TALLINN UNIVERSITY OF TECHNOLOGY
School of Business and Governance

Igor Mohhov

ENDOGENOUS CREDIT CYCLES AND MACROECONOMIC STABILITY: AN AGENT-BASED APPROACH WITH ADAPTIVE BANK LEARNING

Master's thesis

Programme Economic Analysis

Supervisor: Nicolas Reigl, PhD

Tallinn 2026

I hereby declare that I have compiled the thesis independently
and all works, important standpoints and data by other authors
have been properly referenced and the same paper
has not been previously presented for grading.

The document length is 12968 words from the introduction to the end of the conclusion.

Igor Mohhov 12.05.2026

(date)

TABLE OF CONTENTS

ABSTRACT

Traditional macroeconomic frameworks often struggle to capture the endogenous systemic risks that lead to financial crises. While modern Heterogeneous Agent New Keynesian (HANK) models incorporate agent diversity, they frequently rely on exogenous shocks to drive fluctuations. This thesis develops a multi-sector Agent-Based Model (ABM) to investigate how macroeconomic instability emerges from bottom-up interactions between households, production firms, and a commercial bank.

The primary innovation of this research is the behavioral specification of the banking sector as an adaptive learner. Utilizing a rolling-window logistic regression, the bank endogenously updates its perception of firm default risk based on historical data. Crucially, the author introduces an information-filtering rule, specifically a 10-unit debt threshold, within the bank's training set. This mechanism addresses the risk of informational distortion by excluding zero-debt firms that would otherwise lead the bank to systematically underestimate bankruptcy probabilities.

Following a generative social science methodology, the model is validated through Monte Carlo simulations and compared against empirical stylized facts. An ablation study demonstrates that the bank's learning process acts as a significant volatility amplifier. The results show that institutional risk aversion, when calibrated through targeted data filtering, can transform localized firm-level distress into aggregate contractions. This research suggests that credit cycle stability is highly sensitive to the bank's internal heuristics and data-processing rules. By providing a computational laboratory to explore these dynamics, the study offers new insights for the design of resilient macroprudential policies that account for the endogenous nature of financial risk.

Keywords: Agent-Based Modeling, Financial Accelerator, Adaptive Learning, Credit Cycles, Macroeconomic Stability.

INTRODUCTION

The global financial crisis of 2008 and the subsequent economic shocks caused by the COVID-19 pandemic have exposed significant limitations in traditional macroeconomic modeling. While the dominant paradigm of Dynamic Stochastic General Equilibrium (DSGE) models has largely transitioned into the Heterogeneous Agent New Keynesian (HANK) literature, which incorporates financial frictions and agent diversity, these frameworks still largely rely on exogenous shocks to drive business cycle dynamics.

This thesis addresses these limitations by utilizing an Agent-Based Model (ABM). This approach provides a more principled foundation for Stock-Flow Consistency (SFC) and explores the endogenous emergence of credit cycles and macroeconomic instability, where shocks arise from the bottom-up interactions of agents rather than external perturbations.

The core research problem lies in the "micro-macro gap": the inability of aggregate models to explain how individual financial distress scales into systemic crises through decentralized feedback loops. While financial frictions are recognized as a primary driver of economic volatility, the mechanism through which a banking sector "learns" and reacts to risk in real-time is often oversimplified.

The novelty of this work stems from the implementation of an adaptive learning bank. Unlike models with static risk weights, the bank in this simulation employs a rolling-window logistic regression to update its perception of default risk. Crucially, this research introduces an information filter, specifically a 10-unit debt threshold, to the bank's training data. This modification ensures that the bank's institutional memory is not skewed by zero-debt firms, thereby preventing the systematic underestimation of bankruptcy probabilities.

The aim of this thesis is to develop a multi-sector agent-based model to investigate how decentralized credit rationing and adaptive risk assessment contribute to the emergence of macroeconomic stylized facts, such as persistent unemployment and business cycle fluctuations. To achieve this goal, the research addresses three primary questions:

1) How does the bank's adaptive learning process and information-filtering rule influence the volatility and longevity of credit cycles?
2) To what extent do decentralized search frictions in the labor market interact with firm-level credit constraints to drive aggregate instability?
3) Can a bottom-up model built on boundedly rational heuristics replicate the statistical patterns and relative volatilities observed in empirical data?

This research adopts a generative social science methodology, where the model "grows" an economy from the bottom up. The research object consists of interacting agents: households, consumption goods firms (C-firms), capital goods firms (K-firms), and a commercial bank. Behavioral rules are calibrated using foundational literature, including Carroll's (1997) buffer-stock saving and Dosi et al.'s (2010) Schumpeter meeting Keynes (K+S) framework. The model is validated using a Monte Carlo approach, utilizing an average of 89 simulation runs to generate synthetic data for comparison against empirical benchmarks from the FRED database.

A central contribution of this work is the comparative analysis presented in the results section, which is structured as an ablation study. In the context of computational modeling, an ablation study refers to the systematic removal of a specific heuristic or component to isolate its individual impact on the overall system's behavior. In this case, the mechanism being tested is the bank's informational filtering process.

By comparing the "filtered history" model against a baseline configuration where the bank trains its risk parameters on the entire population of firms, the author demonstrates that systemic volatility is significantly amplified by the bank's learning process. Specifically, when the bank incorporates broader and unfiltered datasets into its risk estimation, the accelerated updating of its risk parameters produces sharper fluctuations in credit availability.

These fluctuations subsequently destabilize the broader macroeconomy. Ultimately, the findings suggest that institutional risk aversion, when properly calibrated through targeted data filtering, acts as a primary determinant of macroeconomic stability.

The thesis is organized into three main parts:

1) The first part provides a comprehensive literature review establishing the theoretical grounding for agent behaviors, synthesizing theories of monopolistic competition, financial frictions, and the distinction between HANK and ABM frameworks.
2) The second part details the methodology and technical specification of the model, describing the mathematical implementation of the agents and the bank's adaptive credit-scoring mechanism.
3) The third part presents the results of the simulation and validation, testing for the emergence of stylized facts and exploring parameter sensitivity through the aforementioned ablation analysis.

The author wishes to express their sincere gratitude to their supervisor for the invaluable guidance, patience, and mentorship provided throughout the development of this computational framework. Further thanks are extended to those who provided support for the technical execution and computational infrastructure required for the simulation runs. Any remaining errors or omissions are strictly the author's own.

In the preparation of this thesis, Large Language Model (LLM) tools were utilized exclusively for the purposes of language refinement, grammatical correction, and enhancing the stylistic clarity of the technical prose. The underlying conceptual framework, model architecture, simulation design, and all analytical conclusions remain entirely the work of the author.

1. THEORETICAL FOUNDATIONS OF MACRO-FINANCIAL ABMs

The thesis is grounded in Agent-Based Computational Economics (ACE), which conceptualizes the economy as an evolving system of autonomous interacting agents (Tesfatsion & Judd, 2006). This paradigm departs from the Representative Agent (RA) framework by emphasizing heterogeneity and non-equilibrium dynamics (Kirman, 1992). As noted in the foundational toolkit by Delli Gatti *et al*. (2018), ABMs allow researchers to move beyond the "straitjacket" of equilibrium, focusing instead on the processes of coordination and the emergence of aggregate phenomena from micro-level rules.

1.1. Agent-Based Computational Economics (ACE) vs. DSGE Models

Unlike traditional Dynamic Stochastic General Equilibrium (DSGE) models, this approach explains "stylized facts" from the bottom up (Epstein, 2012). These include the fat-tailed distributions of firm sizes (Axtell, 2001) and the persistence of volatility, manifested as long-range autocorrelation in absolute returns (Lux & Marchesi, 1999).

Leigh Tesfatsion (2002) formalized the field by defining ACE as the "computational study of economic processes modeled as evolving systems of autonomous interacting agents."

To distinguish Agent-Based Modeling from the traditional neoclassical framework, four interconnected concepts serve as the structural pillars of the methodology:

1) **Heterogeneity**: Unlike the "representative agent" assumption, ABM presumes that agents differ fundamentally in their attributes, preferences, and information sets. This allows for the study of distributional effects and acknowledges that the aggregate economy is driven by the diversity of its participants, ranging from small-scale households to multinational conglomerates, each reacting differently to the same economic signal.

2) **Bounded Rationality**: Departing from the assumption of "homo economicus" possessing perfect foresight, ABM agents are characterized by bounded rationality. As agents operate within complex environments with limited computational capacity, they rely on adaptive heuristics and inductive learning. This shifts the focus from finding an optimal "global" solution to observing how agents satisfice and adapt to local information.

3) **Emergence**: This is the hallmark of the ABM approach, where macroscopic patterns (such as market crashes, wealth inequality, or innovation waves) arise endogenously from the bottom up. Emergence implies that the aggregate system possesses properties that are not present in any single agent, making the system "greater than the sum of its parts" and often leading to non-linear outcomes that traditional linear models fail to capture.

4) **Calibration and Validation**: While theoretical ABMs prioritize "generative sufficiency," modern applications emphasize Calibration or the process of tuning model parameters so that simulated outputs match real-world empirical moments. This is the precursor to Validation, which tests the model's predictive power against out-of-sample data, ensuring the artificial economy is not just a theoretical construct but a statistically grounded representation of reality.

1.2. Literature Review of the Framework

The evolution of Agent-Based Modeling (ABM) in economics represents a fundamental shift from the traditional "representative agent" paradigm toward a framework that embraces heterogeneity, local interactions, and emergent phenomena. By simulating the individual actions of autonomous agents, researchers can observe how macroscopic patterns arise from the bottom up, often bypassing the restrictive assumptions of general equilibrium.

1.2.1. Foundations of Emergent Behavior and Bounded Rationality

The intellectual genesis of economic Agent-Based Modeling (ABM) is rooted in the critique of Walrasian equilibrium and the recognition of the economy as a complex adaptive system. This foundation was established through seminal works that demonstrated how simple micro-level rules could precipitate sophisticated, and often unintended, macroscopic phenomena.

The trajectory of the field is frequently traced back to Schelling (1971), whose "Dynamic Models of Segregation" provided a definitive proof of concept for emergent behavior. Schelling utilized a localized interaction framework to demonstrate that aggregate social patterns, specifically residential segregation, do not necessarily reflect the explicit preferences or "micromotives" of individual actors. By showing that even a marginal preference for neighborhood similarity could trigger a tipping point into total systemic segregation, Schelling established a core ABM tenet that macroscopic outcomes are frequently decoupled from individual intentions. This highlighted the necessity of modeling the interaction structure itself, rather than merely the individual components in isolation.

Further advancing the critique of the "representative agent" framework, Kirman (1993) introduced a model of stochastic interaction in "Ants, Rationality, and Recruitment." Drawing parallels between the foraging behavior of ant colonies and the decision-making processes of financial market participants, Kirman demonstrated how "herding" behavior emerges from simple recruitment mechanisms. This work was pivotal in showing that aggregate volatility and structural shifts can arise endogenously through agent imitation and communication, rather than being the result of exogenous shocks. Kirman’s "ants" model provided a formal mechanism to explain why markets often deviate from fundamental values, positioning local interaction as a primary driver of aggregate instability.

Parallel to the study of emergence, Arthur (1994) fundamentally challenged the classical assumption of deductive rationality through the "El Farol Bar Problem." Arthur posited that in complex environments, agents face an "expectational" problem where the optimal strategy depends on the actions of others, creating a feedback loop that renders traditional deductive logic impossible. Instead, Arthur proposed that agents employ inductive reasoning, utilizing a shifting "toolbox" of adaptive heuristics and mental models. This conceptualization of the agent as an adaptive learner, rather than a static optimizer, transformed the modeling of expectations. Arthur’s work underscored that in non-equilibrium settings, the economy behaves as an "organic" and evolving system, where the constant adaptation of expectations prevents the system from ever reaching a final, settled state.

Together, these foundational contributions shifted the focus of economic inquiry toward the study of bounded rationality and non-linear emergence. They established that by relaxing the constraints of perfect information and global coordination, researchers could capture the self-organizing, and occasionally self-destabilizing, properties of real-world economic systems.

1.2.2. Agent-Based Computational Economics

The formalization of these methodologies into a distinct research paradigm was significantly advanced by Tesfatsion (2002), who defined Agent-Based Computational Economics (ACE) as the computational study of economic processes modeled as evolving systems of autonomous, interacting agents. Departing from the static nature of traditional equilibrium analysis, ACE is characterized by three core methodological pillars:

1) **Endogenous Bottom-Up Dynamics**: Rather than imposing aggregate constraints, systemic outcomes are treated as emergent properties driven by the localized behaviors and interactions of heterogeneous individuals.
2) **Agent Autonomy**: The framework assumes decentralized decision-making, where agents operate without a "central controller" or global coordination mechanism, reflecting the fragmented nature of real-world markets.
3) **Generative Explanation**: As articulated by Epstein (2012) within the field of Generative Social Science, the analytical focus shifts toward generative sufficiency. Epstein's principle, which states that a phenomenon is only explained if it can be grown, argues that a macroscopic regularity is only fully understood when a researcher can demonstrate its emergence from a population of decentralized agents.

The methodological and normative imperative for this framework was underscored by Farmer and Foley (2009) in their critique of mainstream modeling following the 2008 financial crisis. They contended that standard Dynamic Stochastic General Equilibrium (DSGE) models, predicated on representative agents and linear dynamics, were inherently incapable of accounting for the complex interdependencies and cascading failures characteristic of modern financial systems. They argued that the field necessitates ABM to capture the nonlinear feedback loops and "black swan" events that define systemic crises. By blending classical and constructive mathematics, agent-based modeling serves as an alternative, potentially more appropriate mathematical framework for the social sciences (Borrill & Tesfatsion, 2011).

Furthermore, the growing demand for instructional programs and educational tools centered on agent-based modeling reflects its utility as a pedagogical approach for studying complex systems

(Macal & North, 2010). Siebers et al. (2010) emphasize that implementing individual behaviors into agent-based models and observing the resulting cumulative effects provides a highly practical and educational experience.

Consequently, ACE provides a robust "in silico" laboratory for policy experimentation. By simulating the path-dependent nature of economic development and institutional complexities, researchers can evaluate how micro-level behavioral rules and interaction structures dictate the long-term resilience or fragility of the macro-economy. This generative approach bridges the gap between individual-level heuristics and aggregate economic phenomena, offering a more granular lens for structural analysis.

1.2.3. The Transition to Data-Driven and Validated Models

Historically, agent-based models were frequently criticized as "toy models" or "opaque black boxes" that, while theoretically insightful, lacked the empirical grounding necessary for quantitative forecasting. However, the recent literature has undergone a "structural turn," emphasizing empirical validation, high-dimensional calibration, and statistical rigor.

A pivotal contribution to this transition was Franke and Westerhoff (2012), who formalized the use of the Simulated Method of Moments (SMM) to rigorously compare competing agent-based frameworks. By establishing a "structural contest" between models of financial markets, they demonstrated that ABMs could be disciplined by empirical data as effectively as traditional econometric models. Their work shifted the focus from qualitative "pattern spotting" toward the objective measurement of how well a model’s simulated moments (such as volatility clustering or fat tails) match observed time-series data.

This empirical evolution has culminated in the contemporary frontier of Data-Driven Economic Agent-Based Models, exemplified by Saxena *et al*. (2026). Leveraging the "Big Data" revolution, their work moves beyond stylized behavioral rules by integrating massive, granular datasets, including longitudinal labor market transitions and high-frequency firm-level interactions. By grounding agent behaviors in real-world micro-data, these models bypass the limitations of the "as-if" assumptions of neoclassical theory, offering significantly higher predictive fidelity for assessing structural shifts and the heterogeneous impacts of economic shocks.

Despite these advancements, the inherent complexity of high-fidelity ABMs, which is characterized by non-linearities and emergent properties, presents substantial challenges for traditional estimation techniques. Because these models often lack an analytically tractable likelihood function, they are frequently classified as "likelihood-free" or "black-box" systems.

This synergy between high-resolution micro-data and advanced inferential machinery marks the maturation of the field. It transforms ABM from a purely theoretical instrument into a rigorous, evidence-based toolkit capable of providing precise, data-validated insights for modern macroeconomic policy.

1.2.4. Theoretical Grounding of Agent Behaviors

As previously established, achieving high-fidelity simulations necessitates that agent behaviors are grounded in empirical observations or modeled to reflect real-world decision-making processes. The table below delineates the foundational literature from which current behavioral frameworks are derived, providing the necessary justifications and contextual commentary for their inclusion in modern agent-based architectures.

Table 1. Foundational Literature for Agent Behavioral Specification and Model Calibration

Reference	Core Behavioral Contribution
Carroll (2009, 1997, 1992)	Precautionary savings and the "Buffer-Stock" model.
Deaton (1989)	Consumption smoothing under liquidity constraints.
Dixit & Stiglitz (1977)	Monopolistic competition
Christiano *et al*. (2010)	Financial Factors in Economic Fluctuations
Gertler, Kiyotaki (2010)	Financial Intermediation and Credit Policy in Business Cycle Analysis
Mortensen and Pissarides (1994)	Search and Matching Theory.
Bernanke, Gertler, and Gilchrist (1999)	The "Financial Accelerator" Mechanism.
Dosi *et al*. (2010)	The "Keynes meets Schumpeter" (K+S) Framework
Sargent (1993), Evans & Honkapohja (2001)	Adaptive Learning in Macroeconomics.
Caiani *et al*. (2016)	Development of the benchmark Agent-Based Stock-Flow Consistent (AB-SFC) macroeconomic framework.

Source: compiled by the author

The household sector's behavior finds its roots in the Buffer-Stock model proposed by Carroll (1992, 1997, 2009) and the liquidity constraints analyzed by Deaton (1989). These works provide the micro-founded justification for agents to maintain wealth buffers and update their income expectations adaptively. By rejecting the assumption of frictionless credit access, the model adopts the behavioral logic of consumption smoothing under uncertainty. This establishes a framework where agents do not merely consume based on current income, but manage a "permanent income" proxy to hedge against the stochastic nature of the labor market. Furthermore, Deaton's establishing of the behavioral logic for agents who cannot borrow is critical for exploring how credit constraints at the micro level lead to aggregate consumption volatility and "hand-to-mouth" dynamics.

For the supply side, the model adopts the monopolistic competition framework of Dixit and Stiglitz (1977) as its structural cornerstone. This provides the theoretical justification for a departure from "price-taking" behavior, granting firms the agency to set prices based on their specific market conditions. To bridge this industrial organization framework with the dynamic macro-behavior of the agents, the model utilizes the "Keynes meets Schumpeter" (K+S) logic developed by Dosi et al. (2010). This literature provides the rationale for replacing hyper-rational optimization with rule-of-thumb heuristics for production planning and capital investment. By following this evolutionary approach, the model's firms navigate market uncertainty through adaptive adjustment parameters rather than global foresight.

The interactions between the real economy and the financial sector are underpinned by theories of asymmetric information and balance sheet constraints. The "Financial Accelerator" mechanism of Bernanke, Gertler, and Gilchrist (1999) serves as the theoretical basis for the risk premium within the firm-bank relationship, where a firm's leverage ratio directly dictates its cost of capital. These financial frictions are further contextualized by Gertler and Kiyotaki (2010) and Christiano et al. (2010), who emphasize how an intermediary's net worth limits credit supply. This provides the foundation for a credit-rationing logic where a bank's equity and a firm's probability of default create a feedback loop that can amplify economic fluctuations through endogenous credit cycles.

The labor market is grounded in the Search and Matching logic of Mortensen and Pissarides (1994). By moving away from a Walrasian clearing market, the model incorporates the informational frictions and coordination failures inherent in job procurement. This decentralized matching process, where agents sample only a subset of the potential market, reflects real-world search intensities and market frictions. This theoretical backing justifies the persistent existence of both involuntary unemployment and unfilled vacancies within the simulation, ensuring that the labor market dynamics are structurally linked to the broader macroeconomic environment.

Finally, the model incorporates the principles of Adaptive Learning as established by Sargent (1993) and Evans and Honkapohja (2001). By rejecting the assumption of Rational Expectations, which requires agents to possess perfect knowledge of the economy's structural equations, this framework assumes that agents act as "econometricians." This provides the theoretical justification for agents to use statistical learning and recursive updating to form expectations. In this context, learning is an endogenous process where agents' beliefs and the actual state of the economy co-evolve, allowing the model to capture transition dynamics and the potential for expectations-driven fluctuations that are absent in static equilibrium models.

In their 2016 paper, Caiani *et al*. propose a methodological framework for Agent-Based Stock-Flow Consistent (AB-SFC) macroeconomics, emphasizing that logical and accounting consistency are as vital as empirical validation.

Below are the rules and practices they outline for model building, accounting, and validation:

1) Accounting consistency must be implemented at the most granular level, starting from the explicit representation of individual agents' balance sheets and decentralized transactions.
2) Every flow of payment must originate from a source and reach a destination; similarly, every financial asset held by one agent must be recorded as a liability for another.
3) Models should avoid common inconsistencies, such as new firms entering with capital and liquidity that appear from nowhere, as these act as unjustified exogenous shocks that undermine model logic.
4) Every model must provide a clear-cut description of the specific financial assets (e.g., deposits, reserves, or cash advances) used to clear every type of transaction.
5) Models should realistically reflect how banks create "inside money" by expanding their balance sheets (granting loans) rather than acting as mere intermediaries of pre-existing savings.

To ground the model in a rigorous methodological framework, This thesis follows the guidelines established by Caiani et al. (2016), which emphasize that logical and accounting consistency are foundational to macroeconomic modeling. By adopting their stock-flow consistent (SFC) approach, model ensures that every transaction is recorded, maintaining the conservation of stocks and flows across all decentralized agent balance sheets. This eliminates "financial black holes" and ensures that money is treated endogenously, where every financial asset held by one agent is explicitly matched by a corresponding liability elsewhere in the system.

The integrity of the model is further maintained through the two-tier validation process advocated by Caiani *et al*. (2016). First, internal accounting validation was conducted using macroeconomic identities and social accounting matrices to verify that the code contains no "leakages" that would violate stock-flow consistency. Second, the model was subjected to external empirical validation, where the resulting artificial time series were measured against established macroeconomic stylized facts, such as GDP, investment and consumption volatility and correlations. This was complemented by checking for microeconomic regularities, such as firm size distributions, to ensure that the model's bottom-up dynamics remain consistent with observed economic phenomena.

1.3. Macroprudential Policy and Systemic Risk in ABMs

The transition from microprudential to macroprudential regulation marks a fundamental shift in the design of financial stabilization policies. While microprudential regulation focuses on ensuring the soundness of individual financial institutions in isolation, macroprudential frameworks recognize that the aggregate behavior of the financial system can generate systemic risk even if individual institutions appear resilient.

Systemic risk in macro-financial systems is largely endogenous. As highlighted by Minsky's (1992) Financial Instability Hypothesis, extended periods of economic stability and low default rates induce agents to take on more risk, thereby transforming stability into instability.

In Agent-Based Models (ABMs), this phenomenon is captured by mapping the feedback loops between the balance sheets of heterogeneous financial intermediaries and the real economy. When banks assess risk using adaptive, short-term forecasting heuristics, they tend to over-expand credit during booms and abruptly restrict lending during downturns. This amplifies business cycle fluctuations and creates procyclical leverage.

Kaszowska-Mojsa (2020) show that financial system interconnectedness can propagate localized shocks. If banks share similar information sets or use the same heuristic to value collateral, aggregate exposure becomes highly correlated.

As Adrian and Shin (2010) document empirically, financial intermediaries adjust their balance sheets actively over the cycle, increasing value-at-risk-based leverage during upturns and selling assets during downturns, which depresses asset prices.

Unlike traditional Dynamic Stochastic General Equilibrium (DSGE) frameworks, which typically rely on a representative agent and focus on equilibrium states, ABMs serve as virtual laboratories to test combinations of macroprudential instruments in and out of equilibrium.

They do this by allowing researchers to evaluate:

1. Lender-based instruments: Such as countercyclical capital buffers and dynamic loan-loss provisioning, which force banks to accumulate reserves during credit booms.
2. Borrower-based instruments: Such as caps on loan-to-value (LTV) and debt-to-income (DTI) ratios, which prevent highly indebted households and firms from destabilizing the financial sector.

While much of the existing macroprudential ABM literature focuses on standard capital requirements or network contagion, this thesis addresses an empirical and institutional gap: the role of information-filtering heuristics within the banking sector.

By analyzing how banks process past defaults, specifically the sample selection bias that occurs when banks exclude zero-debt firms from their risk models, this framework models how credit availability expands during economic upturns. This misjudgment fuels capital accumulation, only to collapse when the accumulation of bad debt triggers a systemic downward adjustment in capital.

1.4. Research Gap and Thesis Contribution

While the literature on macro-financial Agent-Based Models (ABMs) has made significant progress in formalizing the interactions between heterogeneous firms, banks, and the real economy, a critical gap remains in the micro-foundations of institutional learning and risk assessment. Traditional frameworks often assume that banks observe and evaluate the entire population of firms symmetrically when estimating bankruptcy probabilities.

However, this assumption overlooks a fundamental issue of endogenous selection into the debtor population. In the tradition of the Stiglitz-Weiss (1981) framework, credit markets are characterized by a selection process in which only a specific subset of firms, specifically those actively seeking and maintaining credit, enter the pool of potential defaulters.

When banks calibrate default probabilities using a comprehensive dataset that includes zero-debt firms, they introduce a significant informational distortion. Since firms without debt obligations are by definition unable to default, their inclusion in risk metrics creates a class imbalance that artificially dilutes the perceived probability of failure. This leads to a systematic underestimation of default risk during economic expansions, fueling the pro-cyclical credit cycles central to Minskyan theory. Furthermore, the sensitivity of these risk estimations to the historical horizon, particularly the rolling-window specification, is frequently left unaddressed in the literature.

This thesis addresses these gaps by extending the Assenza *et al*. (2015) framework. The model introduces a 10-unit debt threshold that filters out zero-debt firms from the bank's training sample. This prevents the dilution of the bankruptcy probability function and curtails the systematic underestimation of risk.

By showing how information filtering affects credit availability, the model captures the trade-off between systemic stability and cyclical persistence. It demonstrates how eliminating extreme boom-bust cycles lowers GDP persistence to a more empirically consistent level.

By bridging the gap between micro-level sample selection and macro-level financial instability, this thesis provides a new perspective on how institutional data processing heuristics shape macroeconomic resilience.

2. MODEL SPECIFICATION AND METHODOLOGY

This model is adapted from Assenza *et al*. (2015). While prioritizing replication, certain modifications and assumptions were introduced to address original ambiguities and ensure numerical stability.

2.1. General ABM approach

The methodology of this study follows the generative social science paradigm, which posits that to understand a macro-level phenomenon, one must first "grow" it in a controlled *in silico* environment following Tesfatsion framework. Unlike traditional top-down equilibrium models, this research utilizes an Agent-Based Modeling (ABM) approach to explore the emergence of macroeconomic patterns from the bottom up.

2.1.1. Micro-Specification and Behavioral Calibration

The first stage of the methodology involves the granular specification of agents (households, firms, and a bank) whose decision-making processes are not governed by hyper-rational optimization, but by adaptive heuristics. These behaviors are strictly grounded in the empirical and theoretical patterns identified in the foundational literature discussed in Section 1.2.4. By calibrating micro-rules (such as the "Buffer-Stock" consumption rule or adaptive price-setting) to reflect real-world human and organizational behavior, the model ensures that the building blocks of the system are empirically plausible.

2.1.2. Simulation and the Monte Carlo Approach

Once the micro-foundations are established, the model is subjected to extensive computational experimentation. Because ABMs are inherently stochastic and driven by random variables in job search, price shocks, and bank learning a single run is insufficient to draw robust conclusions. Therefore, the methodology employs a Monte Carlo approach, running the simulation a lot of times across various parameter spaces. This process generates a massive synthetic dataset that represents the potential "histories" of the artificial economy.

2.1.3. Validation via Stylized Facts

The validity of the model is determined through a "pattern-oriented" validation process. The synthetic data generated by the simulations are compared against a set of well-documented stylized facts from the real-world economy.

If the model endogenously produces statistical distributions and correlations that mirror these real-world phenomena without being explicitly programmed to do so, it provides a strong epistemological basis to believe that the model's internal mechanisms are functionally similar to those of the actual economy.

2.1.4. Emergence and Predictive Utility

The core philosophy behind this methodology is the concept of emergence. In complex systems, macro-level behaviors are often more than the sum of their parts; they are "emergent properties" that cannot be predicted simply by looking at individual micro-definitions.

By successfully replicating known stylized facts, the model gains structural credibility. This credibility allows us to use the simulation as a laboratory to explore "what-if" scenarios and uncover latent dynamics and relationships or systemic risks that are not yet known or easily observable in real-world data. In this sense, the model moves beyond a mere descriptive tool and becomes a predictive instrument capable of identifying the hidden causal drivers of macroeconomic stability and crisis.

2.2. Description of the Baseline Model

The theoretical framework developed in this thesis is built upon the agent-based macroeconomic model proposed by Assenza *et al*. (2015). While the model maintains the core structural logic of the original work, it introduces several critical modifications designed to enhance financial stability and model longevity. Each of these deviations from the baseline is described separately in the sections that follow.

The model incorporates five distinct agent classes, as illustrated in Figure 2.1: a commercial bank, two household cohorts (workers and capitalists), and two types of production firms—those specializing in consumption goods (C-firms) and those producing capital goods (K-firms).

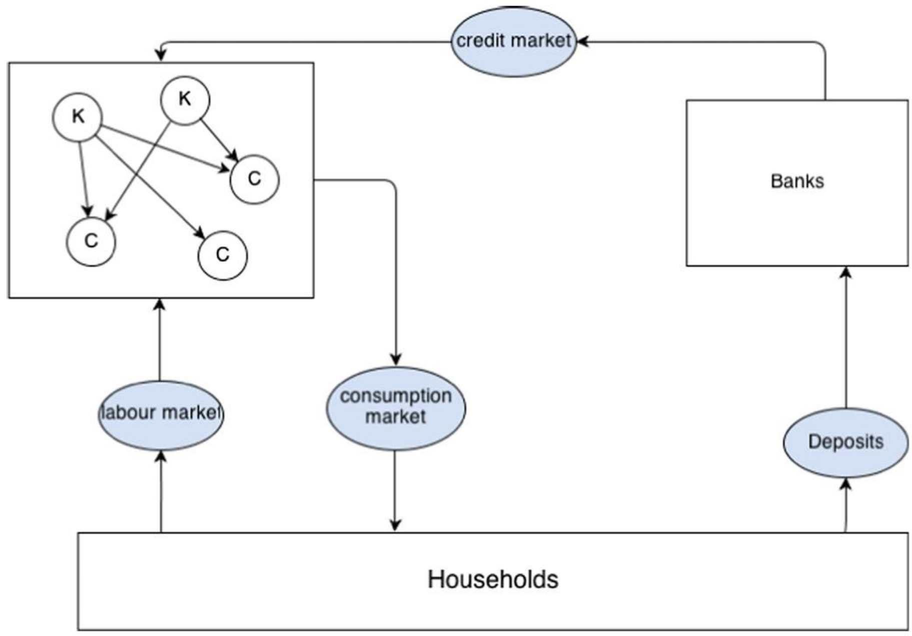

Figure 2.1. Model agents
Source: Assenza *et al*. (2015, 8)

Economic interactions occur across four primary markets: the labour market, the consumption goods market, the capital goods market, and the credit market. A fundamental assumption of the model is that all agents maintain their liquid assets exclusively as deposits within the commercial bank.

The household cohorts are distinguished by their primary source of income. Workers derive their revenue by supplying labour to the market, whereas capitalists' income is comprised of dividend distributions from their respective firms. The ownership structure is defined by a strict one-to-one correspondence: each capitalist owns exactly one firm, and conversely, each firm is owned by a single capitalist.

The agent population remains stationary throughout the simulation, with no market entry or exit. In the event of a firm's insolvency, any remaining liquidity is transferred to the commercial bank. The bank subsequently absorbs the firm's outstanding debt as a loss against its own equity. To ensure the continuity of the production sector, the owner utilizes their personal wealth to reinstate the firm following its bankruptcy.

The model's operations run in the following sequence at each time step:

1) Bankruptcy Processing: The system resolves the bankruptcies and closures from the previous step.
2) Firm Planning: Firms adjust their prices and production quantities, and plan their investment expenditures.
3) Risk Re-estimation: The bank updates its bankruptcy probability function using the rolling-window logistic regression.
4) Labor Demand Calculation: Firms calculate their labor demand based on planned production.
5) Credit Market Resolution: Firms apply for and receive credit from the banks, which is subject to credit rationing and the bank's risk assessment.
6) Labor Market Resolution: The decentralized search-and-matching process occurs, resolving employment and vacancies.
7) Production: Firms execute production using labor and capital.
8) Capital Goods Market: Firms purchase capital goods for future production.
9) Consumption Goods Market: The consumption goods market is resolved, determining total aggregate sales.
10) Accounting Finalization: The system finalizes all transactions, wage distributions, interest payments, and updates the balance sheets to maintain stock-flow consistency.

All of the model's initial values and constants are listed in Appendix 1.

2.2.1. Households and labour market

Household consumption behavior in this model is operationalized through a synthesis of adaptive expectations and liquidity constraints, directly reflecting the Buffer-Stock logic proposed by Carroll (1997) and the consumption smoothing principles of Deaton (1989).

Agents do not possess perfect foresight; instead, they follow a simplified Permanent Income Hypothesis where expectations are formed through Adaptive Learning (Sargent, 1993). Each agent c maintains a proxy for permanent income, $\overline{Y}_{c,t}$, which is updated adaptively based on a memory parameter $\xi \in (0,1)$. When a household receives a wage (or zero if unemployed) or dividends in case of the capitalist household, it updates its permanent income as follows:

$$\overline{Y}_{c,t} = \xi \, \overline{Y}_{c,t-1} + (1 - \xi) Y_{c,t} \quad (1)$$

Where Y current income (wage or dividends). The consumption budget for the period is then determined by both the agent's perceived permanent income and a fraction χ of their current financial wealth (deposits) D:

$$C_{c,t} = \overline{Y}_{c,t} + \chi \mathrm{D}_{c,t} \quad (2)$$

To ensure financial solvency and reflect the real-world credit constraints discussed by Deaton (1989), the model imposes a strict liquidity constraint: the budget cannot exceed the agent's current liquid wealth and is strictly non-negative.

The labor market is characterized by decentralization and imperfect information, moving away from a Walrasian clearing house toward the Search and Matching framework of Mortensen and Pissarides (1994). Workers supply labor inelastically and, if unemployed, engage in a stochastic search process governed by the search parameter Ze. This parameter represents the "frictions" or informational boundaries inherent in real-world markets where agents lack global knowledge of all vacancies.

The matching mechanism functions as follows:

1) Sampling: Each unemployed worker randomly samples Ze firms from the total population of consumption (C) and capital (K) firms.
2) Preference: The worker evaluates the labor demand of these sampled firms. To represent a drive for job security or ease of hiring, the worker identifies the firm with the highest unmet labor demand in their sample.
3) Matching: A match is formed if the chosen firm still has a positive labor vacancy.

Upon hiring, the firm’s demand is decremented, and the worker’s status is updated to "employed". This approach captures the "friction" inherent in real-world labor markets, where workers do not have global knowledge of all vacancies and must rely on local sampling.

2.2.2. Consumption firms

The supply side of the model operationalizes the behavioral principles of bounded rationality and adaptive learning (Sargent, 1993). Firms do not solve a global intertemporal optimization problem; instead, they function as "satisficing" agents that adjust production and pricing based on local signals and heuristics, as established in the K+S framework of Dosi et al. (2010).

Firms operate under conditions of fundamental uncertainty and do not have perfect information regarding the market demand curve. Instead, they form adaptive expectations regarding future demand $Y_{i,t}^{e}$ based on past sales $S_{i,t-1}$ and unfilled orders or "queues" $Q_{i,t-1}$ which combined make actual demand at that step.

$$Y_{i,t}^{e} = Y_{i,t-1}^{e} + \rho(S_{i,t-1} + Q_{i,t-1} - Y_{i,t-1}^{e}) \tag{3}$$

Where $\rho \in (0, 1)$ is the quantity adjustment parameter. Based on this expected demand, the firm sets its planned production level. Simultaneously, firms follow a stochastic rule-of-thumb for price adjustment, consistent with the monopolistic competition setting of Dixit and Stiglitz (1977). The firm adjusts its price $P_{i,t}$ using a stochastic rule-of-thumb relative to the market average price P_{t-1}. If the firm has unsold inventory and its price is above average, it calculates new price with a random shock $P_{i,t-1}(1 - \eta_{i,t})$ with η being a positive parameter drawn from a time invariant uniform distribution with support (0, 0.1). If the firm has a queue (excess demand) and its price is below average, it increases prices.

Production technology follows a Leontief (fixed-coefficient) function where output $Y_{i,t}$ is constrained by both the capital stock and the labor force:

$$Y_{i,t} = \min(\alpha N_{i,t}, k K_{i,t}) \tag{4}$$

where k represents capital productivity and α denotes labor productivity. Furthermore, N is the number of workers employed, and K is the amount of capital possessed by, the firm.

Firms calculate their labor demand by determining the number of workers required to meet planned production, given their current capital constraints. If the current staff exceeds the desired labor, the firm fires the redundant workers; if it is lower, the firm enters the labor market to recruit.

Firms also calculate their liquidity requirements for hiring and capital investment. If their internal liquidity is insufficient to meet these needs, referred to as the financial gap, they request a loan from the bank to cover the shortfall. When applying for credit, firms present their financial data to the bank, specifically their leverage ratio, which the bank utilizes to assess the probability of bankruptcy. The leverage (λ) is calculated as follows:

$$\lambda = \frac{L+F}{E+L+F} \tag{5}$$

Where:

E represents the equity of the firm;

L is the firm's current outstanding debt;

F is the financial gap the firm needs to cover with the new loan.

2.2.3. Capital firms

Capital goods firms (K-firms) represent the upstream sector of the economy, providing the physical capital required by the consumption sector. While K-firms share some core accounting and behavioral logic with C-firms, their production technology and market interactions are distinct. Specifically, K-firms produce capital goods using labor as the sole primary input, representing a labor-intensive upstream industry.

Similar to consumption firms, K-firms operate under conditions of imperfect information and must form adaptive expectations regarding the demand for capital goods. Expected demand $Y_{i,t}^{e}$ is updated recursively based on realized sales $S_{i,t-1}$ and unfilled orders (queues) $Q_{i,t-1}$ from the previous period. The pricing mechanism is also similar to C-firms. The firm's planned production $Y_{i,t}^{p}$ is then determined by the gap between expected demand and current inventory $\Delta_{i,t}$:

$$Y_{i,t}^{p} = Y_{i,t}^{e} - \Delta_{i,t} \tag{6}$$

Production in the capital sector follows a linear technology constrained by labor productivity

$$Y_{i,t} = \alpha N_{i,t} \tag{7}$$

where $N_{i,t}$ is the size of the labor force currently employed by the firm. To meet the planned production level, the firm calculates its desired labor $N^{d} = \frac{Y_{i,t}^{p}}{\alpha}$. To meet the planned production level, the firm determines its desired labor force. If the current number of employees is below this requirement, the firm posts vacancies to be filled within the labor market (the dynamics of which are detailed in the household sector description). Conversely, if the current workforce exceeds the desired level, the firm dismisses redundant workers until the labor supply matches the target. Displaced workers re-enter the labor market immediately; they remain eligible for recruitment within the same simulation step, provided they are successfully matched with an open position elsewhere.

Capital goods firms (K-firms) calculate their financial gap similarly to consumption goods firms (C-firms), with the notable exception that they do not account for capital investment, as capital goods are not utilized in their specific production process.

2.2.4. Bank

The financial sector in this model is represented by a single commercial bank that acts as the central intermediary between liquid savings and firm investment. The bank is not a passive agent; instead, it operationalizes the Adaptive Learning principles of Sargent (1993) and the Financial Accelerator mechanisms identified by Bernanke, Gertler, and Gilchrist (1999).

The bank evaluates the creditworthiness of firms by estimating the probability of default ϕ as a function of the firm's leverage ratio λ. Rather than using a static probability, the bank employs a rolling-window logistic regression. Every period, the bank updates its internal models by analyzing the historical relationship between leverage and bankruptcy events over a specific horizon. The estimated probability of failure for firm i at time t is given by:

$$\phi_{i,t} = \frac{1}{1+e^{-(\beta_0+\beta_1\lambda_{i,t})}} \tag{8}$$

where β_0 and β_1 are the intercept and coefficient learned from the historical data. This mechanism allows for endogenous credit cycles, as the bank’s perception of risk evolves with the macro-economic environment.

Specifically, the rolling-window specification utilizes a maximum window size of 10,000 periods for consumption (C) good firms and 5,000 periods for capital (K) good firms. Testing various alternative window sizes indicates that the specific length of the window does not materially alter the model's dynamics or outcomes. The only exception occurs when the window is set to an extremely small size (i.e., less than the number of firms operating in the model), where a lack of sufficient historical observations leads to unstable parameter estimates.

The bank sets the interest rate r for each loan based on a no-arbitrage condition relative to the risk-free policy rate r^{rf}. The bank assumes an expected survival time $T_{i,t} = \frac{1}{\phi_{i,t}}$ for the firm. The interest rate is derived to ensure the expected return on a risky loan, considering the installment repayment rate θ, is at least equal to the policy rate. The bank solves for r such that:

$$r = \mu \left[\frac{1+\frac{r^{rf}}{\theta}}{\frac{1-(1-\theta)^{T+1}}{\theta}} \right] - \theta \tag{9}$$

Beyond interest rate pricing, the bank manages systemic risk through credit rationing, a behavior informed by the work of Gertler and Kiyotaki (2010) and Christiano et al. (2010). The bank may impose credit rationing to protect its own equity E^{bank}. The maximum credit available to a firm is constrained by the bank's total exposure and a sensitivity parameter ζ:

$$CreditLimit = \frac{\zeta E_t^{bank} - \phi_{i,t} L_{f,t-1}}{\phi_{i,t}} \tag{10}$$

This rule implies that as a firm's probability of default $\phi_{i,t}$ increases, or as the bank’s equity is depleted by loan losses, the available credit for new financing gaps shrinks. This creates a financial accelerator effect where distressed firms are unable to obtain the liquidity necessary for production or investment.

The bank also distributes a portion of its profits as dividends, which are shared equally among all capitalist households. Following the same policy as the firms, the bank distributes 20% of its current-step profits.

2.3. Deviations from Assenza *et al.* (2015)

While the core mechanics of the simulation remain rooted in the original framework, several adjustments were necessary to refine the bank’s behavior and ensure the model's stability. These modifications range from critical structural deviations to minor parameter assumptions required for the model’s successful execution. In some instances, these changes do not represent a fundamental departure from the theoretical intent of Assenza et al. (2015), but were instead essential implementation choices needed to resolve ambiguities or improve the longevity of the credit cycle. The following sections detail these changes separately, distinguishing between major behavioral shifts and the minor technical assumptions adopted for this thesis.

2.3.1. Deviations and Clarifications in Firm Behavior

A primary ambiguity in the original Assenza et al. (2015, pp 13) framework is how firms respond to liquidity constraints, specifically in scenarios where the bank partially or entirely refuses the credit requested for production and capital investment plans. The original authors state:

"Actual production may be smaller than desired if the firm does not succeed in achieving the 'appropriate' level of capital... but also... if the firm does not succeed in achieving the 'appropriate' level of employment... and/or [she] does not get enough credit."

While this suggests that labor and investment are scaled back during liquidity shortages, the exact modeling mechanism remains unclear. In this thesis, the author adopts a hierarchical approach to liquidity management. When hiring and investing, firms first prioritize their existing wage obligations. If current liquidity is lower than the planned wage bill, labor usage is scaled down accordingly. If liquidity is sufficient for the current staff, the firm then determines if it can afford to hire additional workers. Only after these labor requirements are met does the firm evaluate whether its remaining liquidity is sufficient for planned capital investments; if not, investment is scaled down.

Interestingly, both the original paper and this thesis exclude the liquidity required for interest and installment payments when calculating the initial "financial gap." Consistently, these payments are also omitted when recalculating actual hiring and investment levels in liquidity-constrained situations.

A further clarification, which may have been present in the Assenza model but was not explicitly stated, concerns the dividend policy. Following the baseline model, firms distribute 20% of their current-step profits, provided those profits are positive. However, this thesis explicitly assumes that dividends are only paid out if the firm possesses sufficient liquidity to cover the payment after meeting its operational obligations.

Capital goods firms (K-firms) also maintain a price floor of 2 units per unit of production. This threshold is calibrated to the model's labor requirements: since two workers are required to produce one unit of capital and the wage rate is fixed at one unit, the marginal cost of production is consistently equal to two units. While the model could theoretically allow for lower prices, particularly in scenarios where a firm might attempt to liquidate a large accumulation of inventory, this price floor was implemented to ensure model stability and maintain computational simplicity.

2.3.2. Deviations and Clarifications in Bank Behavior

As discussed in the results chapter, the primary challenge in most simulation runs is the solvency of the bank. The most effective, yet structurally simple, deviation from the baseline model involves modifying the data the bank uses for its bankruptcy prediction training. The bank's central problem is a consistent underestimation of bankruptcy probabilities, as defaults occur relatively rarely, causing the training dataset to become skewed toward survival.

Various methods were researched to address this bias, including weighting firm data by debt volume and manipulating interest rates or the base money supply. However, refining the bank's bankruptcy history, specifically how it selects data for training, proved to be the most non-invasive and efficient solution.

Furthermore, a minor modification was made to the bank's behavioral rules regarding dividend distribution. The bank now only distributes dividends if its equity exceeds its initial starting amount. This change is intended to address periods of financial distress: when the bank's equity is low, its ability to generate profit is severely restricted by credit rationing. In such scenarios, distributing dividends would be counterproductive, as it would further deplete the capital buffer needed to restore lending capacity.

2.3.3. Other deviations

To ensure that price fluctuations reflect market dynamics more naturally, the average market price is calculated as a weighted average of all firms' prices, using their respective sales volumes as weights. This adjustment is critical to the decision-making process of individual firms, as their pricing strategies are heavily reliant on the prevailing market average.

2.4. Stylized Facts for Validation

This thesis employs the same stylized facts used in the baseline paper by Assenza *et al*. (2015) to validate the model's performance. These include the standard deviations and first-lag autocorrelations for GDP, consumption and investment. To isolate the cyclical components for statistical analysis, a Hodrick-Prescott (HP) filter was applied to all time series.

The empirical data was sourced from the Federal Reserve Economic Data (FRED) database using the following series:

1) GDPC1 (Real Gross Domestic Product)

2) GPDIC96 (Real Private Domestic Investment
3) PCECC96 (Real Personal Consumption Expenditures)

The sample covers the period from Q1 1955 to Q4 2013 at a quarterly frequency, which was chosen to reflect stable business cycle dynamics while avoiding structural anomalies. To isolate the cyclical components, the time series are transformed into natural logarithms, and the trend is removed using the Hodrick-Prescott filter with a standard smoothing parameter of lambda equal to 1600. The validation approach is both qualitative and quantitative; it qualitatively assesses the model's ability to replicate established stylized business cycle facts, such as the relative volatility of consumption versus investment, while quantitatively comparing the simulated relative standard deviations and autocorrelations against the empirical data.

While the author replicated the statistics for the real-world data using these sources, the results were nearly identical but not perfectly aligned with those reported in the original 2015 study. These minor discrepancies are likely attributable to subsequent data revisions by the Bureau of Economic Analysis (BEA) or changes in the FRED database's historical series since the original paper's publication. All real-world data are expressed in real terms.

Regarding the artificial data, whereas Assenza et al. base their statistics on an average of 20 runs, this thesis utilizes a substantially larger sample of 89 simulation runs to enhance the robustness of the findings. Additionally, while the original authors analyze data from step 1,000 to 3,000, this thesis adopts a different "burn-in" cutoff. Because the model occasionally reaches insolvency near step 1,600 or even earlier, a cutoff point of step 500 was chosen, as the system exhibits stabilized behavior from that point forward.

To test the robustness of the bank's information-filtering mechanism, a sensitivity analysis was incorporated into the simulation methodology. Instead of testing a limited number of discrete points, the debt threshold parameter values were drawn from a continuous uniform distribution between 0 and 15. This approach allows the model to explore the full parameter space without bias. The stability outcomes of these runs were then evaluated using a binary indicator, where simulations reaching the 3,000-step limit were classified as stable, and those terminating prematurely due to insolvency or collapse were classified as crashed.

2.5. Computational Environment and Simulation Design (Monte Carlo Setup)

The model is implemented in Python and structured as a Monte Carlo experiment consisting of N independent realizations. To manage the high-dimensional data generated at each simulation step and avoid memory constraints, the architecture utilizes a real-time data persistence layer. State variables and agent-level attributes are pushed to a centralized MySQL database at the end of every step. This relational storage approach facilitates efficient post-simulation analysis via SQL queries and ensures scalability by offloading the state history from system RAM.

To optimize throughput, the simulation employs run-level parallelization across a hybrid infrastructure of Google Cloud Platform (GCP) virtual machines and local high-end workstations. Each node independently executes self-contained runs and pushes results to the central database. This distributed execution provides asynchronous scalability, where nodes operate at maximum clock speed without throttling, and ensures resilience, as the failure of a single node does not jeopardize the broader experiment. While the current setup focuses on inter-node parallelization, the modular architecture is designed to support future intra-node multi-core optimization using libraries.

In accordance with best practices for computational reproducibility, all stochasticity is strictly managed. Each independent run is initialized with a unique, documented seed, and the Pseudo-Random Number Generator (PRNG) state is recorded at every step. This audit trail allows for the exact reconstruction of any specific realization for debugging or sensitivity analysis. Following the simulation phase, data is processed via a hybridized workflow: Python (pandas, numpy) is used for high-dimensional data extraction and aggregation, while Excel is utilized for final visualization and reporting.

3. SIMULATION RESULTS

3.1. Baseline Simulation Overview

The primary challenge identified during the simulation of this model is the eventual insolvency of the bank, leading to a general economic collapse or, at the very least, invalid results. Figure 3.1 illustrates the dynamics of bank equity alongside the weighted average interest rate for loans during a representative simulation run.

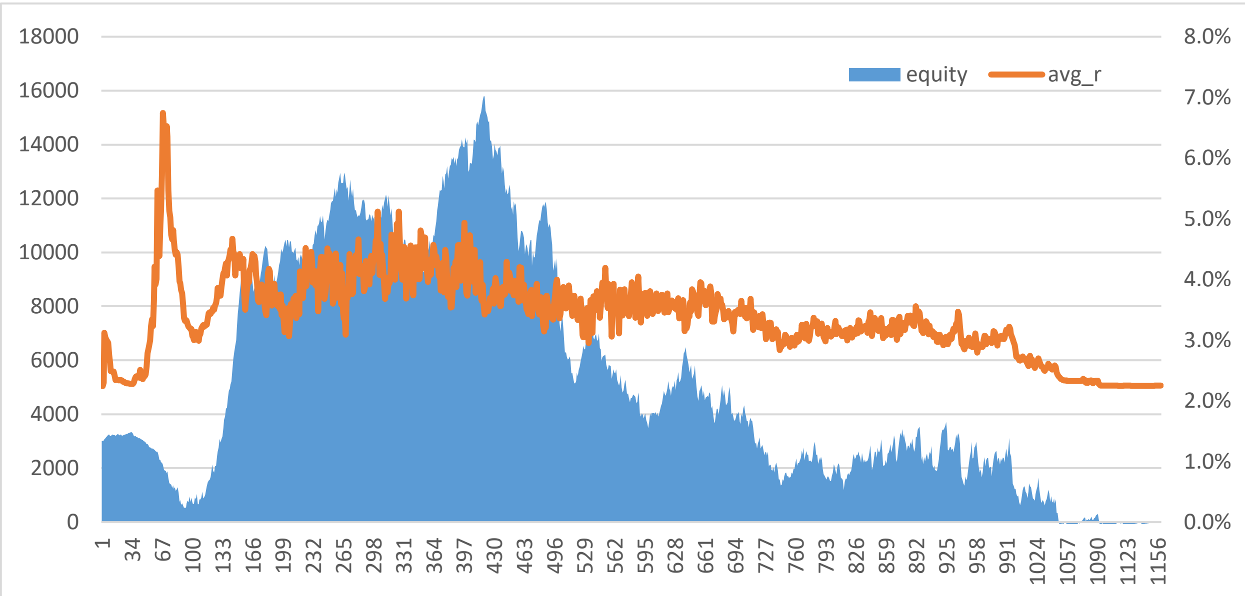

Figure 3.1. Bank equity and average weighted interest rate
Source: author's calculations

During the initial 100 steps, bank equity undergoes a brief contraction before entering a sustained growth phase, peaking between 10,000 and 16,000 units. This is followed by a period of volatility and gradual shrinkage. Around step 760, the system reaches a temporary equilibrium, stabilizing for approximately 250 steps before ultimately descending into insolvency.

The interest rate dynamics reveal the bank's shifting ability to manage risk. Between steps 60 and 100, the bank successfully counteracts rising firm bankruptcies by increasing interest rates, which facilitates an equity recovery. However, from step 490 onward, this mechanism fails to stabilize the system. Although interest rates generally hover between 2% and 5%, significantly above the model's 1% minimum, the resulting interest revenues are insufficient to offset the magnitude of bankruptcy losses.

A critical driver of this insolvency is the bank's endogenous credit rationing. As equity erodes, the bank adopts more aggressive rationing to protect its remaining capital. This creates a dual negative effect: it curtails the bank's primary revenue stream (interest income) and accelerates the failure of struggling firms that lose access to liquidity. This feedback loop ensures that losses rise faster than the bank can generate income, eventually triggering a total collapse.

Once equity reaches zero or turns negative, the bank becomes unable to issue new loans due to its internal regulatory constraints. Consequently, the bank's position is further degraded as old loans are paid out and new bankruptcies continue to drain its remaining value.

Once the bank ceases all credit issuance due to negative equity, a catastrophic chain reaction ensues. Unemployment surges to approximately 90%, and the total production of consumption goods collapses from its peak of 1,100–1,200 units, a level consistent with the average maximum production reported by Assenza et al. (2015), to a mere 100 units. This residual production level is not arbitrary; it precisely reflects the consumption floor of capitalist households.

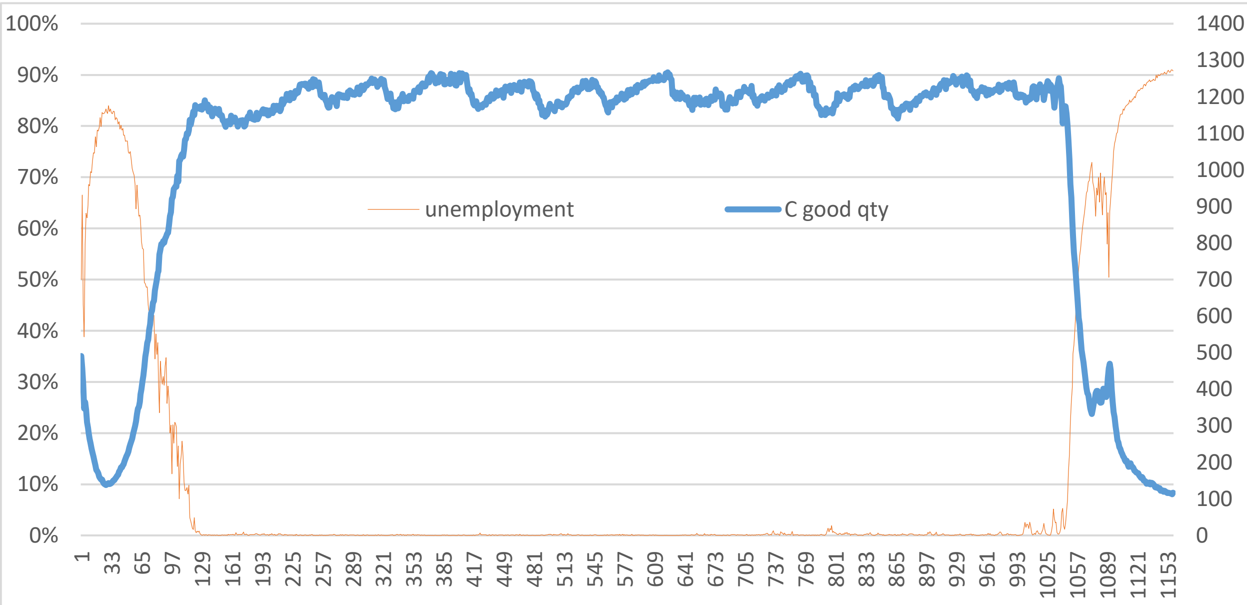

Figure 3.2. Unemployment rate and consumption good sold quantity
Source: author's calculations

The data illustrates a terminal demand-side shock: as employment vanishes, worker household consumption is effectively erased, leading to a precipitous drop in aggregate demand. This contraction further erodes the profits of consumption-good firms, triggering a terminal wave of bankruptcies among firms that still held outstanding debt before the bank stopped issuing loans, eventually resulting in the ultimate collapse of the model. This phenomenon is depicted in Figure 3.2, which illustrates the sharp divergence between the volume of consumption goods sold and the skyrocketing unemployment rate.

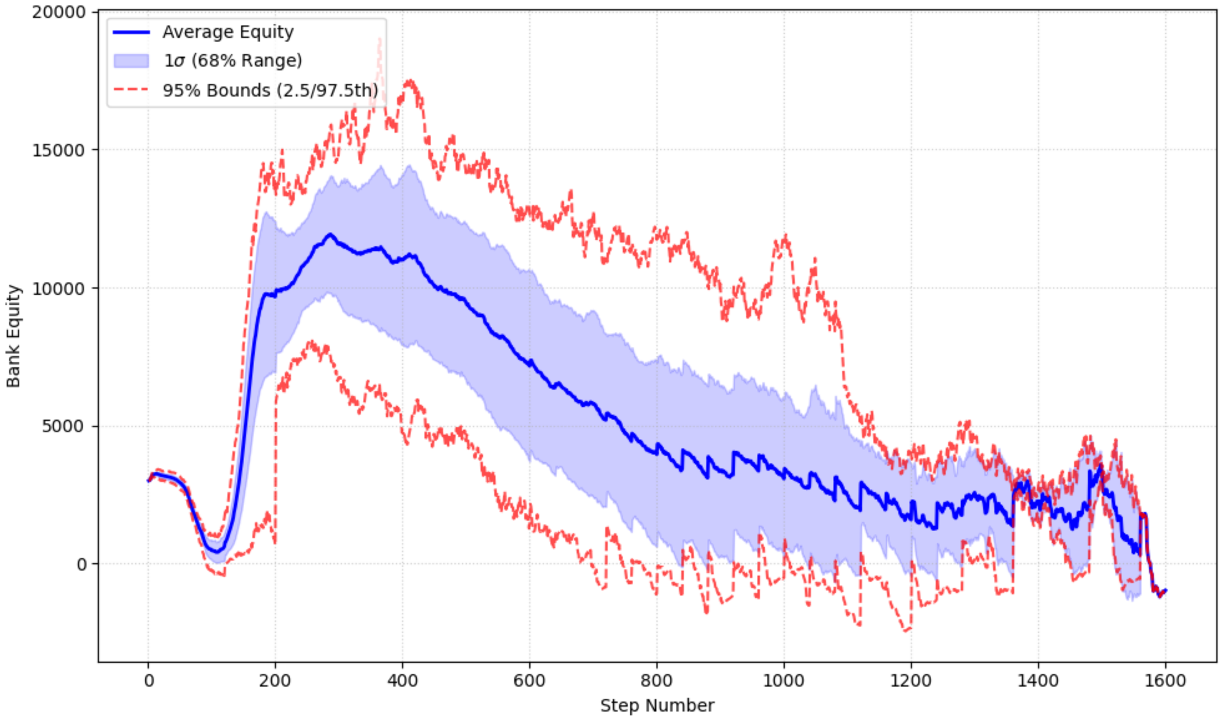

Figure 3.3. Average bank equity, one standard deviation, and 95% confidence intervals
Source: author’s calculations

To assess the robustness of these findings, Figures 3.3 and 3.4 illustrate the mean, one standard deviation, and the 2.5th and 97.5th percentiles for these statistics across 89 simulation runs. The longest simulation reached 1,600 steps, indicating that all other runs resulted in negative bank equity prior to that point. This exhaustion of runs is visible through the jagged fluctuations (or "jumps") in the latter half of the graph. These fluctuations occur because, as runs terminate due to insolvency, the sample size used to calculate the statistics decreases, causing the remaining surviving runs to exert a disproportionate influence on the average.

As shown in Figure 3.3, even with a larger sample of simulations, bank equity eventually shrinks toward zero or becomes negative following an initial growth period during the first 200 steps. This initial phase exhibits remarkably low deviation across all simulations, suggesting a consistent early-stage performance. However, while the bank initially learns and adapts effectively enough to remain solvent or profitable, it eventually gravitates toward a state where revenues are insufficient to offset bankruptcy losses. Ultimately, the bank fails to sustain a long-term equilibrium and inevitably descends into insolvency itself.

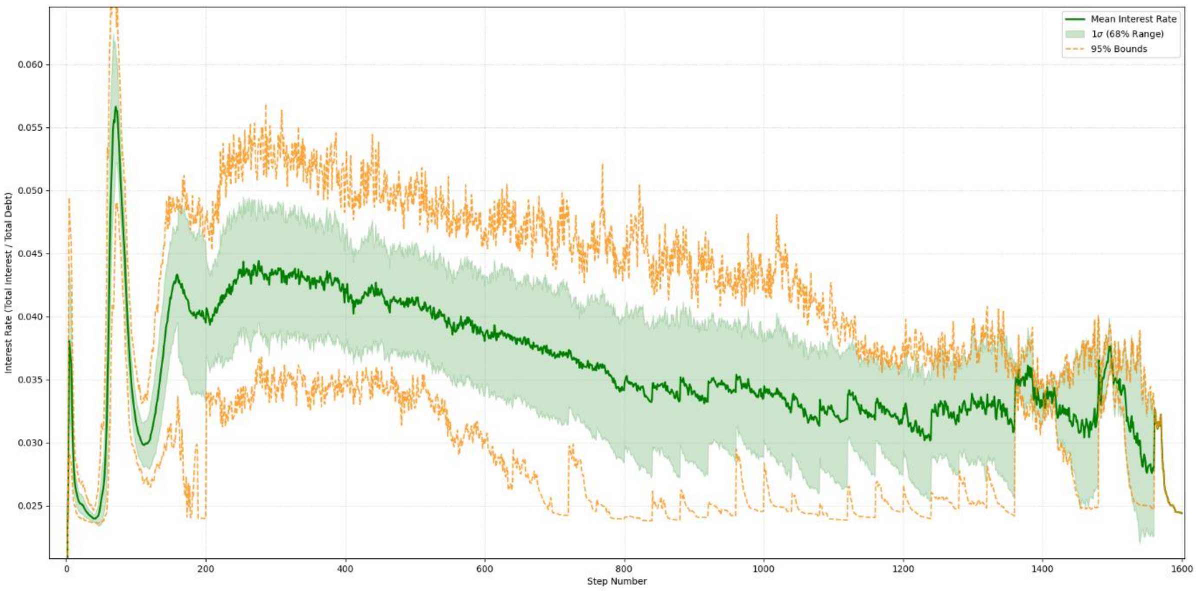

Figure 3.4. Average of interest rate, 1 standard deviation and 95% bounds
Source: author's calculations

The robustness of the interest rate behavior is further illustrated in Figure 3.4. Following the initial 200 steps, the average interest rate peaks at approximately 4.5% around step 210, before gradually declining to 3% over the remainder of the simulation. The upper 2.5th percentile exhibits a similar trajectory, decreasing from 5.5% to 4%.

An analysis of the predictive equation parameters within the bank's bankruptcy model for C-sector firms (based on the same simulation shown in Figures 3.1 and 3.2) reveals a significant trend (Figure 3.5). The bank's internal model rapidly evolves toward a near 0% estimated probability of bankruptcy. Consequently, the only factors maintaining the actual average interest rate at a relatively stable level are the presence of older loans issued at higher rates and the model's enforced minimum bankruptcy probability of 1%.

These results provide insight into potential modifications that could enhance the model's longevity and validity. The primary reason for the underestimation of bankruptcy probability is the emergence of two distinct groups of firms: those that are profitable and those that are not. The successful firms rarely seek bank credit, thereby failing to generate interest revenue for the bank. Conversely, struggling firms frequently require loans but often default shortly after receiving them.

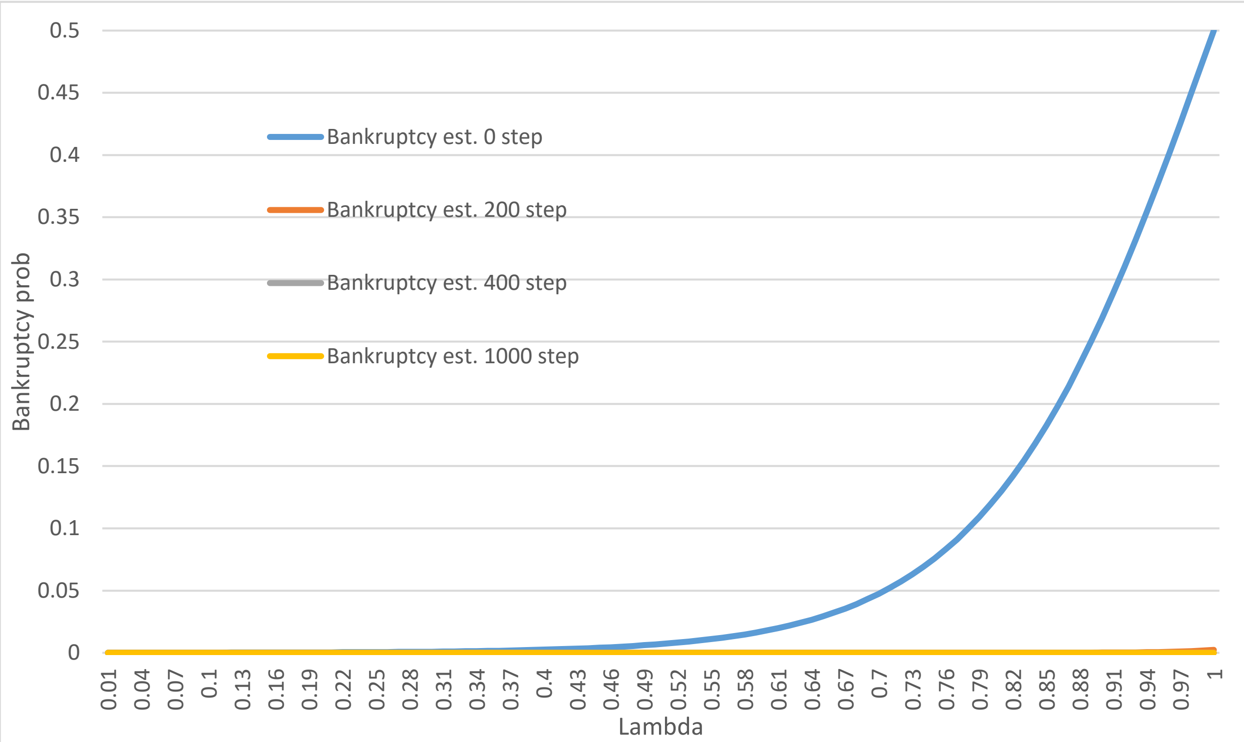

Figure 3.5. Evolution of bank's bankruptcy model for C-sector firms
Source: author's calculations

This dynamic creates a situation where, despite a seemingly low overall percentage of failing firms, defaults occur consistently at every step. This occurs because the most successful firms no longer require external financing and cease borrowing, leaving only struggling firms to continually request new loans. Lacking sufficient internal liquidity, largely due to a lack of prior profits to recapitalize, these struggling firms rely on larger loan amounts and almost immediately face bankruptcy. Consequently, because the bank's risk assessment model only considers the aggregate bankruptcy rate, it fails to account for this risk clustering among the active borrowers, thereby underestimating the vulnerability of its loan portfolio.

To address this issue, the author proposes a revised rule for the bank's bankruptcy history: when estimating bankruptcy probability, the bank will only consider firms that held an outstanding debt of at least 10 units at the end of the step. While this threshold is arbitrary, preliminary testing indicated that a lower minimum (such as 1 unit) yielded no significant difference in results. The impact of this change on bank equity and average interest rates is illustrated in Figures 3.6 and 3.7.

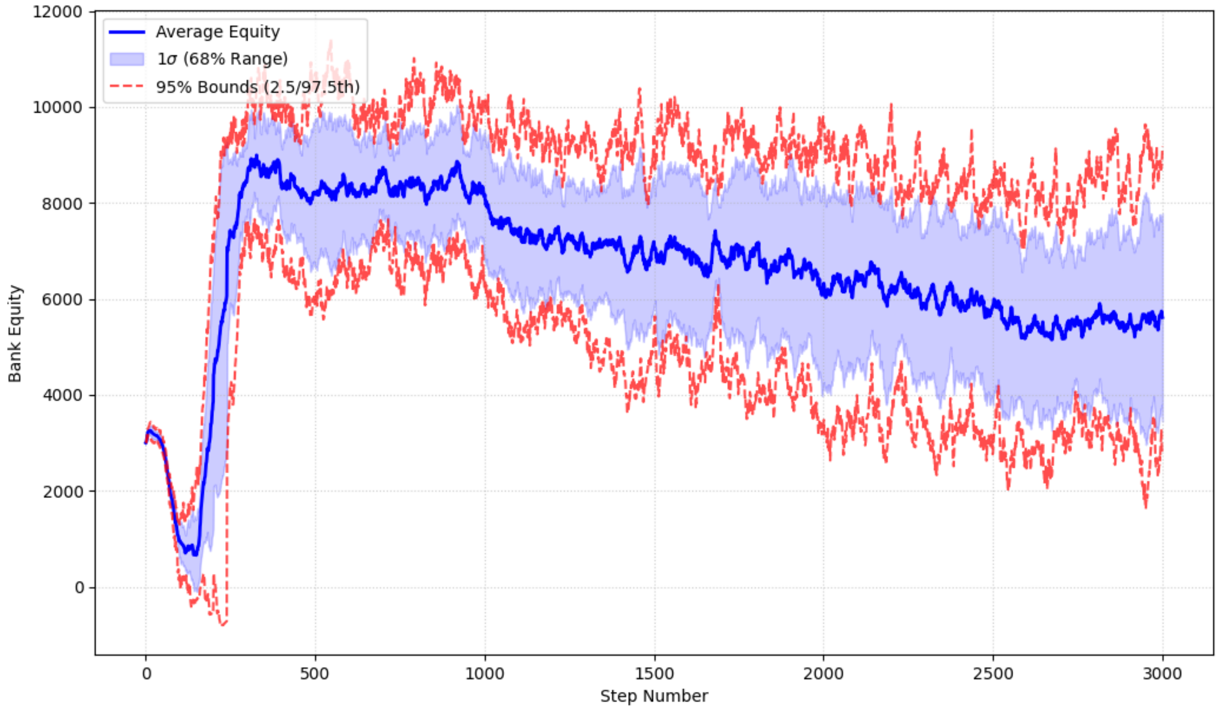

Figure 3.6. Average bank equity, one standard deviation, and 95% confidence intervals after authors change
Source: author's calculations

As illustrated in Figure 3.6, following the proposed modification to the bank's bankruptcy history, bank equity exhibits a similar trajectory to the baseline model for the initial 200 steps. However, unlike the baseline, the subsequent equity contraction occurs at a significantly slower pace. This improved stability allows the simulation to consistently reach 3,000 steps; notably, between steps 2,500 and 3,000, bank equity appears to reach a steady state, showing no further signs of shrinkage.

Figure 3.7 demonstrates that the average interest rate stabilizes after step 200, maintaining a level of approximately 4% for nearly the entire 3,000-step duration of the simulation. The 97.5th and 2.5th percentiles remain relatively narrow and clustered around the mean, fluctuating between just under 5% and slightly over 3%. It is worth noting that these values are expressed in real terms, as the model abstracts from general inflation. Following an initial 200-step adjustment period, the prices of consumption and capital goods exhibit distinct stabilization and volatility profiles: while the consumption good price remains bound within a narrow, stable band, acting as an operational buffer, the price of capital goods displays significantly higher volatility and pronounced cyclical swings. These recurring spikes during boom phases are consistent with the Assenza framework, reflecting heightened competition for capital and liquidity reallocation. Furthermore, given that the nominal wage rate is held constant and production costs comprise solely of wages and capital, it is reasonable to abstract from inflation when evaluating interest rates and firm costs.

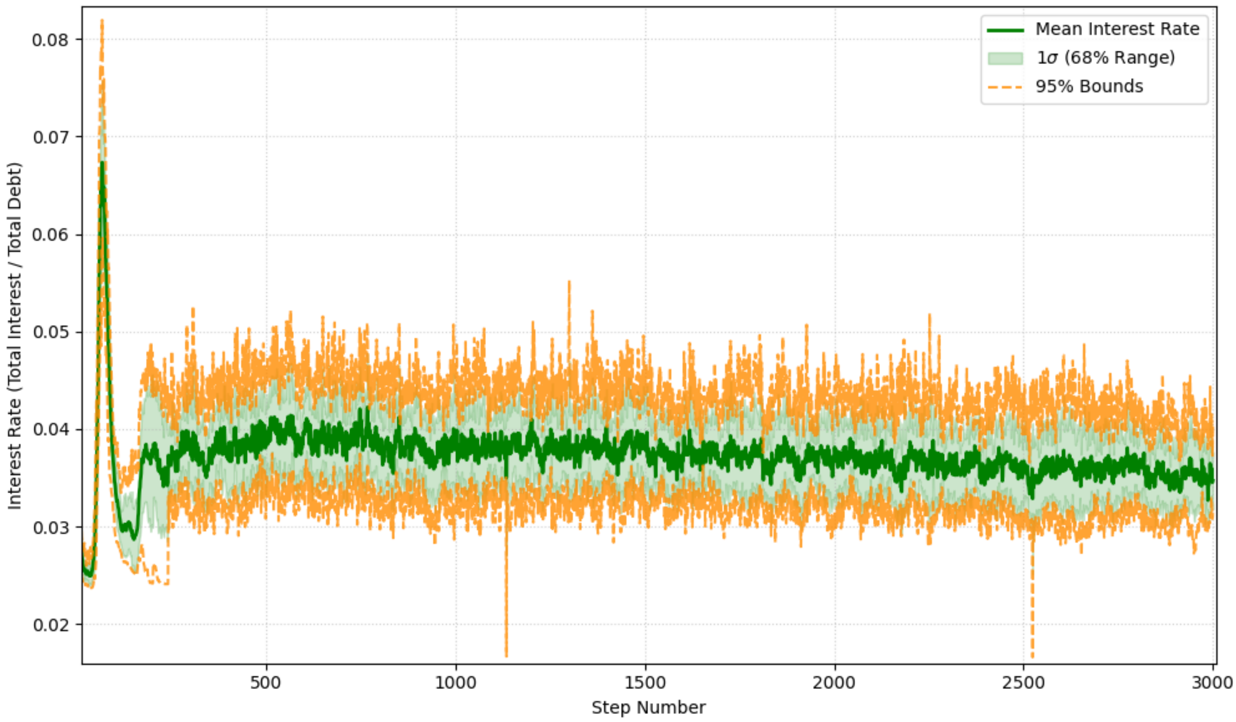

Figure 3.7. Average of interest rate, 1 standard deviation and 95% bounds after authors change
Source: author's calculations

To fully interpret the simulation results, it is essential to distinguish between different sources of variation. Each individual run exhibits its own internal fluctuations, but the means across different simulations can also vary significantly. When calculating the aggregate average over multiple runs, this inter-simulation variance contributes to the overall standard deviation, as seen in Figures 3.6 and 3.8.

Figure 3.8 illustrates the bank equity for ten distinct simulations, with a Hodrick-Prescott (HP) filter applied to highlight the underlying trends (the original unfiltered data is shown in faded colors). The results clearly demonstrate that the model is path-dependent; despite the implemented modifications, the bank still descends into insolvency in certain runs, while achieving stability in others.

For the HP filter, a lambda value of 100,000 was selected. This relatively high smoothing parameter was chosen to effectively isolate the long-term trend from the cyclical component. The resulting gap between the filtered trend and the raw data confirms that the simulated series possesses significant short-term volatility.

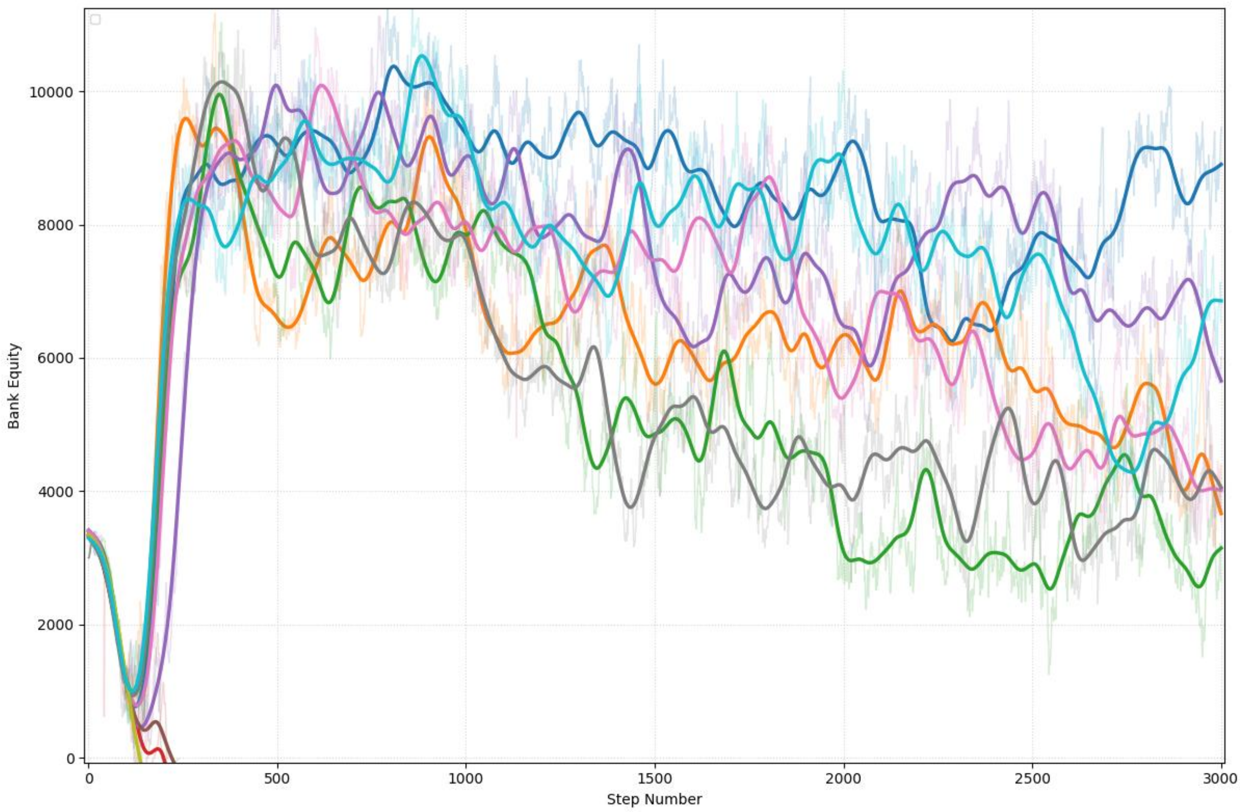

Figure 3.8. Individual Equity Trajectories with HP-Filtered Trends
Source: author’s calculations

The Figure 3.9 illustrates the simulated trajectories of the price of consumption goods (Price C, blue line) and the price of capital goods (Price K, orange line) over approximately 3,000 time steps. The simulation reveals distinct behavioral patterns and cyclical dynamics between the two sectors.

Price C (blue line) starts at 5 at the very beginning of the simulation (Step 1) before falling sharply and stabilizing within the first 80 steps to a baseline range between 2.4 and 2.7. Price K (orange line) starts at 3.2, drops rapidly to a trough between steps 30 and 130, and then begins its long-term cyclical trajectory.

After the initial stabilization, the price of capital goods displays significantly higher volatility and more pronounced cyclical swings than the consumption goods price. The orange line features regular, recurring peaks (e.g., around steps 232, 848, 1541, and 2157).

In contrast, Price C remains relatively stable and bound within a narrow band for the remainder of the simulation. This indicates that the consumption sector acts as an operational buffer while the capital goods sector absorbs the larger structural shocks of the investment cycle.

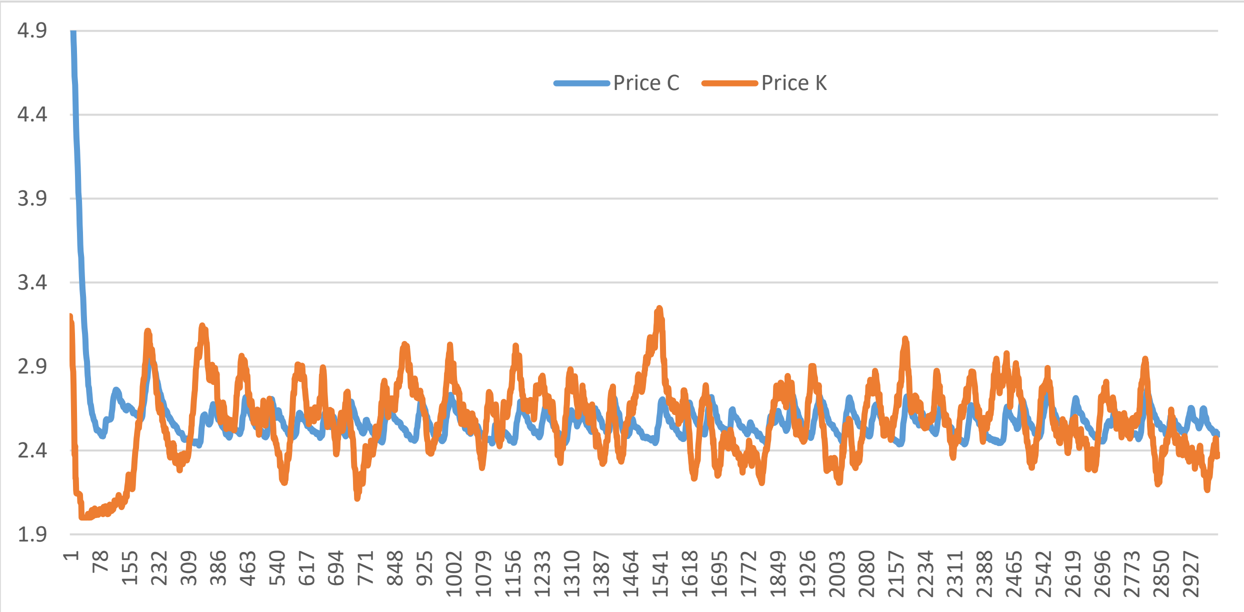

Figure 3.9. Consumption and Capital good prices
Source: author's calculations

The higher cyclical amplitude of Price K is consistent with the Assenza framework, where the boom phase is characterized by a reallocation of liquidity and high demand for capital goods. When firms compete for capital goods, the price spikes, only to drop during the subsequent contraction and deleveraging phase.

Comparing these results to the base model by Assenza *et al*. (2015), it is clear that while the crisis has been successfully recreated (without the banks history changes), the outcomes in this version tend to drive the bank toward inevitable insolvency in a different way. The Assenza model demonstrates crisis development in the following manner:

1) The Boom & Price Spike: As the economy nears full employment, a tight labor market drives up the price of capital goods (K-goods).

2) Liquidity Shift: To afford these expensive K-goods, consumption firms (C-firms) drain their cash reserves and take on massive debt. Liquidity shifts rapidly from the C-sector to the K-sector.
3) Credit Rationing: As C-firm debt and risk rise, banks hit their lending limits and begin rationing credit.
4) The Contraction: Deprived of credit, C-firms, the largest part of the economy, are forced to cut production and investment. This contraction ripples through the system, triggering a general downturn.
5) The Recessionary Reset: During the slump, K-prices fall and firms deleverage. Unemployment rises, but liquidity eventually returns to the C-sector, setting the stage for a new cycle.

Regarding the assumption of full employment in the first point, the model does not include a strict, pre-determined NAIRU (Non-Accelerating Inflation Rate of Unemployment), though it features an analogous endogenous adjustment mechanic. There is no mathematical barrier preventing the unemployment rate from dropping to near zero. However, the labor market functions stochastically: unemployed workers search for jobs by applying to firms at random. Consequently, some firms with open positions might not receive applicants, while others with fewer or no open vacancies might attract job seekers. Because unemployed workers continually search, virtually all firms eventually manage to fill their vacancies. This dynamic keeps the economy in a generally labor-abundant state, except during temporary peaks when nearly all workers are already employed.

3.2. Empirical Validation

This section evaluates the model's ability to replicate the stylized facts of a modern economy. By comparing the statistical properties of the simulated data against the empirical benchmarks derived from FRED, author assess whether the agent-based framework captures the underlying volatility and persistence observed in real-world macroeconomic time series.

3.2.1. Volatility and Standard Deviations

The model's structural validity is assessed by comparing synthetic output against empirical stylized facts derived from FRED (Federal Reserve Economic Data) and the baseline results from

Assenza *et al*. (2015). Tables 3.1 and 3.2 present the first-lag autocorrelations and standard deviations for key macroeconomic aggregates, filtered using the Hodrick-Prescott (HP) trend with smoothing parameters (lambda) of 1,600 and 100,000.

Table 3.1 examines the persistence of economic cycles. Empirical data (FRED) shows high persistence across all variables, with coefficients near 0.80. The authors model without the behavioral change replicates GDP and Consumption persistence well, even exceeding empirical values at lambda=100 000 (0.95). however, like the Assenza baseline, it fails to capture the persistence of Investment, which remains near zero.

Table 3.1. Comparison of the first lag autocorrelations for the chosen series

	Models					
Time series	FRED data	Assenza et al. (2015)	Author's model without change		Author's model with change in bank behavior	
			HP lambda 1600	HP lambda 100 000	HP lambda 1600	HP lambda 100 000
GDP	0,8485	0,6831	0,8310	0,9591	0,1884	0,2822
Investment	0,7952	0,5547	-0,1157	0,0085	-0,1280	-0,0386
Consumption	0,8176	0,6778	0,7105	0.9576	0,4467	0,8298

Source: Author's calculations and Assenza et al. (2015, 21)

The introduction of the bank's information-filtering rule has a divergent effect. While it successfully maintains high persistence for Consumption (0.82 at lambda=100 000), it significantly reduces the autocorrelation of GDP (dropping to 0.28). This suggests that while the bank's more accurate risk assessment stabilizes household behavior, it introduces more frequent, low-amplitude fluctuations in aggregate production.

To further investigate the temporal structure of the simulated economy, Figures 3.10 present the average correlograms (Autocorrelation Function, ACF) for the cyclical components of GDP over 20 lags. These visualizations provide a more granular view of the "memory" within the system compared to the single-lag persistence values reported in Table 3.1.

A comparison between the simulated GDP correlogram (Figure 3.10) and the empirical benchmark derived from real-world data (Figure 3.11) reveals important insights into the model's cyclical properties. While the model exhibits lower initial persistence, the broader structural behavior of the autocorrelation function shows a notable alignment with empirical observations.

As illustrated in Figure 3.11, real-world GDP typically displays high, positive autocorrelation in the short term, which gradually decays as the lag increases. In the simulated environment (Figure 3.8), this decay is visibly accelerated; the "filtered history" mechanism employed by the bank prevents the accumulation of long-term momentum, causing the correlation to drop more sharply in the first three steps.

Despite this difference in initial decay speed, both the simulated and empirical series exhibit a strikingly similar transition point: they both cross the zero-axis at approximately lag 4 or 5. Beyond this point, both correlograms enter negative territory and remain there for the majority of the subsequent lags. In macroeconomic terms, this transition into negative autocorrelation is indicative of a mean-reverting cyclical process. It suggests that the "cycles" in both the model and the real economy have a characteristic wavelength where a period of expansion is typically followed by a contraction (and vice versa) within a specific temporal window.

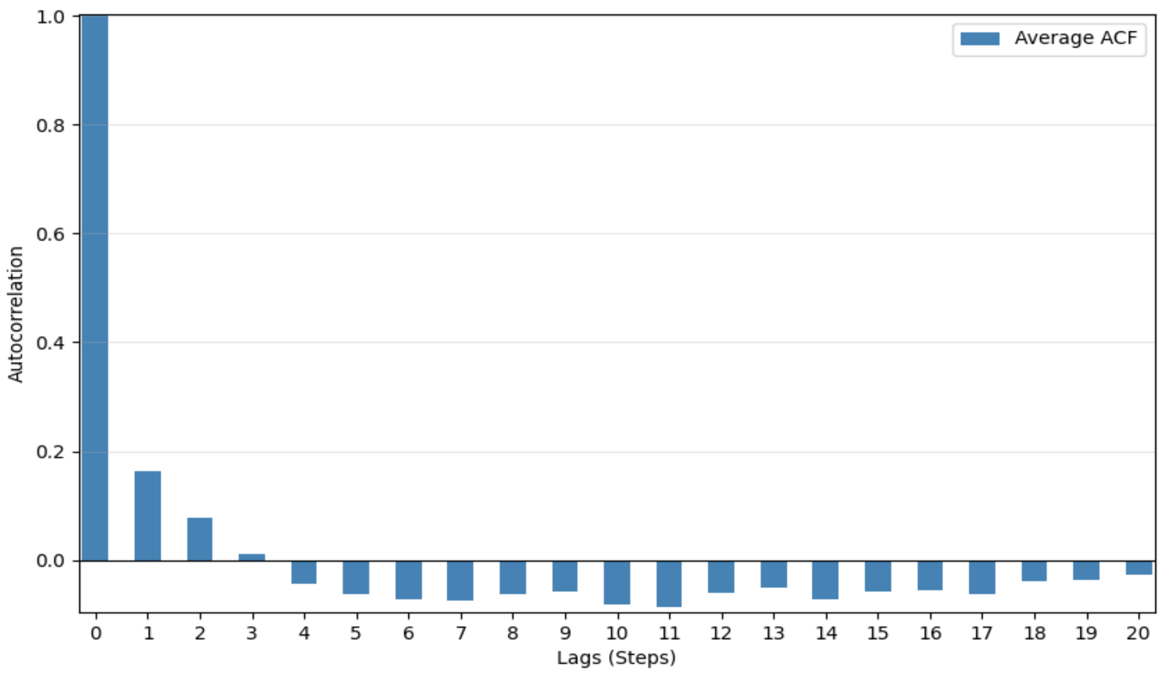

Figure 3.10. GDP correlogram simulated data
Source: author’s calculations

The fact that the model successfully replicates this "zero-crossing" at the 4–5 lag interval suggests that, while the bank’s behavioral changes have reduced the "memory" of the system, the underlying structural constraints of the SFC-ABM framework still produce a cyclical frequency that is qualitatively consistent with observed business cycle dynamics. This provides a degree of validation for the model's internal feedback loops, as it demonstrates that the system is not merely producing random noise, but is instead oscillating around a steady-state trend in a manner comparable to real-world aggregate output.

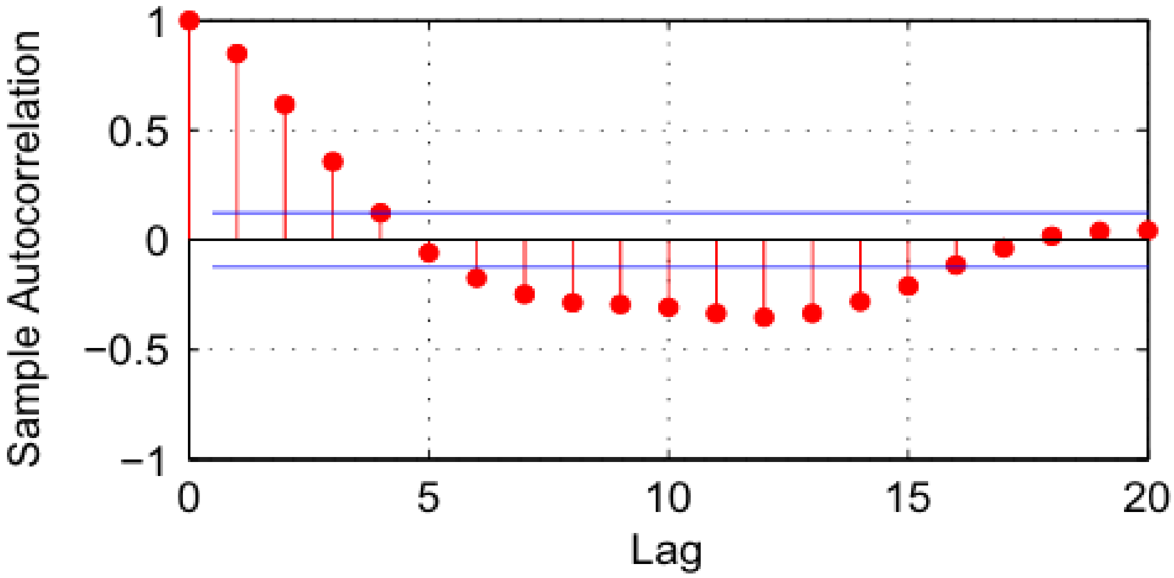

Figure 3.11. GDP correlogram real world data
Source: Assenza et al. (2015, 22)

Table 3.2 highlights the magnitude of fluctuations. A central stylized fact is that Investment volatility is significantly higher than that of GDP and Consumption.

All versions of the model, including the author's modifications, successfully replicate the hierarchy where Investment is the most volatile component. There is a massive reduction in absolute volatility in the model with the bank behavior change. While the baseline and "no change" models produce GDP standard deviations near 1.4–3.6 (close to real-world data), the version with the 10-unit debt filter drops to 0.0028.

This indicates that the information filter is extremely effective at dampening systemic shocks. By preventing the bank from underestimating risk, the system avoids large-scale "boom-bust" credit cycles, resulting in an exceptionally stable but "quiet" macroeconomic environment.

The model remains qualitatively valid as it respects the relative volatility rankings. This confirms that the multi-sector interactions (C-firms and K-firms) are correctly specified. The high autocorrelation in Consumption (0.8298) in the refined model is a strong validation point. It suggests that the bank's improved credit scoring allows for a more stable flow of credit to firms, which in turn stabilizes employment and household spending, matching the "smoothness" seen in real-world data.

Table 3.2. Comparison of the standard deviations for the chosen series

	Models					
Time series	Real world data	Assenza et al. (2015)	Author's model without change		Author's model with change in bank behavior	
			HP lambda 1600	HP lambda 100 000	HP lambda 1600	HP lambda 100 000
GDP	1,6613	1,4369	1,40568	3,60988	0,0027	0,0028
Investment	7,5422	15,2645	17,6982	18,8030	0.2207	0,2302
Consumption	1,2854	1,1784	1,48143	0.0387	0,0094	0,0180

Source: Author's calculations and Assenza et al. (2015, 21)

The change in bank behavior suggests a clear trade-off. By filtering the training data, the bank becomes much more sensitive to risk, which significantly alters not only the real business cycle but also the underlying financial cycle. While data filtering leads to lower output volatility (Table 3.2) and dampens the persistence of GDP (Table 3.1), it has a profound effect on credit growth, bank leverage, and the incidence of banking crises. Specifically, instead of fewer, large-scale crises driven by the gradual buildup of systemic leverage, the economy experiences more frequent, small-scale, and short-lived adjustments. Consequently, the bank's increased risk sensitivity stabilizes credit growth and prevents extreme leverage spikes, mitigating the probability of severe banking crises, but at the cost of long-term cyclical persistence.

To evaluate the robustness of the bank's information-filtering mechanism and confirm that the choice of the 10-unit threshold is not an artifact of specific parameter selection, a comprehensive sensitivity analysis was performed. A total of 1,325 simulations were executed, where the debt threshold values were drawn from a continuous uniform distribution between 0 and 15.

Given the stochastic nature of the agent-based framework, a binary outcome variable was defined to assess long-term viability:

1) Stable (Y=1): Simulations that maintain operational capacity and reach the terminal limit of 3,000 steps.

2) Crashed (Y=0): Simulations that terminate prematurely due to systemic failures, such as banking sector insolvency before reaching the 3,000-step limit.

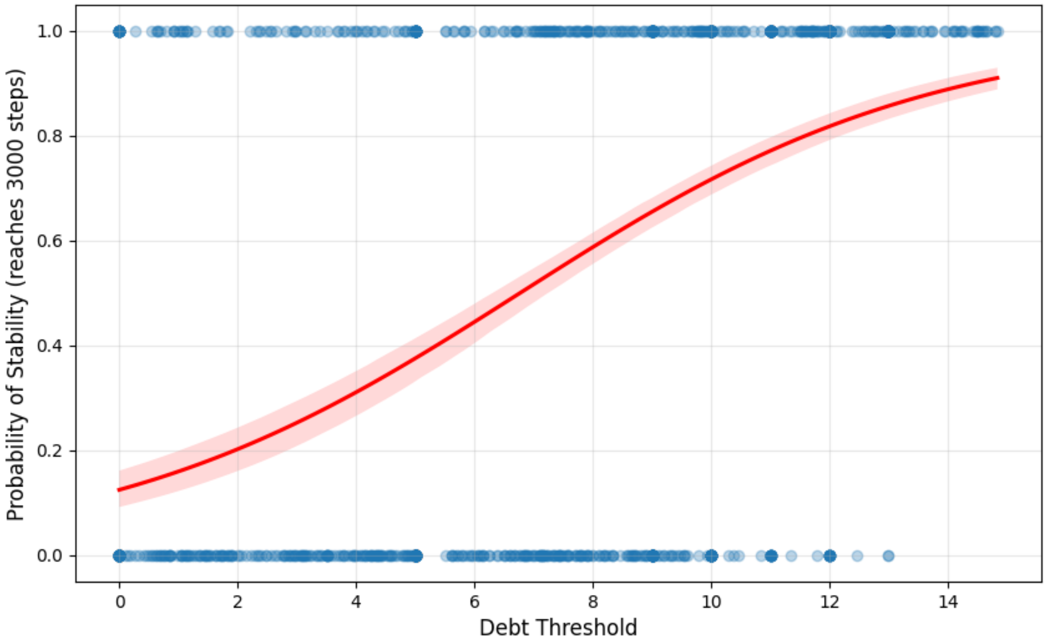

Figure 3.12. Stability of the model and Threshold

Source: Author's calculations

Figure 3.12 illustrates the relationship between the debt threshold and the stability of the model. The results of the sensitivity analysis reveal a distinct non-linear phase transition in systemic resilience. The probability of survival exhibits a classical Sigmoid trajectory across the tested continuous parameter space.

To quantify the relationship between the debt threshold and systemic stability, a binary logistic regression was estimated. The dependent variable is the binary stability indicator (Y), and the independent variable is the continuous threshold value.

The regression results confirm that the debt threshold is a highly significant predictor of systemic stability ($p < 0.001$). The odds ratio indicates that for each 1-unit increase in the debt threshold, the odds of the system remaining stable and reaching the 3,000-step limit increase by 33.3%.

Table 3.3. Logistic Regression Results (Dependent Variable: Stability)

Variable	Coefficient	Standard Error	z-statistic	p-value	Odds Ratio
Intercept	-1.9452	0.149	-13.062	< 0.001	—
Threshold	0.2875	0.018	16.33	< 0.001	1.333

Source: Author’s calculations

The estimated logistic regression model (Table 3.3) demonstrates a robust fit for the sample size of 1,325 observations, yielding a Log-Likelihood of -732.41 and a McFadden Pseudo R^2 of 0.1944. The likelihood ratio test confirms a high overall significance ($p < 0.001$). The regression results reveal that the debt threshold is a highly significant predictor of systemic stability. The odds ratio of 1.333 indicates that for every 1-unit increase in the threshold, the odds of the system remaining stable through the 3,000-step limit increase by 33.3%.

3.3. Policy Implications

The findings of the ablation study (Section 3.1) and the comparative validation (Section 3.2) yield several significant implications for macroprudential oversight and financial regulation. The model suggests that the quality of institutional memory and the granularity of data processing are as critical to systemic stability as traditional capital requirements

A key result of the simulation is that the bank’s underestimation of risk was not caused by a lack of data, but by the inclusion of irrelevant data. In the baseline model, the presence of high-liquidity firms with zero debt "diluted" the bank’s perception of default risk, creating a statistical blind spot.

For regulators, this implies that standardized risk-weighting metrics may be misleading if they do not account for the specific leverage profile of the active borrowing population. Macroprudential policies should encourage financial institutions to adopt "exposure-focused" internal ratings. As demonstrated by the 10-unit threshold, focusing risk evaluations on agents with substantial debt allows the banking sector to avoid excessive credit growth fueled by distorted informational sets, thereby reducing the likelihood of endogenous systemic collapse.

The model illustrates a "financial accelerator" where adaptive learning can lead to rapid credit rationing when localized distress is detected. In the refined model, while the economy was more stable, the bank's increased sensitivity to defaults led to lower GDP persistence.

Because the model exhibits a strong link between credit availability and the hiring capacity of firms (Section 2.3), the results underscore that labor market stability is inseparable from financial stability. In the simulation, decentralized search frictions meant that once a credit-rationing event occurred, the labor market was unable to clear, leading to persistent unemployment.

This provides a rationale for targeted credit support during financial contractions. Since the model shows that firms prioritize existing wage bills over new investment when credit-constrained, policy interventions aimed at providing liquidity for payroll, rather than just capital expenditure, may be more effective at maintaining aggregate demand and preventing the "unemployment traps" observed in the baseline simulation.

Given the highly stylized nature of the model and the limited empirical benchmarking, any policy conclusions drawn from this computational laboratory should be interpreted with caution.

The findings suggest that macroprudential policies should pay attention to the sample population banks use to build their internal risk models. Policies encouraging banks to monitor meaningful risk exposure could prevent the buildup of hidden vulnerabilities.

The relatively "quiet" nature of the refined model suggests that the 10-unit filter may be overly restrictive. Future extensions of this research should explore dynamic thresholds that adapt to the average leverage of the economy, potentially restoring a more realistic amplitude to the business cycle without sacrificing stability.

CONCLUSION

This thesis set out to investigate the endogenous drivers of macroeconomic instability through the lens of agent-based modeling. By extending the framework established by Assenza et al. (2015), this research explored how the intersection of decentralized labor markets and adaptive banking heuristics governs the longevity and volatility of the credit cycle. The central finding of this study is that macroeconomic stability is fundamentally sensitive to the quality of institutional memory within the banking sector.

The primary innovation, the implementation of an adaptive learning bank using a rolling-window logistic regression, demonstrated that credit supply is not merely a passive reflection of the real economy but an active amplifier of its trends. Most notably, the introduction of a 10-unit debt threshold for the bank’s historical training data proved to be a critical stabilizer. By filtering out zero-debt firms that would otherwise dilute the bank’s risk perception, this mechanism corrected a persistent sample selection bias, preventing the systematic underestimation of bankruptcy probabilities that often leads to catastrophic system collapses in baseline models.

The research demonstrates that the bank’s adaptive learning process acts as a powerful feedback mechanism for credit cycle volatility. In the baseline model, the bank’s inability to distinguish between "safe" firms (those with no debt) and "risky" borrowers led to a delusional risk-assessment state, fueling unsustainable credit booms that inevitably ended in systemic collapse. However, the introduction of the information-filtering rule significantly increased the longevity of the credit cycle. By focusing only on firms with meaningful debt exposure, the bank maintained a more realistic perception of default risk. This prevented the buildup of massive leverage, transforming the "boom-bust" cycle into a series of minor, manageable fluctuations.

The model reveals a deep synergy between decentralized search-and-matching frictions and firm-level credit constraints. When firms face liquidity shortages due to credit rationing, the "hierarchical" decision-making process implemented in this thesis, prioritizing existing wage bills over new hiring and investment, serves as a transmission belt from the financial sector to the real economy. Credit constraints do not just reduce production; they freeze the labor market. The inability of firms to secure loans for expansion leads to unfilled vacancies and persistent unemployment, which in turn reduces aggregate demand. This interaction proves to be a primary driver of aggregate instability, as financial distress in the banking sector directly inhibits the labor market's ability to clear.

Regarding the model's ability to replicate empirical data, the results show that a bottom-up framework built on boundedly rational heuristics can successfully capture the relative volatility rankings observed in the real world. The model confirms the stylized fact that Investment is the most volatile component of GDP, followed by Consumption. While the refined model (with the 10-unit filter) produced absolute volatility levels lower than those found in FRED data, it achieved a high degree of Consumption persistence (0.8298), which closely aligns with empirical benchmarks. This suggests that the heuristics used, such as buffer-stock saving and adaptive risk assessment, are sufficient to grow an economy that mimics the qualitative behavior of real-world macro-aggregates.

The validation process confirmed that while there is a trade-off between model longevity and empirical amplitude, the structural integrity of the multi-sector framework remains robust. The "baseline" versions of the model produced volatilities that matched FRED data more closely in the short term but were characterized by premature failures. In contrast, the version incorporating the bank behavior change achieved remarkable systemic resilience. This suggests that the "quiet" nature of the refined model is a byproduct of successful risk mitigation, providing a stable computational laboratory for studying long-term credit dynamics.

The findings suggest that macroprudential policy should focus on the heterogeneity of the bank's information set. Policies that encourage institutions to focus on "meaningful exposure" could prevent the delusional credit expansions that precede crises. However, the low absolute volatility in the refined model suggests that the 10-unit filter may be overly restrictive. Future research could explore dynamic thresholds that adapt to the average leverage of the economy, potentially restoring a more realistic amplitude to the business cycle without sacrificing system longevity.

While the proposed model effectively captures the interaction between institutional information-filtering heuristics, credit availability, and macroeconomic stability, certain limitations must be acknowledged. First, it abstracts from wealth and income distribution dynamics by simplifying household heterogeneity, which limits the ability to evaluate the impact of inequality on aggregate demand. Second, the assumption of a constant nominal wage prevents the generation of endogenous cost-push inflation and wage-price feedback loops. Third, the ten-unit debt threshold used to filter zero-debt firms is fixed, whereas banks might realistically adapt their data-filtering criteria in response to changing conditions. Finally, the closed economy framework excludes external channels of financial contagion and exchange rate volatility that can influence domestic capital investment.

In conclusion, this thesis demonstrates that the "micro-macro gap" can be effectively explored through agent-based simulations. By identifying and correcting the informational biases in banking behavior, this work contributes to a more nuanced understanding of how financial frictions drive aggregate instability. Ultimately, the stability of the macroeconomy depends not only on the health of individual balance sheets but on the accuracy of the lenses through which institutions view and process risk.

KOKKUVÕTE

ENDOGEENSED KREDIIDITSÜKLID JA MAKROMAJANDUSLIK STABIILSUS: AGENDIPÕHINE LÄHENEMISVIIS KOOS PANGA ADAPTIIVSE ÕPPIMISEGA

Igor Mohhov

Selle magistritöö eesmärk oli uurida makromajandusliku ebastabiilsuse endogeenseid tegureid agent-põhise modelleerimise vahendusel. Laiendades Assenza *et al.* (2015) loodud raamistikku, uuriti käesolevas töös, kuidas detsentraliseeritud tööturgude ja adaptiivsete pangandusheuristikate ristumine määrab krediiditsükli kestust ja volatiilsust. Uurimuse keskne järeldus on, et makromajanduslik stabiilsus on fundamentaalselt tundlik pangandussektori institutsionaalse mälu kvaliteedi suhtes.

Töö peamine uuendus – libiseva akna logistilist regressiooni kasutava adaptiivõppega panga rakendamine – näitas, et krediidipakkumine ei ole pelgalt reaalmajanduse passiivne peegeldus, vaid selle suundumuste aktiivne võimendaja. Eelkõige osutus kriitiliseks stabilisaatoriks 10-ühikulise võlaläve kehtestamine panga ajaloolistele treeningandmetele. Filtreerides välja võlata ettevõtted, mis muidu panga riskitaju moonutaksid, korrigeeris see mehhanism püsivat valiminihet ja hoidis ära pankrotitõenäosuse süstemaatilise alahindamise, mis baasmudelites viib sageli katastroofilise süsteemse kokkuvarisemiseni.

Uurimus näitab, et panga adaptiivne õppeprotsess toimib krediiditsükli volatiilsuse võimsa tagasisidemehhanismina. Baasmudelis viis panga suutmatus eristada „turvalisi“ ettevõtteid (neid, kellel ei ole võlgu) ja „riskantseid“ laenuvõtjaid eksliku riskihindamise seisundini, mis toitis jätkusuutmatuid krediidibuume ja lõppes paratamatult süsteemse kokkuvarisemisega. Teabefiltreerimise reegli kasutuselevõtt suurendas aga märkimisväärselt krediiditsükli kestust. Keskendudes ainult ettevõtetele, kellel on märkimisväärne võlakoormus, säilitas pank realistlikuma ettekujutuse maksejõuetuse riskist. See hoidis ära ulatusliku finantsvõimenduse tekke, muutes „buum-languse“ tsükli väiksemate ja hallatavate kõikumiste seeriaks.

Mudel paljastab sügava sünergia detsentraliseeritud otsingu- ja sobitushõõrdumiste ning ettevõtte tasandi krediidipiirangute vahel. Kui ettevõtted seisavad krediidiratsioneerimise tõttu silmitsi likviidsuspuudusega, toimib käesolevas töös rakendatud „hierarhiline“ otsustusprotsess – kus olemasolevaid palgafonde eelistatakse uutele palkamistele ja investeeringutele – ülekandemehhanismina finantssektorist reaalmajandusse. Krediidipiirangud ei vähenda mitte ainult tootmist, vaid külmutavad ka tööturu. Ettevõtete suutmatus saada laenu laienemiseks viib täitmata töökohtade ja püsiva tööpuuduseni, mis omakorda vähendab kogunõudlust. See vastastikmõju osutub koondmajanduse ebastabiilsuse peamiseks põhjustajaks, kuna finantsraskused pangandussektoris pärsivad otseselt tööturu võimet tasakaalustuda.

Mis puudutab mudeli võimet empiirilisi andmeid reprodutseerida, siis tulemused näitavad, et piiratud ratsionaalsusel põhinevatel heuristikatel rajanev alt-üles raamistik suudab edukalt tabada reaalses maailmas täheldatud suhtelise volatiilsuse pingerida. Mudel kinnitab stiliseeritud fakti, et investeeringud on SKP kõige volatiilsem komponent, millele järgneb tarbimine. Kuigi täiustatud mudel (10-ühikulise filtriga) andis absoluutse volatiilsuse tasemed, mis on madalamad kui FRED-i andmetes, saavutas see kõrge tarbimise püsivuse (0.8298), mis on väga lähedal empiirilistele võrdlusnäitajatele. See viitab sellele, et kasutatud heuristikad – nagu puhvervaru säästmine ja adaptiivne riskihindamine – on piisavad, et kasvatada majandust, mis jäljendab reaalsete makroagregaatide kvalitatiivset käitumist.

Valideerimisprotsess kinnitas, et kuigi mudeli kestvuse ja empiirilise amplituudi vahel on kompromiss, on mitut sektorit hõlmava raamistiku struktuurne terviklikkus endiselt tugev. Mudeli „baasversioonid“ tekitasid volatiilsust, mis vastas lühiajaliselt rohkem FRED-i andmetele, kuid mida iseloomustas enneaegne kokkuvarisemine. Seevastu versioon, mis hõlmas panga käitumise muutust, saavutas märkimisväärse süsteemse vastupidavuse. See viitab sellele, et täiustatud mudeli „vaikne“ olemus on eduka riskide maandamise kõrvalsaadus, pakkudes stabiilset arvutuslikku laborit pikaajaliste krediididünaamikate uurimiseks.

Uurimistulemused viitavad sellele, et makroprudentsiaalne poliitika peaks keskenduma panga infokogumi heterogeensusele. Poliitikad, mis julgustavad institutsioone keskenduma „tähenduslikule riskipositsioonile“, võiksid ära hoida kriisidele eelnevaid ekslikke krediidipaisumisi. Kuid madal absoluutne volatiilsus täiustatud mudelis viitab sellele, et 10-ühikuline filter võib olla liiga piirav. Tulevased uuringud võiksid uurida dünaamilisi piirmäärasid, mis kohanduvad majanduse keskmise finantsvõimendusega, taastades potentsiaalselt realistlikuma amplituudi majandustsüklile ilma süsteemi kestvust ohverdamata.

Kuigi kavandatud mudel tabab tõhusalt institutsiooniliste teabe filtreerimise heuristikate, krediidi kättesaadavuse ja makromajandusliku stabiilsuse vahelist vastastikmõju, tuleb tunnistada teatud piiranguid. Esiteks ei võeta selles arvesse rikkuse ja sissetulekute jaotuse dünaamikat, kuna leibkondade heterogeensust on lihtsustatud, mis piirab võimalust hinnata ebavõrdsuse mõju kogunõudlusele. Teiseks takistab püsiva nominaalpalga eeldus endogeense kulu-inflatsiooni ja palga-hinna tagasisideahelate teket. Kolmandaks on nullvõlaga ettevõtete filtreerimiseks kasutatav kümneühikuline võlalävi fikseeritud, kuigi reaalsuses võivad pangad oma andmete filtreerimise kriteeriume muutuvate tingimuste järgi kohandada. Lõpuks välistab suletud majanduse raamistik välised finantsnakkuse kanalid ja vahetuskursi volatiilsuse, mis võivad mõjutada siseriiklikke kapitaliinvesteeringuid.

Kokkuvõtteks näitab see magistritöö, et „mikro-makro lõhet" saab agent-põhiste simulatsioonide abil tõhusalt uurida. Tuvastades ja korrigeerides pangandusliku käitumise teabepõhiseid moonutusi, aitab see töö kaasa nüansirikkamale arusaamale sellest, kuidas finantshõõrdumised põhjustavad koondmajanduse ebastabiilsust. Lõppkokkuvõttes ei sõltu makromajanduse stabiilsus mitte ainult üksikute bilansside tervisest, vaid ka nende vaatenurkade ehk „läätsete" täpsusest, mille kaudu institutsioonid riske vaatavad ja töötlevad.

LIST OF REFERENCES

Adrian, T., & Shin, H. S. (2010). Liquidity and leverage. Journal of financial intermediation, 19(3), 418-437.

Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *The American economic review*, *84*(2), 406-411.

Assenza, T., Gatti, D. D., & Grazzini, J. (2015). Emergent dynamics of a macroeconomic agent based model with capital and credit. *Journal of Economic Dynamics and Control*, *50*, 5-28.

Axtell, R. L. (2001). Zipf distribution of US firm sizes. *science*, *293*(5536), 1818-1820.

Bernanke, B. S., Gertler, M., & Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. Handbook of macroeconomics, 1, 1341-1393.

Borrill, P. L., & Tesfatsion, L. (2011). Agent-based modeling: the right mathematics for the social sciences?. In The Elgar companion to recent economic methodology. Edward Elgar Publishing.

Caiani, A., Godin, A., Caverzasi, E., Gallegati, M., Kinsella, S., & Stiglitz, J. E. (2016). Agent based-stock flow consistent macroeconomics: Towards a benchmark model. Journal of Economic Dynamics and Control, 69, 375-408.

Carroll, C. D., Hall, R. E., & Zeldes, S. P. (1992). The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings papers on economic activity*, *1992*(2), 61-156.

Carroll, C. D., & Samwick, A. A. (1997). The nature of precautionary wealth. *Journal of monetary Economics*, *40*(1), 41-71.

Carroll, C. D., & Toche, P. (2009). *A tractable model of buffer stock saving* (No. w15265). National Bureau of Economic Research.

Christiano, L., Rostagno, M., & Motto, R. (2010). *Financial factors in economic fluctuations* (No. 1192). ECB working paper.

Deaton, A. (1989). Saving and liquidity constraints.

Delli Gatti, D., Fagiolo, G., Gallegati, M., Russo, A., & Richiardi, M. G. (2018). Agent-based models in economics: A toolkit. Cambridge University Press.

Dixit, A. K., & Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *The American economic review*, *67*(3), 297-308.

Dosi, G., Fagiolo, G., & Roventini, A. (2010). Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles. *Journal of economic dynamics and control*, *34*(9), 1748-1767.

Epstein, J. M. (2012). *Generative social science: Studies in agent-based computational modeling*. Princeton University Press.

Evans, G. W., & Honkapohja, S. (2001). Learning and expectations in macroeconomics. Princeton University Press.

Farmer, J. D., & Foley, D. (2009). The economy needs agent-based modelling. *Nature*, *460*(7256), 685-686.

Franke, R., & Westerhoff, F. (2012). Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. Journal of Economic Dynamics and Control, 36(8), 1193-1211.

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

Kaszowska-Mojsa, J., & Pipień, M. (2020). Macroprudential policy in a heterogeneous environment—An application of agent-based approach in systemic risk modelling. Entropy, 22(2), 129.

Kirman, A. P. (1992). Whom or what does the representative individual represent?. *Journal of economic perspectives*, *6*(2), 117-136.

Kirman, A. P. (1993). Ants, rationality, and recruitment. *The Quarterly Journal of Economics*, *108*(1), 137-156.

Lux, T., & Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. Nature, 397(6719), 498-500.

Mortensen, D. T., & Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The review of economic studies*, *61*(3), 397-415.

Minsky, H. P. (1992). The financial instability hypothesis (No. 74). Working paper.

North, M. J., Macal, C. M., Aubin, J. S., Thimmapuram, P., Bragen, M., Hahn, J., ... & Hampton, D. (2010). Multiscale agent-based consumer market modeling. Complexity, 15(5), 37-47.

Saxena, M., Pangallo, M., Caccioli, F., & del Rio-Chanona, R. M. (2026). Machine Spirits: Speculation and Adaptation of LLM Agents in Asset Markets. arXiv preprint arXiv:2604.18602.

Sargent, T. J. (1993). *Bounded rationality in macroeconomics: The Arne Ryde memorial lectures*. Oxford University Press.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of mathematical sociology*, *1*(2), 143-186.

Siebers, P. O., Macal, C. M., Garnett, J., Buxton, D., & Pidd, M. (2010). Discrete-event simulation is dead, long live agent-based simulation!. Journal of Simulation, 4(3), 204-210.

Stiglitz, J. E., & Weiss, A. (1981). Credit rationing in markets with imperfect information. The American economic review, 71(3), 393-410.

Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial life*, *8*(1), 55-82.

Tesfatsion, L., & Judd, K. L. (Eds.). (2006). *Handbook of computational economics: agent-based computational economics* (Vol. 2). Elsevier.

U.S. Bureau of Economic Analysis. (2026). Real Gross Domestic Product (GDPC1) [Data set]. Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/GDPC1

U.S. Bureau of Economic Analysis. (2026). Real Private Domestic Investment (GPDIC96) [Data set]. Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/GPDIC96

U.S. Bureau of Economic Analysis. (2026). Real Personal Consumption Expenditures (PCECC96) [Data set]. Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/PCECC96

APPENDICES

Appendix 1. Parameters and initial conditions

Table A.1. Comparison of the standard deviations for the chosen series

Parameter	Description	Value
T	Number of periods	3000
H	Number of workers	3000
Fc	Number of C-firms	200
Fk	Number of K-firms	50
ξ	Income memory weight	0,96
χ	Wealth consumption ratio	0,05
Eh	Initial households' personal assets	2,0
Ze	Firms visited by unemployed worker	5
Zc	C-firms visited by a consumer	2
τ	Dividend payout ratio	0,2
ρ	Quantity adjustment	0,9
η	Max price adjustment	0,1
w	Fixed wage rate	1
Df,1	Initial liquidity	10
κ	Capital productivity	0,333
γ	Investment probability	0,25
δ	Capital depreciation	0,02
ω	Desired capacity utilization	0,85
Pc	Initial Price of C good	5
α	Labor productivity	0,5
Pk	Initial Price of K good	3,2
Eb,1	Initial bank equity	3000
r	Risk-free interest rate	0,01
μ	Bank's gross mark-up	1,2
ζ	Bank's loss parameter	0,002
θ	Debt installment (repayment) rate	0,05

Source: Compiled by author

Appendix 2. Non-exclusive licence

A non-exclusive licence for reproduction and publication of a graduation thesis[1]

I Igor Mohhov

1. Grant Tallinn University of Technology free licence (non-exclusive licence) for my thesis ENDOGENOUS CREDIT CYCLES AND MACROECONOMIC STABILITY: AN AGENT-BASED APPROACH WITH ADAPTIVE BANK LEARNING

supervised by Nicolas Reigl,

1.1 to be reproduced for the purposes of preservation and electronic publication of the graduation thesis, incl. to be entered in the digital collection of the library of Tallinn University of Technology until expiry of the term of copyright;

1.2 to be published via the web of Tallinn University of Technology, incl. to be entered in the digital collection of the library of Tallinn University of Technology until expiry of the term of copyright.

2. I am aware that the author also retains the rights specified in clause 1 of the non-exclusive licence.

3. I confirm that granting the non-exclusive licence does not infringe other persons' intellectual property rights, the rights arising from the Personal Data Protection Act or rights arising from other legislation.

12.05.2026 (date)

[1] *The non-exclusive licence is not valid during the validity of access restriction indicated in the student's application for restriction on access to the graduation thesis that has been signed by the school's dean, except in case of the university's right to reproduce the thesis for preservation purposes only. If a graduation thesis is based on the joint creative activity of two or more persons and the co-author(s) has/have not granted, by the set deadline, the student defending his/her graduation thesis consent to reproduce and publish the graduation thesis in compliance with clauses 1.1 and 1.2 of the non-exclusive licence, the non-exclusive license shall not be valid for the period*