uniform and discriminatory protocols.

The equilibrium in which all investors are informed (*the informed benchmark*) follows the same construction. The difference is that prices and bids are contingent on the realized quality shock, replacing $\bar{\kappa}$ with $\kappa(s)$, $\bar{P}$ with $P(s)$ and $\bar{B}$ with $B(s)$ in equation (11).

3.2 Equilibrium with Asymmetric Information.

We now study the two-state model with asymmetric information, $0 < n < 1$. We begin by deriving optimal informed bid functions. These turn out to be invariant in the protocol.

Optimal informed bid functions in the two-state model. Since informed investors face no uncertainty regarding the realized state, they submit bids *only* at the realized marginal

price. For any protocol $j \in \{DP, UP\}$, informed investors' first-order condition is

$$M_j^I(s) = \frac{\kappa(s)u'\Big(W - P_j(s)B_j^I(s)\Big)P_j(s)}{(1-\kappa(s))u'(W + (1-P_j(s))B_j^I(s))(1-P_j(s))} = 1. \tag{12}$$

where $\kappa(s)$ is the default probability associated with state s. Under log utility, the bid function can be solved in closed form and is the same for both protocols. Specifically,

$$B_j^I(s) = \frac{1 - \kappa(s) - P_j(s)}{P_j(s)(1-P_j(s))}W. \tag{13}$$

for any protocol j. Informed investors place a fraction of initial wealth in bonds, and this fraction decreases in the marginal price and the probability of default. The fact that bid functions are decreasing in these elements turns out to be critical for determining equilibrium, as we show next. Since uninformed bid functions depend on the protocol, we analyze each protocol in turn.

3.2.1 Equilibrium in the uniform-price auction

We begin by analyzing the uniform protocol. A striking property of this auction is that, at least in the two-state model, uninformed investors can always do as well as informed investors. In particular, if informed bid quantities are strictly ordered by price, in a uniform auction an uninformed investor can form a *replicating* bidding strategy that achieves the same final portfolio as informed investors. This is shown in the next proposition.

Proposition 2 (Properties of optimal bidding under UP auctions)**.** *In the uniform auction,*

(i) If the short-sale constraint is always slack, the bidding problem of an uninformed investor can be recursively solved state by state, starting from the lowest-price state.

(ii) If the optimal bidding strategy of an informed investor is such that informed quantities are ordered by price (i.e., if $P(s_k) < P(s_{k+1})$ then $B^I(s_k) > B^I(s_{k+1})$ for all k), then an uninformed investor can form a replicating bidding strategy that leads to the same portfolio

as that of an informed investor, state by state. This replication strategy is:

$$B^U(s_1) = B^I(s_1) \quad \textit{and} \quad B^U(s_k) = B^I(s_k) - \mathcal{B}^U(s_{k-1}) \qquad \textit{for all } k > 1 \tag{14}$$

The intuition for this result is simple: if informed investors submit an increasing quantity schedule, then uninformed investors can acquire the same final quantities by bidding only the *incremental* informed quantity at each marginal price. Since all uninformed bids at prices above the marginal price are accepted, the cumulative sum of accepted bids is equal to the informed bid quantity in each state. Since all accepted bids are executed at the marginal price, uninformed investors also pay the same as informed investors. This last step immediately implies that perfect replication fails in the discriminatory auction where infra-marginal bids are executed at the bid price.

In the two-state model, informed bid quantities are always decreasing in bid prices. Hence, perfect replication is always feasible under the uniform protocol. This means that the uniform auction equilibrium is independent of the share of informed investors (as long as $n > 0$), and we can characterize it as follows.

Proposition 3 (UP equilibrium in the two-state model)**.** *In the two-state model with a uniform auction, uninformed investors can always perfectly replicate the portfolio of informed investors, and hence equilibrium is invariant in n. Prices and quantities are*

$$P^*_{UP}(s) = 1 - \frac{\kappa(s)}{1 - \psi(s)\frac{D}{W}} \qquad \textit{and} \qquad B^*_{UP}(s) = \frac{\psi(s)D}{P^*_{UP}(s)} \qquad \forall n \in (0, 1].$$

3.2.2 Equilibrium in the discriminatory-price auction

We next turn to the discriminatory protocol. Informed investors' optimal bid functions are the same as before. This is because they only submit bids at the marginal price, and thus do not care at which price bids at above-marginal prices are executed. In contrast, uninformed investors cannot replicate the informed portfolio in the discriminatory-price auction, and thus choose different bid quantities.

Since $P(s_1) > P(s_2)$, uninformed investors' state-contingent bond expenditures are

$$X^U_{DP}(s_1) = P_{DP}(s_1)B^U_{DP}(s_1) \qquad \text{and} \qquad X^U_{DP}(s_2) = X^U_{DP}(s_1) + P_{DP}(s_2)B^U_{DP}(s_2)$$

and risk-free asset holdings are therefore equal to $w^U_{DP}(s) = W - X^U_{DP}(s)$.

The optimal bidding strategy aligns marginal utility after default with that after repayment. The optimality condition for bids at $P_{DP}(s_2)$ (accepted only in state 2) is

$$\begin{aligned} &\kappa(s_2)u'\Big(w^U_{DP}(s_2)\Big)P_{DP}(s_2) \\ =&(1-\kappa(s_2))u'\Big(w^U_{DP}(s_2) + \mathcal{B}^U_{DP}(s_2)\Big)(1 - P_{DP}(s_2)). \end{aligned} \tag{15}$$

where $\mathcal{B}^U_{DP}(s)$ is definied as in (2).

Bids at $P_{DP}(s_1)$ are accepted in both states and so the willingness to pay is a weighted average of marginal utility across states. As such, the optimality condition for $B^U_{DP}(s_1)$ is

$$\begin{aligned} &\sum_{s=s_1}^{s_2} Pr(s)\kappa(s)u'\Big(w^U_{DP}(s)\Big)P_{DP}(s_1) \\ =&\sum_{j=s_1}^{s_2} Pr(s)(1-\kappa(s))u'\Big(w^U_{DP}(s) + \mathcal{B}^U_{DP}(s)\Big)(1 - P_{DP}(s_1)) \end{aligned} \tag{16}$$

where $Pr(s)$ denotes the probability of state s. (We can obtain the familiar formulation in terms of marginal rates of substitution from (8) by rearranging this equation.)

The key property of this system is that bids at the high price are evaluated using an expectation over marginal utility across both states, adjusted for the fact that high-price bids are always executed. This creates cross-state linkages in the optimal strategy, which alter the marginal willingness to pay for bonds at the low price. The optimal bidding strategy cannot be solved recursively state by state.

In the case of quality shocks, these linkages can be interpreted in terms of the classic *winner's curse*: if an uninformed investor bids at the high price associated with the low default probability, he pays the high price even if the bond turns out to be of low quality. Quantity shocks operate somewhat differently because they do not directly affect investor

valuations. Instead, prices are low in the high-supply state *only* to the extent that investors must bear more risk per capita in equilibrium. Hence, uninformed investors are worried not about bond quality but about paying too much for a bond they could have acquired more cheaply (i.e., there is price dispersion for bonds of equal quality). To distinguish this effect from the winner's curse, we label it *buyers' remorse*.

3.2.3 The tradeoff between protocols:

Since the government faces a revenue target, we measure the efficiency of the auction using the quantity of bonds sold, which we refer to as the (conditional) *debt burden*

$$DB(s) = \sum_i n_i \mathcal{B}^i(\kappa, \psi). \tag{17}$$

This is a natural state variable for the government since it measures the future obligations incurred to raise the required revenue in a given state. Our overall efficiency measure is the ex-ante *expected debt burden* across states of the world,

$$DB = \sum_s Pr\{s\} DB(s) \tag{18}$$

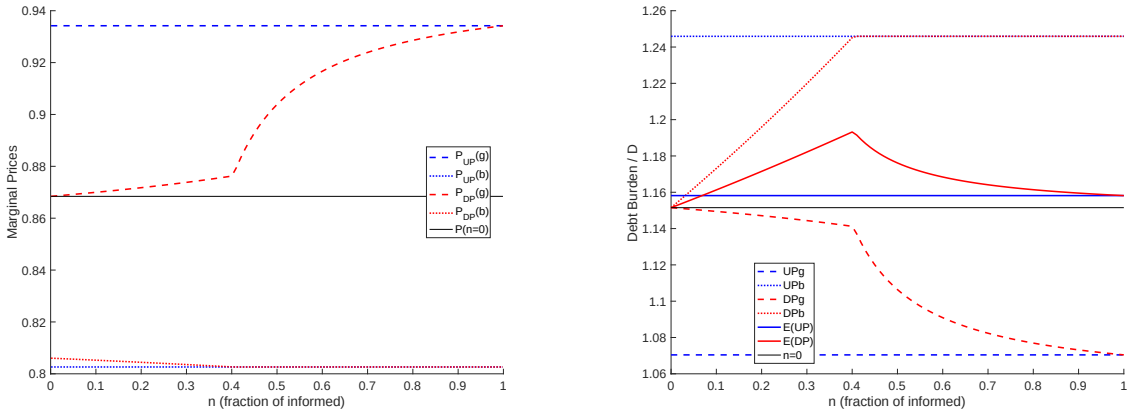
We show that the discriminatory auction performs better when the winner's curse is minor, and there are sufficiently many bids at infra-marginal prices, while the uniform auction performs better when the winner's curse is severe (for instance when debt levels, default risk, and uncertainty about default risk are sufficiently high).

Proposition 4 (Comparing uniform and discriminatory protocols)**.** *In the two-state model with asymmetric information regarding the quality shock, the following statements hold:*

1. *In the limit as $n \to 0$, both the expected debt burden and its standard deviation are strictly higher in the uniform protocol than in the discriminatory auction.*

2. *If $n < 1$, and n, $Pr(s_2)$, and $\kappa(s_1) - \kappa(s_2)$ are sufficiently large and D/W is not too large, the expected debt burden is lower in the uniform protocol than in the discriminatory protocol*

3. *if $n = 1$, the expected debt burden is the same for both protocols in all states of the world.*



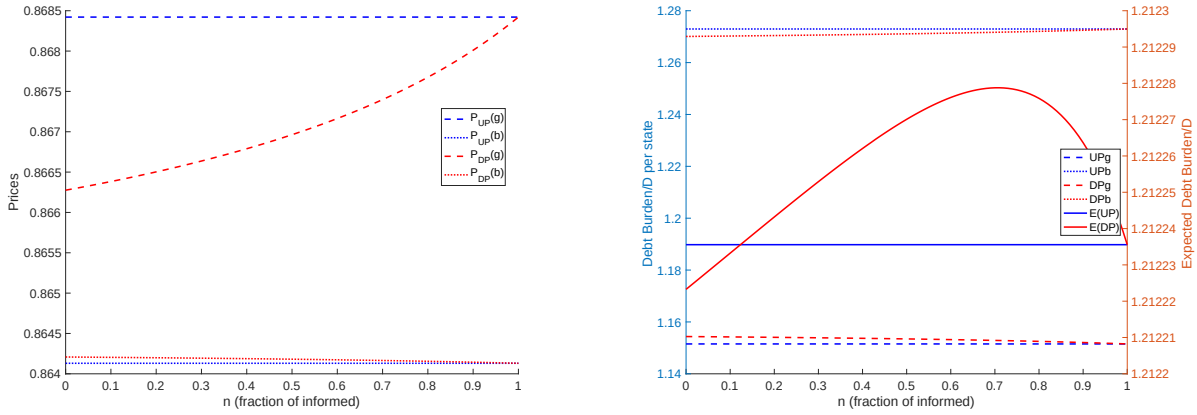
(a) Marginal Prices

(b) Debt Burden Relative to Amount Raised

Figure 1: Marginal Prices and Debt Burden: Quality shocks. Parameters: $u(c) = \log(c)$, $W = 250$, $D = 60$, two equally likely default probabilities $\kappa_b = 0.15$, $\kappa_g = 0.05$, and a known supply shock $\psi = 1.0$.

Figure 1 illustrates the proposition numerically for unobserved quality shocks and Figure 2 for unobserved quantity shocks. Regardless of the source of the shock, when there are few informed investors, the discriminatory protocol achieves a lower expected debt burden, but the uniform auction performs better when there are more informed investors, intensifying the winner's curse in the case of quality shocks and the buyer's remorse in the case of quantity shocks. When all investors are informed, there is no difference between the protocols because all investors bid only at the marginal price.

The amount of public debt: Since investors are risk-averse, more public debt implies compensating investors for holding more risk through lower prices, hence a higher debt burden relative to the amount raised. This is clear in Figure 3. There is an additional effect of the debt level in terms of which protocol is more expensive in terms of debt burden. The higher the debt level, the larger the range of n for which DP protocols represent a lower debt burden. For the benchmark of $D = 60$ (solid in the figure) DP is preferred when $n < 0.05$. When debt duplicates (dashed in the figure) DP is preferred when $n < 0.20$. This result is also obtained in situations in which net wealth declines (W reduces) or risk aversion increases.



(a) Marginal Prices

(b) Debt Burden Relative to Amount Raised

Figure 2: Marginal Prices and Debt Burden: Quantity shocks. Parameters: $u(c) = \log(c)$, $W = 250$, $D = 60$, two equally likely supply shocks $\psi_l = 1.0$, $\psi_h = 1.1$, and a known default probability $\kappa = 0.1$. Note: The expected debt burden is very small and represented on the right axis of the second panel.

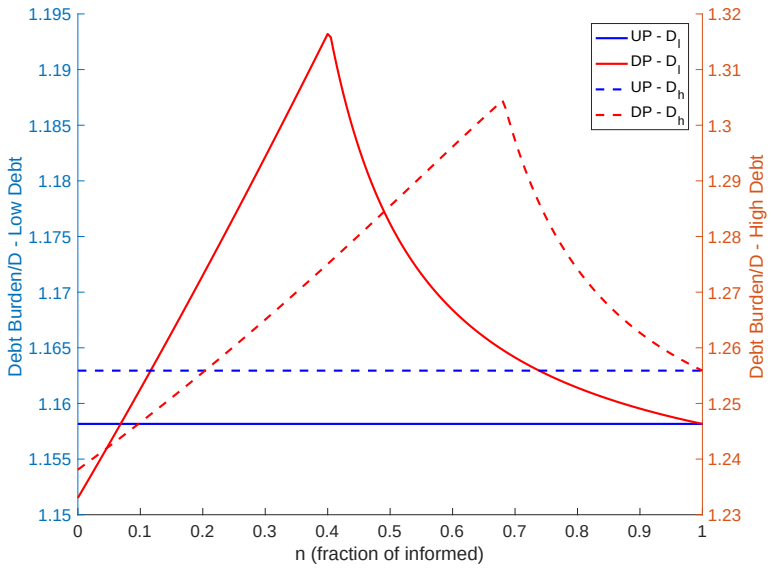
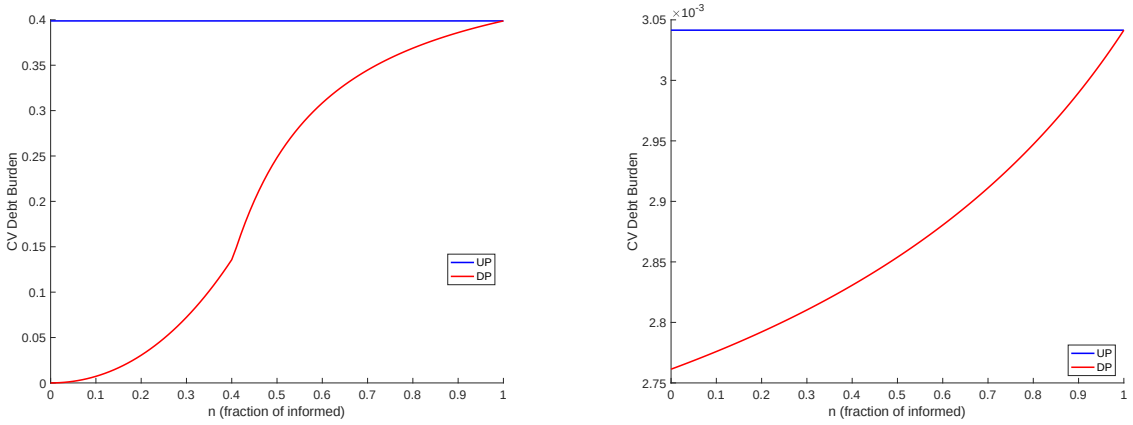


Figure 3: Debt Burden Comparison between UP and DP and different debt levels, with quality shocks. Parameters: $u(c) = \log(c)$, $W = 250$, two equally likely supply shocks $\psi_l = 1.0$, $\psi_h = 1.1$, and a known default probability $\kappa = 0.1$. $D_l = 60$ (benchmark) and $D_h = 120$.

Ex-ante government uncertainty about debt burden: Governments not only care about the size of the future debt burden needed to raise funds but also about its volatility. In Figure 4 we show the coefficient of variation of the debt burden across the two states for both protocols, for each type of shock. It shows that the debt burden under a DP protocol is always less uncertain than under a UP protocol, regardless of the shock type. This result comes from two forces. First, under DP the uninformed buy bonds submitted at high prices in both states. Second prices under DP are more compressed than under UP. This creates a silver lining for DP auctions: though inducing more of a debt burden on average for a region of n, it is less uncertain. This may turn into a relevant tradeoff for a government that worries about debt burden intertemporally.



(a) Quality Shocks

(b) Quantity Shocks

Figure 4: Coefficient of Variation of the Debt Burden: Parameters: $u(c) = \log(c)$, $W = 250$, $D = 60$. Quality shocks: $\kappa_b = 0.15$, $\kappa_g = 0.05$, and a known supply shock $\psi = 1.0$. Quantity shocks: $\psi_l = 1.0$, $\psi_h = 1.1$, and a known default probability $\kappa = 0.1$.

3.3 Overcoming the tradeoff: Partially Discriminating Auctions

The previous section demonstrated the central tradeoff between UP and DP auctions: while charging above marginal price allows the government to extract some surplus, it also distorts the allocation by exposing uninformed investors to the winner's curse (or the buyers' remorse), thereby discouraging bidding and depressing prices overall.

We now show that a simple change to the auction format can weaken this tradeoff and allow the government to capture infra-marginal surplus while sharply reducing the winner's curse. Specifically, we propose a convex combination of uniform and discriminatory protocols we call the *partially discriminating auction* (PD). The key idea is for the government to execute *infra-marginal bids at bid (infra-marginal) prices* and *marginal bids at marginal prices*. The uniform auction instead executes all bids at the marginal price and the discriminatory executes all bids at bid prices.

Partially Discriminating Protocol. Define a *quantity tier* $\mathcal{T}$, chosen by the government in advance of the auction without reference to individual bids placed at auction. As before, investors can submit any number of bids consisting of a price and a quantity. Different from before, if an investor has submitted bids $\mathcal{B}$ at a price above the state-contingent marginal price in a given state, *the government executes* $\mathcal{T}$ *of these bids at the bid price, and the remainder at the marginal price*. As long as the tier is chosen such that uninformed investors are willing to bid more than $\mathcal{T}$, the quantity tier removes the marginal disincentive to submit high bids while maintaining the ability to capture inframarginal surplus.[13]

Remark 1. *The partially discriminating auction (PD) is a convex combination of the DP and UP protocols. Setting* $\mathcal{T} = 0$ *yields the uniform protocol, while setting* $\mathcal{T}$ *sufficiently large recovers the discriminatory protocol. Hence, it has desirable robustness properties: at worst, the equilibrium behavior will be equivalent to that under the standard UP or DP protocols.*

Uninformed bidding strategy in the partially discriminating auction. As before, uninformed investors pick a bidding strategy $\{B(s_1), B(s_2)\}$, taken as given the quantity tier $\mathcal{T}$. Assume that, as is optimal, the quantity tier is such that $B(s_1) > \mathcal{T}$ in equilibrium (else the partially discriminating auction is equivalent to the discriminatory auction). We will refer to such choices of $\mathcal{T}$ as *interior*.

[13]The basic mechanism should be familiar from the literature on non-linear pricing, as surveyed in Armstrong (2016). We argue that it is also useful in the context of auctioning risky bonds to risk-averse buyers.

In this protocol, uninformed investors' state-contingent bond expenditures are

$$\begin{aligned} X^U_{PD}(s_1) &= P_{PD}(s_1)B^U_{PD}(s_1) \\ X^U_{PD}(s_2) &= P_{PD}(s_1)\mathcal{T} + P_{PD}(s_2)\Bigg[\Big(B^U_{PD}(s_1) - \mathcal{T}\Big) + B^U_{PD}(s_2)\Bigg] \end{aligned} \tag{19}$$

and risk-free asset holdings thus are $w^U_{PD}(s) = W - X^U_{PD}(s)$. Expenditures in the high-price state are unaffected by the tier. In the low-price state, the tier determines the fraction of bids that is executed at the high price.

If $B(s_1) > \mathcal{T}$, the marginal bid at the high price is executed at the low price if state 2 is realized. Hence the optimality condition for an interior bid at $P_{PD}(s_2)$ is

$$\begin{aligned} &\kappa(s_2)u'\Big(w^U_{PD}(s_2)\Big)P_{PD}(s_2) \\ =&(1-\kappa(s_2))u'\Big(w^U_{PD}(s_2) + \mathcal{B}^U_{DP}(s_2)\Big)(1 - P_{PD}(s_2)). \end{aligned} \tag{20}$$

while the optimality condition for an interior bid at $P_{PD}(s_1)$ is

$$\begin{aligned} &\sum_{s=s_1}^{s_2} Pr(s)\kappa(s)u'\Big(w^U_{PD}(s)\Big)P_{PD}(s) \\ =&\sum_{j=s_1}^{s_2} Pr(s)(1-\kappa(s))u'\Big(w^U_{PD}(s) + \mathcal{B}^U_{DP}(s)\Big)(1 - P_{PD}(s)). \end{aligned} \tag{21}$$

The critical difference to the system of equations (15)- (16) for the discriminatory protocol is that the marginal bid at the high price is *executed at the low price*, should the low-price state be realized. This breaks cross-state linkages in the optimal bidding strategy, thereby eliminating the winner's curse *on marginal units*. In particular, using (20) in (21) yields an optimality condition for $P_{PD}(s_1)$ that depends *only* on marginal utility in state 1,

$$\begin{aligned} &\kappa(s_1)u'\Big(w^U_{PD}(s_1)\Big)P_{PD}(s_1) \\ =&(1-\kappa(s))u'\Big(w^U_{PD}(s_1) + \mathcal{B}^U_{DP}(s_1)\Big)(1 - P_{PD}(s_1)). \end{aligned} \tag{22}$$

Similar to the uniform auction, it follows that this system can be solved state by state.

The interior optimality conditions (21) and (20) are not generally sufficient for an optimal bidding strategy. This is because the constraint set is not convex (see equation (19). Hence uninformed investors may not submit bids at high price if the tier is too high. This effect is similar to the discouragement of uninformed bids in the standard discriminatory auction. However, since marginal utility is strictly decreasing, there is positive surplus from bidding an interior quantity at a price below the risk-neutral price. This implies that there exists a sufficiently small quantity tier such that it is optimal for uninformed investors to bid a positive quantity at the high price. We call such quantity tiers *interior*.

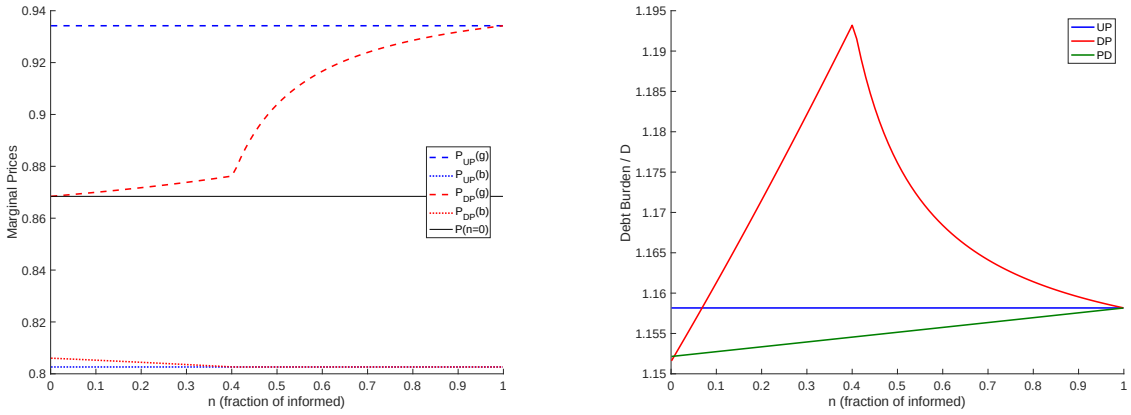
Given such an interior tier, the partially discriminating auction does strictly better than both other protocols: it improves upon the uniform auction by capturing infra-marginal rents and upon the discriminatory auction by weakening the winner's curse.

Proposition 5 (Lower Debt Burdens in the PD Auction)**.** *For any $n \in (0,1)$ and all interior choices of $\mathcal{T}$, the ranking of debt burden across protocols is as follows:*

(i) Compared to the uniform price auction, for any interior choice of $\mathcal{T} > 0$ the partially discriminating auction generates the same debt burden in the high-price state s_1 and a strictly lower debt burden in the low-price state s_2.

(ii) Compared to the discriminatory auction, the partially discriminating auction generates a strictly lower debt burden in the high-price state s_1 for any interior $\mathcal{T}$, and a strictly lower debt burden in the low-price state s_2 if $\mathcal{T}$ is sufficiently large.

Illustration. Figure 5 shows that the expected debt burden is substantially lower in the partially discriminating protocol. We set a quantity tier equal to $\mathcal{T} = 4$, which ensures that uninformed investors submit strictly positive bids at the high price for all $n \in (0,1)$. PD equilibrium prices are shown in the left panel, and the average debt burden is shown in the right panel. Because the PD protocol eliminates the winner's curse on the margin, prices are closer to those of UP. However, since there is an extraction of infra-marginal rents by the quantity tier the amount of bonds sold is sharply lower than either UP or DP, particularly when there are few informed investors.



(a) Marginal Prices

(b) Expected Debt Burden relative to Amount Raised

Figure 5: Marginal Prices and Debt Burden for UP, DP and PD protocols: Quality shocks. Parameters: $u(c) = \log(c)$, $W = 250$, $D = 60$, two equally likely default probabilities $\kappa_b = 0.15$, $\kappa_g = 0.05$, and a known supply shock $\psi = 1.0$.

Even though not depicted, the PD protocol, given this particular tier, also reduces the uncertainty of debt burden. While PD always makes the debt burden less uncertain than UP, because noninformed investors buy bonds at a given price in both states, its comparison with DP depends on the tier imposed. When PD is closer to DP (a high tier) its debt burden is less uncertain than under DP, when PD is closer to UP (a low tier) its uncertainty gets closer to that of UP, which we have shown is larger than that under DP.

4 Sophisticated Ignorance

Now we move into a situation of symmetric ignorance about one shock (no investor has information about its realization), maintaining asymmetric information about the other. If no investor knows about the quality shock, then we still obtain two states of the world. The reason is that investors bid based on the expected default probability $\bar{\kappa}$, and there are only two possible prices depending on the quantity shock realization. Hence, equilibrium has the same characterization as the previous section. The other situation, in which no investor knows about the quantity shock, and there is asymmetric information about the quality shock, brings, however, new challenges. The reason is that there are more than

two prices in equilibrium. Further, these prices may mix quantity and quality shocks, which opens the doors for short-selling and bid-overhang constraints to bind, preventing perfect replication in the uniform protocol.

Consider again a binary quality shock, $\kappa \in \{\kappa_b, \kappa_g\}$ with $\kappa_b > \kappa_g$, whose realization is known by an exogenously determined share n of investors. We place no additional restrictions on quantity shocks, and assume that no investor knows its realization. Since some investors know about the quality shock, prices must depend on quality, and even no investor knows about the quantity shock to condition her bids on, prices still depend on quantity through market clearing. Hence, this information structure will generate multiple marginal prices for each quality shock. We refer to the set of prices associated with each quality shock as the quality-contingent *price schedules* $P(\kappa, \cdot)$.

Given two sources of uncertainty, a given marginal price may be associated with different combinations of quality and quantity shocks. For example, a low marginal price could arise either because the bond is likely to default or because its realized supply is large. Because the value of the bond depends on the underlying realization, uninformed investors face an inference problem – what is the appropriate marginal valuation for bonds *conditional* on the bid being executed?[14]

We first establish a benchmark in which there is no inference problem because information is symmetric. In this case, both protocols induce equivalent outcomes – a corollary of Proposition 1.

Corollary 1 (Symmetric information benchmark)**.** *Assume that no investor is informed about the quantity shock. Then for both uniform and discriminatory protocols,*

(i) If all investors are uninformed about the quality shock, $n = 0$, there exists a single marginal price schedule $\bar{P}(\psi)$ that is decreasing in ψ and satisfies

$$M^U(\kappa_g, \psi) = M^U(\kappa_b, \psi) = 1 \qquad \text{for } j \in \{b, g\}.$$

(ii) If all investors are informed about the quality shock, $n = 1$, there are two distinct price

[14]This inference problem is similar to standard models of asymmetric information in financial markets. The main difference here is that bids are *commitments* to buy at the execution price.

schedules which satisfy $P_j(\kappa_g, \psi|n=1) > \bar{P}(\psi) > P_k(\kappa_b, \psi|n=1)$ *for any* ψ *and* j.

We now characterize equilibrium outcomes in the presence of asymmetric information about the quality shocks and no information about the quantity shocks. The most dramatic changes relative to the two-state model obtain in the uniform auction.

4.1 Equilibrium in the Uniform-Price Auction

In a situation with asymmetric information about the quality shock and no information about the quantity shock, consider first the uniform auction. In contrast to the two-state case, perfect replication may fail. Hence, we first establish conditions under which uninformed investors can replicate the portfolio of the informed using the replication strategy in (14), in which case equilibrium prices and quantities are as if every investor is informed. Then, we characterize the equilibrium when perfect replication is not possible.

4.1.1 Conditions for equilibrium with perfect replication

Perfect replication requires that the replicating bidding strategy does not involve bidding negative quantities, which happens if the bids of informed investors monotonically decline with marginal prices. With asymmetric information about the quality shocks, this condition is violated when some prices on the high-quality price schedule are smaller than some prices on the low-quality schedule. The next Proposition formalizes the conditions for replication to occur.

Proposition 6 (Conditions for perfect replication under quality shocks)**.** *Conditions are:*

1. *the quality-contingent price schedules are distinct,* $P(g,\psi) > P(b,\psi'), \quad \forall \psi, \psi'$
2. *the short-sale constraint never binds, which is ensured by* $B^I(\kappa_b, 1) - \mathcal{B}^I(\kappa_g, \psi_M) \geq 0$,
3. *the bid-overhang constraint does not bind, which is ensured by* $n > 1 - \frac{1}{\psi_M}$.

The bid-overhang constraint is sensitive to the number of informed investors because only informed bids are explicitly conditioned on the quality shock. If too many bids are uninformed, the government can hit its revenue target using only uninformed bids for

some realizations of the quantity shock. This occurs when $n \leq 1 - \frac{1}{\psi_M}$, and the bid-overhang constraint binds for all values of $\psi > n$ on the high-quality schedule, forcing an *overlap* in the high-quality and low-quality price schedules that renders replication infeasible. More precisely, there now exists pairs of quantity shocks $\{\psi_g, \psi_b\}$ such that the marginal price in state (κ_g, ψ_g) is equal to the marginal price in state (κ_b, ψ_b): there is a combination of high-quality bonds in high supply and low-quality bonds in low supply that induce the same marginal price. Since uninformed investors know neither the quality shock or the quantity shock, the realized price does not uniquely identify the underlying state. This leads to a conditional inference problem that affects the bidding decisions.

4.1.2 Equilibrium without perfect replication

How prices are determined in equilibrium if there is no perfect replication? To render this problem tractable, we assume henceforth that the quantity shock is continuously and uniformly distributed, $\psi \sim U[1, \psi_M]$. Assume, as above, log preferences, then the optimal *total* bid in any state ψ given an expected default probability $\tilde{\kappa}$ is

$$\mathcal{B} = \frac{1 - \tilde{\kappa} - P}{P(1-P)} W \tag{23}$$

where $\tilde{\kappa} = \bar{\kappa}$ for the uninformed and $\tilde{\kappa} = \{\kappa_g, \kappa_b\}$ for the informed.

Prices are then determined as follows. Take any two states $s = [\kappa_g, \psi_g]$ and $s' = [\kappa_b, \psi_b]$ for which a binding constraint forces a common price, $P = P(s) = P(s')$. The respective auction-clearing conditions for these two states are

$$n\left(\frac{1 - \kappa_g - P}{1 - P}\right) + (1 - n)\left(\frac{1 - \tilde{\kappa} - P}{1 - P}\right) = \frac{D}{W}\psi_g, \tag{24}$$

$$\text{and} \quad n \max\left[\left(\frac{1 - \kappa_b - P}{1 - P}\right), 0\right] + (1 - n)\left(\frac{1 - \tilde{\kappa} - P}{1 - P}\right) = \frac{D}{W}\psi_b. \tag{25}$$

Since the share of informed investors is fixed and the price is the same across both states, the two endogenous variables determined by these equations are the common price P and uninformed investors' inferred default probability $\tilde{\kappa}$.

Since the uniformed default inference is an equilibrium object, it must be consistent

with the slopes of the price functions at P. In particular, the measure of supply shocks within Δ of the price P is given by $\mathcal{M}(P, \kappa, \Delta) := \overline{\psi} - \underline{\psi}$ where $P(\kappa, \underline{\psi}) = P + \Delta$ and $P(\kappa, \overline{\psi}) = P - \Delta$. Taking the limit as Δ approaches zero then yields the inference

$$\tilde{\kappa}(P) = \lim_{\Delta \to 0} \frac{\mathcal{M}(P, g, \Delta)\kappa_g + \mathcal{M}(P, b, \Delta)\kappa_b}{\mathcal{M}(P, g, \Delta) + \mathcal{M}(P, b, \Delta)}. \tag{26}$$

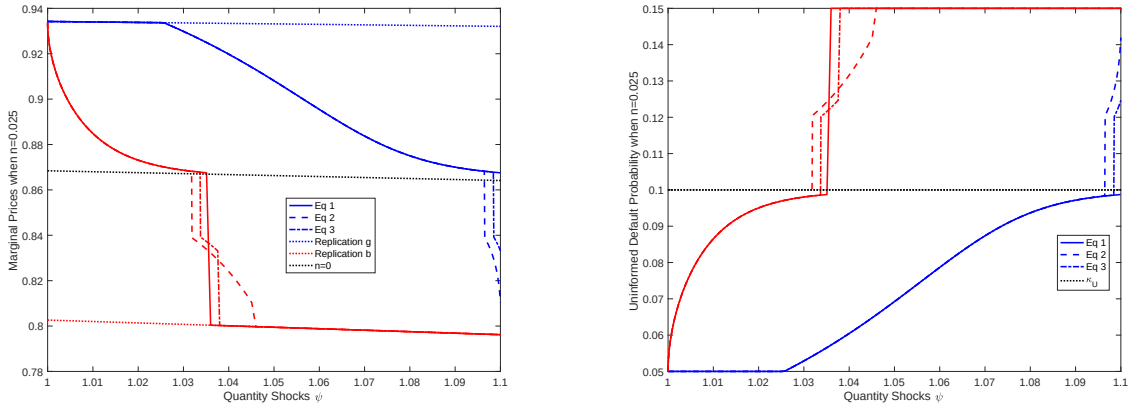
This means that the flatter price schedule receives more weight.

Solution method for the uniform protocol. Solving the investor's portfolio problem formally requires evaluating all of the marginal probability-weighted utilities in the acceptance set in order to determine the appropriate marginal condition (8). The special feature of the UP auction is that we only need to ensure that the *local optimality conditions* are satisfied everywhere if the short-sale constraint does not bind. If, instead, the short-sale constraint binds in a range of states, then the bids over the binding range are *zero*. Hence, for the UP protocol, we incrementally solve the equilibrium conditions as ψ increases. When the bid-overhang constraint does not bind, and the price schedules are distinct, this is fairly straightforward. When the price schedules overlap, either directly or because of the bid-overhang constraint, the uninformed investors have an additional inference problem, which depends on the local slope of the two price schedules. The solution, in this case, is obtained by solving an ODE system as in (26). We describe our solution methods in more detail in Appendix D.

Remark on the possibility of indeterminacy: When the bid-overhang constraint binds, there may be equilibrium indeterminacy. There is a fundamental discontinuity in the bidding behavior of the informed depending on whether their short-sale constraint binds or not. Initially, at high prices, it must be slack, and the uninformed attach probability one to the high-quality shock, κ_g. The two price schedules may coincide, and the first point at which this happens is when the bids of the uninformed are sufficient to meet demand on the low-quality schedule. But the short-sale constraint will typically begin to bind on the informed thereafter, and this can allow for another equilibrium with a discontinuously lower price in which the short-sale is slack. The local equilibrium is conditioned upon

the pair quantity shocks (ψ_g, ψ_b) at which the prices match; i.e. $P(\kappa_g, \psi_g) = P(\kappa_b, \psi_b)$. Conditional on the binding pattern, the equilibrium is locally unique. The common price at which we switch from one binding pattern to another is not generally unique. As a result, there is a wide range of indeterminacy.

Illustration: Figure 6 numerically computes a set of equilibria for which the bid-overhang constraint binds, this is $n < 1 - \frac{1}{\psi_M}$, and prevents perfect replication. This situation forces overlapping price schedules, and we illustrate indeterminacy by showing three possible equilibria: one in which the short sale constraint always binds (solid lines) and two in which the short sale constraint binds initially for a subset of common prices and that differ in the point at which they stop binding (dashed cases). These last two cases illustrate the sources of indeterminacy. The first panel of figure 6 plots the marginal prices for these cases. For reference, we also plot the equilibrium with perfect replication (or $n > 1 - \frac{1}{\psi_M}$) and the equilibrium with no information ($n = 0$) in dashed lines. The second panel plots the associated inferred default probabilities that sustain the equilibrium overlapping prices.



(a) Marginal Prices

(b) Uninformed Expected Default Probability

Figure 6: Asymmetric information about quality shocks when $n = 0.04$. Parameters: $u(c) = \log(c)$, $W = 250$, $D = 60$, two equally likely default probabilities $\kappa_b = 0.15$, $\kappa_g = 0.05$, and unknown quantity shocks, $\psi \in [1.0, 1.1]$.

4.2 Equilibrium in the Discriminatory-Price Auctions

Next, we turn to the discriminatory-price protocol. As we discussed, the investor's portfolio problem inherently involves a simultaneous equation system since the set of all bid quantities and prices jointly determine bond expenditures and consumption conditional on the state and default outcome. Solving this system is relatively straightforward if the quality-contingent price schedules are distinct so that the inference problem of uninformed investors is trivial and the number of possible realizations of the quantity shock is small so the set of equations is not too large. It is much more difficult when there is overlap in the price schedule so that marginal prices are not fully revealing of the quality shock.[15] To be consistent with the uniform-price environment analyzed before, we develop a new solution method that allows for continuous quantity shocks.

Solution method for the discriminatory protocol. We solve this problem using a new solution method in which we approximate the price and bid functions using, respectively, linear interpolations and step-functions on a common set of grid points in (κ, ψ) space. This particular form of approximation method has the advantage that if markets clear at the grid points, then they also clear for all ψ in between the grid points, holding fixed the quality shock. We then evaluate the investors' FOC conditions at points in the state space. Thus we are solving discriminatory-price auctions using standard approximation function methods. We defer the details to Appendix D.

Illustration. Figure 7 shows equilibrium price schedules for different values of n at the benchmark parameters from our previous numerical example used for uniform auctions. The symmetric full information case, $n = 1$, is shown by the extreme solid schedules. The high-quality price schedule declines as n falls, while the low-quality schedule is unchanged because the uninformed only bid for low-quality bonds. This leads to an inefficient concentration of risk that stems from the uninformed not participating in buying high-quality bonds, which initially drives down their price when there are less informed investors bidding for them. When n becomes low enough, and the price of high-quality

[15] This mirrors issues in competitive equilibrium when prices may not be fully revealing; see Allen (1981).

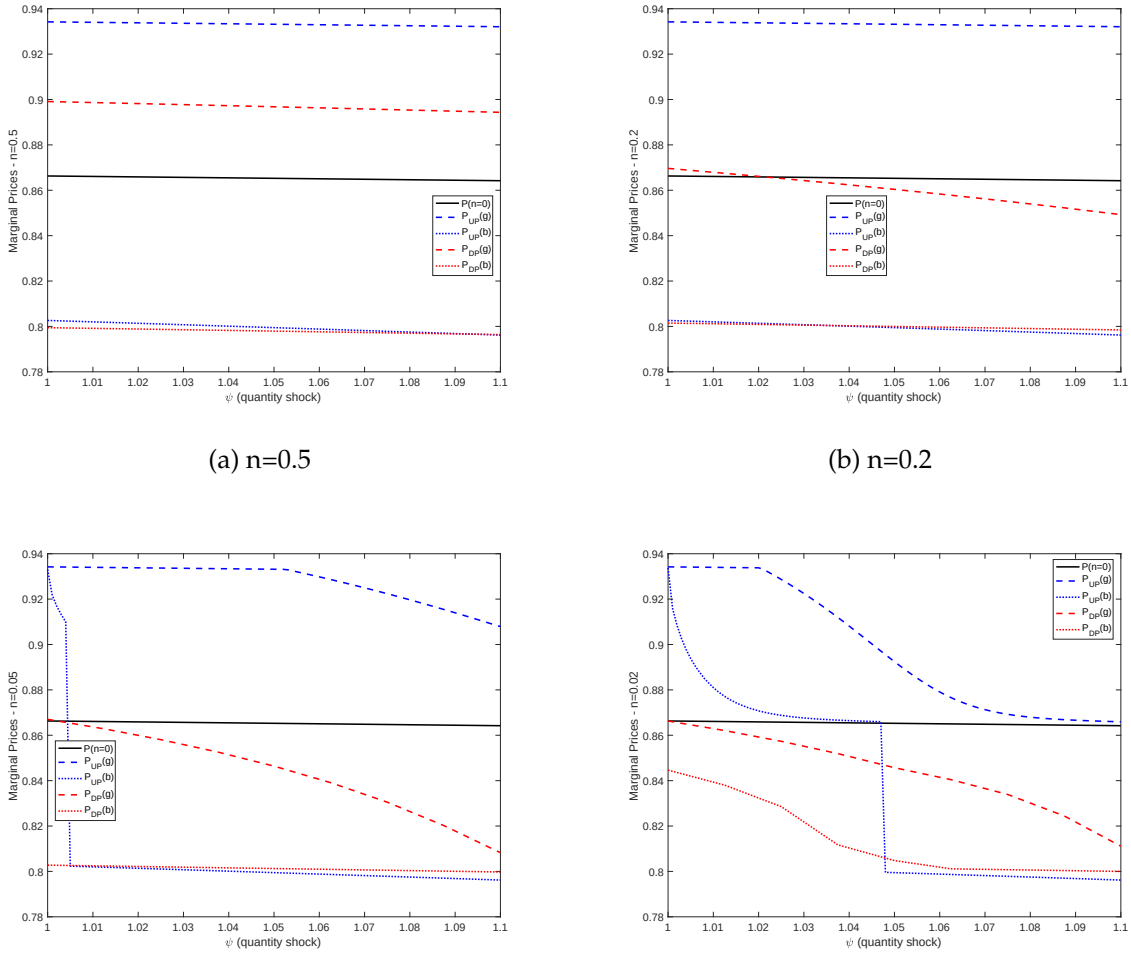
bonds has declined enough, the uninformed also start bidding for them, which raises the low-quality schedule because their high-price purchases (which are also executed in the bad state) reduce the amount that must be sold to the informed. However, investor's remorse is so strong that as n falls and the uninformed make a larger share of demand for the high-quality bond, its price falls below that of the bond when no one is informed (the black schedule that represents the case $n = 0$).

Finally, note that the price function only overlaps for quite low values of $n < 0.05$, as it was the case under uniform auctions. In the narrow band between $0.05 \to 0$, the prices are very sensitive to n. The schedules with $n = 0.02$ that we used as an illustration of the non-replicating equilibrium under the uniform auction are graphed in the last panel.

The comparison of total debt burden relative to the amount raised between UP and DP protocols come from taking weighted averages of the bonds sold to informed and uninformed over quality and quantity shocks, given n. Figure 8 shows that the ranking we identified in the situation with perfect information about the quantity shocks remains, with two differences. First, the ex-ante debt burden relative to the revenues obtained is higher due to additional uncertainty and lower prices implied by $\psi \in [1, 1.1]$ instead of a certain $\psi = 1$. Second, there is a gradual convergence of prices as $n \to 0$, forced by the overlapping prices induced by the bid-overhang constraint in the UP case. The comparison of the coefficient of variation across these two protocols also applies.

5 Endogenous information acquisition

We have characterized the performance of the two protocols in the presence of quantity and quality shocks, assuming the degree of asymmetric information, n, is exogenous. We have shown that equilibrium prices in both UP and DP protocols overlap in a relatively small range of asymmetric information when most investors are uninformed and only when there is asymmetric information about quality shocks and no information about quantity shocks. Except in this case and small range of asymmetric information, prices do not overlap and are ex-post perfectly revealing of the state, with DP generating a larger debt burden but a smaller variance. But given these properties, what are the gains



(a) n=0.5

(b) n=0.2

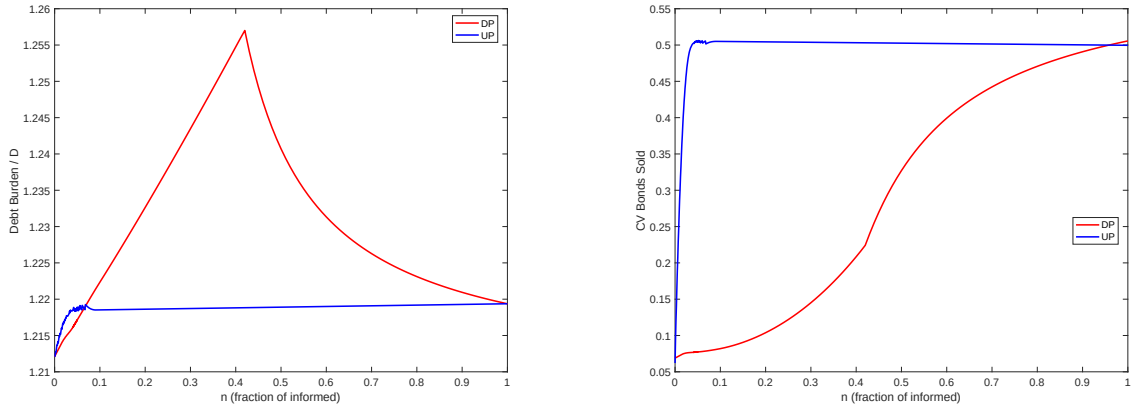
(c) n=0.05

(d) n=0.02

Figure 7: UP and DP marginal prices under different n values. Asymmetric information about quality shocks. Parameters: $u(c) = \log(c)$, $W = 250$, $D = 60$, two equally likely default probabilities $\kappa_b = 0.15$, $\kappa_g = 0.05$, and unknown quantity shocks, $\psi \in [1.0, 1.1]$.

of information? What is the endogenous n if acquiring information is costly?

Endogenizing the extent of asymmetric information implies extending the definition of equilibrium by allowing information acquisition at a utility cost C. Denoting by V^i the investor i's payoff that solves the portfolio choice problem as discussed above, the extended definition is,



(a) Debt Burden / D

(b) Coefficient of Variation Debt Burden

Figure 8: Ex-ante Debt Burden and its coefficient of variation. Parameters: $u(c) = \log(c)$, $W = 250$, $D = 60$, two equally likely default probabilities $\kappa_b = 0.15$, $\kappa_g = 0.05$, and unknown quantity shocks, $\psi \in [1.0, 1.1]$.

Definition 4 (Information acquisition problem)**.** *The information acquisition problem is*

$$\max_{i \in \{I,U\}} V^i - \mathbb{I}_I C. \tag{27}$$

where $\mathbb{I}_I = 1$ *if* $i = I$ *and* $\mathbb{I}_I = 0$ *if* $i = U$.

Figure 9 shows the payoffs for the informed and the uninformed for both of our protocols in our numerical benchmark with asymmetric information about quality shocks and no information about quantity shocks. When $n = 0$, both protocols generate similar gains of information, from adjusting quantities to buying more high-quality bonds and less low-quality bonds at the uninformed equilibrium price. As n increases, the relative payoff to the informed under UP drops sharply, dissipating when the uninformed can perfectly replicate the informed portfolio. For DP, in contrast, the winner's curse effect is initially so strong that the informed payoffs rise with n, before competition among informed investors eventually takes over. As a result, the relative benefits of information (the gap between the payoffs of informed and uninformed) are hump-shaped with DP.

This taxonomy of relative benefits of information is quite robust to the specific parameters. If the cost of information was high (above 5.545 in our numerical illustration), there

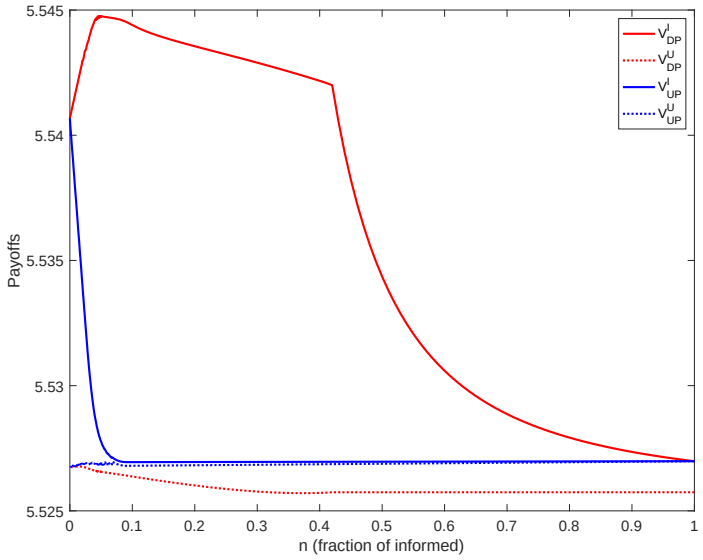


Figure 9: Payoffs for informed and uninformed investors in both protocols for different n values. Parameters: $u(c) = \log(c)$, $W = 250$, $D = 60$, two equally likely default probabilities $\kappa_b = 0.15$, $\kappa_g = 0.05$, and unknown quantity shocks, $\psi \in [1.0, 1.1]$.

would be no informed ($n = 0$) under both protocols, and they would generate the same outcome for both protocols. If the cost was sufficiently high to preclude information acquisition under UP but not under DP (between 5.542 and 5.545), DP would display either few or many informed, but in both cases, it would imply both higher debt burden and variance than UP, making UP protocols strictly preferred. When the information cost is low enough (between 5.527 and 5.542) there are informed investors under both protocols.

In this last case, there is a particularly interesting set of implications. First, DP auctions attract many more informed investors than UP auctions. Second, DP auctions induce a higher, but not necessarily more volatile, debt burden. Finally, and perhaps more surprisingly, DP auctions fully reveal the state, but UP auctions do not, as price schedules overlap. The extent of state revelation from prices can then provide clues about whether information asymmetry is exogenous or endogenous in auctions, which will qualify the comparison across auction protocols. In the next section, we test this possibility in the context of Mexican sovereign bond auctions.

5.1 Empirical evidence of endogenous information.

We have just shown that, when information acquisition is endogenous, discriminatory auctions reveal more information about quality shocks than uniform auctions when there is no ex-ante information about quantity shocks. We now evaluate this prediction in the context of the Mexican bond market, where the government changed the protocol from discriminatory to uniform in 2017.

We study auctions of Mexican Federal Treasury Bills (Cetes), which are domestic- denominated zero-coupon pure discount bonds with typical maturities of 28, 91, 182, and 364 days. These are the leading instruments in Mexican money markets and the main source of federal government debt funding since 1978. The primary market for Cetes consists of a closed system (allocations and prices are only disclosed after the auction is closed) public auctions conducted by the Mexican central bank, which acts as a financial broker. Auctions take place weekly, almost always on Tuesdays from 10 a.m. to 11 a.m.; results are disclosed at 11:30 a.m. and published on the Bank of Mexico website two days later, usually on Thursdays. Cetes were auctioned following a discriminatory price protocol until October 2017, when the Bank of Mexico switched to a uniform price protocol. Our sample consists of data from May 2003 to June 2024. In what follows, we present results for 28-day Cetes, but similar results are obtained for the other maturities. In this sample we observe 736 DP auctions and 345 UP auctions of 28-day Cetes.

In general, it is not trivial to identify how much information about the quality shock is revealed in primary markets. Hence we base our analysis on the empirical strategies devised in Cole et al. (2022) and Cole et al. (2024), where we exploit the relation between prices in primary and secondary markets.[16] Our previous work did not consider variation in the auction protocol and thus did not conduct the same tests. We measure the reaction of bond prices in secondary markets right after the disclosure of bond prices in primary markets. If agents can infer quality shocks from primary prices, then secondary markets should react more strongly. We use secondary yields at each business day closing

[16]Here we treat secondary markets as independent of the functioning of auctions and use them as a laboratory to see how information flows from one to the other. In Cole et al. (2024), we explore how the presence of secondary markets affects the equilibrium in discriminatory auctions and find that they tend to, if anything, increase the value of information.

(2:00pm) as reported by the Bank of Mexico through PIP (Proveedor Integral de Precios) and Valmer (Valuacion Operativa y Referencias de Mercado).

The reaction of secondary markets at close on auction days (i.e., a few hours after the results of auctions are disclosed) allows us to infer how much information about the quality of bonds can be inferred from the unexpected movements in primary market prices. For example, if price schedules at auction do not overlap, then investors can infer the realized quality shock from the auction price; if the price schedules do overlap then auction prices provide only a noisy signal about the state.

Our empirical setting allows us to gauge the relative importance of quantity and quality shocks. If secondary prices react strongly to auction outcomes under both DP and UP protocols, this would suggest that quantity shocks are relatively important shocks, as they are always revealed through auction results. Conversely, if secondary markets react strongly to auctions under the DP protocols but less so for the UP protocol, this would suggest that quality shocks are relatively more important.

We perform two complementary exercises. First, we obtain the elasticity of unexpected changes in secondary prices with respect to unexpected changes in auction prices. We do this by regressing the unexpected log change of secondary prices ($\Delta \log \text{Sec}_t$) on the unexpected log change on primary prices ($\Delta \log \text{Prim}_t$) on auction days t.[17] We also distinguish between the elasticities that are obtained under DP or UP protocols.

Table 1 shows this elasticity is 0.81 when 28-day maturity Cetes are auctioned using DP protocols, and highly statistically significant. This implies that when bonds are auctioned with DP, a 1% unexpected increase in primary prices is associated with a 0.81% increase in the secondary market price of the day the auction takes place. When the same bonds were auctions with UP protocols this elasticity reduced quite drastically, almost by half, to 0.42, both statistically significantly different than 0 and different than the elasticity with DP protocols. This result is the first evidence that UP protocols convey less information than DP auctions, as predicted by our setting.

In a second exercise, we construct a measure of auction price informativeness by com-

[17] We obtain the unexpected changes as deviations between the observed prices and the prices that are expected based on the secondary prices that were realized during the week before the auction. In the case of primary prices, we also use the previous auction price to construct their expectation.

Table 1: Elasticity of secondary market prices to information released at auction

Dependent Variable	$\Delta \log \text{Sec}_t$
$\Delta \log \text{Prim}_t$ (DP)	0.806***
	(0.021)
$\Delta \log \text{Prim}_t$ (UP)	0.418***
	(0.027)
Constant	-0.019***
	(0.0015)
Number Auctions	1,056
R^2	0.62

puting the reduction in the conditional variance of secondary market prices that arise from observing primary market outcomes during auction days. Specifically, we compute the share of unexplained variance of auction-day secondary market yields that can be accounted for by auction prices in addition to pre-auction secondary market yields. This statistic is called *marginal* R^2, and it is formally given by

$$\Delta R^2 = \frac{R^2_{(S_{t-1},P_t)} - R^2_{(S_{t-1})}}{1 - R^2_{(S_{t-1})}},$$

where $R^2_{(S_{t-1})}$ is the R^2 of regression of secondary yields reported at market close on auction days on three lags of pre-auction secondary market yields, and $R^2_{(S_{t-1},P_t)}$ is the R^2 of the same regression but including primary market yields observed the same day at the auction, as an additional regressor. Table 2 shows that DP auctions are more informative events than their UP counterparts in the sense that their price are more relevant in explaining the variance observed in secondary market prices. While 72% of the variance of secondary prices that cannot be explained by the previous week's realizations can be explained by auction prices under DP protocols, only 29% can be explained by auction prices under UP protocols.

Table 2: Marginal R^2. 28-day Cetes

Auction Protocol	**DP**	**UP**
Marginal R^2	0.723	0.291
Number Auctions	*735*	*345*

An additional implication of our theory is that, ceteris paribus, the ex-ante variance of prices under UP protocols is higher than under DP protocols. Directly comparing

auction price variances during periods with different protocols may be misleading given that fundamentals may display different volatility properties: a lower variance of prices during DP protocol may be just the result of a low variance in default probabilities prior to 2017. To control for changes in the environment, we compare variances relative to the price variance of other government bonds that did not change the protocol in the sample.

Bondes D (Bonos de Desarrollo del Gobierno Federal) are issued by the Mexican government in maturities that range from 1- to 5-years, and in contrast to Cetes, accrue interest in pesos every month. Our strategy is to consider auction days in which both 1-year Cetes and 1-year Bondes D were auctioned. This comparison is then controlled by information sets (public regimes) about Mexico during auction days and standardized by maturity considerations. While 1-year Cetes were auctioned monthly until Covid (and weekly since June 2020) and using DP protocol only until October 2017, 1-year Bondes D were auctioned weekly and using a DP protocol throughout. Bondes D were discontinued in September 2021 and replaced by Bondes F, which are fungible in that they allow for reopenings. Given these restrictions, our sample consists of 38 auctions in which both bonds were auctioned the same day under a DP protocol and 102 auctions in which Bondes D was auctions with a DP protocol and Cetes with UP protocols.

Table 3 shows the results. While in the period after October 2017, fundamentals were such that the coefficient of variation of Bondes D declined dramatically, suggesting that underlying fundamentals became less volatile, the coefficient of variation for Cetes, which changed the protocol from DP to UP, indeed increased.

Table 3: Coefficient of Variation

Bond	**One-year Cetes**	**One-year Bondes D**	**Number Auctions**
Both DP protocol	0.0141	0.00026	38
Cetes UP protocol	0.0143	0.00022	102
Change	+1.6%	-17.4%	

6 Conclusion

Governments raise funds by auctioning debt in primary markets. The efficient functioning of these markets is more urgent in periods of instability and high government debt like most countries are currently experiencing. We offer a theoretical framework for analyzing how information asymmetries among risk-averse investors interact with auction protocols to determine the costs and riskiness of government funding. The choice between standard protocols is determined by tradeoff between surplus extraction and bidder participation. A new protocol circumvents this tradeoff and has the potential to decrease the level and volatility of government financing costs. Our empirical analysis confirms that sovereign debt auctions are information events and suggests that the auction protocol is an important factor in information transmission.

References

Aguiar, M. and M. Amador (2014). Sovereign debt. *Handbook of International Economics 4*, 647–687.

Aguiar, M., M. Amador, H. Hopenhayn, and I. Werning (2019). Take the short route: Equilibrium default and debt maturity. *Econometrica 87*(2), 423–462.

Aguiar, M., S. Chatterjee, H. Cole, and Z. Stangebye (2016). Quantitative models of sovereign debt crises. *Handbook of Macroeconomics 2*, 1697–1755.

Aguiar, M. and G. Gopinath (2006). Defaultable debt, interest rates and the current account. *Journal of International Economics 69*(1), 64–83.

Alfaro, L. and F. Kanczuk (2009). Optimal reserve management and sovereign debt. *Journal of International Economics 77*(3), 1105–1131.

Allen, B. (1981). Generic existence of completely revealing equilibria for economies with uncertainty when prices convey information. *Econometrica*, 1173–1199.

Allen, J., A. Hortacsu, E. Richert, and M. Wittwer (2024). Entry and exit in treasury auctions. *Bank of Canada Staff Working Paper 2024-29*.

Alves Monteiro, R. and S. Fourakis (2024). Sovereign debt auctions with strategic interactions. *IMF Working paper*.

Arellano, C. (2008). Default risk and income fluctuations in emerging markets. *American Economic Review 98*(3), 690–712.

Arellano, C., Y. Bai, and G. Mihalache (2024). Deadly debt crises: Covid-19 in emerging markets. *Review of Economic Studies 91*(3), 1243–1290.

Armstrong, M. (2016). Nonlinear pricing. *Annual Review of Economics 8*, 583–614.

Ausubel, L. M., P. Cramton, M. Pycia, M. Rostek, and M. Weretka (2014). Demand reduction and inefficiency in multi-unit auctions. *Review of Economic Studies 81*(4), 1266–1400.

Bai, Y., P. J. Kehoe, P. Lopez, and F. Perri (2025). A neoclassical model of the world financial cycle. *NBER Working Paper 33441*.

Barbosa-Alves, M., J. Bianchi, and C. Sosa-Padilla (2024). International reserve management under rollover crises. *NBER Working paper 32393*.

Barro, R. J. (2003). Optimal management of indexed and nominal debt. *Annals of Economics and Finance 4*, 1–15.

Bianchi, J., J. C. Hatchondo, and L. Martinez (2018). International reserves and rollover risk. *American Economic Review 108*(9), 2629–2670.

Bigio, S., G. Nuño, and J. Passadore (2023, March). Debt-maturity management with liquidity costs. *Journal of Political Economy Macroeconomics 1*(1).

Bocola, L. and A. Dovis (2019). Self-fulfilling debt crises: A quantitative analysis. *American Economic Review 109*(12), 4343–4377.

Boyarchenko, N., D. Lucca, and L. Veldkamp (2021). Taking orders and taking notes: Dealer information sharing in treasury markets. *Journal of Political Economy 129*, 607–645.

Brenner, M., D. Galai, and O. Sade (2009). Sovereign debt auctions: Uniform or discriminatory? *Journal of Monetary Economics 56*(2), 267–274.

Chatterjee, S. and B. Eyigungor (2012). Maturity, indebtedness, and default risk. *American Economic Review 102*(6), 2674–2699.

Chaumont, G. (2020). Sovereign debt, default risk, and the liquidity of government bonds. Working Paper, University of Rochester.

Choi, J., Q. Q. Dang, R. Kirpalani, and D. J. Perez (2024). Exorbitant privilege and the sustainability of us public debt. *NBER Working Paper 32129*.

Choi, J., R. Kirpalani, and D. J. Perez (2024). Us public debt and safe asset market power. *NBER Working Paper 30720*.

Cole, H., D. Neuhann, and G. Ordonez (2022). Asymmetric information and sovereign debt: Theory meets mexican data. *Journal of Political Economy 130*(8), 2055–2109.

Cole, H., D. Neuhann, and G. Ordoñez (2022). Sovereign debt auctions in turbulent times. In *AEA Papers and Proceedings*, Volume 112, pp. 526–530.

Cole, H. L., D. Neuhann, and G. Ordonez (2024). Information spillovers and sovereign debt: Theory meets the eurozone crisis. *Review of Economic Studies 92*, 197–237.

Eaton, J. and M. Gersovitz (1981). Debt with potential repudiation: Theoretical and empirical analysis. *The Review of Economic Studies 48*(2), 289–309.

Engelbrecht-Wiggans, R. (1998). Revenue equivalence in multi-object auctions. *Economics Letters 26*, 15–19.

Fudenberg, D., M. Mobius, and A. Szeidl (2007). Existence of equilibrium in large double auctions. *Journal of Economic theory 133*(1), 550–567.

Gupta, S. and R. Lamba (2024). Treasury auctions during a crisis. *Working paper, Cornell and NYU*.

Horn, S., D. Mihalyi, P. Nickol, and C. Sosa-Padilla (2024). Hidden debt revelations. *NBER Working Paper 32947*.

Hortacsu, A. and D. McAdams (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. *Journal of Political Economy 118*(5), 157–184.

Hortacsu, A. and D. McAdams (2018). Empirical work on auctions of multiple objects. *Journal of Economic Literature 56*(1), 157–184.

International Monetary Fund (2024, October). Fiscal monitor: Putting a lid on public debt. Prepared by IMF staff, published on October 23, 2024.

Jiang, Z., H. Lustig, S. Van Nieuwerburgh, and M. Z. Xiaolan (2024). The u.s. public debt valuation puzzle. *Econometrica 92*(4), 1309–1347.

Malvey, P. F., C. M. Archibald, and S. T. Flynn (1995). Uniform-price auctions: evaluation of the treasury experience. *Office of Market Finance, US Treasury*.

Maskin, E. S. and J. G. Riley (1985). Auction theory with private values. *The American Economic Review 75*(2), 150–155.

Morelli, J. M., P. Ottonello, and D. J. Perez (2022). Global banks and systemic debt crises. *Econometrica 90*(2), 749–798.

Passadore, J. and Y. Xu (2022, July). Illiquidity in sovereign debt markets. *Journal of International Economics 137*, 103618.

Reny, P. J. and M. Perry (2006). Toward a strategic foundation for rational expectations equilibrium. *Econometrica 74*(5), 1231–1269.

Wittwer, M. and J. Allen (2025). Market power and capital constraints. *Bank of Canada Staff Working Paper 2023-51*.

A Proofs

A.1 Proof of Lemma 1

The result follows from the convexity of marginal utility and the fact that $P(s) \in [0, 1 - \kappa(s))$, which holds under risk-averse preferences. □

A.2 Proof of Lemma 2

Consider first the bid-overhang constraint. Because $P(s_1) > P(s_2)$, this constraint may bind in state s_2 only. For informed investors, $\mathcal{E}^I(s) = s$ for all s. For uninformed investors, $\mathcal{E}^U(s_1) = s_1$ and $\mathcal{E}^U(s_2) = \{s_1, s_2\}$. From bidding optimality, it follows that $B_j^I(s) \geq B_j^U(s)$ for any state s and any protocol j. Hence $X^U(s_1) \leq X^I(s_1)$ and, by market clearing, $X^U(s_1) < \psi(s_1)D$ if $n > 0$. In the two-state model with quality shocks, $\psi(s_1) = \psi(s_2)$, while in the two-state model with quantity shocks, $\psi(s_1) < \psi(s_2)$ in order to ensure $P(s_1) > P(s_2)$. Hence $X^U(s_1) < \psi(2)D$, and the bid-overhang constraint holds.

Next, consider the short-sale constraint. We have already argued that $B_j^I(s) \geq B_j^U(s)$ for any state s and any protocol j. Now suppose for a contradiction that the short-sale constraint on informed investors' binds in state s. Then $B_j^I(s) = B_j^U(s) = 0$, violating market clearing. □

A.3 Proof of Proposition 1

Consider first variant with quality shocks. If no investor is informed about the realized quantity shock, no bids can be made contingent on the quality shock. Hence marginal prices cannot be contingent on the state, and optimal uninformed bids are determined as if there were a single state in which the default probability is the unconditional probability $\bar{\kappa}$. Hence optimal bids for any investor solve the same first-order condition as in (13), with $\kappa(s)$ replaced by $\bar{\kappa}$. Since all investors submit bids at only one price, the differential treatment of bids at infra-marginal prices across protocols does not affect equilibrium outcomes.

Next, consider the equilibrium with quantity shocks. Assume for a contradiction that bids and prices are not state-contingent. Then since $\psi(1) \neq \psi(2)$, market clearing cannot simultaneously hold in both states. Hence prices must be state-contingent. Since infra-marginal uninformed bids are executed at different prices in the two protocols, equilibrium outcomes must also differ by protocol. □

A.4 Proof of Proposition 2

(i) Consider the bidding problem of an uninformed investor whose information set contains K discrete states of the world $s = (\kappa, \psi)$ with associated default probability $\kappa(s)$. Without loss, order these states in descending order of price, $P(s_1) > P(s_2) > \ldots > P(s_K)$. Further, assume for now that the short-sale constraint is slack. The acceptance set for any state then contains all states with lower marginal prices.

Rewriting (8), condition (9) at state s_K yields

$$F^i(s_K) \equiv Pr(s_K)\Big[\kappa(s_K)m^i(s_K,1)P(s_K) - (1-\kappa(s_K))m^i(s_K,1)(1-P(s_K))\Big] = 0,$$

where $Pr(s_K)$ denotes the probability of state s_K. Repeating for state s_{K-1}, the state with the second-lowest marginal price, we have

$$Pr(s_{K-1})\Big[\kappa(s_{K-1})m^i(s_{K-1},1)P(s_{K-1}) - (1-\kappa(s_{K-1}))m^i(s_{K-1},1)(1-P(s_{K-1}))\Big] + F^i(s_K) = 0.$$

Since $F^i(s_K) = 0$ by the previous condition, the system of equations can be solved recursively and state by state.

(ii) Follows directly from the construction of the replicating strategy above. The construction fails if informed bids are not ordered by price. In such cases, uninformed investors could replicate the informed portfolio only by bidding negative quantities in some states, which is ruled out by the short-sale constraint.

□

A.5 Proof of Proposition 3

Guess and verify that informed quantities are ordered $B^I(s_1) < B^I(s_1)$ so that uninformed investors can perfectly replicate the informed portfolio. Then the market-clearing condition in state s is $B^I(s)P(s) = \psi(s)D$. Using the informed bid function (13) yields the stated quantities and prices. Since $\psi(1) < \psi(2)$, quantities are indeed strictly ordered by price, verifying the conjecture. □

A.6 Proof of Proposition 4

We have already characterized the equilibrium of the uniform auction in closed form. In particular, for any $n > 0$, the marginal price in state s is $P^*_{UP}(s) = 1 - \frac{\kappa(s)}{1-\psi(s)\frac{D}{W}}$ and the expected marginal price is $\mathbb{E}[P^*_{UP}(s)] = 1 - \frac{\bar{\kappa}}{1-\psi(s)\frac{D}{W}}$. With this in hand, we now prove each statement in turn.

1. Consider the discriminatory auction in the limit as $n \to 0$. I By market clearing $\lim_{n\to 0} X^U_{DP}U(s) = \psi D$. Since all uninformed bids at $P(s_1)$ are also accepted in state s_2, $\lim_{n\to 0} B^U_{DP}(s_2) = 0$. Evaluating the first-order condition for optimal bids in this limit yields that the optimal uninformed bid in state s is $B^*_{DP}(s_1) = \frac{\psi D}{P^*_{DP}}$ where

$$P^*_{DP} = \mathbb{E}[P^*_{UP}(s)] = 1 - \frac{\bar{\kappa}}{1-\psi(s)\frac{D}{W}}.$$

Hence the quantity-weighted average price is the same in both states, and so is the quantity of bonds sold. Hence, the volatility of bonds sold is zero in the discrimina-

tory auction but strictly positive in the uniform auction. Furthermore, by Jensen's inequality, we have

$$\mathbb{E}[B^*_{UP}(s)] = \mathbb{E}[\frac{\psi D}{P^*_{UP}(s)}] > \frac{\psi D}{\mathbb{E}[P^*_{UP}(s)]} = \frac{\psi D}{P^*_{DP}} = \mathbb{E}[B^*_{DP}(s)]. \tag{28}$$

2. We prove this statement by constructing an equilibrium in which the winner's curse induces the short-sale constraint on $B^U_{DP}(s_1)$ to bind. (Note that a binding short-sale constraint is not necessary.) Guess and verify that $B^U_{DP}(s_1) = 0$. Under this conjecture, $B^U_{DP}(s_2) = B^I_{DP}(s_2)$ and so $B^{U*}_{DP}(s_2) = B^*_{UP}(2)(s_2)$. It follows that $P^*_{DP}(s_2)$. Since only informed investors bid at the high price, the equilibrium marginal price is $P^*_{DP}(s_1) = 1 - \frac{\kappa(s_1)}{1-\psi\frac{D}{nW}}$. Hence for any $n < 1$, $P^*_{DP}(s_1) < P^*_{UP}(s_1)$. This implies $B^*_{DP}(s_1) > B^*_{UP}(s_2)$ and $\mathbb{E}[B^*_{DP}(s)] > \mathbb{E}[B^*_{UP}(s_)]$. We now need to verify that the short-sale constraint binds. A sufficient condition is $P^*_{DP}(s1) > 1 - \hat{\kappa}$, or

$$\hat{\kappa} > \frac{\kappa(s_1)}{1-\psi\frac{D}{nW}} \quad \Leftrightarrow \quad nPr(s_2)(\kappa(s_2) - \kappa(s_1)) > \frac{\psi D}{W}$$

3. Follows trivially from the fact that informed investors bid at the marginal price only. Hence the treatment of bids at inframarginal prices is irrelevant.

□

A.7 Proof of Proposition 5

We first characterize the solution to the partially discriminating auction. The optimal bidding strategy of the informed is the same for all three protocols. If τ is not binding, then the optimal bidding strategy of uninformed investors is determined by the two equations

$$B^U(s_1): \qquad \kappa(s_1)P_{PD}(s_1)u'\Big(\tilde{W}_{PD}(s_1)\Big) = (1-\kappa(s_1))(1-P_{PD}(s_1))u'\Big(\tilde{W}_{PD}(s_1) + B_{PD}(s_1)\Big) \tag{29}$$

$$B^U(s_2): \qquad \kappa(s_2)P_{PD}(s_2)u'\Big(\tilde{W}_{PD}s_2)\Big) = (1-\kappa(s_2))(1-P_{PD}(s_2))u'\Big(\tilde{W}_{PD}(s_2) + \sum_s B_{PD}(s)\Big). \tag{30}$$

The first-order condition for $B^U_{PD}(s_1)$ has the same form as that for $B^I(s_1)$, hence $B^U_{PD}(s_1) = B^I_{UP}(s_1)$ for any n, which implies $B^U_{PD}(s_1) < B^U_{DP}(s_1)$. Hence it follows that

$$\mathcal{D}_{PD}(s_1) = \mathcal{D}_{UP}(s_1) < \mathcal{D}_{DP}(s_1) \qquad \text{for all } n \in (0,1).$$

Next consider the low state. Informed expenditures invariant in the protocol,

$$X^I(s_2) = \frac{1-\kappa(s_2)-P_j(s_2)}{(1-P_j(s_2))}W = \left[1 - \frac{\kappa(s_2)}{1-P_j(s_2)}\right]W,$$

which is strictly decreasing in $P_j(s_2)$.

The market-clearing condition in any protocol is $(1-n)X_j^U(s_2) + nX^I(s_2) = \psi D$, and so

$$X_j^U(\psi) = G(P_j(s_2)) \equiv \frac{1}{1-n}\psi(s_2)D - \frac{n}{1-n}\left[1 - \frac{\kappa(s_2)}{1-P_j(s_2)}\right]W,$$

where G is a strictly increasing function of $P_j(s_2)$ that is invariant in the protocol. Hence uninformed investors' risk-free-asset position is invariant in the protocol conditional on the price,

$$\tilde{W}_j^U(s_2) = W - G(P_j(s_2))$$

and is *strictly decreasing* in $P_j(s_2)$. Hence for any protocol, the first-order condition can be written as

$$\kappa(s_2)P_j(s_2)u'\Big(\tilde{W}_j^U(s_2))\Big) = (1-\kappa(s_2))(1-P_j(s_2))u'\Big(\tilde{W}_j^U(s_2) + \mathcal{B}_j^U(s_2)\Big).$$

Under log utility, this can be solved out for the total quantity of bonds purchased by uninformed investors,

$$\mathcal{B}_j^U(s_2) = \left(\frac{1-\kappa(s_2)}{P_j(s_2)} - 1\right)\frac{\tilde{W}_j^U(s_2)}{\kappa(s_2)} \tag{31}$$

where the right-hand side is the same for any protocol.

Now compare UP and PD protocols. Let $\tau_{UP} = 0$ and $\tau_{PD} > 0$. By the rules of the two protocols,

$$X_j^U(s_2, \tau_j) = (P_j(s_1) - P_j(s_2))\tau_j + P_j(s_2)\mathcal{B}_j^U(s_2). \tag{32}$$

Combining this with equation (31) yields

$$W - \left(\frac{1-P_j(s_2)}{\kappa(s_2)}\right)\tilde{W}_j(s_2) = \Big(P_j(s_1) - P_j(s_2)\Big)\tau_j, \tag{33}$$

where the LHS is the same for UP and PD protocols and is strictly increasing in the price $P_j(s_2)$. The RHS is strictly positive for the PD auction and zero for the UP auction. Hence the PD auction achieves a lower debt burden state s_2. Since the two protocols generate the same debt burden in state s_1, the PD auction achieves a lower debt burden in expectation.

Next, compare the PD and DP auctions given s_2. Uninformed expenditures satisfy

$$X_{PD}^U(s_2, \tau) = P_{PD}(s_1)\tau + P_{PD}(s_2)(\mathcal{B}_{PD}^U(s_2) - \tau). \tag{34}$$

$$X_{DP}^U(s_2) = P_{DP}(s_1)B_{DP}^U(s_1) + P_{DP}(s_2)(\mathcal{B}_{DP}^U(s_2) - B_{DP}^U(s_1)). \tag{35}$$

where the first term is "sunk" expenditures given bids made at price $P_j(s_1)$. Since in-

formed demand is the same across both protocols, the PD auction leads to a lower debt burden if and only if $\mathcal{B}^U_{PD}(s_2) < \mathcal{B}^U_{DP}(s_2)$. Hence want to show that, for τ sufficiently large, $\mathcal{B}^U_{PD}(s_2) < \mathcal{B}^U_{DP}(s_2)$. A sufficient condition is $X^U_{PD}(s_2, \tau) > X^U_{DP}(s_2)$ for τ sufficiently large and some price $P_j(s_2)$ that is common to both protocols. We will establish the condition for a specific τ first. Since $B^U_{DP}(s_1) < B^U_{PD}(s_1)$ and $P_{PD}(s_1) > P_{PD}(s_1)$, there exists $\tau^* < B^U_{PD}(s_1)$ such that sunk expenditures are equal across protocols. This tier satisfies $\tau^* = \frac{P_{DP}(s_1)}{P_{PD}(s_1)} B^U_{DP}(s_1) < B^U_{DP}(s_1)$. Now assume some common price $P_{PD}(s_2) = P_{DP}(s_2) = P^*$. Then expenditures satisfy

$$\begin{aligned} X^U_{PD}(s_2, \tau^*) &= P_{PD}(s_1)\tau^* + P^*(\mathcal{B}^U_{PD}(s_2) - \tau^*) \\ &= P_{DP}(s_1)B^U_{DP}(s_1) + P^*(\mathcal{B}^U_{PD}(s_2) - \tau^*) \\ &< P_{DP}(s_1)B^U_{DP}(s_1) + P_{DP}(s_2)(\mathcal{B}^U_{DP}(s_2) - B^U_{DP}(s_1)). \\ &= X^U_{DP}(s_2) \end{aligned}$$

For any candidate P^*, given τ^* the PD auction thus generates higher expenditures (and thus lower debt burdens) from the uninformed. Since bids at the high state are independent of τ in the PD auction, the same ranking holds for any interior $\tau > \tau^*$. □

A.8 Proof of Proposition 6

To increase the number of bids and clear the market, the price schedule conditional on the quality shock κ, $P(\kappa, \psi)$, must be declining in the quantity shock ψ. The optimal bids of the informed investor will satisfy $M^U(\kappa, \psi) = 1$ since the short-sale constraint does not bind, and be strictly decreasing in the price (thus increasing in ψ). Because the high-quality schedule lies strictly above the low, the following bidding strategy of the uninformed will replicate the realized bids of the informed:

i $B^U(\kappa_g, \psi) = B^I(\kappa_g, \psi) \forall \psi$,

ii $B^U(\kappa_b, 1) = B^I(\kappa_b, 1) - B^U(\kappa_g, \psi_M)$

iii $B^U(\kappa_b, \psi) = B^I(\kappa_b, \psi) \forall \psi > 1$.

Condition (2) ensures that the bid of the uninformed is positive at $B^U(\kappa_b, 1)$. Since the total bids of the informed and the uninformed are the same by construction, market clearing implies that $X^I(\kappa, \psi) = X^U(\kappa, \psi) = \psi D$, and hence the condition for the bid-overhang constraint is $(1-n)X^U(\kappa, \psi) < \psi D$. This condition must be satisfied for κ_g, and binds first at the top of the bad quality schedule with $\psi = 1$, hence we must ensure that

$$(1-n)X^U(\kappa_g, \psi_M) = (1-n)\psi_M D < D.$$

But this is true by assumption given (3). □

B Demand Shock Version

An alternative version of our model replaces the supply shock ψ with a demand-side shock η, where $1-\eta$ denotes the share of investors who ended up at the auction ready to participate. Assume that $\eta \in [0, \eta_{max}]$ with density $g(\eta)$. Those who do not show up at the auction can be thought of as having suffered a liquidity shock or an alternative investment opportunity shock, either of which made them uninterested in investing in the government's risk bond. The state (suppressing ρ) is now (κ, η), and we can think of bids and prices as a function of the new state just as we did for the supply shock. For simplicity, we will assume that this shock affects both the informed and the uninformed equally. Hence the market clearing condition is now

$$(1-\eta)\sum_{i=I,U} n^i X^i(\kappa, \psi) = D, \qquad \text{for all } \eta. \tag{36}$$

For an individual who finds herself at the auction (A), the likelihood of this happening is decreasing in η,

$$Pr\{A\} = \int_0^{\eta_{max}} (1-\eta) g(\eta) d\eta,$$

and hence arrival at the auction per se constitutes information for the individual about the shock. Using Bayes' rule we get that

$$h(\eta) := Pr\{\eta|A\} = \frac{Pr\{A|\eta\} g(\eta)}{Pr\{A\}}.$$

The density $h(\eta)$ is what an investor will use to consider the probability distribution over the states s conditional on being at the auction.

By construction, supply and demand shocks have an equal impact on the market clearing condition if $\psi = (1-\eta)^{-1}$. Over any interval in ψ values where $[\psi_1, \psi_2]$ where $\psi_i \in [1, \psi_{max}]$ and $\psi_2 \geq \psi_1$, The probability of the corresponding interval in η values is given by

$$\int_{\psi_1}^{\psi_2} h\left(1 - \frac{1}{\psi}\right) \frac{1}{\psi^2} d\psi.$$

Since ψ was assumed to be uniform, it would induce the same probability if it was equal to $\psi_2 - \psi_1$. This leads to

Proposition 7. *If the following conditions hold, we can map our equilibrium with supply shocks into an equilibrium in the demand shock model where the corresponding value is $\tilde{\eta} = 1 - \frac{1}{\psi}$, and prices and bids on the corresponding values are equal; i.e. $P(\theta, \tilde{\eta}) = P(\theta, \psi)$:*

1. *The maximum values coincide: $\psi_{max} = (1-\eta_{max})^{-1}$*
2. *h induces a uniform distribution over the corresponding ψ values: $h\left(1 - \frac{1}{\psi}\right)\frac{1}{\psi^2} = 1$.*

C Auction Equilibrium vs. Competitive Equilibrium

Now that we have defined an auction equilibrium, we can compare its structure with that of a competitive equilibrium, defined as a tuple consisting of an equilibrium price and demand functions such that there is no excess supply of bonds at the equilibrium price.

It is evident that a DP auction equilibrium cannot be a competitive equilibrium: there is a single market-clearing price in the latter, while bonds trade at multiple distinct prices in the former. In UP auctions, however, bonds are always sold at a single price. In many cases, a UP auction equilibrium will, therefore, turn out to be isomorphic to a standard competitive equilibrium with heterogeneous information. We establish this link here.

The decision problem in a competitive equilibrium is identical except for the fact that the short-sale constraint applies to *total* purchases of bonds in each state s, rather than bid-by-bid. That is, in competitive equilibrium, we replace the *auction* short-sale constraint $B^i(\kappa, \psi) \geq 0$ with an *competitive* short-sale constraint of the form $\mathcal{B}^i(\kappa, \psi) \geq 0$ for all states (κ, ψ).

To check whether a UP auction equilibrium is isomorphic to a competitive equilibrium, it is thus sufficient to verify whether or not the short-sale constraints imply each other. To go from the auction to the competitive short-sale constraint we can construct the associated total risky bond purchases just by summing over the in-the-money bids. This implies that if the auction short-sale constraint does not bind in any state, then neither does the competitive one.

To go from the competitive to the auction short-sale constraint, we can construct the associated state-by-state bids given the total bond purchases in a competitive equilibrium, using the difference between the bond purchases at s and those at the next highest price s', i.e.

$$B^i(s) = \mathcal{B}^i(s) - \mathcal{B}^i(s'), \tag{37}$$

where $P(s') = \min \{P(s'') > P(s)\}$ for all $s'' \in \mathcal{S}$.

We can then distinguish three scenarios. First, if the state-by-state bids are nonnegative in the UP auction equilibrium so that the sort-sale constraints do not bind, then the short-sale constraint cannot bind for total purchases and there is an associated competitive equilibrium. Second, if there is a state in which the auction short-sale constraint binds, there is an associated competitive equilibrium only if the short-sale constraint on total bond purchases also binds and total bids are zero. Third, there is no associated competitive equilibrium when there are states for which the nonnegativity constraint does not bind for total purchases in the competitive equilibrium, but do bind for individual bids in particular states in the UP auction.

Even if the differences in short-sale constraints in competitive equilibrium and the auction auction do not matter, so that the auction equilibrium has an associated competitive equilibrium, the bid-overhang constraint still can break the mapping between the two. Hence the set of UP auction equilibria is a subset of the set of competitive equilibria.

D Solution Methods

D.1 Uniform-Price Auctions

When the bid-overhang constraint does not bind, and the no-short-sale constraint does not bind on the uninformed, we can solve for the price level in closed form as in (**??**). The total bond purchases are then determined by (23). When the no-short-sale constraint binds, one simply imposes on the net purchases, which implies that the net is zero at that point. When the bid-overhang constraint binds, we are seeking a solution to the two-equation systems given by (24) and (25) at the particular binding point (ψ_g, ψ_b). We must also ensure that the slope of the price functions at these points in the state space is consistent with (26), and this determines how the state points must change. Overall, the solution in the region where the bid-overhang constraint binds works like a somewhat complicated differential equation system as a result. However, it can be solved using a version of Euler's method.

Taking the points in the state space as given, (κ_g, ψ_g) and (κ_b, ψ_b), there are two basic cases depending upon whether or not the short-sale constraint binds on the informed at ψ_b:

1. If the short-sale constraint binds, then

$$n\frac{1-\kappa_g-P}{1-P} = \frac{D}{W}\left[\psi_g - \psi_b\right]$$

which implies that

$$P = 1 - \frac{\kappa_g}{1-A}$$

where

$$A = \frac{D}{nW}\left[\psi_g - \psi_b\right]$$

Then, given P, we can solve for beliefs as follows

$$(1-n)\frac{1-\tilde{\kappa}-P}{1-P} = (1-n)\left[1 - \frac{\tilde{\kappa}}{1-P}\right] = \frac{D}{W}\psi_b$$

which implies that

$$\tilde{\kappa} = \left[1 - \frac{D}{(1-n)W}\psi_b\right](1-P)$$

2. Assuming that the short-sale constraint does not bind:

$$n\left[\frac{1-\kappa_g-P}{1-P} - \frac{1-\kappa_b-P}{1-P}\right] = n\frac{\kappa_b-\kappa_g}{1-P} = \frac{D}{W}\left[\psi_g - \psi_b\right]$$

This implies that

$$\frac{nW}{D}[\kappa_b - \kappa_g] = [\psi_g - \psi_b](1-P)$$

Initial Binding Point: Assume we start smoothly from the first binding then $\psi_b = 1$ and $\tilde{\kappa} = \kappa_g$. Then, we must have

$$(1-n)\frac{1-\kappa_g - P}{1-P} = \frac{D}{W}$$

which is exactly the binding point for the bid overhang constraint.

Multiplicity: Given a point on the blended schedule with SS binding, we can ask if we could have a no SS outcome as follows:

1. Compute potential price

$$1 - P = \frac{nW}{D}\frac{[\kappa_b - \kappa_g]}{[\psi_g - \psi_b]} \implies P = 1 - \frac{nW}{D}\frac{[\kappa_b - \kappa_g]}{[\psi_g - \psi_b]}$$

2. Compute the $\tilde{\kappa}$ term to hit the high price

$$n\left(\frac{1-\kappa_g - P}{1-P}\right) + (1-n)\left(\frac{1-\tilde{\kappa} - P}{1-P}\right) = \frac{D}{W}\psi_g$$

3. Verify that this works by checking that $1 - \kappa_b - P > 0$ and that $\tilde{\kappa} < \kappa_b$.

Position Updating: The slope of the price functions, and hence the change in the two positions in the state space must generate the appropriate belief in (26). So

$$\frac{\Delta\psi_g\kappa_g + \Delta\psi_b\kappa_b}{\Delta\psi_g + \Delta\psi_b} = \tilde{\kappa}.$$

So if

$$\alpha = \frac{\Delta\psi_b}{\Delta\psi_g + \Delta\psi_b}$$

Then

$$\alpha = \frac{\tilde{\kappa} - \kappa_g}{\kappa_b - \kappa_g}.$$

And

$$\frac{\Delta\psi_b}{\Delta\psi_g} = \frac{\alpha}{1-\alpha}$$

Hence, if we fix the grid on ψ_b, then

$$\Delta\psi_b = \frac{\alpha}{1-\alpha}\Delta\psi_g$$

When we solve for the equilibria without the SS constraint, we will need to step out along the ψ_b grid, because these are the values that are changing quickly. This inverts everything. Given

$$\frac{\Delta\psi_b}{\Delta\psi_g} = \frac{\alpha}{1-\alpha}.$$

Then, if we fix the grid on ψ_b, then

$$\Delta\psi_g = \frac{1-\alpha}{\alpha}\Delta\psi_b$$

D.2 Discriminatory-Price Auctions

The bond bid function for a bidder is given by $b(s)$, where $s \in \{\kappa, \psi\}$. For each quality shock, the function is composed of an initial bid value $b(\kappa, 1)$ and an incremental bid for each supply shock $b(\kappa, \psi)$. The cumulated bids of an investor at quantity shock $\hat{\psi}$ *conditional* on the quality shock κ is given by

$$B(\kappa, \hat{\psi}) = b(\kappa, 1) + \int_1^{\hat{\psi}} b(\kappa, \psi)\partial\psi$$

The corresponding total expenditures of an investor in a DP auction are

$$EX(\kappa, \hat{\psi}) = b(\kappa, 1)P(\kappa, 1) + \int_1^{\hat{\psi}} P(\kappa, \psi)b(\kappa, \psi)\partial\psi$$

while in a UP auction, it is given by

$$EX(\kappa, \hat{\psi}) = B(\kappa, \hat{\psi})P(\kappa, \hat{\psi})$$

Solving the investor's portfolio problem formally requires one to evaluate all of the marginal probability-weighted utilities in the acceptance set in order to determine the appropriate marginal condition (**??**).

Fortunately, with the UP protocol, this requirement unravels from the bottom of the price schedule, and we only need to determine the local condition if the no-short-sale constraint does not bind. While, if it does bind then the bids over the binding range are zero. Unfortunately, with the DP protocol this is explicitly not the case. Hence solving the investor's portfolio problem inherently involves a large simultaneous equation system since all of the bids are jointly determining the bond expenditures and returns and hence consumptions conditional on the state and default outcome.

For this reason, to make progress in solving DP, we will approximate the bid and price functions as follows: Assume that we have a set of **K equally spaced grid points** in $[1, \bar{\psi}]$, which includes both endpoints. Denote each grid point by ψ_k. The distance between grid points by $\Delta = \frac{\bar{\psi}-1}{K-1}$. We will assume that **the incremental bid function** $b(\kappa, \psi)$ **is a step function** with steps at the grid points. In particular,

$$b(\psi) = b(\psi_k) \text{ for all } \psi_{k-1} < \psi \leq \psi_k$$

We also assume that the **price function is linear between grid points**. In particular, defining

$$\Lambda(\psi) = \frac{\psi - \psi_{k-1}}{\Delta} \quad \text{for all } \psi_{k-1} < \psi \leq \psi_k$$

the price between any two grid points is given by

$$P(\psi) = (1 - \Lambda(\psi))P_{k-1} + \Lambda(\psi)P_k$$

Total expenditures of an investor given a quality shock κ at $\psi \in (\psi_{k-1}, \psi_k]$, for all $k \in \{1, .., K\}$ are given by

$$EX(\kappa, \psi) = b(\kappa, 1)P(\kappa, 1) + \sum_{k'<k} b_{k'} \left[\frac{P_{k'-1} + P_{k'}}{2}\right] \Delta$$
$$+b_k \left[\frac{P_{k'-1} + P(\psi)}{2}\right] \Lambda(\psi)\Delta$$

For the informed, since they know κ, these characterize total expenditures at each ψ, which we denote as $\overline{EX}^I(\kappa, \psi) = EX^I(\kappa, \psi)$. If all investors are informed, for instance, market clearing implies

$$\overline{EX}^I(\kappa, \psi) = \psi D$$

Note that since expenditures are linear here in ψ, it follows that if market clearing is occurring at the grid points, then it will occur everywhere.

For the uninformed, total expenditures are defined by combining expenditures in both schedules. This is because in state (κ, ψ) the bids in-the-money are all those made at prices above $P(\kappa, \psi)$ and also above $P(\kappa^c, \psi') = P(\kappa, \psi)$. If there is no overlap (this is $P(g, \bar{\psi}) > P(b, 1)$, then total expenditures are

$$\overline{EX}^U(g, \psi) = EX^U(g, \psi) \quad and \quad \overline{EX}^U(b, \psi) = EX^U(b, \psi) + EX^U(g, \bar{\psi})$$

If there is overlap (this is $P(g, \bar{\psi}) \leq P(b, 1)$, then total expenditures are

$$\overline{EX}^U(\kappa, \psi) = EX^U(\kappa, \psi) + EX^U(\kappa^c, \psi'),$$
$$\text{where } P(\kappa, \psi) = P(\kappa^c, \psi')$$

where κ^c is the complementary quality shock.

Now we need to construct marginal utilities in case of default and repayment to compute the FOCs. To accumulate marginal utilities, which are not linear, we can improve the approximation by constructing a finer grid between grid points. Let's define a finer grid given by $\{\nu\}$ and assume that it contains all the old grid points plus at least the midpoints of the old grid. Given the linearity in bids and prices, we can populate the expenditures in each finer grid point simply by **interpolating** between grid points.

Then, **(assuming log utilities)** we can define our probability-weighted marginal utilities in case of default for investor $i \in \{I, U\}$ in each as follows:

$$MUD^i(\kappa, \nu) = \left(W - \overline{EX}^i(\kappa, \nu)\right)^{-1} f(\kappa)\kappa(\kappa)$$

and, defining RX symmetrically to EX above, but with $1 - P$ instead of P, we can define our probability-weighted marginal utilities in case of repayment for investor $i \in \{I, U\}$ in

each ν as

$$MUR^i(\kappa,\nu) = \left(W + \overline{RX}^i(\kappa,\nu)\right)^{-1} f(\kappa)(1-\kappa(\kappa))$$

Then the sum of the marginals over the in-the-money sets is given by

$$\overline{MUD}^i(\bar{\kappa},\bar{\nu}) = \sum_{(\kappa,\nu):P(\kappa,\nu)\leq P(\bar{\kappa},\bar{\nu})} MUD^i(\kappa,\nu)$$

$$\overline{MUR}^i(\bar{\kappa},\bar{\nu}) = \sum_{(\kappa,\nu):P(\kappa,\nu)\leq P(\bar{\kappa},\bar{\nu})} MUR^i(\kappa,\nu)$$

Forming FOCs and MCs: We will only check that the focs and market clearing conditions hold at the grid points. Given our construction of cumulated marginal, focs are given by:

$$-\overline{MUD}^i(\kappa,\psi_k)P(\kappa,\psi_k) + \overline{MUR}^i(\kappa,\psi_k)(1-P(\kappa,\psi_k)) = 0 \qquad \forall i,\kappa,\psi_k$$

Notice that, based on the text, this is the same as in equation **??**

$$M^i(\kappa,\psi_k) \equiv \frac{\overline{MUD}^i(\kappa,\psi_k)}{\overline{MUR}^i(\kappa,\psi_k)} = \frac{1-P(\kappa,\psi_k)}{P(\kappa,\psi_k)} \qquad \forall i,\kappa,\psi_k$$

and market clearing is given by

$$n\overline{EX}^I(\kappa,\psi_k) + (1-n)\overline{EX}^U(\kappa,\psi_k) = \psi_k D \qquad \forall \kappa,\psi_k$$