When Better Looks Hide Worse Outcomes: Willingness to Pay and Underwriting in Corporate Credit

Virgilio Zurita*

January 27, 2026

Abstract

Corporate credit often looks safest just before defaults surge. We build a dynamic model in which firms differ in willingness to pay (W2P) and ability to pay (A2P), while lenders underwrite using noisy signals about W2P. Paradoxically, higher cross-sectional volatility in firms' W2P lowers credit spreads by 13 basis points because underwriting improves the composition of the approved pool, making borrowers look safer to banks, even as population default risk rises. In contrast, A2P shocks widen spreads and move the credit risk premium in a conventional way. This financial information channel produces nonlinear dynamics and episodes in which spreads compress at issuance yet are followed by high default outcomes, a pattern that is difficult to reconcile with shocks to firms' fundamentals.

JEL Classification: G12, G14, G21, G33

Keywords: credit spreads, financial frictions, information frictions, banks, financial intermediaries, default risk

*College of Business and Analytics, Southern Illinois University, virgilio.zurita@siu.edu.

1 Introduction

Corporate credit often looks safest right before it performs worst. Spreads compress, issuance booms, and investors treat credit as nearly riskless, yet a few years later defaults surge and realized losses far exceed the premia priced at issuance. Headlines about a debt-financed AI bubble sharpen that concern, while recent liability-management tactics have kept firms technically current even as underlying credit quality deteriorates, blurring which claims will ultimately be made whole.[1] When lenders rely on imperfect signals to screen borrowers, the loans that get funded can look safer even as risk rises in the broader borrower population. These episodes raise an important question: how can the market's required compensation for bearing risk be so low precisely when realized default losses are about to be so high?

The traditional asset-pricing view attributes these patterns to fundamentals—firms' *ability to pay*, the capacity to repay borrowed funds and meet contractual obligations. Investors price expected default losses plus a risk premium for exposure to aggregate risks, with cash-flow volatility, leverage, and recoveries determining a fundamentals-based default boundary. Yet a well-known gap remains: for corporate bonds, observed spreads are too large to be explained by expected default alone, leaving a sizable residual (Huang, Nozawa and Shi (2025)). This puzzling discrepancy motivates us to examine forces beyond capacity—those that govern whether, when, and on what terms obligations are ultimately honored. Do credit spreads reflect such forces even when conventional capacity measures move little?

In this paper we argue that a key missing state variable in this discussion is firms' *willingness to pay* (W2P), distinct from their *ability to pay* (A2P). W2P captures how much firms value remaining current on their debt, conditional on their fundamentals—reputational concerns, franchise value, implicit and explicit guarantees, and the perceived costs of distress and default. This notion is in the spirit of strategic default and the "cannot pay versus will not pay" distinction documented for households and mortgages by Gerardi, Herkenhoff,

[1] "Loan defaults are looking worse below the surface," Financial Times, June 4, 2025.

Ohanian and Willen (2018) and Ganong and Noel (2023), and for corporate and cross-country settings by Davydenko and Strebulaev (2007) and Favara, Schroth and Valta (2012). A2P captures the usual fundamentals—cash flows, leverage, and exposure to aggregate shocks. We show that taking W2P into account, and allowing both W2P and A2P to vary over time and across firms, delivers a powerful information-based channel for credit spreads and expected returns.

Our starting point is the observation that modern credit markets are information-intensive. Lenders do not simply look at a firm's balance sheet or income statement; they form beliefs about how the firm will behave in different states of the world, using a mix of hard information, soft information, and patterns in past repayment behavior (see, for example, Greenwood, Hanson and Jin (2023), Gomes, Grotteria and Wachter (2019) and Greenwood, Hanson, Shleifer and Sorensen (2022)). We model this by allowing lenders to observe noisy signals about firms' W2P, with the quality of information itself varying over time. Time-varying cross-sectional volatility in W2P and A2P summarizes how heterogeneous borrowers are along these two dimensions (e.g., see Herskovic, Kind and Kung (2023) and Herskovic, Kelly, Lustig and Van Nieuwerburgh (2016)). The quality of financial information summarizes how precisely lenders can infer W2P from the signals they see. Heterogeneity in firms' W2P has a dual role: it increases the share of borrowers with very weak repayment incentives, and it strengthens the role of underwriting. Together, W2P, A2P, and financial information quality determine how credit risk is allocated and priced.

We embed W2P and A2P into a dynamic framework with a representative investor and a stochastic discount factor (SDF) that prices all cash flows. A continuum of firms raises external funding through nominally identical debt contracts, but they differ in their fundamentals (A2P) and in W2P. Lenders form posterior beliefs about each firm's W2P and A2P given the available information and choose whether to extend credit and at what spread. In this framework, the key model-implied measures are the risk-neutral and physical default probabilities that emerge from these credit-market information frictions, and how they

co-move with the SDF.

Our quantitative analysis highlights a novel and seemingly counterintuitive result: greater uncertainty in firms' willingness to pay *tightens* credit spreads on impact. By widening the gap between high- and low-incentive borrowers, higher dispersion increases the gains from borrower selection. Because the noise in lenders' signals is unchanged, greater dispersion makes those signals more informative, so lenders tighten underwriting and approvals shift toward safer borrowers, improving the composition of the origination pool. Specifically, a one-standard-deviation increase in the cross-sectional dispersion of W2P reduces origination spreads by about 13 basis points and lowers both physical and risk-neutral default probabilities of newly originated loans. This spread compression reflects a composition effect: new originations tilt toward borrowers with stronger repayment incentives. At the same time, unconditional (population) default risk rises because greater dispersion puts more mass in the low-W2P tail, even as the approved pool becomes safer. By contrast, a comparable increase in the dispersion of A2P generates the conventional pattern: credit spreads *widen* on impact—by around 9 basis points in our baseline specification—and higher spreads go hand in hand with higher expected default losses.

The key distinction is how these shocks load on the stochastic discount factor. Ability-to-pay shocks primarily work through a familiar risk-based channel: they move credit spreads by changing the sensitivity of default outcomes to aggregate risk, and they show up as changes in the credit risk premium. Willingness-to-pay shocks, by contrast, operate mainly through a risk-neutral channel—that is, through shifts in state-price-weighted default probabilities rather than through the credit risk premium. They reallocate probability mass across firms and states where lenders' beliefs about W2P are most fragile, changing risk-neutral default probabilities in a way that compresses spreads initially, while default risk rises in the broader borrower population. Intuitively, these default losses are not especially concentrated in bad aggregate states where the SDF is high, so the credit risk premium moves only modestly. In our simulations, W2P shocks generate large, persistent movements in credit spreads with

modest movements in the credit risk premium, whereas A2P shocks move the credit risk premium strongly and produce more standard spread dynamics.

Quantitatively, the model is disciplined by three sets of targets. First, we choose preferences and aggregate endowment dynamics so that the representative investor's SDF is consistent with standard asset-pricing benchmarks. Second, we parameterize the long-run dispersion and persistence of A2P and W2P to match empirical moments of default probabilities and cross-sectional heterogeneity in observable firm fundamentals. Third, we specify the information structure by matching average loan terms and the dynamics of the signal-to-noise ratio for W2P. With these ingredients in place, the model delivers average credit spreads of about 120 basis points, default rates around 30 basis points, and a credit risk premium close to 100 basis points, all broadly consistent with the data, while also generating episodes in which spreads are persistently high relative to expected default losses.

Our contribution is threefold. First, we introduce a framework in which willingness to pay (W2P) and financial information quality are first-order state variables for credit risk, alongside traditional fundamentals. A central implication is that time-varying cross-sectional volatility in W2P can move spreads down by improving the composition of the approved pool under underwriting. Second, this information-based channel delivers unconventional but empirically relevant dynamics and can generate episodes in which unusually low credit spreads are followed by elevated default outcomes: in our model, such patterns arise naturally in high-dispersion W2P states, but they cannot be generated by traditional fundamentals-driven high-dispersion A2P states. Third, we show that a moment-disciplined parameterization of W2P dispersion and information frictions simultaneously matches key features of credit spreads, default rates, and credit risk premia while delivering a clean separation between W2P forces and conventional risk premia. W2P shocks compress spreads, whereas A2P shocks primarily move the discount-rate component. While variance decompositions often attribute a substantial share of spread movements to discount-rate components (e.g., Nozawa (2017)), we show that shocks to the cross-sectional volatility of firms' willing-

ness to pay generate large and persistent spread dynamics through a default-expectations channel that is largely orthogonal to the standard risk-premium channel. Together, these results establish shocks to W2P dispersion as a distinct driver of both the level and the dynamics of credit spreads. By shifting underwriting-based selection and beliefs, they compress spreads even as aggregate default risk is rising, providing a unified lens on why credit can be priced tightly ahead of default upswings.

The structure of the paper is as follows. Section 2 discusses related literature on credit spreads, default risk, strategic default, and financial and information frictions. Section 3 discusses W2P in a static framework and highlights its key comparative statics. Section 4 embeds W2P and A2P into the dynamic model. Section 5 presents the quantitative implications for credit spreads, default, and risk premia. Section 6 concludes.

2 Related Literature

Our analysis relates to, and complements, several strands of the literature on credit risk and financial frictions that operate through information in fixed income markets. A large body of structural work emphasizes fundamentals and risk premia as the primary drivers of credit spreads; our contribution is to add a willingness-to-pay (W2P) component that operates through lenders' information and beliefs, and to study how time variation in information quality and cross-sectional dispersion in W2P jointly shape risk-neutral default probabilities, spreads, and expected returns. At the same time, a growing empirical literature documents that lenders' information sets, beliefs, and relationships can have large and persistent effects on credit pricing even holding conventional fundamentals fixed. Our framework nests these insights within a unified asset-pricing model in which W2P, ability to pay (A2P), and information quality jointly determine credit outcomes.

First, we connect to work emphasizing information frictions and belief-driven credit cycles in which credit spreads can become disconnected from contemporaneous fundamentals.

Greenwood, Hanson and Jin (2023) develop a model of reflexivity in credit markets in which extrapolative beliefs and refinancing needs generate periods of unusually low spreads followed by sharp deteriorations in credit quality. Greenwood, Hanson, Shleifer and Sørensen (2022) show that rapid joint booms in credit and asset prices forecast large increases in the probability of a financial crisis, while Gomes, Grotteria and Wachter (2019) document that cyclical dispersion in expected defaults predicts both real activity and credit returns. Related empirical work emphasizes how informational and renegotiation frictions shape corporate credit spreads and distressed-debt pricing (see, for instance, Darmouni (2020), Campello, Ladika and Matta (2019) and Berlin, Nini and Yu (2020)). Albagli, Hellwig and Tsyvinski (2024) show that with dispersed information and asymmetric payoffs, market prices can place disproportionate weight on tail outcomes. Our framework is complementary but different: rather than focusing on dispersion in firm fundamentals or on investors' beliefs per se, we treat cross-sectional dispersion in W2P and the quality of lenders' information about W2P as state variables that directly enter the pricing of defaultable cash flows.

Second, our focus on willingness to pay connects to the literature on strategic default and the "cannot pay" versus "will not pay" distinction (Guiso, Sapienza and Zingales (2013)). Gerardi, Herkenhoff, Ohanian and Willen (2018) distinguish borrowers who cannot pay from those who will not pay and show that this distinction is essential for understanding default behavior in mortgage markets. Ganong and Noel (2023) provide an empirical decomposition of mortgage default into cash-flow and strategic components using detailed micro data. Together with evidence from corporate and cross-country settings that strategic actions and default options affect credit outcomes (Davydenko and Strebulaev (2007); Favara, Schroth and Valta (2012)), these papers establish that willingness to pay matters empirically for default behavior. Our approach differs in emphasizing asset pricing: we build W2P directly into equilibrium pricing of credit, so that heterogeneity in W2P, combined with imperfect information on the lender side, moves risk-neutral default probabilities and credit spreads even when traditional measures of fundamental risk move little.

Finally, we contribute to the literature on default risk and the credit spread puzzle (e.g., see Huang and Huang (2012)). Bai, Goldstein and Yang (2020) argue that even state-of-the-art structural models struggle to explain the level and dynamics of investment-grade credit spreads, while Feldhütter and Schaefer (2018) show that standard default-based models leave a sizable non-default component in spreads. More recently, Huang, Nozawa and Shi (2025) document a global credit spread puzzle across countries and currencies, strengthening the case that something beyond conventional ability-to-pay risk is at work. Related work emphasizes time-varying volatility, ambiguity, intermediary capital constraints, and priced risk factors in corporate credit markets (Du, Elkamhi and Ericsson (2019), Chung, Wang and Wu (2019), Dickerson, Mueller and Robotti (2023), Augustin and Izhakian (2020), Doshi, Ericsson, Jacobs and Turnbull (2013) and Siriwardane (2019)). Our W2P-based mechanism adds an information-driven component to this debate: shocks to cross-sectional volatility in willingness to pay, together with changes in information quality, can generate large and persistent movements in credit spreads and risk-neutral default probabilities even when the conventional risk-premium channel is comparatively stable. This perspective helps organize episodes in which low spreads are followed by unusually high defaults and suggests asset-pricing tests that separate ability-to-pay risk from willingness-to-pay risk.

3 Willingness to Pay and Informational Frictions

In this section, we discuss the economic motivation for distinguishing between *ability to pay* (A2P) and *willingness to pay* (W2P) in corporate credit markets, explaining how informational frictions make W2P particularly difficult to observe and price. We introduce a one-period partial-equilibrium benchmark that illustrates how cross-sectional volatility in W2P compresses credit spreads at origination, even as it raises tail default risk and ex post excess returns, providing the foundation for the dynamic model in Section 4.

3.1 Ability to pay versus willingness to pay

In classical models of credit risk (e.g., Merton (1974); Duffie and Singleton (1999)), the central object is the borrower's *ability to pay* as promised—to repay the borrowed funds and honor the debt contract.[2] We use A2P as a compact label for the firm-specific determinants of cash-flow capacity and economic fundamentals: profitability, leverage, asset volatility, investment opportunities, and the value of collateral (Benmelech, Kumar and Rajan (2022); Ioannidou, Pavanini and Peng (2022)). A firm with high A2P has a strong balance sheet and robust cash flows; a firm with low A2P is financially fragile. Most traditional measures of credit quality—such as interest coverage ratios, leverage ratios, and ratings—are designed to capture A2P. This A2P component of credit risk has several observable proxies. Financial statements provide information about leverage, interest coverage, profitability, and cash-flow volatility. Market data provide information about equity volatility, asset betas, and market-implied default risk (Cao, Goyal, Xiao and Zhan (2023)). Credit ratings summarize a large amount of hard information about fundamentals. It is therefore natural to assume that lenders observe A2P reasonably well, even if they do not see every detail of the firm's operations.

In contrast, we define *willingness to pay* (W2P) as the component of default behavior that is not mechanically determined by cash flows and contractual obligations. W2P reflects the firm's propensity to honor its debt obligations when repayment is costly but feasible. It is shaped by non-pecuniary costs of default, reputational concerns, governance structures, legal and institutional constraints, and the preferences of managers and controlling shareholders.[3] Two firms with identical A2P can make very different default choices: one may work hard to avoid default and protect a reputation for honoring contracts, while the other may opportunistically restructure or walk away if it is legally possible and strategically

[2]Bakshi, Gao and Zhong (2022) provide a review of recent modelling approaches for default risk linked to firms' fundamentals.

[3]Survey evidence on mortgage strategic default similarly finds large non-pecuniary components and a role for enforcement beliefs in willingness-to-default decisions (Guiso, Sapienza and Zingales (2013)).

convenient.

Canonical structural and reduced-form models that focus solely on A2P implicitly assume that default occurs when the firm cannot repay, rather than when it chooses not to repay. In practice, observed restructurings, covenant violations, distressed exchanges, and opportunistic defaults suggest that both dimensions matter.

W2P is rooted in *soft information* and institutional features that are difficult to summarize in a small set of hard variables (Karolyi (2018)).[4] Some aspects of W2P can be indirectly proxied by variables such as governance indices, ownership structure, jurisdiction, and covenant strength. However, these proxies are noisy and incomplete, and building a comprehensive empirical measure of W2P is a substantial undertaking. The key assumption in this paper is that W2P is imperfectly observed by lenders. They can form views about it using a combination of soft and hard information, but these views are noisy and subject to error (e.g., Cici, Gibson and Merrick (2011); Eisfeldt, Herskovic and Liu (2024); Fuster, Plosser, Schnabl and Vickery (2019)). Consistent with this imperfect-observability assumption, He, Huang and Zhou (2023) show that changes in borrower data access and information sharing can reshape screening and credit-market competition. This assumption naturally leads to an information structure with a noisy signal about W2P and to underwriting rules based on that signal.

3.2 A static model of W2P dispersion and spread compression

To build intuition for the role of firms' W2P, and in particular for the cross-sectional volatility of firms' W2P, it is useful to consider a simple one-period example in partial equilibrium. Time is indexed by $t = 0$ (origination) and $t = 1$ (repayment or default). In this setup, all

[4]Examples include the firm's culture and governance, whether management and the board place a high value on maintaining a reputation for paying creditors in full, or whether they view default as an acceptable strategic tool. The legal and regulatory environment in which the firm operates, the strength of creditor rights, the ease of using bankruptcy strategically, and the speed and predictability of court processes. The history of the firm and its controlling shareholders, past restructurings, renegotiations, or aggressive uses of covenants and legal loopholes.

firms have the same ability to pay a, and we focus entirely on willingness to pay w. There are two W2P types: high-W2P firms (type H) with W2P index w_H and default probability PD_H; and low-W2P firms (type L) with W2P index w_L and default probability PD_L, where $\text{PD}_L > \text{PD}_H$.

A competitive financial intermediary (hereafter, the bank or the lender) extends loans to firms. For simplicity, we assume that lenders know the default probabilities PD_H and PD_L associated with each type but do not observe firms' types directly. Instead, they observe noisy signals and accept a pool of borrowers. We summarize pool composition with $\pi_w \in (0, 1)$, the share of firms with high willingness to pay in the approved pool. Therefore, among accepted borrowers, a fraction π_w are type H and a fraction $1 - \pi_w$ are type L. We interpret π_w as a reduced-form measure of W2P composition in the accepted pool and view changes in π_w as reflecting underlying changes in W2P dispersion and the underwriting environment. Let $\text{PD}_{\text{pool}}(\pi_w)$ denote the default probability in the accepted pool

$$\text{PD}_{\text{pool}}(\pi_w) = (1 - \pi_w)\text{PD}_L + \pi_w \text{PD}_H \tag{1}$$

which is decreasing in the share of high-W2P firms π_w. Note that, by construction, $\text{PD}_{\text{pool}}(\pi_w)$ is decreasing in π_w because replacing low-W2P firms with high-W2P firms lowers the average default probability. Figure 1 illustrates how changes in W2P dispersion and underwriting translate into changes in pool composition, holding the underwriting rule fixed. In the two-type version of the model, the accepted pool consists of a fraction π_w of high-W2P firms with default probability PD_H and a fraction $1 - \pi_w$ of low-W2P firms with default probability PD_L. The pool default probability $\text{PD}_{\text{pool}}(\pi_w)$ in equation (1) averages these two default intensities and is therefore decreasing in π_w.

[Insert Figure 1 Here]

Lenders are risk-neutral in this model and fund themselves at a risk-free gross rate R^f. They offer a single promised gross loan rate R_L to all accepted borrowers and receive a gross

recovery $\delta \in [0, R_L)$ per unit of face value if a borrower defaults at $t = 1$. The expected payoff per unit of face value is

$$\mathbb{E}_0\big[R_{L,1}\big] = \big(1 - \mathrm{PD}_{\mathrm{pool}}(\pi_w)\big)R_L + \mathrm{PD}_{\mathrm{pool}}(\pi_w)\delta. \tag{2}$$

Competition among lenders implies a zero-profit condition:

$$R^f = \mathbb{E}_0\big[R_{L,1}\big] = \big(1 - \mathrm{PD}_{\mathrm{pool}}(\pi_w)\big)R_L + \mathrm{PD}_{\mathrm{pool}}(\pi_w)\delta. \tag{3}$$

Solving (3) for the promised loan rate R_L as a function of π_w yields

$$R_L(\pi_w) = \frac{R^f - \mathrm{PD}_{\mathrm{pool}}(\pi_w)\delta}{1 - \mathrm{PD}_{\mathrm{pool}}(\pi_w)}. \tag{4}$$

The credit spread CS in this setup is defined as

$$\mathrm{CS}(\pi_w) = R_L(\pi_w) - R^f. \tag{5}$$

Because $\mathrm{PD}_{\mathrm{pool}}(\pi_w)$ is decreasing in π_w and appears in both the numerator and denominator of $R_L(\pi_w)$, a simple derivative calculation shows that $R_L(\pi_w)$ and hence $\mathrm{CS}(\pi_w)$ are also decreasing in π_w for reasonable parameter values. Intuitively, as π_w increases, the accepted pool contains a larger fraction of very high-W2P firms with extremely low default probabilities, which pulls down the average default probability and leads lenders to charge lower promised rates.

In this model, the pool default probability $\mathrm{PD}_{\mathrm{pool}}(\pi_w)$ in equation (1) averages the default intensities of high- and low-W2P firms and is therefore decreasing in π_w. Given risk-neutral lenders and a fixed risk-free rate R^f, the zero-profit condition (3) then implies that the promised loan rate $R_L(\pi_w)$ and the credit spread $\mathrm{CS}(\pi_w)$ inherit this negative sensitivity to π_w. We deliberately abstract from risk premia: with risk-neutral lenders and a fixed R^f,

the expected excess return on the credit portfolio is zero. The purpose of this example is to isolate the composition effect of W2P: even in a static partial-equilibrium environment with zero risk premia, changes in W2P dispersion π_w and underwriting can lower observed credit spreads at origination. In Section 4, the dynamic analogue of π_w is a time-varying cross-sectional dispersion of firms' W2P. Underwriting then translates higher dispersion into a higher share of high-W2P firms in the approved pool. Changes in the time-varying cross-sectional dispersion of W2P will also affect the covariance between credit payoffs and the stochastic discount factor, generating nontrivial risk premia in addition to the mechanical composition effect highlighted in this setup.

3.3 Ability and willingness: conflicting forces on spreads

The analysis above focuses solely on W2P, holding A2P fixed. To understand how A2P and W2P interact, it is helpful to juxtapose the effects of shocks to the dispersion (cross-sectional volatility) of ability and to the dispersion of willingness.

Consider two comparative statics experiments. First, suppose an increase in the dispersion of A2P across firms, holding W2P fixed. More firms are now close to the default boundary in terms of fundamentals. Even if lenders observe A2P well, the cross-section of firms becomes riskier in the usual sense: small shocks to cash flows or macro conditions can push more firms into default. Conditional on any given underwriting rule, the average default probability in the accepted pool increases. Lenders respond by raising promised loan rates, so that credit spreads widen. This is the conventional channel: higher fundamentals dispersion leads to higher spreads. Second, suppose an increase in the dispersion of W2P across firms, holding A2P fixed. The distribution of W2P now has more mass in the tails: more very high-W2P firms that almost never default, and more very low-W2P firms that default aggressively when it is convenient. When lenders only observe a noisy signal about W2P, and when the underwriting rule is based on this signal, the average default probability in the accepted pool is influenced by both tails. Very high-W2P firms pull down the average

default probability, especially in states where signals are favorable, while very low-W2P firms *raise* tail default risk, especially in adverse states that push them to default.

As a result, a positive shock to W2P dispersion can have two opposing effects. On impact, conditional on the signal and the underwriting rule, the average default probability in the accepted pool may fall, leading to lower promised loan rates and lower observed spreads. Over time, the presence of more very low-W2P firms in the accepted pool raises realized default rates and increases the correlation between credit payoffs and the SDF, leading to higher ex ante risk premia and higher ex post excess returns on the affected cohorts. Dispersion in A2P and dispersion in W2P therefore exert conflicting forces on spreads at origination: A2P dispersion pushes spreads up in a standard way, whereas W2P dispersion can pull spreads down on impact, even while raising tail risk.

More broadly, the discussion in this section delivers three main takeaways. First, credit risk has two distinct dimensions: ability to pay (A2P) and willingness to pay (W2P), which is harder to measure directly. Second, informational frictions about W2P—modeled here as noisy signals and underwriting rules based on those signals—create a wedge between the true distribution of default risk and the perceived distribution that lenders use to price credit. Changes in the dispersion of W2P can alter the composition of borrowers in the accepted pool in ways that reduce average default probabilities at origination but increase tail risk. Third, these mechanisms open the door to episodes in which credit spreads are temporarily low relative to the subsequent default experience and realized excess returns on the corresponding cohorts. In the dynamic model of Section 4, these conflicting forces play out over time. W2P dispersion shocks can compress spreads on impact, even as they raise conditional default risk and risk premia. This is a conditional version of the credit spread puzzle: spreads can be "too low" in exactly those states where W2P risk is most elevated and most poorly understood.

4 The W2P Economy

In this section, we develop a dynamic framework in which credit spreads are shaped by firms' cross-sectional heterogeneity in both willingness to pay (W2P) and ability to pay (A2P), together with a non-linear stochastic discount factor and within an endowment economy. This setting keeps the real side parsimonious while allowing to model time-varying cross-sectional volatility, informational frictions, and the asset pricing implications of W2P and information quality.

4.1 Economic environment and timing

The economy is populated by a representative *household*, which consumes the aggregate endowment, saves, and prices all assets (and their associated risks) in the economy; a unit mass of *firms* (borrowers), indexed by $i \in [0, 1]$, which demand external finance in the form of one-period loans; a competitive sector of *financial intermediaries* (lenders), which raise funds from the household and extend risky loans to borrowers. The aggregate state at date t is summarized by $S_t = \big(\sigma_{w,t}, \sigma_{a,t}, Y_t, \big)$, where $\sigma_{w,t}$ is the cross-sectional volatility (standard deviation across firms) of W2P, $\sigma_{a,t}$ is the cross-sectional volatility of A2P, and Y_t is the aggregate endowment.[5] We define in Table 1 the main variables used in the paper.

[Insert Table 1 Here]

Time is discrete, $t = 0, 1, 2, \ldots$. Within each period t we adopt the standard discrete-time convention and decompose the period into three stages. At the *start of period* t, the state inherited from the previous period, S_{t-1}, is known. New aggregate shocks ε_t are then realized

[5] We follow the endowment-economy literature (e.g., see Bansal and Yaron (2004) and Campbell and Cochrane (1999)) and assume an exogenous aggregate endowment process Y_t. In a frictionless benchmark, aggregate consumption equals the endowment each period, $C_t = Y_t$. In our W2P economy, by contrast, defaults generate pure resource losses, so that aggregate consumption is given by $C_t = Y_t - \text{losses}_t$, where losses_t collects the fraction $(1 - \delta)$ of promised loan payoffs that is destroyed in default, aggregated over all defaulting firms. Preferences and the stochastic discount factor are defined over C_t, while Y_t remains the primitive aggregate endowment that enters the state vector together with the cross-sectional dispersions $(\sigma_{w,t}, \sigma_{a,t})$.

and observed by all agents. These shocks determine the beginning of period aggregate state S_t, which includes aggregate endowment and the time-varying cross-sectional volatilities (see, for example, Campbell, Lettau, Malkiel and Xu (2001) and Herskovic, Kelly, Lustig and Van Nieuwerburgh (2016)). In particular, the innovations to $(\sigma_{w,t}, \sigma_{a,t})$ are components of ε_t, so that a shock to σ_w at date t indicates that $\sigma_{w,t}$ is moved away from its unconditional mean at the beginning of period t.

Within period t, after the shock, all decisions and contracts are made conditional on the realized aggregate state S_t. Firm-level types (a_i, w_i) are drawn from the cross-sectional distribution implied by $(\sigma_{w,t}, \sigma_{a,t})$. Banks observe each firm's A2P a_i but not W2P w_i. Instead, they observe a noisy signal $s_{i,t}$ about w_i and form the information objects. We refer to lenders' approval and pricing rule based on observed signals as underwriting (screening): lenders use the signal to decide which applicants are funded and at what rate. Given S_t and the information structure, banks choose an underwriting rule characterized by a cutoff s_t^* and a loan-rate schedule $R_{L,t}(a_i, s_{i,t})$ for firms that pass screening. Firms that receive loan offers decide whether to borrow. We normalize loan size $B_{i,t}$ to one for exposition, so that the intensive margin is shut down and the main action is in the extensive margin (which firms receive credit and at what rate). The endogenous variables of interest in period t—such as the loan approval rate, the physical and risk-neutral default probabilities for new loans, and the credit spread CS_t on new originations—are therefore functions of the shocks ε_t realized at the start of period t. In particular, a shock to σ_w at date t has a contemporaneous effect on CS_t; this contemporaneous movement is what we refer to as the impact response. Note that, in this framework, $\sigma_{w,t}$ is the dynamic counterpart to the static π_w in Section 3.2, the share of high-W2P firms in the approved pool. We model $\sigma_{w,t}$ as the primitive state variable for the cross-sectional dispersion of W2P; given information quality and underwriting, changes in $\sigma_{w,t}$ endogenously shift approved-pool composition toward high-W2P firms.

At the *end of period* t, aggregate output and firm cash flows are realized, and each borrowing firm decides whether to repay or default. When a firm defaults, the bank recovers

only a fraction δ of the promised payoff, while the remaining fraction $(1-\delta)$ is a pure resource loss, so that the representative household ultimately bears the real cost of default through lower aggregate consumption. Banks are therefore exposed to default risk through their portfolio payoffs, and the aggregate state transitions from S_t to S_{t+1} according to the model's law of motion. The updated state S_{t+1} serves as the starting point for period $t+1$.

[Insert Figure 2 Here]

Figure 2 provides a visual aid to the timing of events induced by shocks to σ_w. The three rows track, respectively, the aggregate state, firms' types and borrowing decisions, and banks' information, underwriting, and portfolio payoffs, while the columns group events into information, contracting, and end of period t payoffs and the transition to S_{t+1}. A shock to $\sigma_{w,t}$ shifts the aggregate state at the start of period t, propagates through firm heterogeneity and banks' underwriting and pricing, and then feeds into next period's state through realized default and loan portfolio payoffs. Solid arrows denote sequences of actions and realized flows within the period, and dashed arrows indicate informational and state-dependence links.

4.2 Firm types and cross-sectional heterogeneity

Each firm i is characterized by two persistent latent types. a_i (*ability to pay, A2P*): a latent index capturing the firm's cash-flow capacity and economic fundamentals (productivity, profitability, collateral quality, etc.). w_i (*willingness to pay, W2P*): a latent index capturing the firm's propensity to honor contracts above and beyond its ability to pay, reflecting governance, legal frictions, and the preferences of owners and managers.

At each date t, the cross-sectional distribution of (a_i, w_i) is smooth and summarized by dispersion parameters $(\sigma_{a,t}, \sigma_{w,t})$ and a correlation coefficient ρ_{aw}. We assume (a_i, w_i) is

jointly Gaussian:

$$\begin{pmatrix} a_i \\ w_i \end{pmatrix} \Bigg| \; S_t \sim \mathcal{N}\left(\begin{pmatrix} \bar{a} \\ \bar{w} \end{pmatrix}, \begin{pmatrix} (\sigma_{a,t})^2 & \rho_{aw}\sigma_{a,t}\sigma_{w,t} \\ \rho_{aw}\sigma_{a,t}\sigma_{w,t} & (\sigma_{w,t})^2 \end{pmatrix} \right). \tag{6}$$

The case $\rho_{aw} = 0$ corresponds to independent A2P and W2P; $\rho_{aw} > 0$ captures economies in which fundamental ability and willingness to pay tend to move together (e.g., well-governed firms are also more profitable), while $\rho_{aw} < 0$ captures environments with negative alignment between A2P and W2P (e.g., some highly profitable firms are more aggressive in exploiting limited liability).[6]

4.3 The signal-to-noise ratio and posterior beliefs

Financial intermediaries observe each firm's A2P a_i perfectly but do not observe W2P w_i directly. Instead, for each firm i in period t they observe a noisy signal

$$s_{i,t} = w_i + \nu_{i,t}, \tag{7}$$

where $\nu_{i,t}$ is an idiosyncratic noise term, independent across firms and over time, with $\nu_{i,t} \sim \mathcal{N}(0, \sigma_\nu^2)$ and independent of (a_i, w_i).

Given (6) and (7), the prior for w_i at date t is normal, $w_i \sim \mathcal{N}(\bar{w}, (\sigma_{w,t})^2)$. Gaussian conjugacy implies that the posterior distribution of w_i given the signal $s_{i,t} = s$ is also normal

$$w_i \mid s_{i,t} = s \sim \mathcal{N}\big(\mu_t(s), \sigma^2_{w|s,t}\big) \tag{8}$$

[6] Appendix A derives the conditional distributions and posterior beliefs when $\rho_{aw} \neq 0$ and shows how correlation affects information quality and underwriting.

with posterior mean

$$\mu_t(s) = \bar{w} + \kappa_t\big(s - \bar{w}\big), \qquad \kappa_t \equiv \frac{(\sigma_{w,t})^2}{(\sigma_{w,t})^2 + \sigma_\nu^2}, \tag{9}$$

and posterior variance

$$\sigma^2_{w|s,t} = \frac{(\sigma_{w,t})^2\sigma_\nu^2}{(\sigma_{w,t})^2 + \sigma_\nu^2}. \tag{10}$$

In this setting, posterior beliefs refer to the conditional distribution of w_i given the signal $s_{i,t} = s$, characterized by the mean $\mu_t(s)$ and variance $\sigma^2_{w|s,t}$. Posterior W2P refers to the posterior mean $\mu_t(s)$, and posterior default probabilities refer to default probabilities computed under this conditional distribution.

It is convenient to summarize information quality using a *signal-to-noise ratio* (SNR),

$$\Lambda_t \equiv \frac{(\sigma_{w,t})^2}{\sigma_\nu^2}, \tag{11}$$

which compares the ex ante dispersion in W2P to the variance of signal noise. With the noise variance σ_ν^2 constant, innovations to $\sigma_{w,t}$ increase signal informativeness, $d\Lambda_t/d\sigma_{w,t} > 0$. Thus a $\sigma_{w,t}$ shock simultaneously raises cross-sectional dispersion in W2P and increases the effective informativeness of the W2P signal, so our quantitative experiments are naturally phrased in terms of shocks to $\sigma_{w,t}$ (and to $\sigma_{a,t}$). Λ_t captures how informative the signal $s_{i,t}$ is about w_i: in high-Λ_t states the posterior mean $\mu_t(s)$ reacts strongly to the signal and the posterior variance $\sigma^2_{w|s,t}$ is small, whereas in low-Λ_t states the signal is noisy and underwriting is blunt. Substituting (11) into (9)–(10) yields

$$\kappa_t = \frac{\Lambda_t}{1 + \Lambda_t}, \qquad \sigma^2_{w|s,t} = \frac{(\sigma_{w,t})^2}{1 + \Lambda_t}. \tag{12}$$

High-Λ_t states correspond to posteriors that react strongly to the signal (large slope κ_t) and are tightly concentrated (small $\sigma^2_{w|s,t}$). Low-Λ_t states correspond to weakly informative signals and diffuse posteriors.

Note that Λ_t is a state variable for information quality. High-Λ_t states feature sharper sorting between safe and risky borrowers along the W2P dimension. This affects not only physical default probabilities but also how default risk loads on the SDF, and therefore how credit spreads decompose into expected losses and risk premia. In Section 5, we contrast high- and low-Λ_t states in terms of credit spreads, default rates, and realized excess returns.

4.4 Private cost of default and default probabilities

Given a promised gross loan rate $R_{L,t}(a_i, s_{i,t})$ and loan size $B_{i,t} = 1$, we model default as the outcome of a binary choice. The firm compares the benefit of walking away from the obligation to the private cost of default. Let $X(a_i)$ denote the deterministic component of firm i's cash flow relevant for servicing the loan, which is increasing in A2P. We specify the private default cost as an affine function of W2P,

$$\chi(w_i) = \chi_0 + \chi_1 w_i, \qquad \chi_1 > 0. \tag{13}$$

Intuitively, higher W2P increases the utility cost of defaulting.

Default occurs whenever the net gain from walking away is positive,

$$R_{L,t}(a_i, s_{i,t}) - X(a_i) > \chi(w_i).$$

Equivalently, there is a default *threshold* $\tau(a_i, s_{i,t}; R_{L,t})$ such that the firm defaults if

$$w_i < \tau(a_i, s_{i,t}; R_{L,t}) \equiv \frac{R_{L,t}(a_i, s_{i,t}) - X(a_i)}{\chi_1} - \frac{\chi_0}{\chi_1}. \tag{14}$$

Combining (8) and (14), the physical default probability of a firm with characteristics (a_i, s)

and contract $R_{L,t}$ is

$$\mathrm{PD}_t(a_i, s; R_{L,t}) = \Pr_t\big(w_i < \tau(a_i, s; R_{L,t}) \,\big|\, s_{i,t} = s\big) = \Phi\left(\frac{\tau(a_i, s; R_{L,t}) - \mu_t(s)}{\sigma_{w|s,t}}\right), \tag{15}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. High-SNR states (large Λ_t) steepen the mapping from s to default probabilities: small changes in s translate into large changes in $\mu_t(s)$ and therefore in $\mathrm{PD}_t(a_i, s; R_{L,t})$.

4.5 Household, recursive preferences and the SDF

The representative household consumes the aggregate endowment and prices all assets. Preferences are of the Epstein and Zin (1989) recursive form. Let C_t denote aggregate consumption (equal to the endowment net of any resource losses from default) and U_t the continuation utility. Preferences are

$$U_t = \left[(1-\beta)C_t^{1-1/\psi} + \beta\left(\mathbb{E}_t\big[U_{t+1}^{1-\gamma}\big]\right)^{\frac{1-1/\psi}{1-\gamma}}\right]^{\frac{1}{1-1/\psi}}, \tag{16}$$

where $\beta \in (0,1)$ is the subjective discount factor, $\gamma > 0$ is the coefficient of relative risk aversion, and $\psi > 0$ is the elasticity of intertemporal substitution (EIS). Note that equation (16) separates risk aversion from intertemporal substitution, so that the model can generate sizable movements in the traditional consumption-based credit risk premium CRP_t in response to shocks that affect aggregate consumption risk (such as shocks to A2P dispersion $\sigma_{a,t}$). In contrast, the W2P dispersion shocks $\sigma_{w,t}$ mainly operate through underwriting and default probabilities.

Let M_{t+1} denote the household's intertemporal stochastic discount factor.[7] For any gross payoff X_{t+1} denominated in units of consumption,

$$1 = \mathbb{E}_t\big[M_{t+1}X_{t+1}\big]. \tag{17}$$

[7] Appendix B further describes M_{t+1} in terms of $(C_t, C_{t+1}, U_t, U_{t+1})$.

The household has access to a one-period risk-free asset with gross return R_t^f, and to a diversified portfolio of bank loans with realized gross payoff $\widetilde{R}_{L,t+1}$. The Euler equations are

$$1 = \mathbb{E}_t\big[M_{t+1} R_t^f\big], \tag{18}$$

$$1 = \mathbb{E}_t\big[M_{t+1} \widetilde{R}_{L,t+1}\big]. \tag{19}$$

The difference between the expected return on the credit portfolio and the risk-free rate defines the model-implied credit risk premium.

4.6 Loan pricing, SNR and state-price-weighted probabilities

Financial intermediaries raise one-period funding from the household at the risk-free rate R_t^f and extend loans of size $B_{i,t} = 1$ to firms. A loan to firm i that is originated at date t promises a gross rate $R_{L,t}(a_i, s_{i,t})$, with realized payoff to the bank at $t+1$ given by

$$\Pi_{i,t+1}^L = \begin{cases} R_{L,t}(a_i, s_{i,t}), & \text{if firm } i \text{ repays,} \\ \delta\, R_{L,t}(a_i, s_{i,t}), & \text{if firm } i \text{ defaults,} \end{cases} \tag{20}$$

where $\delta \in [0, 1)$ is the recovery rate (capturing collateral, bankruptcy recoveries, etc.).

It is convenient to work with the *state-price-weighted* default probability

$$\Pi_t^*(a_i, s) \equiv \frac{\mathbb{E}_t\big[M_{t+1}\, \mathbb{I}_{i,t+1}^{\text{def}} \,\big|\, a_i, s_{i,t} = s\big]}{\mathbb{E}_t\big[M_{t+1} \,\big|\, a_i, s_{i,t} = s\big]} \in [0, 1], \tag{21}$$

where $\mathbb{I}_{i,t+1}^{\text{def}}$ is the indicator of default. Competitive banks choose $R_{L,t}(a_i, s)$ so that the expected, SDF-weighted gross payoff per unit lent equals the risk-free return,

$$\mathbb{E}_t\left[M_{t+1} \Pi_{i,t+1}^L \,\middle|\, a_i, s_{i,t} = s\right] = \mathbb{E}_t\big[M_{t+1} R_t^f\big]. \tag{22}$$

Using (20) and (21), this implies the the following default-loss adjustment in the promised

loan rate

$$R_{L,t}(a_i, s) = \frac{R_t^f}{1 - (1 - \delta)\,\Pi_t^*(a_i, s)}. \tag{23}$$

Here $R_{L,t}(a_i, s)$ denotes the promised loan rate, i.e., the contractual gross rate the borrower agrees to pay in the absence of default. Equation (23) shows that the promised loan rate increases with the state-price-weighted default probability and with the expected loss given default $(1 - \delta)$.

Financial intermediaries use the signal $s_{i,t}$ to implement a monotone approval rule. For each (S_t, a_i), they specify a cutoff s_t^* and only extend loans to firms whose signal exceeds the cutoff:

$$s_{i,t} \geq s_t^*. \tag{24}$$

Let $g_t(s)$ denote the cross-sectional density of signals implied by (6) and (7). We denote by $\lambda_t \in [0, 1]$ the fraction of firms approved for credit at date t,

$$\lambda_t = \int_{s_t^*}^{\infty} g_t(s)\, ds \equiv \mathcal{A}\big(\sigma_{a,t}, \sigma_{w,t}, \Lambda_t, s_t^*\big), \tag{25}$$

where the last equality defines the acceptance function $\mathcal{A}(\cdot)$ and emphasizes its dependence on the SNR Λ_t. High-SNR states make the acceptance function more sensitive to the cutoff s_t^*: a small change in s_t^* sharply changes the composition of the approved pool.

Let PD_t denote the aggregate (pool) default probability for the cohort of loans originated at date t:

$$\mathrm{PD}_t = \frac{\displaystyle\int_{s_t^*}^{\infty} \mathbb{E}_t\big[\mathrm{PD}_t(a_i, s; R_{L,t}) \,\big|\, s_{i,t} = s\big]\, g_t(s)\, ds}{\displaystyle\int_{s_t^*}^{\infty} g_t(s)\, ds}, \tag{26}$$

where PD_t is conditional on origination (equivalently, on passing the underwriting cutoff $s_{i,t} \geq s_t^*$). We also define the (pre-underwriting) population default probability as the

unconditional cross-sectional average across all potential borrowers

$$\mathrm{PD}_t^{\mathrm{pop}} \equiv \int_{-\infty}^{\infty} \mathbb{E}_t\big[\mathrm{PD}_t(a_i, s; R_{L,t}) \,\big|\, s_{i,t} = s\big] g_t(s)\, ds, \tag{27}$$

which differs from PD_t by not conditioning on $s \geq s_t^*$.

Let $\Pi_t^{*,\mathrm{pool}}$ denote the analogous state-price-weighted pool default probability:

$$\Pi_t^{*,\mathrm{pool}} = \frac{\displaystyle\int_{s_t^*}^{\infty} \mathbb{E}_t\big[\Pi_t^*(a_i, s) \,\big|\, s_{i,t} = s\big] g_t(s)\, ds}{\displaystyle\int_{s_t^*}^{\infty} g_t(s)\, ds}. \tag{28}$$

A convenient way to model underwriting is to assume that financial intermediaries target a fixed *state-price-weighted* default probability $\bar{\Pi} \in (0, 1)$ across newly originated loans, reflecting risk-management and regulatory constraints. Formally, the underwriting cutoff s_t^* is chosen so that

$$\Pi_t^{*,\mathrm{pool}} = \bar{\Pi}. \tag{29}$$

Given $(S_t, \Lambda_t, \bar{\Pi})$, this condition determines s_t^* as a function of the informativeness of the signal (through Λ_t) and the cross-sectional distribution of (a_i, w_i). We show in Appendix C that larger Λ_t increases the sensitivity of posterior default probabilities to s, changing the equilibrium cutoff and acceptance rate.

4.7 Aggregate credit payoffs, spreads, and equilibrium

Let $\widetilde{R}_{L,t+1}$ denote the realized gross return on the representative portfolio of loans originated at date t,

$$\widetilde{R}_{L,t+1} = (1 - D_{t+1}) R_{L,t} + D_{t+1} \delta R_{L,t}, \tag{30}$$

where $D_{t+1} \in \{0,1\}$ is the cohort-level default indicator and $R_{L,t}$ is the representative promised rate for the portfolio. Taking expectations conditional on date-t information,

$$\mathbb{E}_t\big[\widetilde{R}_{L,t+1}\big] = (1-\mathrm{PD}_t)R_{L,t} + \mathrm{PD}_t\delta R_{L,t} = \big[1-(1-\delta)\mathrm{PD}_t\big]R_{L,t}. \tag{31}$$

Note that $R_{L,t}$ is the representative promised gross loan rate for the cohort originated at date t, while $\widetilde{R}_{L,t+1}$ is the realized gross payoff between t and $t+1$, and $\mathbb{E}_t[\widetilde{R}_{L,t+1}]$ is the corresponding expected gross return conditional on date-t information.

We define the *credit spread* at date t as the difference between the promised loan rate and the risk-free rate,

$$\mathrm{CS}_t \equiv R_{L,t} - R_t^f, \tag{32}$$

and the *expected excess return* (credit risk premium) on the cohort of loans originated at t as

$$\mathrm{CRP}_t \equiv \mathbb{E}_t\big[\widetilde{R}_{L,t+1}\big] - R_t^f. \tag{33}$$

In Appendix D we decompose CS_t and CRP_t into (i) an expected-loss component driven by PD_t and δ, (ii) a risk-premium component driven by the covariance with M_{t+1}, and (iii) a W2P mispricing component arising from informational frictions and underwriting.

Given an exogenous stochastic process for the aggregate state $S_t = (\sigma_{w,t}, \sigma_{a,t}, Y_t)$, the equilibrium is a collection of stochastic processes

$$\big\{C_t, R_t^f, R_{L,t}, \mathrm{CS}_t, \mathrm{PD}_t, \lambda_t, s_t^*, U_t, M_{t+1}, \Lambda_t\big\}_{t\geq 0},$$

such that: (i) given prices and the stochastic discount factor, the household chooses consumption and asset holdings to maximize Epstein–Zin utility (16) and the Euler equations (18)–(19) hold, (ii) given (S_t, Λ_t, R_t^f) and the belief structure, banks choose underwriting rules and loan rates $(s_t^*, R_{L,t}(a_i, s))$ that satisfy the pricing condition (22) and the target default constraint (29), and (iii) markets clear and aggregate variables are consistent with

the endowment process and default losses.

It is worth noting that the model reflects a rational, information-based underwriting channel: even without any improvement in signal precision, higher dispersion in repayment incentives increases the informational content of lenders' signals and the value of screening, so origination spreads can compress through pool composition even as population default risk rises. This mechanism differs from behavioral explanations (e.g., Carvalho, Gao and Ma (2023)) and from pure liquidity-dislocation stories. It also differs from intermediary-balance-sheet channels because, in our W2P economy, spreads move through underwriting and pool composition rather than through shifts in the economy-wide price of risk (e.g., see He and Krishnamurthy (2013) and He, Kelly and Manela (2017)).

5 Quantitative Implications and Asset-Pricing Results

In this section, we study the quantitative implications of the model. We show how shocks to $\sigma_{w,t}$ (cross-sectional volatility in willingness to pay) generate spread compression at issuance even as downside outcomes can materialize later, providing a W2P-based mechanism that helps reconcile the puzzling mismatch between observed default rates and model-implied credit spreads (Huang, Nozawa and Shi (2025)).

5.1 Shocks, parameters and calibration

We specify laws of motion for the aggregate state with $(\varepsilon_{\sigma_a,t+1}, \varepsilon_{\sigma_w,t+1})$ jointly Gaussian innovations

$$\sigma_{a,t+1} = (1-\rho_{\sigma_a})\,\bar{\sigma}_a + \rho_{\sigma_a}\sigma_{a,t} + \sigma_{\sigma_a}\varepsilon_{\sigma_a,t+1}, \tag{34}$$

$$\sigma_{w,t+1} = (1-\rho_{\sigma_w})\,\bar{\sigma}_w + \rho_{\sigma_w}\sigma_{w,t} + \sigma_{\sigma_w}\varepsilon_{\sigma_w,t+1}, \tag{35}$$

Recursive preferences, combined with stochastic volatility imply that approximations of

higher-order can capture non-linearities in the response of M_{t+1} and of credit payoffs to W2P dispersion and SNR shocks. In the baseline quantitative exercises we feed shocks into $\sigma_{a,t}$ and $\sigma_{w,t}$, keeping σ_ν constant so that Λ_t is an implied function of $\sigma_{w,t}$ and the exogenous noise level. Appendix E introduces shocks to Λ_t as pure information-quality shocks and provides details on the state-space representation for solution method around the steady state.

Our benchmark parameterization proceeds in four blocks. First, we set preferences and endowment dynamics: parameters (β, γ, ψ) and (ρ_y, σ_y) discipline the representative investor's stochastic discount factor.[8] Second, we set cross-sectional heterogeneity and credit risk: we choose the long-run dispersion and persistence of A2P, $(\bar{\sigma}_a, \rho_{\sigma_a})$, to match moments of cross-sectional heterogeneity in observable firm fundamentals, and we choose the long-run dispersion and persistence of latent W2P, $(\bar{\sigma}_w, \rho_{\sigma_w})$, to match moments of default probabilities and credit spreads conditional on fundamentals. Third, we set information frictions and the SNR: we choose the signal noise variance σ_ν^2 to match average loan pricing and underwriting intensity, and we parameterize the dynamics of $\sigma_{w,t}$; with σ_ν^2 held fixed, the signal-to-noise ratio Λ_t (and thus κ_t) then moves with $\sigma_{w,t}$. Fourth, we set recovery and loss given default: the recovery rate δ matches average loss-given-default measures in corporate credit markets. Table 2 reports the selection of parameters in our benchmark model.

[Insert Table 2 Here]

5.2 Underwriting and selection effects of W2P and SNR

The underwriting (screening) channel can be formalized by a local experiment in the loan-market block of the full model: we hold the equilibrium SDF and the endowment process fixed and focus on how the acceptance cutoff s_t^* and the composition of approved borrowers respond to $\sigma_{w,t}$ and Λ_t. High-SNR states (large Λ_t) make the signal s more informative,

[8] Preference and endowment parameters follow long-run-risk and habit benchmark calibrations. Since these aggregate parameters are weakly identified by credit data, we hold them fixed and discipline the model using credit-market moments—average default probabilities, credit spreads, and the credit risk premium—to isolate the W2P-versus-A2P channel.

sharpening the link from s to the posterior W2P distribution $\big(w_i \mid s_{i,t} = s\big)$ and hence to default probabilities.

Proposition 1. *Given $(Y_t, \sigma_{a,t})$ and a target state-price-weighted default rate $\bar{\Pi} \in (0,1)$, consider the information structure (6)–(7) with posterior beliefs characterized by (8)–(12). For a small increase in the SNR Λ_t—for instance, induced by higher W2P dispersion $\sigma_{w,t}$ at fixed noise variance σ_ν^2—the following comparative statics hold:*

(i) The optimal underwriting cutoff s_t^* that solves (29) is weakly lower, so that the approved set expands.

(ii) The acceptance rate λ_t in (25) is weakly higher, and aggregate loan volumes increase.

(iii) Even though the unconditional (population) default probability rises with $\sigma_{w,t}$, the state-price-weighted default probability in the approved pool remains fixed at $\bar{\Pi}$, so that underwriting becomes more selective: high-SNR states allow more aggressive expansion of credit at the same target risk.

The economic intuition is that an increase in SNR Λ_t creates more *spread* in true W2P and at the same time raises the informativeness of the signal s. Lenders can then screen more finely on s: the approval rule $s \geq s_t^*$ selects a pool of borrowers whose posterior W2P is higher on average and more tightly distributed. Condition (29) determines the state-price-weighted default probability in the approved pool, but does not determine its size. In contrast, the population (pre-underwriting) default probability $\mathrm{PD}_t^{\mathrm{pop}}$ in (27) is an unconditional cross-sectional object that does not condition on approval. Higher SNR allows banks to lower the cutoff s_t^*, approve more firms, and still satisfy the risk target $\bar{\Pi}$.[9] A formal proof is provided in Appendix F.

Proposition 1 formalizes the underwriting channel: given $(Y_t, \sigma_{a,t})$ and a target state-price–weighted default bound, a higher signal-to-noise ratio Λ_t sharpens underwriting, raises

[9] Throughout, PD_t and $\mathrm{PD}_t^{\mathbb{Q}}$ refer to originated-loan (pool) default probabilities as in (26) and its risk-neutral analogue.

the acceptance rate, and improves the default profile of the approved pool. When σ_ν^2 is unchanged, higher $\sigma_{w,t}$ raises Λ_t and thus strengthens this underwriting mechanism. To connect the mechanism to the dynamic model, we construct "vintage" objects from the ergodic distribution and examine how ex-ante spreads relate to ex-post default outcomes as information quality varies.

We simulate the benchmark economy for a long horizon and, for each quarter t, construct a cohort-level outcome

$$PD_{t\to t+K} = \frac{1}{K}\sum_{j=1}^{K} PD_{t+j},$$

defined as the average physical default rate of newly originated loans between $t+1$ and $t+K$. We interpret $PD_{t\to t+K}$ as a forward default-vintage (origination-regime) measure, not as realized defaults of the loans originated at t. In the implementation we set $K = 8$ (two years), annualize the default rate, and keep only vintages with strictly positive $PD_{t\to t+K}$. In this vintage sample, average CS_t is on the order of a few hundred basis points with substantial dispersion, and the mean forward origination default rate is a few hundred basis points with a high cross-vintage standard deviation. The correlation between CS_t and $PD_{t\to t+K}$ is 0.56, so the model delivers a realistic baseline link between current spreads and future default rates.

[Insert Figure 3 Here]

Figure 3 plots this cloud of vintages, with each point colored by the contemporaneous signal-to-noise ratio Λ_t at origination. On the horizontal axis we plot the annualized origination spread CS_t at date t. The vertical axis uses $\log PD_{t\to t+K}$. Darker colors correspond to higher Λ_t, that is, states in which the W2P signal is especially precise relative to noise. Along most of the cloud, higher spreads line up with higher future default rates, and the relationship appears concave in levels on the log scale: spreads are most informative about default risk when they are relatively tight and become less informative once they are already wide. In this core region, vintages with higher Λ_t tend to lie closer to the steep part of

the curve. Better information sharpens underwriting and amplifies the responsiveness of default outcomes to small changes in spreads, consistent with the comparative statics in Proposition 1.

Interestingly, the color dimension also highlights the "wrong-way" vintages that matter for our W2P mechanism. A non-trivial mass of points lies in the top left of the figure—low spreads but high subsequent default rates—and these vintages are disproportionately dark, meaning they occur in states with high Λ_t. High Λ_t here reflects precise signals about borrowers' willingness-to-pay, not about their fundamental ability-to-pay. Lenders put a large posterior weight κ_t on the W2P signal and approve many borrowers who look excellent on W2P but are weak on fundamentals. Note that these loans receive low spreads at origination and later default at elevated rates. In that sense, the Λ_t-colored scatter makes Proposition 1 tangible in the full model: the same information structure that allows aggressive but safe credit expansion in benign states can also generate "puzzle" vintages where lenders underprice risk with great confidence.

In Figure 4, we make the two-regime comparison more transparent by splitting the same cloud of vintages into low- and high-dispersion states. Panel A restricts to quarters in the bottom quartile of σ_w at origination, while Panel B keeps only quarters in the top quartile.

[Insert Figure 4 Here]

In Panel A (low σ_w), there is a positive relationship between CS_t and $PD_{t\to t+K}$, but the cloud of points is relatively tight and there are few extreme outliers. In Panel B (high σ_w), high-dispersion states makes creening more aggressive and therefore more fragile to misalignment between W2P and A2P. The relationship remains upward sloping on average but becomes much more dispersed: for a given initial spread, subsequent default outcomes vary widely, and some low-spread vintages later experience surprisingly high default rates. These are precisely the high-dispersion, high-SNR regimes in which Proposition 1 predicts that underwriting will be aggressive but fragile: lenders rely heavily on a very sharp W2P

signal, so any misalignment between W2P and A2P produces pools with low spreads but elevated future default rates. Figure 3 and Figure 4 show that the underwriting mechanism is strongest—and most prone to mistakes—in high-dispersion, high-SNR states.

To contrast the differences between cross-sectional volatility in willingness to pay and ability to pay, we repeat the origination spread and forward default outcome scatter of Figure 4, but condition on fundamentals-driven A2P dispersion by splitting the sample into low-σ_a (bottom quartile) and high-σ_a (top quartile) states. Figure 5 plots these results. In both panels, the cloud remains strongly upward sloping: low spreads coincide with low subsequent default outcomes, and higher spreads coincide with higher subsequent default outcomes. The correlation between origination spreads and the next-8-quarter average default outcome is high and nearly unchanged across regimes (about 0.61 in both low- and high-σ_a states). This stands in sharp contrast to the analogous conditioning on σ_w, where the relationship becomes much flatter and "wrong-way" realizations are more prevalent (correlations around 0.31 in low-σ_w states and 0.19 in high-σ_w states). Taken together, the evidence suggests that the wrong-way spread and default patterns are not a feature of time-varying fundamentals risk; they are specific to the W2P dispersion information channel.

[Insert Figure 5 Here]

5.3 Impact response to a transitory W2P-dispersion shock

In this section, we investigate how shocks to the cross-sectional dispersion of firms' willingness to pay ($\sigma_{w,t}$) shift Λ_t, κ_t, credit spreads and risk premiums over time, and how those shocks move the economy into and out of the regimes highlighted by the vintage scatters. Intuitively, the shock raises the mass of both very-high-W2P and very-low-W2P borrowers, while leaving average A2P and signal noise unchanged. In this impulse–response experiment, the innovation to $\sigma_{w,t}$ has two linked effects: it raises dispersion in W2P types and, with σ_ν^2 unchanged, it increases the informativeness of lenders' signals and the effectiveness of

underwriting.

To study the information-channel dynamics, we implement an impulse–response exercise: $\sigma_{w,t}$ follows a persistent AR(1) with coefficient 0.9, is hit by a one-standard-deviation innovation at $t = 0$ and then evolves according to its calibrated law of motion, with all other shocks shut down. Figure 6 plots in Panel A the resulting responses of the information block—the signal-to-noise ratio Λ_t and the posterior weight κ_t—following the σ_w shock.

[Insert Figure 6 Here]

A positive σ_w shock simultaneously raises the dispersion of borrowers' W2P and the precision of the W2P signal. In the model, signal-to-noise $\Lambda_t = \sigma_{w,t}^2/\sigma_\nu^2$ jumps on impact by about 0.0165 and then decays gradually, reaching roughly half of its initial effect around quarter 8. The posterior weight on the W2P signal, $\kappa_t = \Lambda_t/(1 + \Lambda_t)$, moves almost one-for-one with Λ_t: it rises by about 0.012 on impact and likewise halves by quarter 8. These information variables govern the underwriting cutoff and the composition of the approval pool. As κ_t rises, lenders put more weight on the W2P signal and approve a safer set of borrowers. Panel B of Figure 6 reports the response of the population (pre-underwriting) default probability, $\mathrm{PD}_t^{\mathrm{pop}}$, defined in (27). Following a positive $\sigma_{w,t}$ shock, $\mathrm{PD}_t^{\mathrm{pop}}$ increases on impact (about 5 bps) and then mean-reverts as the transitory dispersion shock fades. Intuitively, higher W2P dispersion thickens the left tail of repayment incentives in the borrower population, raising unconditional default risk even before conditioning on origination. In contrast, the next figures show that underwriting adjusts so that the default risk of newly originated loans falls despite the increase in population risk.

The key implication is therefore a wedge between *population* and *pool* risk: the dispersion shock raises $\mathrm{PD}_t^{\mathrm{pop}}$, but, because the same shock increases signal precision, underwriting and selection shift the composition of originations toward safer borrowers. Consistent with this mechanism, the physical default probability of newly originated loans, PD_t, drops sharply on impact: relative to the benchmark, PD_t falls by about 0.0021 (roughly 21 basis points

of annual default probability). The risk-neutral default probability of the originated pool, $\mathrm{PD}_t^{\mathbb{Q}}$, also falls by about 21 basis points. Figure 7 summarizes these pricing-relevant default dynamics: Panel A reports the impulse response of $\mathrm{PD}_t^{\mathbb{Q}}$, while Panel B reports the risk-neutral–physical default gap, $\mathrm{PD}_t^{\mathbb{Q}} - \mathrm{PD}_t$, which moves very little on impact (about half of a basis point).

[Insert Figure 7 Here]

We next study the effects of shocks to the cross-sectional volatility of firms' W2P on credit spreads and risk premia. The model impulse responses of Figure 8 show that the effect on loan pricing is striking. On impact, the credit spread CS_t on new originations falls by about 0.00133 in levels, that is roughly 13 basis points annualized. The response is very persistent: by quarter 8 the spread effect has shrunk to about -6 basis points, and the half-life of the initial impact is around eight quarters. Averaged over the first eight quarters, the credit spread is about 9 basis points lower than in steady state; averaged over the first five years (20 quarters), the spread remains about 6 basis points lower; even over the full 10-year window, the average spread reduction is still around 3 basis points.

[Insert Figure 8 Here]

Figure 8 highlights the core quantitative result of the paper: an increase in the dispersion of W2P leads lenders to cut loan rates, and it does so in a way that is both economically large on impact and quantitatively persistent. Credit spreads combine expected losses and a discount-rate component (e.g., Nozawa (2017)). Our model highlights a novel channel—time-varying cross-sectional volatility in W2P under imperfect information—that moves spreads primarily through default expectations and underwriting outcomes, with little effect on the conventional credit risk premium.

To better highlight the role of W2P-based information shocks, we contrast this experiment with an increase in the cross-sectional dispersion of firms' ability to pay, a fundamentals

shock to σ_a. Figure 9 reports the responses of CS_t and the expected excess return CRP_t to a one-standard-deviation increase in σ_a.

[Insert Figure 9 Here]

A one-standard-deviation increase in σ_a leaves the information block essentially flat: by construction, Λ_t and κ_t do not move, and the physical default probability of new originations remains almost unchanged. Instead, the main effect is on risk-neutral default probability and the risk premium. On impact, a σ_a shock raises $\mathrm{PD}_t^{\mathbb{Q}}$ by 15 basis points, while PD_t remains essentially unchanged. Credit spreads widen by 9 basis points, and the expected excess return CRP_t also rises by roughly the same amount. In other words, the σ_a shock generates the conventional risk-based response: spreads move almost one-for-one with the risk premium, and the entire effect is driven by priced risk, not by changes in underwriting or physical default risk.

By contrast, under the σ_w shock the risk-premium component moves little. On impact, CRP_t increases by only about 1–2 tenths of a basis point, and the response then decays quickly; this is an order of magnitude smaller than the roughly 13-basis-point fall in the spread itself. Importantly, this σ_w exercise delivers a very different decomposition: most of the spread movement comes from the level of the loan rate relative to the risk-free rate, driven by better screening and a safer borrower pool, rather than from higher compensation for risk. This is precisely the mechanism highlighted by our model in Section 4: cross-sectional W2P dispersion, for a given level of signal noise, makes W2P signals more informative and allows lenders to reprice credit in a way that lowers spreads while also lowering both physical and risk-neutral default probabilities.

Taken together, these impulse responses deliver the central mechanism of the paper: a transitory increase in W2P dispersion leads to lower contemporaneous credit spreads because underwriting becomes more selective and the approved origination pool becomes safer, so

both physical and risk-neutral default probabilities of new loans fall on impact.[10] At the same time, greater dispersion raises unconditional borrower default risk by putting more mass on low-W2P firms, increasing the scope for low-spread, high-default episodes. In contrast, ability shocks primarily move credit spreads through the standard risk-premium channel. This W2P underwriting channel thus provides a microfounded mechanism that helps reconcile key features of credit spreads, default probabilities, and credit risk premia.

A natural benchmark in asset-pricing is that large spread movements reflect large movements in priced risk. Our model implies a different, and empirically useful, separation: shocks to $\sigma_{w,t}$ reprice credit primarily through underwriting-driven changes in the composition of the accepted pool and the associated risk-neutral expected losses, while the conventional discount-rate (risk-premium) component remains comparatively stable. This delivers a clean identifying restriction: spreads can compress because who gets funded changes and risk-neutral default expectations move, not because the market price of risk falls. Consequently, the model predicts "wrong-way" episodes (low spreads at issuance followed by elevated defaults) with large movements in underwriting and default expectations, but only modest movements in the credit risk premium.

5.4 Persistent W2P dispersion, SNR, and repricing

The previous section emphasizes the short-run effects of a one-time σ_w shock on credit spreads and default outcomes. We next interpret the same impulse–response experiment as a persistent "high-σ_w" regime. Because both σ_w and Λ_t follow AR(1) processes, a one-standard-deviation innovation to $\sigma_{w,t}$ moves the economy into a high-dispersion, high-SNR state that keeps signal precision and pool quality elevated for several quarters, not just on impact, and thereby affects spreads, default probabilities, and risk premia over a longer horizon. With σ_ν^2 constant in this analysis, an innovation to $\sigma_{w,t}$ raises $\Lambda_t = \sigma_{w,t}^2/\sigma_\nu^2$, so a

[10] Related mortgage evidence shows that selection at origination maps into later defaults (Gupta and Hansman (2022)).

given signal separates high- and low-W2P borrowers more cleanly and underwriting becomes more discriminating.

On the information side, Figure 6 shows that the initial jump in Λ_t and κ_t decays only gradually. Even eight quarters after the shock, Λ_t and κ_t are still roughly half of their impact deviations from steady state, and after five years they remain clearly above zero. This implies that the underwriting rule is persistently tighter, so that a sizable share of newly originated loans over the next several years are drawn from a safer portion of the cross section. Consistent with this, both physical and risk-neutral default probabilities stay depressed relative to steady state for an extended period. Over the first two years, the average reduction in physical default probability is on the order of 15 basis points, and over the full 10-year horizon the average reduction is still around 5 basis points; the risk-neutral default probability displays a very similar pattern.

This persistent improvement in pool quality translates into a persistent repricing of credit spreads. While the impact effect in Figure 8 is about -13 basis points, the credit spread remains substantially below steady state for many quarters: on average, CS_t is about 9 basis points lower than baseline over the first eight quarters, about 6 basis points lower over the first five years, and about 3 basis points lower when averaged over the full 40-quarter window. In other words, a single realization of a high-σ_w shock leaves a persistent footprint in the term structure of expected lending rates, with a long-lived compression of credit spreads that is quantitatively meaningful at asset-pricing horizons.

The comparison with the σ_a shock in Figure 9 highlights that this persistent repricing is specific to the W2P channel. Because σ_a affects only fundamentals dispersion and leaves Λ_t and κ_t unchanged, it does not produce a sustained reallocation in the physical default probability of new loans. Its persistent effect is instead concentrated in risk-neutral default probabilities and risk premia, with spreads widening for several years because lenders require more compensation for systematic default risk. In contrast, the σ_w shock generates a persistent regime in which financial intermediaries expect lower default rates but nev-

ertheless charge lower spreads, precisely because better information allows them to screen out weak borrowers without mechanically raising risk compensation. This "high-σ_w" regime interpretation ties directly to the vintage evidence documented above: loans originated in high-dispersion periods correspond to the parts of Figures 3 and 4 where initial spreads are compressed yet subsequent default rates are elevated. In this sense, the IRF experiment and the scatter plots of current spreads versus future default rates are two complementary views of the same W2P underwriting mechanism.

5.5 Information precision and the strength of the W2P mechanism

We next study how information precision interacts with W2P dispersion to determine the strength of the underwriting channel. To isolate financial information effects, consider two economies that differ only in the SNR process: a low-SNR economy with $\bar{\Lambda}$ small and shocks to Λ_t muted, and a *high-SNR* economy with larger $\bar{\Lambda}$ and more volatile Λ_t. In both economies we expose the system to identical shocks to $\sigma_{w,t}$ and compute generalized impulse responses of CS_t, PD_t, CRP_t, and λ_t. In the low-SNR economy, lenders have a coarse view of W2P and can only weakly condition on the signal. W2P dispersion shocks generate relatively small changes in s_t^* and λ_t. Spreads move mainly through changes in unconditional default probabilities and through the covariance of credit payoffs with M_{t+1}. In the high-SNR economy, signals are very informative. Empirically, relationship-based lending is naturally associated with higher information precision: personal lending relationships reduce spreads and expand credit supply, especially in downturns (Karolyi (2018); Paravisini, Rappoport and Schnabl (2023)). W2P dispersion shocks generate large changes in the composition of the borrowing pool: underwriting becomes sharper, and spreads move strongly at impact. The spread-compression effect is strongest in high-SNR states, and ex post surprises in defaults and excess returns are largest

We vary financial information precision by changing the variance of the W2P noise term, σ_ν^2, holding all other parameters at their benchmark values. A lower σ_ν^2 raises the signal-

to-noise ratio Λ_t for any given dispersion $\sigma_{w,t}$, while a higher σ_ν^2 dampens Λ_t. For each alternative value of σ_ν^2 we recompute the impulse responses to a one-standard-deviation shock to σ_w. In low-SNR economies (high σ_ν^2), the response of Λ_t and κ_t to a σ_w shock is muted, underwriting reacts only weakly, and the spread compression documented in Figure 8 becomes smaller. As information becomes noisier, the model converges toward a standard "quantity of risk" story: σ_w still changes cross-sectional dispersion but lenders cannot use W2P to screen effectively, so spreads hardly move. In high-SNR economies (low σ_ν^2), the opposite occurs. The same σ_w shock generates larger movements in Λ_t and κ_t, stronger shifts in the composition of the approval pool, and a more pronounced drop in credit spreads. The W2P mechanism is therefore strongest when signals about borrowers' willingness-to-pay are precise.

In the model of Section 4, the signal-to-noise ratio for W2P satisfies

$$\Lambda_t = \frac{\sigma_{w,t}^2}{\sigma_\nu^2},$$

so that changing σ_ν^2 rescales Λ_t for any given dispersion $\sigma_{w,t}$. We consider three calibrations that differ only in the value of σ_ν: a low-SNR economy with $\sigma_\nu = 1.0$, the benchmark economy with $\sigma_\nu = 0.5$, and a high-SNR economy with $\sigma_\nu = 0.25$. For each case we recompute the impulse responses to a one-standard-deviation increase in σ_w and compare the resulting movements in Λ_t, κ_t, default probabilities, and credit spreads.

On the information side, the response of Λ_t is as expected. Relative to the benchmark, the impact response of Λ_t in the high-SNR economy is roughly four times larger, while in the low-SNR economy it is roughly one quarter as large. The posterior weight on the W2P signal, $\kappa_t = \Lambda_t/(1+\Lambda_t)$, responds in the same direction but less than proportionally, because it saturates as Λ_t increases: the impact response of κ_t is about twice as large in the high-SNR economy as in the benchmark, and roughly one third as large in the low-SNR economy. In all three cases the shape of the impulse response is the same, with a half-life of about eight

quarters; only the amplitude changes.

The consequences for default probabilities and spreads follow directly. In the benchmark calibration, a one-standard-deviation σ_w shock lowers the physical default probability of new loans by about 21 basis points on impact and by about 10 basis points at quarter 8; the origination spread CS_t falls by about 13 basis points on impact and by about 6 basis points at quarter 8. In the high-SNR economy ($\sigma_\nu = 0.25$), the same σ_w shock produces roughly twice as large movements: the impact drop in physical default probability is about 42 basis points and remains about 20 basis points after eight quarters, while the spread compression is about 25 basis points on impact and about 12 basis points at quarter 8. In the low-SNR economy ($\sigma_\nu = 1.0$), the mechanism is weaker. The impact reduction in physical default probability is about 7 basis points, the quarter-8 effect is about 3 basis points, and the corresponding spread compression is only about 4 basis points on impact and about 2 basis points at quarter 8.

The persistence of these effects is similar across calibrations; what changes is their magnitude. Risk premia move little in all cases. For the benchmark calibration, the impact response of the expected excess return CRP_t to a σ_w shock is on the order of 0.2 basis points, and it is well below 1 basis point in both the low- and high-SNR economies as well. Quantitatively, the σ_w experiment therefore delivers a "component separation": the spread response is an order of magnitude larger than the risk-premium response. This pattern is closely aligned with empirical decompositions that separate an expected-loss component from a residual risk-premium component, as in the "excess bond premium" decomposition in Gilchrist and Zakrajšek (2012). In our mechanism, the underwriting rule reallocates origination toward safer borrowers, so the expected-loss piece of CS_t moves sharply, while the priced-risk component (and hence CRP_t) remains comparatively stable. In contrast, the movements in CS_t range from roughly 4 basis points in the low-SNR case to roughly 25 basis points in the high-SNR case. Thus, variation in information precision operates almost entirely through underwriting and the level of loan rates relative to the risk-free rate, not

through compensation for risk.

Figure 10 summarizes these patterns visually. Panel A shows that the response of the credit spread to a σ_w shock has the same shape across all three calibrations, but its amplitude ranges from only a few basis points in the low-SNR economy to about 25 basis points on impact in the high-SNR economy. Panel B shows an analogous scaling for the physical default probability of new loans. The persistence of the responses is very similar across SNR levels, confirming that information precision primarily changes the strength, rather than the duration, of the W2P mechanism.

[Insert Figure 10 Here]

The comparison across low-, benchmark-, and high-SNR economies in Figure 10 highlights the information content of W2P shocks from an asset-pricing perspective. On the information side, changing σ_ν^2 mechanically rescales the impulse response of $\Lambda_t = \sigma_{w,t}^2/\sigma_\nu^2$ while leaving its shape unchanged. Economically, this means that the same underlying shock to W2P dispersion can correspond either to a very noisy environment, in which signals barely move the posterior, or to a relatively transparent environment, in which the same dispersion shock produces much sharper revisions in perceived W2P. On the underwriting side, the posterior weight κ_t responds very little in the low-SNR economy: when information is too noisy, lenders cannot meaningfully tilt acceptance decisions based on s, so pool composition, credit spreads, and expected excess returns barely move. By contrast, in the high-SNR economy a given innovation to σ_w translates into larger movements in pool quality and sizable compressions in credit spreads, because lenders trust the signal enough to screen aggressively but still face residual uncertainty about fundamentals. The high-SNR economy loads the W2P shock more strongly into state-contingent payoffs: small changes in dispersion generate large swings in model-implied risk premia and realized excess returns on the loan portfolio, even though underlying endowment risk is unchanged.

These experiments therefore show that the quantitative importance of the W2P mecha-

nism is highest when information is precise enough to discipline spreads yet remains imperfect, so that some low-spread vintages still embed substantial downside risk. This sensitivity analysis also helps interpret the empirical vintage evidence: in a high-SNR, high-σ_w regime, the model generates the largest ex post surprises in defaults and excess returns relative to spreads, consistent with the scatter plots where the lowest spreads coincide with the highest future defaults under given underwriting and where "wrong-way" vintages are most pronounced.

5.6 W2P and the credit spread puzzle

We next investigate the implications of the W2P mechanism for the classic credit–spread puzzle. The goal is to confront the benchmark calibration with simple long-run moments that jointly summarize average credit spreads (CS), physical default probabilities and recoveries, and the implied credit-risk premium (CRP). Table 3 reports, for each object, an empirical range drawn from the investment-grade corporate bond literature and the corresponding moment generated by the W2P model. All entries are in annualized basis points (bps).

[Insert Table 3 Here]

On the spread side, the empirical targets are guided by the composite U.S. Baa–Aaa spread in Chen, Collin-Dufresne and Goldstein (2009), who report a mean of about 109 bps and a standard deviation of about 41 bps over a long sample, and by the BBB-level spreads for U.S. and other developed markets in recent work by Huang, Nozawa and Shi (2025), which are in the 150–170 bps range. These studies motivate the CS data band of 109–166 bps for the mean and 35–41 bps for the standard deviation in Table 3. Our calibration yields a model-implied mean spread of 121 bps and a spread standard deviation of 36 bps. Both lie squarely inside the empirically plausible bands, indicating that the W2P model reproduces the level and unconditional volatility of IG credit spreads without resorting to extreme parameter values.

For physical default probabilities, we rely on long-run Moody's (2024) and Standard & Poors (2024) corporate default statistics, as well as the Baa default experience in Chen, Collin-Dufresne and Goldstein (2009). These sources indicate that IG issuers default rarely in physical terms: average annual default rates are on the order of 0.1–0.4% (10–40 bps), with spikes to 0.5–1.0% in systemic downturns. This motivates the PD data band of 20–40 bps for the mean and 30–60 bps for the standard deviation. The W2P calibration targets a mean physical default probability of 28 bps with a standard deviation of 47 bps, which again fall comfortably inside the empirical ranges and are consistent with an IG, non-distressed portfolio.

Following the empirical spread-decomposition literature, we define the credit risk premium as the non-default component of the spread: the credit spread net of actuarial expected losses. Specifically, we measure expected losses using physical default probabilities and loss given default, and define

$$\mathrm{CRP}_t \equiv \mathrm{CS}_t - (1-\delta)\,\mathrm{PD}_t,$$

where we use a recovery rate $\delta = 40\%$ (loss given default of 60%) in both data and model. This CRP_t measure is the model counterpart of residual spread components used in the literature to capture priced risk in credit markets (e.g., see Gilchrist and Zakrajšek (2012), Feldhütter and Schaefer (2018) and Kang and Pflueger (2015)). In our framework in Section 4, CRP_t is also the expected excess return on newly originated loans, making it the natural model measure to compare to proxies for the priced-risk component of corporate spreads—the empirical credit risk premium. Benchmark variance decompositions show that both expected credit losses and expected excess returns contribute materially to the cross-sectional variance of credit spreads (Nozawa (2017)). Consistent with our mechanism, σ_w shocks in the model primarily operate through borrower incentives and default expectations (the loss component), while σ_a shocks load more directly on the discount-rate (CRP)

component.

The empirical range for CRP in Table 3 combines the spread moments above with the fact that expected default losses explain only a minority of IG spread levels. Chen, Collin-Dufresne and Goldstein (2009) and Elton, Gruber, Agrawal and Mann (2001) show that, for Baa-rated bonds, expected losses account for at most one quarter of observed spreads, leaving a residual non-default component of roughly 80–130 bps. Recent structural evidence also supports sizable non-default components in investment-grade spreads: Huang, Nozawa and Shi (2025) decompose global corporate bond spreads into a pure default-model credit component and an illiquidity component induced by secondary-market search frictions. Dealer inventory capacity is also a key driver of corporate bond liquidity and expected returns (Goldberg and Nozawa (2021)).[11] Similar conclusions emerge from studies that decompose spreads into expected-loss and risk-premium components (see, for example, Nozawa (2017) and Feldhütter and Schaefer (2018)). Motivated by this literature, we treat 80–130 bps as a plausible CRP mean band and 30–40 bps as a conservative band for its unconditional standard deviation, reflecting that in benchmark decompositions the risk premium often accounts for a substantial share of the spread variance for IG credit.

The model generates an average CRP of 104 bps, almost exactly in the middle of the 80–130 bps data range. This is a key quantitative outcome: the model simultaneously matches the level of IG spreads and the level of the non-default component that underlies the credit spread puzzle. The model-implied standard deviation of CRP is about 23 bps, below the 29–40 bps range implied by spread decompositions, but still of the same order of magnitude. In other words, the W2P framework produces a time-varying credit risk premium that is non-trivial, but somewhat smoother than in the data. This is a common pattern in structural models, and we view it as a reasonable compromise given the tight joint targets for spreads and physical default rates.

[11] Haddad, Moreira and Muir (2021) document that stressed credit markets can feature large non-default (liquidity) dislocations in investment-grade pricing that reverse after policy backstops.

Two features are worth highlighting from the perspective of the credit spread puzzle. First, the W2P model preserves micro-foundations throughout. Credit spreads, default probabilities, and risk premia are not calibrated directly as reduced-form processes; instead they are endogenous outcomes of optimal contracting between households, banks, and firms in an endowment economy with Epstein–Zin preferences, noisy signals of willingness to pay, and a state-contingent underwriting cutoff. Compared to reduced-form affine term structure or intensity models, this structure allows us to discipline the joint behavior of CS, PD, and CRP using a small set of economically interpretable parameters, such as the dispersion of A2P and W2P and the signal-to-noise ratio.

Second, the W2P channel—captured by the cross-sectional dispersion of willingness to pay, $\sigma_{w,t}$, and banks' posterior beliefs—plays a central role in resolving the credit spread puzzle. In standard models, generating a 100 bps non-default component of IG spreads typically requires either implausibly high physical default probabilities or extreme risk aversion. In our framework, the same magnitude of CRP arises because lenders know that some borrowers have high W2P (non-pecuniary costs of default, reputational concerns, legal frictions), but observe only a noisy signal of this latent trait. As we show elsewhere in Section 5, increases in the signal-to-noise ratio Λ_t or in the dispersion $\sigma_{w,t}$ allow banks to screen more effectively on W2P and expand credit while simultaneously lowering average spreads. The levels in Table 3 show that, even after allowing spreads to decline in response to stronger W2P-based screening, the model still delivers an average CRP in the 100 bps range documented in the data.

In all, the results in Table 3 support two claims. First, the W2P model matches the key unconditional moments of credit spreads, default probabilities, and credit risk premia that underlie the credit spread puzzle. Second, it does so with explicit micro-foundations built around cross-sectional W2P heterogeneity and informational frictions, rather than by imposing a reduced-form risk-premium process. This suggests that W2P—and in particular the interaction between $\sigma_{w,t}$, the signal-to-noise ratio, and banks' screening rule—is a

promising structural channel for explaining why credit spreads are persistently high relative to expected default losses.

6 Conclusion

Episodes in which corporate credit looks exceptionally safe just before default risk materializes are a recurring feature of modern markets. This paper develops a dynamic asset-pricing framework that helps make sense of such patterns. We distinguish firms' willingness to pay (W2P) from their ability to pay (A2P) and model how lenders, operating under financial information frictions, use noisy signals about W2P to screen borrowers and set credit spreads. In equilibrium, W2P, A2P and information quality are state variables that jointly shape risk-neutral and physical default probabilities, expected returns and the cross section of credit spreads.

Our quantitative results highlight a striking and unconventional mechanism. In the model, modest increases in the cross-sectional volatility of firms' W2P compress credit spreads at origination, even as unconditional borrower default risk rises, while underwriting becomes more selective so that both physical and risk-neutral default probabilities of new originations fall. By contrast, increases in the dispersion of A2P generate the familiar pattern of wider spreads and higher expected default losses. The W2P channel therefore produces large, persistent movements in credit spreads with relatively modest effects on excess returns, whereas A2P shocks primarily move the risk premium in a standard way. Despite its parsimonious structure, the model matches key moments of corporate credit markets—average spreads, default rates and credit risk premia—without resorting to extreme assumptions about preferences or disaster risk. More broadly, our analysis shows that decomposing credit risk into W2P and A2P, and modeling how information about W2P enters the stochastic discount factor through state-price-weighted probabilities, offers a useful lens on puzzling combinations of low spreads and high realized defaults.

Appendix

A Relaxing independence: correlated A2P and W2P

Starting from the joint normal distribution (6), we derive the conditional distributions and show how correlation affects conditional SNR and screening (see, for example, Christoffersen, Jacobs, Jin and Langlois (2018)).

A.1 Posterior of W2P given (a, s)

When (a_i, w_i) are jointly Gaussian with correlation ρ_{aw}, the conditional distribution of w_i given a_i is

$$w_i \mid a_i \sim \mathcal{N}\Big(\bar{w} + \rho_{aw}\tfrac{\sigma_{w,t}}{\sigma_{a,t}}(a_i - \bar{a}),\ \ \sigma^2_{w|a,t}\Big), \qquad \sigma^2_{w|a,t} = (\sigma_{w,t})^2(1-\rho^2_{aw}).$$

Updating with $s_{i,t} = w_i + \nu_{i,t}$ and $\nu_{i,t} \sim \mathcal{N}(0, \sigma^2_\nu)$ yields the posterior

$$w_i \mid (a_i, s_{i,t}) \sim \mathcal{N}\Big(\mu_t(a_i, s_{i,t}),\ \sigma^2_{w|a,s,t}\Big),$$

with

$$\mu_t(a_i, s) = \bar{w} + \rho_{aw}\tfrac{\sigma_{w,t}}{\sigma_{a,t}}(a_i - \bar{a}) + \kappa^{(a)}_t\left[s - \left(\bar{w} + \rho_{aw}\tfrac{\sigma_{w,t}}{\sigma_{a,t}}(a_i - \bar{a})\right)\right],$$

and

$$\sigma^2_{w|a,s,t} = \frac{\sigma^2_{w|a,t}\,\sigma^2_\nu}{\sigma^2_{w|a,t} + \sigma^2_\nu}, \qquad \kappa^{(a)}_t = \frac{\sigma^2_{w|a,t}}{\sigma^2_{w|a,t} + \sigma^2_\nu}.$$

Defining the conditional SNR

$$\Lambda^{(a)}_t \equiv \frac{\sigma^2_{w|a,t}}{\sigma^2_\nu} = \frac{(\sigma_{w,t})^2(1-\rho^2_{aw})}{\sigma^2_\nu},$$

we obtain

$$\kappa^{(a)}_t = \frac{\Lambda^{(a)}_t}{1 + \Lambda^{(a)}_t}, \qquad \sigma^2_{w|a,s,t} = \frac{\sigma^2_{w|a,t}}{1 + \Lambda^{(a)}_t}.$$

Positive correlation $\rho_{aw} > 0$ reduces conditional dispersion in W2P and lowers $\Lambda^{(a)}_t$ relative to the independence case. Negative correlation $\rho_{aw} < 0$ raises conditional dispersion and increases $\Lambda^{(a)}_t$. Thus the correlation parameter ρ_{aw} provides an additional lever for the strength of the W2P screening channel, over and above the unconditional dispersions $(\sigma_{a,t}, \sigma_{w,t})$ and the signal noise variance σ^2_ν.

A.2 Implications for default risk and credit spreads

Given the posterior distribution of w_i conditional on $(a_i, s_{i,t})$, the default probability and pricing formulas (15)–(23) carry over with (κ_t, Λ_t) replaced by their conditional counterparts $(\kappa_t^{(a)}, \Lambda_t^{(a)})$. In financial markets with strong positive correlation $\rho_{aw} > 0$, conditional SNR is lower and the W2P signal adds less information beyond A2P, so the incremental impact of W2P dispersion on spreads is muted. In financial markets with weak or negative correlation, conditional SNR is higher and W2P dispersion has a large effect on screening, spreads, and the credit-spread puzzle.

B Recursive preferences and the SDF

Under Epstein and Zin (1989) preferences (16) we define the certainty-equivalent continuation value

$$\mathcal{R}_t(U_{t+1}) \equiv \left(\mathbb{E}_t\left[U_{t+1}^{1-\gamma}\right]\right)^{\frac{1}{1-\gamma}}.$$

The recursion can be written compactly as

$$U_t^{1-1/\psi} = (1-\beta)C_t^{1-1/\psi} + \beta\,\mathcal{R}_t(U_{t+1})^{\,1-1/\psi}.$$

Let $\theta \equiv \frac{1-\gamma}{1-1/\psi}$. The one-period stochastic discount factor that prices $t{+}1$ payoffs satisfies the standard Epstein–Zin form

$$M_{t+1} = \beta^\theta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{1}{\psi}\theta} \left(\frac{U_{t+1}}{\mathcal{R}_t(U_{t+1})}\right)^{\theta-1}, \tag{36}$$

where the ratio $U_{t+1}/\mathcal{R}_t(U_{t+1})$ captures shocks to continuation utility relative to its certainty equivalent.

In the numerical solution, (C_t, U_t) are included among the model's controls and state variables solved by perturbation. Along simulated paths, we evaluate (36) using the simulated (C_t, C_{t+1}, U_{t+1}) and the corresponding certainty-equivalent aggregator $\mathcal{R}_t(U_{t+1})$. We then obtain the risk-free rate from $1/R_t^f = \mathbb{E}t[Mt+1]$ and compute state-price-weighted and risk-neutral quantities from the corresponding M_{t+1}-weighted conditional expectations.

C Posterior beliefs, SNR, and approval rule

We provide additional details on the Gaussian updating and on how firm-level default probabilities aggregate into the pool objects λ_t and PD_t. Starting from the prior $w_i \sim \mathcal{N}(\bar{w}, (\sigma_{w,t})^2)$ and the signal $s_{i,t} = w_i + \nu_{i,t}$ with $\nu_{i,t} \sim \mathcal{N}(0, \sigma_\nu^2)$, Gaussian conjugacy delivers (8)–(12). The signal-to-noise ratio $\Lambda_t = (\sigma_{w,t})^2/\sigma_\nu^2$ controls both the slope of the posterior mean and the tightness of the posterior variance.

Given (S_t, Λ_t), the cross-sectional density of signals is $g_t(s)$, and the acceptance rate and pool default probability are

$$\lambda_t = \int_{s_t^*}^{\infty} g_t(s)\, ds, \tag{37}$$

$$\mathrm{PD}_t = \frac{\displaystyle\int_{s_t^*}^{\infty} \mathbb{E}_t\big[\mathrm{PD}_t(a_i, s; R_{L,t}) \,\big|\, s_{i,t} = s\big] g_t(s)\, ds}{\lambda_t}. \tag{38}$$

The state-price-weighted counterpart of the pool default probability is

$$\Pi_t^{*,\mathrm{pool}} = \frac{\displaystyle\int_{s_t^*}^{\infty} \mathbb{E}_t\big[\Pi_t^*(a_i, s) \,\big|\, s_{i,t} = s\big] g_t(s)\, ds}{\lambda_t}.$$

The target condition (29) can be written compactly as $\Pi_t^{*,\mathrm{pool}} = \bar{\Pi}$. The comparative statics in Proposition 1 are obtained by differentiating this condition with respect to (s_t^*, Λ_t) around the steady state and using the Gaussian structure to sign the derivatives.

D Spread decomposition and probability measures

This appendix summarizes the probability measures and pricing objects used to decompose credit spreads into an expected-loss component and a priced-risk component. Let $\mathbb{P}_t(\cdot)$ denote physical probabilities conditional on time-t information, and let M_{t+1} be the one-period stochastic discount factor. Define the one-step risk-neutral expectation operator $\mathbb{E}_t^{\mathbb{Q}}[\cdot]$ by

$$\mathbb{E}_t^{\mathbb{Q}}[X_{t+1}] \;\equiv\; \frac{\mathbb{E}_t[M_{t+1}\, X_{t+1}]}{\mathbb{E}_t[M_{t+1}]}. \tag{39}$$

Since $1/R_t^f = \mathbb{E}_t[M_{t+1}]$, any payoff X_{t+1} with time-t price equal to one satisfies

$$1 = \mathbb{E}_t[M_{t+1}X_{t+1}] \quad \Longleftrightarrow \quad 1 = \frac{1}{R_t^f}\,\mathbb{E}_t^{\mathbb{Q}}[X_{t+1}]\,.$$

Let $D_{t+1} \in \{0,1\}$ be the default indicator for the cohort of loans originated at t. The physical and risk-neutral default probabilities are

$$\mathrm{PD}_t \equiv \mathbb{E}_t[D_{t+1}]\,, \qquad \mathrm{PD}_t^{\mathbb{Q}} \equiv \mathbb{E}_t^{\mathbb{Q}}[D_{t+1}]\,.$$

We also use the state-price-weighted default probability

$$\Pi_t^* \;\equiv\; \mathbb{E}_t[M_{t+1}\,D_{t+1}]\,, \tag{40}$$

which relates to $\mathrm{PD}_t^{\mathbb{Q}}$ through

$$\mathrm{PD}_t^{\mathbb{Q}} \;=\; \frac{\Pi_t^*}{\mathbb{E}_t[M_{t+1}]} \;=\; R_t^f\,\Pi_t^*.$$

When conditioning on the accepted pool, we use the corresponding pool objects (e.g., $\mathrm{PD}_t^{\mathrm{pool}}$, $\mathrm{PD}_t^{\mathbb{Q},\mathrm{pool}}$, and $\Pi_t^{*,\mathrm{pool}}$) defined analogously.

Applying the pricing identity to the loan payoff (30) implies

$$1 = \frac{1}{R_t^f}\,\mathbb{E}_t^{\mathbb{Q}}\left[\widetilde{R}_{L,t+1}\right] = \frac{1}{R_t^f}\left[1-(1-\delta)\mathrm{PD}_t^{\mathbb{Q}}\right] R_{L,t},$$

and therefore

$$\frac{R_{L,t}}{R_t^f} = \frac{1}{1-(1-\delta)\mathrm{PD}_t^{\mathbb{Q}}}.$$

Equivalently, using Π_t^*,

$$\frac{R_{L,t}}{R_t^f} = \frac{1}{1-(1-\delta)R_t^f\,\Pi_t^*}.$$

For small default probabilities,

$$R_{L,t} - R_t^f \;\approx\; R_t^f\,(1-\delta)\,\mathrm{PD}_t^{\mathbb{Q}}.$$

This expression highlights that the credit spread is controlled by the risk-neutral default probability, not by the physical PD_t. Related empirical decompositions separate expected losses from discount-rate components in credit spreads (e.g., Nozawa (2017)). The gap $\mathrm{PD}_t^{\mathbb{Q}} - \mathrm{PD}_t$ captures the contribution of risk premia and of W2P-driven underwriting-based selection

to spreads. In the quantitative implementation we compute both PD_t and $\mathrm{PD}_t^{\mathbb{Q}}$, as well as realized excess returns on credit portfolios, to compare the model-implied decomposition with the empirical credit-spread puzzle.

E State-space representation

We summarize the state-space objects used to generate the quantitative results in Section 5. The aggregate state vector is

$$S_t \equiv \big(\log Y_t,\ \sigma_{a,t},\ \sigma_{w,t},\ \Lambda_t,\ U_t\big),$$

together with any auxiliary states required by the higher-order approximation.

The control variables include consumption C_t and the portfolio shares in the risk-free asset and the credit asset. Given S_t, market clearing and the recursive-preference recursion determine C_t and the continuation-value update for U_t, which together imply the one-period stochastic discount factor M_{t+1} (and hence the risk-free rate).

Conditional on S_t (and the pricing objects implied by M_{t+1}), lenders update posterior beliefs from the signal, solve for the underwriting cutoff s_t^*, and determine the loan rate and acceptance rate. Cross-sectional moments of the accepted pool (e.g., physical and risk-neutral default probabilities and the implied spread) are computed by integrating over the cross-sectional distribution implied by the Gaussian prior and signal noise.

We solve the model using a fourth-order perturbation around the deterministic steady state and compute impulse responses to the exogenous innovations $\varepsilon_{\sigma_w,t}$, $\varepsilon_{\sigma_a,t}$, and $\varepsilon_{\Lambda,t}$. Reported spread and return objects are annualized in basis points.

F Proof of Proposition 1

For clarity, we suppress the time index t and work at a fixed aggregate state (Y, σ_a, σ_w) and a fixed pricing kernel (equivalently, fixed R_f and fixed state prices). We distinguish two comparative-statics objects. First, a change in the signal-to-noise ratio Λ that reflects *information quality* (e.g., a change in σ_ν^2 holding the cross-sectional distribution of types fixed). Second, a change in the W2P dispersion σ_w, which changes the *population* distribution of types and (with σ_ν fixed) also changes $\Lambda = \sigma_w^2/\sigma_\nu^2$. Parts (i) and (ii) are comparative statics

in Λ holding the population distribution fixed, and part (iii) is a comparative static in σ_w holding σ_ν fixed.

F.1 Setup

Let $f(s)$ denote the cross-sectional density of signals s in the borrower population, and let $p(s;\Lambda)$ denote the state-price-weighted default probability of a firm with signal s under SNR Λ. Under the Gaussian structure (6)–(7) and the posterior (8)–(12), $f(s)$ is strictly positive and smooth on $\mathbb{R}$. $p(s;\Lambda)$ is strictly decreasing in s (better signals mean higher posterior W2P and lower default risk). For each s, $p(s;\Lambda)$ is strictly decreasing in Λ (higher SNR shrinks the posterior variance and tilts the posterior toward higher W2P conditional on a given s).

Banks choose a cutoff s^* and approve all borrowers with $s \geq s^*$. The acceptance rate and pool state-price-weighted default probability are

$$\lambda(s^*) = \int_{s^*}^{\infty} f(s)\, ds, \tag{41}$$

$$\Pi^{*,\text{pool}}(s^*,\Lambda) = \frac{\displaystyle\int_{s^*}^{\infty} p(s;\Lambda) f(s)\, ds}{\displaystyle\int_{s^*}^{\infty} f(s)\, ds}. \tag{42}$$

The target-risk constraint is

$$\Pi^{*,\text{pool}}(s^*,\Lambda) = \bar{\Pi}. \tag{43}$$

F.2 Monotone likelihood ratio and comparative statics

Under the Gaussian structure, the posterior $w \mid s$ satisfies a monotone likelihood ratio property: the likelihood ratio of high vs. low W2P types is increasing in s. This implies that $p(s;\Lambda)$ is strictly decreasing in s and that the s-to-posterior-default-probability relationship becomes more sensitive as Λ increases. Differentiating $p(s;\Lambda)$ with respect to Λ yields

$$\frac{\partial p(s;\Lambda)}{\partial \Lambda} = \phi(s;\Lambda)\, \frac{\partial \tau(s;R_L)}{\partial \Lambda} + \psi(s;\Lambda)\, \frac{\partial \kappa(\Lambda)}{\partial \Lambda},$$

where ϕ and ψ are smooth functions with $\psi(s;\Lambda) < 0$; under the maintained Gaussian benchmark, the κ-channel dominates so that $\partial p(s;\Lambda)/\partial \Lambda < 0$ in the relevant region.[12]

Holding the population density $f(s)$ fixed, differentiate the target constraint (43) with

[12] Intuitively, higher SNR raises $\kappa(\Lambda)$ and shrinks the posterior variance, shifting mass away from low-W2P realizations and lowering default probabilities at any given signal s.

respect to (s^*, Λ). The implicit function theorem yields

$$\frac{ds^*}{d\Lambda} = -\frac{\dfrac{\partial \Pi^{*,\text{pool}}}{\partial \Lambda}}{\dfrac{\partial \Pi^{*,\text{pool}}}{\partial s^*}}. \tag{44}$$

The denominator is

$$\frac{\partial \Pi^{*,\text{pool}}}{\partial s^*} = -\frac{f(s^*)}{\lambda(s^*)} \left[p(s^*; \Lambda) - \Pi^{*,\text{pool}}(s^*, \Lambda) \right] < 0,$$

because $f(s^*) > 0$ and $p(s^*; \Lambda) > \Pi^{*,\text{pool}}(s^*, \Lambda)$: the marginal borrower at the cutoff has strictly higher default probability than the average borrower in the approved pool. The numerator is non-positive under the maintained information-quality comparative static (i.e., higher Λ weakly reduces default risk at a given signal)

$$\frac{\partial \Pi^{*,\text{pool}}}{\partial \Lambda} = \frac{1}{\lambda(s^*)} \int_{s^*}^{\infty} \frac{\partial p(s; \Lambda)}{\partial \Lambda} f(s)\, ds \leq 0,$$

since $\partial p(s; \Lambda)/\partial \Lambda \leq 0$ in the relevant region. Combining these signs with (44) yields

$$\frac{ds^*}{d\Lambda} \leq 0,$$

with strict inequality whenever the density of marginal borrowers at the cutoff is positive. This proves part (i) of Proposition 1.

For part (ii), note that

$$\frac{d\lambda}{d\Lambda} = \frac{\partial \lambda}{\partial s^*} \frac{ds^*}{d\Lambda} = -f(s^*) \frac{ds^*}{d\Lambda} \geq 0,$$

since $f(s^*) > 0$ and $ds^*/d\Lambda \leq 0$. Thus the acceptance rate is weakly increasing in SNR.

F.3 Population default probability

Part (iii) has two statements. First, (43) directly implies that the pool state-price-weighted default probability is determined by the target, $\Pi^{*,\text{pool}}(s^*, \Lambda) = \bar{\Pi}$. Second, define the physical *population* (unconditional, pre-underwriting) default probability as the cross-sectional average of the physical default probability (15) across all applicants (integrating over the supports of $g(a)$ and $f(s)$):

$$\text{PD}_t^{\text{pop}} \equiv \int \int \text{PD}_t(a, s; R_{L,t})\, g(a)\, f(s)\, da\, ds, \tag{45}$$

where $g(a)$ is the cross-sectional density of a in the population and $f(s)$ is the cross-sectional density of signals.

Given $(Y_t, \sigma_{a,t})$, consider a small increase in $\sigma_{w,t}$ holding σ_ν^2 fixed. For a given (a, s) and contract schedule $R_{L,t}$, the default rule is characterized by the threshold (14), and the physical default probability is (15). Writing

$$z(a, s) \equiv \frac{\tau(a, s; R_{L,t}) - \mu_t(s)}{\sigma_{w|s,t}},$$

a direct differentiation of (15) with respect to the posterior standard deviation $\sigma_{w|s,t}$ gives

$$\frac{\partial \mathrm{PD}_t(a, s; R_{L,t})}{\partial \sigma_{w|s,t}} = \varphi\big(z(a, s)\big) \, \frac{\mu_t(s) - \tau(a, s; R_{L,t})}{\sigma_{w|s,t}^2}, \tag{46}$$

where $\varphi(\cdot)$ is the standard normal density. In the empirically relevant region of small default probabilities (a left-tail event), $\tau(a, s; R_{L,t}) < \mu_t(s)$, so the right-hand side of (46) is strictly positive.

Moreover, holding σ_ν^2 fixed, the posterior dispersion $\sigma_{w|s,t}$ is strictly increasing in $\sigma_{w,t}$ under (8)–(12) (equivalently, $\sigma_{w|s,t}^2 = \sigma_{w,t}^2 \sigma_\nu^2 / (\sigma_{w,t}^2 + \sigma_\nu^2)$). Therefore, in this left-tail region, $\mathrm{PD}_t(a, s; R_{L,t})$ is increasing in $\sigma_{w,t}$ for each (a, s). Integrating over (a, s) in (45) implies that $\mathrm{PD}_t^{\mathrm{pop}}$ rises with $\sigma_{w,t}$.

Finally, since parts (i) and (ii) show that higher Λ weakly lowers s^* and raises λ, banks can expand credit while maintaining the target pool risk $\bar{\Pi}$. □

References

[1] Albagli, Elias, Hellwig, Christian and Aleh Tsyvinski. 2024. Information Aggregation with Asymmetric Asset Payoffs. Journal of Finance, 79, 4, 2715–2758.

[2] Augustin, Patrick and Yehuda Izhakian. 2020. Ambiguity, Volatility, and Credit Risk. Review of Financial Studies, 33, 4, 1618–1672.

[3] Bai, Jennie, Goldstein, Robert and Fan Yang. 2020. Is the Credit Spread Puzzle a Myth? Journal of Financial Economics, 137, 2, 297–319.

[4] Bakshi, Gurdip, Gao, Xiaohui and Zhaodong Zhong. 2022. Decoding Default Risk: A Review of Modeling Approaches, Findings, and Estimation Methods. Annual Review of Financial Economics, 14, 391–413.

[5] Bansal, Ravi and Amir Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. Journal of Finance, 59, 4, 1481–1509.

[6] Benmelech, Efraim, Kumar, Nitish and Raghuram Rajan. 2022. The secured credit premium and the issuance of secured debt. Journal of Financial Economics, 146, 1, 143–171.

[7] Berlin, Mitchell, Nini, Greg and Edison Yu. 2020. Concentration of Control Rights in Leveraged Loan Syndicates. Journal of Financial Economics, 137, 1, 249–271.

[8] Campbell, John and John Cochrane. 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. Journal of Political Economy, 107, 2, 205–251.

[9] Campbell, John, Lettau, Martin, Malkiel, Burton and Yexiao Xu. 2001. Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk. Journal of Finance, 56, 1, 1–43.

[10] Campello, Murillo, Ladika, Tomislav and Rafael Matta. 2019. Renegotiation Frictions and Financial Distress Resolution: Evidence from CDS Spreads. Review of Finance, 23, 3, 513–556.

[11] Cao, Jie, Goyal, Amit, Xiao, Xiao and Xintong Zhan. 2023. Implied Volatility Changes and Corporate Bond Returns. Management Science, 69, 3, 1375–1397.

[12] Carvalho, Daniel, Gao, Janet and Pengfei Ma. 2023. Loan spreads and credit cycles: The role of lenders' personal economic experiences. Journal of Financial Economics, 148, 118–149.

[13] Chen, Hui. 2010. Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure. Journal of Finance, 65, 6, 2171–2212.

[14] Chen, Long, Collin-Dufresne, Pierre and Robert Goldstein. 2009. On the Relation Between the Credit Spread Puzzle and the Equity Premium Puzzle. Review of Financial Studies, 22, 9, 3367–3409.

[15] Chen, Qian, Li, Yang, Zhao, Qiuyun and Fan Liu. 2024. Information Quality and Cost of Credit Bond Financing. Finance Research Letters, 59, 104708.

[16] Chung, Kee, Wang, Junbo and Chunchi Wu. 2019. Volatility and the Cross-Section of Corporate Bond Returns. Journal of Financial Economics, 133, 2, 397–417.

[17] Christoffersen, Peter, Jacobs, Kris, Jin, Xisong and Hugues Langlois. 2018. Dynamic Dependence and Diversification in Corporate Credit. Review of Finance, 22, 2, 521–560.

[18] Cici, Gjergji, Gibson, Scott and John Merrick. 2011. Missing the marks? Dispersion in corporate bond valuations across mutual funds. Journal of Financial Economics, 101, 1, 206–226.

[19] Darmouni, Olivier. 2020. Informational Frictions and the Credit Crunch. Journal of Finance, 75, 4, 2055–2094.

[20] Davydiuk, Tetiana, Marchuk, Tatyana and Samuel Rosen. 2024. Market Discipline in the Direct Lending Space. Review of Financial Studies, 37, 4, 1190–1264.

[21] Davydenko, Sergei and Ilya Strebulaev. 2007. Strategic Actions and Credit Risk: Evidence from Debt Restructurings. Journal of Finance, 62, 6, 2633–2671.

[22] Dickerson, Adam, Mueller, Philipp and Cesare Robotti. 2023. Priced Risk in Corporate Bond Returns. Journal of Financial Economics, 148, 3, 676–706.

[23] Doshi, Hitesh, Ericsson, Jan, Jacobs, Kris and Stuart Turnbull. 2013. Pricing Credit Default Swaps with Observable Covariates. Review of Financial Studies, 26, 8, 2049–2094.

[24] Du, Du, Elkamhi, Redouane and Jan Ericsson. 2019. Time-Varying Asset Volatility and the Credit Spread Puzzle. Journal of Finance, 74, 6, 2589–2639.

[25] Duffie, Darrell and Kenneth Singleton. 1999. Modeling Term Structures of Defaultable Bonds. Review of Financial Studies, 12, 4, 687–720.

[26] Eisfeldt, Andrea, Herskovic, Bernard and Shuo Liu. 2024. Interdealer Price Dispersion and Intermediary Capacity. Working Paper.

[27] Elton, Edwin, Gruber, Martin, Agrawal, Deepak and Christopher Mann. 2001. Explaining the Rate Spread on Corporate Bonds. Journal of Finance, 56, 1, 247–277.

[28] Epstein, Larry, and Stanley Zin. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. Econometrica,

57, 4, 937–969.

[29] Favara, Giovanni, Schroth, Enrique and Philip Valta. 2012. Strategic Default and Equity Risk Across Countries. Journal of Finance, 67, 6, 2051–2095.

[30] Feldhütter, Peter and Stephen Schaefer. 2018. The Myth of the Credit Spread Puzzle. Review of Financial Studies, 31, 8, 2897–2942.

[31] Fuster, Andreas, Plosser, Matthew, Schnabl, Philipp and James Vickery. 2019. The Role of Technology in Mortgage Lending. Review of Financial Studies, 32, 5, 1854–1899.

[32] Ganong, Peter and Pascal Noel. 2023. Why Do Borrowers Default on Mortgages? Quarterly Journal of Economics, 138, 2, 1001–1065.

[33] Gerardi, Kristopher, Herkenhoff, Kyle, Ohanian, Lee and Paul Willen. 2018. Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default. Review of Financial Studies, 31, 3, 1098–1131.

[34] Gilchrist, Simon and Egon Zakrajsek. 2012. Credit Spreads and Business Cycle Fluctuations. American Economic Review, 102, 4, 1692–1720.

[35] Goldberg, Jonathan and Yoshio Nozawa. 2021. Liquidity Supply in the Corporate Bond Market. Journal of Finance, 76, 2, 755–796.

[36] Gomes, Joao, Grotteria, Marco and Jessica Wachter. 2019. Cyclical Dispersion in Expected Defaults. Review of Financial Studies, 32, 4, 1275–1308.

[37] Greenwood, Robin, Hanson, Samuel and Lawrence Jin. 2023. Reflexivity in Credit Markets. Journal of Finance, forthcoming.

[38] Greenwood, Robin, Hanson, Samuel, Shleifer, Andrei and Kasper Sorensen. 2022. Predictable Financial Crises. Journal of Finance, 77, 2, 863–921.

[39] Guiso, Luigi, Sapienza, Paola and Luigi Zingales. 2013. The Determinants of Attitudes toward Strategic Default on Mortgages. Journal of Finance, 68, 4, 1473–1515.

[40] Gupta, Arpit and Christopher Hansman. 2022. Selection, Leverage, and Default in the Mortgage Market. Review of Financial Studies, 35, 2, 720–770.

[41] Haddad, Valentin, Moreira, Alan and Tyler Muir. 2021. When Selling Becomes Viral: Disruptions in Debt Markets in the COVID-19 Crisis and the Fed's Response. Review of Financial Studies, 34, 11, 5309–5351.

[42] He, Zhiguo, Huang, Jing and Jidong Zhou. 2023. Open Banking: Credit Market Competition When Borrowers Own the Data. Journal of Financial Economics, 147, 2, 449–474.

[43] He, Zhiguo, Bryan Kelly and Asaf Manela. 2017. Intermediary Asset Pricing: New Evidence from Many Asset Classes. Journal of Financial Economics, 126, 1, 1–35.

[44] He, Zhiguo and Arvind Krishnamurthy. 2013. Intermediary Asset Pricing. American Economic Review, 103, 2, 732–770.

[45] Herskovic, Bernard, Kelly, Bryan, Lustig, Hanno and Stijn Van Nieuwerburgh. 2016. The common factor in idiosyncratic volatility: Quantitative asset pricing implications. Journal of Financial Economics, 119, 2, 249–283.

[46] Herskovic, Bernard, Kind, Thilo and Howard Kung. 2023. Micro Uncertainty and Asset Prices. Journal of Financial Economics, 149, 1, 27–51.

[47] Huang, Jing-Zhi, and Ming Huang. 2012. How Much of the Corporate-Treasury Yield Spread Is Due to Credit Risk? Review of Asset Pricing Studies, 2, 2, 153–202.

[48] Huang, Jing-Zhi, Nozawa, Yoshio and Zhan Shi. 2025. The Global Credit Spread Puzzle. Journal of Finance, 80, 1, 101–162.

[49] Ioannidou, Vasso, Pavanini, Nicola and Yushi Peng. 2022. Collateral and Asymmetric Information in Lending Markets. Journal of Financial Economics, 144, 1, 93–121.

[50] Kang, Johnny, and Carolin Pflueger. 2015. Inflation Risk in Corporate Bonds, Journal of Finance, 70, 1, 115–162.

[51] Karolyi, Stephen Adam. 2018. Personal Lending Relationships. Journal of Finance, 73, 1, 5–49.

[52] Kuehn, Lars-Alexander, Schreindorfer, David and Florian Schulz. 2023. Persistent Crises and Levered Asset Prices. Review of Financial Studies, 36, 6, 2571–2616.

[53] Merton, Robert. 1974. On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. Journal of Finance, 29, 2, 449–470.

[54] Moody's Investors Service. 2024. Annual Default Study: Corporate Default and Recovery Rates, 1920–2023.

[55] Nozawa, Yoshio. 2017. What Drives the Cross-Section of Credit Spreads? A Variance Decomposition Approach. Journal of Finance, 72, 5, 2045–2072.

[56] Paravisini, Daniel, Rappoport, Veronica and Philipp Schnabl. 2023. Specialization in Bank Lending: Evidence from Exporting Firms. Journal of Finance, 78, 4, 2049–2085.

[57] Siriwardane, Emil. 2019. Limited Investment Capital and Credit Spreads. Journal of Finance, 74, 5, 2303–2347.

[58] S&P Global Ratings. 2024. Default, Transition, and Recovery: 2024 Annual Global Corporate Default and Rating Transition Study.

Table 1: Variable definitions

We report the definition of the main variables used in the paper.

Symbol	Definition
Y_t	Aggregate endowment (output) at time t
C_t	Aggregate consumption at time t
S_t	Aggregate state, $S_t = (Y_t, \sigma_{a,t}, \sigma_{w,t})$
a_i	Firm i's ability-to-pay (A2P) index
w_i	Firm i's willingness-to-pay (W2P) index
$\sigma_{a,t}$	Cross-sectional volatility (standard deviation across firms) of A2P at t
$\sigma_{w,t}$	Cross-sectional volatility (standard deviation across firms) of W2P at t
ρ_{aw}	Correlation between A2P and W2P
$s_{i,t}$	Noisy W2P signal observed by banks for firm i
$\nu_{i,t}$	Idiosyncratic noise in the W2P signal
Λ_t	Signal-to-noise ratio, $\Lambda_t \equiv \sigma^2_{w,t}/\sigma^2_\nu$
$\mu_t(s)$	Posterior mean of w_i conditional on $s_{i,t} = s$
$\sigma^2_{w\|s,t}$	Posterior variance of w_i conditional on $s_{i,t} = s$
κ_t	Posterior weight on the W2P signal in $\mu_t(s)$
s^*_t	Approval cutoff in signal space
λ_t	Approval/acceptance rate of loans in period t
R^f_t	One-period gross risk-free rate at time t
$R_{L,t}$	Representative promised gross loan rate on loans originated at t
$\widetilde{R}_{L,t+1}$	Realized gross return on the cohort of loans originated at t
D_{t+1}	Cohort-level default indicator for loans originated at t
$\Pi^L_{i,t+1}$	Payoff to the intermediary from firm i's loan at $t+1$
$\bar{\Pi}$	Target state-price-weighted default probability (or loss)
PD_t	Physical default probability of the loan pool
$\mathrm{PD}^{\mathbb{Q}}_t$	Risk-neutral (state-price-weighted) default probability of the pool
M_{t+1}	Stochastic discount factor between t and $t+1$
U_t	Epstein–Zin continuation utility at time t
CS_t	Credit spread on newly originated loans, $R_{L,t} - R^f_t$
CRP_t	Expected excess return (credit risk premium) on loans, $\mathbb{E}_t[\widetilde{R}_{L,t+1}] - R^f_t$

Table 2: Parameters of the Model

We report parameter definitions and benchmark values for the W2P economy developed in Section 4 and quantitatively implemented in Section 5.

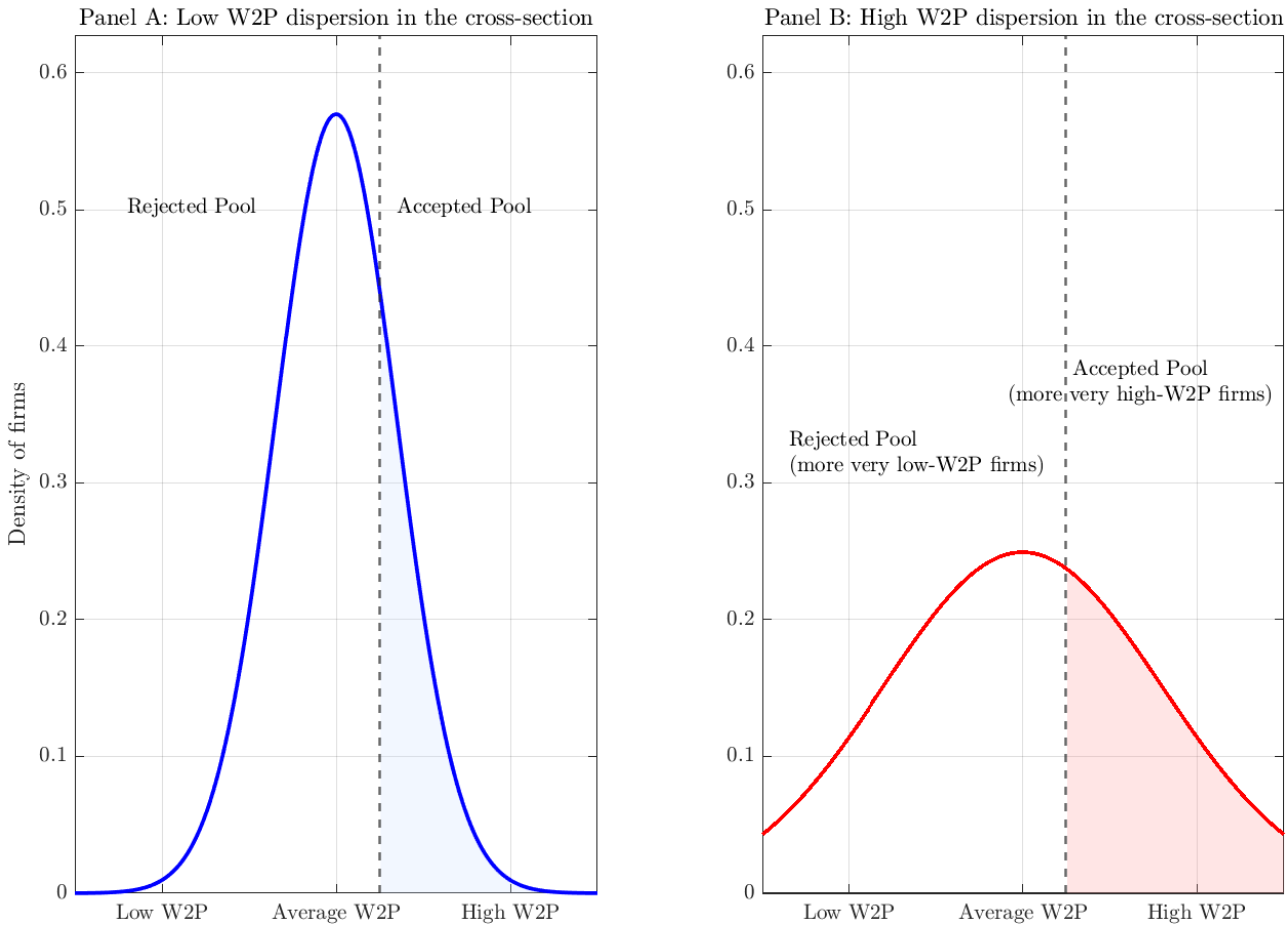
Parameter	Value	Definition
Preferences and recovery		
β	0.99	Intertemporal discount factor
γ	8	Relative risk aversion
ψ	1.5	Elasticity of intertemporal substitution
δ	0.4	Recovery rate
Information		
σ_ν	0.5	Standard deviation of W2P signal noise
Shock processes		
ρ_y	0.9	Persistence of output/endowment shock
ρ_{σ_a}	0.9	Persistence of A2P dispersion shock
ρ_{σ_w}	0.9	Persistence of W2P dispersion shock
σ_y	0.01	Standard deviation of output shock
σ_{σ_a}	0.01	Standard deviation of A2P dispersion shock
σ_{σ_w}	0.01	Standard deviation of W2P dispersion shock
Steady-state real side		
$\bar{y}$	0	Steady-state log output/consumption (normalization, $\bar{Y} = \bar{C} = 1$)
$\bar{\sigma}_a$	0.20	Steady-state dispersion of A2P
$\bar{\sigma}_w$	0.20	Steady-state dispersion of W2P
$\bar{U}$	0	Steady-state continuation utility (normalization)
Steady-state credit market		
$\bar{\Pi}$	0.02	Risk-neutral default probability
$\overline{PD}$	0.003	Physical default probability
$\bar{\lambda}$	0.40	Loan approval rate
$\bar{\Lambda}$	0.16	Steady-state SNR σ_w^2/σ_ν^2
$\bar{\kappa}$	0.138	Steady-state posterior weight on W2P signal

Table 3: Data Historical Moments and Model Simulated Moments

We report the data historical moments (lower and higher bounds) and the model simulated moments for credit spreads (CS), default probability (PD) and credit risk premium (CRP). Column 1 reports the lower end data moments. Column 2 reports the higher end data moments. Column 3 reports the simulated moments using the model in Section 4. For each variable, the first row reports the mean (in annualized basis points) and the second row reports the standard deviation (in annualized basis points). The super index indicates the source of the value for the historical moment: (a) is obtained from Chen, Collin-Dufresne and Goldstein (2009); (b) is obtained from Huang, Nozawa and Shi (2025); (c) is obtained from Chen (2010); (d) is obtained from S&P Global Ratings (2024), "Annual Global Corporate Default and Rating Transition Study"; (e) is obtained from Moody's Investors Service (2024), "Annual Default Study: Corporate Default and Recovery Rates"; (f) is inferred by combining the spread variance decompositions in Nozawa (2017) and Feldhütter and Schaefer (2018) with the spread volatility in Chen, Collin-Dufresne and Goldstein (2009).

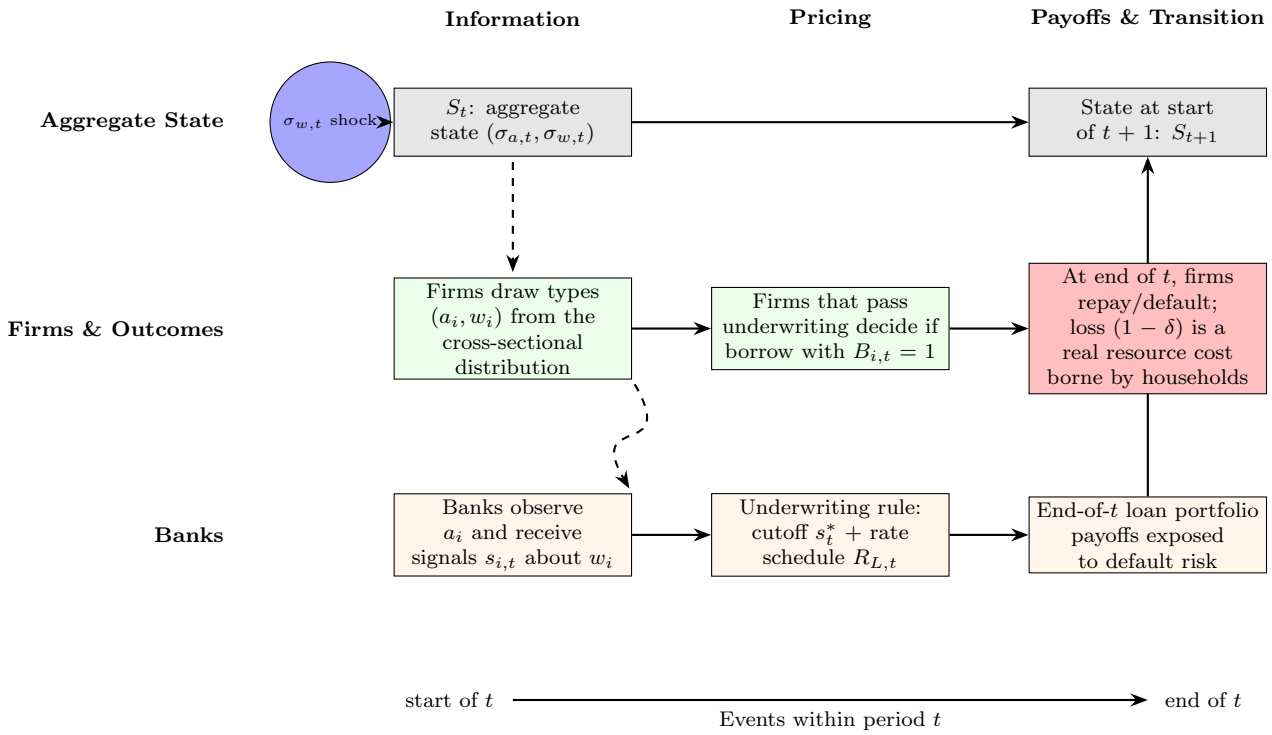
		Data (Low) (1)	Data (High) (2)	Model W2P (3)
Credit Spread (CS)	Mean	$109^{(a)}$	$166^{(b)}$	121
	Std.Dev.	$35^{(c)}$	$41^{(a)}$	36
Default Probability (PD)	Mean	$20^{(d)}$	$40^{(a)}$	28
	Std.Dev.	$30^{(e)}$	$60^{(e)}$	47
Credit Risk Premium (CRP)	Mean	$80^{(a)}$	$130^{(b)}$	104
	Std.Dev.	$29^{(f)}$	$40^{(f)}$	23

Figure 1: Cross Sectional W2P Dispersion and Underwriting



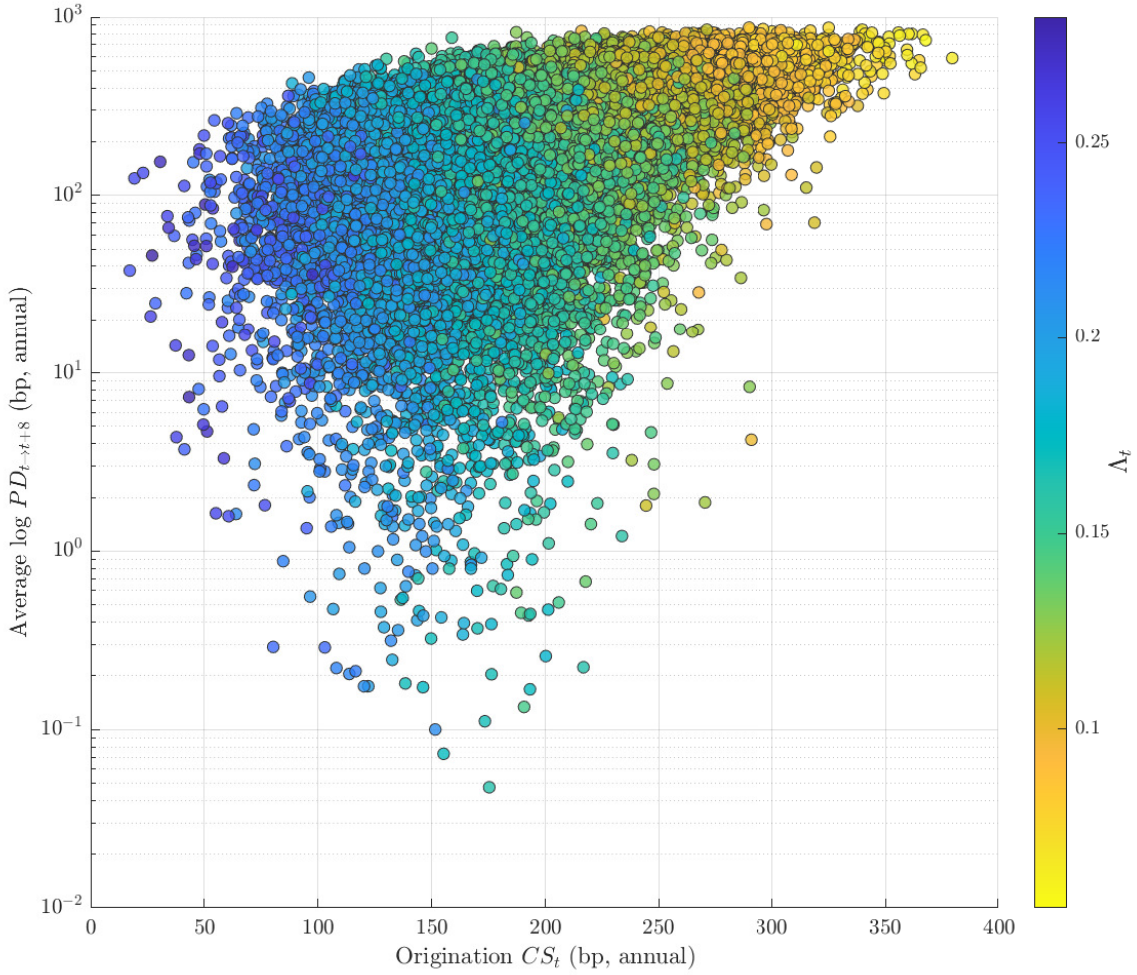
We illustrate the underwriting channel in the partial-equilibrium model by plotting the cross-sectional distribution of firms' latent willingness to pay (W2P) together with a fixed underwriting cutoff. Panel A shows a low-dispersion W2P state: most approved firms are close to the average W2P. Panel B shows a high-dispersion W2P state; with the same underwriting rule, more mass lies in the right tail, underwriting admits more extremely high-W2P borrowers (a higher approved-pool share π_w), the default probability in the accepted pool $\mathrm{PD}_{\mathrm{pool}}(\pi_w)$ falls, and contemporaneous credit spreads compress. The vertical line marks the effective W2P cutoff implied by underwriting, and the shaded region represents accepted firms.

Figure 2: Timing of Information, Underwriting, and Payoffs



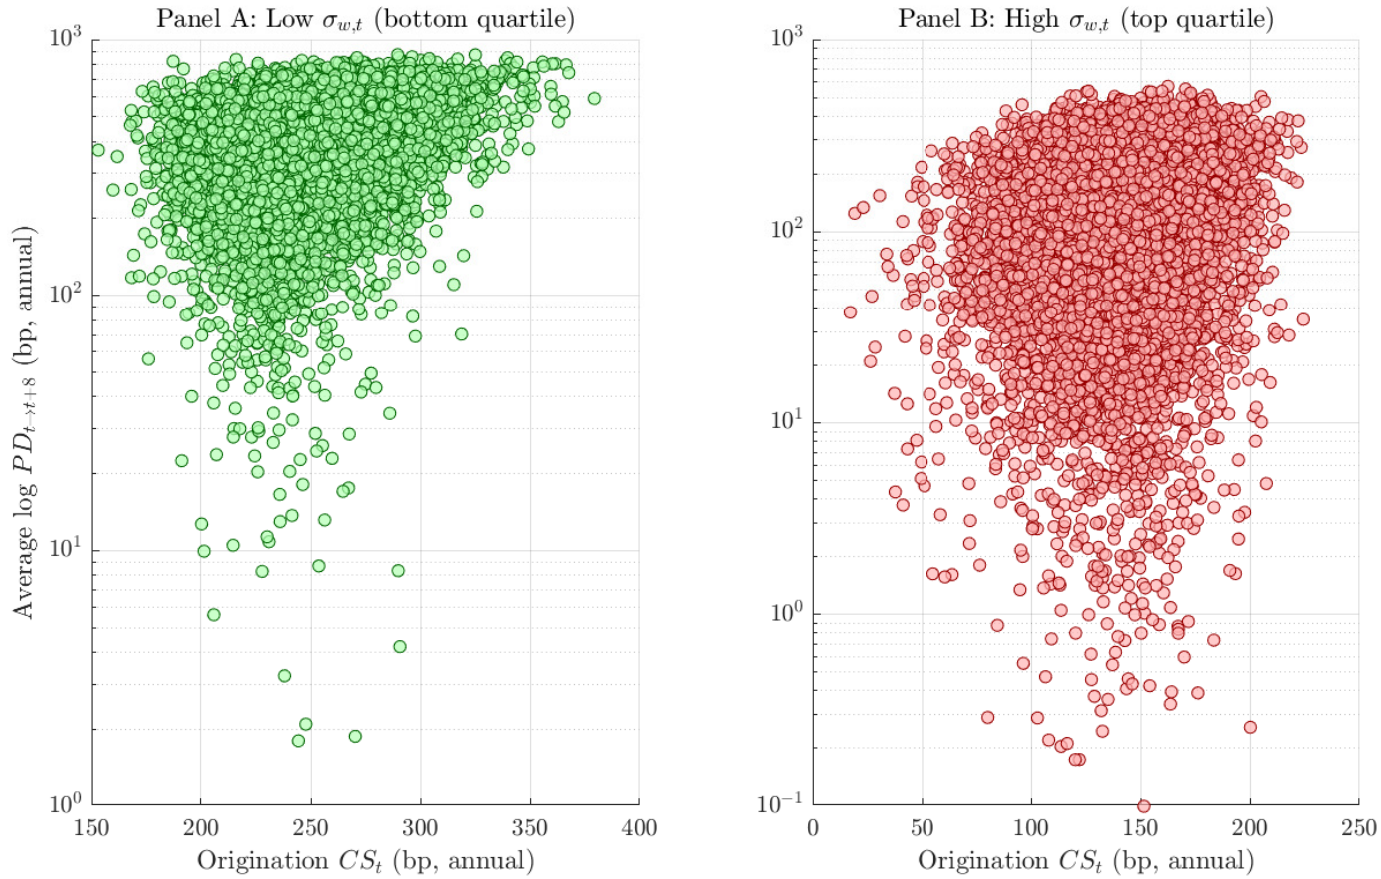
We plot the timing of financial information arrival, contracting, and payoffs within period t in the economy driven by shocks to the cross-sectional volatility of firms' willingness to pay $\sigma_{w,t}$. The top lane tracks the aggregate state S_t, the middle lane shows firms' borrowing and default decisions, and the bottom lane shows banks' signals, underwriting rule s_t^*, and portfolio payoffs, together with the feedback from end-of-t outcomes into S_{t+1}. Solid arrows denote sequences of actions and realized flows within the period. Dashed arrows indicate informational and state-dependence links.

Figure 3: Vintage Credit Spreads and Default Rates by SNR



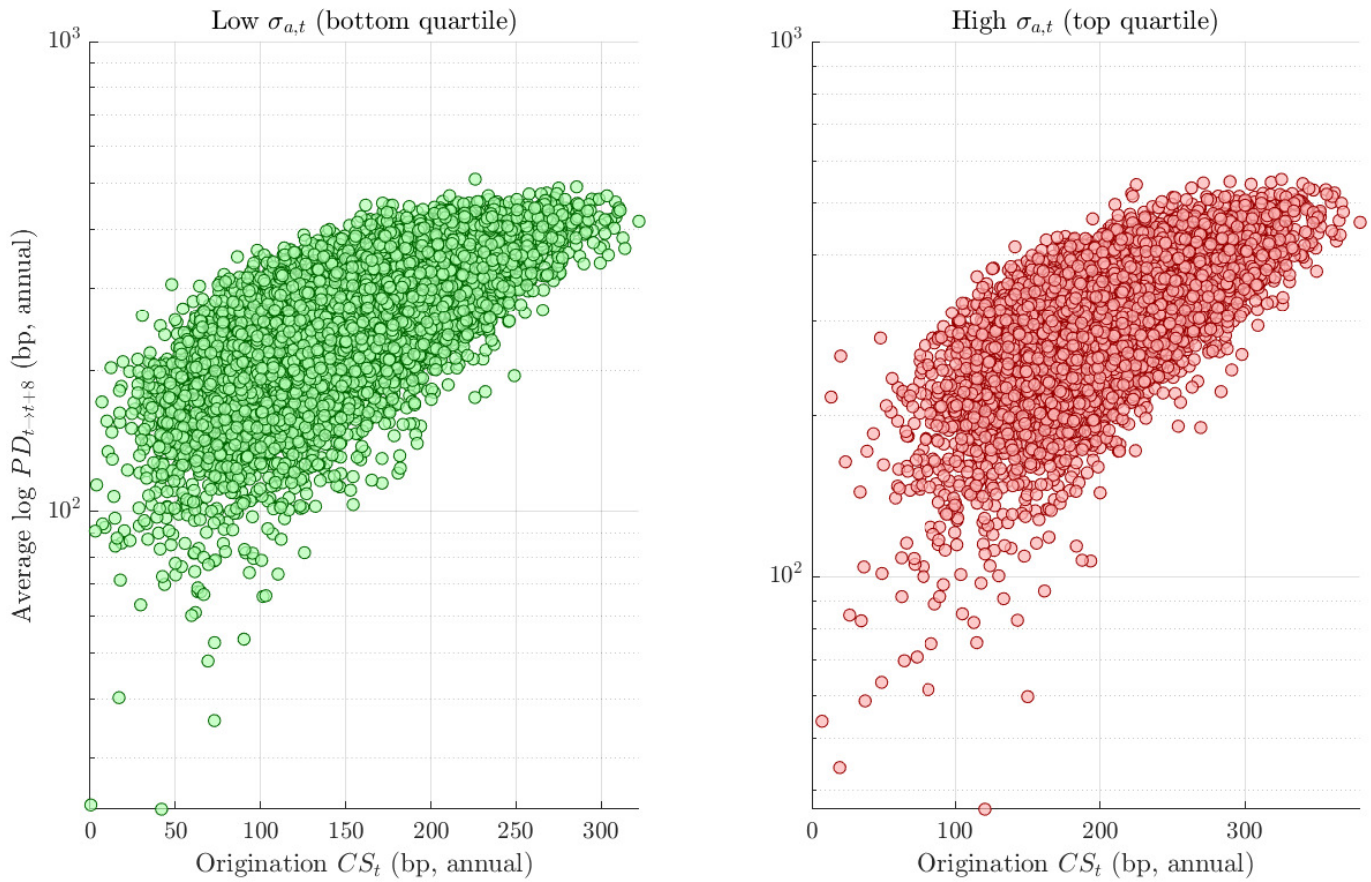
We plot the cloud of loan vintages in the model, with the horizontal axis showing the credit spread at origination CS_t and the vertical axis showing the average default rate $PD_{t\to t+K}$ on newly originated loans over the subsequent horizon $(t+1,\ldots,t+K)$. Each point is colored by the contemporaneous SNR, the signal-to-noise ratio Λ_t at origination. Darker colors correspond to higher Λ_t. Across vintages, higher CS_t is associated with higher subsequent $PD_{t\to t+K}$. Conditional on $PD_{t\to t+K}$, vintages with higher Λ_t are associated with lower CS_t. Both axes are in annualized basis points.

Figure 4: Vintage Spreads and Default Rates by W2P Dispersion at Origination



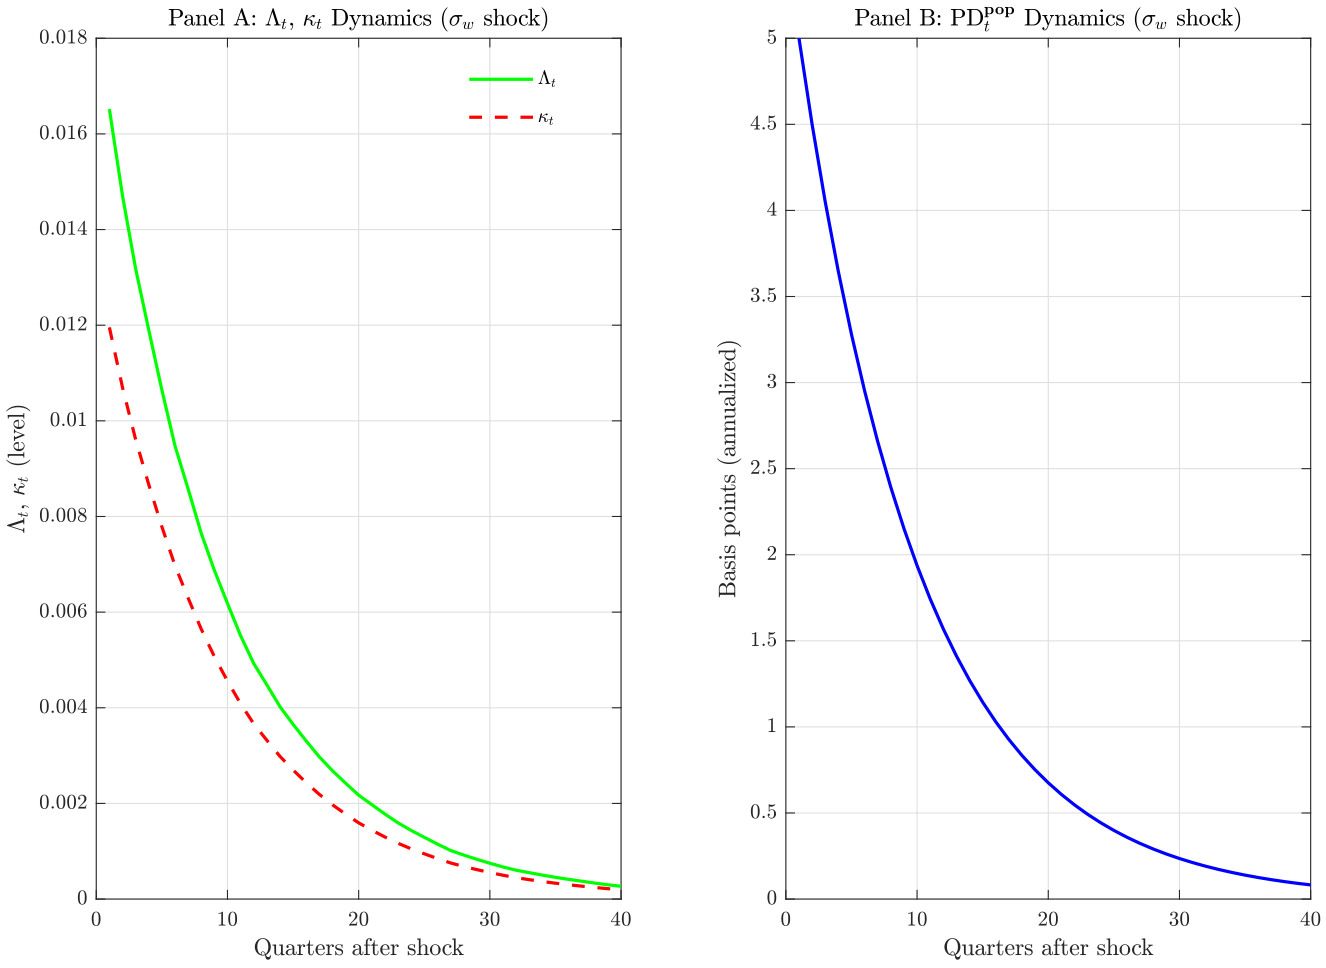
We plot the vintage-level relationship between origination spreads CS_t and the subsequent default outcome $PD_{t\to t+K}$, where $PD_{t\to t+K}$ is the average default rate on newly originated loans over quarters $t+1$ through $t+K$. We split the sample by W2P dispersion at origination. Panel A includes quarters in the bottom quartile of $\sigma_{w,t}$ (low-dispersion states). Panel B includes quarters in the top quartile of $\sigma_{w,t}$ (high-dispersion states). In low-dispersion states, the spread–default relationship is less dispersed and close to monotone. In high-dispersion states, the relationship is more dispersed, and vintages with relatively low CS_t for a given $PD_{t\to t+K}$ are associated with higher contemporaneous underwriting intensity (higher Λ_t). Both axes are in annualized basis points.

Figure 5: Vintage Spreads and Default Rates by A2P Dispersion at Origination



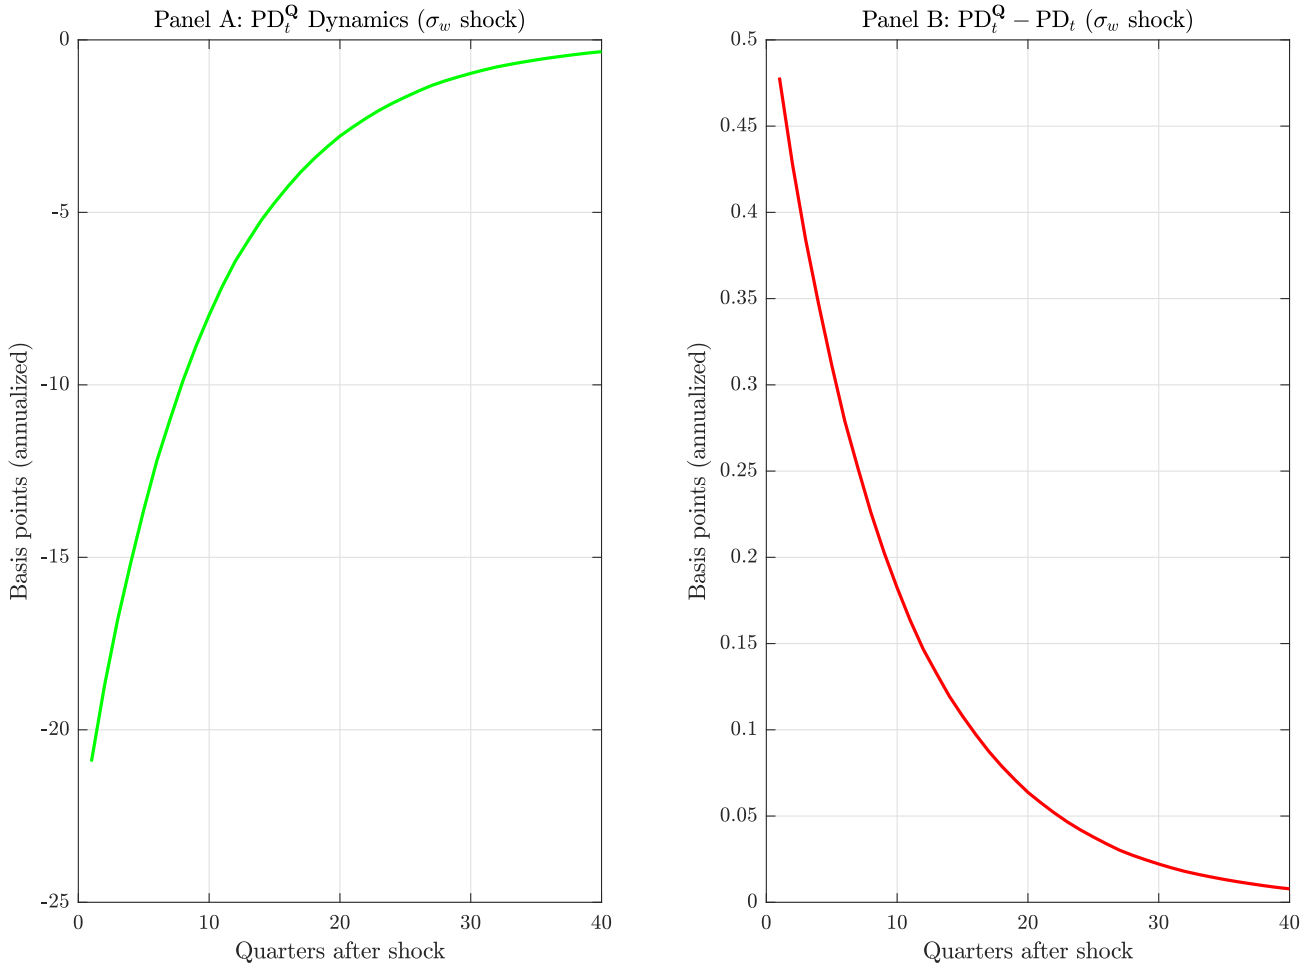
We plot the vintage-level relationship between origination spreads CS_t and the subsequent default outcome $PD_{t\to t+K}$, where $PD_{t\to t+K}$ is the average default rate on newly originated loans over quarters $t+1$ through $t+K$. We split the sample by A2P dispersion at origination. Panel A includes quarters in the bottom quartile of $\sigma_{a,t}$ (low-dispersion states). Panel B includes quarters in the top quartile of $\sigma_{a,t}$ (high-dispersion states). In low-dispersion states, the spread–default relationship is less dispersed and close to monotone. In high-dispersion states, the relationship remains strongly monotone: higher CS_t is associated with higher $PD_{t\to t+K}$, and the overall cloud shifts in the familiar direction. Consistent with this, the correlation between CS_t and $PD_{t\to t+K}$ is high and similar across low- and high-$\sigma_{a,t}$ states. Both axes are in annualized basis points.

Figure 6: Dynamics of SNR, κ_t and PD_t^{pop}



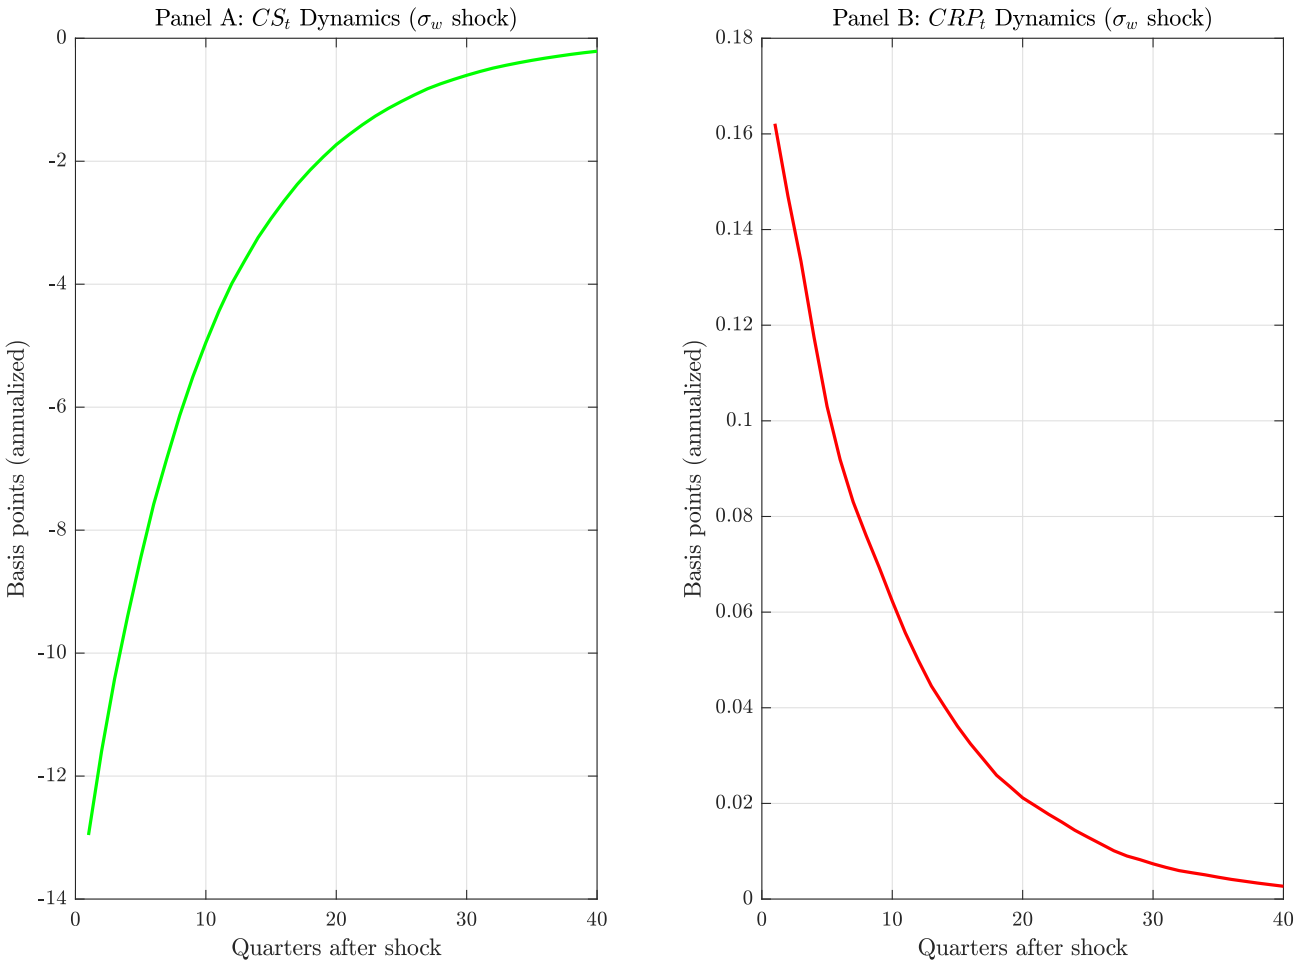
We plot the model impulse responses to a one-standard-deviation shock to $\sigma_{w,t}$, the cross-sectional dispersion of firms' willingness to pay (W2P). Panel A reports the responses of the signal-to-noise ratio Λ_t and the posterior weight κ_t, which measures the loading of bank signals on the W2P component. Panel B reports the response of the population (pre-underwriting) default probability $\mathrm{PD}_t^{\mathrm{pop}}$. The shock occurs in period 0, and responses are reported over the subsequent 40 quarters.

Figure 7: Dynamics of $\mathrm{PD}_t^{\mathbb{Q}}$ and the Risk-Neutral–Physical Default Gap



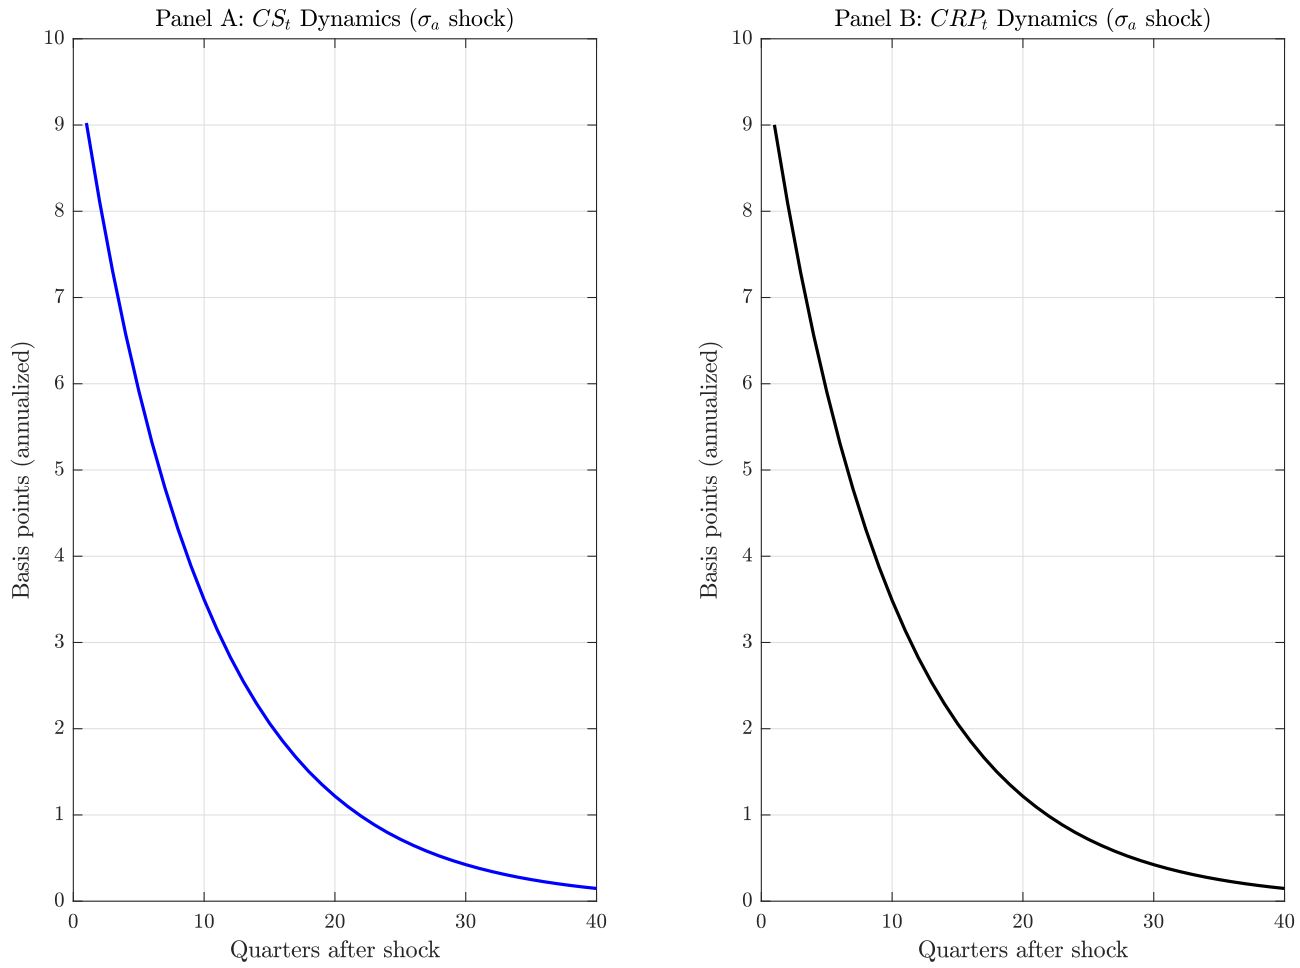
We plot the model impulse responses to a one-standard-deviation shock to $\sigma_{w,t}$, the cross-sectional dispersion of firms' willingness to pay (W2P). Panel A reports the response of the risk-neutral default probability $\mathrm{PD}_t^{\mathbb{Q}}$, and Panel B reports the response of the risk-neutral–physical default gap $\mathrm{PD}_t^{\mathbb{Q}} - \mathrm{PD}_t$. The shock occurs in period 0, and responses are reported over the subsequent 40 quarters. Both series are expressed in annualized basis points.

Figure 8: Dynamics of CS and CRP



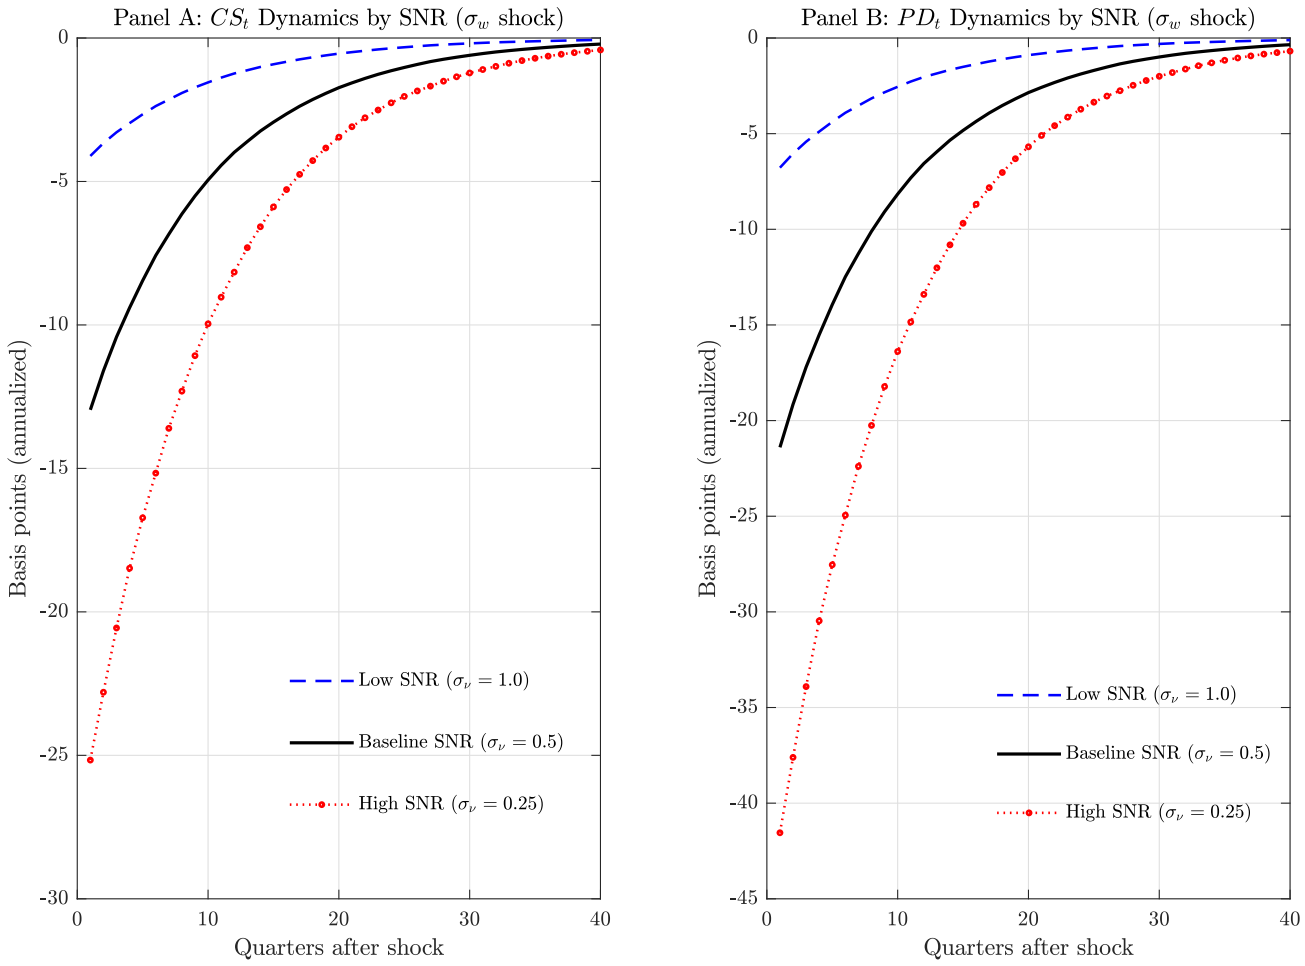
We plot the model impulse responses of credit spreads and the credit risk premium to a one-standard-deviation shock to $\sigma_{w,t}$, the cross-sectional dispersion of firms' W2P. Panel A shows the response of the model credit spread CS_t, and Panel B shows the response of the expected excess return (credit risk premium) CRP_t on newly originated loans. The shock occurs in period 0 and the responses are reported for the following 40 quarters. Both series are expressed in annualized basis points.

Figure 9: Dynamics of CS and CRP (σ_a shocks)



We plot the model impulse responses of credit spreads and the credit risk premium to a one-standard-deviation shock to $\sigma_{a,t}$, the cross-sectional dispersion of firms' fundamentals (ability to pay). Panel A shows the response of the credit spread CS_t, and Panel B shows the response of the credit risk premium CRP_t. The shock hits in period 0 and the responses are reported for 40 quarters in annualized basis points.

Figure 10: Dynamics of CS and PD by SNR Regimes



We plot the impulse responses of credit spreads and the physical default probability to a one-standard-deviation shock to $\sigma_{w,t}$ across three information regimes. Panel A reports the response of CS_t and Panel B reports the response of PD_t, with each panel overlaying low-, baseline-, and high-SNR cases (low, medium, and high Λ_t). The shock occurs in period 0, and responses are reported over the subsequent 40 quarters. Across regimes, differences in Λ_t map into differences in the magnitude of the responses, while the persistence is similar. In higher-SNR regimes, the $\sigma_{w,t}$ shock is associated with a larger decline in CS_t and PD_t. Vertical axes are in annualized basis points.