JEL Codes: D82, G23, G32

Review of Economic Studies (2025) **x**, x–xx doi:10.1093/restud/rdaaxxx

Repurchase Options in the Market for Lemons

SAKI BIGIO
Department of Economics, University of California, Los Angeles and NBER, USA
LIYAN SHI
Tepper School of Business, Carnegie Mellon University and CEPR, USA

First version received August 2020*; Editorial decision December* 2020*; Accepted June* 2025 (*Eds.*)

We study repurchase options (repo contracts) in a competitive asset market with adverse selection. Gains from trade emerge from a liquidity need, but private information about asset quality prevents the full realization of trades. In equilibrium, a single repo contract pools all assets. The embedded repurchase option mitigates adverse selection by improving the volume of trades relative to outright sales. However, liquidity provision can be inefficiently low as lenders compete to attract high-quality assets via high haircuts and low rates. The equilibrium has a closed form and aligns well with empirical patterns across Mortgage-Backed Securities repos.

Key words: Repurchase Agreement, Collateralized Debt, Private Information, Optimal Contracts

1. INTRODUCTION

Many financial contracts, such as collateralized debt, factoring, payroll lending, and pawning, are *de facto* asset-sale contracts that embed a repurchase option.[1] Strip away legal differences, and all of these contracts, which for the rest of the paper we refer to as repos, have a common structure: they all specify a collateral asset, a loan amount, and a loan principal. Because these contracts are non-recourse or rarely settled in court, borrowers have the option to default at little or no cost. A default option is, *de facto*, an option to repurchase the asset. The extensive use of repo-like contracts raises several important questions: why are repos so prevalent? Why don't borrowers simply sell their assets and repurchase similar ones later? What factors determine the key terms of repos: the loan-to-collateral value ratio (i.e., the haircut) and the implied interest rate?

This paper argues that repos emerge as a natural response to adverse selection in financial markets. To formalize this insight, we consider a market where trading

1. These contracts are not a modern financial innovation. An early form of collateralized debt contract, known as the *Pignus*, was already stipulated in Roman law. Holmstrom (2014) considers pawning dating back to Tang-dynasty China, *circa* 650 AD, as the precursor of these contracts.

The editor in charge of this paper was Veronica Guerrieri.

is motivated by liquidity needs and agents have a common valuation of the assets. Specifically, borrowers seek to raise funds for investment projects by trading pre-existing assets. However, because borrowers possess private information about the quality of their assets, a classic lemons problem arises when only asset sales are possible (Akerlof, 1970): those with high-quality assets refrain from selling due to the low price, ensuing market unraveling.

Lenders can mitigate adverse selection by offering a repurchase option. Indeed, a borrower unwilling to sell a high-quality asset at a low price may be willing to do so if offered the option to repurchase the asset, even if the repurchase price entails a slight premium over the selling price. Of course, borrowers with lemons will take advantage by selling their assets and failing to repurchase them, implicitly defaulting. However, rightly set, the repurchase premium can compensate lenders for default losses and induce high-quality assets to participate in the market.

The discussion thus far illustrates why lenders may offer repos, but what should we expect of repos offered in a competitive market? Competition among lenders can significantly influence repo terms and market outcomes. Lenders can compete by targeting high-quality collateral assets through various strategies: adjusting repo terms (such as haircuts and interest rates), employing quantity screening, or even offering menus of repo contracts. What repo contracts will emerge in equilibrium? Are outcomes efficient? The rest of the paper is devoted to answering these questions.

We consider a general setting where asset qualities are drawn from an arbitrary distribution, encompassing continuous- and discrete-type distributions. Lenders compete by offering repo menus. We adopt the Miyazaki-Wilson-Spence—henceforth MWS—equilibrium concept (Miyazaki, 1977; Wilson, 1977; Spence, 1978), which allows for contract withdrawals. Under this equilibrium concept, any additional contract must be unprofitable after potential withdrawals of existing contracts made unprofitable by its introduction.

The MWS equilibrium concept is appealing for several reasons: First, the equilibrium is conveniently characterized by a sequence of Pareto problems, thus leading to Pareto-efficient outcomes. Second, Netzer and Scheuer (2014) provides a microfoundation for this concept through a trading game. This microfoundation aligns with real-world practices where lenders post contracts that may be withdrawn after observing market conditions. Finally, this concept is widely used, for example, in labor-market settings in Stantcheva (2014) and competitive security design settings in Asriyan and Vanasco (2023).

Despite the richness of contractual possibilities, the repo market equilibrium is strikingly simple: the equilibrium features a *single* contract that *pools* all assets. This contract maximizes the borrower's payoff from exercising the repurchase option, subject to the lender's break-even condition. This maximization problem and, thus, the equilibrium admits an explicit solution. In particular, given a continuous distribution of collateral quality $F(\cdot)$ and borrower investment return r, the contract's repurchase price is the $r/(1+r)$ distribution quantile. In turn, the sales price is the expected value of the minimum between the collateral value and the repurchase price—i.e., the lender's break-even price that incorporates the borrowers' default decisions. This simplicity is conducive to a transparent analysis and makes the theory portable to many applications.

The repo market outcome has several noteworthy features. First, unlike environments with only outright sales, quantity screening does not emerge in equilibrium. This is because repo terms serve as a more effective screening device. With outright sales, high-quality types may wish to distort their traded quantities to separate from lower-quality types instead of pooling with them. In contrast, with repos, lenders can attract higher-quality types by instead adjusting the repurchase price, which can compensate for a reduction in the loan amount with a lower interest charge.

A second noteworthy feature is that the repo-market equilibrium is robust. Namely, the same pooling repo contract emerges regardless of whether the assets are divisible, quantity screening is possible, or contracts are exclusive. This robustness follows naturally because repurchase options are a more desirable tool to attract higher-quality types than quantity screening, making these details of market structure irrelevant. This robustness contrasts outcomes in sales-only environments, which are sensitive to the trading protocols and market structure (see, e.g., Rothschild and Stiglitz, 1976; Attar et al., 2011).

A single pooling repo contract serving the whole market is also robust under alternative equilibrium concepts, such as competitive-search environments (Guerrieri et al., 2010) and Riley (1979)'s equilibrium. This robustness suggests that pooling is induced by the repurchase option, not the equilibrium refinements or trading protocols. While pooling is robust, the actual repo terms may differ across environments. The competitive search and Riley (1979)'s equilibria coincide in that their equilibrium sales and repurchase prices both equal the lowest asset quality. That is, lenders fully protect themselves against collateral quality risk by extending the minimum loan amount, thereby ensuring zero default. This repo outcome displays the highest possible haircut and zero interest premium. We adhere to the MWS concept because of its more realistic repo-rate predictions. However, the repo terms across all equilibrium concepts converge under several conditions, suggesting a fundamental consistency in market behavior in generating high haircuts.

We investigate the efficiency of repo markets, both ex-ante and in the interim. As is well known, the MWS concept leads to Pareto-efficient outcomes in the interim stage for a given set of contractual possibilities, guaranteeing that repo and outright sales markets are both interim efficient. Still, introducing repos generally favors high-quality types and may even produce Pareto gains relative to outright sales alone, particularly when adverse selection is severe. However, repos are inefficient from an ex-ante perspective, i.e., before the asset qualities are known to their owners. A planner, facing the same interim incentive and participation constraints as lenders, also offers a single pooling repo contract but one that maximizes cross-subsidization. In contrast, lenders in the repo market compete to cream-skim high-quality assets by offering higher haircuts and lower rates, which leads to excessively high haircuts and ultimately undermines liquidity.

To demonstrate the applied nature of our theory, we use it to rationalize observed patterns of haircuts and repo rates. We focus on a particular repo market with ample empirical evidence suggesting that private information is particularly relevant: the bilateral repo market for different Mortgage-Backed Securities (MBS) tranches. We extend the baseline model to capture this market's institutional details. In this extension, borrowers obtain private signals about the default probability of a common underlying mortgage pool. This informational advantage translates into varying collateral quality risk across different MBS tranches. As more junior tranches absorb default losses before senior tranches, their asset quality is effectively more dispersed from lenders' perspective.

The calibrated model successfully replicates the increasing pattern of haircuts and repo rates for more junior tranches observed in the data. This success does not compromise parsimony or realism: the calibration only concerns two parameters that govern the information structure about the underlying mortgage pool, and the best-fitting values imply a mild borrower informational advantage. Remarkably, the model also captures that the bulk of adjustment against collateral quality risk is through haircuts. Our repo formula explains why we should expect the range of haircuts to be orders of magnitude larger than the range of repo rates. Through the same logic, the model also makes sense of large haircut spikes and liquidity erosion in episodes when information frictions become acute.

In summary, this paper makes several contributions: First, we provide a novel theoretical framework to explain the prevalence of repos in these markets. Our analysis yields a strikingly simple closed-form solution. We show that while repos enhance liquidity by bringing high-quality assets to the asset pool, competition induces excessively high haircuts that undermine overall liquidity provision. The robustness of its features to various trading environments and equilibrium concepts provides a unified explanation for repo market outcomes. Finally, our application to the MBS market offers new insights into the determinants of repo terms.

Related literature. Originally motivated by trades in used-goods markets, Akerlof's model is also foundational in finance, suggesting that asymmetric information is also a fundamental friction hindering liquidity (see Myers and Majluf, 1984; Gorton and Pennacchi, 1990; Eisfeldt, 2004). The study of asymmetric information in macroeconomics and finance branches into two areas: enriching the contract space and enriching the setting. This paper contributes to the former area while having implications for the latter class of models.

The literature that enriches the contract space has primarily focused on security design. In Demarzo and Duffie (1999), borrowers sell securities backed by current and future cash flows before learning future cash flows. Biais and Mariotti (2005) considers a similar setting, but securities are only backed by future cash flows, although the contract includes a traded quantity. A common finding is that collateralized debt, akin to repos, is ideal in both settings, although for different reasons: as a signaling device in Demarzo and Duffie (1999) and as a commitment device in Biais and Mariotti (2005). In Asriyan and Vanasco (2023), the market is competitive and non-exclusive as in Biais and Mariotti (2005), but cash flows are non-deterministic. Our paper departs from these papers because the market trades after observing private information.[2] Thus, while the information structure is similar to Biais and Mariotti (2005), we study a market equilibrium without commitment to a trading contract. The testable implications differ: the extent of asymmetric information manifests in a wedge between the implied interest rate and the return on investment, a source of ex-ante inefficiency. In fact, the ex-ante efficient contract in our setting, which eliminates that wedge, coincides with the security-design solution in Biais and Mariotti (2005).

The market for repo contracts studied also relates to leasing contracts, a form of repurchase options. Hendel and Lizzeri (2002) introduced leases into the sales-only model of Hendel and Lizzeri (1999) where customers self-select into a homogeneous-good primary market and a secondary market plagued by adverse selection. Hendel and Lizzeri (2002) demonstrate that a monopolist can improve customer selection into markets by offering leases. Whereas we focus on market competition, Hendel and Lizzeri (2002) study the problem of a monopolist who wants to affect market selection.

2. For instance, in Demarzo and Duffie (1999), the assumption is: “After designing the security, but before the sale of the security to outside investors, the issuer or underwriter handling the sale receives information relevant to the payoff of the security.”

The second branch embeds lemons markets into richer settings that feature: dynamic trade and learning (Daley and Green, 2012; Guerrieri and Shimer, 2014; Fuchs and Skrzypacz, 2019), nonexclusive contracts (Attar et al., 2011; Kurlat, 2016; Asriyan and Vanasco, 2023), competitive and random search (Guerrieri et al., 2010; Lester et al., 2019), information acquisition (Gorton and Ordoñez, 2014), multidimensional screening (Chang, 2017; Guerrieri and Shimer, 2018), chains of trades (Dang et al., 2013; Gottardi et al., 2019), equilibrium with signals (Kurlat and Scheuer, 2020), and models where trade under private information functions as a form of down payment (Bigio, 2015).[3] With few exceptions, these studies primarily focus on asset sales. One exception is Madison (2024), which studies security design under random search, or Ozdenoren et al. (2023) in dynamic settings. The repo market here provides a simple and portable formula that is easy to introduce in richer environments. Beyond tractability, the properties of the repo market studied may lead to different implications. For example, Ozdenoren et al. (2023) show that when repos emerge from security design, multiplicity in equilibrium dynamics disappears. It would be interesting to explore whether similar stability is also achieved in the competitive setting studied here.

The paper also complements previous work that rationalizes the use of repos. Our model captures the use of repos driven by borrower liquidity needs. Some special repo transactions are driven by the demand for the underlying security, such as for short-selling motives. For example, Duffie (1996) analyzes how these "special" repos, often U.S. Treasury securities, can experience negative spreads. Another argument for using repos is that it is more convenient to pledge collateral than to sell the asset if there are gains from keeping the asset. In Monnet and Narajabad (2017) and Parlatore (2019), repos minimize endogenous transaction costs that result from search frictions. Alternatively, repos can provide insurance in incomplete-markets or limited-commitment settings. Fostel and Geanakoplos (2015) study a two-period two-state Arrow-Debreu equilibrium where trading contracts on state-contingent securities require posting collateral (collateralized debt). That paper shows that the set of equilibria is indeterminate but always includes a zero-default equilibrium. In turn, Gottardi et al. (2019) study a similar setting with limited commitment but introduce a third period to allow retrading and state-contingent repurchase options. Their model shows how lenders reuse collateral to transfer wealth to states where the collateral becomes more valuable.

Outline. The organization of the paper is as follows. We lay out the environment in Section 2. Section 3 characterizes the repo market equilibrium. Section 4 studies variations to the environment and discusses the robustness of the repo equilibrium. Section 5 presents the efficiency results—we study the planner or, equivalently, security design solution and compare the repo market outcomes to markets where only outright sales are allowed. Section 6 presents an application of the theory in the MBS repo market.

3. Adverse selection appears prominently in models where collateral relaxes asymmetric information problems regarding, for example, project risks (Stiglitz and Weiss, 1981; Besanko and Thakor, 1987), costly state verification problems (Townsend, 1979), or moral hazard problems (Holmstrom and Tirole, 1997).

2. MARKETS FOR REPO CONTRACTS

2.1. *Environment*

Consider a two-period economy $t=1,2$. There is a unit mass of borrowers and a large number of lenders. All agents are risk neutral and maximize their $t=2$ payoff. At $t=1$, each borrower is endowed with a *divisible* asset and has access to an investment project. At $t=2$, assets pay dividends, and projects yield a return. Assets are heterogeneous: they differ in the dividend payout, denoted by $\theta \in \Theta \subset \mathbb{R}_+$. Projects are identical: for each dollar of cash invested in the project at $t=1$, the project generates a return $r>0$ with certainty at $t=2$.[4] Borrowers need to raise external financing for investment because project returns are non-pledgeable, as in Hart and Moore (1994). The borrowers' only option is to use their assets to obtain funding from lenders. Lenders provide funds competitively, and their cost of funding is normalized to zero.

Information. Information regarding the asset quality θ is asymmetric: it is only known to its owner, the borrower. Let the support of asset qualities be $\Theta = [\underline{\theta}, \bar{\theta}]$, where $\underline{\theta}$ and $\bar{\theta}$ are the lowest and highest quality types, respectively. Denote its cumulative distribution function by $F(\theta): \Theta \to [0,1]$. With discrete qualities, $F(\theta)$ is a step function, whereas $F(\theta)$ is strictly increasing for continuous qualities. The distribution is common knowledge.

Contracts. Lenders each offer *a set of repo contracts.* A *repo contract* specifies a pair of prices $\boldsymbol{p} = \{p_s, p_r\} \in [\underline{\theta}, \bar{\theta}] \times [\underline{\theta}, \bar{\theta}]$.[5] The first entry, p_s, is the sale price per unit of asset. It specifies the price for the first leg of the contract: at $t=1$, the lender pays the price p_s per unit to the borrower, and the transacted asset is transferred to the lender. The second entry, p_r, is the repurchase price per unit to be repaid by the borrower at $t=2$. We assume limited commitment on the borrower's side. Failure to repurchase is a default, in which case the lender seizes the asset without further recourse. If the borrower honors the contract and repays, the asset is returned. Although borrowers cannot commit, lenders commit to return the asset upon repayment.[6] Note that *outright-sales contracts* are included: for any contract with $p_r = \bar{\theta}$, both parties understand that no borrower would repurchase at that price.

Contracts also specify the quantity of the asset traded, denoted by $q \in [0,1]$. Thus, a contract is summarized by $\{\boldsymbol{p}, q\}$. An equivalent interpretation of q is the trading probability of indivisible assets.

Repo and collateralized debt terminology. Repo contracts can also be interpreted as non-recourse collateralized debt. Under this interpretation, the sales price p_s represents the loan size, and the repurchase price p_r represents the loan principal. Under either interpretation, repos, or collateralized debt, the implied interest rate is $\tilde{r} \equiv p_r/p_s - 1$.

4. The environment is equivalent to one where trade occurs due to the relative impatience of borrowers, as in Biais and Mariotti (2005) and Guerrieri and Shimer (2014). Consider the same two-period economy, where borrowers discount the period-2 payoff by a factor β, but lenders do not. To equate the two settings, one can set the borrower's discount factor to the inverse of the gross investment return, i.e., $\beta = 1/(1+r)$.

5. The domain $[\underline{\theta}, \bar{\theta}] \times [\underline{\theta}, \bar{\theta}]$ contains all relevant contracts. Hence, the restriction of prices to this domain is without loss of generality.

6. A legal environment can induce the lender's commitment. For example, asset custody by a third party can enforce the returning the asset to the borrower upon repayment.

Finally, in the jargon of repo markets, we define $h \equiv 1 - p_s/\mathbb{E}[\theta]$ as the repo's haircut, relative to the notional value, its unconditional average $\mathbb{E}[\theta]$. Likewise, $1-h$ represents the loan-to-collateral value in the context of collateralized debt.[7]

2.2. *Agents' payoffs*

Borrower payoff. For a borrower with an asset of quality θ, the value of her end-of-period wealth from entering a contract with per-unit prices $\boldsymbol{p}$ and quantity q is:

$$V\big(\boldsymbol{p},q;\theta\big) = \max\Big\{\underbrace{\theta + q\big((1+r)p_s - p_r\big)}_{\text{repay}}, \quad \underbrace{\theta + q\big((1+r)p_s - \theta\big)}_{\text{default}}\Big\}.$$

In each scenario, repayment or default, the end-of-period wealth is the sum of the asset's value plus the payoff contingent on that decision. Naturally, the traded quantity q scales the payoff. The payoff is given by the investment return, $(1+r)p_s$, minus the repayment cost. If the borrower defaults, the cost is the loss of the traded portion of the asset, $q\theta$. If the asset share is repurchased, the cost is the repurchase price p_r scaled by q.

The borrower chooses the contract, among the set offered, that maximizes her end-of-period wealth or, equivalently, her payoff from trade.[8] Formally:

$$\max_{\{\boldsymbol{p},q\}\in\mathcal{C}} V\big(\boldsymbol{p},q;\theta\big) = \theta + \max_{\{\boldsymbol{p},q\}\in\mathcal{C}} v\big(\boldsymbol{p},q;\theta\big),$$

where the payoff from entering contract $\{\boldsymbol{p},q\}$ is given by:

$$v(\boldsymbol{p},q;\theta) \equiv q\left[(1+r)p_s - \min\{\theta,p_r\}\right]. \tag{2.1}$$

Because θ is independent of the choice of contract $\{\boldsymbol{p},q\}$, maximizing wealth V is equivalent to maximizing the payoff (2.1). We characterize the equilibrium focusing on the trading value $v(\boldsymbol{p},q;\theta)$ with the understanding that this choice also maximizes the borrower's wealth. Finally, we let $v(\theta)$ denote the equilibrium payoff for a borrower of type θ, i.e., a borrower whose asset quality is θ.

Repos break single-crossing property. Repo contracts are distinguished from outright sales and change the borrower and seller payoffs. To illustrate this point, consider a market with only outright sales. A seller of quality θ obtains a payoff from selling a q fraction of her asset at a unit price p_s:

$$v_a(p_s,q;\theta) = q\left[(1+r)p_s - \theta\right]. \tag{2.2}$$

Recall that asset-sale contracts are a particular case of repo contracts where the repurchase price is set so that the borrower defaults. Thus, the asset-sale payoff function (2.2) drops the "min" operator from the repo payoff function (2.1). This payoff is strictly decreasing in type θ. Furthermore, the asset-sale payoff function (2.2) satisfies the

7. In practice, there are legal distinctions between defaults in repo and collateralized debt contracts, which we do not consider here.

8. The equivalence is shown by taking $\theta + (1+r)p_s$ outside the max operator and using the identity $\max\{-a,-b\} = -\min\{a,b\}$.

single-crossing property, $\partial v_a(p_s,q;\theta)/\partial\theta\partial q = -1$, which is critical for potential quantity screening. In contrast, the repo payoff is *weakly* decreasing in θ. It is only "weakly" decreasing because the payoff becomes flat for borrowers who exercise the repurchase option. As a result, the single-crossing property does not hold with repos:

$$\frac{\partial v(\boldsymbol{p},q;\theta)}{\partial\theta\partial q} = \begin{cases} -1, & \text{if } \theta < p_r \\ 0, & \text{if } \theta \geq p_r. \end{cases}$$

Lender profit. A lender derives profit from a contract $\{\boldsymbol{p},q\}$ against a quality-θ asset:

$$\pi(\boldsymbol{p},q;\theta) = q[\min\{\theta,p_r\} - p_s]. \tag{2.3}$$

The profit is the difference between the repayment decision of type θ minus the sales price, scaled by the transacted quantity, q. A lender's overall profit sums across the set of offered contracts and the borrowers attracted to each contract.

2.3. *Equilibrium*

As a baseline, for most of the paper, we adopt the Miyazaki-Wilson-Spence—henceforth MWS—equilibrium concept. An MWS equilibrium is defined as follows:

Definition ***1 (Equilibrium)*** *An equilibrium is a set of offered contracts such that:*

(i) borrowers choose the optimal contract among the set;
(ii) each lender earns non-negative profit from the set of contracts they offer;
(iii) there is no other contract set that generates positive profits upon being added to the original set and continues to do so even after any original contracts rendered unprofitable by its inclusion is withdrawn.

Items (i-ii) correspond to standard notions of optimality and free entry; item (ii) allows for potential cross-subsidization across contracts offered by the same lender. Item (iii) is the defining equilibrium criterion of the MWS equilibrium.[9] An equilibrium criterion is warranted in this setting because, as is known since Rothschild and Stiglitz (1976), standard notions of competition, which require only conditions (i) and (ii), are problematic in markets with private information. Indeed, any (zero-profit) incentive-compatible set of contracts, e.g., as obtained from a mechanism that maximizes funding or the trading value of any type, could satisfy conditions (i) and (ii) but fail to correspond to equilibrium with an appropriately defined entry decision.[10]

We consider the MWS equilibrium criterion a natural choice for several reasons. First, it has a realistic foundation as a sequential game (Netzer and Scheuer, 2014, see the discussion in Section 4.1). This foundation captures the intuitive idea that if a competitor tries to poach good trading counterparts, a counteracting strategy is to dump bad trading partners by withdrawing contracts. Second, the MWS criterion is a

9. Item (iii) is replaced to obtain Riley's reactionary equilibrium for one in which there is no other set of contracts which would make positive profits, even after additional contracts are introduced—see the formalization by Engers and Fernandez (1987).

10. Once we find a contract set that satisfies (i) and (ii) nothing precludes the entry of additional contracts that attract profitable types rendering the original set unprofitable.

commonly used equilibrium concept: see, for example, Stantcheva (2014) or Asriyan and Vanasco (2023). Third, as we show next, this equilibrium is cast into a sequence of Pareto problems, which renders tractability and Pareto efficiency. As shown by Miyazaki (1977) and Spence (1978) for finite types and Gemmo et al. (2020) for a continuum of types, the equilibrium is the solution to the following recursion of problems:

Problem *1 (MWS Problem)* *For any asset quality type $\theta \in \Theta$, the MWS problem is:*

$$v^{\star}(\theta)=\max_{\{\boldsymbol{p}(\theta'),q(\theta'):\forall\theta'\leq\theta\}} v(\boldsymbol{p}(\theta),q(\theta);\theta) \tag{2.4}$$

subject to

$$v\left(\boldsymbol{p}(\theta'),q(\theta');\theta'\right)\geq v\left(\boldsymbol{p}(\theta''),q(\theta'');\theta'\right),\forall\theta'',\theta'\leq\theta \tag{IC}$$

$$\mathbb{E}\left[\pi\left(\boldsymbol{p}(\theta'),q(\theta');\theta'\right)\middle|\theta'\leq\theta\right]\geq 0 \tag{ZPC}$$

$$v\left(\boldsymbol{p}(\theta'),q(\theta');\theta'\right)\geq v^{\star}\left(\theta'\right),\forall\theta'<\theta. \tag{RC}$$

Problem 1 is a *recursive* sequence of Pareto sub-problems indexed by θ. Because condition (RC) applies to all types below θ, the recursion is ordered from the lowest to the highest types. In sub-problem θ, we denote the solution by $\left\{\boldsymbol{p}(\theta';\theta),q(\theta';\theta);\forall\theta'\in[\underline{\theta},\theta]\right\}$. This solution generates a payoff for the terminal type θ in that sub-problem, $v^{\star}(\theta)$, which we call the *reservation value* and is an input for subsequent sub-problems. The solution to the last sub-problem, corresponding to the highest type $\bar{\theta}$, which we denote by $\{\boldsymbol{p}(\theta),q(\theta):\forall\theta\in\Theta\}$, is the basis for the MWS equilibrium. The resulting equilibrium payoff for borrowers, $v(\theta)=v\left(\boldsymbol{p}(\theta),q(\theta);\bar{\theta}\right)$, obtained in the final step, is, in general, different from their reservation value, $v^{\star}(\theta)$.

In sub-problem θ, we maximize the highest type's payoff subject to several constraints: Condition (IC) the incentive-compatibility constraint, which guarantees that each type prefers its assigned contract. Condition (ZPC) corresponds to lender profitability—item (ii). Finally, the last condition (RC) states that lower types must at least obtain their reservation values.

It is worth reflecting on why the MWS equilibrium, i.e., a contract set satisfying Definition 1, is the solution to Problem 1, and vice versa. We borrow the arguments provided by Spence (1978) for the discrete-type case without reproducing a complete proof here. A key argument to establish the representation of an MWS equilibrium as a recursion of problems is to show that an MWS equilibrium contract set must offer all types a value above their reservation values obtained from the recursion. If any type $\theta<\bar{\theta}$ would receive less than their reservation value, $v^{\star}(\theta)$, a lender could enter and offer the contract set that solves θ's sub-problem. This entry contract would at least attract θ and maintain profitability regardless of any contract withdrawals from the candidate equilibrium set.[11] The existence of such profitable entry would violate condition (iii). Thus, any contract set satisfying Definition 1 must satisfy the constraints that appear in the final problem, $\bar{\theta}$'s sub-problem, i.e., (IC), (ZPC), *and* (RC). A similar argument explains why maximizing the highest types is also necessary.[12] A final argument to

11. This reasoning relies on profits being increasing in type in any sub-problem, as shown in Appendix A.

12. If a given candidate MWS equilibrium does not do so, the solution to problem $\bar{\theta}$ could be offered, leading to a withdrawal of the highest types and, thereby, attracting all types while making profits.

demonstrate the equivalence requires showing that the solution to sub-problem $\bar{\theta}$ satisfies condition (iii) and is, therefore, indeed an MWS equilibrium.[13]

3. EQUILIBRIUM CHARACTERIZATION

This section characterizes the repo market equilibrium for arbitrary asset quality distributions. We derive a simple analytic solution for continuous distributions.

3.1. *A relaxed MWS problem*

We begin with some observations about the MWS Problem (1). The first observation is that, due to the repurchase option, the set of incentive-compatible contracts exhibits patterns distinct from those in outright-sales markets.

Lemma ***1 (Incentive-Compatible Contracts)*** *With any incentive-compatible contracts, the borrower payoff $v(\theta)$ is weakly decreasing in θ. Moreover, there exists a unique threshold $\theta^d \in \Theta$ such that:*

(i) Types $\theta < \theta^d$ strictly prefer to default, i.e., $p_r(\theta) > \theta$. Their payoff $v(\theta)$ is strictly decreasing in type.
(ii) Types $\theta \geq \theta^d$ weakly prefer to repurchase, i.e., $p_r(\theta) \leq \theta$. They obtain the same payoff.

Proof. See Appendix A.2. ||

Lemma 1 can be attributed to how the repurchase option shapes the (IC) constraints. we first note the existence of a default threshold, above which the borrowers opt for a contract and exercise the repurchase option. To illustrate this, consider the contrary: there exists a lower-quality borrower who signs a contract and repurchases and another with a higher-quality asset who defaults. If the lower-quality borrower mimics the higher-quality one by defaulting, she will obtain a strictly higher payoff than her counterpart, given her lower default cost. Hence, the fact that the lower-quality borrower does not default reveals an inconsistency: the higher-quality could adopt the lower-quality borrower's strategy, repurchasing the asset, and achieve the same payoff. In other words, the higher-quality borrower could strictly improve her payoff by not defaulting.

The repurchase option alters the property of the borrower payoff $v(\theta)$. Among the borrowers that default, a low type has incentives to mimic a high type, making the (IC) constraint potentially binding upward. This force implies that their payoff is strictly decreasing among these types, just as with outright-sales contracts. However, for borrowers who do not default, a high type now also has incentives to mimic a low type, making the (IC) constraint potentially binding downward. Thus, the borrower payoff becomes flat in type. Indeed, conditional on not defaulting, the borrowers would

13. Consider the solution to the MWS problem as a candidate equilibrium. Suppose another contract set attempts entry and attracts a loss-making subset of types—those cross-subsidized in the candidate equilibrium. If the new contract set makes profits, neither the original nor the new contract set would be withdrawn. However, that would mean that the orignal set was not maximizing Problem 1—it should have incorporated both sets in the first place. Conversely, if the sub-group attracted is a profit-making subset, doing so would lead to the withdrawal of the original contract set, inducing losses for the new contract set.

choose a contract that maximizes their non-default payoff, $q[(1+r)p_s - p_r]$. This payoff is independent of the asset quality θ, so all non-default types must obtain the same payoff.

The repurchase option also affects the property of the borrower's reservation value $v^\star(\theta)$, but in a different manner, as we describe below.

Lemma ***2 (Reservation Value)*** *The borrower's reservation value* $v^\star(\theta)$ *is weakly increasing in type. For the highest type, its payoff coincides with the reservation value,* $v(\bar{\theta}) = v^\star(\bar{\theta})$.

Proof. See Appendix A.3. ‖

Contrasting Lemma 2 with the outcomes under asset-sales markets is instructive. Absent repos, higher types may have lower reservation values; with repos, the reservation value increases with type. To understand why, recall the sequential nature of the MWS problem (2.4). For an increasing sequence of quality types, we first solve the sub-problem of type θ_1 and then solve the sub-problem of type $\theta_2 > \theta_1$. One possible contract we can offer type θ_2 is the contract we offer to θ_1 in her sub-problem, but adjusting the repurchase price to $p_r(\theta_2;\theta_2) = \min\{p_r(\theta_1;\theta_1), \theta_1\}$. Offering the same contracts to all lower types in θ_1's sub-problem enables θ_2 to at least obtain θ_1's payoff while ensuring all constraints are satisfied. Thus, θ_2's reservation value must be at least as high as θ_1's. Naturally, when we solve the highest type's problem, its value $v^\star(\bar{\theta})$ is the eventual this type's equilibrium payoff, $v(\bar{\theta})$.

Lemma 1 and 2 have an important implication: the (RC) condition is always slack in Problem 1. Indeed, the reservation value $v^\star(\theta)$ is (weakly) increasing, while the incentive-compatible payoff $v(\theta)$ is (weakly) decreasing, and coincide at $\bar{\theta}$. This ensures that, in the highest type's sub-problem, each type's payoff exceeds its reservation value. Hence, the MWS equilibrium solves the final sub-problem without the (RC) constraint:

Corollary ***1 (Relaxed MWS Problem)*** *The repo equilibrium solves:*

$$\max_{\{\boldsymbol{p}(\theta), q(\theta): \forall \theta \in \Theta\}} v\left(\boldsymbol{p}\left(\bar{\theta}\right), q\left(\bar{\theta}\right); \bar{\theta}\right) \tag{3.5}$$

subject to the (IC) *and* (ZPC) *constraints,* $\forall \theta \in \Theta$.

3.2. *Repo equilibrium: a single pooling contract*

Corollary 1 suggests that the highest type obtains a payoff exceeding the lowest type's reservation, $r\underline{\theta}$, ensuring full participation. Building on these observations, we conclude that the equilibrium induces full participation without involving quantity screening.

Lemma ***3 (Full Participation and Full Trades)*** *All borrowers sign a contract that trades their entire quantity, i.e.,* $q(\theta) = 1$.

Proof. See Appendix A.4. ‖

Full participation may be unsurprising when sellers can signal their quality by trading a limited fraction of their assets. The novelty is that full participation is guaranteed regardless of the possibility of quantity screening. In fact, the repurchase option is the sole force that induces full participation. Even if assets were indivisible and rationing is impossible, as in the Akerlof setting, repurchase options bring otherwise non-participating high-quality assets to the market. Indeed, when repurchase options are available, quantity screening is never deployed, as we establish in Lemma 3.

Intuitively, quantity screening does not emerge in equilibrium because repurchase options are a more effective screening device. With outright sales, high-quality types may wish to distort their traded quantities to separate from lower-quality types. In contrast, with repos, lenders can attract high-quality types by adjusting the repurchase price instead. This adjustment can compensate for a reduction in the loan amount with a lower interest charge without distorting the traded quantities.

To elaborate, consider contracts that attract non-default and default types. Neither type of contract will use quantity screening. For non-default types, suppose that at least one contract features $q<1$. We can construct an alternative contract with a slightly higher q and lower p_r. Since both adjustments benefit the non-default types, we can reduce the sales prices, p_s, by a larger magnitude than the reduction in the repurchase price such that the non-default payoff remains the same. The default types will prefer to stay with their existing contracts since they do not benefit from the reduced repurchase price as the non-default types. Thus, this deviation that increases q increases lender profit without disturbing the (IC) constraints. Hence, all non-default contracts must feature full trades.

Contracts designated to default types will not involve fractional trades either. To understand why, we study the dual of the relaxed MWS problem in Corollary 1—a lender profit-maximization problem formalized in Appendix A.4. Given any default threshold, delivering the maximum highest-type payoff requires maximizing lender profits among default types, respecting their (IC) constraints. This is a standard mechanism-design problem whose solution pools default types into a full-trade contract.[14]

Given that all borrowers trade their entire quantity, from now on, we reduce the contract notation to only the price dimension $\boldsymbol{p}$. Furthermore, given any equilibrium default threshold, borrowers who default select the highest sales-price contract, while borrowers who do not default select the contract with the highest non-default payoff. These choices are intuitive: borrowers who default only care about the sales price, while borrowers who intend to repurchase face a trade-off between a higher loan size and a lower implied interest rate. Therefore, it is sufficient to consider two contracts, one for default types and another for non-default types. Formally, we have the following lemma.

Lemma *4 **(Borrower Contract Choice)*** *In any repo equilibrium:*

(i) The default types $\theta<\theta^d$ choose the contract with the highest sales price $\boldsymbol{p}^d$.
(ii) The non-default types $\theta\geq\theta^d$ choose the contract that delivers the highest non-default payoff $\boldsymbol{p}^n$.

While Lemma 4 reduces the equilibrium to at most two contracts, these contracts can potentially be the same, leading to a complete pooling result. The following result establishes that this is indeed the case.

Proposition *1 **(Pooling)*** *The repo equilibrium features a single pooling contract $\boldsymbol{p}^*$. This contract delivers the highest non-default payoff among all zero-profit pooling*

14. Recall we have dropped the (RC) constraint. The repurchase option once again plays a pivotal role: in its absence, the (RC) constraint could lead to separations through fractional trades.

contracts:

$$\max_{\boldsymbol{p}}\{(1+r)p_s - p_r\} \tag{3.6}$$

$$s.t. \quad p_s = \mathbb{E}[\min\{\theta, p_r\}]. \tag{3.7}$$

Moreover, the default threshold satisfies $\theta^d = p_r^*$.

Proof. See Appendix A.5. ‖

Proposition 1 establishes that the repo equilibrium features a single contract pooling all assets. This property can be explained intuitively. Suppose—toward a contradiction—that there were two distinct contracts: one intended for defaulters and another for non-defaulters. By Lemma 4, to separate types, the contract intended for defaulters must offer a strictly *higher* sales price than the one intended for non-defaulters. However, that price gap admits a strictly profitable deviation: one can slightly increase the sales and repurchase prices of the non-default contract (by ϵ and $(1+r)\epsilon$, respectively). This profitable deviation contradicts the premise that the original separating contracts maximized the highest type's payoff or, equivalently, the non-default payoff, as characterized by Corollary 1. Consequently, the only possibility is a single pooling contract.

Once we establish a pooling equilibrium, Proposition 1 reformulates the MWS problem into finding a zero-profit pooling contract that maximizes the non-default payoff, which is the highest type's payoff. The default threshold naturally coincides with the contract's repurchase price. Accordingly, equation (3.7) restates the zero-profit condition: the left-hand side is the funds extended to the borrowers, and the right-hand side is the expected recovery values or repayments from the asset pool.

All in all, the repurchase option plays a dual role in establishing a pooling equilibrium. On the one hand, it breaks the *single-crossing* property that would otherwise hold for outright sales, thereby eliminating the possibility of separation through quantity screening (Lemma 3). On the other hand, it is also critical to attract defaulters and non-defaulters to the same contract (Proposition 1).

3.3. *Analytic solution for continuum types*

For any continuum-type distribution F, the equilibrium repo solving the problem in Proposition 1 has a simple analytic solution. To solve for the repo terms, we consider the marginal effect of reducing the repurchase price p_r by one dollar on the non-default payoff in (3.6). For lenders to break even, the sales price p_s must adjust according to the zero-profit condition (3.7). The repayment from each non-default type goes down by one dollar, a total of $1-F(p_r)$ dollars. Hence, the sales price lenders can offer goes down by $1-F(p_r)$ dollars. The highest non-default payoff is obtained when the marginal benefit of liquidity is equal to the marginal repayment cost: $(1+r)(1-F(p_r))=1$. Rearranging this condition yields the formula below.

Corollary *2 **(Analytic Solution)** Consider a continuum-type distribution, i.e.,* $f(\theta)>0$, $\forall\theta\in[\underline{\theta},\overline{\theta}]$. *The equilibrium default rate* $d=\frac{r}{1+r}$, *and the equilibrium repo prices* *are:*

$$p_r^* = F^{-1}\left(\frac{r}{1+r}\right) \tag{3.8}$$

$$p_s^* = \mathbb{E}\left[\min\left\{\theta, F^{-1}\left(\frac{r}{1+r}\right)\right\}\right]. \tag{3.9}$$

Proof. See Appendix A.6. ‖

While Corollary 2 presents formulas valid for continuous distributions, the solution generalizes to arbitrary distributions by simply replacing $F^{-1}(\frac{r}{1+r})$ by an appropriately defined $\frac{r}{1+r}$ quantile of the distribution. This generalization holds because the non-default payoff in Proposition 1 is increasing in p_r when $F(p_r) \leq \frac{r}{1+r}$.[15] If, for example, the distribution exhibits a mass point at $\underline{\theta}$ such that $F(\underline{\theta}) \geq \frac{r}{1+r}$, the repo terms become $\{\underline{\theta}, \underline{\theta}\}$. Likewise, the formula also extends to discrete-type distributions. For example, in the two-type case, $\theta \in \{\underline{\theta}, \bar{\theta}\}$, studied in Online Appendix B.5, the repo terms are $\{\underline{\theta}, \underline{\theta}\}$ if the fraction of low types exceeds $\frac{r}{1+r}$; otherwise, the repo terms are $\{\bar{\theta}, \mathbb{E}[\theta]\}$.

The simple formula in Corollary 2 is appealing as it renders a transparent analysis. This formula also translates into simple expressions for haircuts and repo rates:

$$h = 1 - \frac{p_s^*}{\mathbb{E}[\theta]} = 1 - \frac{\mathbb{E}\left[\min\left\{\theta, F^{-1}\left(\frac{r}{1+r}\right)\right\}\right]}{\mathbb{E}[\theta]} \tag{3.10}$$

$$\tilde{r} = \frac{p_r^*}{p_s^*} - 1 = \frac{F^{-1}\left(\frac{r}{1+r}\right)}{\mathbb{E}\left[\min\left\{\theta, F^{-1}\left(\frac{r}{1+r}\right)\right\}\right]} - 1. \tag{3.11}$$

These expressions convey useful information to understand the margins of adjustment of the repo terms:

$$\frac{F^{-1}(\frac{r}{1+r})}{\mathbb{E}[\theta]} = (1-h)(1+\tilde{r}) \quad \text{and} \quad \frac{\mathbb{E}[\theta | \theta \leq F^{-1}(\frac{r}{1+r})]}{F^{-1}(\frac{r}{1+r})} = \frac{r-\tilde{r}}{r}. \tag{3.12}$$

Clearly, the haircut is tied to the value of the $\frac{r}{1+r}$ quantile of the θ distribution relative to the unconditional quality average. For a small borrower return r, the $\frac{r}{1+r}$ quantile is in the left tail of the quality distribution. Thus, a large haircut indicates that the left-tail quantile is far from its mean. In turn, the repo rate is tied to the conditional mean below the $\frac{r}{1+r}$ quantile relative to the quantile's value. Hence, a low repo rate implies that the left tail below the quantile decays fast. Taking these observations together, we conclude that an increase in the quality dispersion shifts the left tail away from its mean, generating a more significant haircut. If the increase in dispersion does not substantially increase the decay rate at the left tail, it will manifest in a minor adjustment in the repo rate, albeit a significant impact on the haircut. We come back to these margins of adjustments in the application. We provide the comparative statics of repo terms to changes in r and $F(\cdot)$ in the Online Appendix B.4.

15. We define the quantile function $Q(x) \equiv \inf\{\theta \in \Theta : F(\theta) \geq x\}$. The general formula $p_r = Q(\frac{r}{1+r})$ is the solution to the problem in Proposition 1.

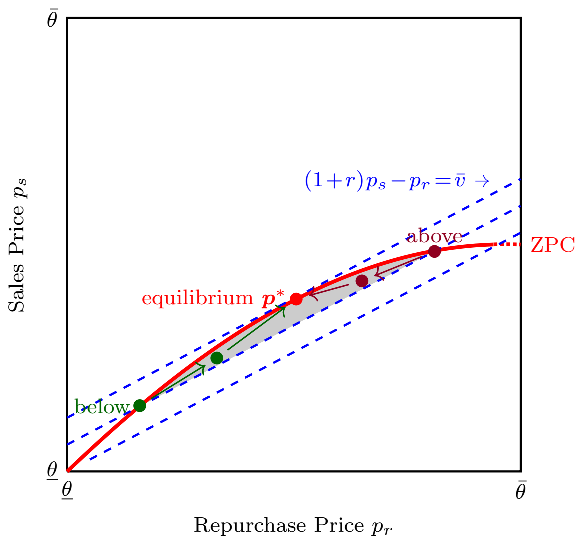

FIGURE 1
Equilibrium repo

Graphical illustration. Corollary 2 renders an intuitive graphical analysis. Figure 1 is constructed in a $[\underline{\theta},\bar{\theta}] \times [\underline{\theta},\bar{\theta}]$ box, the space of possible contracts. Sales prices are depicted on the vertical axis, and repurchase prices are on the horizontal axis. The solid red curve plots the zero-profit curve, representing the pooling contracts that satisfy the zero-profit condition (3.7). This curve is increasing and strictly concave. The dashed lines present iso-value curves conditional on non-default, with higher curves representing higher non-default values. The highest non-default payoff is obtained when an iso-value line with a slope of $1/(1+r)$ is tangent to the zero-profit curve, which has a slope of $1-F(p_r)$. This tangency point coincides with the formula in Corollary 2.

This graph also illustrates why the pooling contract in Corollary (2) is the equilibrium. Any candidate equilibrium involving a zero-profit pooling contract away from the tangency point would be broken by introducing another contract offering a higher non-default payoff, thus cream-skimming the high-quality non-default types. The deviating contract is profitable regardless of whether the original contract withdraws. Such a deviating contract remains profitable regardless of whether the original contract is withdrawn. Conversely, any contract raising the non-default payoff above the tangent line would attract all non-default types, causing the original equilibrium repo contract to become unprofitable and withdrawn, ultimately leading to losses. Thus, only the tangent contract survives as an equilibrium. Following similar logic, the graph also illustrates our earlier argument that two separating contracts—one for defaulters and another for non-defaulters—cannot coexist, demonstrating that the unique pooling contract at the tangency point is the sole equilibrium.

4. ROBUSTNESS AND EXTENSIONS

4.1. *Robustness to trading environments*

We first present the repo equilibrium under alternative competitive market environments and demonstrate the robust features of repo markets.

Contract non-exclusivity. The functioning of the assets market with asymmetric information tends to vary depending on whether contracts are exclusive or not (see, e.g., Attar et al., 2011; Kurlat, 2016; Asriyan and Vanasco, 2023). Notably, Attar et al. (2011) establishes that when contracts are non-exclusive, the asset-sales equilibrium is pooling, akin to the classic analysis by Akerlof (1970), despite the availability of quantity screening devices. This pooling outcome contrasts markets where contracts are exclusive, in which sellers will separate through quantity screening.

Contract exclusivity, however, does not affect repo markets. Our baseline MWS environment can be considered a market with exclusive contractual arrangements. Even when agents can participate non-exclusively in many contracts, the repo market outcome remains unaltered. This robust feature is intuitive: Non-exclusive competition blocks the cream skimming of higher-quality sellers being traded in small quantities, as sellers can always trade their remaining quantities with other buyers. In repo markets, cream skimming through quantity is undesirable, as mentioned in Lemma 3. Instead, an adjustment through repurchase prices or, equivalently, the effective loan interest rate is more desirable. Hence, the price adjustment is unaffected by exclusivity.

Competitive search with capacity constraints. We also study repo contracts in competitive search markets with asymmetric information, adopting the framework developed by Guerrieri et al. (2010). We leave the full details of the competitive search market in Online Appendix C.1. Under competitive search, lenders can freely post contracts without cost but face *capacity constraints*: each lender can service at most one borrower.[16] As a result, rationing can occur if one side of the market exceeds the other. The short side of the market gets perfectly matched, while the long side is rationed. We thus adapt our earlier interpretation of the quantity of divisible assets as part of the contract to the potential rationing of trades for indivisible assets as an outcome of market clearing. The capacity constraint is essential: it affects the off-equilibrium beliefs and, in turn, the equilibrium outcome.

With only outright sales, as established by Guerrieri et al. (2010), the equilibrium is fully separating, regardless of the quality distribution. With a continuum of types, the trading probability for type θ is $q(\theta)=(\underline{\theta}/\theta)^{\frac{1+r}{r}}$.[17] In turn, the total funds raised are $\mathbb{E}\left[(\underline{\theta}/\theta)^{\frac{1+r}{r}}\theta\right]$. When repo contracts become available, the equilibrium outcome is no longer separating but pooling again.

It is worth noting that pooling arises for the same reasons it does in the MWS setting: the repurchase option breaks the single-crossing property, which is crucial to allow for separation. Intuitively, with asset sales only, separation requires distorting

16. Lender capacity can be interpreted as lender size. Thus, the competitive search environment describes a market with small lenders. In contrast, our MWS formulation fits markets with large lenders.

17. Separation operates as follows. Given the capacity constraint, the off-equilibrium belief is that a deviating contract will attract the assets with the most gain and, hence, those that can bear the lowest trading probability—the lowest types. Thus, lenders form pessimistic beliefs and refrain from offering a pooling contract, anticipating it will attract low types.

traded quantities downward as we increase the type. However, this also means that the value of trading is distorted downward, which is too immediate to see by comparing incentive constraints. However, by offering a repurchase price equal to the value of the previous types, a repo allows higher types to obtain at least the same value as lower types. The possibility of offering the same value as to lower types means that if an equilibrium were to separate a type on the margin, a repo contract could be offered that attracts the separating type without making losses. This property is robust to the equilibrium refinement.

Despite featuring a pooling contract, the equilibrium refinement is different and, therefore, so is the equilibrium outcome: the refinement affects the lender's beliefs off-equilibrium. Lenders form pessimistic beliefs that the low types have more to gain and will be the first to show up. Such beliefs force the equilibrium contract to adjust to the lowest possible price, i.e., $\boldsymbol{p}=\{\underline{\theta},\underline{\theta}\}$. While repos always improve payoffs relative to outright sales for all asset types in the competitive search market, they deliver the worst contract among zero-profit pooling contracts.[18]

Contract additions. A key feature of the MWS equilibrium is that contracts can be withdrawn but not added. We also consider alternative contract adjustment processes, allowing for contract additions, leading to the reactionary equilibrium in Riley (1979). With outright sales, Riley's equilibrium induces separation, coinciding with the competitive search equilibrium. Similarly, the repo market outcomes in the two environments also coincide. That is, lenders set repo terms to fully protect themselves against collateral quality risk. They extend a minimum loan amount (maximum haircut) and charge a zero repo rate.

All in all, repo markets feature a single contract that pools all market participants across equilibrium concepts. While the precise terms differ, and the outcome under MWS Pareto dominates the alternative outcomes, there are conditions under which all outcomes converge. For example, all outcomes converge when the mass of worst types exceeds $\frac{r}{1+r}$. Furthermore, for a reasonably low return r, the $\frac{r}{1+r}$ quantile falls in the far left tail of the quality distribution, implying a high haircut relative to the average collateral value, making outcomes across all concepts more similar. These insights highlight a fundamental consistency in that repo markets are characterized by high haircuts and low or zero interest rates.

Microfoundation: the Netzer and Scheuer (2014) trading game. The repo market outcome under the MWS concept can also be obtained by formulating a three-stage trading game, as proposed by Netzer and Scheuer (2014). This trading game unfolds as follows:

(i) *Entry stage*: Each lender offers a set of contracts.
(ii) *Withdrawal stage*: After observing the contracts offered, lenders can withdraw all their contracts incurring a cost κ.
(iii) *Selection stage*: Borrowers choose contracts from the remaining set.

18. The competitive search and the Riley equilibria share an unappealing feature: with an arbitrarily small amount of zero-quality assets, i.e., $\underline{\theta}=0$, a "no-trade" equilibrium is reached. In this scenario, both repos and outright sales generate zero liquidity, a situation that generally does not occur under the MWS concept.

As the withdrawal cost $\kappa \to 0$, the equilibrium converges to the repo contract described in Proposition 1 and Corollary 2.[19] Thus, the Netzer and Scheuer (2014) trading game is a microfoundation for the MWS concept with contract withdrawals. Netzer and Scheuer (2014) proposes that withdrawal costs can be interpreted as a proxy for legal-administrative or reputational costs, or as a commitment device. Our finding that the MWS outcome generally Pareto dominates both the competitive search and Riley's outcomes provides another rationale for the Netzer and Scheuer (2014) trading protocol. This protocol may emerge as a form of market design or platform competition strategy. By delivering Pareto-dominant outcomes, this protocol can attract borrowers and lenders to specific markets or platforms.[20]

4.2. *Extensions*

Next, we discuss two extensions that modify the model's predictions about the equilibrium default rate.

Private values. Borrowers may hold a higher valuation for collateral than lenders. This gap can result from differences in tax advantages, expertise, or costs in seizing collateral after default. To introduce private values, let ψ stand for the lender's recovery value relative to the borrowers. We consider a small valuation gap, $\psi \in [1-r,1]$. The zero-profit condition (3.7) is modified to:

$$p_s = \psi \mathbb{E}[\theta | \theta < p_r] F(p_r) + p_r (1 - F(p_r)) \tag{4.13}$$

The characterization in Proposition 1 extends to this case: the repo contract maximizes the non-default payoff (3.6) subject to the modified zero-profit condition (4.13). A more general version of the repo formula in Corollary 2 follows. Specifically, the repurchase price solves:

$$F(p_r) + (1-\psi) p_r f(p_r) = \frac{r}{1+r}. \tag{4.14}$$

This formula adds the term $(1-\psi)p_r f(p_r)$ to the baseline formula (3.8). Naturally, as the difference in valuations vanishes, $\psi \to 1$, we recover the baseline formula. In the Online Appendix B.1, we show that the modified zero-profit condition (4.13) is concave if the density $f(\cdot)$ does not decay too fast. Moreover, the solution to (4.14) is interior if the density of lowest type $f(\underline{\theta})$ is not too large. Otherwise, the solution is at a corner, $\{\underline{\theta}, \underline{\theta}\}$, in which case the equilibrium outcomes coincide across the different environments considered—MWS, competitive search, and Riley's. Interestingly, this corner-solution outcome under MWS with private valuations also matches the zero-default, maximal-haircut outcome from an alternative framework with heterogeneous valuations studied by Fostel and Geanakoplos (2015). This equivalence is particularly notable given that the core friction in that model is limited enforcement rather than private information.

19. Absent withdrawal costs, multiple equilibria can be sustained by offering some contracts, which will be withdrawn on the equilibrium path but maintained if there are deviations in the offering stage. An arbitrarily small withdrawal cost rules out these strategies.

20. For example, a trading platform could implement an upfront participation fee, refundable only if lenders actively participate in lending. This fee effectively serves as a withdrawal cost. Such a mechanism can ensure the MWS outcome and attract borrowers to the platform.

Several observations follow. First, lower recovery values also translate into smaller loan amounts—higher haircuts—which protect lenders from costlier defaults. Second, with differences in valuation, the default rate becomes sensitive to changes in the asset-quality distribution. Finally, as the lender recovery value decreases, i.e., ψ decreases, the repurchase price decreases to the point that there may be nodefaults.

Aggregate risk. We now consider private information regarding the assets' exposure to aggregate risk. We introduce an aggregate risk event, which occurs with probability λ. If the event is unrealized, all assets have a value of 1. If the event is realized, an asset drops in value to $\theta \in [\underline{\theta}, \bar{\theta}]$, with $\bar{\theta}=1$. Hence, θ captures exposure to the aggregate event. As in the baseline, only borrowers know θ, although their distribution $F(\cdot)$ is common knowledge. Repo contracts are signed before the event is realized; repayment is decided after. In this case, type-θ's payoff accounts for the possibility of the aggregate event:

$$(1+r)p_s - [(1-\lambda)p_r + \lambda \min\{\theta, p_r\}]. \tag{4.15}$$

Accordingly, the zero-profit condition (3.7) becomes:

$$p_s = (1-\lambda)p_r + \lambda \mathbb{E}[\min\{\theta, p_r\}]. \tag{4.16}$$

Again, the repo equilibrium maximizes the non-default payoff subject to the zero-profit condition (4.16). Assuming that the aggregate risk is high enough, $\lambda > \frac{r}{1+r}$, the repo formula generalizes to:

$$p_r = F^{-1}\left(\frac{1}{\lambda}\frac{r}{1+r}\right)$$

Naturally, we recover the original formula (3.8) as $\lambda \to 1$.

Under this extension, if the equilibrium is repeated over time, there is a clear distinction between the time series and cross-sectional default rates. While the unconditional default rate is the same as in the baseline model, $\lambda F(p_r) = \frac{r}{1+r}$, the fraction of periods with zero defaults is $1-\lambda$. However, when the aggregate event is realized, the default rate will be higher than the baseline's: $\frac{1}{\lambda}\frac{r}{1+r} > \frac{r}{1+r}$. Thus, with aggregate events, repo defaults are concentrated in time, with potentially long spans without observed defaults.

5. EFFICIENCY

In this section, we study *ex-ante efficiency* and *interim efficiency*, following the Holmstrom and Myerson (1983) terminology. The latter notion examines the individual gains from trade after types are drawn, while the former examines overall liquidity provision.

Interim efficiency. When repo markets open in the interim stage, the equilibrium is Pareto efficient. This interim-efficiency property immediately follows from the characterization of the MWS equilibrium as a sequence of Pareto problems in Problem 1. This characterization holds for any contractual possibility, restricting to outright sales or including repos. In either environment, after borrowers know their asset's quality, adjusting prices to make some borrowers better off without making others worse off is impossible.

While both outright sales and repo environments are interim efficient, we should ask whether introducing repos can lead to Pareto improvements over outright sales. As argued above, with repos, lenders attract high-quality borrowers through the repurchase option, which gives them higher payoffs than through quantity screening. Thus, repos improve the payoffs of high-quality borrowers. As for lower types, an improvement in their payoffs is not guaranteed in general.

The lack of a simple solution for the outright-sales equilibrium—which may feature full or partial separation—precludes sharp conditions for Pareto improvements. Nevertheless, a two-type example suffices to show that interim Pareto improvements are possible: Consider again the two-type case $\theta \in \{\underline{\theta}, \bar{\theta}\}$ discussed in Section 3.3. When the fraction of low types exceeds $\frac{r\bar{\theta}}{(1+r)\bar{\theta}-\underline{\theta}}$, the sales-only outcome is separating: the low type sells its full unit at a price $\underline{\theta}$ and obtains a payoff $r\underline{\theta}$, whereas the high type sells a distorted fraction and obtains a lower payoff. In this case, introducing repos leads to Pareto improvements: given the fraction of low types exceeds $\frac{r}{1+r}$, a repo contract $\{\underline{\theta}, \underline{\theta}\}$ arises. Thus, the high type is strictly better off with a payoff $r\underline{\theta}$, while the low type is unaffected. However, if the sales-only outcome is pooling, introducing repos may hurt the low type. More generally, similar Pareto improvements from repos can also be found numerically for continuum-type distributions, particularly those that induce substantial adverse selection.

Ex-ante efficiency: liquidity-maximizing repo. To formalize an ex-ante efficiency notion, we consider a prior date, $t=0$, in which borrowers are identical and have not drawn their asset quality yet. From this perspective, efficiency concerns the maximization of their ex-ante expected payoff or, equivalently, aggregate liquidity in the interim stage.

While the repo market is interim efficient, it generically fails to be ex-ante efficient. To formalize this, consider a planner who maximizes the borrower's ex-ante expected payoff and faces the same interim-stage incentive and participation constraint as lenders. This planner offers a menu of repo contracts $\{\boldsymbol{p}(\theta), q(\theta) : \forall \theta \in \Theta\}$, which must also be overall budget balanced.[21] The planner's problem is:

$$\max_{\{\boldsymbol{p}(\theta), q(\theta) : \forall \theta \in \Theta\}} \mathbb{E}[v(\boldsymbol{p}(\theta), q(\theta); \theta)] \tag{$\mathcal{P}$}$$

subject to the (IC) constraint, the participation constraint, $v(\boldsymbol{p}(\theta), q(\theta); \theta) \geq 0$, and the zero-profit condition (ZPC).

The planner's problem differs from MWS problem in Corollary 1 in one critical dimension: whereas the planner maximizes the expected borrower payoff, the market solution maximizes the payoff of high-quality non-default types. Intuitively, the difference is that the planner does not cream-skim high-quality borrowers and instead maximizes overall liquidity. The following proposition establishes that the ex-ante efficient outcome is also a pooling repo contract, but with different terms from the market repo.

Proposition *2 **(Liquidity-Maximizing Repo)** Ex-ante efficiency features a pooling contract, $\boldsymbol{p}^p$, that maximizes liquidity:*

(i) If adverse selection is minimal, i.e., $(1+r)\mathbb{E}[\theta] \geq \bar{\theta}$, $p_s^p = \mathbb{E}[\theta]$ and $p_r^p = \bar{\theta}$.

21. Because borrowers cannot commit to participating, it does not matter whether the planner offers the contracts at $t=0$ or at $t=1$.

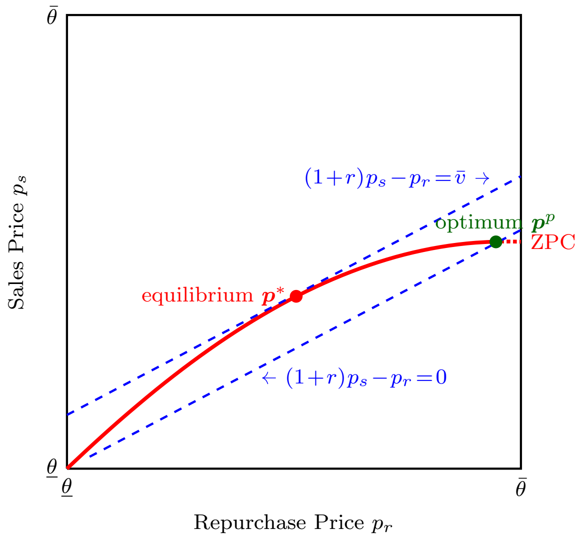

FIGURE 2
Optimal repo contract

(ii) *If adverse selection is substantial, i.e.,* $(1+r)\mathbb{E}[\theta] < \bar{\theta}$, *the repurchase price solves* $(1+r)\mathbb{E}\left[\min\{\theta, p_r^p\}\right] = p_r^p$ *and* $p_s^p = (1+r)p_r^p$.

Proof. See Appendix A.7. ||

Proposition 2 suggests that the liquidity provision in repo markets is generically inefficiently low. Like the market equilibrium, the planner also pools all borrowers together. Unlike the market, where lenders compete for high-quality borrowers by lowering interest rates and, thus, funding, the planner is not tempted by competition. The planner is only concerned about total funding and, therefore, seeks maximum cross-subsidization from high to low types, potentially pushing high types to their participation constraint. That is, the liquidity-maximizing repo features a lower haircut and a higher rate. This comparison indicates that the repo market induces inefficiently high haircuts, undermining liquidity.

The ex-ante optimal solution exploits repos to induce participation, never deploying quantity screening. When adverse selection is minimal, that is, when the highest-quality borrower is willing to sell at a price equal to the unconditional average quality, the planner offers that price and achieves the first best. When the highest type is unwilling to sell at that price, the planner utilizes the repurchase option, lowering the repurchase price just enough to attract highest-quality type to the pool. This latter scenario is depicted in Figure 2.

The ex-ante optimal repo contract here coincides with an optimal ex-ante security design in which lenders offer a $t=0$ contract (see Biais and Mariotti, 2005). Security design has an embedded commitment to a set of contracts, which the market solution does not. In the market outcome, borrowers *cannot commit* to past contracts after they learn their qualities. Under the security-design approach, borrowers *commit* to an ex-ante menu of contracts. However, they cannot commit to choosing specific contracts or participating after learning their types. Hence, while the security design solution must also respect interim incentive and participation constraints, it maximizes liquidity, given an ex-ante commitment to contracts—i.e., a security design. The lack of commitment to ex-ante contracts gives rise to inefficient liquidity provision.

Despite their inefficiencies, repos are a valuable financial innovation: they can increase liquidity relative to outright sales markets when adverse selection is severe. We compare the provision of liquidity between repos and sales in greater detail in Online Appendix B.6.

Discussion: interventions in the repo market. Adverse selection in financial markets was a significant motivation behind the repo and asset-backed security policy interventions during the 2008-2009 financial crisis (Gorton et al., 2020). Motivated by these arguments, Philippon and Skreta (2012) and Tirole (2012) study policy interventions in asset markets plagued by adverse selection, though they focus primarily on sales contracts. As we demonstrated earlier, repos induce full participation, but liquidity provision can still be ex-ante inefficient. This suggests the desirability of policy interventions to enhance liquidity.

Since sub-optimal liquidity manifests through high haircuts and low repo rates, we can conjecture that taxes can affect these terms. The challenge for analyzing tax interventions formally is that one must consider how policies alter the nature of the equilibrium: Would these policies preserve the pooling outcome? Do we need to specify off-equilibrium policies, and if so, what form of government commitment should we consider? Are linear, uniform, or non-linear taxes enough to implement a desired outcome? While there are precedents in studying optimal taxes in "rat race" labor markets (see Stantcheva, 2014), investigating these tools in repo markets falls beyond the scope of this paper.

6. APPLICATION: REPOS FOR MBS

Our theory applies to asset markets featuring repo-like financing where asymmetric information may exist between borrowers and lenders. These markets include repo markets as well as other collateralized debt markets where informational asymmetries are relevant.

This section examines whether our theory can generate observed haircuts and repo rates for different MBS tranches. Several empirical studies suggest that private information is prevalent in both the primary, secondary, and repo MBS markets.[22] Hence, this market is a natural setting to test our theory. To do so, we recast the model to capture institutional details of the MBS repo market and map the model to the empirical patterns of repo rates and haircuts. We, therefore, adhere to the MWS concept because, as shown above, other equilibrium notions yield corner-solution outcomes that fail to generate meaningful predictions about repo rates.

6.1. *MBS reformulation*

We consider an environment where investment funds (borrowers) use repo financing to purchase MBS tranches. As in the baseline model, we assume risk neutrality.

22. See for example Glaeser and Kallal (1997); Agarwal et al. (2012); An et al. (2011); Becker et al. (2023).

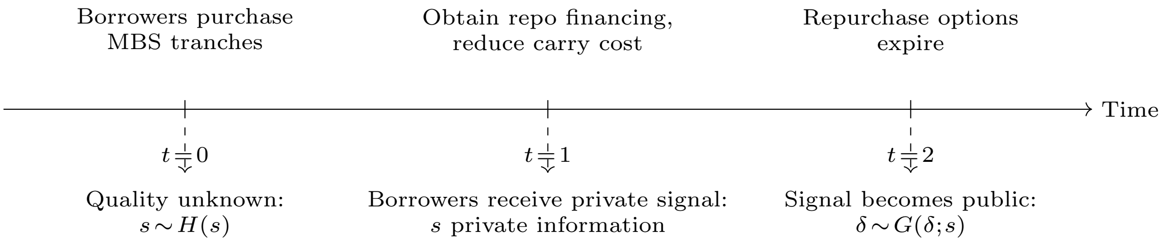

FIGURE 3
Timing in the MBS markets

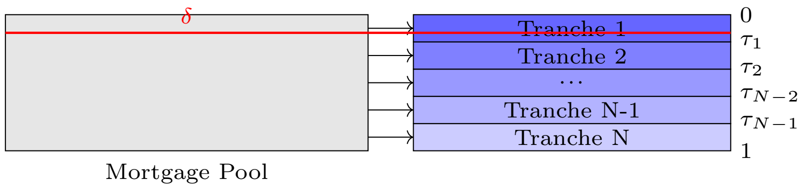

FIGURE 4
Mortgage pool and tranches

Timing and payoffs. The timing now comprises three dates, $t \in \{0,1,2\}$ described in Figure 3. At $t=0$, borrowers purchase different MBS tranches at market prices. At $t=1$, they obtain repo financing to fund their investment with a return of r, i.e., to reduce carry costs. As we explain below, the repo market is subject to asymmetric information as in the baseline model. At $t=2$, repurchase options expire.

A mortgage pool consists of a continuum of individual mortgages. Each mortgage is a promise to pay a dollar by $t=2$. Let $\delta \in [0,1]$ represent the fraction of mortgage defaults and, thus, the default loss in the pool. Without loss of generality, we treat δ as the default loss without reference to recovery values. The pool's realized payoff is $1-\delta$.[23]

A mortgage pool is sliced into $N \geq 1$ tranches as illustrated in Figure 4. Each MBS tranche $i \in \{1,...,N\}$ has an assigned seniority, with tranche N being the most senior. The tranching structure $\{\tau_i\}_{i=0}^{N}$ is a strictly increasing sequence, with $\tau_0 = 0$ and $\tau_N = 1$. This tranching determines the payoffs $\{y_i(\delta)\}_{i=1}^{N}$ of the tranches for a given default probability δ. Specifically, the payoff of tranche i is:

$$y_i(\delta) = \max\{\tau_i - \max\{\tau_{i-1}, \delta\}, 0\}.$$

The tranching structure distributes the pool's payoff $1-\delta$ across tranches: the sum of the MBS tranche payoffs equals the total pool's payoff, $\sum_{i=1}^{N} y_i(\delta) = 1-\delta$. In turn, each tranche either pays out in full, $\tau_i - \tau_{i-1}$, a partial amount, $\tau_i - \delta$, or zero. The seniority structure suggests that if tranche i does not receive its full payoff, its junior tranches pay zero; conversely, if a tranche i gets a positive payoff, its senior tranches are paid in full.

23. We capture default risk as the primary driver of mortgage pool quality during the period leading up to the 2007 subprime crisis. While prepayment risk, rather than default risk, has become the main concern in the current period, our analysis extends to MBS quality issues related to prepayment risk.

Information and prices. The default probability $\delta \sim G(\delta; s)$, where s denotes a state. At $t=0$, s is unknown. However, the realization of s, and hence the default distribution, are drawn *i.i.d.* across pools from a known distribution $H(s)$. The realization of s governs the distribution of defaults in the mortgage pool, affecting each tranche's expected payoff.

The critical informational assumption is that borrowers learn the pool's state s at $t=1$ when the repo market trades are carried out. At $t=2$, when the repo market settles and the market for MBS tranches reopens, the state s becomes public information; hence, the default distribution $G(\delta; s)$ is known. We model the default risk δ as random and not fully resolved at $t=2$ to generate a continuous distribution of tranche values; otherwise, the tranche values would feature mass points at zero and the full tranche payoff.

The interpretation of the environment is as follows. The realization of a specific distribution, $G(\cdot; s)$, can be considered the posterior default-rate distribution of the mortgage pool, calculated from signals s. The borrowers' anticipated knowledge reflects faster data collection or better risk modeling. Hence, the setup captures a situation where borrowers have an informational advantage regarding the MBS tranches, while lenders must rely on imperfect priors. The $t=0$ uncertainty over the s reflects the sequential release of information.

We first describe the evolution of market prices for the MBS tranches at $t=0$ to $t=2$. The $t=2$ price of tranche i is the expected payoff $y_i(\delta)$, relative to the tranche size $\tau_i - \tau_{i-1}$, according to the posterior distribution of default risk:

$$p_{i,2}(s) = \int_0^1 \frac{y_i(\delta)}{\tau_i - \tau_{i-1}} dG(\delta; s). \tag{6.17}$$

The $t=0$ price of tranche i is the expectation of the $t=2$ price over the prior distribution:

$$p_{i,0} = \int p_{i,2}(s) dH(s).$$

We now derive the repo market terms for each tranche. The possible realizations of s induce a quality distribution for each tranche from the lenders' perspective. The intrinsic quality of tranche i is given by its $t=2$ price, which is perfectly anticipated by borrowers but unknown to lenders at $t=1$. Thus, $\theta_i = p_{i,2}(s)$ is a random variable with an unconditional mean $\mathbb{E}_i[\theta_i] = p_{i,0}$. We denote the cumulative distribution function of θ_i by $F_i(\theta_i)$, in accordance with the MBS pricing equation (6.17) and $H(s)$. The tranche structure naturally translates into a *first-order stochastic dominance* in the quality distribution by seniority. This pattern results from how default losses in the mortgage pool are absorbed: for any realization of the state s, a senior tranche is less impacted than a junior one.

Let $\{p_{i,s}, p_{i,r}\}$ represent the sales and repurchase prices of tranche i in the repo market, respectively. Recall that the posterior default risk is not fully resolved at $t=2$, ensuring that the quality distribution $F_i(\theta_i)$ is continuous. Applying Corollary 2, we obtain the repo terms: $p_{i,r} = F_i^{-1}\left(\frac{r}{1+r}\right)$ and $p_{i,s} = \mathbb{E}_i\left[\min\left\{\theta_i, F_i^{-1}\left(\frac{r}{1+r}\right)\right\}\right]$. Following equations (3.10) and (3.11), the repo haircut h_i is the discount of the tranche's repo sales price from its $t=0$ notional value, and the repo rate $\tilde{r}_i$ is the percentage difference between the repurchase and sales prices. These terms are: $h_i = 1 - p_{i,s}/p_{i,0}$ and $\tilde{r}_i = p_{i,r}/p_{i,s} - 1$.

In this setting, agents do not discount payoffs, nor do lenders require a risk-free rate as a cost of funding. This is without loss of generality. Allowing for a risk-free rate leads to the same expressions for haircut and repo rates above. However, as shown in Online Appendix B.3, the return r should be interpreted as borrower excess return and the repo rate $\tilde{r}_i$ as a spread over the lender's cost of funds. When mapping our model to the data, we interpret the borrower investment return as an excess return and map the repo rates to repo spreads.

6.2. *Evidence and model fit*

We now discuss how our model fits some empirical patterns of bilateral MBS repo markets.[24] Despite being recognized as a major short-term funding source, there is no comprehensive public repo-market data repository. Most evidence is either anecdotal or obtained from specific investor segments. One notable exception is the empirical evidence in Auh and Landoni (2022) obtained from transaction-level repo data reported by several hedge funds from 2004 to 2007. That study presents detailed information regarding the MBS tranche structure and tranche credit ratings. Notably, that study presents regression analysis that isolates the effects of tranche seniority (by rating) on haircuts and repo rates. We rely on this study as a source of empirical counterparts.

We calibrate the model to match the haircuts and repo rates across the tranches. We set the tranche number N to 5, corresponding to tranches rated AAA, AA, A, BBB, <BBB in the data. Tranche 5 corresponds to the AAA tranche, tranche 4 to the AA tranche, and so on. We set the seniority structure to $\{\tau_i\}_{i=0}^5 = \{0, 0.13, 0.194, 0.278, 0.389, 1\}$.[25] We take the haircut and repo rate for the most senior (AAA) tranche and reconstruct the haircuts and rates for the junior tranches by adding the marginal effects of tranche seniority reported in the paper.[26] From AAA to <BBB, the haircuts are $\{4, 5, 8, 15, 20\}$ percent, and the annualized repo rates are $\{4, 5, 10, 27, 35\}$ basis points above the LIBOR rate, respectively.

Parameterization. We specify a logistic-normal parametric structure. In particular, the default probability at $t=2$ is a logistic transformation of a normal variable x: $\delta \equiv \exp(x)/(1+\exp(x))$, where $x \sim N(s, \sigma^2)$. This transformation implies that $x = \log(\delta/(1-\delta))$. The mean of the posterior distribution is drawn from another normal distribution, $s \sim N(\mu, \kappa^2)$. The ex-ante distribution of default risk at $t=0$ is also a logistic transformation of a normal random variable, $x \sim N(\mu, \sigma^2+\kappa^2)$. Thus, the borrower's informational advantage is measured by the factor $\kappa^2/(\sigma^2+\kappa^2)$, reflecting their reduced

24. As explained by Copeland et al. (2014), repo markets are divided into two segments: bilateral and tri-party repos. Hedge funds are active borrowers in bilateral repo markets and are considered highly specialized informed agents. In contrast, tri-party repos involve an intermediary that manages the collateral, sets collateral rules, and centralizes settlements. In the tri-party segment, security dealers typically act as borrowers, while money market mutual funds and other security dealers serve as lenders. Given the market-based contracting nature and participation of better-informed agents, our theory fits well with the bilateral repo market.

25. These tranching values are reported in Table II of Auh and Landoni (2022). A similar average tranche structure is reported in Ashcraft and Schuermann (2008).

26. The average haircut and repo rate for the AAA tranche are reported in Table II of Auh and Landoni (2022). By controlling for lender × mortgage pool × time fixed effects, the regressions reported in that paper isolate the effects of tranche seniority from lender characteristics, idiosyncratic pool features, and macroeconomic outcomes. The marginal effects on repo rates are reported in Table IV, and the marginal effects on haircuts are reported in Table V. We take the values from columns (2) of these tables.

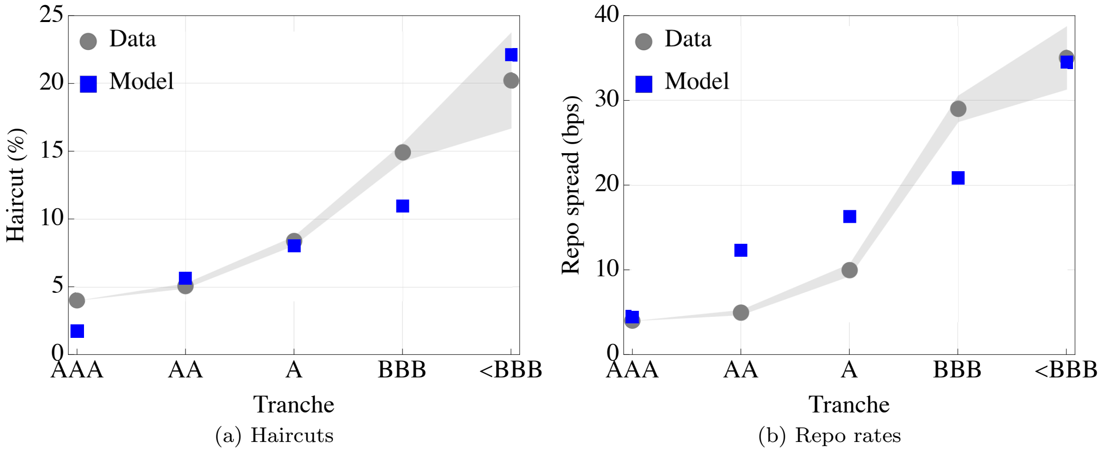

FIGURE 5

Model fit

Notes: Panel (a) plots the model-implied haircuts and the empirical point estimates of haircuts by tranche. Panel (b) plots the model-implied repo rates and the empirical point estimates of repo rates by tranche. The bands in both panels report the 95% confidence intervals of the estimates.

uncertainty or improved precision of the default risk. Overall, this structure only requires three parameters, $\{\mu,\sigma,\kappa\}$.

Model calibration. We set the model period to one quarter consistently with the sample's average maturity of repo contracts.[27] We set the borrower return, r, to be 5% annually (1.25% quarterly) over the lenders' capital cost. We estimate the parameters related to the borrowers' information advantage $\{\sigma,\kappa\}$ to fit the ten data moments—the haircuts and repo rates for the five tranches. For that, we search for the parameter combination that minimizes the squared distance between the model-implied moments and their empirical counterparts.[28] For any pair of $\{\sigma,\kappa\}$, we set a corresponding value of μ to match an ex-ante expected default loss of 2%, ensuring a reasonable unconditional mortgage default rate.[29] The estimated parameters are $\mu = -6.59$, $\sigma = 2.43$, and $\kappa = 1.01$. These values imply a superior borrower information precision of 15%.

Model fit. Figure 5 presents the model fit regarding haircuts and rates for each tranche. Panels (a) and (b) show that the model exhibits higher haircuts and repo rates as the tranche rating falls, aligning well with the empirical patterns. These patterns are

27. We obtain similar model fits and predictions when the repo maturity is reduced to one month.

28. Specifically, we minimize the loss function $\sum_1^5 \left((h_i^{model} - h_i^{data})/h_1^{data}\right)^2 + \sum_1^5 \left((\tilde{r}_i^{model} - \tilde{r}_i^{data})/\tilde{r}_1^{data}\right)^2$. We normalize the difference between the model-implied and data moments for each tranche by the corresponding data moment for the first tranche. This normalization ensures scale independence between the haircut and the repo rate moments. However, it penalizes the lack of fit across tranches in levels.

29. We choose 2% as the default rate, a realistic value, considering that the empirical average delinquency rate of Single-Family Residential Mortgages booked in Commercial Banks in the sample 2004-2007 was 3.0%-7.0%, before adjusting for recovery values and refinancing. See the Fred series DRSFRMACBS.

consistent with the idea that lower-rated tranches feature more significant informational asymmetries.

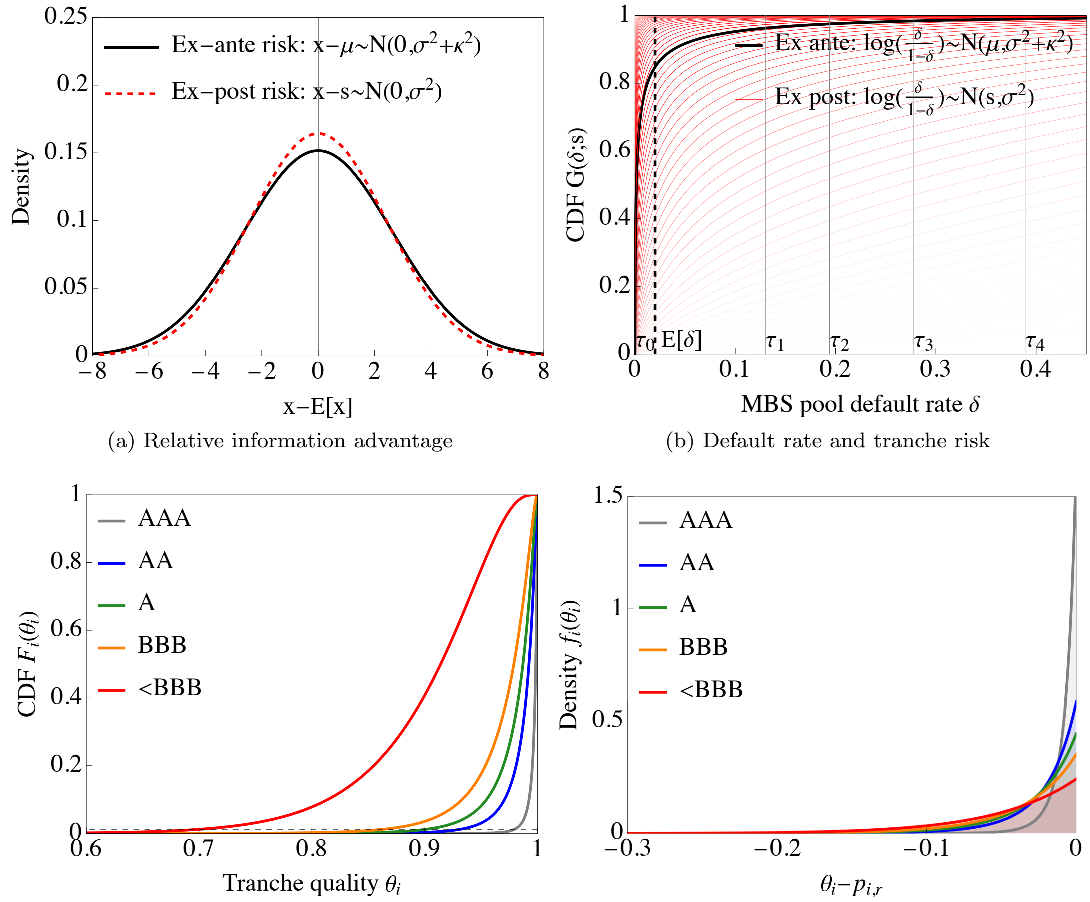

(a) Relative information advantage

(b) Default rate and tranche risk

(c) Tranche quality: whole distribution

(d) Tranche quality: left tail

FIGURE 6

Model-implied information asymmetry

Note: Panel (a) plots the ex-ante and ex-post densities of $x = \log(\delta/(1-\delta))$. Panel (b) plots the ex-ante and ex-post CDFs for the pool's default loss δ. For the ex-post distribution, the color intensities reflect the likelihood of realization. Panel (c) plots the CDF, $F_i(\theta_i)$, for each tranche, where the tranche quality θ_i is measured by $t=2$ price $p_{i,2}(s)$. The black dashed horizontal line marks the $r/(1+r)$ quantile. Panel (d) plots the density in the left tail below the quantile, and the horizontal axis is aligned relative to the quantile.

To further clarify the underlying forces, we present plots regarding the model-implied information asymmetry and predictions for the repo markets in Figure 6. Panel (a) depicts the ex-ante and ex-post distribution of the underlying states centered around the mean. The solid black line represents the ex-ante uncertainty (the lenders'), which follows a Gaussian density with variance $\sigma^2+\kappa^2$; the dashed red represents the ex-post (the borrowers') uncertainty with a variance of σ^2. The calibration suggests a relative borrower information precision of 15%, indicating that the model does not require significant informational asymmetries to replicate the empirical patterns.

The information asymmetries in assessing the pool's default risk translate into varying degrees of tranche risks. Panel (b) plots the CDFs of the default loss in the underlying MBS pool to show these effects. The thick black line represents the ex-ante CDF of δ, with an unconditional expected default loss of 2% captured by the black dashed vertical line, while the red thin lines represent the ex-post CDFs under various draws of s. Panel (b) also marks the tranche thresholds, which is convenient to understand their risk exposures. For example, for the AAA tranche, the default loss in the pool must surpass $\tau_4 = 0.389$ for it to experience losses. From an ex-ante perspective, such a high default loss is unlikely; from an ex-post perspective, while some extreme draw of high s will make it a material risk, such draws are unlikely. As a result, even if the borrowers possess private information on s, the AAA tranche is well protected against risk. In contrast, the most junior tranche <BBB is the first to absorb any losses, and when default surpasses $\tau_1 = 0.13$, it is wiped out. Thus, the pool's risk and information advantages become important predominantly for lower tranches.

These different tranche risks translate into different outcomes for $t = 2$ prices. Panel (c) plots the CDFs of quality distribution across tranches, which confirms that the distribution of a more senior tranche first-order stochastic dominates those of junior tranches.[30] Another noticeable feature is that prices become more dispersed as the tranche rating falls: the AAA tranche price is concentrated around its ex-ante mean. By contrast, the <BBB tranche is more dispersed. Panel (d) zooms into the left tail by plotting the density for qualities below the $\frac{r}{1+r}$ quantile. For the AAA tranche, the density decays fast at the left tail. The density declines more slowly for lower-rated tranches.

Recall from equation (3.12) that the haircut is tied to the $\frac{r}{1+r}$ quantile relative to the mean, whereas the repo rate is connected to the conditional average under the quantile relative to the quantile. The increasing haircut and repo rate pattern is thus explained by the increasing distance from the quantile to the mean and from the conditional mean to the quantile as the tranche rating falls. While lenders adjust haircuts and repo rates to protect themselves against asset quality risk, the lion's share of the adjustment is through haircuts. Indeed, the range of haircut values is two orders of magnitude larger than the range of values for repo rate. The model reproduces this outcome by capturing how the market competes for high-quality assets via higher haircuts and lower rates. The greater sensitivity of haircuts relative to rates is considered a puzzle. As Dang et al. (2013) put it, "the existence of repo haircuts is a puzzle, as standard finance theory would suggest that risk simply be priced." The repo market formula derived from our model rationalizes this pattern.

Repo Default. While the model can successfully reproduce the MBS repo rates and haircut patterns, a word of caution is warranted regarding its repo default-rate predictions. The predicted repo default rate is $\frac{r}{1+r}$, approximately 1% per contract. While we do not possess default data, such a rate may be implausibly high under normal market conditions. Nevertheless, the extensions in Section 4.2 help moderate the model-implied default rates. Specifically, incorporating private values leads to a lower or even zero default rate, whereas introducing aggregate risk results in periods characterized by either no defaults or concentrated defaults. It is possible to further reconcile these

30. The price distribution of tranches is in the ballpark of the range of values of ABX indices for those tranches during the years of our sample; see, for example, Ashcraft and Schuermann (2008), Table 16.

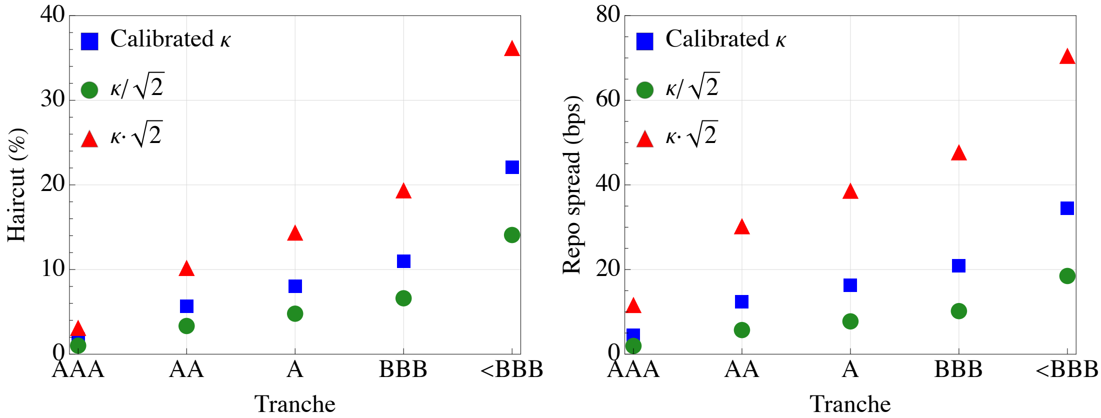

(a) Haircuts (b) Repo rates

FIGURE 7

Effect of information asymmetry on repo terms

Notes: The counterfactual exercises change the borrowers' information advantage by adjusting the parameter κ while keeping the overall variance $\sigma^2+\kappa^2$ fixed at the baseline level.

predictions by reinterpreting the lenders' beliefs as risk-neutral probabilities that place greater weights on worse states. This reinterpretation can potentially explain the gap between a low observed default rate and a higher risk-neutral default probability. Such an interpretation is particularly natural when aggregate risk is present, as the risk-neutral probabilities would reflect not only the exposure to aggregate events but also the aggregate discount factor. Overall, risk-averse behavior, combined with differences in valuation and exposure to aggregate risk, can substantially dampen the model's predicted default rates.

Counterfactuals. We can exploit the model to evaluate how the repo market responds to changes in the extent of informational asymmetries. To that end, we adjust the parameter κ to change the borrower's information advantage while maintaining the overall variance $\sigma^2+\kappa^2$ at the baseline level. We compare the baseline to the cases where we either halve or double the variance κ^2, which results in either halving or doubling the borrower's relative information precision. As Figure 7 shows, doubling the borrower's information advantage increases haircuts and rates for all tranches. However, the funding level for the AAA tranche is relatively unaffected, while the haircut for the <BBB tranche can reach 40%, substantially reducing the funding level.

The greater sensitivity of haircuts among lower-quality tranches raises questions about the desirability of securitization in light of our theory. In an additional exercise, we find that the overall level of funding is insensitive to the securitization and tranching structure. Through our model's lens, each tranche's funding level is tied to the $\frac{r}{1+r}$ quantile of the quality distribution, which corresponds to the same tail event for the overall mortgage pool. Hence, the overall funding is determined by the uncertainty and information asymmetry of the mortgage pool. Tranching merely redistributes funding across tranches and has a neutral effect on overall liquidity provision.

Discussion. Our model offers a private information perspective to understanding repos, contrasting with other theories based on heterogeneity, search, and default costs. While these alternative explanations complement ours, it is important to highlight evidence pointing to private information's relevance.

A positive relationship between haircuts and rates has been documented in multiple markets: Julliard et al. (2022) and Suzuki and Sasamoto (2024) report this relationship from transaction-level repo data in the UK and Japan respectively. Both studies confirm the greater sensitivity of haircuts. Second, consistent with the information-based approach, many repos have zero haircuts when borrowers and lenders belong to similar financial institutions. However, haircuts on the repos on the same securities rise when borrowers are hedge funds with presumably better information. Julliard et al. (2022) finds this pattern in the UK and Baklanova et al. (2019) in the US bilateral repo market. Third, Suzuki and Sasamoto (2024) and Auh and Landoni (2022) also report that repo haircuts are near zero even among long-term bonds with substantial price volatility, suggesting that the motive for haircuts is not necessarily risk-hedging. This contrasts with haircuts on corporate bonds and stocks, suggesting that haircuts are less driven by risk. Indeed, Julliard et al. (2022) conclude their study by stating, "We find evidence in favor of an adverse selection explanation of haircuts, but little evidence in support of lenders' liquidity position or default probabilities affecting haircuts," whereas Suzuki and Sasamoto (2024) open their study, stating that their empirical findings "are consistent with the theory that haircuts function as a screening mechanism to mitigate adverse selection." Additional suggestive evidence on the relevance of asymmetric information is that the same hedge funds often buy multiple tranches of the same MBS pool (as noted by Auh and Landoni, 2022).[31]

Finally, our model can help interpret the differences between the erosion of haircuts in the bilateral and tri-party repo markets during crises: Gorton and Metrick (2012) show that haircuts increased dramatically in the bilateral repo market during 2007-2008, reaching levels beyond 40% without corresponding increases in default rates. In contrast, Krishnamurthy et al. (2014) find that haircuts for tri-party repos backed by similar assets remained stable. This difference could result from the institutional differences in both markets described in Copeland et al. (2014). As is evident from the counterfactual exercise, haircuts will rise dramatically with risk when information is asymmetric but not when information friction is less pronounced, as in the tri-party segment.

7. CONCLUSION

This paper presents a model of repo-like financial contracts in markets subject to adverse selection. Our repo market equilibrium features a strikingly simple pooling repo contract characterized in closed form, which clarifies how underlying information asymmetries regarding collateral affect repo haircuts and rates and undermine market liquidity. Unlike popular environments, the closed forms are readily portable to other settings.

The theory only touches the surface of the broad questions in the introduction. The specific use of collateralized or repos typically depends on other economic or institutional forces outside the scope of this paper. Further research should explore market forces that induce the same assets to be traded in simultaneous markets whose trading protocols differ.

31. This evidence suggests economies of scale in information acquisition about underlying mortgage pools and subsequent trades across multiple tranches to capitalize on the informational advantage.

A. PROOFS

A.1. *Preliminary results*

We first establish several results that hold in general for MWS equilibrium, regardless of whether it is asset sales or repo markets.

We first show that in the solution to Problem 1, the contracts break even overall and involve cross-subsidization from high to low types.

Lemma 5 ***(Zero Profit)*** *In Problem 1, in the solution to a sub-problem for type θ, the contracts earn a zero overall profit, i.e., equation (ZPC) is binding.*

Proof. In the sub-problem for type θ, suppose the overall lender profit is strictly positive. Then consider a new set of contracts, which improve the sales price of the existing contracts by $\varepsilon/q(\theta)$ for some small $\varepsilon>0$, i.e., $\tilde{p}^s(\theta;\theta)=p^s(\theta;\theta)+\frac{\varepsilon}{q(\theta;\theta)}$. These contracts ensure that the (IC) and (RC) conditions are still satisfied. However, the adjustment increases the payoff of the terminal type θ by $(1+r)\varepsilon$, a contradiction to the existing contracts being a solution. ||

Lemma 6 ***(Cross-Subsidization)*** *In Problem 1, in the solution to a sub-problem for type θ, lender profits from the terminal quality, θ, are positive, i.e., $\pi(\boldsymbol{p}(\theta;\theta),q(\theta;\theta);\theta)\geq 0$.*

Proof. In the sub-problem for type θ, suppose lenders incur a strict loss from the terminal type, i.e., $\pi(\boldsymbol{p}(\theta;\theta),q(\theta;\theta);\theta)<0$. Then, lenders need to collectively generate strictly positive profits from a subset of previous types to break even. We show that if this is the case, we could have improved the reservation value of the preceding type.

Consider the sub-problem of the preceding type $\theta'<\theta$. In this problem, we propose a solution that modifies the contracts for types $\tilde{\theta}\in[\underline{\theta},\theta']$ from sub-problem θ by adjusting to a slightly higher sales price: $p_s(\tilde{\theta};\theta')=p_s(\tilde{\theta};\theta)+\varepsilon/q(\tilde{\theta};\theta)$. This proposed set of contracts would deliver payoffs to these types that is ε above their payoffs in sub-problem θ, i.e., $v(\tilde{\theta};\theta')=v(\tilde{\theta};\theta)+\varepsilon$, thus leaving their (IC) constraints undisturbed. Furthermore, this candidate set satisfies all constraints of sub-problem θ': 1) it meets the (IC) constraint; 2) since lenders make strictly positive profits from these types in sub-problem θ, they would still make a non-negative profit here for a sufficiently small value of ε, thereby satisfying the (ZPC) condition; 3) the (RC) constraint is also satisfied and, in fact, slack, with $v(\tilde{\theta};\theta')>v(\tilde{\theta};\theta)\geq v^\star(\tilde{\theta})$, $\forall\tilde{\theta}\leq\theta'$. Critically, this suggests that the payoff for the terminal type θ' exceeds its reservation value $v(\tilde{\theta};\theta')>v^\star(\theta')$, a contradiction to $v(\theta';\theta')=v^\star(\theta')$. Hence, in sub-problem θ, lender profits from the terminal type must be positive. ||

A.2. *Proof of Lemma 1*

Consider any two borrower types $\theta_1<\theta_2$. Denote their preferred contract by $\{\boldsymbol{p}_1,q_1\}$ and $\{\boldsymbol{p}_2,q_2\}$, where $\boldsymbol{p}_1\equiv\{p_{s1},p_{r1}\}$ and $\boldsymbol{p}_2\equiv\{p_{s2},p_{r2}\}$. The (IC) constraint can be specified as:

$$q_1[(1+r)p_{s1}-\min\{\theta_1,p_{r1}\}]\geq q_2[(1+r)p_{s2}-\min\{\theta_1,p_{r2}\}] \tag{IC-12}$$

$$q_2[(1+r)p_{s2}-\min\{\theta_2,p_{r2}\}]\geq q_1[(1+r)p_{s1}-\min\{\theta_2,p_{r1}\}] \tag{IC-21}$$

Monotone default. We show that the default decision is monotonic in types. Note that (IC-12) and (IC-21) imply that

$$q_1[\min\{\theta_2,p_{r1}\}-\min\{\theta_1,p_{r1}\}]\geq q_2[\min\{\theta_2,p_{r2}\}-\min\{\theta_1,p_{r2}\}] \tag{A.1}$$

Suppose type θ_2 strictly prefers to default. Then $p_{r2}>\theta_2$, which implies that the lower quality would default if switching contracts, $p_{r2}>\theta_1$. In this case, condition (A.1) becomes $q_1[\min\{\theta_2,p_{r1}\}-\min\{\theta_1,p_{r1}\}]\geq q_2(\theta_2-\theta_1)>0$. For it to hold, it must be $p_{r1}>\theta_1$. That is, type θ_1 also strictly prefers to default if θ_2 defaults. This further implies that the quantity traded is weakly decreasing among the default types: $q_1\geq q_2$.

Suppose type θ_1 (weakly) prefers not to default. Then $p_{r1}\leq\theta_1$, which implies that $p_{r1}<\theta_2$. Condition (A.1) becomes $q_2[\min\{\theta_2,p_{r2}\}-\min\{\theta_1,p_{r2}\}]\leq 0$. For it to hold, it must be $p_{r2}\leq\theta_2$. That is, type θ_2 weakly prefers not to default.

Decreasing equilibrium payoff $v(\theta)$. Next, we show that $v(\theta_1)\geq v(\theta_2)$. From (IC-12):

$$v(\theta_1)=q_1[(1+r)p_{s1}-\min\{\theta_1,p_{r1}\}]\geq q_2[(1+r)p_{s2}-\min\{\theta_1,p_{r2}\}]$$
$$\geq q_2[(1+r)p_{s2}-\min\{\theta_2,p_{r2}\}]=v(\theta_2).$$

If both types strictly prefer to default $\theta_1<\theta_2<\theta^d$, $p_{r2}>\theta_2>\theta_1$. The second inequality above is strict. Thus, $v(\theta_1)>v(\theta_2)$.

If both types weakly prefer not to default $\theta_2>\theta_1\geq\theta^d$, the repurchase price satisfy $p_{r1}\leq\theta_1<\theta_2$. From (IC-21):

$$v(\theta_2)=q_2[(1+r)p_{s2}-p_{r2}]\geq q_1[(1+r)p_{s1}-\min\{\theta_2,p_{r1}\}]=q_1[(1+r)p_{s1}-p_{r1}]=v(\theta_1).$$

Thus, the two non-default types obtain the same value, $v(\theta_1)=v(\theta_2)$.

A.3. *Proof of Lemma 2*

Consider two types $\theta_1<\theta_2$. In the MWS Program 1, assume we have the solution of the sub-problem for type θ_1 and turn to solve the sub-problem for type θ_2. A possible contract we can offer to type θ_2 is to take the contract for type θ_1 from the problem of θ and adjust the repurchase price to $p_r(\theta_2;\theta_2)=\min\{p_r(\theta_1;\theta_1),\theta_1\}$. This contract and the existing solution from the previous step satisfy the (IC) and (RC) constraints. Lemma 6 implies that in sub-problem θ_1, lenders profit from the terminal type θ_1. In sub-problem θ_2, they would also profit from θ_2, guaranteeing an overall positive profit. This proposed plan will allow type θ_2 to obtain at least the same payoff as type θ_1, but the solution must at least weakly dominate this plan for θ_2.

If types are discrete, considering two consecutive types, $\theta_1<\theta_2$ suffices. Therefore, $v^\star(\theta_2)\geq v^\star(\theta_1)$. If types are continuous, we can extend our reasoning above by an argument via contradiction. Suppose there is an interval $[\theta_1,\theta_2]$ such that $v^\star(\theta)$ is strictly decreasing. Then in the sub-problem θ_2, we can offer the same contract as above to all types in $(\theta_1,\theta_2]$. Again, the (IC) conditions are satisfied. In this case, we must consider the (RC) constraints of the types in this interval, which is automatically satisfied since by assumption $v^\star(\theta)$ is decreasing. However, it also implies that $v^\star(\theta_2)\geq v^\star(\theta_1)$, which leads to a contradiction.

A.4. *Proof of Lemma 3*

No quantity screening of non-default types. We show that types that do not default in a solution to the MWS program must trade their full amount. Suppose by contradiction that there exists a non-default type $\theta\in\left[\theta^d,\bar{\theta}\right]$ such that its traded quantity is $q<1$. Next, we show that we can find a deviation that respects the (IC) constraint but produces positive lender profits.

Let $\{\boldsymbol{p},q\}$ denote the contract for type θ. Consider a deviating contract, $\{\tilde{\boldsymbol{p}},\tilde{q}\}$, for it:

$$\tilde{q}=q+\varepsilon$$
$$\tilde{p}_s=\frac{q}{q+\varepsilon}p_s+\frac{\varepsilon}{q+\varepsilon}\frac{1}{1+r}\underline{\theta}$$
$$\tilde{p}_r=\frac{q}{q+\varepsilon}p_r+\frac{\varepsilon}{q+\varepsilon}\underline{\theta}.$$

The deviating contract preserves the payoff of type θ:

$$\tilde{q}[(1+r)\tilde{p}_s-\tilde{p}_r]=q[(1+r)p_s-p_r].$$

This deviating contract respects the (IC) constraint of all other types $\theta'\neq\theta$, which we establish by considering three potential cases:

(i) Consider $\theta'\geq p_r>\tilde{p}_r$. Then type θ' would not default with either $\{\boldsymbol{p},q\}$ and $\{\tilde{\boldsymbol{p}},\tilde{q}\}$. Thus, the two contracts deliver equal payoffs. Since the (IC) constraint of θ' mimicking θ holds under the original contract, it must also hold for the deviating contract.

(ii) Consider $\tilde{p}_r<\theta'<p_r$. Type θ' would default with $\{\boldsymbol{p},q\}$ but not with $\{\tilde{\boldsymbol{p}},\tilde{q}\}$. The new contract delivers a lower payoff: $\tilde{q}[(1+r)\tilde{p}_s-\tilde{p}_r]=q[(1+r)p_s-p_r]<q[(1+r)p_s-\theta']$.

(iii) Consider $\theta'\leq\tilde{p}_r<p_r$. Then type θ' would default with both $\{\boldsymbol{p},q\}$ and $\{\tilde{\boldsymbol{p}},\tilde{q}\}$. However, the deviating contract has a sufficiently low repurchase price and would deliver a lower default payoff: $\tilde{q}[(1+r)\tilde{p}_s-\theta']=q[(1+r)p_s-\theta']-\varepsilon(\theta'-\underline{\theta})\leq q[(1+r)p_s-\theta']$.

In all three cases, other types weakly prefer their original contracts $\{\boldsymbol{p}(\theta'),q(\theta')\}$ over the new contract $\{\tilde{\boldsymbol{p}},\tilde{q}\}$:

$$q(\theta')\left[(1+r)p_s(\theta')-\min\{\theta',p_r(\theta')\}\right]\geq\tilde{q}\left[(1+r)\tilde{p}_s-\min\{\theta',\tilde{p}_r\}\right].$$

Thus, introducing the deviating contract keeps all payoffs as the original set.

Next, we show that this contract makes a profit (in combination with the original contracts for all other types). This is the case since:

$$\tilde{q}(\tilde{p}_r-\tilde{p}_s)=q(p_r-p_s)+\frac{\varepsilon}{q+\varepsilon}\frac{r}{1+r}\underline{\theta}>q(p_r-p_s).$$

This implies that we can increase profits for any contract featuring quantity screening. In the case of discrete types, this is enough. In the case of a continuum, we can consider a deviating set of contracts aimed at non-default types that trade limited quantities. As long as the fraction of non-default types is non-zero, this deviation will increase profits.

One non-default contract. We further show that a single contract serves all non-default types. To show this, we first establish that lender profit must be equalized from all non-default types. Suppose this is not the case. Consider a scenario where lenders make a higher profit from contract $\{p_{s1},p_{r1}\}$ for type θ_1 than contract $\{p_{s2},p_{r2}\}$ for type θ_2, i.e., $p_{1r}-p_{s1}>p_{r2}-p_{s2}$. Given that all non-default types must obtain the same payoff, i.e., $(1+r)p_{s1}-p_{r1}=(1+r)p_{s2}-p_{r2}$, it follows that $p_{s1}>p_{s2}$ and $p_{r1}>p_{r2}$. However, this implies that lenders can improve their profits by dropping the lower profit contract $\{p_{s2},p_{r2}\}$ and offering the more profitable contract $\{p_{s1},p_{r1}\}$ to type θ_2 instead, provided that θ_2 would not default. This is indeed guaranteed because $p_{r1}\leq\theta_2$. Otherwise type θ_2 would have deviated to contract $\{p_{s1},p_{r1}\}$ and obtained a higher payoff by defaulting.

Since the two contracts earn the same profit and deliver the same non-default payoff, they must be the same contract. As in our main text, we let $\boldsymbol{p}^n$ represent this contract.

No quantity screening of default type. Next, we show that a unique contract is offered to default types, trading their entire quantity. In this step, we utilize the *dual* of the relaxed MWS problem presented in Corollary 1:

$$\max_{\{\boldsymbol{p}(\theta),q(\theta);\theta\in\Theta\}}\mathbb{E}[\pi(\boldsymbol{p}(\theta),q(\theta);\theta)]$$

subject to the (IC) constraint $\forall\theta\in\Theta$ and $v(\boldsymbol{p}(\bar{\theta}),q(\bar{\theta});\bar{\theta})>\bar{v}$ for some $\bar{v}$.

We first establish that the non-default contract must satisfy $p_r^n\geq\max_{\theta<\theta^d}\theta$. If not, meaning $p_r^n<\max_{\theta<\theta^d}\theta\leq\theta^d$, lenders can *increase* the repurchase price for non-default types p_r^n by ε, where ε is sufficiently small such that the non-default types remain not defaulting. Likewise, lenders can simultaneously *reduce* the sales price for default types by $p_s(\theta')$ by $\varepsilon/q(\theta')/(1+r)$, $\forall\theta'<\theta^d$. These adjustments reduce the payoffs of all types uniformly by ε and, therefore, the (IC) constraints are undisturbed. Yet, lender profit improves, which is a contradiction to the profit maximization of the dual program. In fact, it must be the case that the repurchase price $p_r=\theta^d$. Further, it follows that if any of the default types, $\theta<\theta^d$, were to contemplate switching to $\boldsymbol{p}^n$, they would default and obtain a payoff $(1+r)p_s^n-\theta$.

In the next step, we take the default threshold θ^d and the non-default contract $\boldsymbol{p}^n$ and focus on maximizing profits subject to the (IC) constraints for default types $[\underline{\theta},\theta^d)$. With continuum types, the (IC) constraints for default types can be formulated as: $\forall\theta<\theta^d$,

$$v(\theta)=v(\underline{\theta})-\int_{\underline{\theta}}^{\theta}q(\theta')d\theta' \tag{IC-D}$$

$$v(\theta)\geq(1+r)p_s^n-\theta. \tag{IC-D-ND}$$

$$q'(\theta)\leq 0. \tag{Monotonicity}$$

Here, the (IC-D) constraints are within default types, and the (IC-D-ND) constraint states that these types prefer their contract over the contract of non-default types.

Incorporating the (IC-D) constraint, we compute the lender profit from type θ as:

$$\pi(\boldsymbol{p}(\theta),q(\theta);\theta)=q(\theta)(\theta-p_s(\theta))=\frac{1}{1+r}(rq(\theta)\theta-v(\theta)) \tag{A.2}$$

$$=\frac{1}{1+r}\left(rq(\theta)\theta-v(\underline{\theta})+\int_{\underline{\theta}}^{\theta}q(\theta')d\theta'\right).$$

The profit maximization objective can be rewritten as:

$$\max_{\{p(\theta),q(\theta)\}}\int_{\underline{\theta}}^{\theta^d}\pi(\boldsymbol{p}(\theta),q(\theta);\theta)dF(\theta)=\max_{\{q(\theta)\},v(\underline{\theta})}\frac{1}{1+r}\int_{\underline{\theta}}^{\theta^d}\left(rq(\theta)\theta-v(\underline{\theta})+\int_{\underline{\theta}}^{\theta}q(\theta')d\theta'\right)dF(\theta).$$

Apply integration by parts:

$$\int_{\underline{\theta}}^{\theta^d}\int_{\underline{\theta}}^{\theta}q(\theta')d\theta' dF(\theta)=-\int_{\underline{\theta}}^{\theta^d}\int_{\underline{\theta}}^{\theta}q(\theta')d\theta' d\left(F\left(\theta^d\right)-F(\theta)\right)=\int_{\underline{\theta}}^{\theta^d}q(\theta)\frac{F\left(\theta^d\right)-F(\theta)}{f(\theta)}dF(\theta).$$

Thus, the profit maximization problem becomes:

$$\max_{\{q(\theta)\}}\int_{\underline{\theta}}^{\theta^d}q(\theta)\left(r\theta+\frac{F(\theta^d)-F(\theta)}{f(\theta)}\right)dF(\theta)-\min_{v(\underline{\theta})}v(\underline{\theta})F(\theta^d). \tag{A.3}$$

The solution features $q(\theta)=1$, $\forall\theta<\theta^d$. This implies pooling of all default types. Hence, there can be at most one contract among defaulters, represented by $\boldsymbol{p}^d$.[32]

A.5. *Proof of Proposition 1*

We can now restrict our attention to two contracts $\boldsymbol{p}^d$ and $\boldsymbol{p}^n$, and we proceed to show that they are in fact the same contract. Building on the results in Appendix A.4, the payoff for default types are $v(\theta)=(1+r)p_s^d-\theta$, $\forall\theta<\theta^d$. The (IC-D-ND) constraint becomes:

$$(1+r)p_s^d-\theta\geq(1+r)p_s^n-\theta. \tag{A.4}$$

Profit maximization implies delivering to the default types as low payoffs as possible, such that the constraint above must be binding. That is, the sales prices are equal: $p_s^d=p_s^n$. Since $p_r^n\geq\max_{\theta<\theta^d}\theta$, it is without loss of generality to set the same repurchase prices in both contracts $p_r^d=p_r^n$.[33] Therefore, the two contracts are the same: $\boldsymbol{p}^d=\boldsymbol{p}^n=\boldsymbol{p}$.

Now, we are only left to choose the threshold type θ^d and the contract $\boldsymbol{p}$. The sub-problem for the highest quality type would imply the contract must maximize the non-default payoff as stated in equation (3.6). From the previous step, profit maximization implies that $p_r=\theta^d$ and hence the lender zero-profit calculation in equation (3.7).

A.6. *Proof of Corollary 2*

Recall that the equilibrium repo contract maximizes the non-default payoff $(1+r)p_s-p_r$ subject to the zero-profit condition (3.7). With a continuum-type distribution, we first establish that this problem is strictly concave, which rests on the concavity of p_s with respect to p_r. The first-order derivative of p_s, given by (3.7), with respect to p_r is:

$$\frac{\partial p_s}{\partial p_r}=\frac{\partial}{\partial p_r}\left[\int_{\underline{\theta}}^{p_r}\theta dF(\theta)+p_r(1-F(p_r))\right]=1-F(p_r).$$

The second-order derivative is negative, $\frac{\partial^2 p_s}{\partial p_r^2}=-f(p_r)<0$, which guarantees concavity.

Substituting (3.7) into the non-default payoff and taking the first-order condition with respect to p_r:

$$(1+r)(1-F(p_r))=1,$$

which implies a default rate $d=F(p_r)=\frac{r}{1+r}$ and the repurchase price in equation (3.8). Using the zero-profit condition (3.7), we obtain the sales price in equation (3.9).

32. This result can be generalized to discrete types, when there are holes and mass points in the distribution.

33. One could, in principle, increase the repurchase price p_r^d for defaulters, resulting in two distinct contracts. However, such adjustments result in an equilibrium that is both payoff-equivalent and allocation-equivalent to the pooling contract.

A.7. *Proof of Proposition 2*

First, we observe that if we incorporate the zero-profit condition (ZPC), the objective of the planner is equivalent to maximizing aggregate liquidity, $\mathbb{E}[p_s(\theta)q(\theta)]$. Hence, the dual of the planner's problem is to maximize profits, subject to a level of liquidity provision, while respecting the (IC) constraint and the participation constraint $v(\boldsymbol{p}(\theta),q(\theta);\theta)\geq 0$.

The proof of Lemma 1 on monotone borrower payoff and monotone default decision follows from the (IC) constraint alone. Hence, these properties are also guaranteed in the planner's solution. Furthermore, the proof of Lemma 3, which guarantees no quantity screening, involves the dual problem of the relaxed MWS problem. The planner's problem has a similar structure, leading to no quantity screening. By the same token, Lemma 4 also holds in the planner's solution, i.e., the planner's solution also involves at most two contracts. As a result, the incentive-compatibility in equation (A.4) implies that the sales prices are equal: $p_s^d = p_s^n$. Thus, the planner's solution also features a pooling contract.

Return to the original planner's problem, we note that, according to the zero-profit condition (3.7), the funding level p_s is increasing in p_r. Thus, the solution is $p_r = \bar{\theta}$ if the participation of non-default types $(1+r)p_s - p_r \geq 0$ is not a problem. This solution is feasible if $(1+r)\mathbb{E}[\theta] \geq \bar{\theta}$. If this condition is violated, the participation of non-default types is binding. In this case, the solution pushes the non-default payoff to zero, $(1+r)\mathbb{E}[\min\{\theta,p_r\}] - p_r = 0$. Recall that the non-default payoff is increasing in p_r for $p_r < p_r^*$ and decreasing for $p_r > p_r^*$. Furthermore, $(1+r)\mathbb{E}[\min\{\theta,\underline{\theta}\}] - \underline{\theta} \geq 0$ and $(1+r)\mathbb{E}\left[\min\left\{\theta,\bar{\theta}\right\}\right] - \bar{\theta} = (1+r)\mathbb{E}[\theta] - \bar{\theta} < 0$. Therefore, there exists a unique solution to $(1+r)\mathbb{E}[\min\{\theta,p_r\}] - p_r = 0$, with $p_r^p > p_r^*$.

Acknowledgements. We would like to thank Andrew Atkeson, Andrea Attar, Jeremy Bejarano, Jack Bigio, Simon Board, Patrick Bolton, Briana Chang, Mohammad Davoodalhosseini, Andres Drenik, Andrea Eisfeldt, Piero Gottardi, Veronica Guerrieri, Barney Hartman-Glaser, Pablo Kurlat, Moritz Meyer-ter-Vehn, Ken Miyahara, Facundo Piguillem, Joao Ramos, Pablo Ottonello, Tano Santos, Joseph Stiglitz, Andre Sztutman, Victoria Vanasco, Liangjie Wu, and William Zame for useful conversations. We are especially thankful to Florian Scheuer, whose sharp questions helped us clarify important features of this paper. Bigio thanks the hospitality of the Federal Reserve Bank of San Francisco. Also, we would like to thank seminar participants at EIEF, Stanford, UCLA, Shanghai University of Finance and Economics, World Congress of the Econometric Society, Tor Vergata Workshop: Markets, Information and the Macroeconomy, Rice University, University of Hong Kong, Essex-UCL Workshop on Decentralized Financial Markets, and Office of Financial Research.

Data Availability Statement. The code used to conduct the numerical analysis in this article is available on Zenodo at https://doi.org/10.5281/zenodo.16539136.

REFERENCES

Agarwal, S., Y. Chang, and A. Yavas (2012). Adverse selection in mortgage securitization. *Journal of Financial Economics 105*(3), 640–660.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*(3), 488–500.

An, X., Y. Deng, and S. A. Gabriel (2011). Asymmetric information, adverse selection, and the pricing of cmbs. *Journal of Financial Economics 100*(2), 304–325.

Ashcraft, A. B. and T. Schuermann (2008). Understanding the securitization of subprime mortgage credit. FRB of New York Staff Report, No. 318.

Asriyan, V. and V. Vanasco (2023, 04). Security design in non-exclusive markets with asymmetric information. *The Review of Economic Studies*, rdad038.

Attar, A., T. Mariotti, and F. Salanié (2011). Nonexclusive competition in the market for lemons. *Econometrica 79*(6), 1869–1918.

Auh, J. K. and M. Landoni (2022). Loan terms and collateral: Evidence from the bilateral repo market. *The Journal of Finance 77*(6), 2997–3036.

Baklanova, V., C. Caglio, M. Cipriani, and A. Copeland (2019). The use of collateral in bilateral repurchase and securities lending agreements. *Review of Economic Dynamics 33*, 228–249. Fragmented Financial Markets.

Becker, J., K. Hendricks, J.-F. Houde, , and D. Raisingh (2023). Asymmetric information and the supply-chain of mortgages: The case of ginnie mae loans. University of Wisconsin Working Paper.

Besanko, D. and A. V. Thakor (1987). Competitive equilibrium in the credit market under asymmetric information. *Journal of Economic Theory 42*(1), 167 – 182.

Biais, B. and T. Mariotti (2005). Strategic liquidity supply and security design. *The Review of Economic Studies 72*(3), 615–649.

Bigio, S. (2015, June). Endogenous liquidity and the business cycle. *American Economic Review 105*(6), 1883–1927.

Chang, B. (2017, 04). Adverse Selection and Liquidity Distortion. *The Review of Economic Studies 85*(1), 275–306.

Copeland, A., A. Martin, and M. Walker (2014). Repo runs: Evidence from the tri-party repo market. *The Journal of Finance 69*(6), 2343–2380.

Daley, B. and B. Green (2012). Waiting for news in the market for lemons. *Econometrica 80*(4), 1433–1504.

Dang, T. V., G. Gorton, and B. Holmström (2013). Haircuts and repo chains. *Working Paper*.

Demarzo, P. and D. Duffie (1999). A liquidity-based model of security design. *Econometrica 67*(1), 65–99.

Duffie, D. (1996). Special repo rates. *The Journal of Finance 51*(2), 493–526.

Eisfeldt, A. L. (2004). Endogenous liquidity in asset markets. *The Journal of Finance 59*(1), 1–30.

Engers, M. and L. Fernandez (1987). Market equilibrium with hidden knowledge and self-selection. *Econometrica 55*(2), 425–439.

Fostel, A. and J. Geanakoplos (2015). Leverage and default in binomial economies: A complete characterization. *Econometrica 83*(6), 2191–2229.

Fuchs, W. and A. Skrzypacz (2019). Costs and benefits of dynamic trading in a lemons market. *Review of Economic Dynamics 33*, 105 – 127. Fragmented Financial Markets.

Gemmo, I., C. Kubitza, and C. Rothschild (2020). Constrained efficient equilibria in selection markets with continuous types. *Journal of Public Economics 190*, 104237.

Glaeser, E. L. and H. D. Kallal (1997). Thin markets, asymmetric information, and mortgage-backed securities. *Journal of Financial Intermediation 6*(1), 64–86.

Gorton, G., T. Laarits, and A. Metrick (2020). The run on repo and the fed's response. *Journal of Financial Stability 48*, 100744.

Gorton, G. and A. Metrick (2012). Securitized banking and the run on repo. *Journal of Financial Economics 104*(3), 425 – 451.

Gorton, G. and G. Ordoñez (2014). Collateral crises. *American Economic Review 104*(2), 343–78.

Gorton, G. B. and G. Pennacchi (1990). Financial intermediaries and liquidity creation. *Journal of Finance 45*(1), 49–71.

Gottardi, P., V. Maurin, and C. Monnet (2019). A theory of repurchase agreements, collateral re-use, and repo intermediation. *Review of Economic Dynamics*.

Guerrieri, V. and R. Shimer (2014). Dynamic adverse selection: A theory of illiquidity, fire sales, and flight to quality. *American Economic Review 104*(7), 1875–1908.

Guerrieri, V. and R. Shimer (2018, May). Markets with multidimensional private information. *American Economic Journal: Microeconomics 10*(2), 250–74.

Guerrieri, V., R. Shimer, and R. Wright (2010). Adverse selection in competitive search equilibrium. *Econometrica 78*(6), 1823–1862.

Hart, O. and J. Moore (1994). A theory of debt based on the inalienability of human capital. *The Quarterly Journal of Economics 109*(4), 841–879.

Hendel, I. and A. Lizzeri (1999). Adverse selection in durable goods markets. *The American Economic Review 89*(5), 1097–1115.

Hendel, I. and A. Lizzeri (2002). The role of leasing under adverse selection. *Journal of Political Economy 110*(1), 113–143.

Holmstrom, B. (2014). Understanding the role of debt in the financial system. BIS Working Papers No. 479.

Holmstrom, B. and R. B. Myerson (1983). Efficient and durable decision rules with incomplete information. *Econometrica 51*(6), 1799–1819.

Holmstrom, B. and J. Tirole (1997). Financial intermediation, loanable funds, and the real sector. *The Quarterly Journal of Economics 112*(3), 663–691.

Julliard, C., G. Pinter, K. Todorov, and K. Yuan (2022). What drives repo haircuts? evidence from the uk market. BIS Working Paper 1027.

Krishnamurthy, A., S. Nagel, and D. Orlov (2014). Sizing up repo. *The Journal of Finance 69*(6), 2381–2417.

Kurlat, P. (2016). Asset markets with heterogeneous information. *Econometrica 84*(1), 33–85.

Kurlat, P. and F. Scheuer (2020, 10). Signalling to experts. *The Review of Economic Studies 88*(2), 800–850.

Lester, B., A. Shourideh, V. Venkateswaran, and A. Zetlin-Jones (2019). Screening and adverse selection in frictional markets. *Journal of Political Economy 127*(1), 338–377.

Madison, F. (2024). Asymmetric information in frictional markets for liquidity: Collateralized credit vs asset sale. *Journal of Economic Dynamics and Control 159*, 104808.

Miyazaki, H. (1977). The rat race and internal labor markets. *The Bell Journal of Economics 8*(2), 394–418.

Monnet, C. and B. Narajabad (2017). Why rent when you can buy? *FEDS Working Paper No. 2017-094*.

Myers, S. C. and N. S. Majluf (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics 13*(2), 187–221.

Netzer, N. and F. Scheuer (2014). A game theoretic foundation of competitive equilibria with adverse selection. *International Economic Review 55*(2), 399–422.

Ozdenoren, E., K. Yuan, and S. Zhang (2023, 03). Dynamic asset-backed security design. *The Review of Economic Studies*, rdad022.

Parlatore, C. (2019). Collateralizing liquidity. *Journal of Financial Economics*.

Philippon, T. and V. Skreta (2012, February). Optimal interventions in markets with adverse selection. *American Economic Review 102*(1), 1–28.

Riley, J. G. (1979). Informational equilibrium. *Econometrica 47*(2), 331–359.

Rothschild, M. and J. Stiglitz (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics 90*(4), 629–649.

Spence, M. (1978). Product differentiation and performance in insurance markets. *Journal of Public Economics 10*(3), 427 – 447.

Stantcheva, S. (2014, 03). Optimal income taxation with adverse selection in the labour market. *The Review of Economic Studies 81*(3), 1296–1329.

Stiglitz, J. E. and A. Weiss (1981). Credit rationing in markets with imperfect information. *The American Economic Review 71*(3), 393–410.

Suzuki, K. and K. Sasamoto (2024). Quantitative analysis of haircuts: Evidence from the japanese repo and securities lending markets. SSRN 4390755.

Tirole, J. (2012). Overcoming adverse selection: How public intervention can restore market functioning. *The American Economic Review 102*(1), 29–59.

Townsend, R. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory 21*, 265–293.

Wilson, C. (1977). A model of insurance markets with incomplete information. *Journal of Economic Theory 16*(2), 167 – 207.