BANK OF CANADA
BANQUE DU CANADA

Staff working paper / Document de travail du personnel—2026-26

Last updated: July 24, 2026

From Stress to Strategy: How Banks Balance the Scales

Ruben Hipp
European Central Bank
ruben.hipp@ecb.europa.eu

Javier Ojea-Ferreiro
Financial Stability Department
Bank of Canada
jojeaferreiro@bank-banque-canada.ca

DOI: https://doi.org/10.34989/swp-2026-26 | ISSN 1701-9397

From Stress to Strategy: How Banks Balance the Scales*

Ruben Hipp†
Javier Ojea-Ferreiro‡

This draft: August 1, 2026

Abstract

This paper introduces a partial equilibrium stress-testing model in which heterogeneous banks optimize portfolios subject to regulatory constraints, with endogenous price effects and managerial buffers arising from a Nash equilibrium. Applied to Canada's Big Six banks, the framework is used to illustrate three applications. First, it identifies a sharp nonlinearity: aggregate credit provision is stable if banks sit comfortably above the regulatory minimum but contracts abruptly when close or below. Second, a counterfactual increase in the Domestic Stability Buffer from 3% to 3.5% reduces aggregate lending by approximately 1.2%, an effect comparable in magnitude to a 25 basis point increase in the policy rate. Third, reverse stress testing recovers the macro-financial environments consistent with systemic stress without conditioning on any prior narrative. Two of three identified stress clusters feature above-trend GDP growth: one driven by a house price correction, and one by Canadian dollar appreciation alongside rising oil prices. Both represent vulnerabilities that conventional scenario design would systematically miss.

JEL Classification Codes: C63, G21, G32, C73, H12

Keywords: Partial equilibrium, banking, stress testing, policy evaluation

*We are thankful for outstanding support by Vatya Kishore, and research assistance by Sebastian Bratu, Sascha Clazie-Thomson, and Florent Samson. This work benefited from comments from David Aikman, Thibaut Duprey, Jeongwoo Moon, Josef Schroth, seminar participants at the ECB Banking Supervision Seminar, the Bank of England Financial Stability Department Seminar, Dodd Frank Act Stress Testing group Seminar of the Federal Reserve Board, the International Monetary Fund Seminar and the Federal Reserve of Boston SRA Seminar, SAEe 2025 in Barcelona, Banco de España Seminar and Finance Forum 2026 in Alicante. The views expressed are those of the authors and do not necessarily reflect those of the Bank of Canada or the European Central Bank. All errors and omissions are ours.

†European Central Bank. e-mail: ruben.hipp@ecb.europa.eu

‡Bank of Canada. e-mail: jojeaferreiro@bank-banque-canada.ca

1 Introduction

The banking literature typically models regulatory constraints as binding thresholds, yet banks in practice operate well above their regulatory minima and adjust their portfolios continuously as buffers shrink. This dynamic is central to understanding credit provision under stress but largely absent from existing frameworks. Using supervisory data from Canada's Big Six banks, Figure 1 documents two systematic patterns. The left panel shows that banks with smaller capital buffers (measured as the distance between their capital adequacy ratio and the regulatory minimum) reduce mortgage exposure in the following year, even when comfortably above the regulatory minimum. The right panel shows the converse: banks far from the regulatory minimum display little urgency in building additional buffer. Together, these patterns suggest that portfolio decisions depend continuously on regulatory proximity. Banks do not wait for a constraint to bind before responding; they anticipate and manage it, adjusting their asset and liability composition in ways that reflect both return objectives and the option value of maintaining distance from the regulatory floor. This behaviour is consistent with the empirical regularity, documented by Orozco and Rubio (2025) and Juselius et al. (2025) among others, that banks maintain substantial managerial buffers above the regulatory minimum not merely as a precaution, but as an active instrument of balance sheet management.

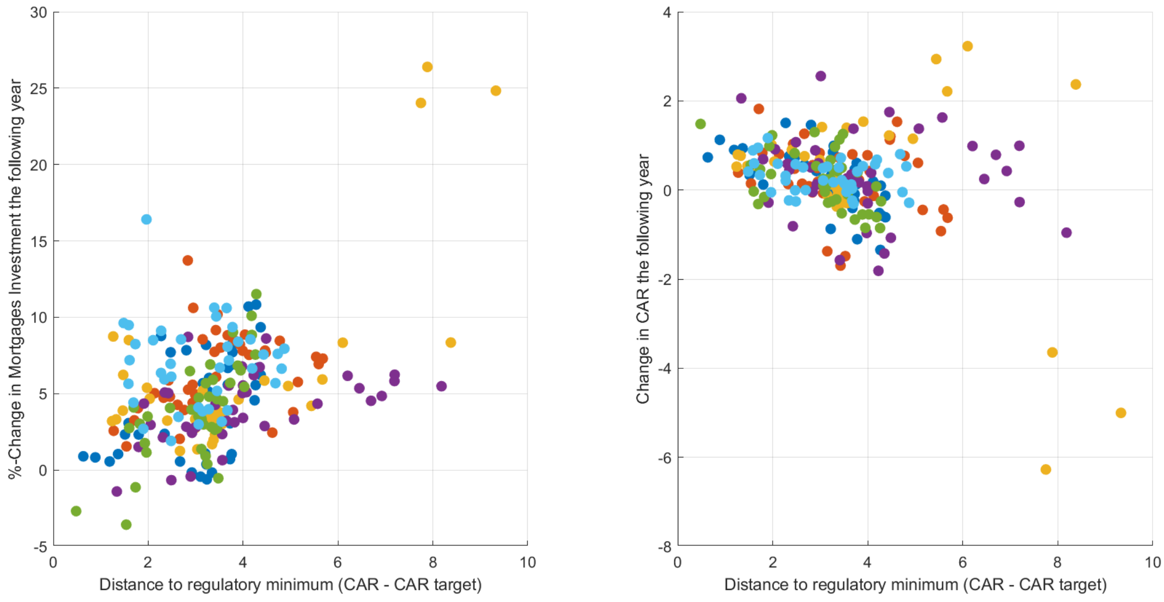

Figure 1: Relationship between distance to the regulatory minimum and bank reactions one year later, for the Big Six banks in Canada (source: OSFI M4 and BCAR reports). Each point represents a bank-quarter observation. Different colours indicate different banks. Left panel: distance to the CAR target versus the percentage change in mortgage holdings. Right panel: distance to the CAR target versus the subsequent change in the regulatory ratio.

This observation poses a fundamental challenge for conventional stress-testing frameworks. Most top-down models project bank outcomes conditional on a fixed or rule-adjusted balance sheet. This modelling choice buys tractability, but in adverse conditions it brings two problems. First, correlation structures shift: the reduced-form relationships estimated in normal times may break down precisely when they are most needed, a concern familiar from the literature on model instability under stress (Aikman et al., 2023). Second, and more fundamentally, banks do not hold their portfolios fixed. Under a prolonged downturn, institutions reallocate assets, adjust funding structures, manage their regulatory buffers, and respond to the actions of their peers. The latter is feedback that static models, by construction, cannot capture. The patterns in Figure 1 suggest that these strategic responses are not idiosyncratic noise, but rather systematic functions of the regulatory environment.

This paper develops the Banks' Strategy Model (BSM) to ask (i) how strategic portfolio responses by banks propagate and amplify adverse conditions, and (ii) what that implies for the design and evaluation of macroprudential policy. The BSM is a structural partial equilibrium framework in which each of Canada's Big Six banks maximises a utility function trading off expected return on equity against proximity to regulatory constraints.[1] Three features distinguish it from existing top-down frameworks. Rather than treating regulatory constraints as hard bounds that trigger adjustment only when breached, the BSM uses log-barriers: a smooth penalty that increases continuously as buffers shrink, generating endogenous managerial cushions above the regulatory floor without discontinuities or corner solutions. Banks optimize jointly over assets and liabilities, allowing them to react to capital pressure with their funding structure and their asset holdings, rather than relying solely on asset sales. All six banks optimize simultaneously, which generates price pressure from aggregate trading. This pressure feeds back into each institution's capital position, ultimately cleared by a Nash equilibrium each period. The model is dynamic by design, such that small shocks accumulate and erode capital buffers gradually.[2] Together these features allow the BSM to serve three macroprudential purposes within a single consistent structure: system resilience testing, counterfactual policy analysis, and reverse stress testing.

Three applications illustrate what the framework delivers. First, simulating 1,000 macro-financial scenarios reveals a sharp nonlinearity: aggregate credit provision is relatively stable as long as banks hold comfortable buffers, but contracts abruptly once the minimum CET1 ratio among the Big Six falls below approximately 50bps above the regulatory minimum. Second, a counterfactual analysis of an increase in the minimum

[1] The concentrated Canadian banking system, in which the Big Six hold more than 93% of total banking assets, provides an ideal setting for introducing a model in which agents are market movers and therefore make strategic adjustments to their balance sheets.

[2] This is consistent with the empirical observation that systemic solvency crises are rarely the product of a single large shock but rather the endpoint of a slow deterioration in fundamentals. For instance, losses at Lehman Brothers were already visible in aggregate balance sheet data more than a year before its failure in September 2008 (Hurley and Hurley, 2013).

capital requirement from a change in the countercyclical capital buffer from 3% to 3.5% shows a reduction in aggregate lending by approximately 1.2%, an effect comparable in magnitude to a 25 basis point increase in the policy rate. This highlights the trade-off between resilience and credit supply that macroprudential authorities face when calibrating buffer requirements. Third, in a reverse stress test, we recover the macro-financial environments consistent with systemic stress without conditioning on any prior narrative. That is, we simulate 10,000 macroeconomic return paths and pass the balance sheet return components to the BSM. We select the stress outcomes and examine the associated macro variables to characterize the scenarios. In the set of stressed scenarios, we see three distinct clusters: a conventional GDP slowdown, a house price correction, and a Dutch disease scenario in which Canadian dollar appreciation alongside rising oil prices supports aggregate output while eroding business loan quality. The last two clusters state vulnerabilities that conventional scenario design, anchored to recessionary narratives, would systematically miss.

This paper relates to the macroprudential stress-testing literature. A well-established concern is that single deterministic adverse scenarios may not identify the full range of relevant vulnerabilities (Borio et al., 2014; Breuer and Summer, 2020), and that structural approaches are needed to generate multiple plausible stress narratives (Aikman et al., 2024).[3] Aikman et al. (2023) document that reduced-form approaches generate projections that become increasingly unreliable over multi-year horizons as banks adapt their portfolios. Large-scale top-down frameworks for the euro area (Budnik et al., 2019) and Canada (Abdelrahman et al., 2025) extend the range of modelled transmission channels but retain the static balance sheet assumption at the portfolio level — the assumption the BSM is designed to relax, at the cost of less granularity. Within the Bank of Canada's toolkit, FRIDA (MacDonald and Traclet, 2018; Gaa et al., 2019) and MFRAF (Fique, 2017) provide the institutional context the BSM is designed to complement. The case for incorporating such responses is supported by the empirical literature on capital buffer management, which documents that banks maintain cushions above the regulatory minimum and adjust their asset and liability mix in response to their proximity to the constraint (Orozco and Rubio, 2025; Juselius et al., 2025). The paper further contributes to the literature on reverse stress testing (Baudino et al., 2018; Breuer et al., 2009; Baes and Schaanning, 2023), which argues that the identification of vulnerabilities should not be conditioned on the analyst's prior about the source of stress.

The BSM also sits within the structural bank modelling literature that incorporates strategic interaction and endogenous asset pricing. The macro-finance literature on financial intermediaries (Gertler and Kiyotaki, 2010; He and Krishnamurthy, 2013) establishes

[3]For instance, Valderrama (2026), Schuermann (2026), and Aikman (2026) highlight key drawbacks of traditional stress testing frameworks, including limited scenario coverage, overly optimistic modelling assumptions, and insufficient treatment of systemic interactions, while pointing toward improvements through broader scenario exploration and more diagnostic approaches.

that balance sheet constraints interact with asset prices to generate amplification absent from frictionless models, and Brunnermeier and Pedersen (2009) formalize how funding liquidity and market liquidity reinforce each other under stress. The connection between financial conditions and macroeconomic tail risks is further formalized by Adrian et al. (2019), who show that adverse financial shocks significantly widen the left tail of the GDP growth distribution. The fire-sale literature (Shleifer and Vishny, 2011; Caballero and Simsek, 2013) characterizes the conditions under which constrained liquidation depresses prices and imposes losses on other institutions. The BSM captures these mechanisms at the level of individual institutions, deriving equilibrium prices endogenously from the collective order flow of the six banks following Kyle (1985). This approach is well-suited to Canada's concentrated banking system, where the Big Six hold over 93% of total assets (Friedrich et al., 2025) and large institutions have been shown to internalize their market impact (Goldstein et al., 2017). In the agent-based modelling tradition more broadly (Hałaj, 2018; Sydow et al., 2024; Bruneau et al., 2025), fire-sale dynamics are typically triggered by threshold crossings and transmitted through fixed price-impact rules. The closest methodological paper is Halaj and Priazhkina (2021), who model bank portfolio optimisation as a simultaneous strategic game cleared by a Nash equilibrium; the present framework extends that approach by incorporating liabilities into the choice set and replacing hard regulatory constraints with smooth log-barriers that generate endogenous managerial buffers.

The remainder of the paper is structured as follows. Section 2 presents the utility functions, optimization problems, and equilibrium conditions. Section 2.3 describes the data, parameter estimation, and transition dynamics between periods. Section 2.6 outlines the scenario generator used to construct multivariate paths for balance sheet returns. Section 3 presents the three applications. Section 4 concludes.

2 The Banks' Strategy Model

Before presenting the formal model, it is worth stating the three economic tensions it is designed to resolve. The first tension is between return and regulatory distance. A bank wants to deploy capital toward higher-returning assets, but doing so consumes regulatory capital and liquidity, shrinking the buffer between its current position and the regulatory floor. The standard modelling response is a hard constraint, i.e., banks optimize freely until the constraint binds, then stop. This produces a discontinuity that is empirically counterfactual. Figure 1 shows a positive relationship between mortgage investment growth and distance to the regulatory minimum in the previous period. This suggest that banks begin adjusting well before any regulatory constraint binds, e.g., a bank two percentage points above its minimum behaves differently from one four points above, even though neither faces imminent breach. The utility function below captures this graduation through a log-barrier. The second tension concerns the margin of adjustment. A bank facing capital pressure can not only sell assets, but also can restructure its funding by reducing wholesale borrowing, shifting deposit mix, and managing leverage directly. Restricting the choice set to assets alone forces all regulatory adjustment onto asset sales, overstating deleveraging and suppressing the funding channel. Here, both assets and liabilities enter the optimization. The third tension is strategic. Canada's Big Six hold over 93% of domestic banking assets. When several banks simultaneously reduce their security exposure, the resulting selling pressure not only depresses prices, but also creates investment opportunities. We address this by making each bank's expected return a function of aggregate order flow. The Nash equilibrium that results makes this amplification channel endogenous rather than imposed.

2.1 Banking sector

Consider B banks with a portfolio consisting of N assets and liabilities. In Halaj and Priazhkina (2021), the banks optimize expected returns while accounting for regulatory constraints. In contrast, we assume banks are risk neutral in profits and drop the variance term from the mean-variance optimization, such that banks are optimizing only over their portfolio choices subject to their regulatory ratios.[4] Lastly, we assume that banks re-optimize their portfolio at each point in time. Put differently, their optimization per period is done through changes in portfolio composition and not a re-optimization of all weights.

Bank b at time t chooses a portfolio adjustment $x_b \in \mathbb{R}^N$ to solve

$$\begin{aligned} \max_{x_b} \quad u_{b,t}(x_b) &= E_{b,t}[r_{t+1}]'x_b + \lambda_b' log(\tau_{b,t}(x_b) - \tau_{min,t}) - g(x_b), \\ \text{s.t.} \quad \sum_{i \in \mathcal{A}} x_{i,b} &= \sum_{i \in \mathcal{L}} x_{i,b}, \end{aligned} \tag{1}$$

where x_b is a vector of size N containing the within-period changes in each balance sheet category, and constitutes the choice variable of the bank. The state variables entering the objective are the beginning-of-period balance sheet $bs_{b,t}$, the equity level $e_{b,t}$ (which is predetermined within the period and updated only after returns are realized, see Section 2.5). $E_{b,t}[\cdot]$ denotes the expectation of bank b at time t, which depends on the

[4] We depart from mean-variance optimization for both practical and conceptual reasons. Estimating a well-conditioned covariance matrix over N balance sheet categories requires far more observations than are available in supervisory data, and the resulting estimates are sensitive to the sample period chosen. More fundamentally, mean-variance optimization treats the covariance matrix as a stable sufficient statistic for risk. For banks, however, the relevant risk concentrates in the left tail, where asset correlations behave very differently from their unconditional counterparts. For example, positions that appear well-diversified in normal times become strongly correlated precisely when a common macroeconomic or liquidity shock materializes. The regulatory aversion terms in our utility function address this. By penalizing proximity to the constraint regardless of the covariance structure, they generate precautionary behaviour that is robust to correlation breakdown under stress.

price the bank pays for the portfolio adjustment (see Equation (2) below). λ_b is a bank specific (3×1) vector containing the constraint aversion parameters for the regulatory minima in the vector $\tau_{min,t}$. $\tau_{b,t}(x_b)$ denotes the vector of regulatory ratios, and $g(x_b)$ the cost function of portfolio adjustment. The adjustment conditions follows from the balance sheet identity with $\mathcal{A}$ and $\mathcal{L}$ as the set of asset and liability categories, respectively. Unlike Halaj and Priazhkina (2021), the vector x_b contains all possible choices for the balance sheet, i.e., it also contains the liability side, such that banks can optimize their funding too, allowing to increase their balance sheet by leveraging. The liability side has the opposite sign of the asset side.

Expected Return We assume that banks act as market movers: their collective trading generates price pressure that feeds back into expected returns. We model this through an endogenous price-formation mechanism in which today's price of each asset responds linearly to aggregate net demand, following the framework of Kyle (1985). Future prices remain stochastic and are unaffected by current trades. The resulting expected return for bank b at time t admits a first-order approximation of the form

$$E_{b,t}[r_{t+1}] = \mu_{b,t}(X) \approx \bar{\mu}_{b,t} + \dot{\mu}\, X \mathbf{1}_{B\times 1}, \tag{2}$$

where $\bar{\mu}_{b,t}$ is the vector of baseline expected returns in the absence of trading, the $N \times B$ matrix X collects the bank's strategies in its columns, and $\dot{\mu}$ is a diagonal matrix of return sensitivities to aggregate demand.[5] The derivation and exact characterization of $\dot{\mu}$ are provided in Appendix A.

The intuition behind Equation (2) is straightforward. When banks collectively demand more of an asset, their buying pressure drives up its current price, which compresses the expected return: the same future payoff is now purchased at a higher price today. The reverse holds for assets that banks collectively shed. Moreover, if some banks are buying assets while others are selling by an equal amount, aggregate demand stays unchanged and current prices remain at their baseline p_t^0, leaving expected returns unaffected. This price-feedback mechanism is what makes the banks' optimization problems interdependent: the attractiveness of any given asset depends not only on its fundamental return prospects but also on what all other banks are doing simultaneously. On the liability side, the same logic applies in reverse. It is through this channel that price-mediated spillovers emerge endogenously. In normal conditions, this creates competition for lending opportunities: one bank's expansion compresses spreads for all. Under systemic liquidity stress, when constrained banks simultaneously shed partially-liquid assets to improve their LCR, the collective selling pressure without opportunistic buyers depresses prices, amplifying

[5] Note, while we apply this model to a concentrated market, we cannot capture the full supply side. We assume that marginal traders that also act on the supply side will be reflected in the calibration of $\dot{\mu}$ in the empirical application.

capital losses at other institutions. This is the mechanism closest to fire-sale dynamics in the model, and it requires no separate assumption. Finally, it is worth noting that in this logic other agents, such as households, firms or non-bank financial institutions, also can participate in this price finding mechanic. This step would endogenize $\dot{\mu}$ as a result of the respective counterparties. Section 2.2 explains how we approximate these agents with a reduced form logic.

Regulatory ratios The model generates managerial buffers not by assumption but through the curvature of the regulatory penalty term in Equation 1. More precisely, in order to maintain healthy fundamentals and remain flexible when economic conditions deteriorate, banks aim to hold more capital and liquidity buffer than what the regulator requires. This feature is consistent with empirical observations that banks hold managerial buffers and do not optimize their balance sheet at the regulatory constraint. We model this component with a log-utility function, resulting in minus infinity utility when they breach the regulatory minimum $\tau_{min,t}$. Importantly, that also gives the banks the opportunity to trade higher expected profits for higher regulatory ratios, if the gain in regulatory ratios is sufficiently large. In other words, if two assets produce the same expected return, banks prefer the asset that gives them more buffer over $\tau_{min,t}$, which in return gives them more flexibility in future optimization rounds. Figure 2 illustrates the utility from holding the regulatory ratios above the constraint.

Regulatory constraints Regulatory constraints are enforced by the regulator that set the regulatory minima $\tau_{min,t}$. For the banking sector, these minima are often enforced for a risk-weighted capital ratios, leverage ratios, or liquidity coverage ratios. We focus here on the three distinct ratios, the liquidity coverage ratio (LCR), the leverage ratio (LR), and the capital adequacy ratio (CAR). Under Basel III regulations, banks must comply with multiple ratios that ensure adequate capital, leverage, and liquidity levels. However, for managerial purposes, focusing on the LCR, LR, and CAR is covering a big range of risks, as other ratios typically follow similar sensitivities and adjustments in response to balance sheet changes. We assume that other ratios are not actively managed by bank managers, and managerial buffers over the minimum are a result of spillover effects from managerial buffers from the LCR, LR, or CAR.

The CAR is calculated as follows:

$$CAR(x_b) = \frac{\text{CET1 capital}}{\text{total risk-weighted assets}} = \frac{e_b}{\theta_a'(bs_b + x_b)},$$

where e_b is the equity (CET1 capital) position of a bank, and θ_a is an $(N \times 1)$ vector of risk-weighted asset densities. bs_b depicts the balance sheet at the beginning of the period. Note that equity e_b is treated as predetermined within each period. The optimization in

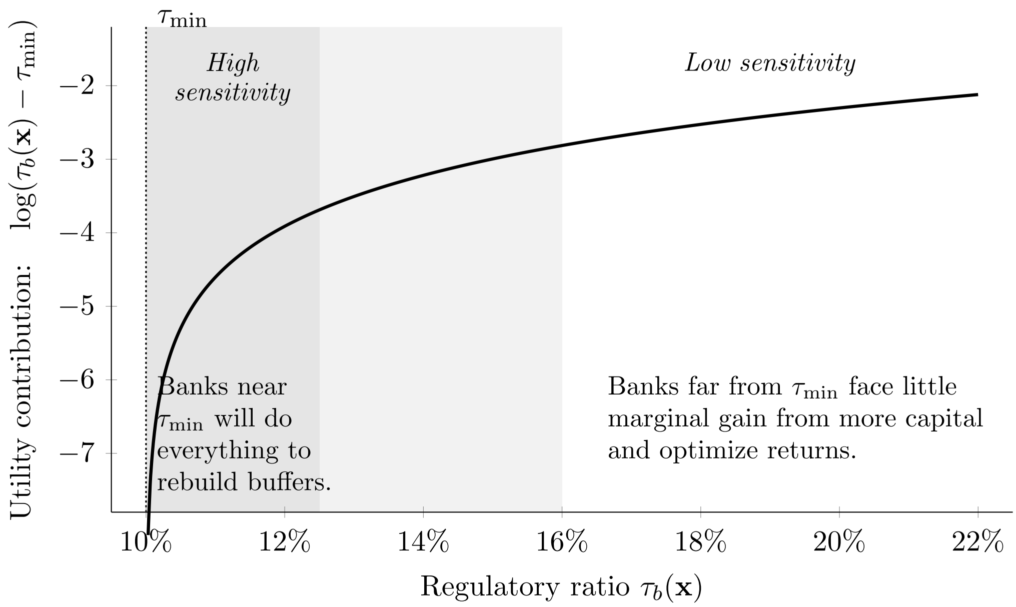

Figure 2: Utility contribution from the regulatory buffer. The curve plots $\log(\tau_b(x_b) - \tau_{\min})$ as a function of the regulatory ratio $\tau_b(\mathbf{x})$, with $\tau_{\min} = 10\%$. Near the regulatory minimum, the marginal utility of additional buffer is large: small improvements in the ratio yield substantial gains, creating a strong incentive to de-risk or raise capital. As the ratio increases, the curve flattens and additional buffer contributes little to no utility, allowing banks to allocate capital toward higher-returning assets. This log-shape generates endogenous managerial buffers: banks maintain cushions above the minimum, but do not accumulate buffer indefinitely.

x_b represents a rebalancing of the portfolio composition—exchanging assets and liabilities at current market prices—which does not alter equity until returns are realized at the end of the period. This is consistent with the sequential timing of the model described in Section 2.5, where equity is updated only after net income, dividends, and taxes have been computed.

The LCR is calculated as follows:

$$LCR(x_b) = \frac{\text{High quality liquid assets}}{\text{30 days stressed net cash outflow}} = \frac{\phi_a'(bs_b + x_b)}{\phi_f'(bs_b + x_b)},$$

with ϕ_a and ϕ_f being $(N \times 1)$ vectors with weights depicting liquidity for HQLA and stressed net-outflow for funding, respectively. Note that ϕ_a will have entries only at the position of assets in vector bs_b, and ϕ_f has negative weights (for inflows) at asset positions and positive weights (for outflows) at liability positions. We assume throughout that the net stressed cash outflow $\phi_f'(bs_b + x_b) > 0$ for all feasible portfolio choices x_b, which holds for all banks across all simulated periods given the composition of their balance sheets and the bounds implied by the quadratic trading cost term.

The LR is calculated as follows:

$$LR(x_b) = \frac{\text{Tier 1 capital}}{\text{total exposure assets}} = \frac{e_b + \Delta Tier1_b}{(bs_b + x_b)},$$

where e_b is the equity (CET1 capital) position of a bank and $\Delta Tier1_b$ is the difference of CET1 capital and Tier 1 capital.

We collect the regulatory ratios in the vector $\tau_b(x_b) = (CAR(x_b), LCR(x_b), LR(x_b))'$ and respectively define the minimum set by the regulator as $\tau_{min,t} = (CAR_{min,t}, LCR_{min,t}, LR_{min,t})'$.

Trading costs Finally, $g(x_b)$ is a cost term disciplining rapid balance sheet adjustments. We assume this cost term to be of quadratic form, i.e.,

$$g(x_b) = x_b' \gamma x_b,$$

with γ being a diagonal matrix with positive entries of size $(N \times N)$. While this quadratic form may not represent the true form of trading-costs, it is standard in the literature on portfolio adjustment and trading frictions (e.g., Gârleanu and Pedersen, 2013), keeps the problem computationally tractable, and regulates trades between banks. The positive set of parameter in γ is then estimated on past observations to match trading behaviour.

Normalization by equity To account for size effects, each bank's strategy is normalized by its equity. This implies that banks optimize returns relative to their own size—specifically, they maximize return on equity rather than absolute returns. This normalization also informs the specification of trading costs. The cost function $g(x_b)$ is assumed to be scale-sensitive, meaning that trading costs are proportional to the bank's equity. Conceptually, this reflects the idea that larger banks have the same cost of equity as smaller banks, implying that they typically benefit from more extensive infrastructure, greater staffing capacity, and broader market access, all of which reduce the per-unit cost of executing larger trades.

The resulting optimization problem With the price impact function from the Taylor approximation, the utility function for bank b reads as

$$u_{b,t}(x_b) \approx \bar{\mu}_{b,t}' x_b + (X \mathbf{1}_{B \times 1})' \dot{\mu} x_b + \lambda_b' log(\tau_{b,t}(x_b) - \tau_{min,t}) - x_b' \gamma x_b, \tag{3}$$

and the optimal strategy in the limiting case of ample regulatory buffers, can be derived by

$$\begin{aligned}\frac{\partial u_t}{\partial x_b}|_{\tau(x_b)\gg\tau_{\min}} &= \bar{\mu}_{b,t} + 2(\dot{\mu}-\gamma)x_b + \dot{\mu}\sum_{i\neq b} x_i + \lambda_b' \frac{\partial \tau_{b,t}(x_b)}{\partial x_b} \underbrace{(\tau_{b,t}(x_b) - \tau_{\min,t})^{-1}}_{\to 0 \text{ for large } \tau(x_b)} = 0,\\ \Rightarrow x_b &= -\frac{1}{2}(\dot{\mu}-\gamma)^{-1}(\bar{\mu}_{b,t} + \dot{\mu}\sum_{i\neq b} x_i),\end{aligned}$$

This expression has a direct economic reading. The first term says: hold the asset with the highest return, discounted by market depth and trading costs. The second term says: partially absorb the imbalance created by your peers — because their selling pressure moves prices in your favour. The log-barrier term vanishes when buffers are large, but re-enters sharply as a bank approaches its constraint, at which point the portfolio tilts toward lower-risk assets regardless of their return. The interaction between these forces is what generates the nonlinear credit dynamics documented in Section 3.

The equilibrium is the result of a quadratic game involving B players, where each bank chooses strategies from a convex set. Each bank can find its optimal strategy by solving a constrained optimization problem, following the Kuhn–Tucker conditions that ensure the constraints are properly accounted for.[6]

2.2 Non-bank sector representation

The price-impact matrix $\dot{\mu}$ summarises how aggregate bank trading moves prices on both sides of the balance sheet. On the asset side, simultaneous credit expansion compresses loan spreads as banks compete for the same pool of borrowers. On the liability side, collective deposit-gathering raises funding costs as banks bid against each other for savings. Both effects enter the expected return through $\dot{\mu}$, consistent with the oligopolistic banking framework of Klein (1971), Monti (1972), and Freixas and Rochet (2008).

Moreover, aggregate lending shifts deposit supply through the money creation mechanism: when the banking sector extends credit, the corresponding funds are redeposited within the system, expanding the aggregate deposit base independently of any individual bank's funding decision (Tobin, 1963; Jakab and Kumhof, 2015). This supply shift enters the price-formation channel in the same way as strategic trading does. The price-relevant aggregate order flow is therefore $X\mathbf{1}_{B\times 1} + \Delta x^{xa}$, where Δx^{xa} denotes the deposit inflows induced by aggregate lending, and the expected return adjusts accordingly:

$$\mu_{b,t}(X) \approx \bar{\mu}_{b,t} + \dot{\mu}\left(X\mathbf{1}_{B\times 1} + \Delta x^{xa}\right). \tag{4}$$

[6] More details on equilibrium conditions can be found in Appendix B.1 or in Equations (19) and (20) of Halaj and Priazhkina (2021).

The deposit supply shift therefore modifies the baseline return from which banks optimize: the no-trade reference point shifts from $\bar{\mu}_{b,t}$ to $\bar{\mu}_{b,t} + \dot{\mu}\,\Delta x^{xa}$, compressing expected returns on deposits and relaxing them on loans relative to the strategic-trading-only baseline.

The induced deposit inflows are governed by a sparse transition matrix $T \in \mathbb{R}^{N \times N}$ that maps aggregate lending into deposit creation:

$$\Delta x^{xa} = T(X\mathbf{1}_{B\times 1}) + \Delta x^{ex}, \tag{5}$$

where Δx^{ex} captures exogenous balance sheet growth along the baseline path. For instance, we assume business loan demand and deposit supply to grow with historical averages (2% and 1.5% each quarter, respectively). The matrix T is non-zero only where loan creation maps to deposit inflows, with rows indicating deposit categories and columns loan categories, and is calibrated by assuming that mortgages origination generates mainly retail deposits and loans creation generates mainly corporate demand deposits. [7] The resulting inflows are distributed across banks according to market shares ω_b, giving $\Delta x_b^{xa} = \omega_b \cdot \Delta x^{xa}$. Banks take Δx_b^{xa} as given when optimizing within the period, consistent with the timing assumption that deposit creation follows rather than precedes the lending decision (Bernanke and Blinder, 1988).

To illustrate this transition matrix, we build T for Canada's Big Six with a non-zero submatrix spanning loan and deposit categories. All remaining rows and columns are zero, reflecting that only loan creation generates induced deposit inflows.

$$T_{\text{loans}\to\text{deposits}} = \begin{array}{l} \\ \text{Retail} \\ \text{Demand Corp.} \\ \text{Term Corp.} \end{array} \begin{array}{c} \begin{array}{ccc} \text{Non-Bus.} & \text{Business} & \text{Mortgages} \end{array} \\ \begin{pmatrix} t_{00} & t_{01} & t_{02} \\ t_{10} & t_{11} & t_{12} \\ t_{20} & t_{21} & t_{22} \end{pmatrix} \end{array},$$

where each entry $t_{ij} \in [0, 1]$ represents the fraction of a unit expansion in loan category j that flows into deposit category i, with $\sum_i t_{ij} \leq 1$ for each j. Non-business loan creation generates mainly retail and demand corporate deposits; business loan creation generates mainly demand and term corporate deposits; mortgage origination generates mainly retail deposits, reflecting the household nature of the borrower. The constraint $\sum_i t_{ij} \leq 1$ allows for leakage outside the Big Six system.

[7] In a closed economy, each unit of loan creation is mirrored by deposit creation within the banking system, implying that each column of T sums to one. In an open economy, this linkage is weaker, as part of the funding may be held by non-residents or leak outside the modelled banking system. To capture this effect, we scale the matrix T by 65%, reflecting the estimated share of deposit creation that remains within the domestic system.

2.3 Balance sheet structure

The balance sheet structure of the model aggregates data from various sources to capture the financial positions of banks comprehensively. The key balance sheet components are categorized into assets, liabilities, and equity, as shown in Table 1. This structure is designed to align with the reporting standards of Canadian D-SIBs (Domestic Systemically Important Banks) and facilitates the modelling of systemic risk dynamics.

Assets		**Liabilities & Equity**	
Category	**Code**	**Category**	**Code**
Cash	O0	Deposits – Retail	F0
Reserves	O2	Deposits – Demand Corporate	F1
Foreign Reserves	O3	Deposits – Term Corporate	F2
Government Securities	S0	Repo	F3
Debt Securities	S1	Other Liabilities	F4
Shares	S2		
Reverse Repo	S3	Total Shareholder Equity	E0
Non-Business Loans	L0		
Business Loans	L1		
Mortgages	L2		
Other Assets	O1		

Table 1: Balance sheet structure of the model. Asset and liability exposures are sourced from the OSFI M4 report, a quarterly disclosure of holdings by Canadian D-SIBs. Equity is adjusted to reflect CET1 capital from the OSFI BCAR report, which also provides risk-weighted assets used to construct the capital adequacy ratio. High-Quality Liquid Asset and stressed cash flow data, used to construct the liquidity coverage ratio, are sourced from OSFI liquidity returns. Codes: S for securities, L for loans, O for other balance sheet items, F for funding, and E for equity.

We source the exposure to the different balance sheet items from the M4 report, which is a quarterly disclosure of holdings by D-SIBs in Canada. We adjust equity to fit the CET1 capital from the BCAR report provided by OSFI (Office of the Superintendent of Financial Institutions). We further obtain the risk-weighted assets (RWA) from the BCAR report, which we use to replicate the CAR. To replicate the LCR, we rely on High-Quality Liquid Assets (HQLA) and inflow/outflow data issued by OSFI. Finally, returns expectations play a key role in driving the banks' investment decisions as shown by Equation 1. Those return expectations rely on past returns, which are obtained from the P3 report. Some summary statistics on key metrics are reported in Table 2.

For modelling purposes, this data spans from 2015Q1 to 2025Q3, where the starting date coincides with the first reporting of the LCR and LR. The interbank rate, which is used to evaluate funding costs and market dynamics, is sourced directly from the Bank of Canada.

	Average	Minimum	Maximum
Total Assets (bn CAD)	1381.95	453.93	2061.61
Leverage Ratio (%)	4.97	4.80	5.16
CAR Ratio (%)	13.16	12.84	13.46
CAR Minus Regulatory Requirement (%)	1.66	1.34	1.96
Liquidity Coverage Ratio (%)	1.38	1.31	1.59
Return on Equity (%)	14.73	10.33	19.40

Table 2: Summary statistics for the last time point 2024-07-31 across all banks.

2.4 Estimation of Parameters

Three sets of parameters enter the banks' utility function and require estimation: the price sensitivity matrix $\dot{\mu}$, the regulatory aversion vectors λ_b, and the trading cost matrix γ. We estimate these jointly using the Simulated Method of Moments (SMM, Duffie and Singleton (1993); Gallant and Tauchen (1996)), matching simulated portfolio strategies to their observed counterparts over the period 2015Q1 to 2025Q3.

The empirical target is the strategic portfolio adjustment of each bank in each quarter. We construct this by taking the observed change in each balance sheet category. The residual captures the discretionary rebalancing decision and serves as the empirical counterpart to the model's optimal strategy x_b. We denote this observed strategic adjustment $\tilde{x}_{b,t} \in \mathbb{R}^N$ for bank b at time t, so that the full vector of empirical moments stacks these across all six banks and all $N = 16$ balance sheet categories into $m_{\text{obs}} \in \mathbb{R}^{6N}$. Regulatory minima $\tau_{\min,t}$ are taken directly from OSFI reporting and vary over time, ensuring that the distance to the constraint at each quarter accurately reflects the regulatory environment each bank faced.

Given a candidate parameter vector $\Theta = (\dot{\mu}, \lambda_b, \gamma)$, the model produces a simulated portfolio strategy $x_b(\Theta)$ for each bank-quarter observation. Baseline expected returns are set equal to the realized returns from the previous period, $\mathbb{E}_{b,t}(r_{t+1}) = \bar{\mu}_{b,t} = r_{b,t}$, reflecting extrapolative expectations: banks project that recent return experience will persist into the next period. This assumption draws on evidence that investors and managers tend to expect trends to continue, rather than making model-consistent forecasts. By adopting a simple adaptive rule for forecasting returns, we capture a form of bounded rationality in banks' decision-making. This approach serves as a reduced-form proxy for how banks update beliefs under uncertainty, and it lies between no forward-looking static expectations and full rational expectations.

The SMM then selects the parameter vector that minimizes the distance between simulated and observed portfolio adjustments,

$$\hat{\Theta} = \arg\min_{\Theta} \left(m_{\text{obs}} - m_{\text{sim}}(\Theta)\right)' W \left(m_{\text{obs}} - m_{\text{sim}}(\Theta)\right),$$

where $m_{\text{sim}}(\Theta)$ collects the simulated counterparts to m_{obs} under parameter vector Θ.

The moment vector has dimension $6N = 96$, which renders a full-covariance weighting matrix infeasible to estimate reliably from the available sample. We therefore restrict W to be diagonal, with entries given by the inverse of the sample variance of each moment. This choice preserves the efficiency gains from down-weighting noisily measured moments while avoiding the estimation error that would accumulate in a dense 96×96 matrix. The minimization is carried out using the particle swarm algorithm (Kennedy and Eberhart, 1995), which is well suited to the high-dimensional, non-convex parameter space arising from the combination of bank-specific aversion parameters and the shared price sensitivity and trading cost matrices.

The three parameter sets are identified from distinct features of the observed adjustments $\tilde{x}_{b,t}$. The price sensitivity matrix $\dot{\mu}$ governs the extent to which banks moderate their trades in response to aggregate order flow: when many institutions move in the same direction simultaneously, a higher $\dot{\mu}$ implies more restraint. The regulatory aversion vector λ_b is identified from the relationship between a bank's distance to its regulatory minimum and the composition of its rebalancing: banks that shift systematically toward lower-risk assets as buffers shrink imply higher aversion. The trading cost matrix γ is identified from the speed of adjustment: banks that close the gap between their current and optimal portfolio slowly, relative to their return incentives, imply higher quadratic adjustment costs.

2.5 Period Transition Dynamics

After the optimization process in the BSM, banks transition to the next period by updating their financial positions based on the newly optimized balance sheet structures. This transition involves the following key steps:

1. Calculation of Net Interest Income (NII): Given an exogenous scenario of returns r_{t+1}, banks receive returns on their assets according to the new balance sheet structure.

$$NII_{b,t+1} = (bs_{b,t} + x_{b,t})'(r_{t+1} + \dot{\mu} X \mathbf{1}_{B\times 1}).$$

   Where bs_t is the balance sheet at time t, and x_t is the changes in balance sheet structure played by bank b. $\dot{\mu} X 1_{B\times 1}$ is the endogenous price changes from the game.

2. Calculation of Net Income (NI):

$$NI_{b,t+1} = NII_{b,t+1} - \text{Fixed Costs}_{b,t+1}.$$

   where fixed costs are estimated in relationship with $equity_{b,t+1}$, i.e., using the coefficients of the regression $\text{FixedCosts}_{b,t} = c_b + \beta_b \cdot equity_{b,t} + \varepsilon_{b,t}$. Note that trading income is included in the return in the assets in the calculation of NII.

3. Dividend payments: Banks distribute dividends to shareholders based on historical dividend payout ratios (ranging from 44 to 51%).[8] This payout ratio is applied to the Net Interest Income to determine the amount allocated for dividends.

$$\text{Dividends}_{b,t+1} = \text{dividend rate}_b \cdot NI_{b,t+1}.$$

4. Taxation: Taxes are subtracted from the remaining income after dividends are paid. The empirical tax rate (ranging from 17% to 21%) is applied to the pre-tax income (i.e., Net Interest Income minus dividends) to determine the tax obligation.

$$\text{Taxes}_{b,t+1} = \text{tax rate}_b \cdot (NI_{b,t+1} - \text{Dividends}_{b,t+1}).$$

5. Retained earnings: The leftover income after taxes are subtracted is allocated to retained earnings. These retained earnings are added to the bank's equity, bolstering the bank's capital base for the next period.

$$\text{Retained Earnings}_{b,t+1} = \Delta e_{b,t+1} = NI_{b,t+1} - \text{Dividends}_{b,t+1} - \text{Taxes}_{b,t+1}.$$

6. Asset side adjustments: securities are growing with return ($r_t^{Sec} bs_{b,t+1}^{Sec}$) and amortizing assets (AA) reduce their effective exposure with repayments of the loans ($-r_t^{AA} bs_{b,t+1}^{AA}$). The residual will be distributed to cash, which banks can use freely next period to allocate in their portfolio decision.[9]

$$\Delta cash_{b,t+1} = \Delta e_{b,t+1} - r_t^{Sec} bs_{b,t+1}^{Sec} + r_t^{AA} bs_{b,t+1}^{AA}.$$

2.6 Exogenous scenario

The BSM takes as its only exogenous input a time path of returns for each balance sheet category: the rate earned on mortgages, business loans, government securities, deposits, and so on across the simulation horizon. For each simulated path w and each quarter h, the vector r_{t+1}^w collects the realized returns across all sixteen balance sheet categories, which enter the NII calculation in Step 1 of the period transition (Section 2.5). All other variables in the scenario, GDP growth, house prices, the exchange rate, oil prices, and credit spreads, travel alongside the simulation but play no direct role in the bank optimization. Their role is interpretive: once the BSM identifies which scenarios produce systemic stress, we examine the associated macro-financial paths to recover the stress narratives in Section 3.

[8] Note that pausing dividend payments in the Canadian banking market is historically an exogenous choice done by OSFI rather then an endogenous reaction of the bank itself.

[9] An illustration of the adjustments can be found in Appendix C.

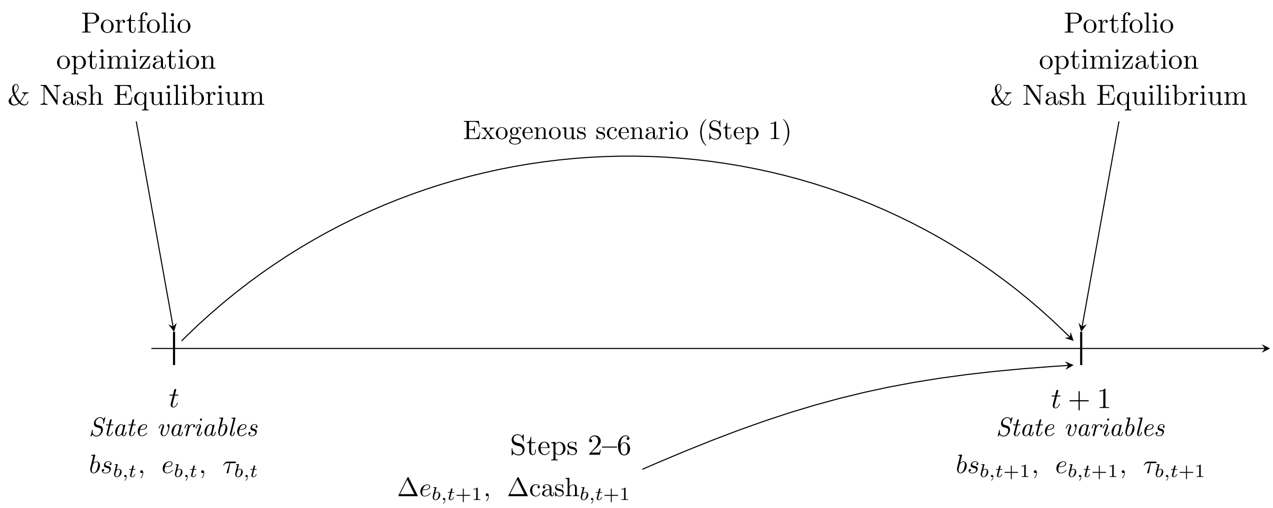

Figure 3: Period structure of the BSM. At each t, banks simultaneously optimize their portfolio for the next period $t+1$ and the Nash equilibrium is computed. The exogenous scenario feeds into the post-trade balance sheet (Step 1); Steps 2–6 update equity and cash, delivering the state at $t + 1$.

We generate these joint paths using the Multivariate Scenario Generator (MSG), summarized in Algorithm 1. The MSG builds on residual bootstrapping: for each variable, we remove autocorrelation and heteroskedasticity through standard filters (log differences for macro aggregates, arithmetic returns for financial series), bootstrap the full residual vector across variables, and reconstruct simulated paths by feeding the resampled residuals back through the inverse filter. Because the MSG samples residual vectors directly from the data rather than fitting a parametric model, the dependence structure requires no estimation. This means the approach avoids imposing assumptions on tail behaviour, such as those embedded in a Gaussian or Student-t copula, and instead lets the data determine how correlations behave across the full distribution. Critically, correlations during historically stressed periods are preserved as they were observed: tighter, asymmetric, and often stronger in the left tail than unconditional estimates would suggest. This feature is central to the BSM applications, where the plausibility of joint stress across institutions depends on whether the scenario generator can reproduce the kind of co-movement that historically accompanied adverse conditions.

Starting from the last observed level of each series, we run the procedure W times over a horizon of H quarters, producing a panel of internally consistent macro-financial paths. Figure 4 illustrates the resulting distribution of forward paths for a representative variable.

3 Applications

The BSM is designed to be used in three complementary ways, each of which we illustrate using Canada's Big Six banks. We begin with the most natural application: simulating

Algorithm 1 Multivariate Scenario Generator

Initialization:

- Set the number of scenarios, W, and the time horizon, H.
- Define the filter and transformation function $f_i(\dots)$ for each variable i in the set of M variables, such that generates a vector ϵ_i i.e. $\epsilon_i = f_i(r_i)$, without trend, changes in variance or autocorrelation.
- Use a smoothing function $G_t(\dots)$ to fix potential outliers, e.g. COVID outbreak, at time t in the matrix ϵ_t, i.e. $G_t(\epsilon_t) = \varepsilon_t$.
- Set initial levels, $r_t^{(0)}$, for all variables.

for $w \in \{1, \dots, W\}$: **do**

for $h \in \{1, \dots, H\}$ **do**

k ← DiscreteUniform(1,T). ▷ T: sample length (quarters); k: sampled quarter index

for $i \in \{1, \dots, N\}$ **do**

$r^w_{i,h} = f_i^{-1}(\varepsilon_{i,k} | r^w_{i,0:h-1})$

end for

end for

end for

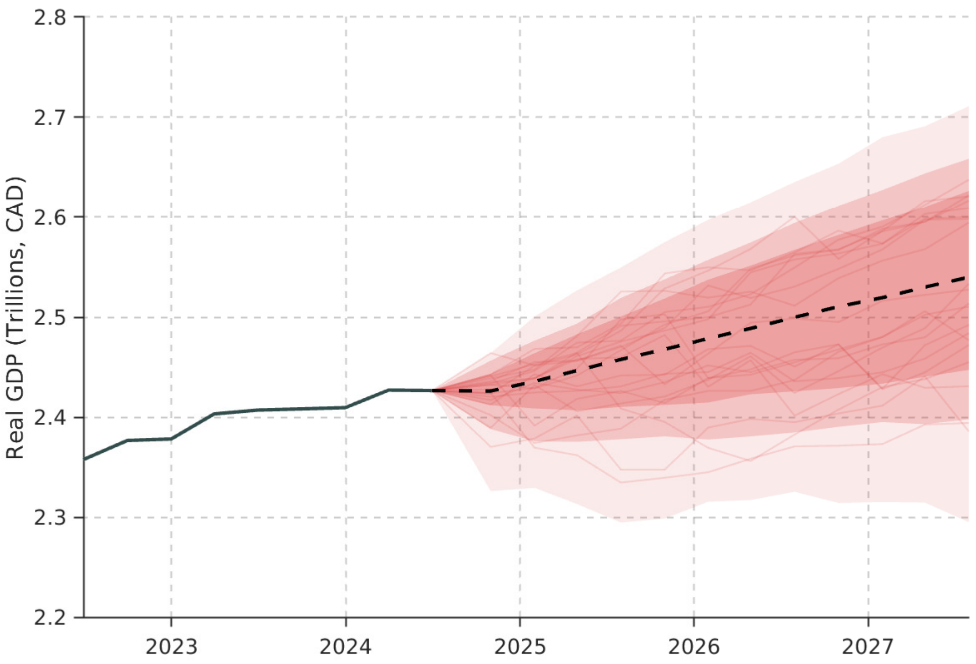

Figure 4: Illustration of the simulated pathways of a variable. Note that the different paths might correspond to different quantiles over time.

the distribution of banking outcomes across a large set of macro-financial scenarios to characterize how aggregate credit provision responds to capital erosion and to identify the threshold at which the system transitions from resilient to actively contracting. We then use the model as a policy laboratory, asking how a specific regulatory change, here an increase in the Domestic Stability Buffer from 3% to 3.5%, affects aggregate credit supply, and quantifying that effect relative to a standard monetary policy benchmark. Finally, we invert the question entirely: rather than asking what outcomes follow from a given scenario, we ask which scenarios are linked to systemic stress, recovering the distribution of macro-financial paths that bring multiple institutions simultaneously close to their regulatory minimum. Together, the three applications constitute a consistent analytical toolkit for a macroprudential authority that seeks to assess the resilience of a concentrated banking system.

3.1 Conducting system resilience tests

We simulate 1,000 macro-financial scenarios starting from 2020Q2 and project outcomes three years forward, tracking the effects on aggregate credit provision, mortgages, and business loans. Figure 5 plots the minimum CAR ratio among the Big Six against aggregate credit provision across all simulations. The central finding is a sharp nonlinearity: credit provision is relatively stable across a wide range of capital outcomes, but deteriorates abruptly once the minimum CAR ratio among D-SIBs falls below approximately 9.5%. This threshold effect is not imposed by assumption — it emerges endogenously from the interaction between regulatory aversion and the price-impact channel in the BSM.

The relationship visible in Figure 5 is notably nonlinear, a result that arises endogenously from the model's structure. Across a wide range of capital outcomes (above 9% CAR), aggregate credit provision varies only modestly: banks with comfortable buffers compete for lending opportunities, and the loss of capital at one institution is largely offset by expansion at others. However, once the minimum CAR ratio among D-SIBs falls below approximately 9.5%, a qualitative shift occurs. Credit provision drops steeply and the dispersion of outcomes widens. As a bank's capital ratio approaches the regulatory minimum, the marginal disutility from further lending increases sharply, causing it to pull back from risk-weighted assets. The resulting selling pressure depresses asset prices through the endogenous price-impact channel, eroding the capital positions of the remaining institutions and limiting their capacity to expand lending and absorb the contraction. The system-wide credit crunch is therefore not necessarily the product of simultaneous stress across all institutions — it can be triggered by one, propagated through prices, and amplified by the inability of healthier banks to compensate.

This mechanism has a concrete implication for the distribution of capital buffers across

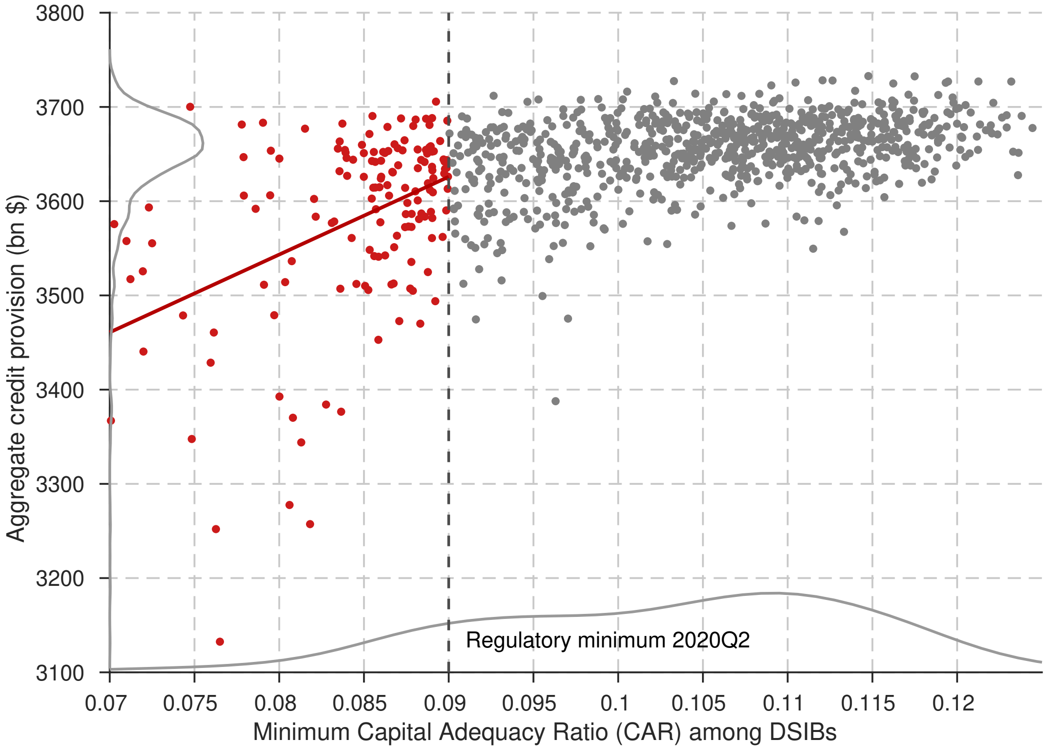

Figure 5: Minimum CAR ratio among D-SIBs against aggregate credit provision across 1,000 simulated scenarios starting from 2020Q2. Red dots indicate the left tail of the CAR distribution, corresponding to scenarios of severe capital stress. The nonlinear relationship reflects the collective deleveraging mechanism of the BSM: above approximately 9.5% CAR, credit provision is relatively stable; below this threshold, simultaneous capital pressure across institutions triggers a self-reinforcing contraction in lending.

the system. A system in which all six banks hold 10.5% CAR is more fragile than one in which five hold 12% and one holds 9%, even if the average is identical: in the former, a common shock pushes all institutions toward the high-sensitivity region simultaneously, whereas in the latter only one institution is vulnerable. This heterogeneity is precisely what the BSM is designed to capture, and what representative-agent models, by construction, cannot. The point is sharpest for the period under study: the simulations begin in 2020Q2, when OSFI had released the DSB in response to the pandemic, compressing the distance between observed capital ratios and the regulatory minimum and making the tipping point identified here directly relevant to conditions that prevailed at the time.

3.2 Conducting policy experiments: Domestic Stability Buffer

A natural question following from the resilience analysis is how regulatory design shapes the system's position relative to the tipping point identified above. In particular, questions around Canada's Domestic Stability Buffer (DSB) are of interest when calibrating its level. It is a macroprudential capital buffer for the country's largest banks (domestic systemically important banks), conceptually analogous to Basel III's countercyclical capital buffer. It is set by the federal banking regulator and can be raised or lowered over the financial cycle to address systemic vulnerabilities, ensuring these banks hold additional capital as risks evolve. In June 2023, OSFI announced an increase in the DSB from 3% to 3.5%, citing the need to strengthen resilience against cyclical vulnerabilities in Canada's banking system. We use this policy change as the basis for a counterfactual experiment, asking how aggregate credit supply would have differed had the higher buffer been in place from the start of the simulation period. Because bank behaviour is derived from explicit utility maximization over a constraint set, changing a parameter produces counterfactual predictions that are internally consistent with the same underlying preference structure. [10]

The exercise's design is straightforward to conduct within the BSM. We generate a set of macro-financial scenarios using the procedure in Section 2.6 and apply the model twice under identical conditions, varying only the regulatory floor. In the baseline, the DSB stands at 3%, implying an effective CAR minimum of $\tau^{\min}_{\text{base}}$. In the counterfactual, the DSB rises to 3.5%, raising the effective floor by 50 basis points to $\tau^{\min}_{\text{cf}} = \tau^{\min}_{\text{base}} + 0.005$. The DSB enters the BSM directly through the log-barrier term in each bank's utility function: a higher floor shifts the high-sensitivity region of the utility curve outward, increasing the marginal disutility of lending for all six institutions simultaneously. The policy impact is then estimated as the difference in average credit provision between the two sets of

[10] This limits, though does not fully resolve, the Lucas critique: unlike reduced-form models whose estimated coefficients may shift arbitrarily under a policy change, the BSM restricts how behaviour can change to what is consistent with the same utility specification. While preference functions may not be invariant to policy changes, the BSM is still suited to evaluate marginal policy changes.

simulations across the full scenario distribution, holding the macro-financial environment fixed. This design isolates the effect of the regulatory change from any difference in economic conditions, giving the experiment a clean counterfactual interpretation.

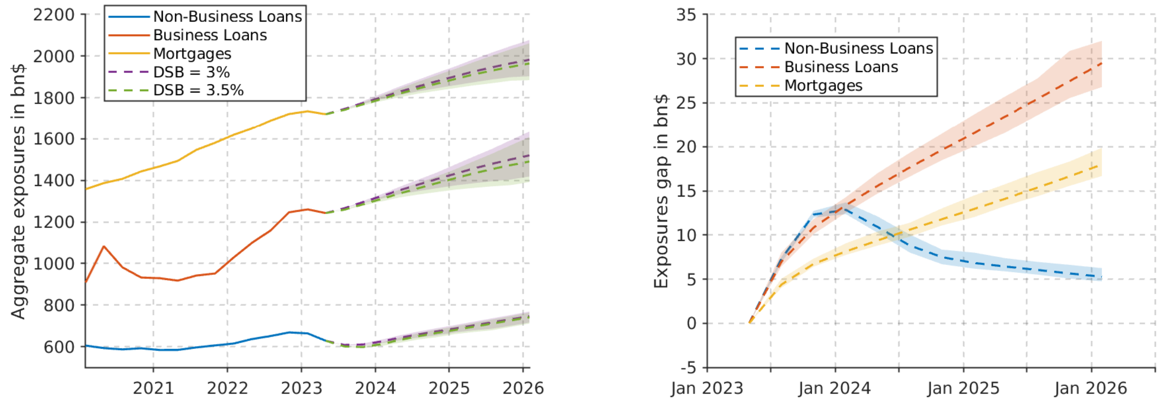

(a) Aggregate credit paths under DSB = 3% (purple) and 3.5% (green).

(b) Difference (3% minus 3.5%) across scenarios.

Figure 6: Aggregate credit provision for the Big Six under different Domestic Stability Buffers (DSB) across the distribution of simulated macro-financial scenarios. The left panel compares the full paths of aggregate lending, while the right panel shows the implied gap, capturing the estimated impact of the June 2023 buffer increase. Credit components include non-business loans (blue), business loans (red), and mortgages (yellow).

Figure 6 reports the results. Raising the DSB to 3.5% reduces aggregate loans by approximately \$54 billion on average, a contraction of 1.2% with a simulation range of −0.7% to −1.9%. The range reflects variation across macro-financial scenarios rather than parameter uncertainty: under benign conditions banks hold ample buffers and the tighter constraint binds loosely, while under adverse conditions the upward shift in the regulatory floor has greater bite, amplifying the contraction precisely when credit supply is already under pressure. The model predicts a reduction that is correlated across institutions by construction, because all six banks face the same shift in the floor simultaneously, limiting the scope for cross-sectional compensation of the kind that attenuates individual-bank shocks in the resilience analysis above. As a benchmark for the magnitude, the estimated contraction is comparable to the credit restriction associated with a 25 basis point increase in the policy rate according to standard macroeconomic transmission models.[11] The finding quantifies the trade-off between resilience and credit supply that the resilience analysis in Section 3.1 identifies qualitatively: raising the buffer moves the system away from the tipping point, but at a cost to aggregate lending that has real economic impact.

[11]We translate the effect with MP2 model from the Bank of Canada (Alpanda et al., 2018).

3.3 Conducting reverse-stress tests

In this application, we implement a reverse stress testing framework to identify the macro-financial environments associated with systemic stress in the banking sector. Rather than conditioning on a small number of pre-defined scenarios, here, we evaluate a broad set of simulated macro-financial paths and isolate those under which a system-level stress event materializes.[12]

For that, we simulate 10,000 macro-financial scenarios using the procedure described in Section 2.6, each projecting outcomes eight quarters ahead.[13] The BSM provides the structural link from those environments to banking outcomes. For each simulated path s_w, we compute a bank-level stress indicator $D_{b,w,t}$, equal to one when a bank's capital ratio falls within 50 basis points of the regulatory minimum ($K = CAR_{min} + 0.005$), and zero otherwise. We denote the average across institutions the system-level share $\bar{D}_{w,t}$, and define the systemic stress event as $\mathcal{S} = \{\bar{D}_{w,t} \geq \bar{d}^*\}$, where $\bar{d}^* = 2/B$ corresponds to at least two banks entering stress.

Before characterizing the macro-financial environments associated with systemic stress, it is useful to clarify how the results should be interpreted. The BSM establishes an empirical association: the identified macro-financial paths are those under which the systemic event materializes. The model does not explain the real-economy mechanisms that generate these environments. Its role is to surface the conditions under which banks become stressed, with economic interpretation provided through the narratives that follow. In contrast to conventional stress testing, which requires specifying a transmission mechanism ex ante, reverse stress testing identifies relevant environments first and allows the underlying mechanisms to be interpreted ex post. The value of the clustering described below lies not in confirming any specific transmission channel, but in identifying macro-financial environments associated with systemic stress without imposing prior assumptions on which combination of conditions should be of concern.

Of the 10,000 simulated scenarios, 280 satisfy this criterion at the two-year horizon, yielding an estimate of the unconditional probability of systemic stress of $\hat{P}(\mathcal{S}) = 0.028$, as defined in equation (10).[14] The top panel of Figure 7 presents the 280 stress scenarios identified by the reverse stress test, plotting the individual macroeconomic paths underlying each scenario. Red lines correspond to scenarios in which GDP falls below its historical trend, while blue lines track scenarios in which GDP remains at or above trend throughout the simulation horizon. Sorting stress scenarios by their GDP trajectory

[12] Appendix D provides a formal comparison between traditional and reverse stress testing approaches within a unified framework.

[13] The bootstrap procedure does not carry a causal interpretation. The correlation between stress outcomes and macro variables characterizes environments conducive to systemic stress.

[14] The reliability of this estimate depends directly on the credibility of the scenario generator. The preserved cross-sectional dependence of the bootstrap ensures that the 280 stress simulations represent internally consistent macro-financial environments rather than combinations of marginal shocks.

reveals that a substantial share of systemic stress episodes occur without an aggregate economic downturn.

We recover $p(s \mid \mathcal{S})$ by examining the macroeconomic paths associated with these 280 scenarios and grouping them into clusters with similar trajectories. The bottom panel summarizes these patterns through cluster averages. It reports, for each cluster, the average trajectory of key macro variables (solid lines) and the 90% confidence interval (dashed lines), with the unconditional distribution shown in black. One cluster (shown in red) is characterized by a pronounced *GDP slowdown*, with output tracking below the corresponding fifth percentile of the unconditional distribution. By contrast, the remaining clusters (blue and green) evolve at or above the unconditional mean, indicating that systemic stress can materialize even in environments where aggregate economic activity continues to expand. The key question is therefore which features of these environments give rise to stress despite the absence of a broad macroeconomic contraction.

Figure 8 addresses this by examining the macro-financial conditions that distinguish the blue and green clusters across three additional variables: Canadian house prices, the CAD–USD exchange rate, and oil prices. In the blue cluster, house prices evolve weakly, with average trajectories lying below those observed in GDP-slowdown scenarios and the lower tail displaying declines similar in magnitude to the fifth-percentile outcomes under GDP contractions. We label this the *house price correction* narrative, in which usually stretched household balance sheets gradually erode mortgage loan quality and compress bank capital buffers even as aggregate output continues to grow. The green cluster tells a different story, combining an appreciation of the Canadian dollar with rising oil prices. We label this the *Dutch disease*. Higher oil prices can lead to a stronger exchange rate, affecting the export channel. Thus, exchange rate appreciation disproportionately impairs firms with US-dollar-denominated revenues, weakening their debt-servicing capacity, while firms that benefit from a stronger Canadian dollar do not generate commensurate gains in bank profitability. This asymmetry of aggregate conditions being benign and sectoral credit quality deteriorating is consistent with a financial variant of Dutch-disease dynamics, in which a resource-driven appreciation yields balance sheet stress that only becomes visible in bank loan books (see Corden and Neary, 1982). Note that the BSM does not model this transmission directly; it simply identifies the macro environment in which stress materializes. Yet, such a scenario would be unlikely to emerge under a conventional stress testing framework based on predefined narratives. This underscores the analytical value of the outcome-based approach.

Figure 7 also conveys information about the relative likelihood of each narrative, which approximates the shape of $p(s \mid \mathcal{S})$ across clusters. The GDP slowdown cluster contains the largest share of stress paths, indicating that conventional macroeconomic deterioration remains the most probable short-term pathway to systemic stress. Table 3 reports how this distribution evolves across forecast horizons. The unconditional probability of

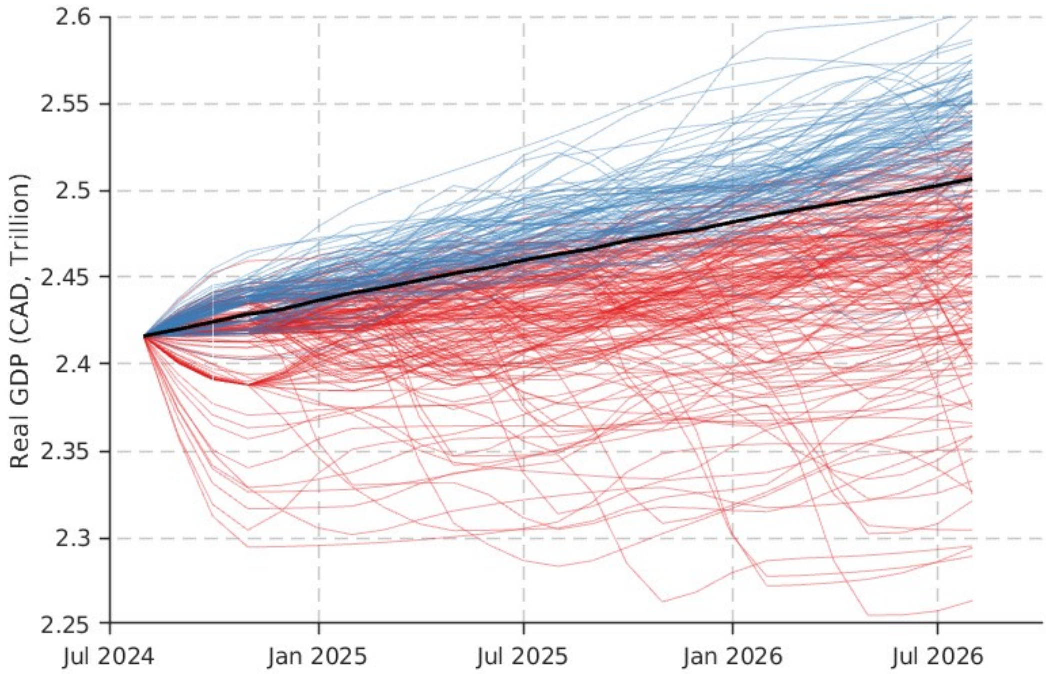

(a) Canadian real GDP simulations associated with stress financial conditions. Reverse stress testing analysis uncovers macro-conditions scenarios that are not always associated to a low-growth GDP paths.

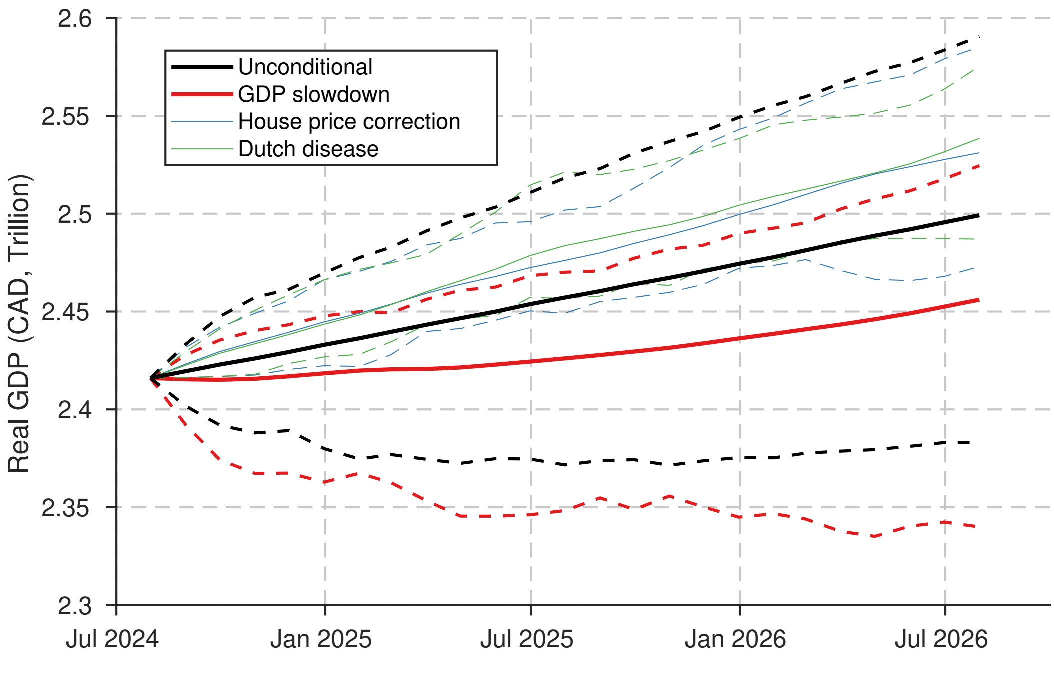

(b) Canadian real GDP. The red cluster tracks below the fifth percentile of the unconditional distribution, defining the GDP slowdown narrative. The blue and green clusters evolve at or above the unconditional mean, indicating that financial distress can arise even in the absence of an aggregate downturn.

Figure 7: Stress scenario clusters identified by the reverse stress test. Solid lines indicate the average trajectory within each cluster; dashed lines indicate the 90% confidence interval. The unconditional distribution across all 10,000 scenarios is shown in black. Three distinct stress narratives emerge: GDP slowdown (red), house price correction (blue), and CAD appreciation shock (green).

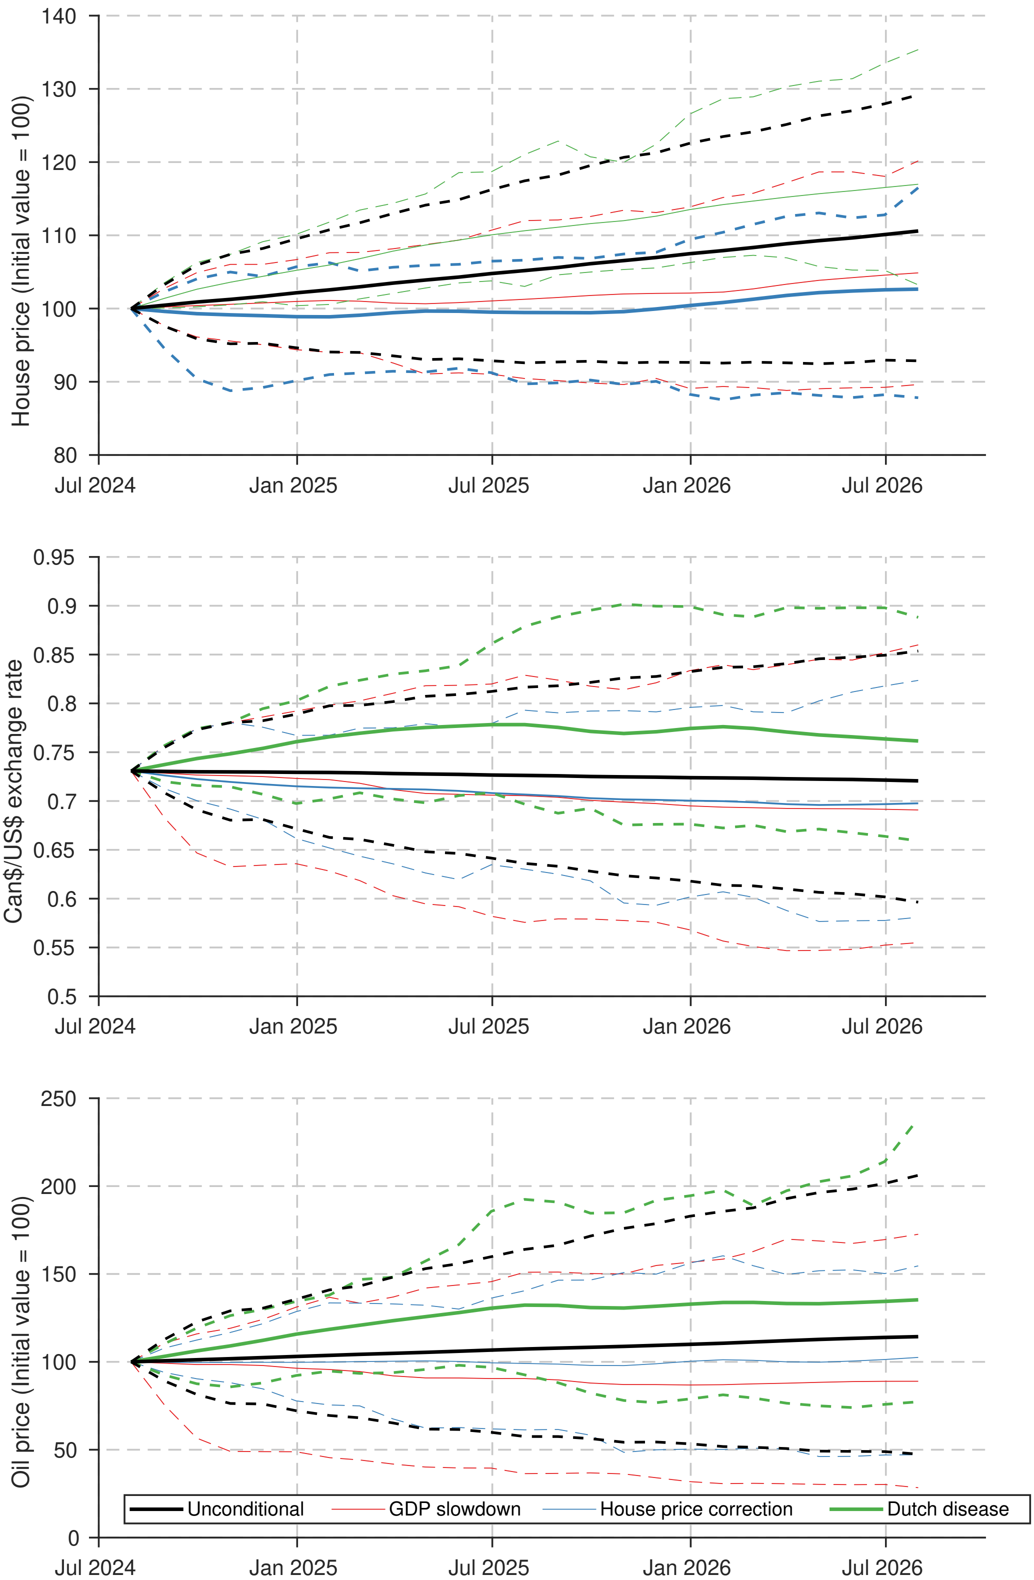

Figure 8: Stress scenario clusters identified by the reverse stress test, showing Canadian house prices (top), the CAD–USD exchange rate (middle), and oil prices (bottom). Solid lines indicate the average trajectory within each cluster; dashed lines indicate the 90% confidence interval. The unconditional distribution across all 10,000 scenarios is shown in black. Three distinct stress narratives emerge: GDP slowdown (red), house price correction (blue), and Dutch disease (green).

stress rises monotonically with the horizon, reflecting the path-dependent accumulation of shocks in the model. The conditional distribution shifts as well: the GDP slowdown narrative accounts for a declining share of stress outcomes at longer horizons, while the CAD appreciation shock becomes progressively more prevalent, rising from 4.2% of stress scenarios at horizon five to 12.1% at horizon eight. This transition implies that short-term systemic fragility is dominated by aggregate demand dynamics, while longer-term vulnerabilities increasingly originate from sectoral imbalances that are slow to materialize but persistent once established. For scenario designers, this finding suggests that standard stress test exercises, which typically operate at a one-year horizon and anchor on recessionary narratives, may systematically underweight the exchange rate and commodity price channels that become dominant sources of risk at medium-term horizons.

Quarters ahead		4	5	6	7	8
Probability of stress $\hat{p}(\mathcal{S})$		0.01%	0.24%	0.89%	1.61%	2.80%
Narrative share conditional on stress	GDP slowdown	100.00%	75.00%	76.40%	70.19%	64.64%
	House price correction	0.00%	20.86%	19.10%	19.25%	23.21%
	Dutch disease	0.00%	4.17%	4.49%	10.56%	12.14%

Table 3: Distribution of stress narratives across forecast horizons. The first row reports the estimated unconditional probability of systemic stress $\hat{p}(\mathcal{S})$, corresponding to the denominator of equation (11). The remaining rows report the share of each narrative within the set of stress scenarios, approximating the shape of $p(s \mid \mathcal{S})$ across clusters. Shares do not sum to 100% at horizons beyond four quarters, as a residual set of stress scenarios does not map cleanly to any of the three identified narratives.

4 Conclusions

We develop the Banks' Strategy Model, a partial equilibrium framework in which heterogeneous banks optimize their balance sheets subject to regulatory constraints, internalise the price impact of their trades, and reach a Nash equilibrium each period. By modelling both assets and liabilities as strategic variables, replacing hard regulatory constraints with smooth log-barriers that generate endogenous managerial buffers, and iterating the equilibrium dynamically, the BSM captures the portfolio behaviour that conventional top-down stress test models suppress by assumption. The result is a framework that is structural enough to support counterfactual analysis, transparent enough to trace mechanisms from assumption to output, and flexible enough to serve multiple distinct purposes without recalibration.

Three applications to Canada's big six banks illustrate what this buys in practice. First, the resilience analysis reveals a sharp nonlinearity in the relationship between bank capital and aggregate credit provision: above roughly 9.5% CAR, credit supply is relatively stable as losses at one institution are offset by expansion at others, but below this threshold a single institution's deleveraging can propagate through the price-impact channel and compress credit system-wide, a vulnerability that was directly relevant during the pandemic period when OSFI released the DSB and compressed observed buffers toward the floor. Second, the policy experiment quantifies the cost of rebuilding that buffer: raising the DSB from 3% to 3.5% reduces aggregate lending by approximately 1.2%, a contraction comparable in magnitude to a 25 basis point increase in the policy rate, and one that is amplified relative to a single-bank calculation because all six institutions face the same shift in the regulatory floor simultaneously. Third, we run a reverse stress test experiment that completes the picture by recovering the distribution of macro-financial environments consistent with systemic stress without conditioning on any prior narrative. Three distinct clusters emerge: a GDP slowdown, a house price correction, and a CAD appreciation shock in which rising oil prices and a stronger exchange rate support aggregate output while impairing the debt-servicing capacity of firms with US-dollar-denominated revenues. The last two of these would not appear in a conventionally designed stress scenario, illustrating the analytical value of a reverse stress test.

The model has three limitations that are in practice solvable but beyond the scope of this work. First, the use of quarterly data constrains the analysis to solvency dynamics and limits the model's ability to capture liquidity stress, which can deteriorate within a quarter and may precede rather than follow capital erosion. Second, the current balance sheet aggregation does not distinguish between hold-to-maturity and mark-to-market assets, which means unrealized losses on securities portfolios do not feed into regulatory ratios in the way they would for a bank with a large trading book. Third, the framework covers the banking sector only, abstracting from asset managers, insurers, and other non-bank intermediaries whose portfolio decisions may reinforce or offset the dynamics we identify.

Each limitation points toward a concrete extension. Introducing differentiated decision frequencies, with trading book positions updated at higher frequency than loan portfolios, would allow the model to capture the transformation of liquidity problems into solvency risk and better reflect the distinction between held-to-maturity and marked-to-market exposures. Allowing regulatory aversion parameters to vary over the cycle would endogenize the precautionary buffer-building behaviour observed empirically, generating time-varying distance to the constraint without imposing it by assumption. Incorporating non-bank financial institutions in a parsimonious way, for instance through a reduced-form representation of their demand for the same assets the banks trade, would extend the price-impact mechanism beyond the banking sector and allow the model to speak to the broader systemic risk implications of concentrated intermediation.

References

Abdelrahman, O., Chen, D. X., MacDonald, C., Mordel, A. and Ouellet Leblanc, G. (2025), Simulating the resilience of the canadian banking sector under stress: An update of the bank of canada's top-down solvency assessment tool, Technical report, Bank of Canada.

Adrian, T., Boyarchenko, N. and Giannone, D. (2019), 'Vulnerable growth', *American Economic Review* **109**(4), 1263–1289.

Aikman, D. (2026), 'Stress testing needs a reset', `https://niesr.ac.uk/blog/stress-testing-needs-reset`. NIESR blog.

Aikman, D., Angotti, R. and Budnik, K. B. (2024), Stress testing with multiple scenarios: a tale on tails and reverse stress test scenarios, ECB Working Paper 2941, European Central Bank.

Aikman, D., Beale, D., Brinley-Codd, A., Hüser, A.-C., Covi, G. and Lepore, C. (2023), *Macro-prudential stress test models: A survey*, International Monetary Fund.

Alpanda, S., Cateau, G. and Meh, C. (2018), 'A policy model to analyze macroprudential regulations and monetary policy', *Canadian Journal of Economics/Revue canadienne d'économique* **51**(3), 828–863.

Baes, M. and Schaanning, E. (2023), 'Reverse stress testing: Scenario design for macroprudential stress tests', *Mathematical Finance* **33**(2), 209–256.

Baudino, P., Goetschmann, R., Henry, J., Taniguchi, K. and Zhu, W. (2018), *Stress-testing banks: a comparative analysis*, Bank for International Settlements, Financial Stability Institute.

Bernanke, B. S. and Blinder, A. S. (1988), 'Credit, money, and aggregate demand', *American Economic Review* **78**(2), 435–439.

Borio, C., Drehmann, M. and Tsatsaronis, K. (2014), 'Stress-testing macro stress testing: Does it live up to expectations?', *Journal of Financial Stability* **12**, 3–15.

Breuer, T., Jandačka, M., Rheinberger, K. and Summer, M. (2009), 'How to find plausible, severe and useful stress test scenarios', *International Journal of Central Banking* **5**, 205–224.
URL: *https://www.ijcb.org/journal/ijcb09q3a7.pdf*

Breuer, T. and Summer, M. (2020), 'Systematic stress tests on public data', *Journal of Banking & Finance* **118**, 105886.

Bruneau, G., Ojea-Ferreiro, J., Plummer, A., Tremblay, M.-C. and Witts, A. (2025), 'The interdependencies of Canadian financial institutions: An application to climate transition shocks', *Latin American Journal of Central Banking* p. 100163.

Brunnermeier, M. K. and Pedersen, L. H. (2009), 'Market liquidity and funding liquidity', *The review of financial studies* **22**(6), 2201–2238.

Budnik, K., Mozzanica, M. B., Dimitrov, I., Groß, J., Hansen, I., Kleemann, M., Sanna, F., Sarychev, A., Siŋenko, N., Volk, M. et al. (2019), *Macroprudential stress test of the euro area banking system*, number 226, ECB Occasional Paper.

Caballero, R. J. and Simsek, A. (2013), 'Fire sales in a model of complexity', *The Journal of Finance* **68**(6), 2549–2587.

Corden, W. M. and Neary, J. P. (1982), 'Booming sector and de-industrialisation in a small open economy', *The economic journal* **92**(368), 825–848.

Duffie, D. and Singleton, K. J. (1993), 'Simulated moments estimation of markov models of asset prices', *Econometrica* **61**(4), 929–952.

Fique, J. (2017), The macrofinancial risk assessment framework (MFRAF), Version 2.0, Technical report, Bank of Canada.

Freixas, X. and Rochet, J.-C. (2008), *Microeconomics of Banking*, 2nd edn, MIT Press, Cambridge, MA.

Friedrich, C., Friedrich, H., Lawrence, N., Cortes Orihuela, J. and Tian, P. (2025), The International Exposure of the Canadian Banking System, Technical report, Bank of Canada.

Gaa, C., Liu, X., MacDonald, C. and Shen, X. (2019), Assessing the resilience of the canadian banking system, Technical report, Bank of Canada.

Gallant, A. R. and Tauchen, G. (1996), 'Simulated moments estimation', *Econometrica* **64**(4), 837–857.

Gârleanu, N. and Pedersen, L. H. (2013), 'Dynamic trading with predictable returns and transaction costs', *The Journal of Finance* **68**(6), 2309–2340.

Gertler, M. and Kiyotaki, N. (2010), Financial intermediation and credit policy in business cycle analysis, *in* 'Handbook of monetary economics', Vol. 3, Elsevier, pp. 547–599.

Goldstein, I., Jiang, H. and Ng, D. T. (2017), 'Investor flows and fragility in corporate bond funds', *Journal of Financial Economics* **126**(3), 592–613.

Hałaj, G. (2018), 'System-wide implications of funding risk', *Physica A: Statistical Mechanics and its Applications* **503**, 1151–1181.

Halaj, G. and Priazhkina, S. (2021), Stressed but not helpless: Strategic behaviour of banks under adverse market conditions, Technical report, Bank of Canada.

He, Z. and Krishnamurthy, A. (2013), 'Intermediary asset pricing', *American Economic Review* **103**(2), 732–770.

Hurley, P. R. and Hurley, R. E. (2013), 'Warning signals of the impending lehman brothers' bankruptcy filing.', *Journal of Applied Financial Research* **1**.

Jakab, Z. and Kumhof, M. (2015), Banks are not intermediaries of loanable funds — and why this matters, Working Paper 529, Bank of England.

Juselius, M., Ponte Marques, A. and Tarashev, N. A. (2025), 'Banks' regulatory risk tolerance'.

Kennedy, J. and Eberhart, R. (1995), Particle swarm optimization, *in* 'Proceedings of ICNN'95 - International Conference on Neural Networks', Vol. 4, IEEE, pp. 1942–1948.

Klein, M. A. (1971), 'A theory of the banking firm', *Journal of Money, Credit and Banking* **3**(2), 205–218.

Kyle, A. S. (1985), 'Continuous auctions and insider trading', *Econometrica* **53**(6), 1315–1335.

MacDonald, C. and Traclet, V. (2018), The framework for risk identification and assessment, Technical report, Bank of Canada.

Monti, M. (1972), Deposit, credit, and interest rate determination under alternative bank objectives, *in* G. P. Szegö and K. Shell, eds, 'Mathematical Methods in Investment and Finance', North-Holland, Amsterdam, pp. 431–454.

Orozco, L. and Rubio, S. (2025), 'Regulatory capital management to exceed thresholds', *Journal of Money, Credit and Banking* **57**(6), 1421–1464.

Rosen, J. B. (1965), 'Existence and uniqueness of equilibrium points for concave n-person games', *Econometrica* **33**(3), 520–534.

Schuermann, T. (2026), 'Mix me the perfect stress test', `https://www.systemicrisk.ac.uk/mix-me-perfect-stress-test`. Systemic Risk Centre blog.

Shleifer, A. and Vishny, R. (2011), 'Fire sales in finance and macroeconomics', *Journal of economic perspectives* **25**(1), 29–48.

Sydow, M., Schilte, A., Covi, G., Deipenbrock, M., Del Vecchio, L., Fiedor, P., Fukker, G., Gehrend, M., Gourdel, R., Grassi, A. et al. (2024), 'Shock amplification in an interconnected financial system of banks and investment funds', *Journal of Financial Stability* **71**, 101234.

Tobin, J. (1963), Commercial banks as creators of "money", *in* D. Carson, ed., 'Banking and Monetary Studies', Irwin, Homewood, IL, pp. 408–419.

Valderrama, L. (2026), 'Stress testing: Pushing the boundary', https://www.systemicrisk.ac.uk/stress-testing-pushing-boundary. Systemic Risk Centre blog.

A Endogenous Price Formation and the Linearised Return

Let $X_{N\times B} = [x_1, x_2, \dots, x_B]$ collect the portfolio strategies of all B banks, where each column $x_b \in \mathbb{R}^N$ contains the balance sheet changes of bank b and each row corresponds to a specific asset. The aggregate net demand across banks is

$$\omega := X\mathbf{1}_{B\times 1} \in \mathbb{R}^N,$$

which captures total buy or sell pressure per asset.

We distinguish between two distinct objects: the future price p_{t+1}, which is a random variable whose distribution is formed in the beliefs of each bank, and the current price $p_t(X)$, which is a deterministic function of the aggregate order flow. Specifically, following Kyle (1985), current prices respond linearly to net demand,

$$p_t(X) = p_t^0 + \Lambda\omega = p_t^0 + \Lambda X\mathbf{1}_{B\times 1},$$

where $p_t^0 \in \mathbb{R}^N$ is the baseline price vector that would prevail in the absence of any trading, and $\Lambda \in \mathbb{R}^{N\times N}$ captures the sensitivity of prices to order flow. For tractability we restrict Λ to be diagonal, so that $\Lambda = \lambda^{\text{kyle}} \cdot I_N$ for a scalar $\lambda^{\text{kyle}} > 0$, reflecting the assumption that each asset's price responds only to its own net demand with no cross-asset spillovers.

Each bank forms expectations over the future price p_{t+1}, which is assumed to be independent of current trades. The expected return for bank b is therefore

$$\mu_{b,t}(X) := \mathbb{E}_{b,t}[r_{t+1}] = \frac{\mathbb{E}_{b,t}[p_{t+1}]}{p_t(X)} - 1,$$

where the dependence on X enters entirely through the deterministic denominator $p_t(X)$. Working element-wise, the expected return on asset i is

$$\mu_{i,b,t}(\omega) = \frac{\mathbb{E}_{b,t}[p_{i,t+1}]}{p_{i,t}^0 + \lambda^{\text{kyle}} \cdot \omega_i}.$$

At the no-trade baseline $\omega = 0$, this reduces to the baseline expected return

$$\bar{\mu}_{i,b,t} := \frac{\mathbb{E}_{b,t}[p_{i,t+1}]}{p_{i,t}^0}.$$

Differentiating with respect to ω_i and evaluating at $\omega = 0$ gives

$$\left.\frac{\partial \mu_{i,b,t}}{\partial \omega_i}\right|_{\omega=0} = -\frac{\mathbb{E}_{b,t}[p_{i,t+1}] \cdot \lambda^{\text{kyle}}}{(p_{i,t}^0)^2} = -\frac{\bar{\mu}_{i,b,t} \cdot \lambda^{\text{kyle}}}{p_{i,t}^0}.$$

Since Λ is diagonal, cross-asset sensitivities vanish: $\partial \mu_{i,b,t}/\partial \omega_j = 0$ for $i \neq j$. Collecting terms across assets, the first-order expansion in vector form yields equation (2), with the return sensitivity matrix given by

$$\dot{\mu} := -\lambda^{\text{kyle}} \cdot \operatorname{diag}\left(\frac{\bar{\mu}_t}{p_t^0}\right),$$

where $\bar{\mu}_t/p_t^0$ denotes element-wise division. Although the baseline return $\bar{\mu}_{b,t}$ is bank-specific due to heterogeneous beliefs, the sensitivity matrix $\dot{\mu}$ depends only on the ratio of expected returns to current prices, which varies little across banks in practice. We therefore treat $\dot{\mu}$ as common across institutions and constant over time, preserving the core price-feedback mechanism while ensuring tractability.

B Equilibrium

B.1 Equilibrium existence

The model defines a simultaneous-move game among B banks. A natural question is whether this game admits a Nash equilibrium. We now argue that an equilibrium exists under empirically relevant conditions.

Intuition. Each bank's utility combines a return component, regulatory barrier terms, and quadratic trading costs. As a bank expands any position, it pushes up the purchase price (reducing expected returns) and incurs increasing trading costs. These forces generate strong concavity in the bank's own strategy. The regulatory log-barriers add curvature that depends on proximity to the constraint: steeply concave near the minimum, flattening further away. Equilibrium existence requires that the overall utility remains concave—which, as we show below, holds comfortably for realistic parameter values.

Hessian decomposition. The Hessian of $u_{b,t}$ with respect to x_b, holding other banks' strategies fixed, decomposes as

$$\frac{\partial^2 u_{b,t}}{\partial x_b \, \partial x_b'} = \underbrace{2(\dot{\mu} - \gamma)}_{\text{price impact + costs}} + \underbrace{\lambda_b^{\text{CAR}} \cdot \mathcal{H}_b^{\text{CAR}}(x_b)}_{\text{capital barrier}} + \underbrace{\lambda_b^{\text{LCR}} \cdot \mathcal{H}_b^{\text{LCR}}(x_b)}_{\text{liquidity barrier}} + \underbrace{\lambda_b^{\text{LR}} \cdot \mathcal{H}_b^{\text{LR}}(x_b)}_{\text{leverage barrier}}. \tag{6}$$

The first term, $2(\dot{\mu} - \gamma)$, is a diagonal matrix with strictly negative entries, since $\dot{\mu}$ has negative diagonal (price impact) and γ has positive diagonal (trading costs). This term alone guarantees strict concavity.

Capital barrier. Define $d(x_b) := \theta_a'(bs_b + x_b)$ as total risk-weighted assets. Since all entries of θ_a are positive (risk-weight densities lie between zero and one), d is increasing in asset holdings. The Hessian of $\log(\mathrm{CAR}(x_b) - \mathrm{CAR}_{\min})$ takes a rank-one form:

$$\mathcal{H}_b^{\mathrm{CAR}}(x_b) \;=\; \frac{e_b}{d^3} \;\cdot\; \frac{\mathrm{CAR} - 2\,\mathrm{CAR}_{\min}}{(\mathrm{CAR} - \mathrm{CAR}_{\min})^2} \;\cdot\; \theta_a\,\theta_a', \tag{7}$$

where $\theta_a\,\theta_a'$ denotes the outer product of the risk-weight vector (a positive semidefinite matrix). The scalar prefactor determines the sign. When $\mathrm{CAR} < 2\,\mathrm{CAR}_{\min}$, the Hessian is negative semidefinite and the barrier reinforces concavity. This condition is easily satisfied: with $\mathrm{CAR}_{\min} \approx 11.5\%$, a bank would need a CAR ratio exceeding 23% to violate it; well above the observed range of 12–14%. Hence, $\mathcal{H}_b^{\mathrm{CAR}}$ is negative semidefinite throughout the empirically relevant region, and the capital barrier unambiguously supports concavity. Violations of this region are multiplied by λ_b^{CAR}, and if λ_b^{CAR} is small enough do not dominate the concavity of the price impact and cost parameter

Liquidity barrier. The LCR barrier requires more care. Writing $\log(\mathrm{LCR}-\mathrm{LCR}_{\min}) = \log(\phi_a'(bs + x) - \mathrm{LCR}_{\min} \cdot \phi_f'(bs + x)) - \log(\phi_f'(bs + x))$, the Hessian decomposes as

$$\mathcal{H}_b^{\mathrm{LCR}}(x_b) \;=\; \underbrace{-\frac{(\phi_a - \mathrm{LCR}_{\min}\cdot \phi_f)(\phi_a - \mathrm{LCR}_{\min}\cdot \phi_f)'}{(\phi_a'(bs{+}x) - \mathrm{LCR}_{\min}\cdot \phi_f'(bs{+}x))^2}}_{\text{negative semidefinite}} \;+\; \underbrace{\frac{\phi_f\,\phi_f'}{(\phi_f'(bs + x))^2}}_{\text{positive semidefinite}}\,. \tag{8}$$

The first (concave) term captures the curvature from approaching the LCR minimum: it scales with $1/(\mathrm{LCR} - \mathrm{LCR}_{\min})^2$ and dominates when the buffer is small. The second (convex) term arises from the denominator of the LCR and scales with $1/f^2$, where $f = \phi_f'(bs + x)$ is the level of net cash outflows. When the LCR buffer is small, the first term dominates the second: a bank near its liquidity minimum faces a steeply concave penalty. When the buffer is large, the second term can dominate—but in that region, the overall contribution $\lambda_b^{\mathrm{LCR}} \cdot \mathcal{H}_b^{\mathrm{LCR}}$ is small in magnitude, and the quadratic terms $2(\dot{\mu} - \gamma)$ absorbs any residual convexity.

Leverage barrier. Define $\ell(x_b) := \mathbf{1}'(bs_b{+}x_b)$ as the total exposure measure (summing all asset holdings with unit weight). Since each entry of the unit weight vector $\mathbf{1}$ is positive, $\ell(x_b)$ is increasing in asset holdings. The Hessian of $\log(\mathrm{LR}(x_b){-}\mathrm{LR}_{\min})$ likewise takes a rank-one form:

$$\mathcal{H}^{\mathrm{LR}}(x_b) = \frac{e_b}{\ell^3} \cdot \frac{\mathrm{LR} - 2, \mathrm{LRmin}}{(\mathrm{LR} - \mathrm{LR}_{\min})^2} \cdot \mathbf{1}\mathbf{1}',$$

where $\mathbf{1}\,\mathbf{1}'$ denotes the outer product of the unit exposure vector (a positive semidefinite matrix). The scalar prefactor determines the sign. When $\mathrm{LR} < 2\,\mathrm{LR}_{\min}$, this Hessian is negative semidefinite, so the leverage ratio barrier reinforces concavity. This condition is easily satisfied: with $\mathrm{LR}_{\min} \approx 3\%$, a bank would need a leverage ratio above 6% to violate it – at or above the upper end of observed Tier 1 leverage ratios. Hence, $\mathcal{H}_b^{\mathrm{LR}}$ is negative semidefinite throughout the empirically relevant region, and the leverage barrier unambiguously supports concavity. Violations of this region are scaled by the penalty parameter λ_b^{LR}; if λ_b^{LR} is sufficiently small, such violations do not overturn the overall concavity driven by price-impact and cost parameters.

Practical assessment. To summarize: the quadratic terms provide uniform strict concavity; the CAR and LR barrier reinforce it for all empirically observed capital ratios; and the LCR barrier is concave near the constraint and negligible far from it. In our estimated parametrization, the overall Hessian (6) is negative definite for all banks across all simulated periods. With each bank's utility strictly concave in its own strategy and the strategy sets convex and effectively compact,[15] existence of a pure-strategy Nash equilibrium follows from Rosen (1965).[16]

B.2 Gauss-Seidel

At the beginning of each period, the order in which banks act is randomly determined to ensure that no single bank consistently moves first, thereby preventing systematic bias in the results due to the sequence of actions. Each bank, when making its decision, observes the strategies of those that have already acted and selects its optimal strategy to maximize utility. Once all banks have made their decisions, the order is reshuffled, and each bank updates its strategy based on the revised market conditions. Through repeated iterations of this process, the system converges toward a Nash equilibrium, independent of the initial order of play.

For example, if bank A makes its investment decision before bank B, in the subsequent iteration, bank A may revise its decision to account for the market impact generated by bank B's actions. Figure 9a illustrates this process by showing how the utility of bank A—represented as a function of adjustments to assets 1 and 2 in its portfolio—shifts in response to bank B's investment decision. The green elliptical contour represents bank A's utility landscape prior to bank B's move, while the blue contour depicts the updated utility after bank B has acted. The black dot marks the optimal investment strategy for bank A before bank B's intervention. If, for instance, bank B chooses to sell asset 1—a move that aligns with bank A's initial strategy as indicated by the negative position along

[15]The balance sheet identity is a linear constraint, and the quadratic trading cost drives utility to $-\infty$ for large $\|x_b\|$, ensuring that the effective strategy set is bounded.

[16]Uniqueness requires the stronger condition of *diagonal strict concavity* (Rosen, 1965). While we do not formally verify this condition, the Gauss–Seidel procedure described in Appendix B.2 converges to the same equilibrium regardless of the initial ordering of banks across all simulated scenarios, suggesting generic uniqueness for our parametrization.

the x-axis—it could further depress the price of asset 1 beyond bank A's expectations. This additional price impact may cause bank A to reconsider and reduce its own sale of asset 1, potentially shifting its optimal decision closer to zero on the x-axis. Notably, the utility function of bank A remains concave with respect to changes in assets 1 and 2, both before and after bank B's move, ensuring the existence of a unique global maximum.

Figure 9b demonstrates how this iterative, randomized decision-making process continues until the maximum change in asset holdings across all banks falls below a predefined threshold, indicating convergence to a stable outcome. While Figure 9a highlights the conditional utility-maximizing decision of bank A given the current strategies of other players, Figure 9b captures the convergence of these iterative decisions toward a Nash equilibrium—where no individual bank can improve its utility by unilaterally altering its strategy. This convergence process bears resemblance to the derivation of long-run probabilities in Markov chains, where the system converges to a stationary distribution regardless of the initial state. Similarly, although a bank's early investment choices are influenced by the prior actions of others, repeated strategic adjustments across rounds lead the system toward an equilibrium that is independent of the initial sequencing.

Figure 9: optimization process

(a) Utility of a bank as function of changes in two assets of the balance sheet before (green) and after (blue) knowing its peers investment strategies.

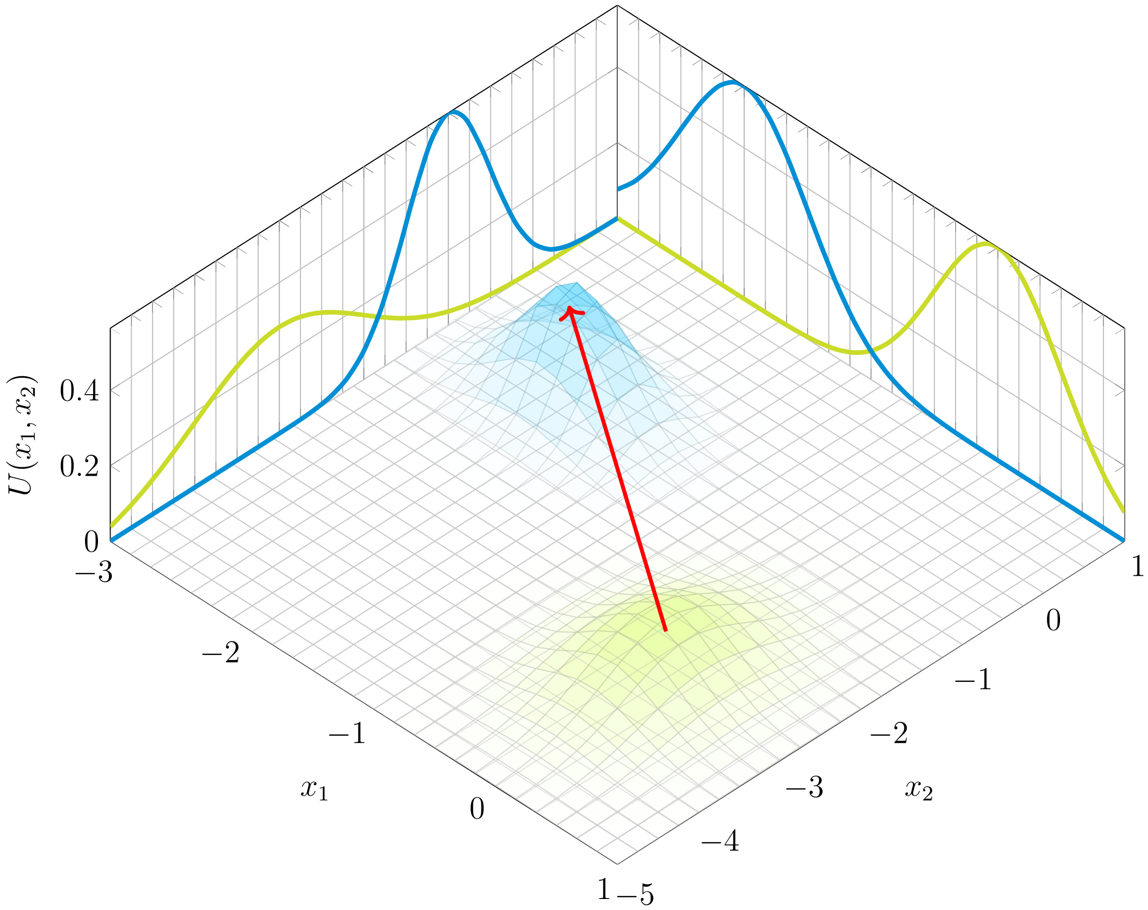

(b) Illustration of the convergence process

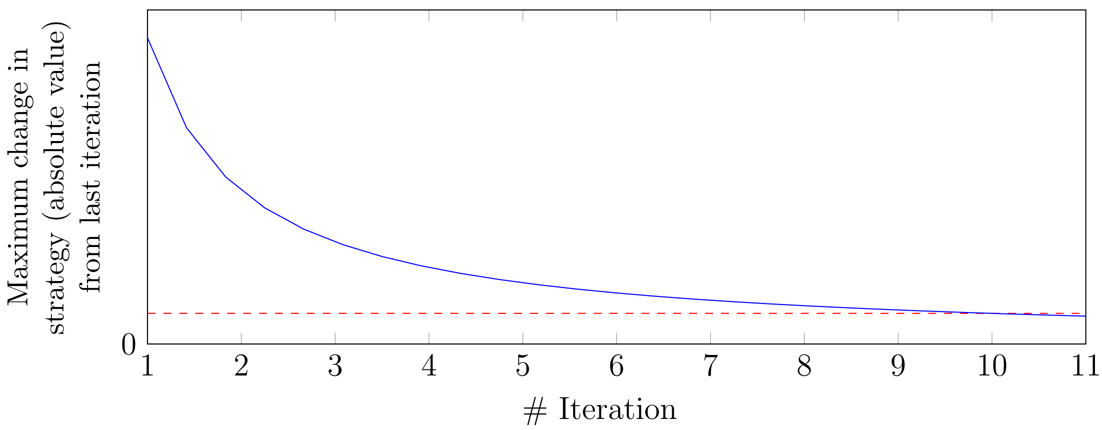

Top figure shows an example of the utility function of a bank before (after) knowing what its peers investment strategies in the green (blue) area. New information of its peers would update the bank investment decision to maximize its utility given the effects of their peers on market prices. This change is reflected by the red arrow. The process continues until some convergence criterion is reached. Bottom figure shows how the iteration of the decisions between the set of banks would lead to smaller changes in each iteration. The process stops when the change is small enough, indicated by the red dashed line in the bottom figure.

C Balance sheet identity: from retained earnings to asset side adjustments

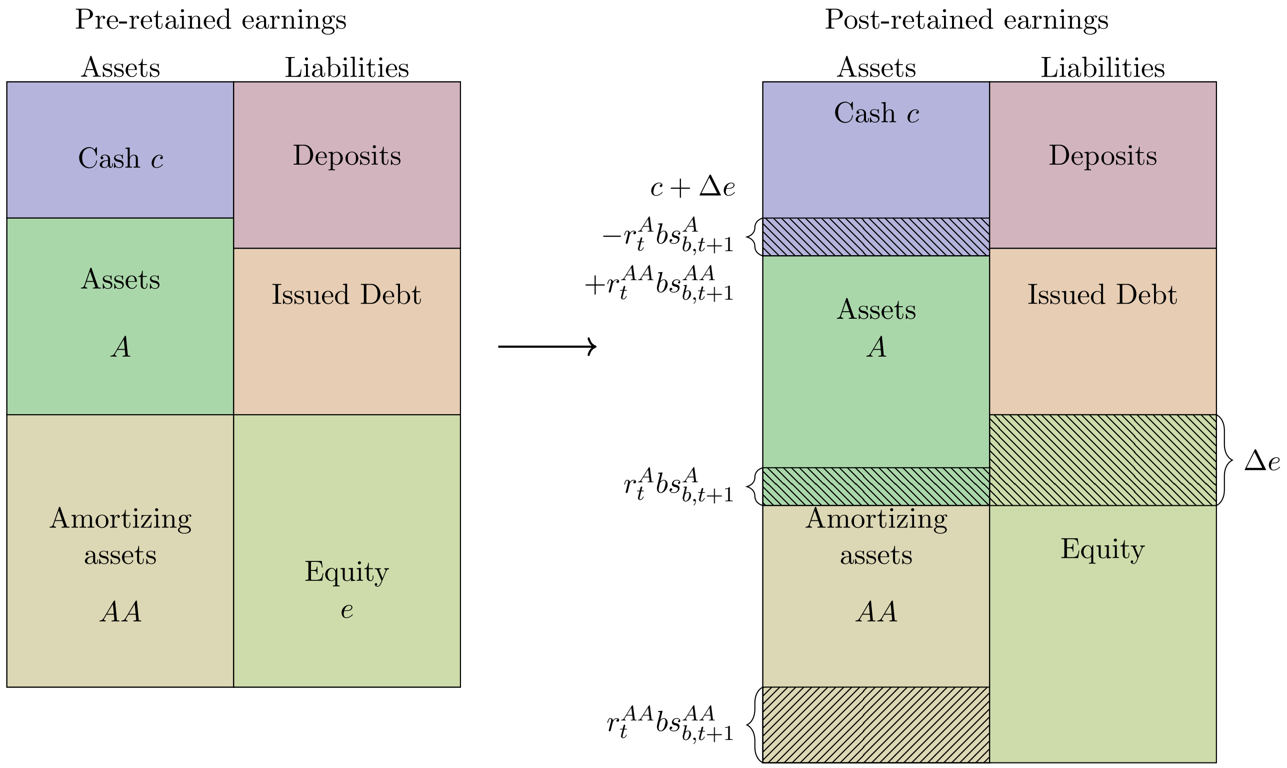

Figure 10: Stylized book and realized balance sheet for a bank after when moving from time t to t+1. North-west (north-east) hatch indicates section of the balance sheet increasing (decreasing)

D Comparison between traditional and reverse stress testing approaches

This section provides a methodological comparison between traditional (forward-looking) stress testing and reverse stress testing within a unified analytical framework. While both approaches rely on the same structural mapping from macro-financial scenarios to bank-level outcomes, they differ fundamentally in how scenarios and stress events are defined and evaluated. By expressing both methodologies using common building blocks, we clarify their conceptual differences and highlight the additional information delivered by reverse stress testing.

To formalize this comparison, let w index a draw from a macro-financial scenario generator, which produces a full macro-financial path s_w. Let $\tau_{b,w,t}$ denote a key metric for bank b at horizon t under a macrofinancial scenario path s_w. Outcomes are generated by our structural banking model $BSM(\dots)$, which maps a macro-financial scenario path m_w into outcomes according to $\tau_{b,w,t} = BSM(s_w)$, capturing behavioural patterns of banks in terms of strategic balance sheet adjustments. The function $BSM(\dots)$ captures not only the direct effect of macroeconomic conditions, but also the interactions among heterogeneous institutions and their endogenous responses to a common macro-financial environment. Given a policy-relevant threshold K, we summarize stress outcomes using the indicator $D_{b,w,t} = \mathbf{1}\{\tau_{b,w,t} \leq K\}$, which serves as the primitive building block for both traditional and reverse stress-testing methodologies.

A conventional forward stress test fixes a macroeconomic scenario $\bar{s}$ and evaluates these indicators institution by institution. By focusing on a few severe but stylized trajectories, such tests may understate nonlinear dynamics, miss interactions across risk factors, and leave large portions of the risk space unexplored. The test is binary: each bank either passes or fails under the prescribed scenario, i.e.

$$D_{b,c,t} = \mathbf{1}\{\tau_{b,\bar{s},t} \leq K\} \quad \text{for b=1, \dots, B.} \tag{9}$$

Importantly, no joint or systemic condition is imposed ex ante; any system-level assessment is derived only by aggregating individual outcomes after the scenario has been evaluated.

Reverse stress testing addresses these limitations by adopting an outcome-based perspective: rather than specifying a scenario and asking how bad outcomes could be, it begins with bad outcomes and asks which scenarios produce them. Reverse stress testing instead begins by defining a system-level stress criterion. Let $\bar{D}_{w,t} = \frac{1}{B}\sum_{b=1}^{B} D_{b,w,t}$ denote the fraction of banks close to binding regulatory constraints, and define the systemic stress event $\mathcal{S} = \{\bar{D}_{w,t} \geq \bar{d}^*\}$, where $\bar{d}^*$ is a threshold chosen to defined the systemic event. The probability of this stressful event implies looking into the full range of potential W macrofinancial scenarios, i.e.

$$P(\mathcal{S}) = \mathbb{E}(\mathbf{1}\{\bar{D}_{w,t} \geq \bar{d}^*\}), \tag{10}$$

rather than a single scenario path $\bar{s}$ as in traditional stress tests.

Apart from getting the probability of a systemic risk event, reverse stress testing allows us to identify which macro-financial environments are consistent with this systemic stress. Using Bayes' rule, the distribution of macro-financial scenarios conditional on systemic stress is given by

$$p(s \mid \mathcal{S}) = \frac{P(\mathcal{S} \mid s)\, p(s)}{P(\mathcal{S})}, \tag{11}$$

where s is a random element representing a macro-financial scenario (a path of macroeconomic and financial variables) with probability density $p(s)$, $P(\mathcal{S} \mid s)$ assigns weight to macro-financial scenarios according to whether they are associated with systemic stress, and $P(\mathcal{S})$ ensures proper normalization. Equation (11) makes explicit what reverse stress testing delivers what equation (9) cannot: a characterization of the set of macro-financial environments related to systemic stress, weighted by their empirical plausibility, without conditioning on any prior narrative. This feature provides a strong power of reverse stress testing as an exploratory exercise mapping vulnerabilities of the financial system.

We implement this approach in section 3.3 by simulating 10,000 macro-financial scenarios using the procedure described in Section 2.6, each projecting outcomes eight quarters ahead.[17] The BSM provides the structural link from those environments to banking outcomes For each simulated path s_w, we compute bank-level stress indicators $D_{b,w,t}$ defining the threshold as being less than 50 bps from the CAR minimum, i.e. $K = CAR_{min} + 0.005$, and the corresponding system-level share $\bar{D}_{w,t}$. The systemic stress event is triggered when at least two banks comply with the stress condition, i.e. defined as $\bar{d}^* = 2/B$. Of the 10,000 simulated scenarios, 280 satisfy this criterion at the two-year horizon, yielding an estimate of the unconditional probability of systemic stress of $\hat{P}(\mathcal{S}) = 0.028$, as defined in equation (10).[18]

[17] The bootstrap procedure does not carry a causal interpretation. The correlation between stress outcomes and macro variables characterizes environments conducive to systemic stress.

[18] The reliability of this estimate depends directly on the credibility of the scenario generator. The preserved cross-sectional dependence of the bootstrap ensures that the 280 stress simulations represent internally consistent macro-financial environments rather than combinations of marginal shocks.