

Loans in the Fog: Cognitive Frictions and Capital Regulation in Bank Lending

June 2, 2025

Abstract

We develop a model of bank lending under cognitive frictions. Banks face cognitive costs when screening borrowers and optimally choose the precision of noisy signals about borrower quality. Higher attention costs reduce signal precision, distorting loan allocation: banks with pessimistic priors under-lend, while those with optimistic priors over-lend. Our closed-form characterization links information-processing frictions to systematic credit misallocation. Regulatory capital constraints limit risk-taking but also weaken banks' incentives to acquire information. This interaction can further reduce precision: a phenomenon we call “risk parity by ignorance.” The framework shows how cognitive costs and capital regulation jointly affect credit allocation and underscores the need to balance prudential limits with informational efficiency.

Keywords: Bank lending; Cognitive frictions; Credit allocation; Rational inattention; Capital regulation.

JEL Classification: D83; G21; G28; D80.

1 Introduction

How do cognitive constraints and information-processing limitations shape bank lending, interact with capital regulation, and affect financial stability? Banks demonstrably do **not** operate with perfect foresight or limitless analytical bandwidth. Whether managing complex structured products prior to the 2008 financial crisis or handling relationship-based SME lending today, loan officers routinely process thousands of heterogeneous signals under tight deadlines and regulatory scrutiny. Campbell *et al.* (2019) provide evidence that spikes in loan-officer workloads or timing around weekends systematically degrade the interpretation

of soft information, leading to poorer loan outcomes. Wang (2020) shows that the median loan officer believes they screen three times more information than their actual capacity allows, and Gao *et al.* (2018) document that loan officers experience regular distractions that impair decision quality. Yet, standard banking theories typically assume lenders face constraints primarily arising from moral hazard (e.g., Holmström & Tirole, 1997) or costly auditing (e.g., Diamond, 1984; Townsend, 1979), largely overlooking intrinsic cognitive or informational limitations. This disconnect matters: misallocated credit due to misjudgment or information overload can significantly destabilize financial systems, as illustrated by the Global Financial Crisis.

Rational inattention (RI), pioneered by Sims (2003), provides a parsimonious framework to analyze these information-processing limits.¹ Agents optimally choose *how much* information to acquire, balancing the benefits of better decisions against cognitive or monetary costs. RI has notably influenced macroeconomics (Maćkowiak & Wiederholt, 2009) and asset pricing (Miao & Su, 2023; Huang & Liu, 2007), and has recently been applied to macro credit-cycle modeling (Gemmi, 2024). However, while recent studies have applied RI to asset pricing and macroeconomic settings, the micro-level interaction between banks' cognitive constraints and capital regulation remains unexplored, leaving critical policy-relevant questions unanswered.

This paper integrates rational inattention into a micro-level loan-choice framework to analyze its interaction with regulatory capital requirements. We embed rational inattention into a canonical two-asset model to study how attention costs distort borrower screening, credit allocation, and the efficacy of capital regulation. Specifically, we address three key questions: (i) How do attention costs influence banks' screening precision? (ii) Under what conditions do cognitive frictions lead to systematic over-lending or under-lending relative to an optimal benchmark? (iii) How does capital regulation interact with banks' cognitive constraints, and what are the resulting regulatory design implications?

We make four main contributions. First, we incorporate rational inattention into a standard loan-choice framework in which a bank optimally selects the risk profile of its loan portfolio by choosing the precision of borrower-specific signals. This explicitly models how cognitive costs affect screening decisions. Second, we show that higher attention costs reduce signal precision, distorting loan allocation; banks with pessimistic priors under-lend, while those with optimistic priors over-lend. The economic mechanism is intuitive: when precise information is costly, banks rely more heavily on their prior beliefs, amplifying initial biases. Third, we explicitly analyze the interaction of cognitive frictions and regulatory

¹For a comprehensive survey of RI applications, including portfolio under-diversification, price stickiness, and political information choice, see Maćkowiak *et al.* (2023).

capital constraints. We show that capital requirements act as a blunt substitute for attention: they curb excessive risk-taking but simultaneously weaken banks’ incentives to acquire precise borrower information. Once capital requirements bind and the marginal benefit of distinguishing among borrowers falls, banks may achieve compliance by reducing information processing, a phenomenon we label *risk parity by ignorance*. That is, rather than repricing or resizing loans, banks cut back on screening precision and allow the capital rule to drive credit allocation. Thus, capital regulation and information regulation emerge as complements: effective policy must simultaneously manage capital buffers and cognitive constraints. Fourth, we outline policy tools that internalize these cognitive frictions, including regulatory capital surcharges tied to screening precision, simplified disclosure requirements, and targeted RegTech investment.

Our results also relate to classic insights in corporate finance. Myers (1977) shows how debt overhang can lead to underinvestment, while Jensen and Meckling (1976) argue that equityholders may overinvest due to risk-shifting incentives. Unlike traditional agency-based explanations focusing on incentive misalignments, our model highlights purely cognitive drivers of these distortions, offering regulators novel angles to address systemic inefficiencies.

Our findings resonate with recent advances in understanding the economic implications of cognitive limitations. Gabaix & Laibson (2022) demonstrate how cognitive constraints (manifested as noisy simulations of future outcomes) lead to systematic behavioral biases such as hyperbolic discounting and myopia; similarly, our model shows how cognitive constraints in banks induce systematic biases in lending decisions. Bordalo *et al.* (2018) highlight how diagnostic expectations arising from cognitive heuristics create excess volatility and predictable reversals in credit cycles; our RI framework provides complementary evidence, linking banks’ screening precision directly to these cognitive mechanisms. Additionally, Caplin & Dean (2015) demonstrate that seemingly irrational decisions can rationally result from costly information acquisition, reinforcing our model’s rational foundations for cognitive constraints. Although our paper is theoretical, the explicit testable predictions and numerical illustrations provided establish a clear foundation for future empirical validation.

A growing empirical literature finds that binding capital requirements alter banks’ portfolios and, at times, weaken credit discipline, though it stops short of measuring information acquisition itself. For instance, Gopalakrishnan *et al.* (2021) show that after the adoption of Basel II risk-weighted capital rules, U.S. banks curtailed lending to lower-rated firms and reallocated credit toward safer exposures. Behn *et al.* (2022) document that German banks using the Basel IRB approach systematically reported lower probability of default, reducing capital charges without improving realized loan performance. Bank of England balance-sheet analysis (Haldane, 2013) shows that large banks’ average risk weights fell by roughly half,

from over 70% in 1993 to below 40% in 2011, despite rising leverage, a pattern consistent with regulatory arbitrage rather than deeper screening. None of these studies provides *direct* evidence on banks' incentives to acquire borrower-specific information; yet their findings are *consistent* with our model's prediction that binding capital constraints reduce the marginal value of screening and encourage mechanical compliance strategies. This gap highlights an opportunity for new empirical work, guided by the testable implications we derive, to measure how capital regulation interacts with banks' information-acquisition choices and to quantify the resulting credit allocation distortions.

The remainder of the paper proceeds as follows. Section 2 lays out the model environment and derives the bank's optimal signal precision. Section 3 shows how rational inattention contributes to credit misallocation. Section 4 examines its interaction with capital regulation. Section 5 presents the regulatory and managerial implications that flow from our results. Finally, Section 6 offers testable empirical implications and concludes.

2 The Model

2.1 Economic environment

Consider a risk-neutral bank that faces a *continuum* of infinitesimal borrowers indexed by $i \in [0, 1]$.² The bank allocates a unit mass of loanable funds between a *safe* and a *risky* technology. For readability, we suppress the index i in what follows; all random variables should be understood as referring to a generic borrower. Payoffs per unit are

- Safe loan: deterministic return R_s .
- Risky loan: R_h with probability $p \in (0, 1)$ and R_ℓ with probability $1 - p$, where $R_\ell < R_s < R_h$.

The *true* default probability p is borrower-specific and unknown ex-ante. The bank's prior belief is

$$p \sim \mathcal{N}(\mu_0, \tau_0^{-1}), \quad (1)$$

where μ_0 is the prior mean and $\tau_0 > 0$ is the prior precision (i.e., inverse of the prior variance).

Let $\alpha = 1$ denote the bank extending a risky loan, and $\alpha = 0$ denote the bank investing in the safe asset. Before choosing $\alpha \in \{0, 1\}$, the bank may acquire a noisy signal about borrower quality:

$$s = p + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2), \quad (2)$$

²Key model symbols are defined in Table B1.

where $\sigma^2 > 0$ is the noise variance (signal precision σ^{-2}). Signal draws $\{s_i\}$ are independent across borrowers. Lower variance yields a sharper signal but incurs an *attention cost*

$$C(\sigma^{-2}) = \frac{\gamma}{2\sigma^2}, \quad \gamma > 0, \quad (3)$$

where γ is the cost parameter. This quadratic cost is the second-order Taylor approximation to the standard entropy-based (mutual-information) cost around the prior; the entropy specification itself is used in canonical RI models (e.g. Sims, 2003; Matějka & McKay, 2015).³ Therefore, the approximation preserves the same micro-foundations while keeping the model tractable.

2.2 Posterior belief

Upon observing s , the bank updates its belief via Bayes' rule:

$$p \mid s \sim \mathcal{N}\left(\mu(\sigma^2), (\tau_0 + \sigma^{-2})^{-1}\right), \quad (4)$$

with posterior mean given by the usual precision-weighted average:

$$\mu(\sigma^2) = \frac{\tau_0 \mu_0 + \sigma^{-2} s}{\tau_0 + \sigma^{-2}}. \quad (5)$$

Conditional expected return from the risky loan, therefore, equals

$$\mathbb{E}[\pi_r \mid s] = \mu(\sigma^2) R_h + (1 - \mu(\sigma^2)) R_\ell. \quad (6)$$

The bank allocates share $\alpha(s)$ to the risky loan whenever the conditional expectation exceeds R_s . Because the signal is Gaussian, the threshold is linear in s .

2.3 Value of information

Before the noise level σ^2 is chosen, the bank anticipates how the signal will affect its lending decision.

Decision rule. After observing a signal s , the bank compares the posterior expected payoff from the risky loan, with a certain payoff from the safe loan. Thus, the bank invests in the

³Replacing the quadratic approximation with the exact Sims (2003) mutual-information cost leaves all of the paper's qualitative results and comparative-static conclusions unchanged. The quadratic form is adopted solely because it yields closed-form expressions for the optimal precision; the full entropy cost produces a transcendental first-order condition that must be solved numerically.

risky loan if the following holds:

$$\mu(\sigma^2)(R_h - R_\ell) + R_\ell \geq R_s. \quad (7)$$

The bank is *indifferent* when (7) holds with equality. Define the *break-even success probability*

$$\xi := \frac{R_s - R_\ell}{R_h - R_\ell} \in (0, 1), \quad (8)$$

i.e. the borrower success rate that makes the risky and safe loans yield the same expected payoff under perfect information. Inequality (7) can therefore be rewritten as $\mu(\sigma^2) \geq \xi$. Because $\mu(\sigma^2)$ is linear in s , this condition defines a unique signal cut-off $s^*(\sigma^2)$:⁴

$$s \geq s^*(\sigma^2) \implies \alpha(s) = 1, \quad s < s^*(\sigma^2) \implies \alpha(s) = 0.$$

In words, if the realized signal exceeds the cut-off, the bank classifies that borrower as “risky” and extends a risky loan ($\alpha(s) = 1$); otherwise, it extends the safe loan ($\alpha(s) = 0$).

Standardizing the cut-off. Under the prior and signal-noise assumptions, we have

$$s \sim \mathcal{N}(\mu_0, \tau_0^{-1} + \sigma^2). \quad (9)$$

Write $\Sigma^2 := \tau_0^{-1} + \sigma^2$ and define the z -score of the cut-off,

$$Z(\sigma^2) := \frac{s^*(\sigma^2) - \mu_0}{\Sigma}. \quad (10)$$

Then $\Pr[s \geq s^*(\sigma^2)] = 1 - \Phi(Z(\sigma^2))$, where $\Phi(\cdot)$ is the standard-normal cdf and $\phi(\cdot)$ its pdf.

Expected gain from information We call the prior *pessimistic* when $\mu_0 < \xi$ (the prior mean lies below the break-even success probability) and *optimistic* when $\mu_0 > \xi$ (the prior mean lies above the break-even success probability).

Without information, a pessimistic bank holds the safe asset (which yields the constant return R_s), whereas an optimistic bank chooses the risky asset (with expected return $\mu_0 R_h + (1 - \mu_0) R_\ell$).

Information is valuable only if it *reverses* that prior decision: it must raise the posterior success probability above ξ for a pessimistic bank, or push it below ξ for an optimistic one.

⁴The closed-form expression for $s^*(\sigma^2)$ is derived in Appendix A.1.

Because

$$\Pr[s \geq s^*(\sigma^2)] = 1 - \Phi(Z(\sigma^2)), \quad \Pr[s \leq s^*(\sigma^2)] = \Phi(Z(\sigma^2)),$$

a signal strong enough to trigger the switch arrives with probability $1 - \Phi(Z)$ under a pessimistic prior and $\Phi(Z)$ under an optimistic prior, where the z -score $Z(\sigma^2)$ is defined in (10).

The resulting *expected monetary gain* is

$$f(\sigma^2) = (R_h - R_\ell) \begin{cases} \int_{s^*(\sigma^2)}^{\infty} [\mu(\sigma^2) - \xi] \phi_\Sigma(s) ds, & \text{if } \mu_0 < \xi \quad (\text{pessimistic prior}), \\ \int_{-\infty}^{s^*(\sigma^2)} [\xi - \mu(\sigma^2)] \phi_\Sigma(s) ds, & \text{if } \mu_0 > \xi \quad (\text{optimistic prior}). \end{cases} \quad (11)$$

Here $\phi_\Sigma(\cdot)$ is the normal density of the signal $s \sim \mathcal{N}(\mu_0, \Sigma^2)$.

- **Pessimistic prior** ($\mu_0 < \xi$): only signals *above* the cut-off ($s \geq s^*$) matter, because they lift the posterior success probability above ξ and justify switching to the risky loan.
- **Optimistic prior** ($\mu_0 > \xi$): only signals *below* the cut-off ($s \leq s^*$) matter, because they push the posterior success probability below ξ and induce a switch to the safe loan.
- Multiplying by $R_h - R_\ell$ converts that probability gain (or loss avoided) into expected monetary terms.

From Eq. (11), it is straightforward to show that $f'(\sigma^2) < 0$ and $f''(\sigma^2) > 0$. Intuitively, as the noise variance σ^2 falls, the signal becomes more informative and the expected monetary gain from information acquisition increases, although at a decreasing rate due to diminishing returns. Conversely, as σ^2 approaches infinity, the signal degenerates into pure noise and $f(\sigma^2)$ approaches zero.

Precision choice. The bank trades this benefit against the quadratic attention cost:

$$\max_{\sigma^2 > 0} V(\sigma^2) = f(\sigma^2) - \frac{\gamma}{2\sigma^2}. \quad (12)$$

Because $f(\sigma^2)$ is strictly decreasing and the cost term is strictly convex, $V(\sigma^2)$ is concave and the first-order condition $f'(\sigma^{2*}) + \gamma/(2\sigma^{4*}) = 0$ pins down a unique optimal precision σ^{-2*} .

Proposition 1 (Optimal precision). *Let $V(\sigma^2) = f(\sigma^2) - \gamma/(2\sigma^2)$, with $f'(\sigma^2) < 0$ and $f''(\sigma^2) > 0$. Then*

$V(\sigma^2)$ is strictly concave on $(0, \infty)$ and admits a unique maximizer $\sigma^{2} > 0$ characterized by*

$$V'(\sigma^{2*}) = f'(\sigma^{2*}) + \frac{\gamma}{2\sigma^{4*}} = 0. \quad (13)$$

The optimal variance can be written

$$\sigma^{2*} = \left[\frac{\gamma}{-2f'(\sigma^{2*})} \right]^{1/2}, \quad f'(\sigma^{2*}) < 0, \quad (14)$$

and it is strictly increasing in the attention-cost parameter:

$$\frac{d\sigma^{2*}}{d\gamma} = \frac{\sigma^{2*}}{2\gamma} > 0. \quad (15)$$

Proof. See Appendix A.2. □

This result captures a simple trade-off: greater signal precision improves lending decisions but becomes increasingly costly to acquire. The optimal noise level σ^{2*} balances these marginal benefits and costs. As attention costs γ rise, the bank rationally chooses to acquire less precise signals, foreshadowing the distortions in credit allocation we explore next.

3 Cognitive Frictions and Credit Allocation

From state-contingent to aggregate lending shares. For any chosen noise level σ^2 the bank lends to the risky project whenever the signal exceeds the cutoff $s^*(\sigma^2)$ derived in Appendix A. Hence, the *ex-ante* (aggregate) share of funds channeled to risky loans is

$$\alpha^*(\sigma^2) = \Pr[s \geq s^*(\sigma^2)] = 1 - \Phi(Z(\sigma^2)). \quad (16)$$

where $Z(\sigma^2)$ is the standardized cutoff defined in (10).⁵

Break-even success probability. A borrower with true success probability p is indifferent between the risky project and the safe project when the expected payoffs from both projects

⁵Because the bank faces a continuum of independent, infinitesimal borrowers (with a unit mass of 1), the strong law of large numbers implies that the realized fraction of projects with $s \geq s^*(\sigma^2)$ equals the probability of that event. We therefore treat this probability as the aggregate share of funds channeled to risky loans.

are equal. Recall from Section 2.3 that this occurs when $p = \xi$, where

$$\xi = \frac{R_s - R_\ell}{R_h - R_\ell} \in (0, 1).$$

Hence, under perfect information, the bank lends risky if and only if $p \geq \xi$, and safe otherwise.

First-best risky-loan share. Under perfect information ($\sigma^2 \rightarrow 0$), the bank lends risky exactly when $p \geq \xi$, where ξ is the break-even threshold above. With a Gaussian prior $p \sim \mathcal{N}(\mu_0, \tau_0^{-1})$, the corresponding first-best risky-loan share is⁶

$$\alpha^{\text{FB}} = 1 - \Phi((\xi - \mu_0) \tau_0^{1/2}). \quad (17)$$

Misallocation metric. To quantify how attention costs distort lending, define the misallocation metric as:

$$\Delta(\gamma) := |\alpha^*(\sigma^{2*}(\gamma)) - \alpha^{\text{FB}}|, \quad (18)$$

i.e. the absolute deviation of the realized risky-loan share from the first-best benchmark, evaluated at the optimal variance $\sigma^{2*}(\gamma)$ obtained in Proposition 1.

We can then prove the following proposition.

Proposition 2 (Credit misallocation under rational inattention). *The model yields two key results:*

(a) **Misallocation effect.** *The misallocation gap*

$$\Delta(\gamma) = |\alpha^*(\sigma^{2*}(\gamma)) - \alpha^{\text{FB}}|$$

is strictly increasing in the attention-cost parameter:

$$\frac{\partial \Delta(\gamma)}{\partial \gamma} > 0 \quad (\text{for all } \gamma > 0). \quad (19)$$

(b) **Allocation bias and corner limits.** *Exactly one of the following two cases obtains:*

Pessimistic prior: *If $\mu_0 < \xi$, then $\alpha^*(\sigma^{2*}(\gamma)) < \alpha^{\text{FB}}$ for every $\gamma > 0$. There exists a finite $\gamma_h \in (0, \infty)$ with $\alpha^*(\sigma^{2*}(\gamma_h)) = 0$; for $\gamma > \gamma_h$ the bank makes no risky loans.*

⁶Throughout, “first-best” refers to the bank’s private full-information optimum. A planner who internalizes systemic risk might choose a different allocation; analyzing that is beyond this paper.

Optimistic prior: If $\mu_0 > \xi$, then $\alpha^*(\sigma^{2*}(\gamma)) > \alpha^{\text{FB}}$ for every $\gamma > 0$. The over-lending gap widens monotonically, and

$$\lim_{\gamma \rightarrow \infty} \alpha^*(\sigma^{2*}(\gamma)) = 1.$$

Proof. See Appendix A.3. □

Higher attention costs make precise signals expensive, causing the bank to rely on increasingly noisy information. If the bank’s prior belief is pessimistic (expecting borrowers to have a low probability of success), it already leans towards under-lending to risky projects. Noisier signals amplify this caution, pushing lending even further below the first-best allocation. Conversely, if the bank’s prior is optimistic, noisier signals exacerbate its tendency to over-lend. In the limit, as attention becomes prohibitively costly, the bank disregards the noisy signals entirely and simply follows its prior belief—allocating either zero or all funds to risky loans. Thus, increasing information frictions always drive the bank’s lending decisions away from the first-best benchmark, never closer.

These results echo classic insights from the corporate finance literature. In particular, our finding that pessimistic priors combined with high attention costs lead to excessive caution mirrors the underinvestment problem highlighted by Myers (1977). Likewise, when optimistic priors drive over-lending in the presence of noisy signals, the pattern resembles the risk-shifting behavior discussed by Jensen and Meckling (1976). However, our model departs from these agency-based explanations by attributing distortions not to conflicting incentives but to cognitive limitations in processing information. This alternative mechanism suggests that improving the informational environment, rather than solely aligning incentives, may be key to correcting inefficiencies in capital allocation.

Figure 1 illustrates Proposition 2 with plausible parameter values: the pessimistic bank retreats from risk as signals become noisy, whereas the optimistic bank eventually allocates its entire portfolio to the risky project.

Table 1 reports the parameter values used in these numerical illustrations.

4 Capital Regulation and Cognitive Frictions

Banks rarely allocate assets in a vacuum: minimum-capital rules restrict their ability to expand risk-weighted exposures. We introduce a simple Basel-style leverage constraint and show how it interacts with attention costs.

Table 1: **Parameter Values Used in Numerical Figures**

Symbol	Value	Description	Empirical Plausibility / Source
R_s	1.03	Gross annual return on a safe loan (approx. 3% yield).	Short-term, senior-secured corporate loans typically yield 2–4%.
R_h	1.10	Gross return on a successful risky loan (10%).	Reflects upper-end SME lending rates, which average 5–8% for established firms but 10–12% for smaller or less-rated borrowers (OECD 2024).
R_ℓ	0	Zero recovery in default (stylized). [†]	Convention in theoretical banking models (e.g., Acharya & Naqvi 2012; Holmström & Tirole 1997).
ξ	0.30	Break-even success probability: $(R_s - R_\ell)/(R_h - R_\ell)$.	Mechanical from returns; 30% intuitive for high-yield lending.
μ_0	0.20 / 0.40 / 0.60	Prior means for borrower success.	S&P default studies: annual success 80–95%; lower values used illustratively to show under-/over-lending extremes.
τ_0	100 / 4	Prior precision: high ($\sigma \approx 0.10$) vs. low ($\sigma \approx 0.50$).	Reflects borrower transparency (new vs. repeat).
γ	0.30	Attention-cost scale parameter.	Chosen to produce clear comparative statics. Results qualitatively unchanged for any $\gamma > 0$.
σ^2	Varied	Noise variance (i.e., inverse of signal precision).	Swept to illustrate comparative statics in Figure 1.
α^{\max}	Varied (0.0–1.0)	Implied ceiling on risky-loan share (Figure 2).	Ranges from no constraint (1.0) to total ban (0.0). Derived from the capital requirement.

[†]Average recovery on senior secured loans is typically 30–50% (Moody’s 2023). We set $R_\ell = 0$ for tractability; allowing partial recovery would shift the break-even probabilities but not alter comparative static insights.

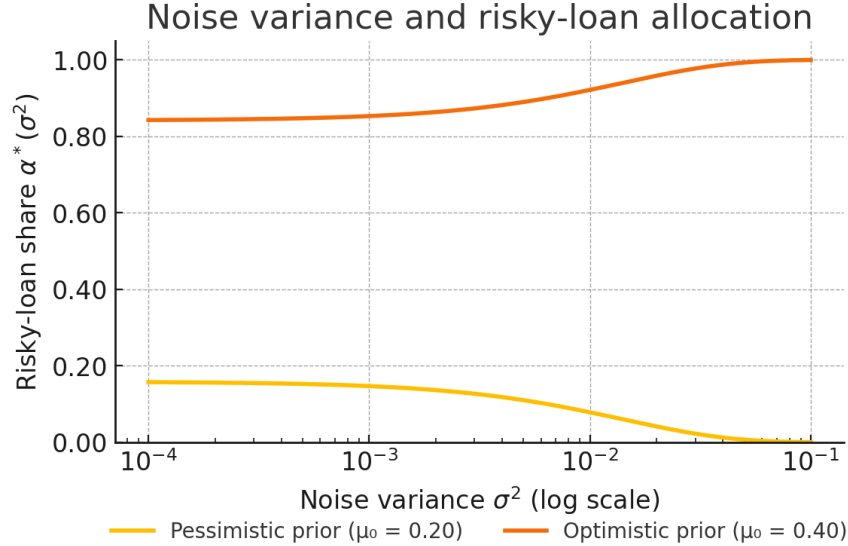


Figure 1: **Signal precision and risky-loan allocation.** The plot shows the ex-ante share of funds channeled to the risky technology, $\alpha^*(\sigma^2)$, as a function of the noise variance σ^2 (log scale). Parameter values are: safe return $R_s = 1.03$, risky up-state $R_h = 1.10$, risky down-state $R_\ell = 0$, break-even success probability $\xi = (R_s - R_\ell)/(R_h - R_\ell) = 0.30$, and prior precision $\tau_0 = 100$ (implying a prior standard deviation of 0.10). We consider two representative beliefs about borrower success: a *pessimistic* prior ($\mu_0 = 0.20 < \xi$) and an *optimistic* prior ($\mu_0 = 0.40 > \xi$). Lower σ^2 corresponds to higher information precision. Consistent with Proposition 2, noisier signals (larger σ^2) push the pessimistic bank toward zero lending, whereas the optimistic bank converges toward universal lending. The parameters lie in plausible ranges for large-firm default probabilities and typical loan spreads and are chosen purely for illustration.

4.1 Regulatory set-up

Let the bank hold fixed equity $K > 0$. Equity is set prior to the lending period and typically adjusted only infrequently due to issuance costs, disclosure lags, and regulatory approvals. Empirically, when facing short-term capital pressures, banks typically adjust by reducing risk-weighted assets rather than raising new equity. For example, Gropp et al. (2019) exploit the 2011 EBA capital exercise and find that treated banks improved their capital ratios by reducing risk-weighted assets by an average of 16 percentage points, rather than by increasing equity, consistent with debt-overhang frictions. Similarly, Memmel and Raupach (2010) show that although liability-side adjustments are more potent, asset-side adjustments occur more quickly. We therefore treat K as fixed within our decision horizon and focus on the bank's dual choice of (i) the risky-loan share α and (ii) screening precision σ^{-2} .⁷

⁷This setup is equivalent to a two-stage game: *Stage 0* (reporting date), the bank sets equity K based on last period's information; *Stage 1*, it observes new borrower signals and allocates loans. Nevertheless, for

Denote by $\alpha^*(\sigma^2) \in [0, 1]$ the *aggregate* share of funds invested in the risky loan (defined in Section 3). Risk-weighted assets (RWA) are

$$\text{RWA} = w \alpha^* + \bar{w} (1 - \alpha^*), \quad 0 < \bar{w} < w \leq 1, \quad (20)$$

where w and \bar{w} are Basel risk weights on the risky and safe loans, respectively.⁸

The bank is *compliant* if its capital ratio exceeds $\theta \in (0, 1)$:

$$\frac{K}{\text{RWA}} \geq \theta \iff \alpha^* \leq \alpha^{\max} := \frac{K/\theta - \bar{w}}{w - \bar{w}}. \quad (21)$$

We assume $K/\theta > \bar{w}$, so $\alpha^{\max} \in (0, 1]$. When (21) binds, the bank faces an *upper bound* on its risky-loan share. The capital requirement in (21) thus imposes an *implied ceiling* on the bank's risky-loan share, α^{\max} , henceforth referred to as the *implied ceiling*. When this implied ceiling binds, the bank faces a strict upper limit on risky exposures; otherwise, it behaves as in the baseline model.

4.2 Optimal behavior with the constraint

Define the capital-adjusted signal threshold $s^\dagger(\sigma^2)$ as follows:

- **Slack constraint** ($\alpha^*(\sigma^2) \leq \alpha^{\max}$). The bank's desired risky-loan share already satisfies the capital requirement, so no tightening is needed: set

$$s^\dagger(\sigma^2) = s^*(\sigma^2).$$

- **Binding constraint** ($\alpha^*(\sigma^2) > \alpha^{\max}$). Choose the smallest threshold $s^\dagger(\sigma^2) > s^*(\sigma^2)$ such that

$$\Pr[s \geq s^\dagger(\sigma^2)] = \alpha^{\max}.$$

Increasing the cut-off from s^* to s^\dagger lowers the aggregate risky-loan share from α^* down to the regulatory ceiling α^{\max} .

completeness, we show in Online Appendix C that even with gradual equity adjustment, the model yields qualitatively similar predictions, as capital remains sticky over typical decision horizons relevant to quarterly stress tests or internal credit cycles.

⁸Setting $\bar{w} = 0$ and $w = 1$ reproduces a simple leverage ratio; any $0 < \bar{w} < w$ captures differential weighting under Basel III.

Hence, if the constraint is slack, $s^\dagger = s^*$; if binding, s^\dagger is chosen such that $\Pr[s \geq s^\dagger] = \alpha^{\max}$. The *capital-constrained* lending rule is therefore

$$\tilde{\alpha}(s; \sigma^2) = \mathbf{1}\{s \geq s^\dagger(\sigma^2)\}. \quad (22)$$

Here $\tilde{\alpha}(s; \sigma^2)$ takes the values 0 or 1 at the borrower level. Taking expectations over the continuum of borrowers gives the aggregate risky-loan share

$$\tilde{\alpha}^*(\sigma^2) = \mathbb{E}[\tilde{\alpha}(s; \sigma^2)] = \Pr[s \geq s^\dagger(\sigma^2)] = \min\{\alpha^*(\sigma^2), \alpha^{\max}\}. \quad (23)$$

Expected gain under the capital constraint. For a given signal noise level σ^2 , let

$$f_c(\sigma^2) := (R_h - R_\ell) \begin{cases} \int_{s^*(\sigma^2)}^{s^\dagger(\sigma^2)} [\mu(\sigma^2) - \xi] \phi_\Sigma(s) ds, & \text{if } \mu_0 < \xi \quad (\text{pessimistic prior}), \\ \int_{-\infty}^{s^*(\sigma^2)} [\xi - \mu(\sigma^2)] \phi_\Sigma(s) ds, & \text{if } \mu_0 > \xi \quad (\text{optimistic prior}). \end{cases} \quad (24)$$

represent the expected monetary benefit of the signal when the capital requirement is in place, where the integration limits $s^*(\sigma^2)$ and $s^\dagger(\sigma^2)$ are the baseline and capital-constrained cut-offs defined earlier.

Equation (24) gives the expected monetary gain from information once the capital requirement is in place. For a pessimistic prior, it is the upper band $s^* \leq s \leq s^\dagger$ that matters; for an optimistic prior the relevant region remains the lower tail $s \leq s^*$ because the capital requirement truncates only the upper tail. Signals above s^\dagger are “cut off” by the capital cap and therefore add no marginal value, while signals below s^* leave the bank in the safe project by construction.

Precision choice under the capital constraint. The bank selects the signal noise level to maximize its net value

$$V_c(\sigma^2) = f_c(\sigma^2) - \frac{\gamma}{2\sigma^2}, \quad \sigma^2 > 0. \quad (25)$$

Here $V_c(\sigma^2)$ represents the *expected monetary payoff from lending, net of attention costs, when the capital requirement may restrict the risky-loan share*. Because $f_c(\sigma^2) \leq f(\sigma^2)$ and inherits the same concavity properties until the capital constraint binds, the first-order condition is

$$f'_c(\sigma_c^{2*}) + \frac{\gamma}{2\sigma_c^{4*}} = 0, \quad (26)$$

which yields a unique optimum satisfying $\sigma_c^{2*} \geq \sigma^{2*}$.

Proposition 3 (Capital constraint and optimal precision). *Let $\alpha^*(\sigma^2) = 1 - \Phi(Z(\sigma^2))$ denote the unconstrained risky-loan share. Then:*

- (i) **Slack constraint.** *If $\alpha^*(\sigma^{2*}) \leq \alpha^{\max}$, the capital constraint does not bind, and the optimal noise level satisfies*

$$\sigma_c^{2*} = \sigma^{2*},$$

so all baseline results carry over unchanged.

- (ii) **Binding constraint.** *If $\alpha^*(\sigma^{2*}) > \alpha^{\max}$, the capital constraint binds and the first-order condition (26) implies*

$$\sigma_c^{2*} > \sigma^{2*};$$

that is, the bank acquires strictly less information than in the unconstrained benchmark.

- (iii) **Effect on misallocation.**

- Optimistic prior ($\mu_0 > \xi$). *When the cap binds, it reduces the over-lending gap but also weakens information acquisition, so the net effect on misallocation is non-monotone in γ .*
- Pessimistic prior ($\mu_0 < \xi$). *The cap binds only if the ceiling is set so low that $\alpha^{\max} < \alpha^*(\sigma^{2*})$. When this happens, it trims the positive-gain region $[s^*, s^\dagger]$, further increasing under-lending and worsening misallocation; otherwise, the capital constraint is slack, and the baseline results are obtained.*

Proof. See Appendix A.4. □

Figure 2 illustrates how an implied ceiling affects credit allocation for a fixed attention cost.⁹ The figure plots the *misallocation bias*,

$$\tilde{\Delta} = \alpha_c^* - \alpha^{\text{FB}},$$

where α_c^* denotes the aggregate share of funds allocated to risky loans *after* applying the implied ceiling α^{\max} (this coincides with the unconstrained share α^* whenever the ceiling is slack). The bias is positive for over-lending and negative for under-lending.

The bias declines as the ceiling tightens from very loose levels, reaches zero when the ceiling matches the bank's first-best share, and becomes negative once the capital requirement

⁹Figure 2 uses a more optimistic and precise prior ($\mu_0 = 0.60$, $\tau_0 = 100$) to reflect a repeat borrower and isolate the effect of capital constraints. Figure 1 instead employs a less optimistic, diffuse prior ($\mu_0 = 0.40$, $\tau_0 = 4$) to emphasize how attention costs distort decisions under borrower uncertainty. These differences do not affect the results and are chosen solely to clarify distinct comparative statics.

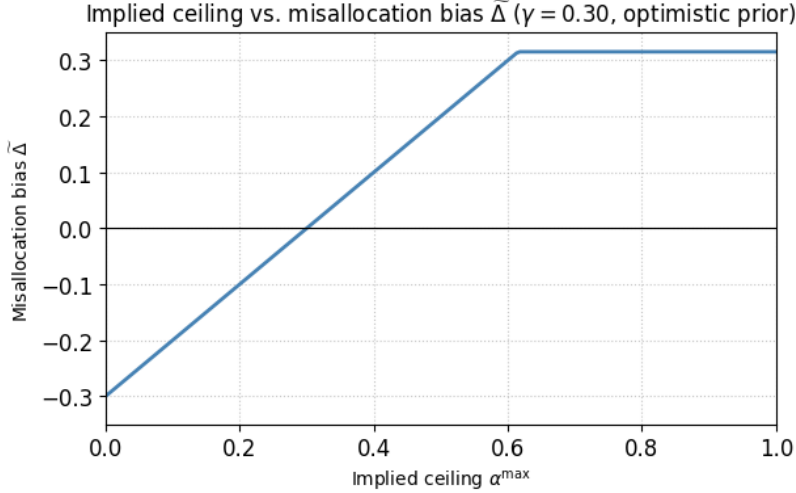


Figure 2: **Implied ceilings and credit misallocation.** The plot shows the *misallocation bias* $\tilde{\Delta} = \alpha_c^* - \alpha^{\text{FB}}$ as a function of the implied ceiling α^{\max} , which reflects the maximum risky-loan share permitted by the capital requirement. Parameter values follow Figure 1, except for the prior and attention cost: optimistic prior $\mu_0 = 0.60 > \xi$, prior precision $\tau_0 = 4$ (standard deviation = 0.50), and attention cost $\gamma = 0.30$. These imply an unconstrained risky-loan share of 0.62 and a first-best share $\alpha^{\text{FB}} = 0.30$. When the implied ceiling is loose ($\alpha^{\max} \gtrsim 0.62$), the capital constraint is slack and over-lending persists ($\tilde{\Delta} > 0$). Tightening the ceiling initially reduces the bias, but once α^{\max} falls below α^{FB} , the capital requirement induces under-lending ($\tilde{\Delta} < 0$).

forces lending below that benchmark. Hence, a single regulatory tool can either improve or worsen allocation depending on how tightly it is set, consistent with Proposition 3.

The minimum capital requirement acts as a *blunt substitute for attention*. When the bank is **optimistic** ($\mu_0 > \xi$), it would like to expand its risky portfolio beyond the implied ceiling α^{\max} . Once the ceiling binds, very high signals no longer translate into larger exposures, so the *marginal* value of precision falls, and the bank rationally opts for noisier signals: $\sigma_c^{2*} > \sigma^{2*}$. Whenever compliance is achieved by relaxing screening effort rather than by repricing or resizing individual loans, we refer to this mechanism as *risk parity by ignorance*. As the ceiling tightens (or attention costs rise), precision is dialed back further; in the extreme, information acquisition can be crowded out altogether, leaving the bank compliant yet uninformed.

With a **pessimistic** prior ($\mu_0 < \xi$), the capital constraint is typically *slack*. If the regulator sets an exceptionally tight ceiling such that it *does* bind, the cap further suppresses the already low risky-loan share and weakens screening incentives, thereby exacerbating under-lending.

The key takeaway is that capital regulation and information regulation are *complements*:

a capital requirement can curb tail risk, but only at the cost of weaker screening incentives. Policies that both lower attention costs (e.g. via RegTech adoption) *and* set appropriate capital buffers are therefore more effective than relying on either tool in isolation, as discussed in Section 5.

5 Regulatory and Managerial Implications

Our results have direct and meaningful implications for regulatory frameworks and managerial policy design. Below, we detail specific policy recommendations derived from our model, emphasizing practical strategies for regulators and financial institutions to mitigate cognitive-driven credit misallocation.

5.1 Supervisory design: Stress-tests for cognition

Regulators currently quantify risks using standardized market and liquidity metrics and routinely assess banks’ governance, internal controls, and operational resilience under Pillar 2. Our model suggests extending this approach by explicitly incorporating banks’ cognitive processing capacity, captured by the attention-cost parameter γ . While γ itself is latent, supervisors can proxy it using readily available indicators such as loan-officer workloads, audit-trail data on review times, employee turnover in risk-assessment units, IT investment per loan officer, or the frequency of monitoring disruptions. Banks identified as having elevated cognitive frictions could then face enhanced Pillar 2 capital surcharges or be required to meet minimum staffing ratios in risk-assessment teams or to deploy certified underwriting-technology upgrades to restore screening precision. As highlighted in Section 4, stress tests should also explicitly model the interaction between capital constraints and cognitive capacity to comprehensively assess bank resilience.

5.2 Capital incentives for precision

Our analysis indicates that capital constraints, while curbing excessive risk-taking, may inadvertently weaken banks’ incentives for precise borrower screening (*see Proposition 3(ii)*). Regulators should therefore introduce a *precision-linked capital adjustment*: banks demonstrating consistently superior screening accuracy (measured by out-of-sample prediction performance relative to peers) would receive lower capital charges or reduced risk-weight floors. This incentive structure directly counteracts the informational disincentive created by binding capital constraints, encouraging continued investment in advanced screening technologies and cognitive infrastructure. Unlike the enhanced Pillar 2 surcharges proposed in Section 5.1,

which act as penalties on banks with high cognitive frictions, the precision-linked capital adjustment functions as a positive incentive, granting lower capital requirements to banks that demonstrate superior screening accuracy.

5.3 Calibrating capital requirements

While Section 5.2 outlines one possible incentive tool (precision-linked adjustments), the broader calibration challenge is to design capital requirements that balance financial stability with robust information acquisition. Regulators should design the architecture of capital standards to strike a balance between prudential safety and information-acquisition incentives. As Section 4 demonstrates, overly stringent capital limits can lead banks to rationally choose lower screening precision, creating a trade-off between financial stability and informational efficiency. One way to operationalize this principle is the precision-linked capital adjustment described in Section 5.2, which rewards banks that demonstrate superior screening accuracy by lowering their capital charges. More generally, calibration might include differentiated capital requirements based on banks' demonstrable information-processing efficiency, thereby harmonizing prudential regulation with incentives for cognitive investment.

5.4 RegTech and fintech adoption

Beyond capital levers, regulators and banks can also address cognitive frictions through technological innovation. Emerging technologies such as cloud-based machine learning platforms for data extraction and credit scoring materially lower banks' information-processing costs and thus the attention-cost parameter γ . A lower γ raises the optimal signal precision, improving lending decisions (*see Proposition 1*). Regulators can accelerate adoption by creating public-private sandboxes and by enabling secure, anonymized sharing of supervisory data with approved technology providers. Such initiatives foster innovation, enhance screening precision, and complement capital regulation by allowing banks to meet prudential requirements without sacrificing informational quality.

5.5 Managerial attention dashboards

Banks should enhance risk governance by explicitly tracking screening precision and cognitive load through an *attention dashboard*. Relevant indicators include model predictive accuracy, override rates, analyst workloads, average time spent per credit file, backlog volume, system response latency, IT investment per analyst, and staff turnover in credit-assessment teams. These metrics can serve as real-time proxies for cognitive frictions. For example, unusually

short review times may signal rushed screening under overload, while excessive review times may indicate system inefficiencies or decision fatigue. Similarly, rising override rates might reflect reduced model trust or inconsistency in judgment, whereas persistently low override rates could indicate disengagement or insufficient time for manual evaluation. Breaches of internal thresholds, such as persistent backlogs or declining model accuracy, should prompt governance-level reviews and corrective actions such as hiring staff, deploying RegTech, or adjusting workflows. Dashboards should be reviewed monthly, with real-time alerts available for critical deviations to support proactive cognitive risk management.

5.6 Disclosure simplification

Complex regulatory reporting burdens banks’ limited cognitive resources. Simplifying regulatory disclosures, standardizing reporting templates, and harmonizing borrower information requirements (such as standardized ESG disclosures) could significantly reduce sector-wide attention costs (γ). Streamlined disclosures, aligned with frameworks such as the Bank of England’s *Strong and Simple* initiative (Bank of England PRA , 2022), would enable banks to allocate cognitive resources more effectively, complementing capital regulation by indirectly fostering improved credit screening precision.

6 Discussion

6.1 Testable Implications

Our model yields several testable empirical implications. Banks facing higher cognitive load, proxied by loan officer backlogs or spikes, analyst turnover, or lower IT investments, should exhibit weaker screening precision.

Moreover, conditional on observable borrower risk, higher attention costs tend to amplify deviations from benchmark lending allocations. These benchmarks can be constructed using publicly available default probabilities or external rating-agency default statistics.

Following a tightening of risk-weighted capital ratios, the misallocation gap (the difference between actual and benchmark risky-loan shares) should initially shrink, reflecting improved alignment. However, if the constraint becomes too tight, this gap may widen again, indicating renewed misallocation due to diminished screening incentives.

Finally, banks that comply with capital requirements by reducing information acquisition rather than raising equity are likely to experience higher default rates on newly originated loans in the one- to three-year period following the change, compared to peers who maintained high screening precision or strengthened capital buffers.

6.2 Conclusion

We develop a tractable model demonstrating how cognitive frictions, modeled via *rational inattention*, reshape banks’ risk-taking and credit allocation decisions, explicitly incorporating the influence of regulatory capital requirements. Our framework embeds these frictions into a micro-level borrower-screening model of bank lending and analyzes their interaction with capital regulation, highlighting the implications for banks’ screening precision and credit outcomes.

Our framework shows that cognitive limitations can systematically distort credit allocation, leading to under-lending or over-lending depending on banks’ prior beliefs. We also find that while capital regulation is essential for curbing excessive risk-taking, it may unintentionally weaken banks’ incentives to acquire borrower-specific information, creating a tension between limiting systemic risk and preserving information-based screening efficiency. Binding capital constraints can give rise to what we term “risk parity by ignorance,” whereby banks comply mechanically with the implied ceiling on risky-loan share by substantially weakening their screening precision and, in the extreme, eliminating it altogether.

For policymakers, these insights suggest that traditional regulatory tools such as capital ratios and liquidity buffers should be carefully calibrated to preserve screening incentives. Explicitly incorporating cognitive constraints into supervisory frameworks can enhance regulatory effectiveness. Investments in technologies such as FinTech and RegTech, which lower information-processing costs, offer dual benefits: improving credit allocation and reducing systemic risk.

As discussed in Section 6.1, future empirical research could operationalize these predictions using granular loan-level data.

While advances in artificial intelligence and machine learning promise to reduce information frictions, cognitive constraints are likely to persist in legacy systems and slower-moving institutions. Understanding how attention limitations interact with technological change remains an important area for future research. A promising avenue is to develop a dynamic model in which attention costs adapt endogenously, reflecting changes in technology or regulatory conditions, and enabling richer policy trade-offs over time. Ultimately, our model offers a structured foundation for integrating cognitive dimensions into modern banking theory, regulatory design, and empirical analysis.

Appendix A: Technical Proofs

A.1 Cut-off Signal $s^*(\sigma^2)$

The bank lends to the risky project whenever the posterior expected payoff exceeds the safe payoff:

$$\mu(\sigma^2)(R_h - R_\ell) + R_\ell \geq R_s. \quad (27)$$

Write $\Delta R := R_h - R_\ell$ and define the break-even success probability

$$\xi := \frac{R_s - R_\ell}{\Delta R} \in (0, 1).$$

Posterior mean as a linear function of the signal. With a normal prior $p \sim \mathcal{N}(\mu_0, \tau_0^{-1})$ and signal $s = p + \varepsilon$, $\varepsilon \sim \mathcal{N}(0, \sigma^2)$, the posterior mean is

$$\mu(\sigma^2) = \mu_0 + \kappa(\sigma^2)[s - \mu_0], \quad \kappa(\sigma^2) := \frac{\sigma^{-2}}{\tau_0 + \sigma^{-2}} = \frac{1}{1 + \tau_0 \sigma^2} \in (0, 1). \quad (28)$$

Solving for the signal cut-off. Substituting the linear form of $\mu(\sigma^2)$ into the lending inequality and solving for s yields the unique cut-off:

$$s^*(\sigma^2) = \mu_0 + \frac{\xi - \mu_0}{\kappa(\sigma^2)} = \mu_0 + (\xi - \mu_0)(1 + \tau_0 \sigma^2) = \xi + \tau_0 \sigma^2 (\xi - \mu_0). \quad (29)$$

Standardized form. Because $s \sim \mathcal{N}(\mu_0, \tau_0^{-1} + \sigma^2)$, the z -score of the cut-off used in the main text is

$$Z(\sigma^2) = \frac{s^*(\sigma^2) - \mu_0}{\sqrt{\tau_0^{-1} + \sigma^2}}.$$

Expression (29) guarantees a strictly increasing cut-off in σ^2 : a noisier signal (larger σ^2) requires a higher realization s to justify lending to the risky project.

A.2 Proof of Proposition 1

The objective is $V(\sigma^2) = f(\sigma^2) - \gamma/(2\sigma^2)$ with $f'(\sigma^2) < 0$ and $f''(\sigma^2) > 0$. The derivative is

$$V'(\sigma^2) = f'(\sigma^2) + \frac{\gamma}{2\sigma^4}.$$

Since $f'(\sigma^2)$ is strictly increasing and $\gamma/(2\sigma^4)$ is strictly decreasing, $V'(\sigma^2)$ crosses zero exactly once, implying a unique maximizer $\sigma^{2*} > 0$. Solving $V'(\sigma^{2*}) = 0$ yields the stated expression for σ^{2*} .

Differentiating the first-order condition with respect to γ gives

$$\left[f''(\sigma^{2*}) + \frac{3\gamma}{2\sigma^{5*}}\right] \frac{d\sigma^{2*}}{d\gamma} + \frac{1}{2\sigma^{4*}} = 0. \quad (30)$$

Because the bracketed term is positive, we obtain $d\sigma^{2*}/d\gamma = \sigma^{2*}/(2\gamma) > 0$, completing the proof. \square

A.3 Proof of Proposition 2

Let $\sigma^{2*}(\gamma)$ be the optimal variance. Recall that $Z(\gamma) = \frac{s^*(\sigma^{2*}(\gamma)) - \mu_0}{\sqrt{\tau_0^{-1} + \sigma^{2*}(\gamma)}}$ with $s^*(\sigma^2) = \xi + \tau_0(\xi - \mu_0)\sigma^2$. The ex-ante risky-loan share is $\alpha^*(\gamma) = 1 - \Phi(Z(\gamma))$.

Part (a). From Proposition 1, $d\sigma^{2*}/d\gamma = \sigma^{2*}/(2\gamma) > 0$. Differentiate Z :

$$\frac{\partial Z}{\partial \sigma^2} = (\xi - \mu_0) \frac{\tau_0^{3/2}}{2\sqrt{1 + \tau_0\sigma^2}}, \quad \frac{dZ}{d\gamma} = \frac{\partial Z}{\partial \sigma^2} \frac{d\sigma^{2*}}{d\gamma}. \quad (31)$$

Hence $\text{sign}(dZ/d\gamma) = \text{sign}(\xi - \mu_0)$. Because $\alpha^{*'}(\gamma) = -\phi(Z(\gamma)) dZ/d\gamma$ and $\phi > 0$, we obtain $\text{sign } \alpha^{*'} = -\text{sign}(\xi - \mu_0)$. For every $\gamma > 0$ the difference $x(\gamma) = \alpha^*(\gamma) - \alpha^{\text{FB}}$ retains this constant sign, so $\Delta'(\gamma) = |x'(\gamma)| = |\alpha^{*'}(\gamma)| > 0$.

Part (b). If $\mu_0 < \xi$ then $\alpha^{*'}(\gamma) < 0$ for all γ and $\alpha^* \downarrow 0$ as $\gamma \rightarrow \infty$; continuity yields a unique γ_h with $\alpha^*(\gamma_h) = 0$. If $\mu_0 > \xi$ the derivative is positive and $\alpha^* \uparrow 1$ as $\gamma \rightarrow \infty$. \square

A.4 Proof of Proposition 3

Recall the unconstrained value function $V(\sigma^2) = f(\sigma^2) - \gamma/(2\sigma^2)$ with first-order condition $f'(\sigma^{2*}) + \gamma/(2\sigma^{4*}) = 0$. Under the capital requirement, the objective becomes $V_c(\sigma^2) = f_c(\sigma^2) - \gamma/(2\sigma^2)$ where

$$f_c(\sigma^2) = (R_h - R_\ell) \int_{s^*(\sigma^2)}^{s^\dagger(\sigma^2)} [\mu(\sigma^2) - \xi] \phi_\Sigma(s) ds, \quad (32)$$

and the upper limit $s^\dagger(\sigma^2)$ satisfies $\Pr[s \geq s^\dagger] = \alpha^{\max}$. Define the first-order condition under the capital requirement as

$$f'_c(\sigma_c^{2*}) + \frac{\gamma}{2\sigma_c^{4*}} = 0. \quad (33)$$

Part (i): Slack constraint ($\alpha^* \leq \alpha^{\max}$). If the capital constraint is not binding at $s^\dagger(\sigma^{2*}) = s^*(\sigma^{2*})$ and hence $f_c(\sigma^2) = f(\sigma^2)$ on a neighborhood of σ^{2*} . Equation (33) reduces

to the unconstrained first-order condition, yielding $\sigma_c^{2*} = \sigma^{2*}$. Part (i) follows.

Part (ii): Binding constraint ($\alpha^* > \alpha^{\max}$). When the capital constraint binds, $s^\dagger(\sigma^2)$ is finite and strictly larger than $s^*(\sigma^2)$. Differentiating (32) gives

$$f'_c(\sigma^2) = f'(\sigma^2) - (R_h - R_\ell)[\mu(\sigma^2) - \xi] \phi_\Sigma(s^\dagger(\sigma^2)) \frac{\partial s^\dagger}{\partial \sigma^2}. \quad (34)$$

where the second term is non-negative because $\partial s^\dagger / \partial \sigma^2 > 0$. Hence $f'_c(\sigma^2) \geq f'(\sigma^2)$ for every σ^2 , with strict inequality when the cap is binding. Since $f'(\sigma^2) < 0$, f'_c is *closer to zero* than f' . The term $\gamma/(2\sigma^4)$ in (33) is unchanged, so the root of the equation shifts to the right:

$$\sigma_c^{2*} > \sigma^{2*},$$

establishing part (ii).

Part (iii): Effect on misallocation

- *Optimistic prior* ($\mu_0 > \xi$). Unconstrained lending already exceeds the first-best share α^{FB} . Imposing the capital requirement lowers the risky-loan share toward α^{\max} (mitigating over-lending), but the rise in σ_c^{2*} weakens screening precision, which increases misallocation. The net effect on $|\tilde{\alpha}^* - \alpha^{\text{FB}}|$ is therefore non-monotone in γ .
- *Pessimistic prior* ($\mu_0 < \xi$). The capital requirement binds *only* if the regulator chooses an implied ceiling tighter than the bank's unconstrained risky-loan share, $\alpha^{\max} < \alpha^*(\sigma^{2*})$. When this happens, the positive-gain region $[s^*, s^\dagger]$ is truncated, so $f_c(\sigma^2) < f(\sigma^2)$ and the first-order condition again shifts rightward. Hence $\sigma_c^{2*} > \sigma^{2*}$, the bank acquires *less* information, and the under-lending gap widens. If instead $\alpha^{\max} \geq \alpha^*(\sigma^{2*})$, the capital constraint is slack, $\sigma_c^{2*} = \sigma^{2*}$, and the baseline allocation obtains.

□

Appendix B: Glossary of Symbols

Table B1: **Glossary of Key Symbols**

Symbol	Meaning / Role
p	Borrower's true success probability (latent)
μ_0	Bank's prior mean belief about p
τ_0	Bank's prior precision (inverse of prior variance)
s	Noisy signal about borrower quality: $s = p + \varepsilon$
ε	Signal noise: $\varepsilon \sim \mathcal{N}(0, \sigma^2)$
σ^2	Variance of signal noise (i.e., inverse of signal precision σ^{-2})
γ	Attention-cost parameter (higher γ implies costlier precision)
$C(\sigma^{-2})$	Cognitive cost of acquiring signal precision: $C = \frac{\gamma}{2\sigma^2}$
$\mu(\sigma^2)$	Posterior mean of p after observing s
R_s	Return on the safe loan
R_h	Return on risky loan if successful
R_ℓ	Return on risky loan if failed
ξ	Break-even success probability: $\xi = \frac{R_s - R_\ell}{R_h - R_\ell}$
$s^*(\sigma^2)$	Signal threshold for lending to risky project
$s^\dagger(\sigma^2)$	Adjusted signal threshold to enforce capital requirement: $\Pr[s \geq s^\dagger] = \alpha^{\max}$
$Z(\sigma^2)$	Standardized cut-off: $Z = \frac{s^* - \mu_0}{\sqrt{\tau_0^{-1} + \sigma^2}}$
$\alpha(s)$	Lending rule: 1 if $s \geq s^*$, else 0
$\alpha^*(\sigma^2)$	Aggregate risky-loan share (equals $\Pr[s \geq s^*]$)
α^{FB}	First-best risky-loan share under perfect information
α^{\max}	Implied ceiling on risky-loan share (derived from the capital requirement)
$\Delta(\gamma)$	Misallocation gap: $ \alpha^* - \alpha^{\text{FB}} $
$f(\sigma^2)$	Expected gain from information (unconstrained case)
$f_c(\sigma^2)$	Expected gain from information under minimum capital requirement
$V(\sigma^2)$	Net benefit of information: $V = f(\sigma^2) - C(\sigma^{-2})$
$V_c(\sigma^2)$	Constrained net benefit: $V_c = f_c(\sigma^2) - C(\sigma^{-2})$

Online Appendix C: Slow Equity Adjustment and Asset-Side Compliance

This appendix shows formally that when new equity arrives with a lag and incurs issuance costs, a bank that suddenly breaches its capital ratio can regain compliance *only* by adjusting its risky-loan share α_t in the current period. Equity issuance affects the constraint no earlier than the next period, so it becomes a secondary (future) adjustment margin.

C.1 Dynamic set-up

Time is discrete, $t = 0, 1, 2, \dots$. At the start of period t , the bank holds equity K_t and chooses

$$(\alpha_t, \Delta K_t) \in [0, 1] \times [0, \infty),$$

where

- α_t = fraction of the loan book invested in the *risky* technology this period;
- ΔK_t = equity announced for issuance at date t (shares sold today, proceeds received at $t+1$).

Timing within period t .

Stage 1: The bank observes its beginning-of-period equity K_t and the current shock (e.g. a change in risk weights or asset values).

Stage 2: It selects $(\alpha_t, \Delta K_t)$.

Lending profits $\Pi(\alpha_t, \sigma^{-2})$ are realized *during* the period. The function Π is strictly increasing in expected return and already incorporates the attention cost analyzed in the main text, so no additional derivations are required here.

Stage 3: If the bank's capital ratio at the *end* of period t violates the regulatory floor θ , it pays a supervisory penalty $\phi > 0$.

Stage 4: The equity issue settles: $K_{t+1} = K_t + \Delta K_t$ (new shares are now on the balance sheet).

Capital-ratio constraint (end of period t). Risk-weighted assets after lending are

$$\text{RWA}_t = w \alpha_t + \bar{w} (1 - \alpha_t), \quad 0 < \bar{w} < w \leq 1,$$

so the regulatory ratio equals K_t/RWA_t . Define

$$\alpha_t^{\max}(K_t) := \max\left\{\alpha \in [0, 1] : \frac{K_t}{w\alpha + \bar{w}(1 - \alpha)} \geq \theta\right\}.$$

In closed form,

$$\alpha_t^{\max}(K_t) = \min\left\{1, \max\left\{0, \frac{K_t/\theta - \bar{w}}{w - \bar{w}}\right\}\right\}.$$

If $\alpha_t > \alpha_t^{\max}(K_t)$ the penalty ϕ is incurred.

Equity-issuance cost. Announcing $\Delta K_t \geq 0$ shares costs

$$\Psi(\Delta K_t) = \frac{\eta}{2} (\Delta K_t)^2, \quad \eta > 0,$$

with no proceeds until $t+1$.

Objective. The bank **maximizes** discounted profit

$$V_t(K_t) = \max_{\alpha_t, \Delta K_t} \left\{ \underbrace{\Pi(\alpha_t, \sigma^{-2}) - \Psi(\Delta K_t) - \phi \mathbf{1}\{\alpha_t > \alpha_t^{\max}(K_t)\}}_{\text{period-}t \text{ profit net of costs}} + \beta V_{t+1}(K_t + \Delta K_t) \right\},$$

where $\beta \in (0, 1)$ is the bank's discount factor. Negative signs appear in front of Ψ and ϕ because they reduce profit; this is therefore an unconstrained maximization problem.¹⁰

C.2 The optimal adjustment rule

Lemma 1 (Instantaneous asset-side compliance). *Define*

$$\alpha_t^{\max}(K_t) = \max\left\{\alpha \in [0, 1] : \frac{K_t}{w\alpha + \bar{w}(1 - \alpha)} \geq \theta\right\}.$$

If the penalty ϕ is large enough that violating the constraint is never optimal, then the bank's optimal risky-loan share is $\alpha_t^ = \alpha_t^{\max}(K_t)$.*

Proof. Any $\alpha_t > \alpha_t^{\max}(K_t)$ triggers the penalty ϕ . Because ϕ dominates the incremental lending profit by assumption, such choices are strictly dominated. Among the penalty-free choices $[0, \alpha_t^{\max}(K_t)]$, Π is increasing in α_t , so the bank attains its maximum at $\alpha_t^{\max}(K_t)$. \square

¹⁰The one-period lending payoff $\Pi(\alpha_t, \sigma^{-2})$ is strictly increasing in the risky share α_t (all else equal) and already incorporates the endogenous information cost analyzed in Sections 2–4 of the main text.

Lemma 2 (Optimal equity issuance). *Given α_t^* , the optimal equity issue solves*

$$\eta \Delta K_t^* = \beta \frac{\partial V_{t+1}}{\partial K}(K_t + \Delta K_t^*), \quad \text{with } \Delta K_t^* \geq 0.$$

Proof. Substituting α_t^* from Lemma 1 into the objective removes the penalty term. The maximization over ΔK_t reduces to

$$\max_{\Delta K_t \geq 0} \{ -\Psi(\Delta K_t) + \beta V_{t+1}(K_t + \Delta K_t) \}.$$

Ψ is strictly convex and differentiable; V_{t+1} is strictly concave in K (equity relaxes tomorrow's constraint). The Kuhn–Tucker conditions yield

$$\eta \Delta K_t^* - \beta \partial_K V_{t+1} = 0, \quad \Delta K_t^* \geq 0,$$

which gives the stated first-order condition. □

Proposition 4 (Asset-side adjustment is *always* required). *In any period when the capital ratio is violated at the start, the bank must set α_t no higher than $\alpha_t^{\max}(K_t)$ to regain compliance. Equity issuance ΔK_t^* can never restore the ratio within the same period because the proceeds arrive with a lag.*

Proof. Equity issued at date t is added to K only at $t+1$, so it cannot affect the end-of-period- t capital ratio. By Lemma 1, the *only* available control that changes the ratio today is α_t . Therefore, the bank must choose $\alpha_t^* = \alpha_t^{\max}(K_t)$ to avoid the penalty ϕ . □

Corollary 1 (When does equity issuance occur?). *The bank sets $\Delta K_t^* = 0$ unless*

$$\beta \frac{\partial V_{t+1}}{\partial K}(K_t) > 0.$$

Issuance is therefore purely forward-looking: it occurs only when the discounted marginal value of additional equity in future periods exceeds the marginal issuance cost $\eta \Delta K_t$. Whether that happens is irrelevant for Propositions 2–3 in the main text, which describe period- t misallocation; those results depend solely on the instantaneously chosen α_t^ .*

Proof. Set $\Delta K_t = 0$ in Lemma 2. If the right-hand side is non-positive, the Kuhn–Tucker complementarity implies the optimum is $\Delta K_t^* = 0$. Issuance arises only if the marginal continuation benefit is strictly positive. □

C.3 Implications

Proposition 4 establishes that *regardless of the level of issuance cost η , the lag in settlement forces the bank to meet today’s capital requirement by adjusting its risky-loan share α_t* . Equity issuance is, at best, a secondary tool for easing *future* constraints; Corollary 1 shows it activates only when the discounted marginal benefit outweighs the cost.

Hence, the fixed-equity assumption in Sections 2–4 of the main text is without loss of generality for analyzing *same-period* lending behavior and misallocation. All comparative-static results carry over unchanged.

References

- Acharya, V. V., & Naqvi, H. (2012). The seeds of a crisis: A theory of bank liquidity and risk taking over the business cycle. *Journal of Financial Economics*, 106(2), 349–366.
- Bank of England Prudential Regulation Authority. (2022). *CP5/22 — The Strong and Simple Framework: Definition of a Simpler-regime Firm*. Available at: <https://www.bankofengland.co.uk/prudential-regulation/publication/2022/april/definition-of-a-simpler-regime-firm>.
- Behn, M., Haselmann, R., & Vig, V. (2022). The limits of model-based regulation. *The Journal of Finance* 77(3), pp. 1571–1612.
- Bordalo, P., Gennaioli, N., & Shleifer, A. (2018). Diagnostic Expectations and Credit Cycles. *Journal of Finance* 73(1), pp. 199–227.
- Campbell, D., Loumioti, M., & Wittenberg-Moerman, R. (2019). Making sense of soft information: Interpretation bias and loan quality. *Journal of Accounting and Economics*, 68(2–3), 101240.
- Caplin, A., & Dean, M. (2015). Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review* 105(7), pp. 2183–2203.
- Diamond, D. (1984). Financial Intermediation and Delegated Monitoring. *Review of Economic Studies* 51(3), pp. 393–414.
- Gabaix, X., & Laibson, D. (2022). Myopia and Discounting. Working paper, Harvard University.
- Gao, J., Karolyi, S. & Pacelli, J. (2018). Screening and Monitoring by Inattentive Corporate Loan Officers. Working paper, Indiana University.

- Gemmi, L. (2024). Rational Overoptimism and Limited Liability. *Journal of Monetary Economics* 143, 103538.
- Gopalakrishnan, B., Jacob, J., & Mohapatra, S. (2021). Risk-sensitive Basel regulations and firms' access to credit: Direct and indirect effects. *Journal of Banking & Finance* 126, 106101.
- Gropp, R., T. Mosk, S. Ongena, & C. Wix (2019). Banks' response to higher capital requirements: Evidence from a quasi-natural experiment. *Review of Financial Studies*, 32(1), 266–299.
- Haldane, A. G. (2013). Constraining discretion in bank regulation. *BIS Central Bankers' Speeches*, Federal Reserve Bank of Atlanta Conference, April 9, 2013.
- Holmström, B., & Tirole, J. (1997). Financial Intermediation, Loanable Funds, and the Real Sector. *Quarterly Journal of Economics* 112(3), pp. 663–691.
- Huang, L., & Liu, H. (2007). Rational Inattention and Portfolio Selection. *Journal of Finance* 62(4), pp. 1999–2040.
- Jensen, M.C., & Meckling, W.H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305–360.
- Maćkowiak, B., Matějka, F., & Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, 61(1), 226–273.
- Maćkowiak, B., & Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3), 769–803.
- Matějka, F., & McKay, A. (2015). Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review* 105(1), pp. 272–292.
- Mommel, C. & P. Raupach (2010). How do banks adjust their capital ratios? Evidence from Germany. *Journal of Financial Intermediation*, 19(4), 509–528.
- Miao, J., & Su, D. (2023). Asset market equilibrium under rational inattention. *Economic Theory* 75, pp. 1–30.
- Myers, S.C. (1977). Determinants of corporate borrowing. *Journal of Financial Economics*, 5(2), 147–175.
- Organization for Economic Co-operation and Development. (2024). *Financing SMEs and Entrepreneurs 2024: An OECD Scoreboard*. OECD Publishing, Paris.

- Sims, C. (2003). Implications of Rational Inattention. *Journal of Monetary Economics* 50(3), pp. 665–690.
- Townsend, J. (1979). Optimal Contracts and Competitive Markets with Costly State Verification. *Journal of Economic Theory* 21, pp. 265–293.
- Wang, J. (2020). Screening Soft Information: Evidence From Loan Officers. *Rand Journal of Economics* 51(4), pp. 1287–1322.