Dynamic Interactions in Economics:

From Micro-Level Games to Macroeconomic Agent-Based Models



Samuel Wiese

Wolfson College

University of Oxford

Dissertation submitted for the degree of

Doctor of Philosophy

Michaelmas 2024

Acknowledgements

Completing this thesis has been a rewarding journey, made possible by the support of numerous individuals and organisations. While I name many of them below, I remain deeply grateful to everyone who has contributed to this work and my development as a researcher over the past four years.

First and foremost, I would like to express my deepest gratitude to my supervisors, Ani and Doyne, for their invaluable guidance, unwavering support, and constructive feedback. Their expertise and encouragement have been instrumental in shaping this work and fostering my growth as a researcher. They have guided me through all aspects of academic life — from cultivating scientific rigour to delivering engaging talks — and their championing of my work has meant a great deal to me. I have benefitted from Ani's profound knowledge, extensive experience, and genuine compassion. I am especially grateful for our countless technical discussions and for the many informal meetings at University Parks and over lunch at Reuben. I am also immensely grateful to Doyne for the opportunities he has provided, including letting me work on two exciting complementary projects, delivering various talks, and spending a summer in Santa Fe. The freedom he has granted me is reflected throughout this thesis.

I am also thankful to my co-authors. Collaborating on our joint publications has been an enriching experience, and their contributions have greatly enhanced the depth and rigour of this research. They have profoundly shaped the way I think about, consume, value, and discuss science. In particular, I would like to thank Torsten, Marco, and Bassel for shaping my experience in Oxford from the very beginning and for teaching me what it means to do good science. I am grateful to Jagoda for our detailed discussions, which always encouraged me to develop my own approaches rather than rely on imposed solutions; this has been instrumental in fostering my critical thinking and independence. Finally, the rewarding collaboration with Joel introduced me to methods for simulation-based inferences, which remains the most exciting area of research I have encountered during my DPhil.

I am also sincerely grateful to my friends and colleagues at the Institute for New Economic Thinking at the Oxford Martin School, which has been my academic home. In particular, I would like to thank Aymeric, Cephas, Ben, Julian, and Valentina for their camaraderie and support through the shared challenges and

successes in the lab. I am thankful to Dorothy for her patience and unwavering championship. I will fondly miss the group meetings, seminars, trips to the Turf Tavern, and, of course, my desk.

I would also like to thank my colleagues at Baly — Simon, Dustin, Chris, Mads, Laurant, and Zhi. Your practical insights and professional mentorship have greatly enriched my doctoral studies, making them far more productive and enjoyable than they otherwise would have been.

To my friends, thank you for your understanding, humour, and support. Wolfson College has truly become a home away from home, and I will always cherish the bar evenings, foosball tournaments, gym sessions, and punting trips. I would like to thank Madhu for his company during late work sessions in the common room and for his unwavering friendship through both the highs and lows. I've written the majority of this thesis sitting in the Wolfson Café next to Freddy, Oana, Anna, Sarah, Theodor, Immi, and Laura. Our routine has been a comforting constant, and I will miss that the most. I also want to thank Tejo, Wan Ting, Csabi, Estelle, and Caroline for the time we spent together in Oxford and for welcoming me when I visited their respective countries. A special thanks to Paul, whose wisdom, guidance, and friendship have been a continual source of support. My trips to London have enriched my life, and I am grateful for the time spent together. I would also like to thank my friends in Leipzig — Max, Vicky, Fabi, and Franz — for the calls, for visiting me in Oxford, and for making me feel like I never left whenever I visit. Finally, I owe my deepest thanks to my parents for their unwavering love, belief, and Sunday Zoom calls; and to my brothers for their invitations to brunch and to the monastery. Doing a DPhil is challenging, but you all made it more bearable.

Lastly, I acknowledge the generous funding and support provided by the Department of Computer Science, which made this possible.

Abstract

We study economic dynamics using two representative classes of models with interacting agents — (i) a microeconomic model of agents repeatedly playing normal-form games and (ii) a large-scale multi-country macroeconomic model. On the microeconomic level, we provide novel results on the likelihood of convergence of best-response dynamics, under which each player selects the pure strategy that maximises their payoff, given the strategies chosen by others. More generally, we show that convergence to equilibrium is rare under a range of well-known learning rules (such as reinforcement learning or fictitious play) when the games' number of players or number of actions is large. This implies that the assumption that systems converge to an equilibrium may be unrealistic. Therefore, turning to the macroeconomic scale, we employ agent-based models, which, unlike traditional macroeconomic models (e.g. DSGEs), do not rely on this assumption. Specifically, we develop a novel framework for constructing and calibrating large-scale, data-driven, agent-based models of entire economies. Our model incorporates individuals, households, firms, banks, governments, and central banks across 38 countries, encompassing markets for goods, labour, credit, and housing. Each agent is initialised using publicly available microdata by aggregating different data sources of varying levels of granularity, ensuring realistic heterogeneity. We calibrate the model using state-of-the-art simulation-based inference techniques: neural posterior estimation and neural ratio estimation. In both cases, we demonstrate that the calibrated model outperforms a simple benchmark time series model and the current state-of-the-art agent-based model in terms of forecasting accuracy. This improvement is rigorously assessed and confirmed as statistically significant using a Bayes factor approach. Our findings suggest that agent-based models when calibrated with real-world data and advanced inference techniques, offer a compelling alternative to traditional models for understanding and predicting economic dynamics. Our agent-based model is built within a user-friendly software platform we have developed, making it easy to create, calibrate, validate, run, and evaluate alternative models.

Declaration

This thesis is submitted to the University of Oxford in support of my application for the degree of Doctor of Philosophy. It has not been submitted, in whole or in part, in any previous application for any degree or qualification. The work presented in this thesis, including the generation and analysis of data, is entirely my own, except where otherwise acknowledged and explicitly outlined below:

Abstract The abstract was written by me.

Chapter 1 The introduction of the thesis was written by me.

Chapter 2 The chapter on the dynamics of random games was written by me. The selection of learning rules was decided jointly with Marco Pangallo and Torsten Heinrich. Bassel Tarbush proposed the periodic best-response dynamic.

Chapter 3 Section 3.1 was written by me. The theoretical results on the frequency of convergence in Section 3.2 are joined work with Alex Scott, Bassel Tarbush, Marco Pangallo, and Yoojin Yang. The simulation-based results of Section 3.3 were performed by me; the writing of that section incorporated feedback from Marco Pangallo and Bassel Tarbush.

Chapter 4 The chapter on the frequency of games was written by me and incorporated feedback from Torsten Heinrich.

Chapter 5 The literature chapter on agent-based modelling was written by me. The selection of agent-based models in Section 5.3.2 was improved by suggestions from Marco Pangallo. The descriptions of Neural Posterior Estimation and Neural Ratio Estimation in Sections 5.4.2.1 and 5.4.2.2 were provided by Joel Dyer.

Chapter 6 The description of the agent-based model was written by me and improved based on feedback from Doyne Farmer, Ani Calinescu, Jagoda Kaszowska-Mojsa, François Lafond, José Moran, Marco Pangallo, and John Muellbauer. Marco Pangallo wrote the code for converting COICOP household consumption weights into ISIC (see Section 6.4.5.1).

Chapter 7 The chapter on economic forecasting using the agent-based model was written by me, incorporating feedback from Doyne Farmer, Ani Calinescu, José Moran, François Lafond, and Marco Pangallo. The simulations, calibration, and forecasting experiments were performed by me; the work on Bayes factor estimation (Section 7.5) was supervised by Joel Dyer, who also devised the approach for proving statistical significance (Section 7.6).

Chapter 8 The conclusion was written by me.

The project on macroeconomic agent-based modelling (Part II of this thesis, spanning Chapters 5 – 7) was guided by Doyne Farmer and Ani Calinescu. These sections roughly correspond to the resulting preprint (Wiese et al., 2024) and all benefitted from suggestions and edits of my coauthors — Jagoda Kaszowska-Mojsa, Joel Dyer, José Moran, Marco Pangallo, François Lafond, John Muellbauer, Ani Calinescu, and Doyne Farmer. More details on publications are provided in Section 1.4.

Contents

List of Figures x

List of Tables xvi

List of Algorithms xix

1 Introduction 1
- 1.1 Motivation . . . 2
- 1.2 Random Games . . . 4
  - 1.2.1 The Frequency of Convergence . . . 7
  - 1.2.2 The Frequency of Games . . . 8
- 1.3 Macroeconomic Agent-based Modelling . . . 9
  - 1.3.1 The Framework . . . 12
  - 1.3.2 Calibration . . . 13
  - 1.3.3 Forecasting . . . 14
- 1.4 Publications . . . 14
- 1.5 Structure . . . 16

I Random Games 18

2 Dynamics on Random Games 19
- 2.1 Introduction . . . 20
- 2.2 Formalisation . . . 22
- 2.3 Best-Response Dynamics . . . 24
  - 2.3.1 Clockwork Best-Response Dynamics . . . 25
  - 2.3.2 Random Best-Response Dynamics . . . 27
  - 2.3.3 Simultaneous Best-Response Dynamics . . . 29
  - 2.3.4 Periodic Best-Response Dynamics . . . 29
- 2.4 Other Learning Dynamics . . . 30
  - 2.4.1 Reinforcement Learning . . . 30
  - 2.4.2 Fictitious Play . . . 32
  - 2.4.3 Replicator Dynamics . . . 34
- 2.5 Conclusion . . . 35

3 Convergence on Random Games **37**
3.1 Introduction . . . 38
3.2 Theoretical Results . . . 40
3.2.1 Multi-Player Games . . . 40
3.2.2 Two-Player Games . . . 55
3.2.3 Games on Networks . . . 62
3.3 Simulation Results . . . 65
3.3.1 Simulation of Best-Response Dynamics . . . 66
3.3.2 Comparison to Other Learning Rules . . . 71
3.4 Conclusion . . . 78

4 The Frequency of Games **80**
4.1 Introduction . . . 81
4.2 The Frequency of Game Ensembles . . . 82
4.2.1 The Distribution of Nash Equilibria . . . 82
4.2.2 Classes of Games . . . 83
4.2.3 The Structure of the Attractor . . . 84
4.3 The Best-Response n-partite Digraph . . . 87
4.4 The Frequency of Convergent Games . . . 88
4.4.1 Convergent Games with a Unique PNE . . . 89
4.4.2 Convergent Games with Multiple PNEs . . . 94
4.5 Conclusion . . . 96

II Macroeconomic Agent-based Modelling **98**

5 Literature **99**
5.1 Introduction . . . 100
5.2 Conventional Macroeconomic Modelling . . . 101
5.2.1 Review of DSGE Modelling . . . 102
5.2.2 Limitations of DSGE Models . . . 105
5.3 Agent-based Modelling in Macroeconomics . . . 106
5.3.1 Overview . . . 106
5.3.2 Notable Examples of Macroeconomic ABMs . . . 109
5.4 Calibrating Agent-Based Models . . . 117
5.4.1 Likelihood-based Approaches . . . 118
5.4.2 Simulation-based Inference . . . 118
5.5 Conclusion . . . 120

6 A Framework for Macroeconomic Agent-based Modelling 121
6.1 Introduction . . . 122
6.2 Model Overview . . . 123
6.2.1 Description of Agents . . . 125
6.2.2 Data . . . 126
6.2.3 Sequence of Events . . . 128
6.2.4 Economic Aggregates . . . 133
6.2.5 Expectations . . . 137
6.3 Key Differences to the IIASA Model . . . 138
6.4 Agents . . . 140
6.4.1 Banks . . . 140
6.4.2 Central Bank . . . 149
6.4.3 Firms . . . 151
6.4.4 Government . . . 166
6.4.5 Households . . . 169
6.4.6 Individuals . . . 185
6.4.7 Rest of the World . . . 189
6.5 Markets . . . 191
6.5.1 Goods Market . . . 191
6.5.2 Labour Market . . . 192
6.5.3 Credit Market . . . 193
6.5.4 Housing Market . . . 194
6.6 Conclusion . . . 195

7 Economic Forecasting 198
7.1 Introduction . . . 199
7.2 Calibration . . . 200
7.2.1 Prior and Posterior Estimation . . . 201
7.2.2 Estimation Details . . . 204
7.3 Running the IIASA Model as a Benchmark . . . 205
7.4 Forecasting Economic Aggregates . . . 206
7.4.1 Average Forecasting Performance . . . 207
7.4.2 Country-Level Forecasting Performance . . . 207
7.5 Bayesian Model Selection . . . 208
7.5.1 Estimating the Bayes Factor . . . 210
7.5.2 Classification Results . . . 211
7.6 Statistical Significance . . . 214
7.7 Conclusion . . . 216

8 Conclusion 219

Bibliography 223

List of Figures

1.1 Quarterly GDP growth in constant prices for Germany, France, and the United States. The French protests of May 1968, the 2007-2008 financial crisis, and the COVID-19 pandemic are clearly visible. . . 2

1.2 Possible best-response digraphs in 2-player, 3-action games. The value on the left of each action profile represents the payoff to the first (row) player, while the value on the right represents the payoff to the second (column) player. Best responses that are part of Nash equilibria are highlighted in blue; if they form part of a cycle, they are shown in red, and those within the attractor of a cycle are marked in orange. The two leftmost games in the figure exhibit three Nash equilibria and no cycles. The games in the middle contain a two-cycle (where each player alternates moves twice), and the games on the right display a three-cycle. 9

2.1 Illustration of a 3-player, 2-action non-degenerate game (left) and its associated best-response digraph (right). The axes shown in the centre states the coordinate system. 23

2.2 The digraphs in panels (a)-(c) above are all identical and correspond to the best-response digraph of the game shown in Figure 2.1. In the panels, we show the first few elements of paths generated according to the best-response dynamic for different initial profiles and playing sequences. 26

2.3 Example trajectories for the Prisoner's Dilemma and Rock-Paper-Scissors under Bush-Mosteller learning. For the Prisoner's Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, none of the trajectories converge. 33

2.4 Example trajectories for the Prisoner's Dilemma and Rock-Paper-Scissors under Fictitious Play. For the Prisoner's Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, a game with a single mixed Nash equilibrium where both players play any action with equal probability, all trajectories converge to that equilibrium. 34

2.5 Example trajectories for the Prisoner's Dilemma and Rock-Paper-Scissors under Replicator Dynamics. For the Prisoner's Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, none of the trajectories converge. 36

3.1 Panels (a) and (b) provide examples of possible best-response digraphs for 3-player, 2-action games. Panel (c) illustrates a possible best-response digraph for a 3-player, 3-action game. PNEs are denoted by filled black dots. If the path reaches one of the action profiles illustrated as a red dot, then the path can never escape to the PNE, regardless of the playing sequence. 41

3.2 Illustration of Algorithms 3 (the clockwork random walk) and 4 (the coupled dynamic). Panel (a) shows the first elements of a possible path generated according to Algorithm 3, and Panel (b) shows the corresponding path generated according to Algorithm 4. The table in Panel (c) provides details, highlighting environments in bold. . . 48

3.3 Illustration of Algorithms 3 (the clockwork random walk) and 5 (the coupled dynamic with sink $\boldsymbol{x}$). Panel (a) shows the first elements of a possible path generated according to Algorithm 3, and Panel (b) shows the corresponding path generated according to Algorithm 5. . 52

3.4 Illustration of possible paths for 2-player, m-action games. We arbitrarily set the initial action profile to be $(1, 1)$ in period 0. In period 1, player 1 either plays action 1 (left branch) or another action (right branch), which we arbitrarily call action 2. Player 2 then responds in period 2, and so on. All red leaves are Nash equilibria, and all blue leaves are profiles that belong to best-response cycles. . 58

3.5 Frequency of convergence of the clockwork sequence best-response dynamic to $2k$-cycles in 2-player games, using the exact formula from Theorem 3. The vertical black line on the right corresponds to Formula (3.46), showing that the probability of convergence is roughly uniform up to $2 \cdot \sqrt{1000}$. 59

3.6 Frequency of convergence of clockwork best-response dynamics on 4-player games with varying numbers of actions and under different networks. The frequency of convergence decreases with the number of actions and with higher network density. 65

3.7 Frequency of convergence under the clockwork best-response dynamic, verifying Formula (3.46) for 2-player games (solid black line) by comparing it to numerical results for 2-, 3-, and 4-player games in the top panel. We also verify Corollary 1 by showing that the analytical approximations (dashed black lines) match the numerical results. In the bottom panel, the horizontal axis is rescaled according to Corollary 1, showing that the numerical data collapses even for a small number of players and actions. 67

3.8 Frequency of convergence to PNE under clockwork best-response dynamics. The solid line corresponds to $\Gamma = 0$, the darker dashed lines to $\Gamma = \frac{n-1}{5}, 2\frac{n-1}{5}, \ldots, (n-1)$, and the lighter dashed lines to $\Gamma = -0.2, -0.4, \ldots, -1$. The frequency of convergence decreases with the number of players and the number of actions, but increases with the level of correlation. 68

3.9 Frequency of convergence to pure Nash equilibria under clockwork, random, and simultaneous best-response dynamics for different numbers of players and actions when payoffs are independent. The frequency of convergence decreases with the number of actions and players for clockwork and simultaneous updating. The random sequence best-response dynamic almost always converges to a PNE in games with a PNE (Amiet et al., 2021). Simultaneous best-response dynamics are less likely to converge than the other two across all observed game ensembles. 69

3.10 Frequency of convergence of clockwork best-response dynamics in 2-action games with uncorrelated payoffs and with varying number of edges in the network of players. We also highlight the frequency of convergence for the corresponding regular network. The frequency of convergence decreases with the number of edges. 70

3.11 Frequency of convergence under the p-periodic best-response dynamic in 3-player games for 2 actions and 3 actions. The probability of convergence to a Nash equilibrium for p-periodic playing sequences increases in p. 71

3.12 Conditional speed of convergence to a PNE for p-periodic best-response dynamics in 3-player games. The average number of time steps required to reach equilibrium increases with p. 72

3.13 Frequency of convergence to PNE under clockwork, random, and simultaneous best-response dynamics compared to the frequency of convergence to (mixed and pure) Nash equilibria under Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics for games with independent payoffs. Except for multi-player random best-response dynamics, the convergence frequency decreases with the number of actions. The clockwork best-response dynamic closely tracks the convergence behaviour of complex learning dynamics. 74

3.14 Frequency of convergence to PNE in 5-action games under clockwork, random, and simultaneous best-response dynamics compared to the frequency of convergence to (mixed and pure) Nash equilibria under the Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics, with varying levels of correlation Γ. The frequency of convergence increases with higher correlation. Both the clockwork and random sequence best-response dynamics track the other learning rules fairly well, the simultaneous best-response dynamic does not. 75

3.15 Frequency of convergence to PNE under clockwork, random, and simultaneous best-response dynamics against the frequency of convergence to (mixed or pure) Nash equilibria under Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics, for varying values of n, m, and Γ. We show the Pearson correlation for each combination, finding the clockwork best-response dynamics are the best predictor for convergence. The identity line is included as a reference. 76

3.16 Trajectories of Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics imposed on best-response digraphs of different 3-player, 2-action games. Blue trajectories converge to Nash equilibria (pure or mixed); red trajectories do not. Blue arrows correspond to pure Nash equilibria, light blue arrows lead there, red arrows correspond to cycles, and orange arrows lead there. 77

3.17 Frequency of convergence of clockwork best-response dynamics compared to Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics on 2-action games with varying numbers of players and different networks of players. The frequency of convergence of the best-response dynamic stays close to the frequency of convergence of the other learning rules. 78

4.1 All possible best-response digraphs in 2-player, 3-action games with three pure Nash equilibria. The value on the left of each action profile represents the payoff to the first (row) player, while the value on the right represents the payoff to the second (column) player. All actions are part of Nash equilibria. 85
4.2 All possible best-response realizations in 2-player, 3-action games with one 3-cycle. The value on the left of each action profile represents the payoff to the first (row) player, while the value on the right represents the payoff to the second (column) player. The best responses that are part of a cycle are shown in red, and those within the attractor of a cycle are marked in orange. 86
4.3 A 3-player, 2-action game with one Nash equilibrium and the corresponding 3-partite graph representation under the clockwork playing sequence. The red arrows correspond to the actions of player 1, the black arrows to the actions of player 2, and the red arrows to the actions of player 3. The Nash equilibrium (I-IV-V) is highlighted by thicker arrows in the digraph. 88
4.4 For 3-player, 2-action games the full graph G^{f} on the left and the condensed graph G^{c} on the right. The PNE corresponds to the environment where every player chooses their first action. 91
4.5 The frequency of randomly drawn games with a unique PNE following Theorem 8. For comparison, the solid black lines denote frequencies obtained by numerically sampling over 500 randomly created games with independent payoffs. The frequency of convergent games decreases with the number of players and the number of actions. . . 94
4.6 The frequency of randomly drawn convergent 2-player games that have a given number of PNEs following Theorem 10. The frequency of convergent games decreases with the number of PNEs, and, generally speaking, with the number of actions. 96

5.1 Estimating the posterior using simulation-based inference. The process involves sampling a parameter set $\boldsymbol{\theta_1}$ from the prior, running the agent-based model to generate simulation outputs $p(\boldsymbol{x} \mid \boldsymbol{\theta_1}, \mathcal{M})$, and using the simulations to estimate the likelihood function $\hat{p}(\boldsymbol{y} \mid \boldsymbol{\theta_1}, \mathcal{M})$ at the observed data $\boldsymbol{y}$. The illustration follows Dyer (2022). 119

6.1 The basic structure of the agent-based model with the most important datasets used to calibrate each agent. These include the inter-country input-output tables from the OECD, World Bank data, Compustat microdata for firms and banks, Household Finance and Consumption Survey data for households and individuals, and policy rates provided by the Bank for International Settlements. 124

6.2 Initial distributions of firm sectors, number of employees, deposits, and debt for Australia, Austria, and Canada in 2013 at a scale of 1:1000, as obtained from OECD and Compustat data. 155
6.3 The maximum number of firms recorded in Compustat from 2013-Q1 to 2017-Q4 by country. 156
6.4 Initial distributions of household income, wealth, and debt for Australia, Austria, and Canada in 2013 at a scale of 1:1000, as obtained from the Household Finance and Consumption Survey. 174

7.1 Forecasts for real GDP, consumption, and investment for Austria from 2015-Q1 to 2018-Q1. The green line represents true data, the yellow line corresponds to an iterated AR(1) multiperiod forecast, the blue lines indicate the average of 1000 simulation runs of our model, and the orange line corresponds to the IIASA model. For this particular country and period of time, our model with NPE performs exceptionally well in predicting consumption; however, all models struggle to predict investment accurately. 200

List of Tables

1.1 An overview of our results related to the existing literature on learning dynamics on random games: BRS19 (Pangallo et al., 2019), PT92 (Opper and Diederich, 1992), EWA13 (Galla and Farmer, 2013), COR24 (Mimun et al., 2024), EWA18 (Sanders et al., 2018), BRS24 (Heinrich et al., 2023), CG21 (Wiese and Heinrich, 2022), RG21 (Amiet et al., 2021), CON23 (Johnston et al., 2023), EVO10 (Gokhale and Traulsen, 2010). We are highlighting clockwork, random, periodic (per.), and simultaneous (sim.) best-response dynamics, as well as Bush-Mosteller (BM) learning, fictitious play (FP), and replicator dynamics (RD). 6

1.2 Comparison of our model to other large macroeconomic agent-based models in the literature regarding covered countries, sectors, heterogeneity, and ability to be calibrated to match time series. 11

3.1 The value of the function $f_{\langle \vec{\mathbf{a}}, s_c \rangle}(t)$ for the first few elements of the paths generated according to the clockwork playing sequence in Panel (a) and Panel (b) of Figure 2.2. The path in Panel (a) reached the PNE in period 3, the path in Panel (b) reached a best-response cycle in period 1. 49

5.1 Comparison of 11 macroeconomic agent-based models against the one we are building Chapter 6 in terms of agents, markets, and relation to data (calibration, initialisation, stylised facts, and forecasting). Our model features all types of agents and markets, and is validated by forecasting economic time series instead of satisfying stylised facts. 110

6.1 Classification of economic activities NACE Rev. 2 Level 1. 125

6.2 Data sources used for model initialisation and calibration: from the OECD, the IMF, the World Bank, the BIS, the ECB, and Compustat (S&P). 128

6.3 Aggregate variables in the model, in particular related to production, prices, and credit. 134

6.4 Key differences of our model to the IIASA Model in terms of scope and agent behaviour. . . . 140
6.5 Variables and parameters in the model related to banks, in particular on balance sheets, the supply of credit, interest rate setting, and regularity requirements. . . . 142
6.6 Variables and parameters in the model related to loans, in particular principal amount, interest rate, and maturity. . . . 143
6.7 Variables and parameters in the model related to the central bank, in particular for setting the policy rate. . . . 150
6.8 Variables and parameters in the model related to firms, in particular production, demand, inventory and stocks, labour, prices, and balance sheets. . . . 153
6.9 Compustat codes for firms, in particular on the number of firms by sector, the number of employees of each firm, and each firm's debt and deposits. . . . 156
6.10 Variables and parameters in the model related to the government, in particular, tax rates, expenditures, deficit, and debt. . . . 167
6.11 Variables and parameters in the model related to households, in particular income, consumption, investment, wealth, debt and housing. 171
6.12 Household Finance and Consumption Survey codes related to households, in particular income, assets, liabilities, household, and consumption. . . . 175
6.13 Variables and parameters in the model related to individuals, in particular income and labour. . . . 185
6.14 Household Finance and Consumption Survey codes related to individuals, in particular personal characteristics and income. . . . 186
6.15 Matching between the Household Finance and Consumption Survey codes for the labour status of an individual to the activity status in the model. . . . 186
6.16 Variables in the model related to the rest of the world, in particular trade and prices. . . . 189
6.17 Variables and parameters in the model related to labour market clearing. . . . 192
6.18 Parameters in the model related to credit market clearing. . . . 194
6.19 Variables in the model related to properties, in particular prices and rent. . . . 194
6.20 Aspects of the agent-based modelling framework that could be improved in the future, with references to other macroeconomic agent-based models/frameworks that incorporate them. . . . 197

7.1 The prior and the corresponding posterior moments estimated by NPE for Austria between 1990-Q1 and 2013-Q1. U($\{0, 1\}$) refers to binary parameters chosen with equal likelihood, U($[0, 1]$) to the continuous uniform distribution on the closed interval $[0, 1]$. The parameters in the prior are related to firm demand predictions, price setting, and target production. 202
7.2 Datasets used for calibration of the IIASA model (Poledna et al., 2023a) with corresponding Eurostat and OECD codes. 206
7.3 The first super-column ("AR1 RMSE") shows the median RMSE of an AR(1) model predicting five different economic aggregates for six different forecasting horizons. The median is taken over the full set of 38 countries, 20 times of initialisation, and 1000 trajectories. The remaining rows of Table 7.3 compare the performance of the agent-based models (the IIASA model and our calibrated models) to the AR(1) by showing the relative improvement in forecasting error (measured as a median across countries and initialisation times) compared to the AR(1). Overall, our model performs significantly better than both the AR(1) and the IIASA model. 208
7.4 Median relative improvements in RMSE of our model's growth rates of nominal GDP over the IIASA model (Poledna et al., 2023a). The values on the left of the '/' correspond to model calibration using NPE, and those on the right correspond to NRE. The model was run for 20 initialisation times between 2013-Q1 and 2017-Q4 in both cases. 209
7.5 Estimated Bayes factors by the time of initialisation and country, from 2013 to 2015. 212
7.6 Estimated Bayes factors by the time of initialisation and country, from 2016 to 2017. 213
7.7 Classification of Bayes factor estimates as reported in Jeffreys (1998), and counts of each Bayes factor class in Tables 7.5 and 7.6. 214

List of Algorithms

1 The s-sequence best-response dynamic on $g_{n,m}$ initiated at $\boldsymbol{a}^0$. . . . 24
2 The s-sequence best-response dynamic on $G_{n,m}$. 28
3 The clockwork random walk. 45
4 The coupled dynamic, which follows the random walk process until an environment is encountered by some player for the second time and becomes deterministic after that. 45
5 The coupled dynamic with sink $\boldsymbol{x}$, which initialises the directed edges to point at $\boldsymbol{x}$ for all players. 51
6 The process for simplifying $\widetilde{D} - \frac{1}{m^{n-1}} \cdot \widetilde{S} \cdot \prod_{i=1}^{n-1} \widetilde{N}_i$ to obtain A. . . . 92

1 Introduction

We study economic dynamics using two representative classes of models with interacting agents — one on the microeconomic level of agents repeatedly playing normal-form games and one on the macroeconomic level of entire countries. This chapter states our research questions, explains the motivation behind our methods (Section 1.1), and provides a high-level overview of our key findings on random games (Section 1.2) and agent-based modelling (Section 1.3). For random games, we examine the probability of convergence under various learning dynamics, where payoffs are drawn from a continuous distribution, and explore the likelihood of generating games with specific structural properties. On the macroeconomic scale, we describe the development and calibration of a macroeconomic model that simulates 38 countries. Finally, we summarise related publications (Section 1.4) and provide an outline of the thesis structure (Section 1.5).

Contents

1.1 Motivation **2**
1.2 Random Games **4**
1.2.1 The Frequency of Convergence 7
1.2.2 The Frequency of Games 8
1.3 Macroeconomic Agent-based Modelling **9**
1.3.1 The Framework 12
1.3.2 Calibration 13
1.3.3 Forecasting 14
1.4 Publications **14**
1.5 Structure **16**

1.1 Motivation

Economies experience significant fluctuations rather than changing gradually over time. Figure 1.1 shows quarterly GDP growth in Germany, France, and the United States between 1960 and 2024. Notable events such as the French protests of May 1968, the 2007-2008 financial crisis, and the COVID-19 pandemic are clearly reflected in the data. Other fluctuations are more difficult to explain; their *causes* motivates this thesis. In Part I, we employ a simple game-theoretic model to show that (economic) systems are often unstable and do not tend towards an equilibrium. However, this is a fundamental assumption in traditional macroeconomic models[1]. Therefore, in Part II, we turn to agent-based modelling as an alternative approach, which does not rely on this assumption.

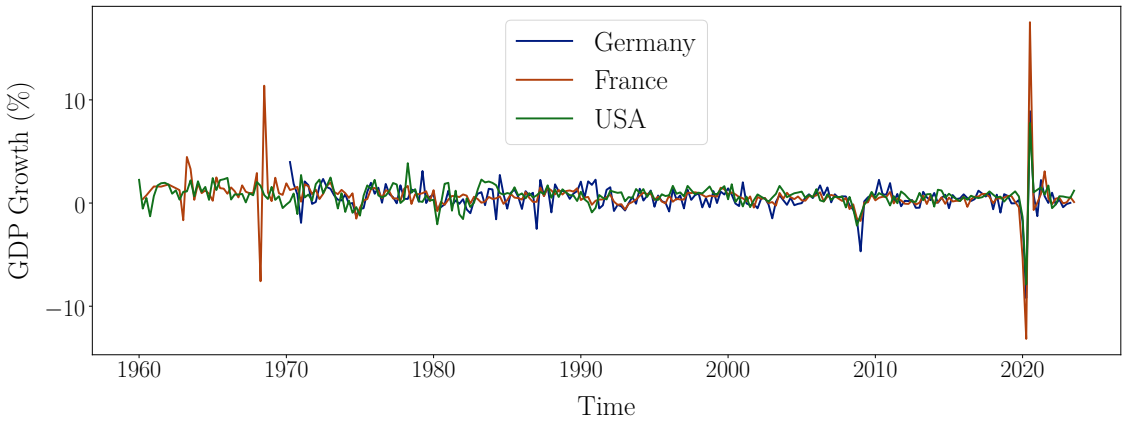


Figure 1.1: Quarterly GDP growth in constant prices for Germany, France, and the United States. The French protests of May 1968, the 2007-2008 financial crisis, and the COVID-19 pandemic are clearly visible.

Part I: Random Games as a Null Model There are two different perspectives on the causes of fluctuations that are not explained by obvious exogenous shocks (Pangallo, 2025). The first perspective argues that economies operate in stable, steady states, subject to constant external shocks (Kydland and Prescott, 1982; Clarida et al., 1999). The properties of these shocks and their transmission through the economy explain economic fluctuations. The second view argues that such steady states are inherently unstable (Minsky, 1977; Lorenz, 1993). This thesis adopts the latter view, supported by a simple game-theoretic model.

Comparing these views on a small scale uses a repeatedly played normal-form game as a null model. A group of players learn from their mistakes and develop

[1]The prevalent approach to economic modelling, *Dynamic Stochastic General Equilibrium* models (DSGEs), assume that the economy evolves towards a "steady state". See Section 5.2 for a review on DSGEs.

expectations about the future behaviour of the other players. For the family of studied learning rules (see Chapter 2), the players either converge to a Nash equilibrium or fail to converge and follow limit cycles or chaos instead.[2] For a simple learning dynamic called *best-response dynamics* where players choose their optimal action myopically, we derive analytical expressions for the likelihood of convergence in 2-player games and provide asymptotic bounds for multi-player games. We then use computer simulations to compare the convergence properties of best-response dynamics with more sophisticated learning algorithms. This analysis is extended to games with correlated payoffs. Through similar methods, we develop analytical results for the likelihood of generating specific classes of games at random. Our findings indicate that games with a large number of players or a large number of actions are less likely to converge to a Nash equilibrium under various dynamics, and that games on which those dynamics converge are rare. Turning to the macroeconomic scale, we therefore study large-scale, complex economies using agent-based models, that do not rely on the equilibrium assumption.

Part II: Large-Scale Macroeconomic Agent-based Modelling The economic upheavals of the past two decades, including the global financial crisis of 2008, the COVID-19 pandemic, and the energy shock following the Russian invasion of Ukraine, have highlighted the limitations of traditional macroeconomic models. *Dynamic Stochastic General Equilibrium* models (DSGEs), which remain widely used in academia and policy institutions, failed to predict or capture the complexities and immediate impacts of these events on global economies (Stiglitz, 2018). These shortcomings have intensified calls from within the economic community for advancements in modelling techniques. In response, two distinct approaches have emerged: enhancing existing DSGE models to incorporate agent heterogeneity through approaches like *Heterogeneous Agent New Keynesian* models (Kaplan et al., 2018) or pursuing a more radical departure from traditional methods (Axtell and Farmer, 2024).[3] In this latter category, *Agent-Based Models* (ABMs) emerged as a compelling alternative, offering a fundamentally different approach to understanding economic dynamics.

We develop a novel framework for constructing and calibrating large, data-driven macroeconomic agent-based models. Our framework enables the direct implementation and comparison of well-established models from the literature, facilitating the study of how different behavioural rules affect economic outcomes. Specifically, we

[2]We, therefore, follow the *descriptive interpretation* of game theory, assuming that the learning dynamics we employ (see Chapter 2) model human behaviour sufficiently well (Wooldridge, 2012).

[3]See Section 5.2.1 for an review of DSGE modelling and Section 5.2.2 for limitations of the approach.

construct and calibrate an agent-based model representing individuals, households, firms, banks, governments, and central banks across 38 countries[4]. The model is calibrated to public microdata and macrodata using state-of-the-art simulation-based inference techniques. We validate the model by comparing its time series forecasting accuracy against benchmark models.

Similarities In both cases — whether in repeatedly played normal-form games with certain learning dynamics or in macroeconomic agent-based models — agents continually adjust their behaviour in response to their environment. This generates agent-level dynamics, while aggregate-level behaviour emerges as a consequence of these interactions, rather than being explicitly defined (Wooldridge, 2009). In the literature, agent-based models have been extensively used in game-theoretic investigations, often incorporating features that are difficult to handle analytically, such as networks, or features that deviate from traditional specifications, such as agent memory (Farmer and Axtell, 2022). A significant body of research has explored the use of ABMs in the prisoner's dilemma game, for example, where Nowak and May (1992) fixed agents in a two-dimensional space and let them play the game with neighbours using parallel updating. This led to the emergence of dynamic, transient patterns that fostered high levels of cooperation. However, Huberman and Glance (1993) demonstrated that these results dissipate when synchronous updating is relaxed. Other studies have examined endogenous group formation (Aktipis, 2004, 2006), multiple equilibria (Axtell et al., 2000), and coalition formation (De Vany, 1996).

The remainder of this chapter summarises the key contributions and findings of the thesis. Regarding random games, we provide detailed results on the frequency of convergence and the likelihood of generating games with specific structures. Regarding agent-based modelling, we discuss our framework's design, calibration methods, and forecasting capabilities. Tables 1.1 and 1.2 offer a concise overview of our contributions, situating them within the broader context of the literature.

1.2 Random Games

The best-response dynamic is a widely used iterative game-playing process where, in each period, players choose pure strategies that are the best responses to the pure strategies selected by all other players in the previous round. While most existing research on best-response dynamics identifies sufficient conditions within

[4]In this thesis, we present results for the 38 OECD member countries. This choice is due to data availability and could easily be extended to additional countries.

a game's payoff structure to ensure convergence to a Nash equilibrium, we take a different approach. Instead of focusing on games with specific structural properties, our investigation centres on the role of the *playing sequence* — the order in which players update their (pure) strategies — in determining the convergence properties of the best-response dynamic.

We study three specific playing sequences: *clockwork* sequences, *random* sequences, and *simultaneous* updating. Under the clockwork playing sequence, players take turns to play one at a time according to a fixed cyclic order. Player 1 goes first, followed by player 2, up to the last player n; then the sequence returns to player 1. The behaviour of the best-response dynamic under this playing sequence has received relatively little attention in existing literature.[5] Under the random playing sequence, players play one at a time, and the next player to play is chosen uniformly randomly from among all players. The random sequence is the most well-studied in the existing literature.[6] Finally, we consider simultaneous updating by all players in each period.[7] It is important to note that many other playing sequences exist.[8]

To investigate the role of the playing sequence in determining the convergence properties of the best-response dynamic, we analyse the probability that the best-response dynamic converges to a pure Nash equilibrium in random n-player games, where every player chooses among m actions. Specifically, we generate a game by drawing all payoffs randomly (from atomless distributions to avoid payoff ties) and determine the probability that the best-response dynamic starting from a random initial pure strategy profile converges to a pure Nash equilibrium of the randomly generated game.

Table 1.1 summarises the existing literature on learning dynamics on random games and relates it to the corresponding section in this thesis. We differentiate between results for 2-player games and multi-player (n-player) games, as well as between games with independent payoffs, correlated payoffs, and games on networks. In terms of dynamics, we present results for clockwork, random, simultaneous, and periodic best-response dynamics, as well as for more complex learning rules, such as the Bush-Mosteller reinforcement learning algorithm (Bush and Mosteller,

[5] Boucher (2017) analyse the clockwork sequence best-response dynamic in potential games.

[6] The random sequence best-response dynamic has been analysed in anonymous games (Babichenko, 2013), near-potential games (Candogan et al., 2013), potential games (Christodoulou et al., 2012; Coucheney et al., 2014; Durand et al., 2019; Durand and Gaujal, 2016; Swenson et al., 2018), and games on a lattice (Blume and others, 1993). "Sink" equilibria are studied in Goemans et al. (2005); Mirrokni and Skopalik (2009).

[7] This case is studied in Quint et al. (1997) for 2-player games and in Kash et al. (2011) for anonymous games.

[8] For example, Feldman and Tamir (2012) study the case in which the sequence of play depends on current payoffs.

1953), fictitious play (Brown, 1951; Robinson, 1951), and replicator dynamics (Maynard Smith, 1982).[9] Part I of this thesis provides additional results and addresses gaps in the existing literature, particularly for games with correlated payoffs and network-based games, which we study using large-scale numerical simulations.

	Clockwork	Random	Sim.	Per.	BM	FP	RD
2-Pl./Ind.	BRS19 Section 3.2.2	BRS19			BRS19	BRS19	PT92 EWA13 BRS19
2-Pl./Corr.	BRS19 COR24	BRS19 COR24 Section 3.3.2			BRS19	BRS19	BRS19 PT92 EWA18 Section 3.3.2
n-Pl./Ind.	BRS24 CG21 Section 3.2.1	BRS24 RG21 CON23			Section 3.3.2		EVO10
n-Pl./Corr.	Section 3.3.2				Section 3.3.2		
Network	Section 3.2.3 Section 3.3.1.3 Section 3.3.2				Section 3.3.2		

Table 1.1: An overview of our results related to the existing literature on learning dynamics on random games: BRS19 (Pangallo et al., 2019), PT92 (Opper and Diederich, 1992), EWA13 (Galla and Farmer, 2013), COR24 (Mimun et al., 2024), EWA18 (Sanders et al., 2018), BRS24 (Heinrich et al., 2023), CG21 (Wiese and Heinrich, 2022), RG21 (Amiet et al., 2021), CON23 (Johnston et al., 2023), EVO10 (Gokhale and Traulsen, 2010). We are highlighting clockwork, random, periodic (per.), and simultaneous (sim.) best-response dynamics, as well as Bush-Mosteller (BM) learning, fictitious play (FP), and replicator dynamics (RD).

We also study a related problem, which is the likelihood of drawing games with certain structural properties when payoffs are drawn randomly. Again, our key question is whether games on which learning dynamics tend to converge are more common than games on which they do not. We use a best-response n-partite digraph to develop an expression for the likelihood of games that are "perfectly predictable" since they have a unique pure Nash equilibrium. We show that these types of games become very rare as the number of players or the number of actions increases.

[9] Our focus on reinforcement learning, fictitious play, and replicator dynamics is driven by the fact that these learning rules have been used to model human behaviour in game-playing experiments (Arthur, 1991; Bush and Mosteller, 1953; Cheung and Friedman, 1997; Erev and Roth, 1998; Friedman, 1996; Sarin and Vahid, 2001; Van Huyck et al., 1995).

Research Questions of Part I To summarise, our research questions for Part I of the thesis are the following:

RQ1. How and under which conditions do simple learning rules converge or not converge?

- RQ1.1. What is the role of the playing sequence in the convergence behaviour of the best-response dynamic?
- RQ1.2. Can we predict the convergence behaviour of complex learning rules using simple best-response dynamics?
- RQ1.3. What is the likelihood of generating games with certain structural properties at random?

Section 1.2.1 summarises our theoretical and numerical results on the frequency of convergence for different learning dynamics in random games, corresponding to Chapter 2 and Chapter 3 (RQ1.1, RQ1.2). In a related context, Section 1.2.2 presents our theoretical results on the likelihood of generating games with specific structural properties at random (RQ1.3), which is discussed in detail in Chapter 4.

1.2.1 The Frequency of Convergence

Our novel theoretical contributions primarily relate to the convergence properties of the clockwork best-response dynamic in random games, where payoffs are drawn independently. Our main result shows that the probability of convergence to a pure Nash equilibrium under the clockwork best-response dynamic is of order $1/\sqrt{m^{n-1}}$, up to a polynomial factor. This finding has two key implications: first, as the number of players n and the number of actions m becomes large ($nm \to \infty$), the probability of convergence to a pure Nash equilibrium approaches zero. Second, since the asymptotic convergence probability depends primarily on the quantity m^{n-1}, the probability of convergence to a pure Nash equilibrium in n-player, m-action games is approximately the same as in 2-player, m^{n-1}-action games when both n and m are large. Our simulations confirm that this asymptotic relationship between n-player, m-action games and 2-player, m^{n-1}-action games holds reasonably well, even for smaller values of n and m.

In the 2-player case, we provide results on game duration and derive an exact expression for the probability that the best-response dynamic reaches a (best-response) cycle of a given length at a specific period. As a special case, we determine the exact probability that the clockwork best-response dynamic converges to a pure Nash equilibrium in 2-player, m-action games and argue that this probability is the same for random playing sequences in the 2-player case. Additionally, we show that

this probability is asymptotically $\sqrt{\pi/m}$ as m grows large. For intermediate cases between 2-player and n-player games, we provide analytical results on the frequency of convergence in network games, focusing on bipartite networks, stars, and rings.

Turning to results obtained from numerical simulations, we find that the clockwork playing sequence is of particular interest because its frequency of convergence to equilibrium closely tracks that of three widely studied learning dynamics: Bush-Mosteller learning, fictitious play, and replicator dynamics. This relationship holds even for games with correlated payoffs.[10]

1.2.2 The Frequency of Games

We then turn to a related question: *How likely is it to draw certain classes of games at random?* We mainly build on the existing literature on the likelihood of games with a certain attractor structure. Consider, for instance, the 2-player, 3-action games depicted in Figure 1.2. Both players can choose from actions 1, 2, and 3. The value on the left of each action profile represents the payoff to the first (row) player, while the value on the right represents the payoff to the second (column) player. Best responses that are part of Nash equilibria are highlighted in blue; if they form part of a cycle, they are shown in red, and those within the attractor of a cycle are marked in orange. The two leftmost games in Figure 1.2 exhibit three Nash equilibria and no cycles. The games in the middle contain a two-cycle (where each player alternates moves twice), and the games on the right display a three-cycle. Suppose we select the action profile $(1, 1)$ as the starting point (where both players initially choose action 1) in the third game from the left in Figure 1.2, and the column player moves first. Their best response to the row player choosing action 1 is to switch to action 2. In response, the row player then chooses action 2, the column player reverts to action 1, and the row player switches back to action 1, leading back to the original starting position.
Using a novel graph-theoretic approach, we obtain an exact formula for the probability of generating a multi-player random game with a unique pure Nash equilibrium. Similar to the frequency of convergence, the likelihood of encountering such perfectly predictable games approaches zero as the number of players or actions increases. Moreover, games with a larger number of pure Nash equilibria (PNEs) are less common than games with fewer PNEs.

Our findings on the frequency of convergence and the frequency of games support the argument that economic fluctuations are endogenous to the economy,

[10] See Berg and Weigt (1999); Galla and Farmer (2013); Goldberg et al. (1968); Stanford (1999); Rinott and Scarsini (2000); Sanders et al. (2018) for work on random games with payoff correlations.

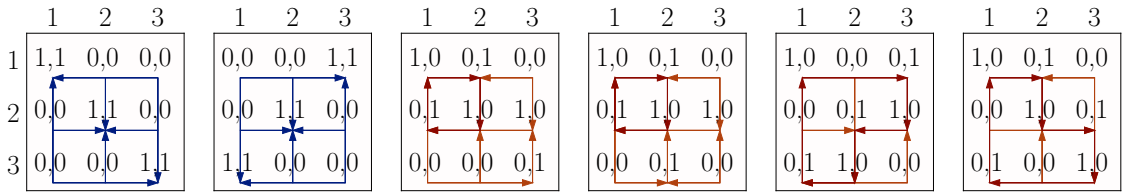


Figure 1.2: Possible best-response digraphs in 2-player, 3-action games. The value on the left of each action profile represents the payoff to the first (row) player, while the value on the right represents the payoff to the second (column) player. Best responses that are part of Nash equilibria are highlighted in blue; if they form part of a cycle, they are shown in red, and those within the attractor of a cycle are marked in orange. The two leftmost games in the figure exhibit three Nash equilibria and no cycles. The games in the middle contain a two-cycle (where each player alternates moves twice), and the games on the right display a three-cycle.

challenging the equilibrium assumption commonly used in standard economic models. In response to this critique, we construct and calibrate a complex agent-based model to better understand economic fluctuations in Part II of this thesis. Agent-based models have emerged as a successful alternative to DSGE models, particularly for capturing the effects of economic fluctuations.[11]

1.3 Macroeconomic Agent-based Modelling

Although macroeconomic ABMs have been around for over two decades (Dawid and Gatti, 2018), they have primarily succeeded in generating qualitative insights, particularly regarding policy (Fagiolo and Roventini, 2016). Typical macroeconomic ABMs are validated by demonstrating their ability to reproduce stylised facts or summary statistics from real-world data. To achieve this, researchers calibrate the key parameters of the model, while initialising the state variables of individual agents randomly. This is because standard macroeconomic ABMs are not designed to replicate a specific economy at a particular point in time, but rather to simulate a plausible, generic economy that reproduces key stylised facts.

In recent years, ABMs have become increasingly data-driven, aiming to faithfully represent specific economies at specific points in time (Pangallo, 2024). This shift requires not only calibrating aggregate parameters such as consumption propensities or capital depreciation rates, but also initialising the state variables of individual agents in a manner consistent with real-world data. The main advantage of this approach is that data-driven ABMs can track and forecast empirical time series, allowing for explicit validation as a faithful representation

[11] For example, Pichler et al. (2020) predicted the COVID-19-related loss in GDP for the British economy. See also Section 5.3.1.2.

of the economy. This increases confidence in the model's counterfactual analyses and policy recommendations. Most applications of data-driven ABMs have focused on specific sectors of the economy, such as housing (Geanakoplos et al., 2012), labour markets (del Rio-Chanona et al., 2021), and the economic impact of disasters (Hallegatte, 2008; Pichler et al., 2022). However, data-driven macroeconomic ABMs also exist (Papadopoulos, 2019; Kaszowska-Mojsa and Pipień, 2020).

A breakthrough in using data-driven ABMs for macroeconomics was achieved by Poledna et al. (2023a). This model, which we refer to as the *IIASA model*[12,13], and use as our benchmark, was the first to compare its out-of-sample forecasting performance to that of statistical and DSGE models. It demonstrated that ABMs could match, and potentially surpass, the performance of traditional models in analysing and predicting economic dynamics. The IIASA model quantitatively replicated the Austrian economy, initialising firm and household variables to be consistent with national accounts and macroeconomic statistics for a given year. All parameters in the model are directly derived from data or estimated from time series, making any parameter calibration unnecessary. However, this approach suggests that a more flexible model specification and calibration process could further improve its forecasting performance.

We propose a novel framework for constructing large data-driven macroeconomic agent-based models. Our model encompasses individuals, households, firms, banks, governments, and central banks across a given set of countries, along with markets for goods, labour, credit, and housing. Each agent type is initialised using microdata, incorporating heterogeneity. As in the IIASA model, we ensure full stock-flow consistency. Table 5.1 compares our model to several prominent large-scale macroeconomic models in terms of scope and detail. This comparison includes the EURACE project, developed respectively at the University of Bielefeld (UNIBI) and first described in Deissenberg et al. (2008), and at the University of Genoa (UNIGE), introduced in Cincotti et al. (2012a). Furthermore, we compare our model to the LAGOM framework (Mandel et al., 2010), the Keynes meets Schumpeter (K+S) family of ABMs developed at Sant'Anna Pisa (Dosi et al., 2010), the IIASA model, and its extension by the Bank of Canada, CANVAS (Hommes et al., 2022).[14]

Unlike the other two models evaluated by their empirical forecasting ability (IIASA[15] and CANVAS), our implementation is modular and allows for flexibility in terms of agents, behavioural rules, and datasets. This flexibility enables the initialisation of a global model with as many countries and sectors as data availability

[12]As it was developed at the International Institute for Applied Systems Analysis (IIASA).
[13]See Section 5.3.2.9 for a detailed description.
[14]We provide a detailed review of large macroeconomic models in Section 5.3.2.
[15]For a detailed comparison to the IIASA model, see Section 6.3.

Model	EURACE	LAGOM	K+S	IIASA	CANVAS	**Ours**
Section	5.3.2.2	5.3.2.3	5.3.2.4	5.3.2.9	5.3.2.10	Chapter 6
Year	2008	2010	2010	2023	2022	2024
Countries	EU-27	1	1	1	1	**global**
Sectors	2	3	2	62	18	**many**
Heterog. Firms	no	no	no	no	no	**yes**
Heterog. Banks	no	no	no	no	no	**yes**
Heterog. Pop.	no	no	no	no	yes	**yes**
Forecasting	no	no	no	yes	yes	**yes**
Code	C++	Java	C++	Matlab	Matlab	**Python**

Table 1.2: Comparison of our model to other large macroeconomic agent-based models in the literature regarding covered countries, sectors, heterogeneity, and ability to be calibrated to match time series.

permits, while also incorporating heterogeneity to the extent allowed by data and theoretical insights. The entire framework is written in Python, making it highly readable and adaptable.

Research Questions of Part II To summarize, our research questions for Part II of the thesis are the following:

RQ2. Can we build a modular, extensible, and user-friendly framework to create, calibrate, validate, run, and evaluate large data-driven macroeconomic agent-based models?

- RQ2.1. Can we use the framework to initialise an agent-based model to a set of countries at a given point in time?
- RQ2.2. Can we use modern inference techniques to calibrate the model?
- RQ2.3. Can we show that the use of realistic behavioural rules, microdata, and systematic calibration leads to better out-of-sample forecasting performance, and can we show that this improvement is statistically significant?

We describe our framework (RQ2), as well as model initialisation (RQ2.1) in high detail in Chapter 6, and provide a summary in Section 1.3.1 below. The methods for simulation-based inference (RQ2.2) are introduced in Section 5.4 and summarised below in Section 1.3.2. Finally, Chapter 7, summarised in Section 1.3.3 below, concludes with our key result — that the use of more realistic modelling rules, microdata, and careful calibration leads to statistically significant improvements in forecasting performance (RQ2.3).

1.3.1 The Framework

We propose a novel framework for building macroeconomic agent-based models.[16] The framework is designed to be fast, reliable, transparent, and flexible, allowing for the replacement of behavioural rules of individual agents or datasets as needed.[17] This unified setup for comparing different model architectures and rules makes development and benchmarking more efficient, as many models share similar components. Alongside the framework, we provide a dashboard for running simulations and analysing the simulation output.[18]

We use the framework to build a large agent-based model representing the 38 OECD countries. Each country is composed of firms, individuals, households, central government entities, banks, and a central bank.[19] Firms produce goods using labour, intermediate inputs, and capital, operating in 18 sectors corresponding to NACE level 1.[20] Individuals are classified as employed, unemployed, or not economically active. Employed individuals supply labour and receive wages, while unemployed or inactive individuals receive social benefits. Households consist of one or more individuals, who receive income from wages, social benefits, rent, and financial assets. They participate in the goods market, with consumption behaviour that varies according to household composition, wealth, debt, and income. Households also engage in the housing market, either owning, renting, or investing in properties, and may apply for mortgages to purchase real estate. Banks provide loans to firms (short-term and long-term) and households (consumption loans and mortgages) under realistic credit limits. Banks offer loans and set interest rates heterogeneously, adjusting rates based on market conditions and a reference rate set by the central bank. The central bank controls the policy rate. Countries not explicitly modelled are summarised as a "rest of the world" agent, interacting via imports and exports. There is an international market for goods and local markets for labour, credit, and housing. The goods and housing markets feature price competition, while the credit and labour markets exhibit competition for interest rates and wages, respectively. More details on the model and agent interactions can be found in Chapter 6. As we are attempting to achieve a high level of realism to reflect real-world dynamics, Table 6.20 provides an overview of what we are explicitly not modelling or what should be improved in the future.

[16] For existing packages, see Farmer and Axtell (2022).

[17] For a preliminary taxonomy of some behavioural rules, see Dawid and Gatti (2018).

[18] The dashboard can be accessed at `model.samuelwiese.com` with username *sam* and password *inet*.

[19] See Section 6.2 for a detailed summary, and, in particular, Figure 6.1 for the key interactions between agents.

[20] See Table 6.1 for the description of sectors.

1.3.2 Calibration

The model is calibrated by using data from various sources, including inter-country input-output tables from the OECD, World Bank data, Compustat microdata for firms and banks, Household Finance and Consumption Survey data for households and individuals, and policy rates provided by the Bank for International Settlements.[21] Most model parameters can be directly estimated from data; some are selected so that simulated data aligns with real-world observations.

1.3.2.1 Likelihood-based Methods

A natural approach for identifying a set of parameters $\boldsymbol{\theta}$ for which the simulated data closely matches the observed data $\boldsymbol{y}$ is to use likelihood-based methods. The likelihood function $p(\boldsymbol{y} \mid \boldsymbol{\theta})$ is combined with prior information about the parameters through Bayes' theorem, yielding the posterior distribution as:

$$p(\boldsymbol{\theta} \mid \boldsymbol{y}) = \frac{p(\boldsymbol{y} \mid \boldsymbol{\theta})}{p(\boldsymbol{y})} p(\boldsymbol{\theta}) \tag{1.1}$$

where $p(\boldsymbol{\theta} \mid \boldsymbol{y})$ is the posterior distribution of parameters $\boldsymbol{\theta}$ given the observed data $\boldsymbol{y}$, and $p(\boldsymbol{y} \mid \boldsymbol{\theta})$ is the likelihood. The term $p(\boldsymbol{\theta})$ represents the prior distribution of the parameters, while $p(\boldsymbol{y})$ is the marginal likelihood, which normalises the posterior.

The strength of Bayesian methods lies in their natural ability to quantify uncertainty via the posterior distribution $p(\boldsymbol{\theta} \mid \boldsymbol{y})$. They also allow for the incorporation of prior knowledge through $p(\boldsymbol{\theta})$ and provide intuitive probabilistic interpretations that are helpful for decision-makers, even those without technical expertise. However, in agent-based models, the likelihood $p(\boldsymbol{y} \mid \boldsymbol{\theta})$ is often intractable because the states of individual agents are typically unobserved. This renders traditional Bayesian inference impractical. To overcome this, we turn to *likelihood-free* approaches, commonly known as *simulation-based inference*, where the model is simulated under different parameter values to identify those parameters $\boldsymbol{\theta}$ most consistent with the observed data $\boldsymbol{y}$.

1.3.2.2 Simulation-Based Inference

Simulation-based inference is particularly difficult for agent-based models. These models are typically computationally expensive to simulate and generate high-dimensional, complex time series data (Wiese et al., 2024). Many existing methods require a large simulation budget to reduce Monte Carlo error in the approximate posterior, primarily due to the lack of information sharing between pointwise estimates of the posterior density (Dyer et al., 2024). Additionally, modelling

[21] See Table 6.2 for an overview of data sources.

the data distribution is particularly challenging for high-dimensional data, often necessitating simplifying assumptions.

Dyer et al. (2024) propose two promising solutions: *Neural Posterior Estimation* and *Neural Ratio Estimation*. These methods, along with other neural network-based techniques for approximate Bayesian inference in agent-based models (Dyer et al., 2023), have demonstrated significantly faster and more accurate performance compared to traditional methods commonly used in the agent-based modelling literature (Dyer et al., 2024). We introduce these methods briefly in Sections 5.4.2.1 and 5.4.2.2, with detailed training procedures provided in Section 7.2.2.

1.3.3 Forecasting

Even a calibrated agent-based model is not immune to the criticism that, due to the wide range of possible behavioural rules[22], a wide range of simulation output can be generated. Therefore, evaluating agent-based models by their forecasting performance instead of their ability to replicate stylised facts provides a stronger test for validity (Fagiolo and Roventini, 2016). Forecasting has been a key feature of conventional macroeconomic models, which makes them useful for policymakers. To assess forecasting performance, we compare our calibrated model against the IIASA model by measuring forecasting errors for economic aggregates over a 3-year time horizon.[23] Our calibrated model outperforms the benchmark IIASA model overall, yielding lower forecasting errors for all considered economic aggregates in 33 out of 38 countries.

We also estimate a Bayes factor[24] (Jeffreys, 1998) to further evaluate the forecasting abilities, and find that of all estimated Bayes factors, 18% favour the IIASA model over ours, while the remaining Bayes factors support our model to varying degrees. Finally, we assess the statistical significance of our model's performance improvement over the baseline IIASA model (Section 7.6) and find that the improvement is statistically significant at the 0.005 significance level under the hypothesis that the IIASA model performs better than ours.

1.4 Publications

Our work on random games has so far culminated in the following publications:

[22] Referred to as the "wilderness of bounded rationality" by Sims (1980).
[23] See Section 7.4 for forecasting results.
[24] See Section 7.5 for estimation details and results.

1. **The Frequency of Convergent Games under Best-Response Dynamics**. *Dynamic Games and Applications*, 12: 689–700, 2022. With Torsten Heinrich.

    - Torsten and I jointly developed the n-partite digraph methodology and co-authored the paper; I formulated and proved the theoretical results and conducted the simulations.
    - The results in this paper correspond to Chapter 4.

2. **Best-Response Dynamics, Playing Sequences, and Convergence to Equilibrium in Random Games**. *International Journal of Game Theory* 52: 703–735, 2023. With Torsten Heinrich, Yoojin Jang, Luca Mungo, Marco Pangallo, Alex Scott, Bassel Tarbush.

    - Torsten, Marco, and I developed the initial concept. Torsten, Yoojin, Marco, Alex, Bassel and I contributed to the theoretical results, and I performed the simulations. All authors contributed to the manuscript.
    - The literature review of the paper corresponds to parts of the literature review in Chapter 2. The results contained in the paper correspond to the theoretical results in Section 3.2 of Chapter 3.

3. **Best-Reponse and Learning Dynamics in Multi-Player Network Games**. *In preparation.* With Marco Pangallo.

    - I conducted the numerical simulations, and Marco and I are co-authoring the manuscript.
    - The literature review corresponds to parts of the literature review in Chapter 2, and the numerical results are detailed in Section 3.3.

Our work on macroeconomic agent-based modelling (Wiese et al., 2024) has just been submitted to the *Journal of Economic Dynamics and Control* as **Forecasting Macroeconomic Dynamics using a Calibrated Data-Driven Agent-based Model**, co-authored with Jagoda Kaszowska-Mojsa, Joel Dyer, José Moran, Marco Pangallo, François Lafond, John Muellbauer, Ani Calinescu, and Doyne Farmer. This corresponds to Part II of the thesis. All authors worked on modelling; I developed the initialisation routines, Joel and I developed the calibration procedure and the forecasting exercises, José and I wrote the Python code of the agent-based model, and I developed the framework. All authors contributed to the manuscript. I presented this work at the *1st Macro ABM Workshop* at the Institute for Advanced Systems Analytics (IIASA), at the *27th Annual Workshop on Economics with Heterogeneous Interacting Agents* (WEHIA), at the *Conference on Complex Systems*

2024 (CCS), and at the 10th meeting of the *German Network for New Economic Dynamics* (GENED) in 2024.

Additional publications, finalised during the first half of the DPhil and not included in this thesis, are:

1. **A Convex Surface with Fractal Curvature**. *Fractals*, 28(4): 2050059, 2020. With Iancu Dima, Rachel Popp, Robert S. Strichartz.

2. **Spectrum of the Laplacian on Regular Polyhedra**. *Communications on Pure and Applied Analysis*, 20(1): 193–214, 2021. With Evan Greif, Daniel Kaplan, Robert S. Strichartz.

3. **Spectrum of the Laplacian on Snowflake Domains and filled-in Julia sets**. *Experimental Mathematics*, 31(3): 1014–1025, 2022. With Robert S. Strichartz.

1.5 Structure

For the reader's convenience, each chapter begins with an abstract and an introduction, which summarises the key ideas, methods, results, and the structure of the chapter. Each chapter also includes a conclusion, which highlights potential weaknesses, future research, and places the chapter in the wider context of the thesis. Part I and Part II of this thesis share the same mathematical notation, respectively, which is introduced at the beginning of each chapter. In both parts, vector-valued variables are consistently denoted in boldface to distinguish them from scalar variables.

In Part I, we focus on individual agents repeatedly playing a normal-form game. We begin with the necessary formalism in Chapter 2, describing the learning dynamics under study, their parameters, and the tests for convergence. Chapter 3 examines the convergence properties of these dynamics in randomly generated games, with a particular focus on the best-response dynamic. We find that the order in which players select their optimal action is crucial for convergence. We present novel formulas for the convergence frequency in 2-player games and network-based games. We also provide extensive numerical results, showing how the convergence properties of simple best-response dynamics relate to more complex learning rules. Chapter 4 explores the frequency of certain classes of games. In particular, we introduce a novel method to derive an explicit formula for the frequency games with a unique pure Nash equilibrium.

In Part II, we shift to the macroeconomic scale by developing a framework for data-driven macroeconomic agent-based modelling. Chapter 5 reviews the history

and limitations of traditional macroeconomic models, compares 11 prominent agent-based models and frameworks, and discusses techniques for model calibration. Chapter 6 describes our novel data-driven agent-based model in detail, specifying agent behavioural rules, market interactions, and parameters. Chapter 7 outlines the calibration of a large-scale, data-driven agent-based model for 38 countries using simulation-based inference techniques and presents forecasting exercises, with a focus on demonstrating the statistical significance of the results.

Chapter 8 presents the conclusion, summarising our key findings on random games and macroeconomic agent-based modelling. We also address potential limitations of our approach and propose avenues for future research, including further model comparisons, modular improvements, and the ambitious goal of developing a global macroeconomic agent-based model.

Part I

Random Games

2

Dynamics on Random Games

Part I of the thesis investigates learning dynamics in randomly generated normal-form games. This chapter introduces the formal framework for generating such games (Section 2.2) by sampling payoffs from a continuous distribution. We establish best-response dynamics as a baseline, where players myopically choose the best response to their current environment. Four variations of best-response dynamics are explored based on the order in which players update their strategies: clockwork, random, simultaneous, and periodic (Section 2.3). These baseline dynamics are then compared to more advanced learning rules (Section 2.4). For each dynamic, we define its parameters and convergence criteria, both of which are key to deriving the analytical and numerical results presented in Chapter 3. Finally, Section 2.5 provides a conclusion.

Contents

2.1 Introduction 20
2.2 Formalisation 22
2.3 Best-Response Dynamics 24
2.3.1 Clockwork Best-Response Dynamics 25
2.3.2 Random Best-Response Dynamics 27
2.3.3 Simultaneous Best-Response Dynamics 29
2.3.4 Periodic Best-Response Dynamics 29
2.4 Other Learning Dynamics 30
2.4.1 Reinforcement Learning 30
2.4.2 Fictitious Play 32
2.4.3 Replicator Dynamics 34
2.5 Conclusion 35

2.1 Introduction

The best-response dynamic is a ubiquitous iterative game-playing process in which, in each period, players myopically select pure strategies that are the best response to the pure strategies that all other players last chose. Most of the existing work on best-response dynamics and learning rules, such as fictitious play, establishes sufficient conditions on a game's payoff structure to guarantee convergence to a Nash equilibrium.[1] We, however, investigate the convergence properties of the best-response dynamic, but rather than focusing on games with a particular structure, our focus is on the role of the *playing sequence* — the order in which players update their (pure) strategies.

We study four specific playing sequences: *clockwork* sequences, *random* sequences, *periodic* sequences, and *simultaneous* updating. Under the clockwork playing sequence, players take turns to play one at a time according to a fixed cyclic order. Player 1 goes first, followed by player 2, up to the last player n; then the sequence returns to player 1. The behaviour of the best-response dynamic under this playing sequence has received relatively little attention in existing literature.[2] Under the random playing sequence, players take turns to play one at a time, and the next player to play is chosen uniformly randomly from among all players. This sequence is the most well-studied in the literature.[3] Periodic sequencing interpolates between clockwork and random sequencing, and under simultaneous updating, all players update their actions at the same time in each period.[4]

We additionally introduce three more complex learning rules: the *Bush-Mosteller reinforcement learning* algorithm (Bush and Mosteller, 1953), *fictitious play* (Brown, 1951; Robinson, 1951), and *replicator dynamics* (Smith et al., 1992). These rules are of particular interest due to their widespread recognition, their encapsulation of different behavioural assumptions regarding learning in games, and their calibration to human gameplay in experimental settings (Bush and Mosteller, 1953; Arthur, 1991; Erev and Roth, 1998; Sarin and Vahid, 2001; Van Huyck et al., 1995; Friedman, 1996; Cheung and Friedman, 1997).

[1] For example, previous work has shown that the best-response dynamic converges to a Nash equilibrium in weakly acyclic games (Fabrikant et al., 2013), potential games (Monderer and Shapley, 1996), and quasi-acyclic games (Friedman and Mezzetti, 2001; Takahashi and Yamamori, 2002).

[2] Boucher (2017) analyses the clockwork sequence best-response dynamic in potential games.

[3] The random sequence best-response dynamic has been analysed in anonymous games (Babichenko, 2013), near-potential games (Candogan et al., 2013), potential games (Christodoulou et al., 2012; Coucheney et al., 2014; Durand et al., 2019; Durand and Gaujal, 2016; Swenson et al., 2018), and games on a lattice (Blume and others, 1993). "Sink" equilibria are studied in (Goemans et al., 2005; Mirrokni and Skopalik, 2009).

[4] This case is studied in Quint et al. (1997) for 2-player games and in Kash et al. (2011) for anonymous games.

Bush-Mosteller reinforcement learning assumes that players repeat actions that have yielded high payoffs in the past. Variants of reinforcement learning models have been calibrated to human gameplay in experiments, such as those in Arthur (1991), Erev and Roth (1998), and Sarin and Vahid (2001). Under Bush-Mosteller learning, each player selects their action by sampling according to a mixed strategy vector, which evolves through reinforcement learning in each period. We assess the convergence of these vectors by checking whether the difference from one period to the next falls below a threshold and becomes indistinguishable from sampling noise. By tracking the evolution of mixed strategy vectors instead of just the actions played, we can assess whether the dynamic converges to mixed Nash equilibria.

Fictitious play requires more sophistication, as it assumes that players construct a mental model of their opponents. Each player believes that the empirical distribution of their opponents' past actions is their mixed strategy and responds optimally to this belief. Classical experiments with human players have shown that fictitious play is a helpful model of learning in various settings. Cheung and Friedman (1997) studied coordination, dominance-solvable, and cyclic 2-player, 2-action games using this model. To assess convergence, we follow Fudenberg and Levine (1998) in tracking the convergence of belief vectors rather than the actions themselves, allowing us to include convergence to mixed equilibria in our analysis.

The replicator equation, commonly used in ecology and population biology, has also been viewed as a learning algorithm, where each population trait corresponds to an action.[5] In our setup, each player samples their actions based on a mixed strategy vector, which evolves according to the replicator equation. Van Huyck et al. (1995) studied two tacit bargaining games and found that players' behaviour was consistent with replicator dynamics. Similar conclusions were drawn by Friedman (1996) in a more extensive set of games. As with Bush-Mosteller learning, we track the convergence of mixed strategy vectors. However, note that the multi-population replicator dynamic does not converge to mixed equilibria in random games. If the dynamic converges to an equilibrium, each player's mixed strategy vector will assign all the mass to a single action.

As mentioned, previous research has analysed the convergence properties of best-response dynamics in games with specific payoff structures (e.g., potential games or aggregate games). In these cases, convergence depends on the game's details, the playing sequence, and the initial conditions. Rather than imposing such restrictions, we assume that the inputs to the game are drawn randomly, allowing for a broader analysis of the dynamics. We start with the formalism required to generate such random games.

[5]We use a multi-population version of the replicator dynamic because our randomly drawn payoff matrices are generally not symmetric.

2.2 Formalisation

This section introduces the key notation for studying learning dynamics on randomly generated games. We formalise the core concepts, including games and best-response structures, and describe the process for generating games at random.

Games A *game* with $n \geq 2$ players and $m \geq 2$ actions is a tuple

$$g_{n,m} := \Big([n], [m], \{u_i\}_{i\in[n]}\Big) \tag{2.1}$$

where $[n] := \{1, ..., n\}$ is the set of players, and each player $i \in [n]$ has a set of actions $[m] := \{1, ..., m\}$ and a payoff function $u_i : [m]^n \to \mathbb{R}$. Games with more than two players ($n > 2$) are referred to as *multi-player* games. We will denote the set of all games with n players and m actions by $\mathcal{G}_{n,m}$.

An *action profile* is a vector of actions $\boldsymbol{a} = (a_1, \ldots, a_n)$ belonging to the set $[m]^n$, listing the action taken by each player. An *environment* for player i is a vector $\boldsymbol{a}_{-i}$ belonging to the set $[m]^{n-1}$ that lists the actions taken by each player except i. A *best-response correspondence* b_i for player i is a mapping from the set of environments for player i to the set of all non-empty subsets of i's actions, defined by

$$b_i(\boldsymbol{a}_{-i}) := \arg\max_{a_i\in[m]} u_i(a_i, \boldsymbol{a}_{-i}). \tag{2.2}$$

A game is *non-degenerate* if for each player i and environment $\boldsymbol{a}_{-i}$, the best-response action is unique. Games in which there are no ties in payoffs are non-degenerate.[6] Throughout this thesis, we focus exclusively on non-degenerate games, meaning that "game" will always refer to a "non-degenerate game." Since best-responses are unique in non-degenerate games, we write $a_i = b_i(\boldsymbol{a}_{-i})$ whenever $a_i \in b_i(\boldsymbol{a}_{-i})$.

An action profile $\boldsymbol{a} \in [m]^n$ is a *pure Nash equilibrium* (PNE) if for all $i \in [n]$ and all $a_i \in [m]$,

$$u_i(\boldsymbol{a}) \geq u_i(a_i, \boldsymbol{a}_{-i}). \tag{2.3}$$

Equivalently, $\boldsymbol{a} \in [m]^n$ is a PNE if each player $i \in [n]$ is playing their best-response action, i.e. $a_i = b_i(\boldsymbol{a}_{-i})$. Denote the set of PNEs of the game $g_{n,m}$ by $\text{PNE}(g_{n,m})$, and let $\#\text{PNE}(g_{n,m})$ denote the number of pure Nash equilibria in the game.

We may also impose an undirected network[7] of n nodes on a game $g_{n,m}$, which determines which players' actions impact which player's payoffs. Unless specified otherwise, we assume the network is fully connected.

[6] There are no ties in payoffs if for all players i, all environments $\boldsymbol{a}_{-i}$, and for all $a_i \neq a_i'$, $u_i(a_i, \boldsymbol{a}_{-i}) \neq u_i(a_i', \boldsymbol{a}_{-i})$.

[7] For an introduction to the concept of games on networks, see Jackson and Zenou (2015).

Best-Response Structure The best-response structure of a non-degenerate game $g_{n,m}$ can be represented by a *best-response digraph* $\mathcal{D}(g_{n,m})$, whose vertex set is the set of action profiles $[m]^n$. The edges are constructed as follows: for each $i \in [n]$ and each pair of distinct vertices $\boldsymbol{a} = (a_i, \boldsymbol{a}_{-i})$ and $\boldsymbol{a}' = (a'_i, \boldsymbol{a}_{-i})$, place a directed edge from $\boldsymbol{a}$ to $\boldsymbol{a}'$ if and only if a'_i is player i's best-response to environment $\boldsymbol{a}_{-i}$, i.e. $a'_i = b_i(\boldsymbol{a}_{-i})$. Edges exist only between action profiles that differ in exactly one coordinate. A profile $\boldsymbol{a}$ is a PNE of $g_{n,m}$ if and only if it is a sink of the best-response digraph $\mathcal{D}(g_{n,m})$.

Figure 2.1 illustrates a 3-player 2-action game (on the left) and its associated best-response digraph (on the right). Player 1 selects rows (along the depth), player 2 selects columns (along the width), and player 3 selects levels (along height). In the left-hand panel, the payoffs of players 1, 2, and 3 are listed in that order. The vertices of the best-response digraph are the action profiles. The unique PNE at the profile $(1, 2, 1)$ is underlined and is a sink of the digraph.

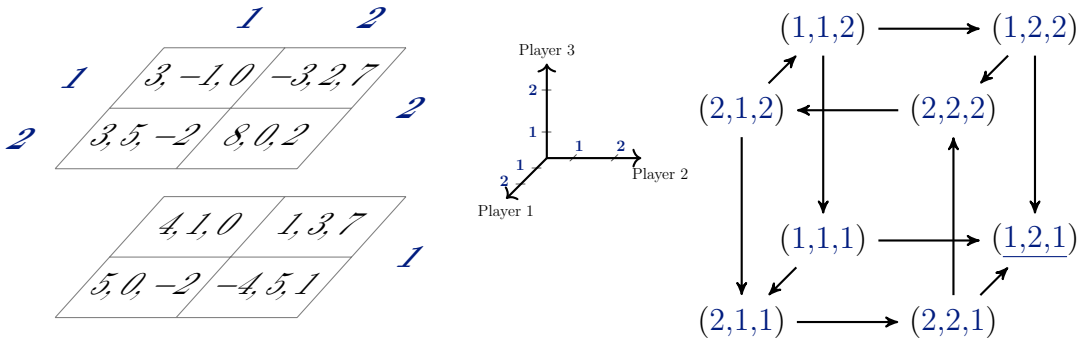


Figure 2.1: Illustration of a 3-player, 2-action non-degenerate game (left) and its associated best-response digraph (right). The axes shown in the centre states the coordinate system.

Creating Random Games We generate random games by drawing all payoffs at random: for each $\boldsymbol{a} \in [m]^n$ and $i \in [n]$, the payoff $U_i(\boldsymbol{a})$ is a random number drawn from an atomless distribution $\mathbb{P}$. These draws are independent across all $i \in [n]$ and $\boldsymbol{a} \in [m]^n$. The distribution $\mathbb{P}$ ensures that any ties in payoffs have zero measure, so any resulting game is almost surely non-degenerate. A random n-player m-action game drawn according to this process is denoted by $G_{n,m} := ([n], [m], \{U_i\}_{i\in[n]})$.

At initialisation, for each pure strategy profile, we draw a vector of payoffs for each player at random from a multivariate normal distribution with mean zero, unit variance, and a covariance matrix with 1s on the diagonal and $\frac{\Gamma}{n-1}$ on the off-diagonal entries. Here, $\Gamma \in [-1, n-1]$ parametrises the degree to which payoffs are correlated. If $\Gamma = 0$, then the payoffs are chosen independently. At one extreme, $\Gamma = n-1$ means that all players receive the same payoff at each action profile, while

at the other extreme, $\Gamma = -1$ corresponds to a zero-sum game. More generally, when $\Gamma > 0$, the game is considered "cooperative," and when $\Gamma < 0$, the game is "competitive" (Rinott and Scarsini, 2000).

Finally, we randomly draw the initial profile $\boldsymbol{A}^0$ from all profiles. Since the game is randomly generated, the choice of initial condition does not affect the results, making it irrelevant. Randomly selecting the initial profile allows us to drop the dependence on the initial profile in our description of the best-response dynamic.

2.3 Best-Response Dynamics

We now consider games played over time, with each player myopically best responding to their current environment. A *playing sequence* function $s : \mathbb{N} \to [n]$ determines whose turn it is to play at each period $t \in \mathbb{N}$, where $\mathbb{N}$ denotes the set of positive integers $\{1, 2, ...\}$. A *path* $\langle \vec{\mathbf{a}}, s \rangle$ is an infinite sequence of action profiles $\vec{\mathbf{a}} = (\boldsymbol{a}^0, \boldsymbol{a}^1, ...)$ and an associated playing sequence function $s : \mathbb{N} \to [n]$ satisfying the constraint that only one player changes her action at a time, $\boldsymbol{a}^t_{-s(t)} = \boldsymbol{a}^{t-1}_{-s(t)}$ for each $t \in \mathbb{N}$. In other words, only the action of the player chosen by the sequence at time t changes, while other players retain their previous actions.

Note that this setup rules out simultaneous updating from our theoretical analysis because we allow only one player to play in any given period. A more general framework could allow subsets of players to update their actions simultaneously in each period — this would extend the playing sequence to non-empty subsets of $[n]$. However, we avoid this generality here. For our theoretical results, we will focus specifically on clockwork and random playing sequences.

The best-response dynamic with playing sequence $s : \mathbb{N} \to [n]$ on a game $g_{n,m}$ initiated at the action profile $\boldsymbol{a}^0$ generates a path $\langle \vec{\mathbf{a}}, s \rangle$ according to Algorithm 1. Namely, set the initial action profile to $\boldsymbol{a}^0$ and, in each period $t \in \mathbb{N}$, player $s(t)$ myopically plays the best-response to their current environment $\boldsymbol{a}^{t-1}_{-s(t)}$.

Algorithm 1 The s-sequence best-response dynamic on $g_{n,m}$ initiated at $\boldsymbol{a}^0$.

`1. For` $t \in \mathbb{N}$`:`

- `(a) Set` $i = s(t)$
- `(b) Set` $\boldsymbol{a}^t_{-i} = \boldsymbol{a}^{t-1}_{-i}$
- `(c) Set` $a^t_i = b_i(\boldsymbol{a}^{t-1}_{-i})$ `where` $b_i(\boldsymbol{a}^{t-1}_{-i}) := \arg\max_{x_i \in [m]} u_i(x_i, \boldsymbol{a}^{t-1}_{-i})$

Algorithm 1 generates a path by travelling along the edges of the best-response digraph $\mathcal{D}(g_{n,m})$ in direction $s(t)$ at step t, starting from the initial profile $\boldsymbol{a}^0$. More precisely, the infinite sequence of actions $\vec{\mathbf{a}}$ is determined as follows: if player $s(t)$ is

already best responding, then $\boldsymbol{a}^{t-1}$ does not point to any vertex $(a'_{s(t)}, \boldsymbol{a}^{t-1}_{-s(t)}) \neq \boldsymbol{a}^{t-1}$, and the following profile in the sequence is $\boldsymbol{a}^{t-1}$ itself, i.e. $\boldsymbol{a}^t = \boldsymbol{a}^{t-1}$; otherwise, if player $s(t)$ is not already playing their best response, then the sequence travels to the vertex where the player plays their best-response action, i.e. set $\boldsymbol{a}^t = (a'_{s(t)}, \boldsymbol{a}^{t-1}_{-s(t)})$ where $(a'_{s(t)}, \boldsymbol{a}^{t-1}_{-s(t)}) \neq \boldsymbol{a}^{t-1}$ is the unique vertex that $\boldsymbol{a}^{t-1}$ points to.

For any path $\langle \vec{\mathbf{a}}, s \rangle$ and any set of action profiles $\mathcal{A} \subseteq [m]^n$, define $\tau_{\langle \vec{\mathbf{a}}, s \rangle}(\mathcal{A})$ as the first period $t \geq 1$ in which some element of the sequence $\vec{\mathbf{a}}$ is in the set $\mathcal{A}$:

$$\tau_{\langle \vec{\mathbf{a}}, s \rangle}(\mathcal{A}) := \inf\{t \in \mathbb{N} : \boldsymbol{a}^t \in \mathcal{A}\}, \tag{2.4}$$

where inf is the infimum operator, and we use the convention that $\inf \emptyset = \infty$ (i.e. we take $\tau_{\langle \vec{\mathbf{a}}, s \rangle}(\mathcal{A})$ to be infinite if no element of the sequence $\vec{\mathbf{a}}$ is in $\mathcal{A}$). The path $\langle \vec{\mathbf{a}}, s \rangle$ *reaches* the set $\mathcal{A}$ (in period t) if $t = \tau_{\langle \vec{\mathbf{a}}, s \rangle}(\mathcal{A})$ and t is finite.[8]

Definition 1. The s-sequence best-response dynamic on game $g_{n,m}$ initiated at $\boldsymbol{a}^0$ *converges* to a PNE if the path $\langle \vec{\mathbf{a}}, s \rangle$ generated according to Algorithm 1 reaches $\mathrm{PNE}(g_{n,m})$.

There are infinitely many possible playing sequences. However, we will focus particularly on *clockwork* (Section 2.3.1), *random* (Section 2.3.2), *simultaneous* (Section 2.3.3), and *periodic* (Section 2.3.4) best-response dynamics, which we will discuss in detail below.

2.3.1 Clockwork Best-Response Dynamics

The *clockwork* playing sequence is defined by

$$s(t) = s_{\mathsf{c}}(t) := 1 + (t-1) \bmod n. \tag{2.5}$$

In other words, player 1 plays in period 1, followed by player 2, then 3, and so on until player n, at which point the sequence returns to player 1.

Definition 1 applies to all playing sequences, but when the sequence is clockwork, we can characterise convergence (and non-convergence) more simply regarding path properties. We refer to one complete rotation through the clockwork sequence as a *round* of play. For example, if a round starts at player i, then each player plays once, and the round is complete when it is again i's turn to play.

For any $k \in \mathbb{N}$, define

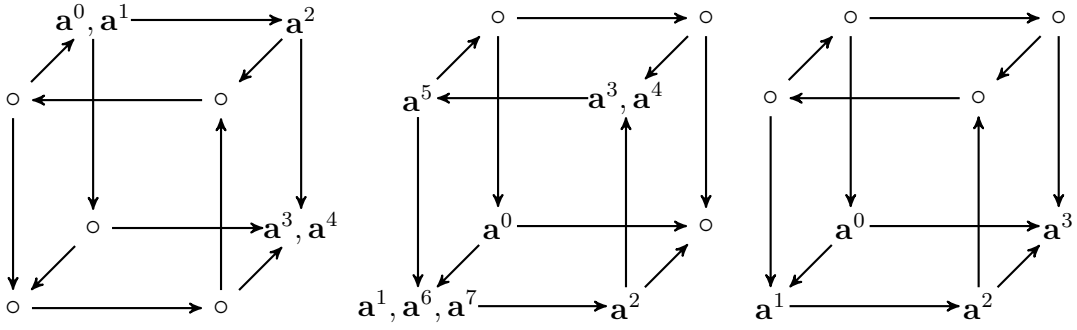
$$T_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle}(k) := \inf \Big\{ t \in \mathbb{N} : \boldsymbol{a}^t = \boldsymbol{a}^{t+nk} \text{ and } \boldsymbol{a}^t \neq \boldsymbol{a}^{t+nk'} \text{ for all } k' \in \mathbb{N} \text{ such that } k' < k \Big\} \tag{2.6}$$

[8] We also say that the path reaches the set $\mathcal{A}$ *by* period t if it reaches $\mathcal{A}$ in period τ with $\tau \leq t$, and the path reaches $\mathcal{A}$ *before* (*after*) period t if it reaches $\mathcal{A}$ in period $\tau < t$ ($\tau > t$).

to be the first period in which an action profile is repeated k rounds later (and not in any earlier round). If $T_{\langle \vec{\mathbf{a}}, s_c \rangle}(k)$ is finite, then the path $\langle \vec{\mathbf{a}}, s_c \rangle$ has the property that from period $T_{\langle \vec{\mathbf{a}}, s_c \rangle}(k)$ onwards, the sequence of nk possibly non-distinct action profiles $\boldsymbol{a}^t, ..., \boldsymbol{a}^{t+nk-1}$ repeats itself forever. We therefore say that the path $\langle \vec{\mathbf{a}}, s_c \rangle$ reaches a *nk-cycle* in period $T_{\langle \vec{\mathbf{a}}, s_c \rangle}(k)$, or equivalently, that the clockwork sequence best-response dynamic converges to a nk-cycle.

Notice that if the action profile is $\boldsymbol{a}^t$ in some period t and no one deviates from this profile in a single round (i.e. $\boldsymbol{a}^t = \boldsymbol{a}^{t+n}$), then $\boldsymbol{a}^t$ must be a PNE. Therefore, if the path $\langle \vec{\mathbf{a}}, s_c \rangle$ reaches a nk-cycle in period $T_{\langle \vec{\mathbf{a}}, s_c \rangle}(k)$ and $k = 1$, then we say that the clockwork sequence best-response dynamic reaches a PNE in that period. However, if the path $\langle \vec{\mathbf{a}}, s_c \rangle$ reaches a nk-cycle in period $T_{\langle \vec{\mathbf{a}}, s_c \rangle}(k)$ and $k > 1$, then we say that the clockwork sequence best-response dynamic converges to a *best-response cycle* (of length nk).

Because the number of action profiles is finite, $T_{\langle \vec{\mathbf{a}}, s_c \rangle}(k)$ must be finite for some k, so the clockwork best-response dynamic must always converge either to a PNE or to a best-response cycle. In dynamical systems language, nk-cycles (including PNE) are *attractors* of the clockwork best-response dynamic. Thus, the earliest period $T_{\langle \vec{\mathbf{a}}, s_c \rangle}$ is the period in which the path reaches a PNE or a best-response cycle.



(a) Sequence: clockwork $\mathbf{a}^0 = (1, 1, 2)$

(b) Sequence: clockwork $\mathbf{a}^0 = (1, 1, 1)$

(c) Sequence: 1-2-1-$\cdots$ $\mathbf{a}^0 = (1, 1, 2)$

Figure 2.2: The digraphs in panels (a)-(c) above are all identical and correspond to the best-response digraph of the game shown in Figure 2.1. In the panels, we show the first few elements of paths generated according to the best-response dynamic for different initial profiles and playing sequences.

Example (Best-Response Dynamics and Convergence)**.** The digraphs in panels (a)-(c) of Figure 2.2 are all identical and correspond to the best-response digraph of the game shown in Figure 2.1. Each vertex of a best-response digraph is a point in $[m]^n$ but, unlike the right-hand panel of Figure 2.1, we no longer show the explicit coordinate of each vertex in our illustrations to avoid clutter. In panels (a)-(c)

of Figure 2.2, we show the first few elements of paths generated according to the best-response dynamic for different initial profiles and playing sequences.

In panel (a), the initial profile is set to $\boldsymbol{a}^0 = (1, 1, 2)$ and the playing sequence is clockwork. The first few elements of the infinite sequence $\vec{\mathbf{a}}$ are shown in the figure. The path stays at $(1, 1, 2)$ in period one because player 1 does not change their action. The path then moves to $\boldsymbol{a}^2 = (1, 2, 2)$ in period two because player 2 plays action 2. In period 3, player 3 plays action 1, which takes the path to $\boldsymbol{a}^3 = (1, 2, 1)$. Once the path reaches this unique PNE profile, it will remain there forever. Furthermore, $T_{\langle \vec{\mathbf{a}}, s_{\mathtt{c}} \rangle}(1) = 3$.

In panel (b), the initial profile is set to $\boldsymbol{a}^0 = (1, 1, 1)$ and the playing sequence is clockwork. This time, the path moves to the bottom left corner on the front face of the cube in period 1. The path then cycles forever among the four profiles on the front face of the cube. Since period $t = 1$ is the first period in which $\boldsymbol{a}^t \neq \boldsymbol{a}^{t+3}$ but $\boldsymbol{a}^t = \boldsymbol{a}^{t+3k}$ for some integer k (namely $k = 2$) we have that $T_{\langle \vec{\mathbf{a}}, s_{\mathtt{c}} \rangle}(2) = 1$. In other words, the path reaches a best-response cycle in period 1. The path reaches a 6-cycle in period 1: once reached, the action profile sequence $\boldsymbol{a}^1, ..., \boldsymbol{a}^6$ is repeated forever. Not all action profiles in the sequence $\boldsymbol{a}^1, ..., \boldsymbol{a}^6$ are distinct.

In panel (c), the initial profile is once again set to $\boldsymbol{a}^0 = (1, 1, 1)$, but the playing sequence is 1-2-1-$\cdots$. This time, the path reaches the PNE in period 3. The playing sequence from period four onwards is irrelevant: once at the PNE, the path will remain forever, regardless of the playing sequence. The examples in panels (b) and (c) illustrate how changes to the playing sequence and the initial profile can affect convergence to a PNE.

2.3.2 Random Best-Response Dynamics

Another playing sequence that deserves highlighting is the *random* playing sequence $s_{\mathtt{r}}$, which is determined as follows: for each $t \in \mathbb{N}$, draw $s_{\mathtt{r}}(t)$ uniformly at random from $[n]$. So, in each period, the player playing is drawn uniformly at random from among all players. In what follows, we will take a playing sequence s to be an element of $\{s_{\mathtt{c}}, s_{\mathtt{r}}\}$ because we will compare our results concerning the clockwork sequence against existing analytical results concerning the random playing sequence (Amiet et al., 2021). In 2-player games, random best-response dynamics are equivalent to clockwork best-response dynamics.

The best-response dynamic on a random game, playing sequence, and initial condition is described by Algorithm 2. We randomly draw the game and initial condition and run Algorithm 1. This process induces a distribution over paths and PNE sets. Our definitions of convergence, given in Section 2.3, apply here. For example, we say that the s-sequence best-response dynamic on game $G_{n,m}$ (and

initial condition $\boldsymbol{A}^0$) converges to a PNE if the path $\langle \vec{\mathbf{A}}, s \rangle$ generated according to Algorithm 2 reaches PNE($G_{n,m}$).

Algorithm 2 The s-sequence best-response dynamic on $G_{n,m}$.

1. `For all` $i \in [n]$ `and` $\boldsymbol{a} \in [m]^n$ `draw` $U_i(\boldsymbol{a})$ `at random according to` $\mathbb{P}$
2. `Draw` $\boldsymbol{A}^0$ `uniformly at random from` $[m]^n$
3. `For` $t \in \mathbb{N}$:
   (a) `Set` $i = s(t)$
   (b) `Set` $\boldsymbol{A}^t_{-i} = \boldsymbol{A}^{t-1}_{-i}$
   (c) `Set` $A^t_i = B_i(\boldsymbol{A}^{t-1}_{-i})$ `where` $B_i(\boldsymbol{A}^{t-1}_{-i}) := \arg\max_{x_i \in [m]} U_i(x_i, \boldsymbol{A}^{t-1}_{-i})$

Step (1) of Algorithm 2 effectively creates a best-response digraph $\mathcal{D}(G_{n,m})$ on the vertices $[m]^n$ according to the following stochastic process: for each $i \in [n]$ and environment $\boldsymbol{a}_{-i}$, select an action a'_i uniformly at random from $[m]$ and then, for each $a_i \neq a'_i$, create a directed edge from $(a_i, \boldsymbol{a}_{-i})$ to $(a'_i, \boldsymbol{a}_{-i})$. Because the payoffs are drawn randomly, the probability of ties is zero due to the atomless distribution, and we can model the probability that a player plays a best-response action. This follows from how the payoffs are drawn: there is a zero probability of ties because $\mathbb{P}$ is atomless, and for each $i \in [n]$, the probability that action $a_i \in [m]$ is a best-response to the environment $\boldsymbol{a}_{-i}$ is given by

$$\Pr\left[U_i(a_i, \boldsymbol{a}_{-i}) \geq \max_{x_i \in [m]} U_i(x_i, \boldsymbol{a}_{-i})\right] = \frac{1}{m}. \tag{2.7}$$

Step (2) selects an initial profile, and step (3) essentially traces a path by travelling along the edges of the best-response digraph in the direction $s(t)$ at step t, starting from the initial profile.

Drawing payoffs independently randomly (from atomless distributions) induces a uniform distribution over the best-response digraphs. Thus, the probability of convergence to a pure Nash equilibrium reduces to calculating the probability that the best-response path, initiated at a random vertex, reaches a sink of the randomly drawn digraph. The main theoretical challenge we face when analysing the best-response dynamic is that it exhibits path dependence: if a player encounters an environment they have seen before, they must play the same action they played initially. We address this by using a coupling argument that links the best-response dynamic to a (memoryless) random walk through the digraph, which is easier to analyse.

In large 2-action games, the random sequence best-response dynamic converges with high probability to a PNE if there is one:

Proposition 1 (Amiet et al. (2021))**.**

$$\lim_{n\to\infty} \Pr\left[s_r\text{-BRD on } G_{n,2} \text{ converges to a PNE} \,|\, \#\text{PNE}(G_{n,2}) \geq 1\right] = 1. \tag{2.8}$$

A generalisation of Proposition 1 to m-action games is non-trivial, and we are unaware of existing analytical results for $m > 2$.[9] However, we conjecture that the random sequence best-response dynamic converges to a PNE with high probability if there is one as $mn \to \infty$. Consistent with this conjecture, in the simulations of Section 3.3, we show that the random sequence best-response dynamic does converge to a PNE with probability close to $1 - \exp(-1)$ when m or n get large, provided that $n > 2$.

2.3.3 Simultaneous Best-Response Dynamics

Under *simultaneous* best-response dynamics, each player chooses their best response myopically simultaneously in each period. Due to the complexity introduced by simultaneous updating, there are no known analytical results regarding convergence to equilibrium for this dynamic in random games. Obtaining such results is non-trivial because the path dependence is more intricate compared to clockwork best-response sequences.

2.3.4 Periodic Best-Response Dynamics

Finally, we introduce an intermediate case between clockwork and random sequences. A playing sequence is *p-periodic* if it consists of a sequence of players of length $p \geq n$ that is repeated forever, with the constraint that each player appears at least once in the repeated sequence. In other words,

$$i_1, i_2, \ldots, i_p, i_1, i_2, \ldots, i_p, i_1, i_2, \ldots, i_p, \ldots \tag{2.9}$$

is a p-periodic sequence, where for each $i \in [n]$, there is some $j \in [p]$ such that $i_j = i$. We generate p-periodic sequences randomly by constructing a sequence of

[9] When $m = 2$, the random digraph $\mathcal{D}(G_{n,2})$ is a random n-cube in which, independently, for each pair of profiles $\boldsymbol{a}$ and $\boldsymbol{a}'$ that differ in exactly one coordinate, there is a directed edge from $\boldsymbol{a}$ to $\boldsymbol{a}'$ with probability $1/2$; otherwise, there is a directed edge from $\boldsymbol{a}'$ to $\boldsymbol{a}$ with complementary probability $1/2$. For such an n-cube, Amiet et al. (2021) show that when n is large, every pure Nash equilibrium belongs to the set of vertices that are reachable by some directed path from the initial action profile $\boldsymbol{a}^0$. This is sufficient to show that, when the number of players is large, the random sequence best-response dynamic converges with probability 1 to a pure Nash equilibrium if there is one: indeed, the random playing sequence ensures that some path to equilibrium is played given sufficient time. The edges in $\mathcal{D}(G_{n,2})$ are oriented in one way or the other independently of each other, but this is no longer true when $m > 2$. In $\mathcal{D}(G_{n,m})$ with $m > 2$, if there is a directed edge from $(a_i, \boldsymbol{a}_{-i})$ to $(a_i', \boldsymbol{a}_{-i})$ for some $a_i \neq a_i'$, then the graph must also have the directed edges $(a_i'', \boldsymbol{a}_{-i})$ to $(a_i', \boldsymbol{a}_{-i})$ for all $a_i'' \neq a_i'$. This dependence and the more complex graph structure renders a generalisation of Proposition 1 to $m > 2$ non-trivial.

$p - n$ integers drawn randomly from $[n]$ and appending $1, \ldots, n$ to this sequence. This results in a sequence of p integers. We then select a random permutation of this sequence and repeat it indefinitely. If $p = n$, we recover clockwork playing sequences. Moreover, fixing n, we recover random playing sequences as $p \to \infty$.

This chapter continues by introducing more advanced learning dynamics, whose converge behaviour will be compared to best-response dynamics under different playing sequences. A key contribution is the formalisation of convergence criteria and the selection of parameters for multi-player games. This was not needed for best-response dynamics.

2.4 Other Learning Dynamics

In this section, we describe the learning dynamics, the convergence criteria chosen, and our choice of parameters for three additional learning dynamics: Bush-Mosteller reinforcement learning (Section 2.4.1), fictitious play (Section 2.4.2), and replicator dynamics (Section 2.4.3). The convergence properties of these learning algorithms have been studied theoretically. It is well-known, for instance, that fictitious play converges to Nash equilibrium in certain games such as potential, zero-sum, and supermodular games (Fudenberg and Levine, 1998). It is also established that evolutionarily stable strategies are locally stable fixed points of replicator dynamics (Hofbauer and Sigmund, 1998). However, there is no general result concerning the probability of convergence of these learning rules to pure Nash equilibria in random games. All the dynamics presented here are discrete or have been converted into discrete time.

There are unavoidable arbitrary choices in the specification of the dynamics, their parameters, and in the criteria that determine convergence to pure or mixed Nash equilibria. However, our findings suggest that the overall behaviour is robust to the specific implementation for all reasonable parameter settings.[10]

2.4.1 Reinforcement Learning

We use the Bush-Mosteller learning algorithm as a representative of reinforcement learning (Bush and Mosteller, 1953), following the specifications in Macy and Flache (2002) and Galla and Farmer (2013).

Each player has an *aspiration* level corresponding to a weighted average of payoffs received during gameplay. Each player then assigns a level of satisfaction to each action, which is positive if the action's payoff exceeds the player's aspiration

[10] Pangallo et al. (2019) tested this for 2-player games; we performed similar robustness tests in the multi-player case.

level and negative otherwise. The probability of playing an action increases if satisfaction is positive and decreases if it is negative.

2.4.1.1 Formal Description

In each period, player $i \in [n]$ chooses an action $x \in [m]$ with probability $p_i^t(x)$. The evolution of the mixed strategy of player i, $\boldsymbol{p}_i^t = (p_i^t(1), ..., p_i^t(m))$, is governed by reinforcement learning, as described below. The learning rule defines a mapping from $\boldsymbol{p}^t = (\boldsymbol{p}_1^t, ..., \boldsymbol{p}_n^t)$ to $\boldsymbol{p}^{t+1}$.

Let $\aleph_i^t$ be the *aspiration level* of player i in period t. It evolves according to:

$$\aleph_i^{t+1} = (1-\alpha)\aleph_i^t + \alpha u_i(x, \boldsymbol{a}_{-i}^t), \tag{2.10}$$

where $(x, \boldsymbol{a}_{-i}^t)$ is the action profile played in period t.

Thus, the aspiration level is a weighted average of the current payoff and the past aspiration level, with older payoffs discounted by $(1-\alpha)$. Here, α represents the rate of *memory loss*. Player i's *satisfaction* with action $x \in [m]$ in period t is defined as:

$$\sigma_i^t(x) = \frac{u_i(x, \boldsymbol{a}_{-i}^t) - \aleph_i^t}{\max_{\boldsymbol{y} \in [m]^n} |u_i(\boldsymbol{y}) - \aleph_i^t|}. \tag{2.11}$$

Note that $\sigma_i^t(x)$ is normalised to lie between -1 and 1. Positive satisfaction occurs if the payoff exceeds the aspiration level.

If player i played action x in period t, the probability of playing x again in period $t+1$ is updated as:

$$p_i^{t+1}(x) = \begin{cases} p_i^t(x) + \beta\sigma_i^t(x)(1 - p_i^t(x)) & \sigma_i^t(x) \geq 0 \\ p_i^{t+1}(x) + \beta\sigma_i^t(x)p_i^t(x) & \sigma_i^t(x) < 0 \end{cases}, \tag{2.12}$$

while the probability of choosing a different action $y \neq x$ is updated as:

$$p_i^{t+1}(y) = \begin{cases} p_i^t(y) - \beta\sigma_i^t(x)p_i^t(y) & \sigma_i^t(x) \geq 0 \\ p_i^t(y) - \beta\sigma_i^t(x)\frac{p_i^t(x)p_i^t(y)}{1-p_i^t(x)} & \sigma_i^t(x) < 0 \end{cases}. \tag{2.13}$$

Here, β represents the learning rate. Positive satisfaction for action x increases the likelihood of selecting x in the next period, while negative satisfaction decreases it. Notably, players do not need to track the actions of others — only their own actions and payoffs. This makes the Bush-Mosteller dynamic a classic example of reinforcement learning with limited information requirements.

2.4.1.2 Convergence Criteria

To assess convergence, we check whether $\boldsymbol{p}^t$ converges to a fixed point of the mapping $\boldsymbol{p}^t \mapsto \boldsymbol{p}^{t+1}$. This method allows us to capture convergence to mixed Nash equilibria, which would be missed if we only examined the actions played. Since the evolution of $\boldsymbol{p}^t$ is stochastic, we must account for random noise in our convergence assessment. As in Pangallo et al. (2019), we adopt the following heuristic for detecting convergence:

1. Consider only the last 20% of time steps to avoid transient effects.

2. Focus on actions with a probability greater than 0.05, averaged over the final time interval.

3. If the average standard deviation of these selected actions exceeds 0.01, the simulation is deemed non-convergent; otherwise, it is deemed convergent. A pure Nash equilibrium is identified if any belief vector $\boldsymbol{p}_i^t$ has a component greater than 0.98.

2.4.1.3 Parameter Values

We performed the simulations with $\alpha = 0.2$ and $\beta = 0.5$, following Pangallo et al. (2019). We observed only minor sensitivity to parameter variations within reasonable ranges in the 2-player and in the multi-player case.

2.4.1.4 Examples

Figure 2.3 shows example trajectories on the *Prisoner's Dilemma* game (Flood and Dresher, 1950) and Rock-Paper-Scissors under Bush-Mosteller learning. For the Prisoner's Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, a game with a single mixed Nash equilibrium where both players play any action with equal probability, none of the trajectories converge.

2.4.2 Fictitious Play

Under fictitious play (Brown, 1951; Robinson, 1951), each player estimates their opponents' mixed strategies based on their past actions, calculates the expected payoff of each action, and selects the action with the highest expected payoff. Variants include *weighted* fictitious play, which gives greater weight to recent actions, and *stochastic* fictitious play, where actions are selected probabilistically.

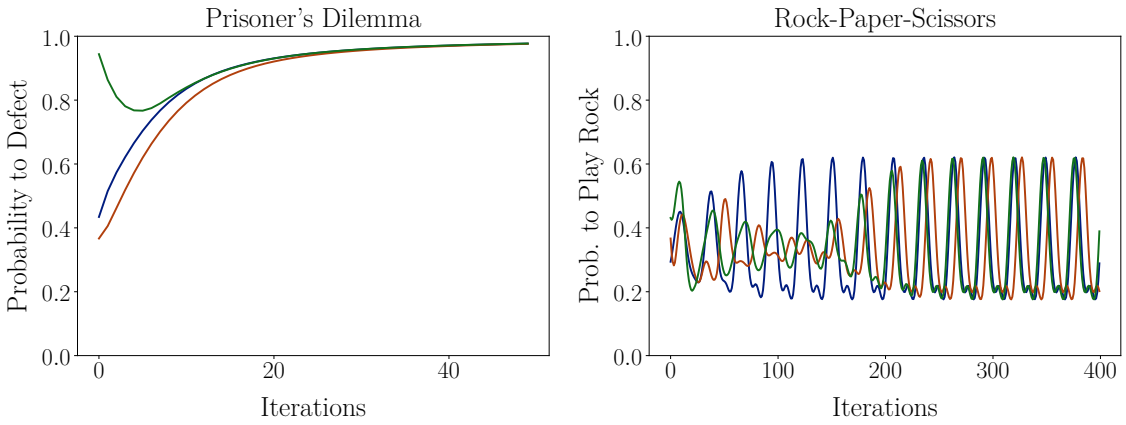


Figure 2.3: Example trajectories for the Prisoner's Dilemma and Rock-Paper-Scissors under Bush-Mosteller learning. For the Prisoner's Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, none of the trajectories converge.

2.4.2.1 Formal Description

In period $t \geq 0$, player j's belief $p_j^t(x)$ that they will play action x in period $t+1$ is given by the fraction of times they chose x in past periods:

$$p_j^t(x) = \frac{1}{t+1} \sum_{\tau=0}^{t} \mathbf{1}[a_j^\tau = x], \tag{2.14}$$

where $\mathbf{1}[a_j^\tau = x]$ is an indicator function equal to 1 if j played action x in period τ and 0 otherwise. In each period, player i then selects the action with the highest expected payoff based on their belief $\boldsymbol{p}_{-i}^t$ about opponents' actions:

$$a_i^{t+1} = \arg\max_{x \in [m]} \sum_{\boldsymbol{x_{-i}} \in [m]^{n-1}} u_i(x, \boldsymbol{x_{-i}}) \prod_{j \in [n] \setminus \{i\}} p_j^t(x_j). \tag{2.15}$$

2.4.2.2 Convergence Criteria

As in Fudenberg and Levine (1998), we assess convergence by tracking the evolution of beliefs $\boldsymbol{p}^t = (\boldsymbol{p}_1^t, ..., \boldsymbol{p}_n^t)$. The convergence criteria are similar to those for reinforcement learning, except that a pure Nash equilibrium is identified if any belief vector $\boldsymbol{p}_i^t$ has a component larger than 0.99.

2.4.2.3 Parameter Values

Fictitious play does not require parameter tuning.

2.4.2.4 Examples

Figure 2.4 shows example trajectories on the *Prisoner's Dilemma* game (Flood and Dresher, 1950) and Rock-Paper-Scissors under fictitious play. For the Prisoner's

Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, a game with a single mixed Nash equilibrium where both players play any action with equal probability, all trajectories converge to that equilibrium.

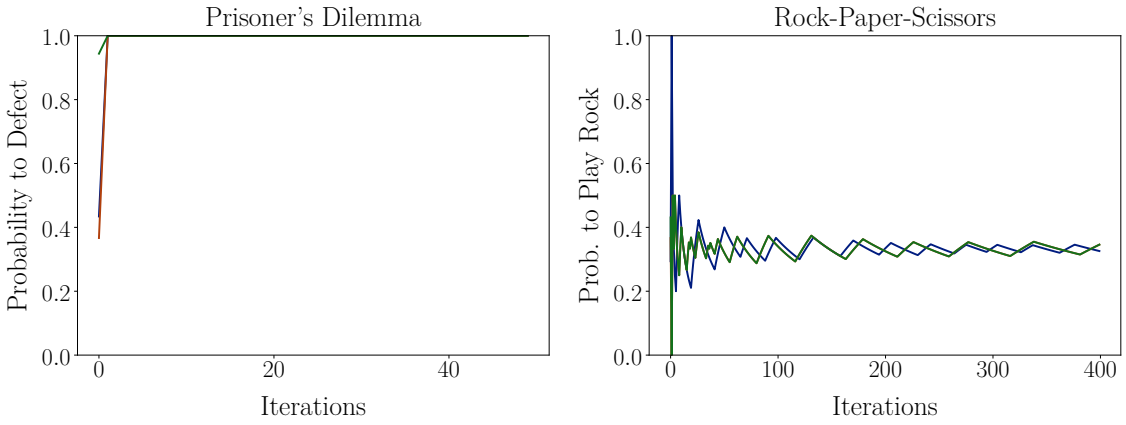


Figure 2.4: Example trajectories for the Prisoner's Dilemma and Rock-Paper-Scissors under Fictitious Play. For the Prisoner's Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, a game with a single mixed Nash equilibrium where both players play any action with equal probability, all trajectories converge to that equilibrium.

2.4.3 Replicator Dynamics

Replicator dynamics are one of the most fundamental models in evolutionary game theory (Maynard Smith, 1982). They describe the evolution of strategies in a population context, where the fitness of a strategy determines its growth rate. Following Börgers and Sarin (1997), we treat replicator dynamics as a learning algorithm for individual players. Our implementation involves a multi-population version to account for the general asymmetry of randomly generated payoff matrices (Taylor and Nowak, 2006; Gokhale and Traulsen, 2010).

2.4.3.1 Formal Description

In each period t, player i chooses an action x with probability $p_i^t(x)$, and the probability vector $\boldsymbol{p}_i^t = (p_i^t(1), ..., p_i^t(m))$ evolves according to the replicator equation.

When all other players act according to $\boldsymbol{p}_{-i}^t$, player i's expected payoff for action x is given by:

$$\tilde{u}_i(x, \boldsymbol{p}_{-i}^t) = \sum_{\boldsymbol{x}_{-i} \in [m]^{n-1}} u_i(x, \boldsymbol{x}_{-\boldsymbol{i}}) \prod_{j \in [n] \setminus \{i\}} p_j^t(x_j). \tag{2.16}$$

The average expected payoff for player i is:

$$\bar{u}_i(\boldsymbol{p}^t) = \sum_{x \in [m]} \tilde{u}_i(x, \boldsymbol{p}_{-i}^t) p_i^t(x). \tag{2.17}$$

In our simulation, the continuous replicator equation is discretised as:

$$p_i^{t+1}(x) = p_i^t(x)\frac{1+\delta\tilde{u}_i(x,\boldsymbol{p}_{-i}^t)}{1+\delta\bar{u}_i(\boldsymbol{p}^t)}. \tag{2.18}$$

2.4.3.2 Convergence Criteria

We track the convergence of $\boldsymbol{p}^t = (\boldsymbol{p}_1^t, ..., \boldsymbol{p}_n^t)$. Simulating replicator dynamics involves challenges due to machine precision limits, particularly when probabilities approach zero. We follow Pangallo et al. (2019) and apply the following criteria:

1. For each player, identify the action with the highest probability at the last time step and verify whether it has been increasing over the last 20% of time steps.

2. Ensure that the probabilities of all other actions have been decreasing over the same period.

3. If these conditions hold for all players, the simulation is considered convergent.

2.4.3.3 Parameter Values

Following Pangallo et al. (2019), we choose $\delta = 0.1$.

2.4.3.4 Examples

Figure 2.5 shows example trajectories on the *Prisoner's Dilemma* game (Flood and Dresher, 1950) and Rock-Paper-Scissors under replicator dynamics. For the Prisoner's Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, a game with a single mixed Nash equilibrium where both players play any action with equal probability, none of the trajectories converge.

2.5 Conclusion

In this chapter, we introduced several learning dynamics; four variations of best-response dynamics — clockwork, random, simultaneous, and periodic — and contrasted these with more sophisticated learning models such as Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics. This will allow us to capture a broad spectrum of dynamics in Chapter 3. Of course, numerous other playing sequences and learning algorithms could be explored. For example, approaches like *regret testing* (Foster and Young, 2006; Germano and Lugosi, 2007)

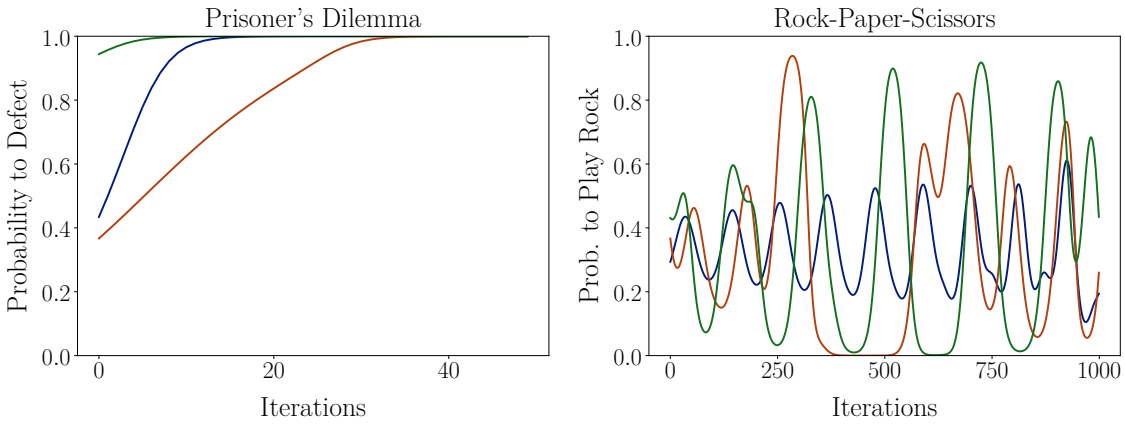


Figure 2.5: Example trajectories for the Prisoner's Dilemma and Rock-Paper-Scissors under Replicator Dynamics. For the Prisoner's Dilemma, all trajectories converge to the PNE where both players defect. For Rock-Paper-Scissors, none of the trajectories converge.

exhibit high-probability convergence to Nash equilibria in general n-player, m-action games. We leave the exploration of additional best-response dynamics and alternative learning rules for future research.

We also formalised the process of creating games at random, set the parameters for the learning dynamics, and defined the corresponding criteria for convergence. This lays the groundwork for studying the convergence behaviour of the introduced learning dynamics on random games. The following chapter will present our key results on the frequency of convergence of these learning dynamics.

3

Convergence on Random Games

This chapter presents novel results on the convergence behaviour of various learning dynamics in randomly generated games. In Section 3.2, we derive theoretical results regarding the convergence frequency of best-response dynamics. One of the central insights is the significant difference in the probability of convergence between the clockwork and random sequence best-response dynamics. For 2-player games, we provide explicit formulas for the convergence frequency under both clockwork and random sequences. Additionally, we introduce novel findings on the convergence frequency of games played on network structures. Section 3.3 focuses on numerical simulations, highlighting the convergence properties of different best-response dynamics, the impact of a network structure between players, and the speed of convergence. Furthermore, we compare the convergence properties of basic best-response dynamics with those of more advanced learning rules, establishing connections between the two. Finally, Section 3.4 provides a conclusion.

Contents

3.1 Introduction 38
3.2 Theoretical Results 40
3.2.1 Multi-Player Games 40
3.2.2 Two-Player Games 55
3.2.3 Games on Networks 62
3.3 Simulation Results 65
3.3.1 Simulation of Best-Response Dynamics 66
3.3.2 Comparison to Other Learning Rules 71
3.4 Conclusion 78

3.1 Introduction

We demonstrate that the playing sequence — the order in which players update their actions — is a critical determinant of whether the best-response dynamic converges to a Nash equilibrium. Specifically, we analyse the probability that the best-response dynamic converges to a pure Nash equilibrium in random n-player, m-action games under three distinct playing sequences: clockwork sequences (where players take turns according to a fixed cyclic order), random sequences, and simultaneous updating by all players. We provide an analytical characterisation of the convergence properties of the clockwork sequence best-response dynamic, with our key asymptotic result showing that this dynamic almost never converges to a pure Nash equilibrium when n and m are large. By contrast, the random sequence best-response dynamic almost always converges to a pure Nash equilibrium when one exists and when n and m are large. The clockwork best-response dynamic deserves special attention: our simulations show that its convergence properties most closely resemble those of three popular learning rules calibrated to human gameplay in experiments — namely, reinforcement learning, fictitious play, and replicator dynamics.

To examine the role of the playing sequence in determining the convergence properties of the best-response dynamic, we focus on the probability of convergence to a pure Nash equilibrium in random n-player games, where each player selects among m pure strategies. In other words, we generate a game by drawing all payoffs at random (from atomless distributions to avoid payoff ties) and assess the likelihood that the best-response dynamic, starting from a randomly selected pure strategy profile, converges to a pure Nash equilibrium of the randomly generated game. This approach abstracts from any specific game structure, allowing us to concentrate solely on the role of the playing sequence. Random games are valuable as null models, representing generic strategic interactions in various settings.

For 2-player games, we provide detailed results on the convergence properties of best-response dynamics, including the game duration and exact frequency of convergence under clockwork and random best-response dynamics. In multi-player games, we show that the probability of convergence under the clockwork best-response dynamic is up to a polynomial factor of order $1/\sqrt{m^{n-1}}$. This has two important implications: first, as the number of players n or the number of actions m increases ($nm \to \infty$), the probability of convergence to a pure Nash equilibrium approaches zero. Second, since the asymptotic convergence probability depends mainly on the term m^{n-1}, we observe that, when n or m is large, the probability of convergence to a pure Nash equilibrium in n-player, m-action games is similar to that in 2-player, m^{n-1}-action games. Our simulations further indicate that this

asymptotic relationship between n-player, m-action games and 2-player, m^{n-1}-action games holds reasonably well even for smaller values of n and m.

Our results for the convergence properties of the best-response dynamic under random and simultaneous updating are primarily simulation-based. Under a random sequence, conditional on the game having a pure Nash equilibrium, we demonstrate that the probability of convergence approaches one as n or m increases.[1] This sharp contrast with the clockwork sequence highlights the significant impact of the playing sequence on the probability of convergence to equilibrium. While most existing literature on best-response dynamics has concentrated on identifying sufficient conditions in the game's payoff structure to guarantee convergence, our findings suggest that the playing sequence is equally critical. To reinforce this observation, we show that the playing sequence also influences the probability of convergence in random games with correlated payoffs.[2] For example, in 2-player games with highly positively correlated payoffs, the best-response dynamic under simultaneous updating is unlikely to converge, while convergence is more likely under clockwork or random sequences. Across all levels of payoff correlation, the random sequence best-response dynamic has the highest probability of converging to a pure Nash equilibrium, while simultaneous updating has the lowest.

Through extensive simulations, we show that the clockwork best-response dynamic's convergence frequency closely tracks that of three prominent learning rules: the Bush-Mosteller reinforcement learning algorithm (Bush and Mosteller, 1953), fictitious play (Brown, 1951; Robinson, 1951), and replicator dynamics (Smith et al., 1992).[3] Although these learning algorithms are generally framed with simultaneous updating, we find that the clockwork sequence best-response dynamic outperforms both random sequence and simultaneous updating dynamics in most simulations, particularly when we vary n, m, or the level of payoff correlation. Moreover, the paths traced by the clockwork sequence in the space of action profiles most closely resemble those of the three learning algorithms, compared with random sequence dynamics.

[1] Amiet et al. (2021) prove this analytically for the case where $m = 2$ and $n \to \infty$.

[2] See Goldberg et al. (1968); Stanford (1999); Berg and Weigt (1999); Rinott and Scarsini (2000); Galla and Farmer (2013); Sanders et al. (2018) for previous work on random games with payoff correlations.

[3] The convergence properties of these learning algorithms have been well studied (Fudenberg and Levine, 1998), but there is no general result regarding their probability of convergence to Nash equilibria in random games.

3.2 Theoretical Results

In this section, we will provide theoretical results on the frequency of convergence of clockwork best-response dynamics. We start with a general result contrasting clockwork and random best-response dynamics (Section 3.2.1). Our key insight is that the playing sequence determines whether the best-response dynamic converges to a pure Nash equilibrium. We then continue with 2-player games in Section 3.2.2, which allow us to provide more granular results on the convergence properties of the clockwork best-response dynamic. In particular, we provide results on game duration and derive an exact expression for the probability that the best-response dynamic reaches a (best-response) cycle of a given length at a particular period. As a special case, we obtain the exact probability that the clockwork best-response dynamic converges to a pure Nash equilibrium in 2-player, m-action games (and we argue that, in the 2-player, m-action case, this probability is the same for random playing sequences). Furthermore, we show that this probability is asymptotically $\sqrt{\pi/m}$ when m is large. Finally, in Section 3.2.3, we derive analytical expressions for the probability of convergence of clockwork best-response dynamics when players are connected in star and bipartite networks.

3.2.1 Multi-Player Games

This section shows a striking difference between the probability of convergence to a PNE of the clockwork sequence vs the random sequence best-response dynamic. Roughly, while the clockwork sequence best-response dynamic never converges to a PNE in large games, the random sequence best-response dynamic always converges in large games (conditional on there being a PNE). Our following example develops some intuition for this result.

Example (Best-Response Dynamics for the Clockwork vs. Random Playing Sequence)**.** In the best-response digraph of Figure 2.2, we saw that the clockwork playing sequence may not converge to the PNE, depending on the initial profile. By contrast, the random playing sequence best-response dynamic converges to the PNE at $(1, 2, 1)$ with probability 1 given sufficient time. The path cannot be stuck in the best-response cycle on the front face of the cube, for example, because there is a positive probability that the path will escape to the PNE.

Figure 3.1 provides further examples of best-response digraphs. Panels (a) and (b) illustrate possible best-response digraphs for 3-player 2-action games. In the digraph of Panel (a), there are two PNEs, which are represented by black dots on the profiles $(1, 1, 1)$ and $(2, 2, 1)$. In other words, these profiles are the PNE of any game whose best-response digraph is the one illustrated in Panel (a). While the random

sequence best-response dynamic converges to one of the PNE with probability 1, the clockwork sequence best-response dynamic will converge if the initial profile is one of the four at the bottom of the cube. Since there is no PNE in the Panel (b) digraph, the best-response dynamic does not converge to a PNE regardless of the playing sequence. Finally, Panel (c) illustrates a possible best-response digraph for 3-player 3-action games. The unique PNE at $(3, 3, 1)$ is represented by a black dot. Note that there is a directed edge from $(3, 1, 1)$ to $(3, 3, 1)$ as well as a directed edge from $(3, 2, 1)$ to $(3, 3, 1)$, but these edges overlap in our illustration, so appear as a single long edge at the bottom of the front face of the cube. This digraph shows us a situation in which the random sequence best-response dynamic may not converge to a PNE even if there is one: indeed, if the path reaches one of the action profiles illustrated as a red dot, then the path can never escape to the PNE regardless of the playing sequence. As the results below imply, cases like this become vanishingly rare in large games.

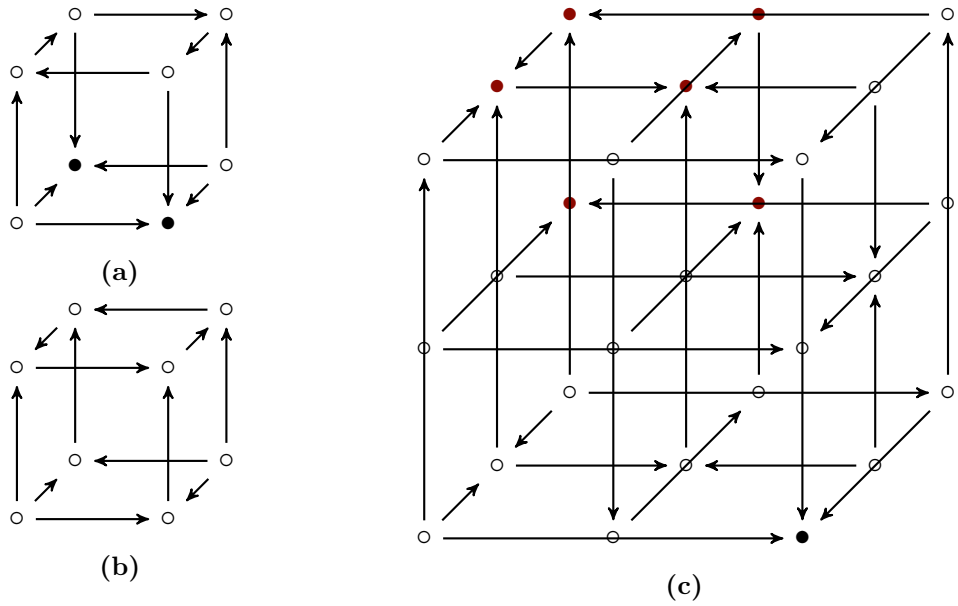


Figure 3.1: Panels (a) and (b) provide examples of possible best-response digraphs for 3-player, 2-action games. Panel (c) illustrates a possible best-response digraph for a 3-player, 3-action game. PNEs are denoted by filled black dots. If the path reaches one of the action profiles illustrated as a red dot, then the path can never escape to the PNE, regardless of the playing sequence.

Next, we present the theoretical results for best-response dynamics in random games.[4] In the multi-player case, we find that best-response dynamics behave very

[4]Note that all the convergence results hold modulo any relabelling of the players. For example, while we described the clockwork playing sequence as ordering the players according to 1-2-$\cdots$-n, our results would equally hold if the sequence had ordered the players according to n-$\cdots$-2-1 or any such permutation.

differently under clockwork vs. random playing sequences. Most of our results on the probability of convergence to equilibrium are asymptotic.

Our main result for the clockwork sequence best-response dynamic in games with $n > 2$ players is given below.

Theorem 1.

$$\lim_{nm\to\infty} \Pr\left[s_{\mathsf{c}}\text{-best-response dynamic on } G_{n,m} \text{ converges to a PNE}\right] = 0. \tag{3.1}$$

So, with high probability, the clockwork sequence best-response dynamic does not converge to a PNE as the number of players or actions gets large. This contrasts sharply with the asymptotic behaviour of the random sequence best-response dynamic.

Theorem 1 is an immediate consequence of the result below, which gives us bounds on the probability of convergence to equilibrium:

Theorem 2.

$$\frac{1}{4\sqrt{n}}\frac{1}{\sqrt{m^{n-1}}} \leq \Pr\left[\begin{array}{c} s_{\mathsf{c}}\text{-best-response dynamic} \\ \text{on } G_{n,m} \text{ converges to a PNE} \end{array}\right] \leq \frac{7n^{3/2}\sqrt{\log m}}{\sqrt{m^{n-1}}}. \tag{3.2}$$

Theorem 2 also gives us the following corollary:

Corollary 1. *The probability that the clockwork sequence best-response dynamic converges to a PNE is, up to a polynomial factor, of order* $1/\sqrt{m^{n-1}}$.

This result gives us a clear "scaling" law: since the asymptotic convergence probability depends essentially only on the quantity m^{n-1}, we have that when n or m are large, the probability of convergence to a pure Nash equilibrium in n-player, m-action games is approximately the same as it is in 2-player m^{n-1}-action games.[5] This scaling is reflected in our simulations even for relatively small values of m and n.

We briefly comment on the approach to derive Theorem 2. As the discussion following Algorithm 2 shows, drawing payoffs independently at random (from atomless distributions) induces a uniform distribution over best-response digraphs. So, the probability of convergence to a pure Nash equilibrium can be reduced to working out the probability that the path generated by Algorithm 2 initiated at a random vertex reaches a sink of the randomly drawn digraph. The main theoretical challenge we face when analysing such paths is that they exhibit some path dependence: if a player encounters an environment they had seen before, they must play the same action they played when the environment was first encountered.

[5]Indeed, consider G_{n_1,m_1} and G_{n_2,m_2} where $n_2 = 2$ and $m_2 = m_1^{n_1-1}$. Then, $\sqrt{m_2^{n_2-1}} = \sqrt{m_1^{n_1-1}}$.

We tackle this issue by relying on a coupling argument in which the clockwork best-response dynamic is coupled to a (memoryless) random walk through the digraph that is easier to analyse.

We start the proof of Theorem 2 by stating two lemmas: Lemma 1 bounds the probability that the clockwork sequence best-response dynamic converges to a pure Nash equilibrium or a best-response cycle only after period t, Lemma 2 bounds the probability that the clockwork sequence best-response dynamic converges to a pure Nash equilibrium by period t.

Lemma 1. *Let $\langle \vec{\mathbf{A}}, s_c \rangle$ be generated according to Algorithm 2 (the s-sequence best-response dynamic). For any $t \in \mathbb{N}$,*

$$\Pr\left[T_{\langle \vec{\mathbf{a}}, s_c \rangle} > t\right] \leq \exp\left(-\frac{\left(\left\lfloor \frac{t}{n} - 1 \right\rfloor\right)^2}{2m^{n-1}}\right). \tag{3.3}$$

Recall that $T_{\langle \vec{\mathbf{A}}, s_c \rangle}$ is the period in which $\langle \vec{\mathbf{A}}, s_c \rangle$ reaches $\mathrm{PNE}(G_{n,m})$ or a best-response cycle.

Lemma 2. *Let $\langle \vec{\mathbf{A}}, s_c \rangle$ be generated according to Algorithm 2 (the s-sequence best-response dynamic). For any $t \in \mathbb{N}$,*

$$\begin{aligned}\left\lceil \frac{t}{n} \right\rceil \frac{1}{m^{n-1}} \left(1 - \frac{n}{m^{n-1}} \frac{\left(\left\lceil \frac{t}{n}^2 \right\rceil\right)}{2}\right) &\leq \Pr\left[\langle \vec{\mathbf{A}}, s_c \rangle \text{ reaches } \mathrm{PNE}(G_{n,m}) \text{ by } t\right] \\ &\leq \frac{t}{m^{n-1}}.\end{aligned} \tag{3.4}$$

We now show how Theorem 2 follows from Lemmas 1 and 2. In what remains of this section, we provide proofs for the lemmas themselves.

Proof of Theorem 2. Let $\langle \vec{\mathbf{A}}, s_c \rangle$ be generated according to Algorithm 2. The probability that the s_c-best-response dynamic on $G_{n,m}$ converges to a PNE is equal to the probability that $\langle \vec{\mathbf{A}}, s_c \rangle$ reaches $\mathrm{PNE}(G_{n,m})$. Let us start with the upper bound. For any $t \in \mathbb{N}$,

$$\begin{aligned}&\Pr\left[\langle \vec{\mathbf{A}}, s_c \rangle \text{ reaches } \mathrm{PNE}(G_{n,m})\right] \\ =&\Pr\left[\langle \vec{\mathbf{A}}, s_c \rangle \text{ reaches } \mathrm{PNE}(G_{n,m}) \text{ by } t\right] + \Pr\left[\langle \vec{\mathbf{A}}, s_c \rangle \text{ reaches } \mathrm{PNE}(G_{n,m}) \text{ after } t\right] \\ \leq&\Pr\left[\langle \vec{\mathbf{A}}, s_c \rangle \text{ reaches } \mathrm{PNE}(G_{n,m}) \text{ by } t\right] + \Pr\left[T_{\langle \vec{\mathbf{a}}, s_c \rangle} > t\right] \\ \leq&\frac{t}{m^{n-1}} + \exp\left(-\frac{\left(\left\lfloor \frac{t}{n} - 1 \right\rfloor\right)^2}{2m^{n-1}}\right).\end{aligned} \tag{3.5}$$

The Inequality (3.5) follows from Lemmas 1 and 2. Now, set

$$t = n\left(\left\lceil\sqrt{2m^{n-1}\log(m^{n-1})}\right\rceil + 1\right). \tag{3.6}$$

Since $x \leq \lceil x \rceil \leq x+1$ and $\sqrt{2m^{n-1}\log(m^{n-1})} > 1$ for $m \geq 2$ and $n \geq 2$, we obtain

$$\begin{aligned} n\left(\sqrt{2m^{n-1}\log(m^{n-1})}+1\right) &\leq t \\ &\leq n\left(\sqrt{2m^{n-1}\log(m^{n-1})}+2\right) \\ &< 3n\sqrt{2m^{n-1}\log(m^{n-1})}. \end{aligned} \tag{3.7}$$

It follows that

$$\frac{t}{m^{n-1}} < \frac{3n\sqrt{2m^{n-1}\log(m^{n-1})}}{m^{n-1}} < \frac{5n^{3/2}\sqrt{\log m}}{\sqrt{m^{n-1}}}, \tag{3.8}$$

and that

$$\exp\left(-\frac{\left(\left\lfloor\frac{t}{n}-1\right\rfloor\right)^2}{2m^{n-1}}\right) \leq \frac{1}{m^{n-1}} < \frac{2n^{3/2}\sqrt{\log m}}{\sqrt{m^{n-1}}}. \tag{3.9}$$

Adding the upper bounds in Inequality (3.8) and Inequality (3.9) yields the desired result.

Let us now turn to the lower bound. For any $t \in \mathbb{N}$,

$$\begin{aligned} &\Pr\left[\langle\vec{\mathbf{A}}, s_{\mathsf{c}}\rangle \text{ reaches PNE}(G_{n,m})\right] \\ \geq &\Pr\left[\langle\vec{\mathbf{A}}, s_{\mathsf{c}}\rangle \text{ reaches PNE}(G_{n,m}) \text{ by } t\right] \\ \geq &\left\lceil\frac{t}{n}\right\rceil\frac{1}{m^{n-1}}\left(1-\frac{n}{m^{n-1}}\frac{\left(\left\lceil\frac{t}{n}\right\rceil\right)^2}{2}\right). \end{aligned} \tag{3.10}$$

Inequality (3.10) follows from Lemma 2. Now, set

$$t = n\left\lfloor\frac{\sqrt{m^{n-1}}}{\sqrt{n}}\right\rfloor. \tag{3.11}$$

Since $m \geq 2$ and $n \geq 2$, we obtain $\frac{1}{2}\sqrt{n}\sqrt{m^{n-1}} \leq t \leq \sqrt{n}\sqrt{m^{n-1}}$. It follows that

$$1-\frac{n}{m^{n-1}}\frac{\left(\left\lceil\frac{t}{n}\right\rceil\right)^2}{2} \geq \frac{1}{2}, \tag{3.12}$$

and that

$$\left\lceil\frac{t}{n}\right\rceil\frac{1}{m^{n-1}} \geq \frac{1}{2\sqrt{n}}\frac{1}{\sqrt{m^{n-1}}}. \tag{3.13}$$

Multiplying the lower bounds in Inequality (3.12) and Inequality (3.13) together yields the desired result. □

We now turn to the proofs of Lemmas 1 and 2. The main challenge posed by paths generated according to the s-sequence best-response dynamic (Algorithm 2) is that they have "memory": whenever player $s_{\mathsf{c}}(t)$ encounters an environment that she has encountered before (i.e. $\boldsymbol{A}^{t-1}_{-s_{\mathsf{c}}(t)} = \boldsymbol{A}^{u-1}_{-s_{\mathsf{c}}(u)}$ and $s_{\mathsf{c}}(t) = s_{\mathsf{c}}(u)$) then, in period t, the player must play the same action that she played when she previously encountered the environment (i.e. $A^t_{s_{\mathsf{c}}(t)} = A^u_{s_{\mathsf{c}}(u)}$). This path dependence complicates the analysis of the clockwork best-response dynamic. We, therefore, study a simpler (random walk) process that is "memoryless" to which we couple a dynamic that induces the same distribution over paths as Algorithm 2. The coupled dynamic follows the random walk process until an environment is encountered by some player for the second time and becomes deterministic after that.

The coupled system is described by Algorithms 3 and 4 below.

Algorithm 3 The clockwork random walk.

1. Draw an initial profile $\boldsymbol{X}^0$ uniformly at random from $[m]^n$
2. For $t \in \mathbb{N}$:
   (a) Set $i = s_{\mathsf{c}}(t)$
   (b) Set $\boldsymbol{X}^t_{-i} = \boldsymbol{X}^{t-1}_{-i}$
   (c) Independently draw X^t_i uniformly at random from $[m]$

Algorithm 4 The coupled dynamic, which follows the random walk process until an environment is encountered by some player for the second time and becomes deterministic after that.

1. Set $R_i(\boldsymbol{a}_{-i}) = 0$ for all $i \in [n]$ and $\boldsymbol{a}_{-i} \in [m]^{n-1}$
2. Set the initial action profile to $\boldsymbol{Y}^0 = \boldsymbol{X}^0$
3. For $t \in \mathbb{N}$:
   (a) Set $i = s_{\mathsf{c}}(t)$
   (b) Set $\boldsymbol{Y}^t_{-i} = \boldsymbol{Y}^{t-1}_{-i}$
   (c) If $R_i\left(\boldsymbol{Y}^{t-1}_{-i}\right) = 0$: set $Y^t_i = X^t_i$ and $R_i\left(\boldsymbol{Y}^{t-1}_{-i}\right) = Y^t_i$
   If $R_i\left(\boldsymbol{Y}^{t-1}_{-i}\right) \neq 0$: set $Y^t_i = R_i\left(\boldsymbol{Y}^{t-1}_{-i}\right)$

We will use $\langle \vec{\mathbf{X}}, s_{\mathsf{c}} \rangle$ and $\langle \vec{\mathbf{Y}}, s_{\mathsf{c}} \rangle$ to denote paths generated according to Algorithms 3 and 4, respectively.

Algorithm 3 is a "clockwork random walk" on the set of action profiles $[m]^n$. The walk starts at some randomly drawn initial profile $\boldsymbol{X}^0$ and, in each period, t moves

in direction $s_{\mathsf{c}}(t)$ to a profile chosen uniformly at random from among the m profiles in that direction. A path generated according to this process does not have memory.

Algorithm 4 describes the coupled dynamic. The process starts at the same initial profile as the clockwork random walk. For each player i and environment $\boldsymbol{a}_{-i}$, we set the initial "response" value $R_i(\boldsymbol{a}_{-i})$ to zero. The crucial step to how the process evolves is (3c): if the response value to the current environment $\boldsymbol{Y}_{-i}^{t-1}$ is zero, then the environment was never encountered before and, in that case, player i's response value is set to X_i^t, the action drawn by the clockwork random walk in period t. If, on the other hand, the response value to the current environment $\boldsymbol{Y}_{-i}^{t-1}$ is non-zero (i.e. the environment was encountered before), then this value is the action that i takes in period t. In other words, $\langle \vec{\mathbf{Y}}, s_{\mathsf{c}} \rangle$ has the same memory property that is characteristic of paths generated according to Algorithm 2.[6]

Recall that Algorithm 2 essentially draws a best-response digraph, selects an initial profile, and then traces a path by travelling along the edges of the digraph starting at the initial profile and moving in direction $s_{\mathsf{c}}(t)$ at step t. The entire best-response digraph is drawn upfront under Algorithm 2. In contrast, Algorithm 4 starts with an empty digraph and then generates its edges in an "online" manner. Nevertheless, the following remark summarises that both algorithms induce the same distribution over paths.

Remark 1. Let $\langle \vec{\mathbf{A}}, s_{\mathsf{c}} \rangle$ and $\langle \vec{\mathbf{Y}}, s_{\mathsf{c}} \rangle$ be generated according to Algorithms 2 and 4 respectively. Then $\langle \vec{\mathbf{A}}, s_{\mathsf{c}} \rangle$ and $\langle \vec{\mathbf{Y}}, s_{\mathsf{c}} \rangle$ have the same distribution.

For any path $\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle$ and for each $t \in \mathbb{N}$ define

$$f_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle}(t) := \min \left(u \leq t : \boldsymbol{a}_{-s_{\mathsf{c}}(u)}^{u-1} = \boldsymbol{a}_{-s_{\mathsf{c}}(t)}^{t-1} \text{ and } s_{\mathsf{c}}(u) = s_{\mathsf{c}}(t) \right). \tag{3.14}$$

So $f_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle}(t)$ is the first period along the path $\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle$ that player $s_{\mathsf{c}}(t)$ encounters the environment $\boldsymbol{a}_{-s_{\mathsf{c}}(t)}^{t-1}$. Notice that if $s_{\mathsf{c}}(t)$ encounters $\boldsymbol{a}_{-s_{\mathsf{c}}(t)}^{t-1}$ for the first time in period t then $f_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle}(t) = t$, and if $s_{\mathsf{c}}(t)$ encountered $\boldsymbol{a}_{-s_{\mathsf{c}}(t)}^{t-1}$ for the first time in some period $u < t$ then $f_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle}(t) < t$. We also define

$$F_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle} := \inf \left(t \in \mathbb{N} : f_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle}(t) < t \right). \tag{3.15}$$

So $F_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle}$ is the first period in which some player encounters an environment that they encountered previously along the path. The value $F_{\langle \vec{\mathbf{a}}, s_{\mathsf{c}} \rangle}$ is bounded above by $1 + nm^{n-1}$ for any path.

[6] The pattern of path-dependence is more complex for simultaneous updating than for the clockwork playing sequence. In the case of simultaneous updating, the choices of the players who have encountered an environment twice are deterministic. However, the choices of the players who have never encountered an environment twice remain random. For example, suppose $\boldsymbol{A}^0 = (1, 1, 1)$ and $\boldsymbol{A}^1 = (2, 1, 1)$. Then, player 1 must repeat action 2 in period 2, but the actions of players 2 and 3 in period 2 are not deterministic. Keeping track of the evolution of the path under simultaneous updating is, therefore, complex.

By construction, the sequences $\vec{\mathbf{X}}$ and $\vec{\mathbf{Y}}$ must agree at least up to (but not including) the period at which some player encounters an environment for the second time. In that period, under Algorithm 4, the player must play the action determined by their response function evaluated at that environment, but under Algorithm 3, the following action may be any of the available actions for that player. Remark 2 summarises the key relationship between the clockwork random walk and the coupled dynamic.

Remark 2.

$$F_{\langle \vec{\mathbf{X}}, s_c \rangle} = F_{\langle \vec{\mathbf{Y}}, s_c \rangle}. \tag{3.16}$$

Example (Illustration of Algorithms 3 and 4)**.** Figure 3.2 illustrates the relationship between $\langle \vec{\mathbf{X}}, s_c \rangle$ and $\langle \vec{\mathbf{Y}}, s_c \rangle$ by plotting the first few elements of $\vec{\mathbf{X}}$ and of $\vec{\mathbf{Y}}$. Panel (a) shows the first few elements of a possible path generated according to the clockwork random walk starting at the profile $\boldsymbol{X}^0 = (1, 1, 1)$. We represent the coupled dynamic in Panel (b), starting with an empty digraph and numbering the directed edges according to the first placed period.

In the clockwork random walk, player 1 chooses action 2 in period 1. This generates the directed edge from $\boldsymbol{Y}^0$ to $\boldsymbol{Y}^1$ in period 1 in Panel (b). The paths are identical up to and including period 6. In period 7, however, player 1 encounters the same environment she had encountered in period 1 (namely, players 2 and 3, each choosing action 1). The first time player 1 encountered this environment, she responded by playing action 2, so she must play action 2 again in period 7. In other words, the path must follow the edge placed in period 1 and therefore $\boldsymbol{Y}^7 = (2, 1, 1)$. This is not true for the clockwork random walk. Since the process is memoryless, it can remain at $\boldsymbol{X}^7 = (1, 1, 1)$ in period 7 and travel to $\boldsymbol{X}^8 = (1, 2, 1)$ in period 8.

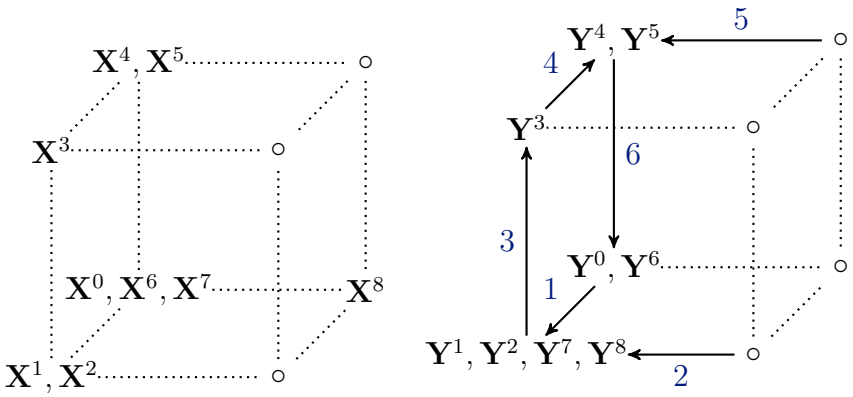
The path in Panel (b) will keep cycling among the action profiles on the left-hand side of the cube forever, whereas the path in Panel (a) is allowed to wander freely. Note here that $F_{\langle \vec{\mathbf{X}}, s_c \rangle} = F_{\langle \vec{\mathbf{Y}}, s_c \rangle} = 7$.

Remark 3. Any path $\langle \vec{\mathbf{A}}, s_c \rangle$ generated according to Algorithm 2 must reach $\mathrm{PNE}(G_{n,m})$ or a best-response cycle before any player encounters an environment for the second time:

$$T_{\langle \vec{\mathbf{A}}, s_c \rangle} < F_{\langle \vec{\mathbf{A}}, s_c \rangle}, \tag{3.17}$$

where $T_{\langle \vec{\mathbf{A}}, s_c \rangle}$ is the earliest period in which the path reaches a PNE or a best-response cycle (see Equation 2.6).

Example (Illustration of Remark 3)**.** Table 3.1 shows the values of the function $f_{\langle \vec{\mathbf{a}}, s_c \rangle}(t)$ for the first few elements of the paths generated according to the clockwork sequence best-response dynamic in Panel (a) and Panel (b) of Figure 2.2. Recall



(a) **(b)**

Period	Player	$\vec{\mathbf{X}}$	$f_{\langle \vec{\mathbf{X}}, s_c \rangle}(t)$	$\vec{\mathbf{Y}}$	$f_{\langle \vec{\mathbf{Y}}, s_c \rangle}(t)$
–	–	$(1, \mathbf{1}, \mathbf{1})$	–	$(1, \mathbf{1}, \mathbf{1})$	–
1	1	$(\mathbf{2}, 1, \mathbf{1})$	1	$(\mathbf{2}, 1, \mathbf{1})$	1
2	2	$(\mathbf{2}, \mathbf{1}, 1)$	2	$(\mathbf{2}, \mathbf{1}, 1)$	2
3	3	$(2, \mathbf{1}, \mathbf{2})$	3	$(2, \mathbf{1}, \mathbf{2})$	3
4	1	$(\mathbf{1}, 1, \mathbf{2})$	4	$(\mathbf{1}, 1, \mathbf{2})$	4
5	2	$(\mathbf{1}, \mathbf{1}, 2)$	5	$(\mathbf{1}, \mathbf{1}, 2)$	5
6	3	$(1, \mathbf{1}, \mathbf{1})$	6	$(1, \mathbf{1}, \mathbf{1})$	6
7	1	$(\mathbf{1}, 1, \mathbf{1})$	1	$(\mathbf{2}, 1, \mathbf{1})$	1
8	2	$(\mathbf{1}, \mathbf{2}, 1)$	8	$(\mathbf{2}, \mathbf{1}, 1)$	2

Figure 3.2: Illustration of Algorithms 3 (the clockwork random walk) and 4 (the coupled dynamic). Panel (a) shows the first elements of a possible path generated according to Algorithm 3, and Panel (b) shows the corresponding path generated according to Algorithm 4. The table in Panel (c) provides details, highlighting environments in bold.

that the path in Panel (a) reaches the pure Nash equilibrium in period 3 and that the path in Panel (b) reaches a best-response cycle in period 1. Furthermore, from Table 3.1, we can see that the value of $F_{\langle \vec{\mathbf{a}}, s_c \rangle}$ is 6 for Panel (a) and 7 for Panel (b). We therefore conclude that, for Panel (a), $T_{\langle \vec{\mathbf{a}}, s_c \rangle} = 3 < 6 = F_{\langle \vec{\mathbf{a}}, s_c \rangle}$, and for Panel (b), $T_{\langle \vec{\mathbf{a}}, s_c \rangle} = 1 < 7 = F_{\langle \vec{\mathbf{a}}, s_c \rangle}$.

Lemma 3 below, which concerns paths $\langle \vec{\mathbf{X}}, s_c \rangle$ that are generated by the clockwork random walk, is useful for proving Lemmas 1 and 2. Under the clockwork sequence, player $i \in [n]$ plays in period $h_i(k) := i + (k-1)n$ for $k \in \mathbb{N}$. For any $i \in [n]$ and any period $t \in \mathbb{N}$, define

$$k_i^*(t) := 1 + \left\lfloor \frac{t-i}{n} \right\rfloor, \tag{3.18}$$

which is the largest $k \in \mathbb{N}$ such that $h_i(k) \leq t$. The environments that player $i \in [n]$ encounters on her turns between (and including) periods 1 and t are

Period	Player	Panel (a) $\vec{\mathbf{a}}$	Panel (a) $f_{\langle \vec{\mathbf{a}}, s_c \rangle}(t)$	Panel (b) $\vec{\mathbf{a}}$	Panel (b) $f_{\langle \vec{\mathbf{a}}, s_c \rangle}(t)$
–	–	$(1, \mathbf{1}, \mathbf{2})$	–	$(1, \mathbf{1}, \mathbf{1})$	–
1	1	$(\mathbf{1}, 1, \mathbf{2})$	1	$(\mathbf{2}, 1, \mathbf{1})$	1
2	2	$(\mathbf{1}, \mathbf{2}, 2)$	2	$(\mathbf{2}, \mathbf{2}, 1)$	2
3	3	$(1, \mathbf{2}, \mathbf{1})$	3	$(2, \mathbf{2}, \mathbf{2})$	3
4	1	$(\mathbf{1}, 2, \mathbf{1})$	4	$(\mathbf{2}, 2, \mathbf{2})$	4
5	2	$(\mathbf{1}, \mathbf{2}, 1)$	5	$(\mathbf{2}, \mathbf{1}, 2)$	5
6	3	$(1, \mathbf{2}, \mathbf{1})$	3	$(2, \mathbf{1}, \mathbf{1})$	6
7	1	$(\mathbf{1}, 2, \mathbf{1})$	4	$(\mathbf{2}, 1, \mathbf{1})$	1
8	2	$(\mathbf{1}, \mathbf{2}, 1)$	5	$(\mathbf{2}, \mathbf{2}, 1)$	2

Table 3.1: The value of the function $f_{\langle \vec{\mathbf{a}}, s_c \rangle}(t)$ for the first few elements of the paths generated according to the clockwork playing sequence in Panel (a) and Panel (b) of Figure 2.2. The path in Panel (a) reached the PNE in period 3, the path in Panel (b) reached a best-response cycle in period 1.

given in the sequence

$$\left(\boldsymbol{X}_{-i}^{h_i(1)-1}, \boldsymbol{X}_{-i}^{h_i(2)-1}, \ldots, \boldsymbol{X}_{-i}^{h_i(k_i^*(t))-1} \right). \tag{3.19}$$

Lemma 3 establishes bounds on the probability that these environments are all distinct.

Lemma 3. *For any $i \in [n]$ and $t \in \mathbb{N}$,*

$$\begin{aligned} 1 - \frac{1}{m^{n-1}} \frac{\left(\left\lceil \frac{t}{n} \right\rceil \right)^2}{2} &\leq \Pr\left[\boldsymbol{X}_{-i}^{h_i(k)-1} \text{ for } k \in \{1, ..., k_i^*(t)\} \text{ are all distinct} \right] \\ &\leq \exp\left(- \frac{\left(\left\lfloor \frac{t}{n} - 1 \right\rfloor \right)^2}{2m^{n-1}} \right). \end{aligned} \tag{3.20}$$

Proof of Lemma 3. For any $i \in [n]$, the environments

$$\boldsymbol{X}_{-i}^{h_i(1)-1}, \boldsymbol{X}_{-i}^{h_i(2)-1}, ..., \boldsymbol{X}_{-i}^{h_i(k_i^*(t))-1} \tag{3.21}$$

are independent because they are disjoint subsets of the draws of the clockwork random walk. Each environment is distributed uniformly on $[m]^{n-1}$. Therefore,

$$\Pr\left[\boldsymbol{X}_{-i}^{h_i(k)-1} \text{ for } k \in \{1, ..., k_i^*(t)\} \text{ are all distinct} \right] = \prod_{k=1}^{k_i^*(t)-1} \left(1 - \frac{k}{m^{n-1}} \right). \tag{3.22}$$

If $k_i^*(t) > 1 + m^{n-1}$, then the Lemma trivially holds, because Equation (3.22) is equal to zero, and $\left\lfloor \frac{t-i}{n} \right\rfloor > m^{n-1}$ which, in turn, implies $\left\lceil \frac{t}{n} \right\rceil > m^{n-1}$, so the lower

bound in the statement of the lemma is negative and the upper bound is positive. We will therefore consider the case in which $k_i^*(t) \leq 1 + m^{n-1}$.

We obtain the following upper bound:

$$\begin{aligned}\prod_{k=1}^{k_i^*(t)-1}\left(1-\frac{k}{m^{n-1}}\right) &\leq \prod_{k=1}^{k_i^*(t)-1}\exp\left(-\frac{k}{m^{n-1}}\right)\\ &\leq \exp\left(-\frac{\left(k_i^*(t)-1\right)^2}{2m^{n-1}}\right)\\ &\leq \exp\left(-\frac{\left(\left\lfloor\frac{t}{n}-1\right\rfloor\right)^2}{2m^{n-1}}\right).\end{aligned} \tag{3.23}$$

The first step follows from $\exp(x) \geq 1 + x$ for all x. The final inequality follows from

$$k_i^*(t) - 1 = \left\lfloor\frac{t-i}{n}\right\rfloor \geq \left\lfloor\frac{t-n}{n}\right\rfloor = \left\lfloor\frac{t}{n}-1\right\rfloor. \tag{3.24}$$

We now turn to the lower bound:

$$\begin{aligned}\prod_{k=1}^{k_i^*(t)-1}\left(1-\frac{k}{m^{n-1}}\right) &\geq 1-\sum_{k=1}^{k_i^*(t)-1}\frac{k}{m^{n-1}}\\ &= 1-\frac{1}{m^{n-1}}\frac{k_i^*(t)^2}{2}\\ &\geq 1-\frac{1}{m^{n-1}}\frac{\left(\left\lceil\frac{t}{n}\right\rceil\right)^2}{2}.\end{aligned} \tag{3.25}$$

The first step is to apply the Weierstrass product inequality[7]. The final inequality follows from the fact that

$$k_i^*(t) = 1 + \left\lfloor\frac{t-i}{n}\right\rfloor \leq 1 + \left\lfloor\frac{t-1}{n}\right\rfloor = \left\lceil\frac{t}{n}\right\rceil. \tag{3.26}$$

□

Proof of Lemma 1. $T_{\langle\vec{\mathbf{A}},s_c\rangle} > t$ is the event that $\langle\vec{\mathbf{A}}, s_c\rangle$ reaches PNE($G_{n,m}$) or a best-response cycle only after period t. Remark 3 implies that $F_{\langle\vec{\mathbf{A}},s_c\rangle} > t$. So

$$\Pr\left[T_{\langle\vec{\mathbf{A}},s_c\rangle} > t\right] \leq \Pr\left[F_{\langle\vec{\mathbf{A}},s_c\rangle} > t\right]. \tag{3.27}$$

By Remarks 1 and 2,

$$\Pr\left[F_{\langle\vec{\mathbf{A}},s_c\rangle} > t\right] = \Pr\left[F_{\langle\vec{\mathbf{Y}},s_c\rangle} > t\right] = \Pr\left[F_{\langle\vec{\mathbf{X}},s_c\rangle} > t\right]. \tag{3.28}$$

Let us focus on the path $\langle\vec{\mathbf{X}}, s_c\rangle$ and on player 1. The environments that player 1 faces between periods 1 and t are given in the sequence

$$\left(\boldsymbol{X}_{-1}^{h_1(1)-1}, \boldsymbol{X}_{-1}^{h_1(2)-1}, ..., \boldsymbol{X}_{-1}^{h_1(k_1^*(t))-1}\right). \tag{3.29}$$

[7] For $(a_i)_{i=1...n} \in [0,1]$: $\prod_{i=1}^n (1-a_i) \geq 1 - \sum_{i=1}^n a_i$.

The event $F_{\langle \vec{\mathbf{X}}, s_c \rangle} > t$ implies that the environments in this sequence are all distinct. Hence

$$\begin{aligned} \Pr\left[F_{\langle \vec{\mathbf{X}}, s_c \rangle} > t\right] &\leq \Pr\left[\boldsymbol{X}_{-1}^{h_1(k)-1} \text{ for } k \in \{1, ..., k_1^*(t)\} \text{ are all distinct}\right] \\ &\leq \exp\left(-\frac{\left(\left\lfloor \frac{t}{n} - 1 \right\rfloor\right)^2}{2m^{n-1}}\right), \end{aligned} \tag{3.30}$$

where the final step follows from Lemma 3. □

To prove Lemma 2, we introduce Algorithm 5, which describes a dynamic coupled with the clockwork random walk. We denote a path generated according to Algorithm 5 by $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_c \rangle$.

Algorithm 5 The coupled dynamic with sink $\boldsymbol{x}$, which initialises the directed edges to point at $\boldsymbol{x}$ for all players.

---

1. Set $R_i(\boldsymbol{a}_{-i}) = 0$ for all $i \in [n]$ and $\boldsymbol{a}_{-i} \in [m]^{n-1}$
2. Set $R_i(\boldsymbol{x}_{-i}) = x_i$ for all $i \in [n]$
3. Set the initial action profile to $\boldsymbol{Z}^0 = \boldsymbol{X}^0$
4. For $t \in \mathbb{N}$:
   (a) Set $i = s_c(t)$
   (b) Set $\boldsymbol{Z}_{-i}^t = \boldsymbol{Z}_{-i}^{t-1}$
   (c) If $R_i\left(\boldsymbol{Z}_{-i}^{t-1}\right) = 0$: set $Z_i^t = X_i^t$ and $R_i\left(\boldsymbol{Z}_{-i}^{t-1}\right) = Z_i^t$
   If $R_i\left(\boldsymbol{Z}_{-i}^{t-1}\right) \neq 0$: set $Z_i^t = R_i\left(\boldsymbol{Z}_{-i}^{t-1}\right)$

Algorithm 5 is identical to Algorithm 4 (the coupled dynamic) except that for some particular profile $\boldsymbol{x}$ the algorithm is initialized with $R_i(\boldsymbol{x}_{-i}) = x_i$ for all $i \in [n]$. Algorithm 5 therefore initializes the digraph with the directed edges $(x_i', \boldsymbol{x}_{-i})$ to $(x_i, \boldsymbol{x}_{-i})$ for all i and $x_i' \neq x_i$, so that the profile $\boldsymbol{x}$ is a sink. In the remaining steps, the algorithm selects a random initial profile and starts tracing a path by travelling along edges that (other than those edges already pointing to $\boldsymbol{x}$ in the initialisation) are generated in an online manner. The paths traced by the clockwork random walk and this coupled dynamic with a sink at $\boldsymbol{x}$ must agree at least up to (but not including) the period at which *either* an environment is encountered by a player for the second time *or* the environment is $\boldsymbol{x}_{-i}$ for some player i.

Example (Illustration of Algorithms 3 and 5)**.** Figure 3.3 illustrates the relationship between $\langle \vec{\mathbf{X}}, s_c \rangle$ and $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_c \rangle$ by plotting the first few elements of $\vec{\mathbf{X}}$ and of $\vec{\mathbf{Z}}$.

Panel (a) shows the first few elements of a possible path generated according to the clockwork random walk starting at the profile $\boldsymbol{X}^0 = (1, 1, 1)$. In Panel (b), we represent the corresponding path generated according to Algorithm 5. This time, rather than starting with an empty digraph, the profile $\boldsymbol{x}$ is made a sink (with the red edges placed in period 0). The remaining directed edges are numbered according to the first place period.

The clockwork random walk takes the path $\vec{\mathbf{Z}}$ to $\boldsymbol{Z}^4 = (1, 1, 2)$ in period 4. While the random walk can continue wandering through the action profiles according to the clockwork sequence, the path $\vec{\mathbf{Z}}$ must end up at $\boldsymbol{Z}^5 = \boldsymbol{x}$ in period 5.

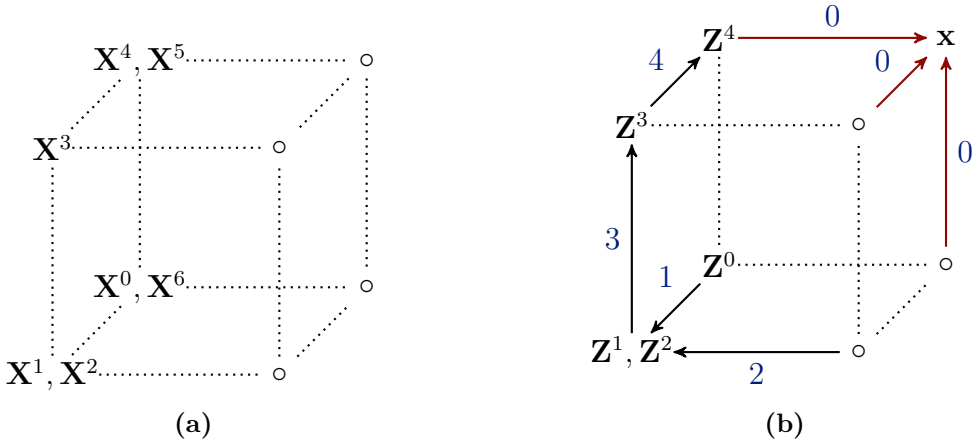


Figure 3.3: Illustration of Algorithms 3 (the clockwork random walk) and 5 (the coupled dynamic with sink $\boldsymbol{x}$). Panel (a) shows the first elements of a possible path generated according to Algorithm 3, and Panel (b) shows the corresponding path generated according to Algorithm 5.

Remark 4. Let $\langle \vec{\mathbf{A}}, s_{\mathsf{c}} \rangle$ and $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle$ be generated according to Algorithms 2 and 5 respectively. Then the distribution of $\langle \vec{\mathbf{A}}, s_{\mathsf{c}} \rangle$ *conditional* on $\boldsymbol{x} \in \mathrm{PNE}(G_{n,m})$ is the same as the distribution of $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle$.

Proof of Lemma 2. For any $t \in \mathbb{N}$,

$$
\begin{aligned}
&\Pr\left[\langle \vec{\mathbf{A}}, s_{\mathsf{c}} \rangle \text{ reaches } \mathrm{PNE}(G_{n,m}) \text{ by } t\right] \\
&= \sum_{\boldsymbol{x} \in [m]^n} \Pr\left[\langle \vec{\mathbf{A}}, s_{\mathsf{c}} \rangle \text{ reaches } \{\boldsymbol{x}\} \text{ by } t \text{ and } \boldsymbol{x} \in \mathrm{PNE}(G_{n,m})\right] \\
&= \sum_{\boldsymbol{x} \in [m]^n} \Pr\left[\langle \vec{\mathbf{A}}, s_{\mathsf{c}} \rangle \text{ reaches } \{\boldsymbol{x}\} \text{ by } t \,\middle|\, \boldsymbol{x} \in \mathrm{PNE}(G_{n,m})\right] \Pr\left[\boldsymbol{x} \in \mathrm{PNE}(G_{n,m})\right] \\
&= \sum_{\boldsymbol{x} \in [m]^n} \underbrace{\Pr\left[\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle \text{ reaches } \{\boldsymbol{x}\} \text{ by } t\right]}_{(3.31.1)} \underbrace{\Pr\left[\boldsymbol{x} \in \mathrm{PNE}(G_{n,m})\right]}_{(3.31.2)} \qquad (3.31)
\end{aligned}
$$

The first step follows from the definition of reaching a pure Nash equilibrium. The final step follows from Remark 4; namely, the probability that $\langle \vec{\mathbf{A}}, s_{\mathsf{c}} \rangle$ reaches $\{\boldsymbol{x}\}$ by period t conditional on $\boldsymbol{x} \in \mathrm{PNE}(G_{n,m})$ is equal to the probability that $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle$ reaches $\{\boldsymbol{x}\}$ by period t. We now analyse the expressions (3.31.1) and (3.31.2).

For (3.31.2), since payoffs are drawn identically and independently according to the atomless distribution $\mathbb{P}$, we have that

$$\Pr\left[\boldsymbol{x} \in \mathrm{PNE}(G_{n,m})\right] = \prod_{i=1}^{n} \Pr\left[U_i\left(\boldsymbol{x}\right) \geq \max_{x_i' \in [m]} U_i\left(x_i', \boldsymbol{x}_{-i}\right)\right] = \frac{1}{m^n}. \tag{3.32}$$

We now find upper and lower bounds on (3.31.1) by relating $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle$ to the clockwork random walk path $\langle \vec{\mathbf{X}}, s_{\mathsf{c}} \rangle$. We start with the upper bound. Notice that $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle$ cannot reach $\{\boldsymbol{x}\}$ by period t unless $\boldsymbol{X}_{-s_{\mathsf{c}}(\tau)}^{\tau-1} = \boldsymbol{x}_{-s_{\mathsf{c}}(\tau)}$ for some $\tau \leq t$. Therefore

$$\begin{aligned}\Pr\left[\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle \text{ reaches } \{\boldsymbol{x}\} \text{ by } t\right] &\leq \Pr\left[\bigcup_{\tau=1}^{t} \left\{\boldsymbol{X}_{-s_{\mathsf{c}}(\tau)}^{\tau-1} = \boldsymbol{x}_{-s_{\mathsf{c}}(\tau)}\right\}\right] \\ &\leq \sum_{\tau=1}^{t} \Pr\left[\boldsymbol{X}_{-s_{\mathsf{c}}(\tau)}^{\tau-1} = \boldsymbol{x}_{-s_{\mathsf{c}}(\tau)}\right] = \frac{t}{m^{n-1}}. \end{aligned} \tag{3.33}$$

The final step follows from the fact that $\boldsymbol{X}_{-s_{\mathsf{c}}(\tau)}^{\tau-1}$ consists of $n-1$ independent random variables, each uniformly distributed on $[m]$.

We now turn to the lower bound. If $F_{\langle \vec{\mathbf{X}}, s_{\mathsf{c}} \rangle} > t$ and $\boldsymbol{X}_{-s_{\mathsf{c}}(\tau)}^{\tau-1} = \boldsymbol{x}_{-s_{\mathsf{c}}(\tau)}$ for some $\tau \leq t$ then $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle$ must reach $\{\boldsymbol{x}\}$ by period t. In other words, if no environments are repeated for any player and the environment is $\boldsymbol{x}_{-i}$ for some player i by period t, then $\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle$ must reach $\{\boldsymbol{x}\}$ by period t. Therefore,

$$\begin{aligned} &\Pr\left[\langle \vec{\mathbf{Z}}, \boldsymbol{x}, s_{\mathsf{c}} \rangle \text{ reaches } \{\boldsymbol{x}\} \text{ by } t\right] \\ \geq\ &\Pr\left[\bigcup_{\tau=1}^{t} \left\{\boldsymbol{X}_{-s_{\mathsf{c}}(\tau)}^{\tau-1} = \boldsymbol{x}_{-s_{\mathsf{c}}(\tau)}\right\} \text{ and } F_{\langle \vec{\mathbf{X}}, s_{\mathsf{c}} \rangle} > t\right] \\ =\ &\Pr\left[\bigcup_{\tau=1}^{t} \left\{\boldsymbol{X}_{-s_{\mathsf{c}}(\tau)}^{\tau-1} = \boldsymbol{x}_{-s_{\mathsf{c}}(\tau)}\right\} \,\middle|\, F_{\langle \vec{\mathbf{X}}, s_{\mathsf{c}} \rangle} > t\right] \Pr\left[F_{\langle \vec{\mathbf{X}}, s_{\mathsf{c}} \rangle} > t\right]. \end{aligned} \tag{3.34}$$

To bound the first term in Inequality (3.34), notice that $\boldsymbol{X}_{-1}^{h_1(k)-1} = \boldsymbol{x}_{-1}$ for some

$k \in \{1, ..., k_1^*(t)\}$ implies that $\boldsymbol{X}_{-s_c(\tau)}^{\tau-1} = \boldsymbol{x}_{-s_c(\tau)}$ for some $\tau \leq t$. Therefore

$$\begin{aligned}
\Pr\left[\bigcup_{\tau=1}^{t}\left\{\boldsymbol{X}_{-s_c(\tau)}^{\tau-1} = \boldsymbol{x}_{-s_c(\tau)}\right\} \,\middle|\, F_{\langle\vec{\mathbf{X}},s_c\rangle} > t\right] &\geq \Pr\left[\bigcup_{k=1}^{k_1^*(t)}\left\{\boldsymbol{X}_{-1}^{h_1(k)-1} = \boldsymbol{x}_{-1}\right\} \,\middle|\, F_{\langle\vec{\mathbf{X}},s_c\rangle} > t\right] \\
&= \sum_{k=1}^{k_1^*(t)} \Pr\left[\boldsymbol{X}_{-1}^{h_1(k)-1} = \boldsymbol{x}_{-1} \,\middle|\, F_{\langle\vec{\mathbf{X}},s_c\rangle} > t\right] \\
&= \sum_{k=1}^{k_1^*(t)} \frac{1}{m^{n-1}} \\
&= \left\lceil \frac{t}{n} \right\rceil \frac{1}{m^{n-1}}. \qquad (3.35)
\end{aligned}$$

The first summation follows from the fact that since all the environments for player 1 are distinct, the events in the union are mutually exclusive. The next step follows from the fact that our process is invariant under symmetry. So for any $k \in \{1, ..., k_1^*(t)\}$ and for all $\boldsymbol{x}_{-1}$ and $\boldsymbol{y}_{-1}$,

$$\Pr\left[\boldsymbol{X}_{-1}^{h_1(k)-1} = \boldsymbol{x}_{-1} \mid F_{\langle\vec{\mathbf{X}},s_c\rangle} > t\right] = \Pr\left[\boldsymbol{X}_{-1}^{h_1(k)-1} = \boldsymbol{y}_{-1} \mid F_{\langle\vec{\mathbf{X}},s_c\rangle} > t\right] \qquad (3.36)$$

which implies that

$$\Pr\left[\boldsymbol{X}_{-1}^{h_1(k)-1} = \boldsymbol{x}_{-1} \mid F_{\langle\vec{\mathbf{X}},s_c\rangle} > t\right] = \frac{1}{m^{n-1}}. \qquad (3.37)$$

This follows from $k_1^*(t) = 1 + \lfloor \frac{t-1}{n} \rfloor = \lceil \frac{t}{n} \rceil$.

To bound the second term in Inequality (3.34), notice that if for each $i \in [n]$ the environments $\boldsymbol{X}_{-i}^{h_i(1)-1}, \boldsymbol{X}_{-i}^{h_i(2)-1}, ..., \boldsymbol{X}_{-i}^{h_i(k_i^*(t))-1}$ are all distinct then $F_{\langle\vec{\mathbf{X}},s_c\rangle} > t$. Therefore

$$\begin{aligned}
\Pr[F_{\langle\vec{\mathbf{X}},s_c\rangle} > t] &\geq \Pr\left[\bigcap_{i\in[n]}\left\{\boldsymbol{X}_{-i}^{h_i(k)-1} \text{ for } k \in \{1, ..., k_i^*(t)\} \text{ are all distinct}\right\}\right] \\
&= 1 - \Pr\left[\bigcup_{i\in[n]}\left\{\boldsymbol{X}_{-i}^{h_i(k)-1} \text{ for } k \in \{1, ..., k_i^*(t)\} \text{ are not all distinct}\right\}\right] \\
&\geq 1 - \sum_{i\in[n]} \Pr\left[\boldsymbol{X}_{-i}^{h_i(k)-1} \text{ for } k \in \{1, ..., k_i^*(t)\} \text{ are not all distinct}\right] \\
&\geq 1 - \frac{n}{m^{n-1}} \frac{(\lceil \frac{t}{n} \rceil)^2}{2}. \qquad (3.38)
\end{aligned}$$

The final step follows from Lemma 3.

Gathering the results in Equation (3.31), Equation (3.32), Inequality (3.33), Inequality (3.35), and Inequality (3.38) together yields the desired conclusion. □

3.2.2 Two-Player Games

When there are $n = 2$ players, we can provide more detailed results on both game duration and the probability of convergence to equilibrium. In particular, we obtain the probability that the clockwork best-response dynamic converges to a pure Nash equilibrium in 2-player, m-action games. Notably, we argue that in the 2-player, m-action case, this probability is the same for random playing sequences. Furthermore, we show that this probability is asymptotically $\sqrt{\pi/m}$ when m is large (where $\pi \approx 3.14$).

The following theorem provides an *exact* expression for the probability that the clockwork sequence best-response dynamic converges to a $2k$-cycle in period t.[8]

Theorem 3. *For any $k \in \{1, ..., m\}$ and $t \in \{1, ..., 2(m-k+1)\}$, the probability that the s_c-best-response dynamic on $G_{2,m}$ converges to a $2k$-cycle in period t is given by*

$$\frac{1}{m} \prod_{i=1}^{t+2(k-1)} \left(1 - \frac{1}{m}\left\lfloor \frac{i}{2} \right\rfloor\right). \tag{3.39}$$

For any $k \in \{1, ..., m\}$, if $t > 2(m-k+1)$ then the probability that the clockwork sequence best-response dynamic reaches a $2k$-cycle is zero.[9] Setting $k = 1$ in Formula (3.39) yields the exact probability that the clockwork sequence best-response dynamic on $G_{2,m}$ converges to a PNE in period t.

Proof of Theorem 3. For a path generated by Algorithm 4 (the coupled dynamic) to reach neither a pure Nash equilibrium nor a best-response cycle by period t it must be the case that, by period $t+1$ (inclusive), no player encounters an environment that they have seen before, and the action taken by player $s_c(t+1)$ in period $t+1$ must not repeat any of the environments encountered by period t by player $s_c(t)$. To put it differently, for each $i \in \{1, 2\}$ it must be the case that the environments

$$\left(\boldsymbol{X}_{-i}^{h_i(1)-1}, \boldsymbol{X}_{-i}^{h_i(2)-1}, ..., \boldsymbol{X}_{-i}^{h_i(k_i^*(t+1))-1}\right) \tag{3.40}$$

are all distinct, and the action $X_{s_c(t+1)}^{t+1}$ taken by player $s_c(t+1)$ in period $t+1$ is distinct from each of the environments encountered by period t by player $s_c(t)$. It follows that the probability that the clockwork sequence best-response dynamic converges to neither a pure Nash equilibrium nor a best-response cycle by period t for $t \in [2m]$ is

$$\prod_{i=1}^{t} \left(1 - \frac{1}{m}\left\lfloor \frac{i}{2} \right\rfloor\right). \tag{3.41}$$

[8] See Pangallo et al. (2019) for an exact formula giving the probability of existence of cycles of any length in 2-player games.

[9] Since the number of action profiles is finite, a path cannot reach a $2k$-cycle only after $2(m-k+1)$ periods.

For the path to reach a pure Nash equilibrium in period t, for each $i \in \{1,2\}$, it must be the case that the environments

$$\left(\boldsymbol{X}_{-i}^{h_i(1)-1}, \boldsymbol{X}_{-i}^{h_i(2)-1}, ..., \boldsymbol{X}_{-i}^{h_i(k_i^*(t+1))-1}\right) \tag{3.42}$$

are all distinct, and the action $X_{s_{\mathsf{c}}(t+1)}^{t+1}$ taken by player $s_{\mathsf{c}}(t+1)$ in period $t+1$ is equal to the environment $\boldsymbol{X}_{-s_{\mathsf{c}}(t)}^{t-1}$ encountered by player $s_{\mathsf{c}}(t)$ in period t. Therefore, the probability that the clockwork sequence best-response dynamic converges to a pure Nash equilibrium in period $t \in [2m]$ is

$$\frac{1}{m}\prod_{i=1}^{t}\left(1-\frac{1}{m}\left\lfloor\frac{i}{2}\right\rfloor\right). \tag{3.43}$$

More generally, for the path to reach a $2k$-cycle in period t, for each $i \in \{1,2\}$ it must be the case that the environments

$$\left(\boldsymbol{X}_{-i}^{h_i(1)-1}, \boldsymbol{X}_{-i}^{h_i(2)-1}, ..., \boldsymbol{X}_{-i}^{h_i(k_i^*(t+2k-1))-1}\right) \tag{3.44}$$

are all distinct, and the action $X_{s_{\mathsf{c}}(t+2k-1)}^{t+2k-1}$ taken by player $s_{\mathsf{c}}(t+2k-1)$ in period $t+2k-1$ is equal to the environment $\boldsymbol{X}_{-s_{\mathsf{c}}(t)}^{t-1}$ encountered by player $s_{\mathsf{c}}(t)$ in period t. Therefore, the probability that the clockwork sequence best-response dynamic converges to a $2k$-cycle for $k \in [m]$ in period $t \in [2(m-k+1)]$ is

$$\frac{1}{m}\prod_{i=1}^{t+2(k-1)}\left(1-\frac{1}{m}\left\lfloor\frac{i}{2}\right\rfloor\right). \tag{3.45}$$

This is precisely Formula (3.39) given in the statement of Theorem 3. Notice that setting $k=1$ in Formula (3.45) recovers Formula (3.43). □

Example (Illustration of Formula (3.41) to Formula (3.45))**.** To illustrate our results, we schematically map out all the possible paths of the clockwork sequence best-response dynamic in 2-player, m-action games in the tree shown in Figure 3.4. In period 0, the initial profile is arbitrarily set to $(1,1)$.

To reach neither a pure Nash equilibrium nor a best-response cycle by period 3, for example, we must travel down the tree along the sequence $(1,1)$, $(1,1)$, $(1,2)$, $(2,2)$, $(2,3)$ or the sequence $(1,1)$, $(2,1)$, $(2,2)$, $(1,2)$, $(1,3)$. Let us look at the first sequence more closely. Up to (and including) period 4, the environments encountered by each player respectively are all distinct. Player 2 must select action 3 in period 4 to ensure that the sequence does not end up revisiting a previously encountered environment. The probability of traveling along $(1,1)$, $(1,1)$, $(1,2)$, $(2,2)$, $(2,3)$ or $(1,1)$, $(2,1)$, $(2,2)$, $(1,2)$, $(1,3)$ is given by Formula (3.41) with $t=3$.

Let us consider the probability of reaching a pure Nash equilibrium in period 4. In Figure 3.4, we must travel along the sequence $(1,1)$, $(1,1)$, $(1,2)$, $(2,2)$,

$(2,3)$, $(2,3)$ or the sequence $(1,1)$, $(2,1)$, $(2,2)$, $(1,2)$, $(1,3)$, $(1,3)$. Let us look at the first sequence more closely. Up to (and including) period 5, the environments encountered by each player respectively are all distinct. Player 1 must select action 2 in period 5 (the environment encountered by player 2 in period 4) to ensure that the pure Nash equilibrium $(2,3)$ was reached in period 4. The probability of traveling along $(1,1)$, $(1,1)$, $(1,2)$, $(2,2)$, $(2,3)$, $(2,3)$ or $(1,1)$, $(2,1)$, $(2,2)$, $(1,2)$, $(1,3)$, $(1,3)$ is given by Formula (3.43) with $t=4$.

Let us consider the probability of reaching a 4-cycle in period 2. In Figure 3.4 we must travel down the tree along the sequence $(1,1)$, $(1,1)$, $(1,2)$, $(2,2)$, $(2,3)$, $(1,3)$ or the sequence $(1,1)$, $(2,1)$, $(2,2)$, $(1,2)$, $(1,3)$, $(2,3)$. Let us look at the first sequence more closely. The action profile is $(1,2)$ in period 2, and this is the first action profile of our 4-cycle. A further $2k-1$ period must pass before looping back to $(1,2)$. Up to (and including) period 5, the environments encountered by each player respectively are all distinct. To "close" the cycle, player 1 must select action 1 in period 5 (this was the environment encountered by player 2 in period 2). The probability of traveling along $(1,1)$, $(1,1)$, $(1,2)$, $(2,2)$, $(2,3)$, $(1,3)$ or $(1,1)$, $(2,1)$, $(2,2)$, $(1,2)$, $(1,3)$, $(2,3)$ is given by Formula (3.45) with $t=k=2$.

As a straightforward corollary of Theorem 3, we can sum Formula (3.39) over all $t \in \{1,...,2(m-k+1)\}$ to obtain the exact probability that the clockwork sequence best-response dynamic converges to a $2k$-cycle:

Corollary 2. *The probability that the s_c-best-response dynamic on $G_{2,m}$ converges to a $2k$-cycle is given by*

$$\frac{1}{m}\sum_{t=1}^{2(m-k+1)}\prod_{i=1}^{t+2(k-1)}\left(1-\frac{1}{m}\left\lfloor\frac{i}{2}\right\rfloor\right). \tag{3.46}$$

Setting $k=1$ in Formula (3.46) yields the probability that the clockwork sequence best-response dynamic on $G_{2,m}$ converges to a PNE. To better understand the behaviour of Formula (3.46), we now study the expression when m is large. To do so, let $\lfloor x \rfloor$ be the floor operator of x and let $\Phi(\cdot)$ denote the standard normal cumulative distribution function:

$$\Phi(x) := \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x}\exp\left(-\frac{z^2}{2}\right)\mathrm{d}z. \tag{3.47}$$

The asymptotics of Formula (3.46) are given below.

Proposition 2. *If $k=o(m^{2/3})$ then, as $m\to\infty$, Formula (3.46) is asymptotically*[10]

$$2\sqrt{\frac{\pi}{m}}\left(1-\Phi\left(\frac{2k-1}{\sqrt{2m}}\right)\right). \tag{3.48}$$

[10] $f(n)=o(g(n))$ denotes $f(n)/g(n)\to 0$ as $n\to\infty$.

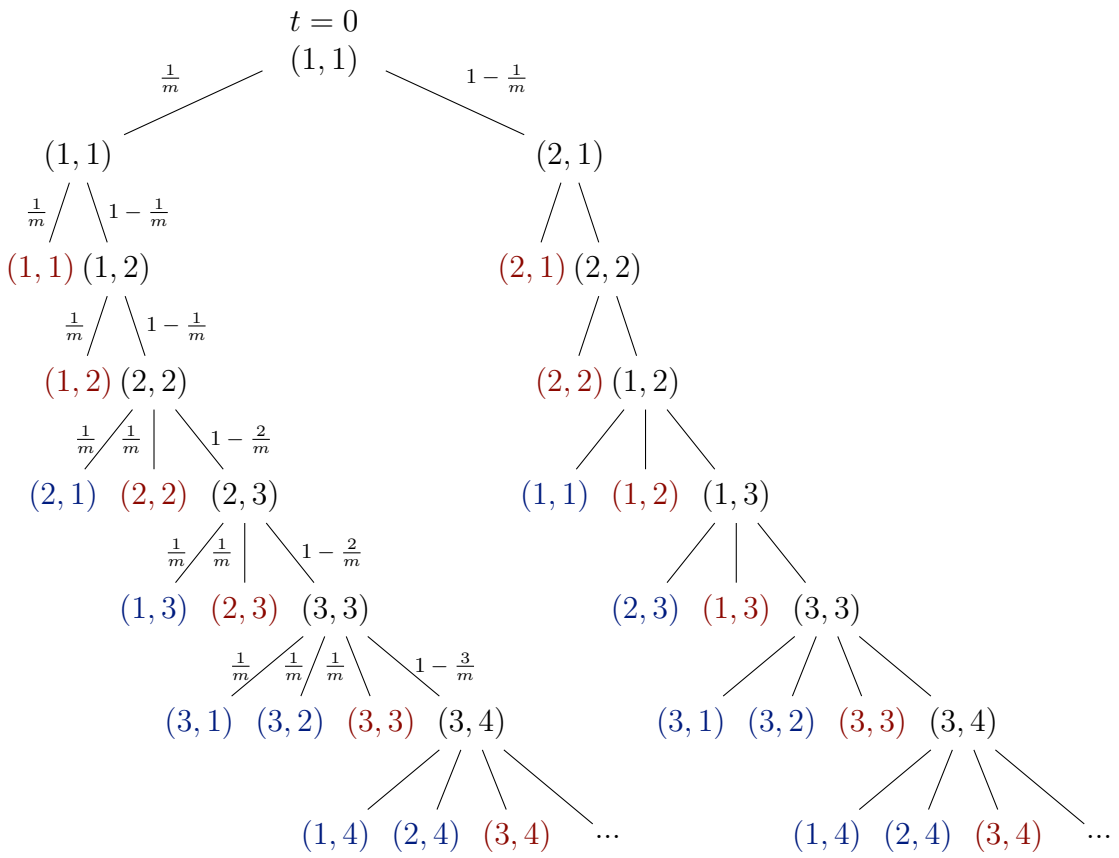


Figure 3.4: Illustration of possible paths for 2-player, m-action games. We arbitrarily set the initial action profile to be $(1, 1)$ in period 0. In period 1, player 1 either plays action 1 (left branch) or another action (right branch), which we arbitrarily call action 2. Player 2 then responds in period 2, and so on. All red leaves are Nash equilibria, and all blue leaves are profiles that belong to best-response cycles.

And if $k = o(\sqrt{m})$ then, as $m \to \infty$, Formula (3.46) is asymptotically $\sqrt{\pi/m}$.

We provide the proof at the end of this section.

The asymptotics given in Proposition 2 help us better understand the clockwork sequence's behaviour in large 2-player games. First, the probability of convergence to a PNE, corresponding to setting $k = 1$, goes to zero when $m \to \infty$. Second, short cycles all have about the same probability. Indeed, for $k = o(\sqrt{m})$ the probability is asymptotically $\sqrt{\pi/m}$. Finally, it is very unlikely that the best-response dynamic converges to a very long cycle: if $k/\sqrt{m} \to \infty$, then the probability that the dynamic converges to a cycle of length at least $2k$ tends to 0.[11]

[11]When $k = o(\sqrt{m})$, the argument of $\Phi(\cdot)$ goes to zero. Since $\Phi(0) = 1/2$ we have that the convergence probability goes to $\sqrt{\pi/m}$ which is independent of k. If, instead, $k/\sqrt{m} \to \infty$ then the argument of $\Phi(\cdot)$ grows large and since $\Phi(\infty) = 1$, the convergence probability goes zero. Our proof of Proposition 2 allows us to derive the asymptotics only for the range $k = o(m^{2/3})$, but this is sufficient to obtain some insight into the behaviour of Formula (3.46).

Our results are illustrated in Figure 3.5, which shows the probability of convergence to cycles of a given length as calculated from the exact formula in Theorem 3. The panels on the left and in the centre plot for $m = 5$ and $m = 10$ show that the probability of convergence is lower for long cycles than for short cycles. The Panel on the right plots Formula (3.46) for $m = 1000$. Since $31 \approx \sqrt{1000}$, we placed a vertical line at $2k = 62$. The probability of convergence is more or less uniform up to that point, which is consistent with our observations above.

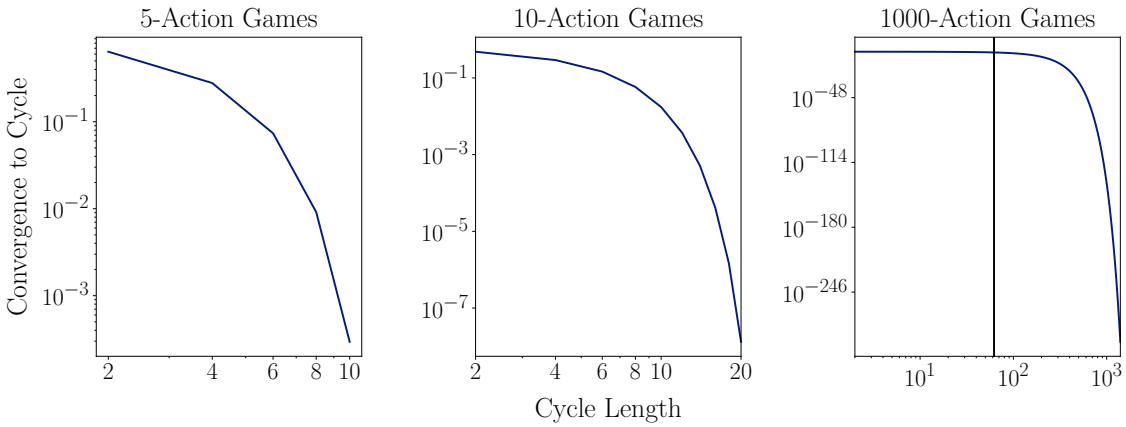


Figure 3.5: Frequency of convergence of the clockwork sequence best-response dynamic to $2k$-cycles in 2-player games, using the exact formula from Theorem 3. The vertical black line on the right corresponds to Formula (3.46), showing that the probability of convergence is roughly uniform up to $2 \cdot \sqrt{1000}$.

We now compare the behaviour of the clockwork sequence best-response dynamic in 2-player games with that of the random sequence best-response dynamic in 2-player games. Specifically, we observe that the probability of convergence to a PNE is the same for both clockwork and random playing sequences in 2-player games. The reason is that, under the random playing sequence, players' actions do not change whenever the sequence asks the same player to play several times in a row. The profiles that are thus visited along the path are the same under both playing sequences, which induces the same probability of convergence to equilibrium. However, the game duration will differ since the random playing sequence introduces delays. The result below enables us to determine the game duration more precisely.

Theorem 4. *The probability that the s_c-best-response dynamic on $G_{2,m}$ does not converge to a PNE or a best-response cycle before period $x\sqrt{2m}$ is $\exp(-x^2/2)$ as $m \to \infty$.*

Proof of Theorem 4. Recall the definitions of $h_i(k)$ and $k_i^*(t)$ preceding Lemma 3, see Equation 3.18. On a path $\langle \vec{\mathbf{X}}, s_c \rangle$ generated by Algorithm 3, the environments

that player $i \in \{1,2\}$ encounters on her turns between (and including) periods 1 and t are given in the sequence

$$\left(\boldsymbol{X}_{-i}^{h_i(1)-1}, \boldsymbol{X}_{-i}^{h_i(2)-1}, ..., \boldsymbol{X}_{-i}^{h_i(k_i^*(t))-1}\right). \tag{3.49}$$

To prove Theorem 4, we now work out the asymptotic behaviour of Equation (3.41). Note that Equation (3.41) can be written as

$$\prod_{i=1}^{t}\left(1-\frac{1}{m}\left\lfloor\frac{i}{2}\right\rfloor\right) = \begin{cases} \frac{m!^2}{\left(\left(m-\frac{t+1}{2}\right)!\right)^2 m^{t+1}} & \text{if } t \text{ is odd} \\ \binom{m-\frac{t}{2}}{m} \frac{(m!)^2}{\left(\left(m-\frac{t}{2}\right)!\right)^2 m^t} & \text{if } t \text{ is even} \end{cases}. \tag{3.50}$$

Using Stirling's formula[12], we obtain that[13]

$$\frac{(m!)^2}{\left(\left(m-\frac{t+1}{2}\right)!\right)^2 m^{t+1}} \sim \left(\frac{m-\frac{t+1}{2}}{m}\right)^{t-2m} \exp\left(-(t+1)\right) \tag{3.51}$$

and

$$\binom{m-\frac{t}{2}}{m} \frac{(m!)^2}{\left(\left(m-\frac{t}{2}\right)!\right)^2 m^t} \sim \left(\frac{m-\frac{t}{2}}{m}\right)^{t-2m} \exp(-t) \tag{3.52}$$

whenever $m - t \to \infty$. Taking a logarithm of the last expression,[14]

$$\begin{aligned} &-t + (t-2m)\ln\left(1-\frac{1}{m}\frac{t}{2}\right) \\ =&-t + (t-2m)\left(-\frac{1}{2}\frac{t}{m} - \frac{1}{8}\frac{t^2}{m^2} + O\left(\frac{t^3}{m^3}\right)\right) \\ =&-\frac{1}{4}\frac{t^2}{m} + O\left(\frac{t^3}{m^2}\right). \end{aligned} \tag{3.53}$$

Provided that $t = o(m^{2/3})$, the second term goes to zero and therefore Approximation (3.52) behaves asymptotically like $\exp\left(\frac{-t^2}{4m}\right)$. An identical argument shows that, under the same conditions, Approximation (3.51) is also asymptotically $\exp\left(\frac{-t^2}{4m}\right)$. Hence,

$$\prod_{i=1}^{t}\left(1-\frac{1}{m}\left\lfloor\frac{i}{2}\right\rfloor\right) \sim \exp\left(\frac{-t^2}{4m}\right). \tag{3.54}$$

This completes the proof of Theorem 4. Note that Approximation (3.54) holds uniformly in the range $\left[1, o(m^{2/3})\right]$. □

[12] As $n \to \infty$, $n! \sim \sqrt{2\pi n} \cdot n^n \exp(-n)$.
[13] $f(n) \sim g(n)$ denotes $f(n)/g(n) \to 1$ as $n \to \infty$.
[14] $f(n) = O(g(n))$ if there is $M > 0$ and N such that $|f(n)| \leq Mg(n)$ for all $n \geq N$.

This result shows that the clockwork sequence best-response dynamic in 2-player games will likely converge to a PNE or a best-response cycle within $\sqrt{2m}$ periods when m is large. The game duration for the random playing sequence should be greater than for the clockwork playing sequence by a factor of 2. The reason is that, under the clockwork playing sequence, the players alternate at the tick of each period, whereas, under the random playing sequence, the number of periods that it takes for the playing sequence to turn to the other player is Geometric($\frac{1}{2}$). Thus, the random playing sequence can be considered a slowing down of the clockwork playing sequence, in which the expected time to play the next step is 2.

We conclude this section with the proof of Proposition 2, which uses Approximation (3.54).

Proof of Proposition 2. Recall the definitions of $h_i(k)$ and $k_i^*(t)$ preceding Lemma 3, see Equation 3.18. On a path $\langle \vec{\mathbf{X}}, s_{\mathsf{c}} \rangle$ generated by Algorithm 3, the environments that player $i \in \{1, 2\}$ encounters on her turns between (and including) periods 1 and t are given in the sequence

$$\left(\boldsymbol{X}_{-i}^{h_i(1)-1}, \boldsymbol{X}_{-i}^{h_i(2)-1}, ..., \boldsymbol{X}_{-i}^{h_i(k_i^*(t))-1} \right). \tag{3.55}$$

To prove Proposition 2, we now turn to the asymptotic behaviour of Formula (3.39). Let $T = T(m)$ satisfy $T = o(m^{2/3})$ and $k = o(T)$. We assume that $T \geq \frac{m^{2/3}}{\ln(m)}$ so that T is not too small, and we split the summation in Formula (3.39) into two ranges: $t \leq T$ and $t > T$. Since the Approximation (3.54) holds uniformly in our first range, we have

$$\frac{1}{m} \sum_{t=1}^{T} \prod_{i=1}^{t+2(k-1)} \left(1 - \frac{1}{m} \left\lfloor \frac{i}{2} \right\rfloor \right) \sim \frac{1}{m} \sum_{t=1}^{T} \exp\left(-\frac{(t + 2(k-1))^2}{4m} \right). \tag{3.56}$$

We now approximate the summation on the right-hand side with an integral. Firstly, note that

$$\begin{aligned} \frac{1}{m} \int_1^{T+1} \exp\left(-\frac{(t + 2(k-1))^2}{4m} \right) \mathrm{d}t &= \sqrt{\frac{2}{m}} \int_{\frac{2k-1}{\sqrt{2m}}}^{\frac{T+1+2(k-1)}{\sqrt{2m}}} \exp\left(-\frac{x^2}{2} \right) \mathrm{d}x \\ &\sim \sqrt{\frac{2}{m}} \int_{\frac{2k-1}{\sqrt{2m}}}^{\infty} \exp\left(-\frac{x^2}{2} \right) \mathrm{d}x \\ &= 2\sqrt{\frac{\pi}{m}} \left(1 - \Phi\left(\frac{2k-1}{\sqrt{2m}} \right) \right), \end{aligned} \tag{3.57}$$

where the first step uses the transformation

$$x = \frac{t + 2(k-1)}{\sqrt{2m}}. \tag{3.58}$$

Furthermore,

$$\frac{1}{m}\int_0^1 \exp\left(-\frac{(t+2(k-1))^2}{4m}\right)\mathrm{d}t \leq \frac{1}{m}, \tag{3.59}$$

which goes to zero faster than Approximation (3.57). Since

$$\int_1^{T+1} f(t)\,\mathrm{d}t \leq \sum_{t=1}^{T} f(t) \leq \int_0^{T} f(t)\,\mathrm{d}t \leq \int_1^{T+1} f(t)\,\mathrm{d}t + \int_0^1 f(t)\,\mathrm{d}t \tag{3.60}$$

for any positive and decreasing function $f(\cdot)$, it follows that

$$\frac{1}{m}\sum_{t=1}^{T}\exp\left(-\frac{(t+2(k-1))^2}{4m}\right) \sim 2\sqrt{\frac{\pi}{m}}\left(1-\Phi\left(\frac{2k-1}{\sqrt{2m}}\right)\right). \tag{3.61}$$

We must show that the summation of Formula (3.39) over the second range is negligible. Since $\exp(x) \geq 1+x$ and $\lfloor x \rfloor > x-1$ for all x, we obtain the following upper bound:

$$\begin{aligned}
&\frac{1}{m}\sum_{t=T+1}^{2(m-k+1)}\prod_{i=1}^{t+2(k-1)}\left(1-\frac{1}{m}\left\lfloor\frac{i}{2}\right\rfloor\right)\\
&\leq \frac{1}{m}\sum_{t=T+1}^{2(m-k+1)}\prod_{i=1}^{T+1+2(k-1)}\left(1-\frac{1}{m}\left\lfloor\frac{i}{2}\right\rfloor\right)\\
&\leq \frac{1}{m}\sum_{t=T+1}^{2(m-k+1)}\exp\left(-\frac{1}{m}\sum_{i=1}^{T+1+2(k-1)}\left(\frac{i}{2}-1\right)\right)\\
&\leq \frac{2m-2k-T+1}{m}\exp\left(-\frac{1}{4m}\left(T+2(k-1)-2\right)^2\right).
\end{aligned} \tag{3.62}$$

This expression is small compared to the other half of the sum. □

3.2.3 Games on Networks

Up to this point, we have assumed that, in general, the chosen actions of each player impact the payoffs of each other player. We now allow this assumption to loosen, meaning that for all randomly generated games, there is a specific player whose payoffs do not depend on the behaviour of another specified player. We formalise this by means of a directed *network of players* with n nodes, that is imposed on a game $g_{n,m}$, and which determines which players' actions impact which player's payoffs. If there is no edge from player p to player q, then the payoffs of q do not depend on the actions of player p. We provide convergence results for clockwork best-response dynamics on bipartite graphs, which include games on stars, rings with an even number of vertices, and hypercubes. Our approach is similar to the proof of Theorem 3, but we will make purely simple combinatorial arguments here. Our key results are obtained from a recurrence relation:

Lemma 4. *The solution to the recurrence relation*

$$\begin{aligned} f(k+1) &= h(k)f(k) + g(k) \\ f(1) &= f_1 \end{aligned} \tag{3.63}$$

is given by

$$f(k) = \prod_{i=1}^{k-1} h(i) \left(f_1 + \sum_{j=1}^{k-1} \frac{g(j)}{\prod_{i=1}^{j} h(i)} \right). \tag{3.64}$$

Proof of Lemma 4. We have

$$\frac{f(k+1)}{\prod_{i=1}^{k} h(i)} - \frac{h(k)f(k)}{\prod_{i=1}^{k} h(i)} = \frac{g(k)}{\prod_{i=1}^{k} h(i)}. \tag{3.65}$$

Define

$$F(k) = \frac{f(k)}{\prod_{i=1}^{k-1} h(i)} \tag{3.66}$$

then

$$\sum_{j=1}^{k-1} (F(j+1) - F(j)) = F(k) - F(1) = \sum_{j=1}^{k-1} \frac{g(j)}{\prod_{i=1}^{j} h(i)} \tag{3.67}$$

which implies

$$\frac{f(k)}{\prod_{i=1}^{k-1} h(i)} = f_1 + \sum_{j=1}^{k-1} \frac{g(j)}{\prod_{i=1}^{j} h(i)} \tag{3.68}$$

from where Equation (3.64) follows. □

We now provide our key result for the frequency of convergence on bipartite networks:

Theorem 5. *The probability that the clockwork best-response dynamic converges on $G_{n,m}$ where the underlying network of players is bipartite of two sets A, B of sizes a and b, $n = a + b$, $a \geq b$ is given by*

$$\left(\prod_{i=1}^{m^b-1} h(i) \right) \left(\frac{1}{m^a} + \frac{m^a - m^b}{m^n} + \sum_{j=1}^{m^b-1} \frac{g(j)}{\prod_{i=1}^{j} h(i)} h(i) \right) \tag{3.69}$$

where $h(r) = \frac{(m^a - m^b + r)r}{m^n}$ *and* $g(r) = \frac{1}{m^a} + \frac{m^a - m^b + r}{m^n}$.

Proof of Theorem 5. Assume, without loss of generality, that the players in set A move first. Similar to the reasoning for the proof of Theorem 3, convergence is achieved with probability $\frac{1}{m^a}$, otherwise $\frac{m^a-1}{m^a}$. Then, the players in set B move and converge with probability $\frac{1}{m^b}$, otherwise $\frac{m^b-1}{m^b}$. This continues with the numerators in the probabilities decreasing by 1 until they vanish.

The probability of convergence is summed up down the playing tree as

$$\frac{1}{m^a} + \frac{m^a - 1}{m^a} \left(\frac{1}{m^b} + \frac{m^b - 1}{m^b} \left(\frac{1}{m^a} + \frac{m^a - 2}{m^a} \left(\frac{1}{m^b} + \frac{m^b - 2}{m^b} (\dots) \right) \right) \right) \tag{3.70}$$

which can be rewritten as $f(m^b)$ where $f(r)$ is given by

$$\begin{aligned} f(r+1) &= \frac{1}{m^a} + \frac{m^a - m^b + r}{m^a}\left(\frac{1}{m^b} + \frac{r}{m^b} f(r)\right) \\ &= \frac{1}{m^a} + \frac{m^a - m^b + r}{m^{a+b}} + \frac{\left(m^a - m^b + r\right) r}{m^{a+b}} f(r). \end{aligned} \tag{3.71}$$

Solving the recurrence relation yields

$$f(r) = \left(\prod_{i=1}^{r-1} h(i)\right)\left(\frac{1}{m^a} + \frac{m^a - m^b}{m^{a+b}} + \sum_{j=1}^{r-1} \frac{g(j)}{\prod_{i=1}^{j}} h(i)\right) \tag{3.72}$$

where h and g are defined as above. Using the fact that the probability of convergence is $f\left(m^b\right)$ yields the desired result. □

In particular, the corollaries below state the probability of convergence under star networks and bipartite networks of equally-sized sets, which include rings with an even number of vertices and hypercubes.

Corollary 3. *The probability that the s_c-best-response dynamic converges on $G_{n,m}$ where the underlying network of players is a star, is given by*

$$\left(\prod_{i=1}^{m-1} h(i)\right)\left(\frac{1}{m} + \sum_{j=1}^{m-1} \frac{g(j)}{\prod_{k=1}^{j} h(k)}\right) \tag{3.73}$$

where $g(j) = \frac{1}{m} + \frac{j}{m^n}$ and $h(k) = \frac{k(m^{n-1} - m + k)}{m^n}$.

Corollary 4. *The probability that the s_c-best-response dynamic converges on $G_{n,m}$ where the underlying network of players is bipartite of two sets of equal sizes k, is given by*

$$\left(\prod_{i=1}^{m^k-1} h(i)\right)\left(\frac{1}{m^k} + \sum_{j=1}^{m^k-1} \frac{g(j)}{\prod_{i=1}^{j} h(i)}\right) \tag{3.74}$$

where $h(r) = \frac{r^2}{m^{2k}}$ and $g(r) = \frac{m^k + r}{m^{2k}}$.

Figure 3.6 summarises the results of this section by showing the convergence rate of clockwork best-response dynamics for 4-player games under different networks: fully connected (see Section 3.2.1), fully connected bipartite, and star networks. As shown earlier, the frequency of convergence decreases with the number of actions, and is higher on networks with lower density. For numerical results on the frequency of convergence on random networks with varying numbers of edges, see Section 3.3.1.3.

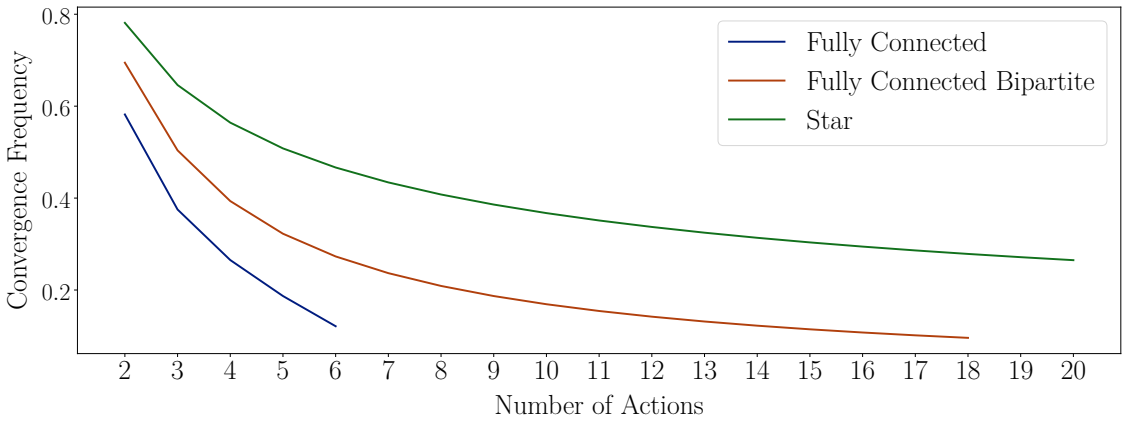


Figure 3.6: Frequency of convergence of clockwork best-response dynamics on 4-player games with varying numbers of actions and under different networks. The frequency of convergence decreases with the number of actions and with higher network density.

3.3 Simulation Results

This section summarises our results obtained from numerical simulations. We run the different types of best-response dynamics and compare their behaviour to the other learning rules: reinforcement learning, fictitious play, and replicator dynamics.

For our theoretical results, we limited ourselves to analysing best-response dynamics in random games where the payoffs are drawn independently across players and action profiles. For our simulations, we allow the payoffs to be correlated across players (Goldberg et al., 1968; Stanford, 1999; Berg and Weigt, 1999; Rinott and Scarsini, 2000; Galla and Farmer, 2013; Sanders et al., 2018). To do so, at initialisation, for each action profile $\boldsymbol{a}$ we draw the vector $U(\boldsymbol{a}) = (U_1(\boldsymbol{a}), \ldots, U_n(\boldsymbol{a}))$ at random from a multivariate normal distribution with mean zero, unit variance, and covariance matrix with 1s on the diagonal and $\frac{\Gamma}{n-1}$ on the other entries. So $\Gamma \in [-1, n-1]$ parametrises the degree to which payoffs are correlated. If $\Gamma = 0$, the payoffs are chosen independently, which recovers the case for which we derived our theoretical results. At one extreme, if $\Gamma = n-1$, then at each action profile, all players receive the same payoff, and at the other extreme, if $\Gamma = -1$, we have a zero-sum game. More generally, when $\Gamma > 0$ the game is "cooperative" and when $\Gamma < 0$ the game is "competitive" (Rinott and Scarsini, 2000). We find that negatively correlated payoffs make convergence less likely and that positively correlated payoffs make convergence more likely.

In all simulations, for each combination of n, m, and Γ, we draw 500 games and simulate 5000 time steps starting from randomly chosen initial conditions. In Section 3.3.1, we first simulate the clockwork sequence best-response dynamic in random n-player, m-action games with independent payoffs (i.e. $\Gamma = 0$). This allows us to verify our theoretical results. Then, we compare the behaviour of the

clockwork sequence best-response dynamic against the best-response dynamic under random and simultaneous updating for different values of n, m, and Γ. In Section 3.3.2, we compare the behaviour of the best-response dynamic (with clockwork, random, and simultaneous updating) against other learning rules (reinforcement learning, fictitious play, and replicator dynamics) for different values of n, m, and Γ. We also vary the structure of the network of players, and show by numerical simulations that the network's density lowers the convergence probability, similar to increasing the number of players or actions.

3.3.1 Simulation of Best-Response Dynamics

This section presents the results of our simulations on best-response dynamics. In Section 3.3.1.1, we focus on clockwork best-response dynamics, numerically confirming our theoretical results from Section 3.2. Section 3.3.1.2 compares the observed frequencies of convergence across different types of best-response dynamics — clockwork, random, and simultaneous. Section 3.3.1.3 presents our findings on games played on networks (see Section 3.2.3 for our theoretical results), while Section 3.3.1.4 discusses the speed of convergence under various settings, in particular using the periodic best-response dynamic (see Section 3.2.2 for theoretical results).

3.3.1.1 Clockwork Best-Response Dynamics

We start by simulating the clockwork sequence best-response dynamic in n-player, m-action games with independent payoffs ($\Gamma = 0$). We find good agreement with our main theoretical results. We show that Corollary 1, which states that the asymptotic probability of convergence to a PNE is, up to a polynomial factor, of order $1/\sqrt{m^{n-1}}$, is also reflected in our simulations for relatively small values of m and n.

The blue markers in Figure 3.7 show the frequency of convergence to a PNE in our simulations for different numbers of players n and numbers of actions m. The error bars correspond to one standard deviation calculated over the average rate of convergence of 100 sampled games for each ensemble. The frequency of convergence to a PNE decreases as the number of players or actions increases. The solid black line in the top panel is the analytical probability of convergence to a PNE in 2-player games, calculated using Formula (3.46). Up to sampling noise, our analytical results align perfectly with the numerical simulations.

Additionally, we explore the validity of our asymptotic results in games with a small number of players and actions. Specifically, Figure 3.7 allows us to verify Corollary 1; namely, that the frequency of convergence to a PNE in a n-player, m-action game is roughly the same as in a 2-player, m^{n-1} action game. While in Section 3.2.1, this result was only proved asymptotically ($mn \to \infty$), we investigate

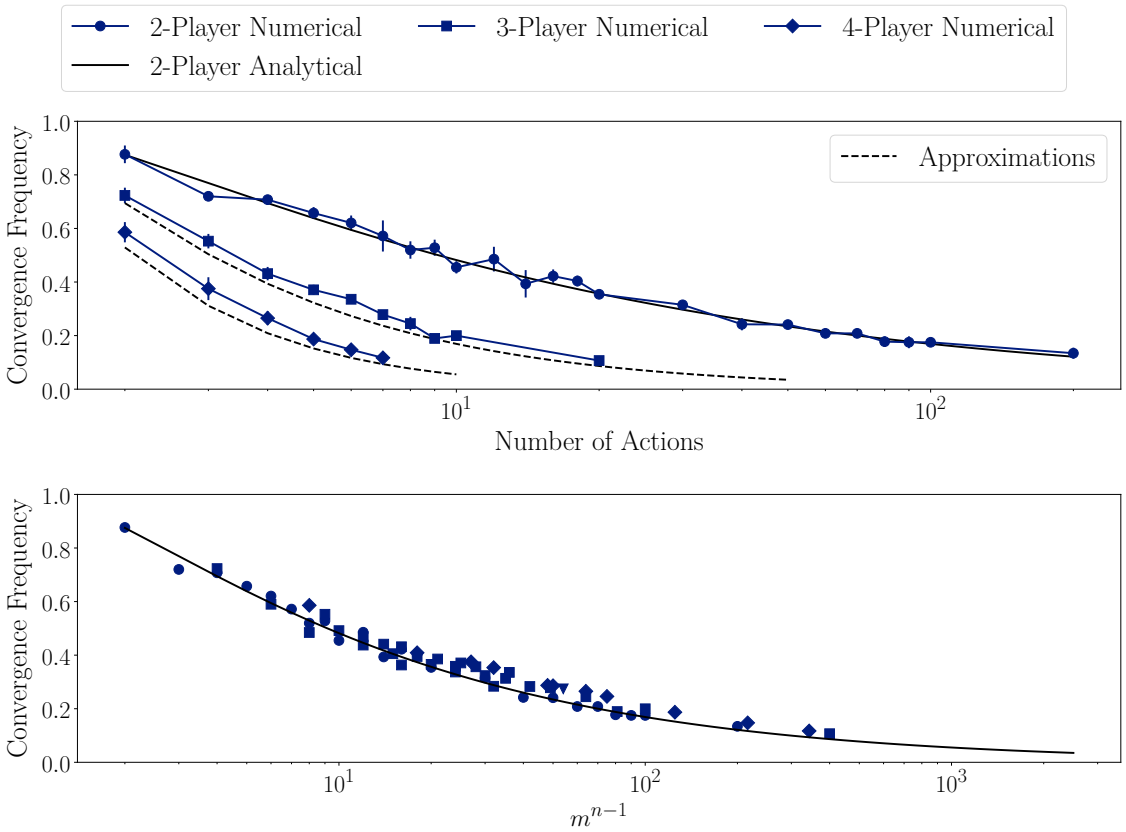


Figure 3.7: Frequency of convergence under the clockwork best-response dynamic, verifying Formula (3.46) for 2-player games (solid black line) by comparing it to numerical results for 2-, 3-, and 4-player games in the top panel. We also verify Corollary 1 by showing that the analytical approximations (dashed black lines) match the numerical results. In the bottom panel, the horizontal axis is rescaled according to Corollary 1, showing that the numerical data collapses even for a small number of players and actions.

the extent to which the result holds for small values of m and n. Using Corollary 1, we approximate the frequency of convergence to a PNE in a n-player, m-action game by replacing the number of actions m in Formula (3.46) by m^{n-1}, which is the equivalent number of actions in the corresponding 2-player game. We plot these approximate frequencies as dashed lines in the top panel of Figure 3.7. As can be seen, there is a good match between the approximation and our simulation results, particularly when the number of actions m is relatively large. The bottom panel of Figure 3.7 gives us another way to illustrate Corollary 1. Here, we rescale the number of actions for n-player games to match the number of actions of the equivalent 2-player game. After rescaling, a point corresponding to m actions in an n-player game is moved on the horizontal axis to several actions given by m^{n-1}. For example, the point giving the convergence frequency for 4-player 10-action games is translated to the right to the horizontal coordinate corresponding to 10^3 actions in a 2-player game. The rescaled markers all lie relatively close to the black line, corresponding

to the analytical probability of convergence to a PNE in 2-player, m-action games.

We continue with results on games with correlated payoffs ($\Gamma \neq 0$). Figure 3.8 shows the frequency of convergence to a PNE under clockwork best-response dynamics. The solid line corresponds to $\Gamma = 0$, the darker dashed lines correspond to positive correlations, $\Gamma = \frac{n-1}{5}, 2\frac{n-1}{5}, \ldots, (n-1)$, and the lighter dashed lines correspond to negative correlations, $\Gamma = -0.2, -0.4, \ldots, -1$. If the payoffs are positively correlated, for any value of n and m, convergence tends to be more likely than for independent payoffs. The reason is that under positively correlated payoffs (especially if the correlation is very strong) there is a proliferation of pure Nash equilibria (Goldberg et al., 1968; Stanford, 1999; Berg and Weigt, 1999; Rinott and Scarsini, 2000). The best-response dynamic is therefore very likely to converge to one of these equilibria.[15] If the payoffs are negatively correlated, for any value of n and m, convergence to a PNE tends to be less likely than with independent payoffs. When $\Gamma \approx -1$ we essentially generate zero-sum games (Goldberg et al., 1968; Rinott and Scarsini, 2000). Such games do not have pure Nash equilibria, so clockwork best-response dynamics can not converge to a Nash equilibrium.[16]

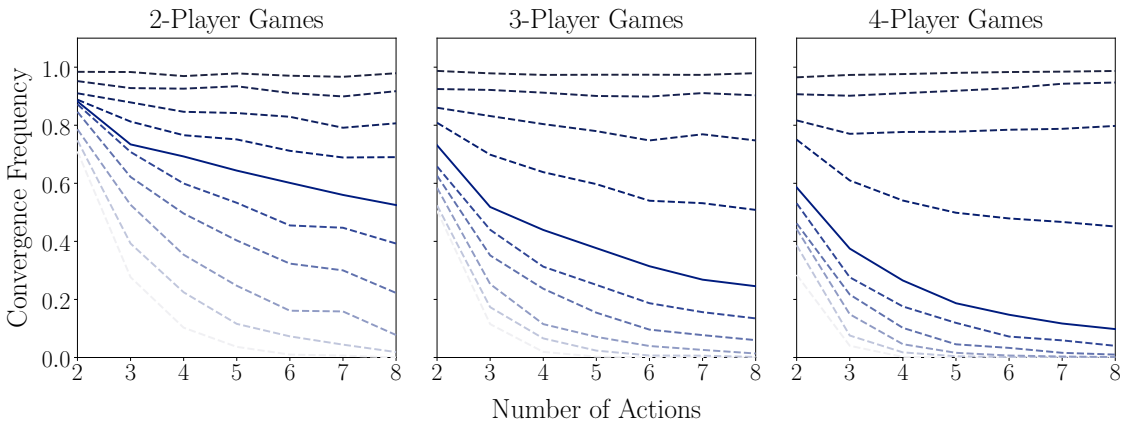


Figure 3.8: Frequency of convergence to PNE under clockwork best-response dynamics. The solid line corresponds to $\Gamma = 0$, the darker dashed lines to $\Gamma = \frac{n-1}{5}, 2\frac{n-1}{5}, \ldots, (n-1)$, and the lighter dashed lines to $\Gamma = -0.2, -0.4, \ldots, -1$. The frequency of convergence decreases with the number of players and the number of actions, but increases with the level of correlation.

[15]This does not occur under 2-player simultaneous best-response dynamics when payoffs are highly correlated. For instance, consider a 2-player, 2-action coordination game (in which payoffs are positively correlated). If players start at one of the two action profiles that are not pure Nash equilibria, they keep bouncing between them forever.

[16]And neither can any version of best-response dynamics, see Figure 3.14.

3.3.1.2 Comparison of Different Best-Response Dynamics

We now turn to compare the convergence frequency of different types of best-response dynamics (clockwork, random, and simultaneous). Figure 3.9 shows the frequency of convergence of 2-, 3- and 4-player games and varying numbers of actions when payoffs are independent. The error bars correspond to one standard deviation to the mean, calculated over average rates of convergence of 100 games each. The frequency of convergence to a PNE decreases in n and m for the clockwork playing sequence and for simultaneous updating. The random playing sequence is different. When there are only $n = 2$ players, the random playing sequence has the same convergence probability as the clockwork playing sequence — as argued in Section 3.2.2. Amiet et al. (2021) proved that the random sequence best-response dynamic always converges to a PNE if there is one when $m = 2$ and $n \to \infty$. As argued in Section 3.2.1, this gives us an unconditional probability of convergence of $1 - 1/e \approx 63\%$. Our simulations show that the result of Amiet et al. (2021) also appears to hold for games with more than two actions provided $n > 2$. The random sequence best-response dynamic almost always converges to a PNE in games with a PNE (even for relatively small values of n and m). The simultaneous best-response dynamic is less likely to converge than the other two across all observed game ensembles.

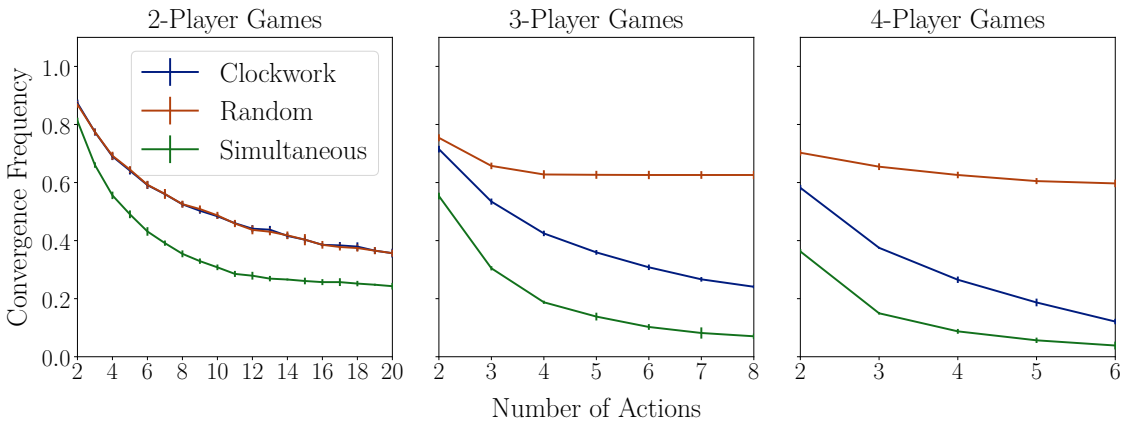


Figure 3.9: Frequency of convergence to pure Nash equilibria under clockwork, random, and simultaneous best-response dynamics for different numbers of players and actions when payoffs are independent. The frequency of convergence decreases with the number of actions and players for clockwork and simultaneous updating. The random sequence best-response dynamic almost always converges to a PNE in games with a PNE (Amiet et al., 2021). Simultaneous best-response dynamics are less likely to converge than the other two across all observed game ensembles.

3.3.1.3 Games on Networks

Figure 3.10 illustrates the rate of convergence for the clockwork best-response dynamics as the underlying network increases in edges, ranging from 1 up to $\frac{n(n-1)}{2}$. The figure highlights the frequency of convergence across a broad set of randomly generated connected networks with varying numbers of players and edges, represented by the solid line. Additionally, we overlay the frequency of convergence for the corresponding regular network. Our focus on connected networks is due to the fact that non-connected networks can be treated as a set of smaller, independent games. Thus, analysing connected networks offers a clearer understanding of how increasing connectivity affects convergence. We observe that as the number of edges in the network increases, the frequency of convergence decreases. Moreover, the convergence frequency curves tend to converge to a common shape, regardless of the number of players, indicating a robust relationship between network structure and convergence.

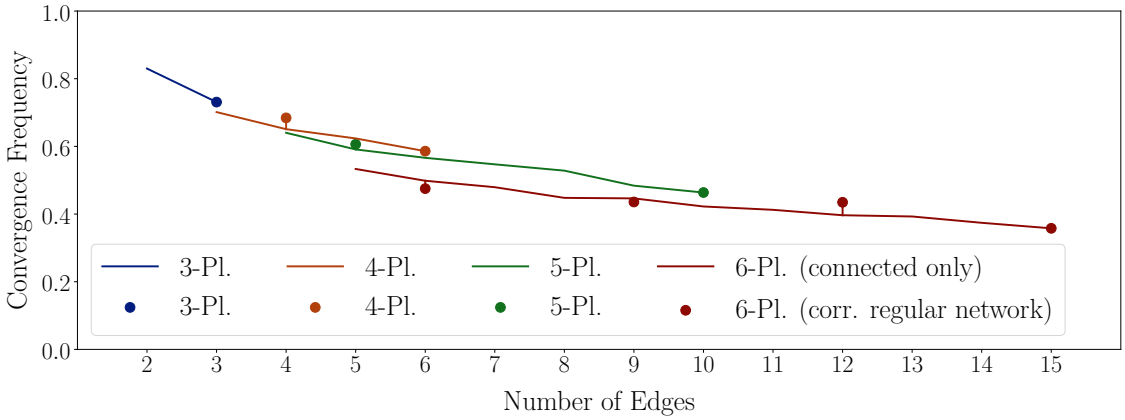


Figure 3.10: Frequency of convergence of clockwork best-response dynamics in 2-action games with uncorrelated payoffs and with varying number of edges in the network of players. We also highlight the frequency of convergence for the corresponding regular network. The frequency of convergence decreases with the number of edges.

3.3.1.4 The Speed of Convergence

We employ the p-periodic best-response dynamic (see Section 2.3.4) to study the time it takes to reach convergence. Figure 3.11 plots the frequency of convergence to a pure Nash equilibrium against the length of the random subsequence (namely $p-n$) for p-periodic best-response dynamics in 3-player games with 2 or 3 actions per player. As expected, the probability of convergence to a Nash equilibrium for p-periodic playing sequences increases in p. The error bars correspond to one standard deviation to the mean, calculated over the average convergence frequency of 100 games.

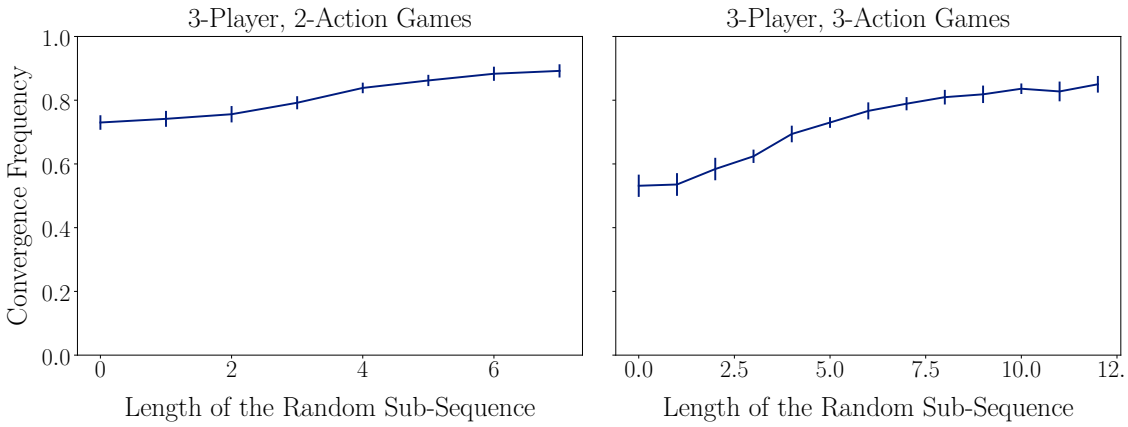


Figure 3.11: Frequency of convergence under the p-periodic best-response dynamic in 3-player games for 2 actions and 3 actions. The probability of convergence to a Nash equilibrium for p-periodic playing sequences increases in p.

Figure 3.12 shows, for varying lengths of the random subsequence, the distribution of the number of time steps required to reach a pure Nash equilibrium, conditional on converging, for p-periodic best-response dynamics in 3-player games with 2 and 3 actions per player. Notably, given that the system converges to a Nash equilibrium, the average number of time steps required to reach equilibrium increases with p. This observation aligns with the findings of Durand and Gaujal (2016), which indicate that, on average, convergence to equilibrium in potential games is quicker under the clockwork playing sequence compared to other playing sequences.

In our analysis, we find that the speed of convergence (conditional on reaching equilibrium) is slower for playing sequences with a higher proportion of random elements (i.e., a larger $p - n$) across the game space. This presents an interesting trade-off between the likelihood of finding equilibria and the speed of convergence to those equilibria, which may offer fertile ground for future research.

3.3.2 Comparison to Other Learning Rules

In this section, we compare the behaviour of best-response dynamics (under clockwork, random, and simultaneous updating) with three classic learning rules: Bush-Mosteller reinforcement learning (Bush and Mosteller, 1953), fictitious play (Brown, 1951; Robinson, 1951), and replicator dynamics (Smith et al., 1992). The main takeaway from our simulation results is that, compared to random and simultaneous updating, the convergence properties of the clockwork best-response dynamic align most closely with those of the three learning rules. This suggests that the clockwork dynamic may offer a more accurate model of learning behaviour in games where these classical learning rules are relevant.

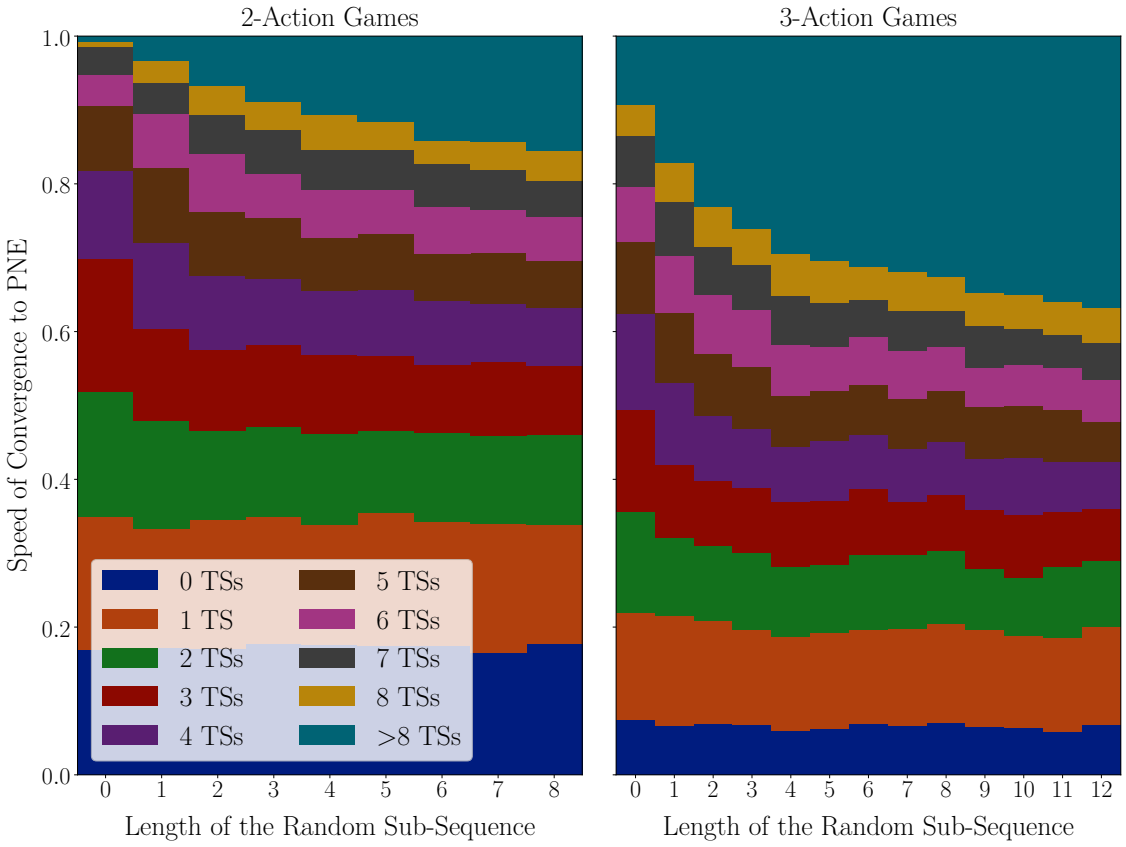


Figure 3.12: Conditional speed of convergence to a PNE for p-periodic best-response dynamics in 3-player games. The average number of time steps required to reach equilibrium increases with p.

We compare the probability of convergence of best-response dynamics (under clockwork, random, and simultaneous updating) against each of the three learning rules. In Figure 3.13, we do this for uncorrelated payoffs so $\Gamma = 0$, $n = 2$ and $n = 3$ players, and for various numbers of actions m. In Figure 3.14, we allow Γ to vary but fix the number of actions to $m = 5$. Finally, in Figure 3.15 we allow n, m, and Γ to all vary. In our comparisons, we distinguish between cases where Bush-Mosteller learning and fictitious play converge to pure vs. mixed Nash equilibria (replicator dynamics only converge to pure Nash equilibria). In Figures 3.13 and 3.14, we represent the frequency of convergence to pure equilibria with a dashed line, while a solid line indicates convergence to any type of equilibrium. Thus, for any given number of actions, the vertical distance between the dashed and solid lines represents the frequency of convergence to mixed equilibria only.

As shown in Figure 3.13, the frequency of convergence to a PNE under the clockwork best-response dynamic closely aligns with the frequency of convergence to a PNE under the three learning rules. In games with three or more players, the random sequence best-response dynamic almost always converges to a PNE, provided

one exists, which explains why the orange line rapidly flattens at $1 - 1/e \approx 0.63$. In contrast, the frequency of convergence to pure (or mixed) Nash equilibria under the three learning rules declines as the number of actions increases. The best-response dynamic under simultaneous updating converges to a PNE far less frequently compared to the three learning rules.[17] When considering convergence to mixed Nash equilibria as well, we observe that the clockwork best-response dynamic converges less frequently than the three learning rules, particularly in the case of fictitious play with two players. However, as the number of actions increases, the clockwork dynamic more closely follows the trend in convergence to mixed or pure Nash equilibria observed in the learning rules, outperforming the random best-response dynamics when there are three or more players. Additionally, the frequency of convergence to equilibrium of the clockwork best-response dynamics is generally closer to that of the three learning rules than that of the simultaneous best-response dynamic.

In Figure 3.14, we fix the number of actions at $m = 5$ and vary Γ.[18] Both the clockwork and random sequence best-response dynamics track the other learning rules fairly well, though the clockwork sequence appears to outperform the random sequence. However, it is worth noting that the best-response dynamic under simultaneous updating converges far less frequently. This observation is consistent with our earlier finding that, even in the presence of many Nash equilibria (as is the case under strongly positively correlated payoffs), this version of the dynamic fails to converge as reliably as the others.

To directly measure the closeness of different types of best-response dynamics to the complex learning rules, Figure 3.15 presents scatter plots comparing the frequency of convergence to a PNE for the best-response dynamic (under each playing sequence) to the frequency of convergence to pure or mixed equilibria for each of the three learning rules. Each dot represents the convergence frequency for a specific combination of n, m, and Γ.[19] We show the Pearson correlation for each combination, finding that while the frequency of convergence to a PNE for the clockwork sequence best-response dynamic does not perfectly align with the frequency of convergence to pure or mixed equilibria for the three learning rules (see Figure 3.13 and Figure 3.14), it generally outperforms the other versions of the

[17]There is a discrepancy between the frequency of convergence to a PNE for the replicator dynamic and that for the clockwork best-response dynamic. This is primarily due to numerical limitations: the replicator dynamic has infinite memory, so a trajectory might hit the machine precision limit without reaching a PNE.

[18]Note that in the 2-player case with $\Gamma = -1$, the frequency of convergence to any equilibrium (pure or mixed) for fictitious play is close to one. This aligns with existing theoretical results on the convergence of fictitious play in 2-player zero-sum games (Fudenberg and Levine, 1998).

[19]The 418 combinations of n, m, and Γ considered are: $n = 2$ with $m = 2, ..., 15$; $n = 3$ with $m = 2, ..., 10$; $n = 4$ with $m = 2, ..., 8$; $n = 5$ with $m = 2, ..., 5$; $n = 6$ with $m = 2, 3$; and $n = 7$ with $m = 2, 3$, each for $\Gamma = -1, -0.8, ..., 0, \frac{n-1}{5}, 2\frac{n-1}{5}, ..., (n-1)$.

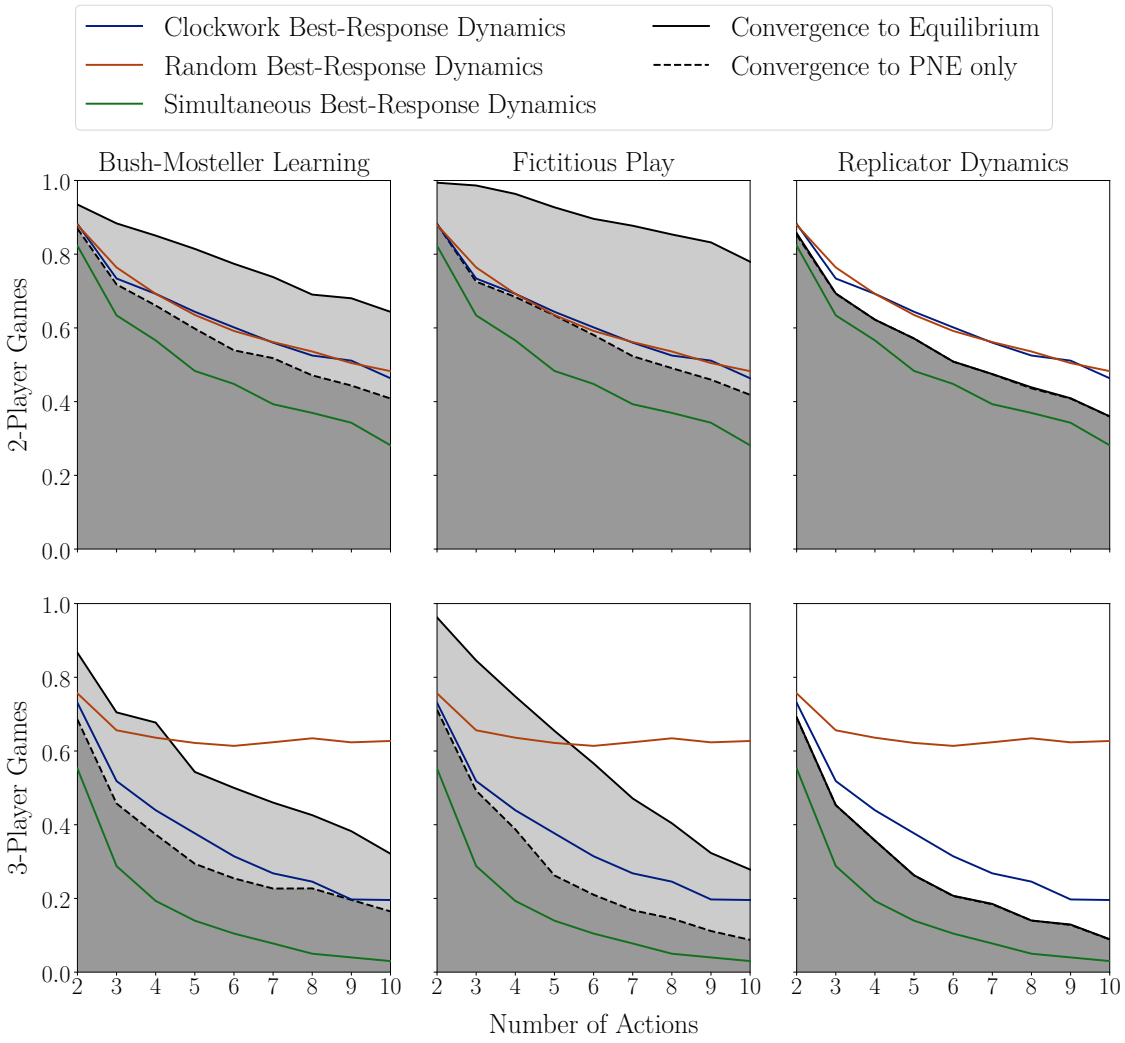


Figure 3.13: Frequency of convergence to PNE under clockwork, random, and simultaneous best-response dynamics compared to the frequency of convergence to (mixed and pure) Nash equilibria under Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics for games with independent payoffs. Except for multi-player random best-response dynamics, the convergence frequency decreases with the number of actions. The clockwork best-response dynamic closely tracks the convergence behaviour of complex learning dynamics.

best-response dynamic discussed here. Had we focused exclusively on convergence to pure equilibria for each of the learning rules only, the clockwork best-response dynamic would have resembled the outcomes of the three learning rules even more closely. The identity line is included as a reference.

To develop more intuition for why the clockwork best-response dynamic is a good predictor for the three complex learning rules in terms of convergence probability and evolution of play, we turn to qualitatively compare the dynamics instead of just the convergence probability. Figure 3.16 illustrates a best-response digraph

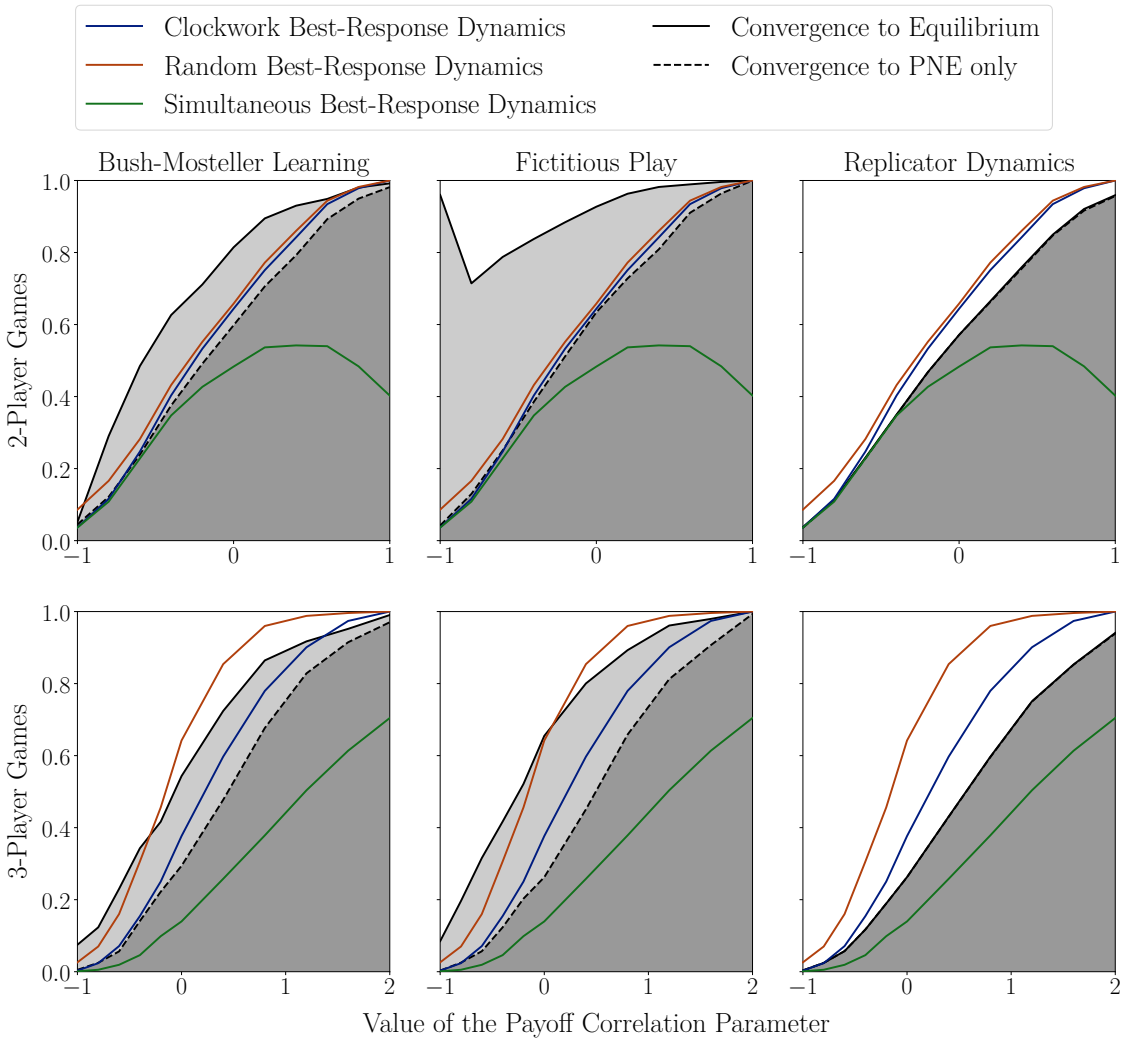


Figure 3.14: Frequency of convergence to PNE in 5-action games under clockwork, random, and simultaneous best-response dynamics compared to the frequency of convergence to (mixed and pure) Nash equilibria under the Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics, with varying levels of correlation Γ. The frequency of convergence increases with higher correlation. Both the clockwork and random sequence best-response dynamics track the other learning rules fairly well, the simultaneous best-response dynamic does not.

alongside the trajectories traced by Bush-Mosteller learning, fictitious play, and the replicator dynamic, starting from various initial conditions.[20] The trajectories illustrate the evolution of mixed strategy vectors for the learning rules, which generally follow the directions of the edges in the best-response digraph. These edges also guide the evolution of play in the clockwork best-response dynamic,

[20]The underlying game has binary payoffs of zero or one. The patterns exhibited by the learning algorithms, as noted by Pangallo et al. (2019), are fairly robust across different payoff values in 2-player games. However, altering payoff values would distort the trajectories in the action profile space.

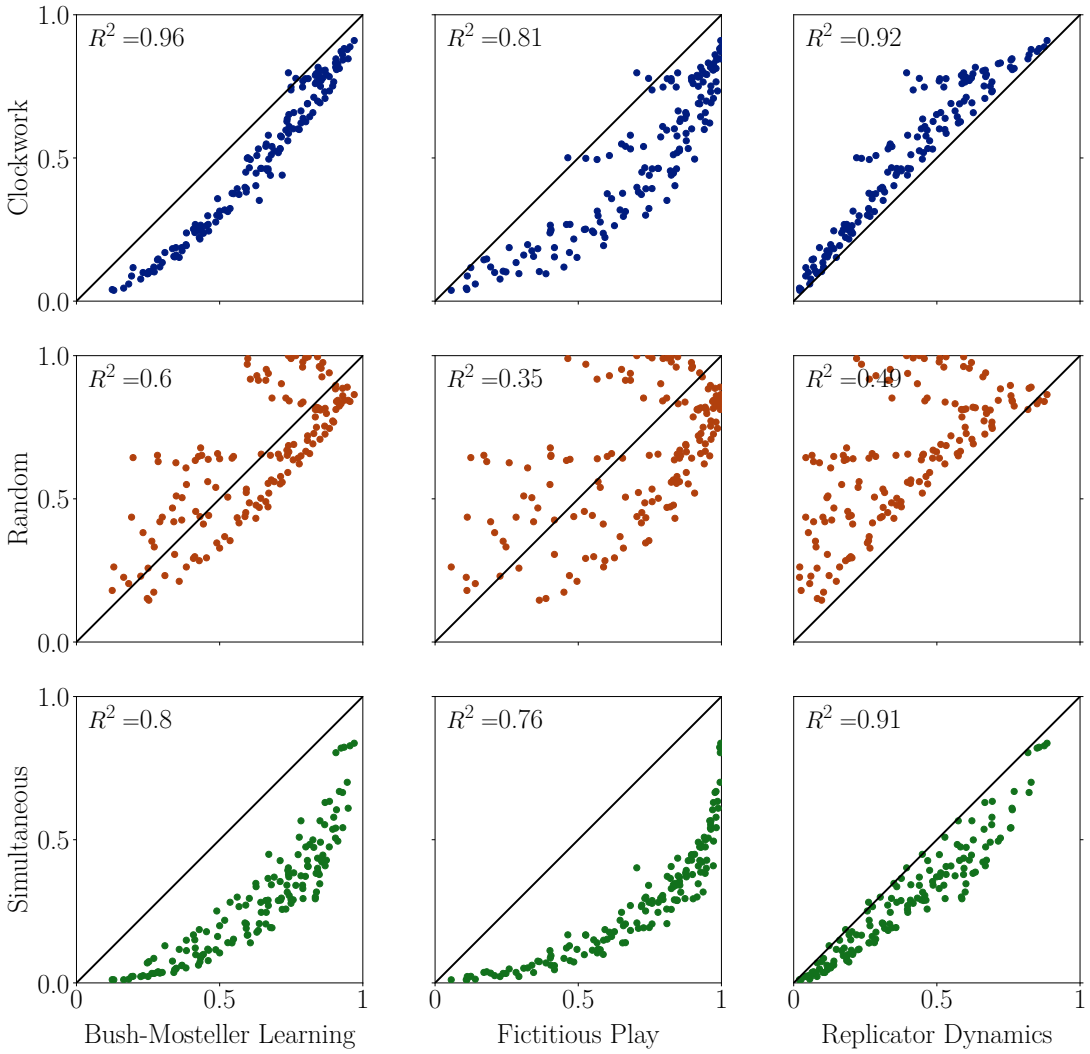
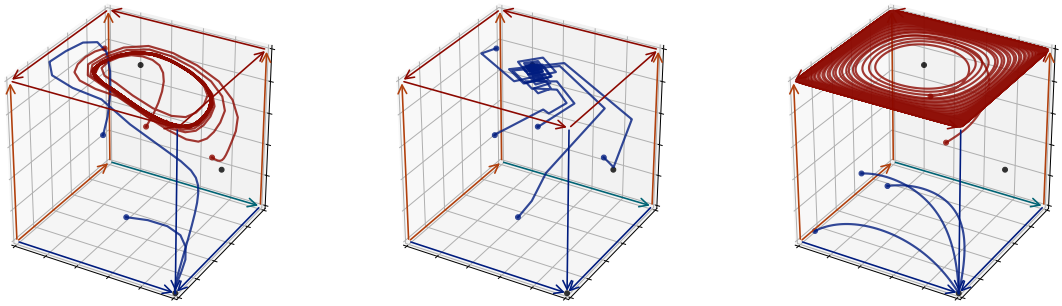


Figure 3.15: Frequency of convergence to PNE under clockwork, random, and simultaneous best-response dynamics against the frequency of convergence to (mixed or pure) Nash equilibria under Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics, for varying values of n, m, and Γ. We show the Pearson correlation for each combination, finding the clockwork best-response dynamics are the best predictor for convergence. The identity line is included as a reference.

but they do not influence the evolution of play in the random sequence. Over time, the random sequence best-response dynamic would eventually converge to the pure Nash equilibrium represented in this digraph. Thus, the paths traced by the clockwork sequence best-response dynamic tend to more closely resemble the trajectories followed by the three learning algorithms compared to those traced by the random sequence best-response dynamic. This similarity holds despite the fact that the three learning algorithms are typically defined with simultaneous

updating.[21] Our assessment of how "close" the paths of the clockwork and random sequence best-response dynamics are to those of the learning algorithms is based on observations across various games.[22] Moreover, our findings align with Pangallo et al. (2019), who found that the prevalence of $2k$-cycles serves as a strong indicator of the frequency of convergence to Nash equilibrium for learning algorithms in 2-player random games. Overall, our results suggest that, to the extent that the learning algorithms are consistent with human gameplay in randomly generated games, the clockwork best-response dynamic could serve as a reasonable first-order approximation for the evolution of play in such settings.



(a) Bush-Mosteller Learning **(b)** Fictitious Play **(c)** Replicator Dynamics

Figure 3.16: Trajectories of Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics imposed on best-response digraphs of different 3-player, 2-action games. Blue trajectories converge to Nash equilibria (pure or mixed); red trajectories do not. Blue arrows correspond to pure Nash equilibria, light blue arrows lead there, red arrows correspond to cycles, and orange arrows lead there.

Finally, we turn to games on networks and provide numerical results comparing the frequency of convergence of best-response dynamics to the frequency of convergence of the more complex learning rules. Figure 3.17 shows the frequency of convergence of clockwork best-response dynamics, Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics on fully connected networks, stars, rings, and wheels. We find that the frequency of convergence of the best-response dynamic stays close to the frequency of convergence of the other learning rules.

[21] The paths traced by the learning algorithms may exhibit features similar to those of the best-response dynamic under both clockwork and simultaneous updating. The extent to which the learning algorithms have "memory" will likely affect how closely the paths resemble the best-response dynamics under either updating method.

[22] Although we have not performed a quantitative analysis of "path closeness", we expect the comparison between clockwork and random sequence best-response dynamics to hold, particularly in large games.

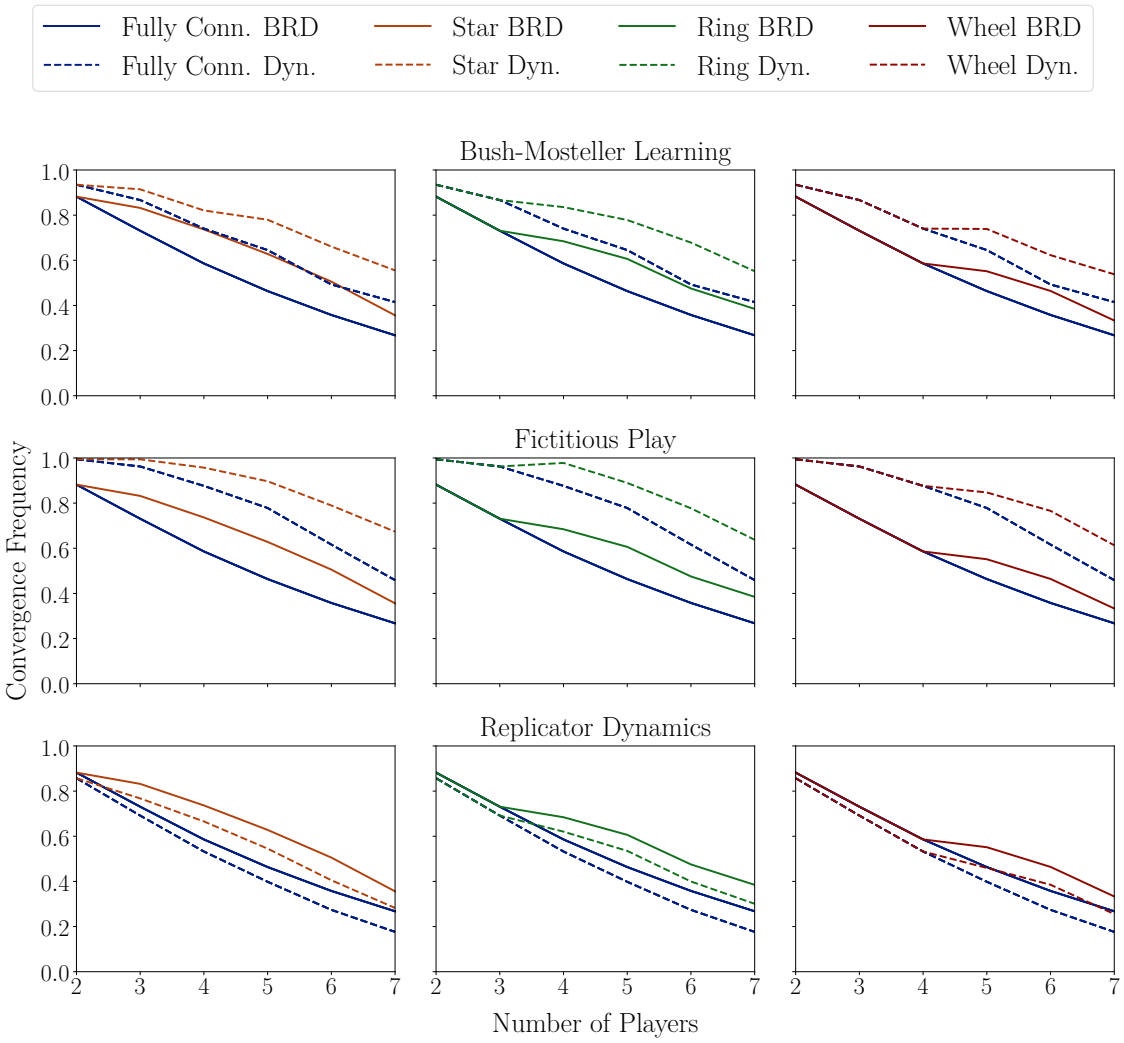


Figure 3.17: Frequency of convergence of clockwork best-response dynamics compared to Bush-Mosteller reinforcement learning, fictitious play, and replicator dynamics on 2-action games with varying numbers of players and different networks of players. The frequency of convergence of the best-response dynamic stays close to the frequency of convergence of the other learning rules.

3.4 Conclusion

This chapter investigated the convergence frequency of various learning dynamics in randomly generated games. Through a combination of theoretical analysis and numerical simulations, we demonstrated that the sequence in which players updated their actions had a significant impact on the likelihood of convergence. Specifically, we found that the clockwork sequence — where players updated in a fixed cyclic order — rarely converged to a Nash equilibrium when the number of players n and actions m was large. In contrast, random sequences almost always led to convergence when a Nash equilibrium existed. Our simulations further revealed that

the convergence properties of the clockwork sequence closely resembled those of popular learning rules, despite these rules typically being framed with simultaneous updating. Additionally, we presented new results on the impact of payoff correlations and network structures on convergence, showing that both factors influenced the probability of reaching an equilibrium. These findings suggest that structural features of the game, such as player interactions and payoff dependencies, play a crucial role in determining convergence behaviour.

Future research could explore the theoretical properties of simultaneous and periodic best-response dynamics; areas where our current methods do not apply. We also did not attempt to extend our theoretical results to games with correlated payoff or to games on networks.

In the next chapter, we turn to the related problem of determining the frequency with which specific game structures arise in randomly generated games. We will be using similar methods; in particular, digraph game representations and the clockwork best-response dynamic.

4

The Frequency of Games

We examine the likelihood of generating random games with specific structural characteristics. Section 4.2 reviews the existing literature on the distribution of Nash equilibria, the prevalence of various game classes, and the likelihood of drawing games with a specific attractor structure at random. Section 4.3 introduces our main analytical framework for studying attractor structures in multi-player games: the best-response n-partite digraph. Using this framework, we present an explicit formula for the frequency of a specific class of multi-player games in Section 4.4. Finally, Section 4.5 offers concluding remarks.

Contents

4.1 Introduction 81
4.2 The Frequency of Game Ensembles 82
4.2.1 The Distribution of Nash Equilibria 82
4.2.2 Classes of Games 83
4.2.3 The Structure of the Attractor 84
4.3 The Best-Response n-partite Digraph 87
4.4 The Frequency of Convergent Games 88
4.4.1 Convergent Games with a Unique PNE 89
4.4.2 Convergent Games with Multiple PNEs 94
4.5 Conclusion 96

4.1 Introduction

In this chapter, we study the likelihood of drawing games with specific structural properties at random. As in the previous chapter, we utilise the clockwork best-response dynamic, where players select their optimal strategy (facing the previous optimal strategies of their opponents) sequentially — player 1 moves first, followed by player 2, and so on, until it is player 1's turn once more. In the 2-player case, where the clockwork sequence is equivalent to the random playing sequence, it is possible to provide an explicit formula for the likelihood of games with a specific attractor structure. This is closely connected to Chapter 3 since the attractor structure summarises the likelihood of convergence or non-convergence of clockwork (or random) best-response dynamics on 2-player games. In the multi-player case, we provide general results on the likelihood of generating games on which the clockwork best-response dynamic converges. As before, those are games for which, starting from any initial strategy profile, no player modifies their strategy after a sufficiently large number of turns.

The frequency of generating games with certain properties at random can provide valuable insights into the predictability and stability of equilibria in economic systems. Such properties are often straightforward to observe in scenarios modelled by low-dimensional games, such as 2-player, 2-action setups. However, in more complex contexts — such as biological interactions (Gokhale and Traulsen, 2010; Duong, 2016), bidding behaviour (Cheng et al., 2020), interactions within supply chains (Babu and Mohan, 2018), trading behaviour in financial markets (Challet et al., 2005), or social behaviour during crises (e.g., the Covid-19 pandemic) — predictability and stability are less obvious.

Distinction to Chapter 3 We want to emphasise the distinction between this chapter, which focuses on the frequency of drawing certain classes of games at random, and Chapter 3, which discusses the frequency of convergence of different learning rules. This chapter specifically analyses one particular playing sequence in n-player games, represented as n-partite networks, and proves some unique results. In contrast, Chapter 3 considers multiple playing sequences and provides general results for each. While one playing sequence (clockwork) is analysed in both chapters, the results presented here remain distinct.

Specifically, in Section 4.4, we will provide an upper bound for the frequency of drawing a game without cycles if the number of players n and the number of actions m is finite, that is

$$\sum_{k=1}^{K} p_{n,m}^{k} \tag{4.1}$$

where K is the maximum number of pure Nash equilibria in n-player m-action games. This frequency is not equivalent to the frequency of convergence under clockwork playing. For instance, in the 2-player case, the frequency of games without cycles is

$$\sum_{k=1}^{m} \frac{2m-k}{m^{2k+2}(k-1)!} \left(\frac{m!}{(m-k)!} \right)^2 \tag{4.2}$$

where m is the number of actions. On the other hand, the frequency of convergence under a clockwork playing sequence (Equation (3.46) with $k = 1$ in Corollary 2) in the 2-player case is

$$\sum_{t=1}^{2m} \frac{1}{m} \prod_{i=1}^{t} \left(1 - \frac{1}{m} \left\lfloor \frac{i}{2} \right\rfloor \right). \tag{4.3}$$

The frequency of games without cycles is bounded by the frequency of convergence since games with both cycles and pure Nash equilibria have a convergence frequency above 0.

We start by reviewing the literature on the frequency of game ensembles, specifically on the distribution of Nash equilibria and on a subset of well-studied classes of games. We then turn to the likelihood of drawing games with a specific attractor structure at random.

4.2 The Frequency of Game Ensembles

This section discusses some key results on the likelihood of drawing games with specific structures at random. Specifically, Section 4.2.1 summarises asymptotic results on the distribution of Nash equilibria and Section 4.2.2 reviews popular classes of games.

4.2.1 The Distribution of Nash Equilibria

We start by noting that in random n-player m-action games, the probability that there is a pure Nash equilibrium is asymptotically $1 - \exp(-1) \approx 0.63$ as either n or m (or both) get large:

Proposition 3 (Rinott and Scarsini, 2000)**.**

$$\lim_{nm \to \infty} \Pr\left[\#\text{PNE}(G_{n,m}) \geq 1 \right] = 1 - \exp(-1). \tag{4.4}$$

In fact, Rinott and Scarsini (2000) prove a much stronger result: they characterise the asymptotic distribution of the number of pure Nash equilibria in random games, showing that $\#\text{PNE}(G_{n,m})$ is asymptotic Poisson(1) as $nm \to \infty$. The probability that a PNE exists in a random game has been studied by Goldberg et al. (1968) in the

2-player case and by Dresher (1970) in the n-player case as $m \to \infty$. Powers (1990) and Stanford (1995) noted that the distribution of the number of PNE approaches a Poisson(1) as $m \to \infty$. Building on Arratia et al. (1989), Rinott and Scarsini (2000) show that the Poisson(1) limit holds as $nm \to \infty$ (i.e. when m or n get large).

Concerning the use of best-response structures as a tool to study convergence frequencies, Pangallo et al. (2019) and Pei and Takahashi (2019) both obtained exact results for the frequency of one or more PNEs in the 2-player case.

4.2.2 Classes of Games

Closely related to the frequency of Nash equilibria, we now turn to the frequency of games that are a member of a certain well-defined class of games with notable structural properties. In this section, we introduce *constant-sum* games (Section 4.2.2.1), *dominance-solvable* games (Section 4.2.2.2), *potential* games (Section 4.2.2.3), and *coordination* games (Section 4.2.2.4).

4.2.2.1 Constant-Sum Games

A game is *constant-sum*, if the sum of the payoffs of all players is the same across all action profiles $\boldsymbol{a}$,

$$\sum_{i=1}^{n} u_i(\boldsymbol{a}) = c \tag{4.5}$$

where $c \in \mathbb{R}$ is that sum. In particular, when $c = 0$, the game is called *zero-sum*. Goldman (1957) considered zero-sum 2-player games and showed that the probability of having a PNE goes to zero as the number of actions grows.

4.2.2.2 Dominance-Solvable Games

An action of a player is *strictly dominated* by another action of the same player if the latter action yields a higher payoff in all environments. Rational players would, therefore, never choose a strictly dominated action. Iterated elimination of strictly dominated actions is the process of removing all strictly dominated actions. A game in which this process leaves only one action for each player is called a *dominance-solvable game*. The remaining action profile is the game's unique pure Nash equilibrium. Alon et al. (2021) studied the frequency of dominance-solvable games and derived an exact formula for the 2-player case.

4.2.2.3 Potential Games

A potential game (Monderer and Shapley, 1996) is a game that admits a potential function $v : [m]^n \to \mathbb{R}$, which maps all action profiles $\boldsymbol{a}$ to real numbers. The key feature of a potential game is that differences in payoffs can be entirely expressed through the potential function: for all players $i \in [n]$, all corresponding actions $a_i, a_i' \in [m]$, and all environments $\boldsymbol{a}_{-i}$,

$$u_i(a_i, \boldsymbol{a}_{-i}) - u_i(a_i', \boldsymbol{a}_{-i}) = v(a_i, \boldsymbol{a}_{-i}) - v(a_i', \boldsymbol{a}_{-i}). \tag{4.6}$$

This means any change in a player's utility from switching actions is mirrored precisely by a change in the potential function. This property ensures that the strategic incentives of the players are aligned with the potential function, making the game more tractable for analysis. Potential games can have multiple Nash equilibria. Collevecchio et al. (2024) study the distribution of the sizes of the basins of attraction of PNEs in 2-player potential games as the number of actions goes to infinity.

4.2.2.4 Coordination Games

A game is a *coordination game* (Foster and Young, 1998) if the players' actions can be indexed such that all action profiles where all players select actions with the same index are strict pure Nash equilibria. This implies that all strategy profiles along the main diagonal correspond to pure Nash equilibria. In addition to the pure Nash equilibria, coordination games often exhibit several mixed Nash equilibria.

4.2.3 The Structure of the Attractor

This section provides an overview of previous analyses of the structure of the attractor of 2-player games, drawing from Pangallo et al. (2019) and our work in Wiese (2020). We start by introducing the formalism for the structure of the attractor by Pangallo et al. (2019) and, similar to the section above, provide some key results on the likelihood of drawing certain games.

In the 2-player case, Pangallo et al. (2019) characterise the set of attractors of best-response dynamics in a given payoff matrix as the *best-response vector*

$$\boldsymbol{v} = (v_m, \ldots, v_1) \tag{4.7}$$

where v_1 is the number of fixed points, v_2 is the number of cycles spanning two actions, etc. It is important to note that the structure of the attractor can be exactly described by the best-response vector. For instance, in Figure 1.2, the two games on the left have the best-response vector $\boldsymbol{v} = (0, 0, 3)$, the middle games

have $\boldsymbol{v} = (0, 1, 0)$, and the two games on the right have $\boldsymbol{v} = (1, 0, 0)$. Note that an attractor described by the best-response vector $\boldsymbol{v}$ takes up

$$A = \sum_{k=1}^{m} v_k k \leq m \tag{4.8}$$

of each player's actions, and the cycles alone take up

$$C = A - v_1. \tag{4.9}$$

Then, the quantity

$$F(\boldsymbol{v}) = \frac{v_1}{v_1 + C} \tag{4.10}$$

measures the relative number of actions on attractors that are part of fixed points versus cycles, which also corresponds to the relative size of the basins of attraction of fixed points under clockwork best-response dynamics, as averaged over all best-response digraphs[1] with the same best-response vector.[2] For instance, Figure 4.1 shows all possible best-response digraphs corresponding to $\boldsymbol{v} = (0, 0, 3)$. Since all actions are part of pure Nash equilibria, $F(\boldsymbol{v}) = 1$. Figure 4.2 shows all possible best-response digraphs corresponding to $\boldsymbol{v} = (1, 0, 0)$. All actions of both players are part of the 3-cycle, so $F(\boldsymbol{v}) = 0$.

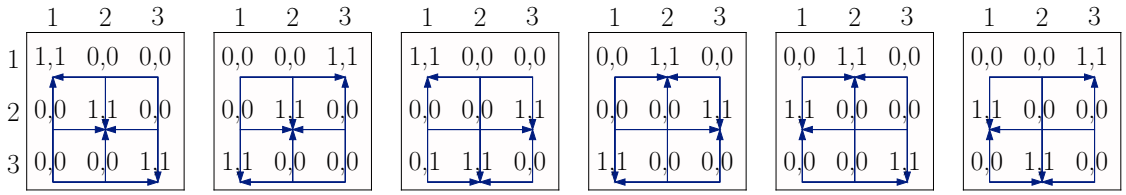


Figure 4.1: All possible best-response digraphs in 2-player, 3-action games with three pure Nash equilibria. The value on the left of each action profile represents the payoff to the first (row) player, while the value on the right represents the payoff to the second (column) player. All actions are part of Nash equilibria.

Let $p(\boldsymbol{v})$ denote the likelihood of drawing a game with the best-response vector $\boldsymbol{v}$ at random. The key result of Pangallo et al. (2019) is their development of a recursive expression for $p(\boldsymbol{v})$, which allows us to write the probability that the clockwork best-response dynamic converges to a PNE (see Theorem 3 for $k = 1$ in Chapter 3) as a weighted sum across all possible best-response vectors,

$$\sum_{\boldsymbol{v}} p(\boldsymbol{v}) F(\boldsymbol{v}). \tag{4.11}$$

[1] See Section 2.2 for an explanation.
[2] See Section 3.1 in the supplementary material of Pangallo et al. (2019).

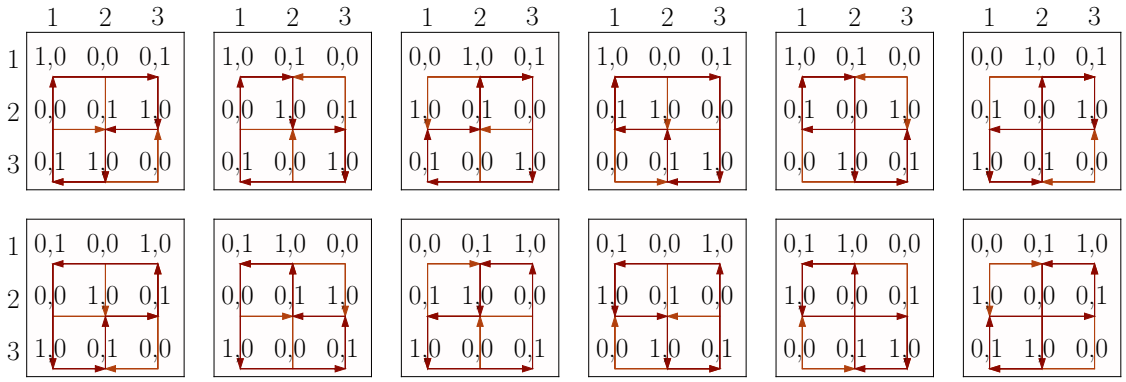


Figure 4.2: All possible best-response realizations in 2-player, 3-action games with one 3-cycle. The value on the left of each action profile represents the payoff to the first (row) player, while the value on the right represents the payoff to the second (column) player. The best responses that are part of a cycle are shown in red, and those within the attractor of a cycle are marked in orange.

Theorem 6 (Pangallo et al. (2019))**.** *The likelihood of drawing a game with the best-response vector* $\boldsymbol{v}$ *is given by*

$$p_1^{2,m}(\boldsymbol{v}) = \frac{1}{m^{2m}} \left(\prod_{k=1}^{m} \prod_{j=1}^{v_k} \frac{p\left(m - \sum_{l=k+1}^{m} v_l l - (j-1)k, k\right)}{j} \right) \times g_m \left(m - \sum_{l \in M} v_l l, 0 \right). \tag{4.12}$$

where

$$p(\tilde{m}, k) = \binom{\tilde{m}}{k}^2 k!(k-1)! \tag{4.13}$$

and

$$g_m(\tilde{m}, d) = m^{2\tilde{m}} - \sum_{k=1}^{\tilde{m}} \frac{p(\tilde{m}, k) g_m(\tilde{m} - k, d+1)}{d+1}. \tag{4.14}$$

The expression for the frequency of games with a certain best-response vector in Theorem 6 was improved in Wiese (2020) by not requiring a recursion:

Theorem 7 (Wiese (2020))**.** *The likelihood of drawing a game with the best-response vector* $\boldsymbol{v}$ *is given by*

$$p_1^{2,m}(\boldsymbol{v}) = \frac{1}{m^{2m}} \left(\frac{m!}{(m-A)!} \right)^2 \left(m^{m-A-1} \right)^2 \frac{2m - A}{(A-1)!} \times \prod_{k \in M} \prod_{j=1}^{v_k} \frac{p\left(A - \sum_{l=k+1}^{m} v_l l - (j-1)k, k\right)}{j}. \tag{4.15}$$

where

$$p(\tilde{m}, k) = \binom{\tilde{m}}{k}^2 k!(k-1)! \tag{4.16}$$

denotes the number of ways to choose a k-cycle from $\tilde{m}$ nodes.

Theorems 6 and 7 can only be applied to 2-player games. The remainder of this chapter will discuss cases for which the frequency of multi-player games with certain attractor structures can be calculated. Our results employ the best-response n-partite digraph representation introduced next.

4.3 The Best-Response n-partite Digraph

We introduce an *n-partite digraph* as an additional representation of the best responses for a given fixed sequence of players. There is a total of nm^{n-1} nodes in n groups, each group corresponding to a player and each node corresponding to an environment of a player. At each node, a player chooses the best response; formally the edges are constructed as follows[3]: for each pair (i, j) of players, where j moves directly after i in the clockwork sequence, and each environment $\boldsymbol{a}_{-i}$, place a directed edge from $\boldsymbol{a}_{-i}$ to another environment $\boldsymbol{a}'_{-j} = (a'_i, \boldsymbol{a}_{-i-j})$, if and only if

$$a'_i = b_i(\boldsymbol{a}_{-i}). \tag{4.17}$$

As we assume that games are non-degenerate, each node in a graph representing a game has an out-degree of 1. A PNE corresponds to a cycle of length n. Each player chooses among m actions at each node; the total number of possible arrangements is $m^{nm^{n-1}}$, each equally likely.

We call the n-partite graph constructed as above but without the Condition (4.17) the *full n-partite graph*. Any n-partite graph corresponding to a given game is a subgraph of the full n-partite graph. We will call a node *free* if its out-degree is m, and *fixed* if its out-degree is 1. An example is provided in Figure 4.3, showing a 3-player, 2-action game with the corresponding 3-partite graph for the playing sequence 1-2-3. The red arrows correspond to the actions of player 1, the black arrows to the actions of player 2, and the red arrows to the actions of player 3. The Nash equilibrium (I-IV-V) is highlighted by thicker arrows in the digraph.
We could have relaxed the assumption of assuming only one specific sequence of players. However, PNEs are stable under any playing sequence. Introducing random or other playing sequences would make the digraph more complicated by adding more arcs without changing the results. Dynamics, where players choose their actions simultaneously (with simultaneous best-response dynamics being the most extreme case), can not be represented by an n-partite graph.

For the remainder of this chapter, we will employ the n-partite digraph to study the likelihood of the class of games that converge under the clockwork best-response dynamic.

[3]For the notation, see Section 2.2.

		Pl. 3		V	VI
	I	Pl. 2	III	(0,0,0)	(0,0,1)
Pl. 1			IV	(1,1,1)	(0,1,0)
	II	Pl. 2	III	(1,0,1)	(1,1,0)
			IV	(0,1,0)	(1,0,1)

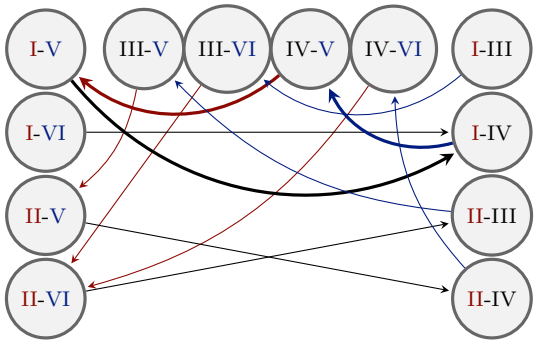


Figure 4.3: A 3-player, 2-action game with one Nash equilibrium and the corresponding 3-partite graph representation under the clockwork playing sequence. The red arrows correspond to the actions of player 1, the black arrows to the actions of player 2, and the red arrows to the actions of player 3. The Nash equilibrium (I-IV-V) is highlighted by thicker arrows in the digraph.

4.4 The Frequency of Convergent Games

We now turn to the class of games for which, starting from any initial strategy profile, no player changes their strategy under the clockwork playing sequence after a sufficiently large number of turns. We will demonstrate that the frequency of *convergent* games with a given number of pure Nash equilibria approaches zero as the number of players or actions increases. These games include the set of dominance-solvable games (Section 4.2.2.2) and the set of coordination games (Section 4.2.2.4).

In general, there are three types of games: *Type A* games, which are convergence games with a unique PNE; *Type B* games, which are convergent games with multiple PNEs; and finally, *Type C* games, which are non-convergent games. Type A games (for instance, the Prisoners' Dilemma) are straightforward to understand and perfectly predictable, as they converge to the PNE. Type B games, which can be thought of as coordination games, allow for multiple possible equilibrium outcomes. An example of a Type C game is Matching Pennies. Unlike Types A and B, Type C games have at least one mixed-strategy Nash equilibrium (MSNE). We will investigate the likelihood of randomly generated games that converge (Type A and Type B) in the ensemble of games with a given number of players and varying numbers of actions available to each player. Our findings will illustrate that Type A and Type B games become increasingly rare as game complexity increases. Thus, using Type A or Type B games to model intricate scenarios, such

as climate change or financial market dynamics, could yield misleading results. Instead, employing Type C games may provide more robust and accurate models despite the inherent complexity involved.

Using the n-partite graph described in Section 4.3, we obtain the frequency of randomly generated games with a unique PNE in the ensemble of n-player, m-action games (Section 4.4.1). We then study convergent games with more than one PNE under the clockwork best-response dynamics (Section 4.4.2). Our analysis reveals that convergent games with fewer PNEs are more common than those with a larger number of PNEs. We provide an exact frequency for convergent 2-player games with any given number of PNEs. Finally, we highlight that, for 2-player games with fewer than 10 actions, games with a unique PNE are more common than convergent games with multiple PNEs; beyond this threshold, the latter become more frequent.

4.4.1 Convergent Games with a Unique PNE

We randomly generate n-player, m-action games by drawing m^n tuples of payoffs from a continuous distribution. This ensures that randomly created games are almost surely non-degenerate. The exact type of distribution does not matter as long as payoffs are uncorrelated because the game's structure only depends on which one in a set of values is the largest. We leave games with correlated payoffs for future research. However, results from the literature (Pangallo et al., 2019) and from Chapter 3 suggest that Type A games are less likely when payoffs are negatively correlated and that Type B games (see below) become overwhelmingly likely with positive correlations.

Let $p^k_{n,m}$ denote the frequency of n-player, m-action convergent games with exactly k PNEs. Our key contribution is the following theorem:

Theorem 8. *The frequency of games with one unique PNE in the ensemble of n-player, m-action games is given by*

$$p^1_{n,m} = r^{n-1} + \frac{m-1}{m-r}\left(\left(\frac{r}{m}\right)^{n-1} - 1\right) \tag{4.18}$$

where $r := \frac{m-1}{m^n} + 1$.

Proof of Theorem 8. Consider the full n-partite digraph (see Section 4.3) for an n-player, m-action game. We order the nodes in the following way: $\boldsymbol{a}_{-i} < \boldsymbol{a}_{-j}$ for different players i and j, if and only if $i < j$, and for the same player i, $\boldsymbol{a}_{-i} < \boldsymbol{a}'_{-i}$ under lexicographical ordering. Denote this full n-partite graph by $G^{\mathrm{f}} = (V^{\mathrm{f}}, E^{\mathrm{f}})$, where V is the set of vertices and E is the set of edges.

The Laplacian matrix of a graph $G = (V, E)$ without multiple edges and self-loops is defined as the square matrix with side length $|V|$ and

$$(L(G))_{ij} = \begin{cases} \delta^+(i) & \text{if } i = j \\ -1 & \text{if } i \neq j,\ (i,j) \in E \\ 0 & \text{if } i \neq j,\ (i,j) \notin E \end{cases} \tag{4.19}$$

where $\delta^+(i)$ is the out-degree of a node i. For G^{f} described above, the Laplacian matrix takes the following form:

$$L\left(G^{\mathrm{f}}\right) = \begin{bmatrix} D & N_1 & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & & \ddots & \ddots & N_{n-1} \\ S & 0 & \cdots & 0 & D \end{bmatrix} \tag{4.20}$$

where $D, S, N_1, \ldots, N_{n-1}$ are square matrices with side length m^{n-1} defined as follows:

- $D = \mathrm{diag}(m)$ is a diagonal matrix with m's on the diagonal
- $N_k = \mathrm{diag}\left(K_1, \ldots, K_{m^{k-1}}\right)$ is a blockmatrix with blockmatrices K_l on the diagonal, where each K_l has side length m^{n-l} and consists of m^2 diagonal matrices $\mathrm{diag}(-1)$, each with side length m^{n-l-1}.
- S is more irregular,

$$(S)_{ij} = \begin{cases} -1 & \text{if } (i \mod m^{n-2}) = \left\lfloor \frac{j-1}{m} \right\rfloor \\ 0 & \text{otherwise.} \end{cases} \tag{4.21}$$

For instance, in the case of 3-player, 2-action games, the Laplacian matrix corresponding to Figure 4.4 (left) is given by

$$L\left(G^{\mathrm{f}}\right) = \left[\begin{array}{cccc|cccc|cccc} 2 & 0 & 0 & 0 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & -1 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 0 & -1 & 0 & -1 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & -1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & -1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & -1 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & -1 & -1 \\ \hline -1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & -1 & -1 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 \\ -1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & -1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 \end{array}\right]. \tag{4.22}$$

There are m^n ways to choose the first PNE, each fixing n nodes. Without loss of generality, we choose the nodes where each player chooses their first action. We condense these n nodes to a single node representing the PNE; see Figure 4.4. The PNE node has an in-degree of $n(m-1)$; we delete all outgoing edges. The resulting $(n+1)$-partite graph consists of $nm^{n-1} - (n-1)$ nodes and will be denoted by $G^{\mathrm{c}} = (V^{\mathrm{c}}, E^{\mathrm{c}})$, where V^{c} is the set of vertices and E^{c} is the set of edges. All nodes except the PNE node are free.

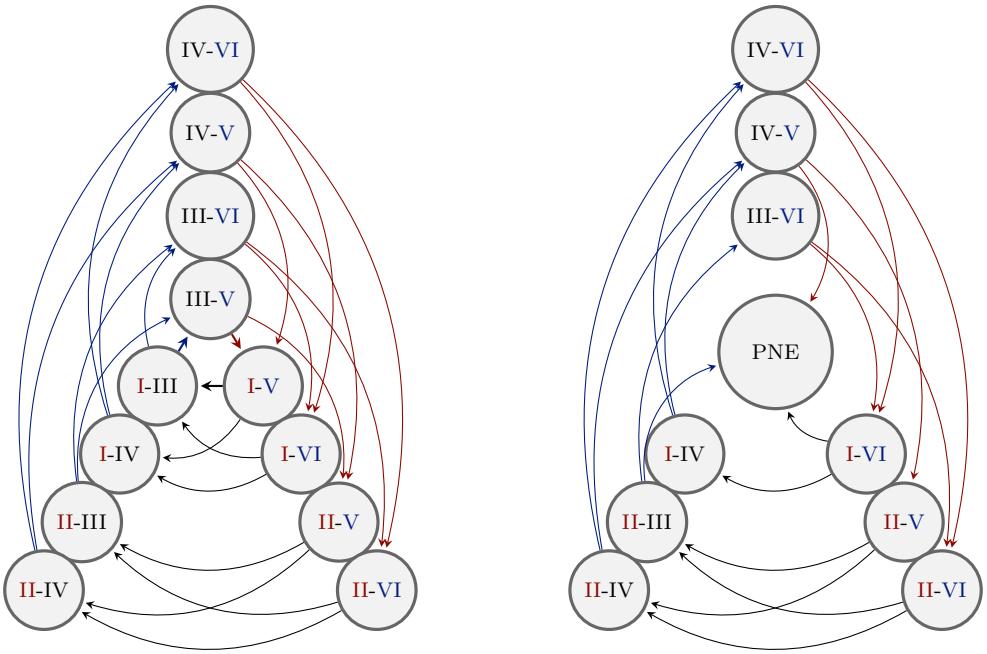


Figure 4.4: For 3-player, 2-action games the full graph G^{f} on the left and the condensed graph G^{c} on the right. The PNE corresponds to the environment where every player chooses their first action.

We apply Kirchhoff's theorem to G^{c} to get the number of spanning trees. This guarantees that the game converges under clockwork best-response dynamics. Kirchhoff's theorem (applied to our problem) states that the number of spanning trees is the determinant of the Laplacian matrix of G^{c} with the first row and column deleted, which corresponds to the PNE-node. For a quadratic matrix A with side length n, we define $\widetilde{A}$ as the quadratic matrix with side length $(n-1)$ obtained from A by deleting the first row and column.

For a general block matrix

$$K = \begin{pmatrix} A & B \\ C & D \end{pmatrix}, \tag{4.23}$$

provided that A is invertible, we have

$$\det K = \det\left(D - CA^{-1}B\right) \det A. \tag{4.24}$$

Applying this identity iteratively to $\widetilde{L(G^{\mathrm{c}})}$ yields

$$\det \widetilde{L(G^{\mathrm{c}})} = m^{\left(m^{n-1}-1\right)(n-1)} \cdot \det\left(\widetilde{D} - \frac{1}{m^{n-1}} \cdot \widetilde{S} \cdot \prod_{i=1}^{n-1} \widetilde{N_i}\right). \tag{4.25}$$

The matrix $\widetilde{S} \cdot \prod_i \widetilde{N_i}$ is given by

$$\left(\widetilde{S} \cdot \prod_i \widetilde{N_i}\right)_{ij} = m - \mathbb{1}_{\left[1, m^{n-\delta(i)}-1\right]}(j) \tag{4.26}$$

where

$$\delta(i) := \underset{p\in[1,n-1]}{\arg\min}\left(\min_{k\in[1,m^p]}\left(\left|i - km^{n-p-1}\right|\right)\right). \tag{4.27}$$

We simplify the matrix $\widetilde{D} - \frac{1}{m^{n-1}} \cdot \widetilde{S} \cdot \prod_{i=1}^{n-1} \widetilde{N_i}$ by elementary row- and column-operations given in Algorithm 6 to obtain a matrix A.

Algorithm 6 The process for simplifying $\widetilde{D} - \frac{1}{m^{n-1}} \cdot \widetilde{S} \cdot \prod_{i=1}^{n-1} \widetilde{N_i}$ to obtain A.

1. For $p \in [1, \ldots, n-1]$:
   (a) For $i \in [1, \ldots, m^{n-1} - 1]$:
      i. If $i = m^{p-1}$ or $\delta(i) \neq n - p$, continue.
      ii. Subtract the m^{p-1}'s row from i.
2. For $k \in [1, \ldots, n^{m-1} - 1]$:
   (a) If for any $p \in [0, \ldots, n-1]$, $k | m^p$, continue.
   (b) Add column k to column $m^{n-\delta(k)-1}$.

The determinant of the matrix A can be written as

$$\det A = m^{m^{m-1}-n} \cdot \det \widehat{A} \tag{4.28}$$

for a matrix $\widehat{A}$ with side length $(n-1)$ and given by

$$\left(\widehat{A}\right)_{ij} = \begin{cases} m + \frac{m-1}{m^{n-1}} - \frac{m-1}{m^{i-1}} & i = j \\ \frac{m-1}{m^{n-i+j-1}} - \frac{m-1}{m^{j-1}} & i > j \\ -\frac{m-1}{m^{j-1}} & i < j \end{cases} \tag{4.29}$$

Adding the i-th column multiplied by $\left(-\frac{1}{m}\right)$ to the $(i+1)$-th column for $i =$

$n-2, n-3, \ldots, 1$, we get a matrix of the following form

$$\begin{bmatrix} D_1 & N & 0 & \cdots & 0 \\ E_2 & D & \ddots & \ddots & \vdots \\ \vdots & 0 & \ddots & \ddots & 0 \\ \vdots & \vdots & \ddots & \ddots & N \\ E_{n-1} & 0 & \cdots & 0 & D \end{bmatrix} \tag{4.30}$$

where

$$D_1 = \frac{m-1}{m^{n-1}} + 1 \tag{4.31}$$

$$E_i = \frac{m-1}{m^{n-i}} - (m-1) \tag{4.32}$$

$$D = m \tag{4.33}$$

$$N = -\frac{m-1}{m^n} - 1 \tag{4.34}$$

To eliminate the N-entries on the upper diagonal, we add the i-th row multiplied by

$$F := -\frac{N}{D} = \frac{\frac{m-1}{m^n} + 1}{m} \tag{4.35}$$

to the $(i-1)$-th row for $i = n-1, \ldots, 2$. Then the matrix is lower-triangular, and the D_1-entry is given by

$$\begin{aligned} \widetilde{D_1} &= D_1 + F \cdot E_2 + \cdots + F^{n-2} \cdot E_{n-1} \\ &= D_1 + \sum_{i=0}^{n-3} F^{i+1} \cdot E_{i+2} \\ &= D_1 + \sum_{i=0}^{n-3} F^{i+1} \left(\frac{m-1}{m^{n-i-2}} \right) - \sum_{i=0}^{n-3} F^{i+1}(m-1) \\ &= D_1 + F \left(\frac{m-1}{m^{n-2}} \right) \sum_{i=0}^{n-3} (Fm)^i - F(m-1) \sum_{i=0}^{n-3} F^i \\ &= D_1 + F \left(\frac{m-1}{m^{n-2}} \right) \left(\frac{(Fm)^{n-2} - 1}{Fm - 1} \right) - F(m-1) \left(\frac{F^{n-2} - 1}{F - 1} \right) \\ &= m(r-1) + m\left(r^{n-1} - r \right) - \frac{r(m-1)}{r-m} \left(\left(\frac{r}{m} \right)^{n-2} - 1 \right) + 1 \end{aligned} \tag{4.36}$$

where $r := \frac{m-1}{m^n} + 1$, and then

$$\det \widehat{A} = m^{n-2} \cdot \widetilde{D_1}. \tag{4.37}$$

Finally, the frequency of games with precisely one PNE is given by

$$p^1_{n,m} = \frac{m^n}{m^{nm^{n-1}}} \det \widetilde{L\left(G^{\mathrm{c}}\right)} = \frac{1}{m^{m^{n-1}-1}} \det A = \frac{1}{m^{n-1}} \det \widehat{A} = \frac{1}{m} \widetilde{D_1} \tag{4.38}$$

where we have multiplied by the number of possible positions of the PNE and divided by the total number of possible arrangements; this completes the proof. □

Note that the frequency $p^1_{n,m} \to 0$ as the number of actions or the number of players goes to infinity and that $p^1_{n,m}$ is decreasing in both n and m. For instance:

$$\begin{aligned} p^1_{2,m} &= \frac{1}{m}\left(2 - \frac{1}{m}\right) \\ p^1_{3,m} &= \frac{1}{m^2}\left(3 - \frac{3}{m} + \frac{3}{m^2} - \frac{3}{m^3} + \frac{1}{m^4}\right) \\ p^1_{4,m} &= \frac{1}{m^3}\left(4 - \frac{4}{m} + \frac{6}{m^3} - \frac{8}{m^4} + \frac{2}{m^5} + \frac{4}{m^6} - \frac{6}{m^7} + \frac{4}{m^8} - \frac{1}{m^9}\right) \end{aligned} \tag{4.39}$$

Our result is different from the one in Stanford (1995), where the frequency of games with one PNE, but possibly also mixed strategy Nash equilibria, converges to $\exp(-1)$ as the number of actions for at least two players goes to infinity.

Figure 4.5 shows the frequency of randomly created games with a unique PNE following Theorem 8. For comparison, the solid black lines denote frequencies obtained by numerically sampling over 500 randomly created games with independent payoffs.[4] The frequency of convergent games decreases with the number of players and the number of actions.

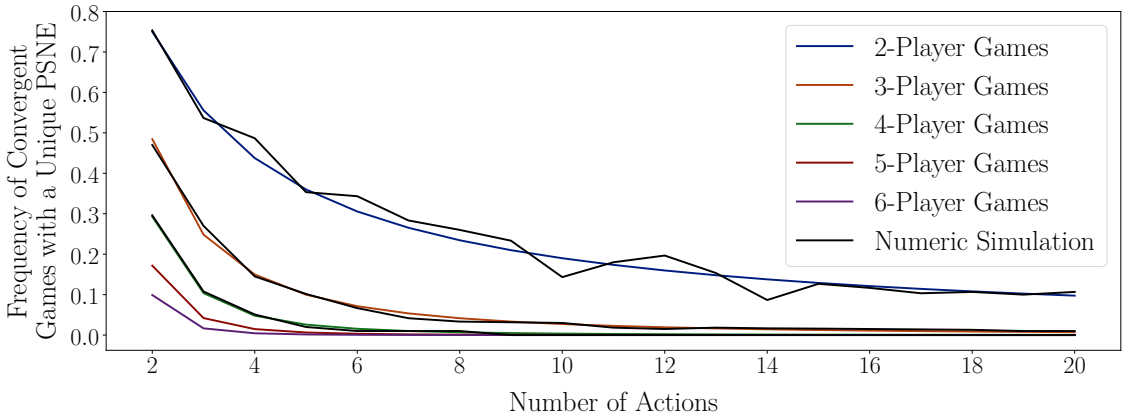


Figure 4.5: The frequency of randomly drawn games with a unique PNE following Theorem 8. For comparison, the solid black lines denote frequencies obtained by numerically sampling over 500 randomly created games with independent payoffs. The frequency of convergent games decreases with the number of players and the number of actions.

4.4.2 Convergent Games with Multiple PNEs

We can bound the frequency of convergent games with more than one PNE from above:

[4] See Section 2.2 for details on creating random games.

Theorem 9. *For $k_1 < k_2$, we have $p^{k_1}_{n,m} > p^{k_2}_{n,m}$.*

Proof of Theorem 9. Consider a full n-partite graph and assign k PNEs, thereby fixing the outgoing edges of kn nodes. We show that the number of possible realisations as a game decreases when adding another PNE.

The number of ways we can add another PNE (which is, in general, very complicated to compute) is bounded from above by $(m^{n-1} - k)\, m = m^n - km$, which is because there are $m^{n-1} - k$ free nodes for each player, each free node has an out-degree of m, and fixing two nodes of an n-cycle fixes the remaining ones. However, adding a PNE decreases the number of possible digraph realisations by a factor of $m^n - 1$ because the n nodes may not form a cycle.

Induction over the number of added PNEs completes the proof. □

Theorems 8 and 9 imply that for every k, $p^k_{n,m} \to 0$ as the number of actions or the number of players goes to infinity. We computed for 3-player, 2-action games that $p^1_{3,2} = \frac{1984}{4096} \approx 48.43\%$, $p^2_{3,2} = \frac{828}{4096} \approx 20.21\%$, $p^3_{3,2} = \frac{56}{4096} \approx 1.37\%$, $p^4_{3,2} = \frac{2}{4096} \approx 0.049\%$.

In 2-player games, we can precisely state the frequency of games with k PNEs.

Theorem 10. *The frequency of 2-player, m-action convergent games with exactly k PNEs in the ensemble is given by*

$$p^k_{2,m} = \frac{2m-k}{m^{2k+2}(k-1)!}\left(\frac{m!}{(m-k)!}\right)^2. \tag{4.40}$$

for $k \leq m$, and is otherwise 0.

Proof of Theorem 10. It was shown by Austin (1960) that the number of chromatic digraphs with m nodes of each type, where each node has an out-degree one, and with a cycle of length $2k$, $1 \leq k \leq m$, is

$$(2m-k)\left(m^{m-k-1}\right)^2\left(\frac{m!}{(m-k)!}\right)^2. \tag{4.41}$$

Factoring out the number of ways to arrange k vertices on a cycle ($(k+1)!$) and the total number of possible arrangements (m^{2m}), we get

$$p^k_{2,m} = \frac{2m-k}{m^{2k+2}(k-1)!}\left(\frac{m!}{(m-k)!}\right)^2. \tag{4.42}$$

□

The frequency of drawing a 2-player convergent game (Type A or Type B) is then given by $\sum_{k=1}^{m} p_{2,m}^{k}$, the frequency of Type B games only is $\sum_{k=2}^{m} p_{2,m}^{k}$. Numerical evidence shows that Type A games are more common than Type B games for $m = 2, \ldots, 9$, and less common for $m \geq 10$. Figure 4.6 shows the frequency of randomly drawn convergent 2-player games with a given number of actions and PNEs following Theorem 10. The frequency of convergent games decreases with the number of PNEs, and, generally speaking, with the number of actions.

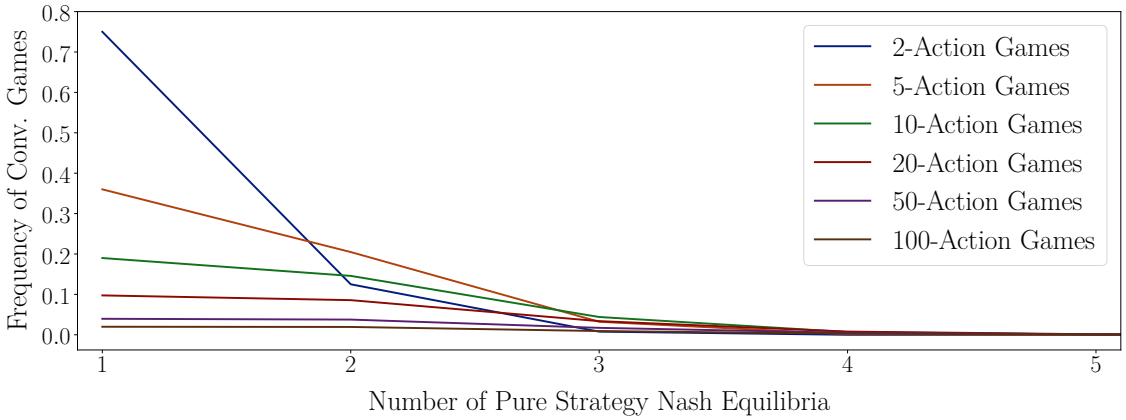


Figure 4.6: The frequency of randomly drawn convergent 2-player games that have a given number of PNEs following Theorem 10. The frequency of convergent games decreases with the number of PNEs, and, generally speaking, with the number of actions.

4.5 Conclusion

In this chapter, we examined the frequency of generating games with specific structures when drawn randomly. By analysing the attractor structures in 2-player and multi-player games, we provided key insights into how often certain types of games — such as those with unique pure Nash equilibria (PNEs) — arise within the broader ensemble of n-player, m-action normal-form games. Our results showed that games with a unique PNE become increasingly rare as the complexity of the game (in terms of the number of players or actions) increases, and we derived an explicit formula for this frequency. We also demonstrated that games with multiple PNEs are more common than those with a unique PNE as the number of players or actions grows. Additionally, we extended existing frameworks by introducing the best-response n-partite digraph as a tool to study attractor structures in multi-player games, and developed upper bounds for the frequency of games that converge under the clockwork best-response dynamic. Our findings contribute to a deeper understanding of the predictability and stability of equilibria in complex strategic environments. This work highlights the challenges in using simple game

structures to model intricate real-world scenarios, as convergent games become rare in more complex setups.

Future research could develop bounds or explicit expressions on the frequency of generating other classes of games at random. A related open question relates to classifying attractors in multi-player games and expressing the likelihood of generating such attractors.

This concludes our exploration of random games, having examined the likelihood of drawing specific game ensembles and the convergence properties of best-response dynamics. Our findings on the predictability and frequency of different types of games contribute to a deeper understanding of strategic interactions in complex systems. With this foundation laid, we now turn to the second part of this thesis, which focuses on data-driven macroeconomic agent-based modelling. Since our findings indicated that complex and competitive games are less likely to converge to an equilibrium, and that games that converge are rare, we turn to large-scale macroeconomic agent-based models, which do not rely on the equilibrium assumption. In the following chapters, we will explore how agent-based models can be initialised, calibrated, and used for economic forecasting.

Part II

Macroeconomic Agent-based Modelling

5 Literature

This chapter surveys the literature on macroeconomic agent-based modelling. To underscore the need for alternative approaches, we begin by discussing the development and limitations of conventional models (Section 5.2). In Section 5.3, we conduct a comprehensive review of agent-based macroeconomic models, comparing 11 prominent models and frameworks in terms of their scope — covering which agents and markets are represented — and their relationship to data, including calibration, initialisation, the ability to match stylised facts, and forecasting ability. We also contrast these models with the framework introduced in Chapter 6. Section 5.4 outlines the calibration techniques used in Chapter 7, comparing likelihood-based approaches with simulation-based inference. In particular, we introduce Neural Posterior Estimation and Neural Ratio Estimation. Section 5.5 concludes the chapter.

Contents

5.1 Introduction 100
5.2 Conventional Macroeconomic Modelling 101
5.2.1 Review of DSGE Modelling 102
5.2.2 Limitations of DSGE Models 105
5.3 Agent-based Modelling in Macroeconomics 106
5.3.1 Overview 106
5.3.2 Notable Examples of Macroeconomic ABMs 109
5.4 Calibrating Agent-Based Models 117
5.4.1 Likelihood-based Approaches 118
5.4.2 Simulation-based Inference 118
5.5 Conclusion 120

5.1 Introduction

The economic upheavals of the past two decades, ranging from the global financial crisis of 2008 to the COVID-19 pandemic and the energy shock triggered by the Russian invasion of Ukraine, have highlighted the limitations of traditional macroeconomic models. *Dynamic Stochastic General Equilibrium* (DSGE) models, which continue to be widely used in both academia and policy-making institutions, failed to foresee or adequately capture the complexities and immediate impacts of these crises on global economies.[1] These shortcomings have amplified calls within the economic community for advancements in modelling techniques. This has given rise to two distinct avenues: either enhancing existing DSGE models to incorporate agent heterogeneity, such as through the development of *Heterogeneous Agent New Keynesian* (HANK) models (Kaplan et al., 2018), or taking a more radical departure from traditional frameworks. In the latter category, *Agent-Based Models* (ABMs) have emerged as a robust and compelling alternative, offering a fundamentally different approach to analysing and understanding economic dynamics (Axtell and Farmer, 2024).

Agent-based models consist of computational objects, known as agents, which interact according to predefined behavioural rules (Miller and Page, 2009), often situated within space and time (Holland and Miller, 1991; Page, 2008). These agents can be either homogeneous, sharing identical characteristics, or heterogeneous, endowed with distinct attributes (Bankes, 2002). The rules governing their behaviour may range from simple heuristics, such as backward induction strategies in an extensive-form game, to more complex mechanisms based on cognitive or psychological insights (Marchi and Page, 2014). When these individual behaviours are aggregated, they give rise to system-wide outcomes that are described as emergent phenomena or generative processes (Cederman, 2005).

One of the earliest agent-based models was Thomas Schelling's famous "Segregation model" (Schelling, 1969), where he initially only used simple tools like coins and graph paper to study the degree of racial segregation in cities. He found that even a slight preference of individuals for living in a specific neighbourhood leads to significant segregation on the macro level. This demonstrated how aggregate behaviour could sharply diverge from individual preferences. Since then, ABMs have been used extensively (Haldane and Turrell, 2018), in diverse fields such

[1] See Blanchard (2018); Stiglitz (2018) for general critiques, further discussed in Section 5.2.2; see Smets and Wouters (2007); Benes et al. (2014) for critiques in the context of the global financial crisis and Guerrieri et al. (2022); Reis (2022) for critiques in the context of the COVID-19 pandemic.

as biology[2], network science[3], and finally economics[4], including large economies and, specifically, business cycles.

Agent-based models allow us to study disequilibrium states, complex interactions between agents, and leverage large datasets. Although the model can rarely be explicitly solved, and finding robust mathematical results is often impossible, this flexibility can offer a broader view of the entire state space. However, their flexibility is also a reason for their limited adoption: First, calibration of large, analytically intractable systems is difficult. The outcome can only be simulated for a specific combination of parameters and initial conditions.[5] Second, even though agents often use simple rules-of-thumb for their decisions, there remains an overwhelming range of decision rules available[6], meaning a wide range of simulation output can be generated. Therefore, we validate our agent-based model (see Section 6) by assessing its forecasting performance (see Section 7.4), instead of its ability to replicate stylised facts.

We begin this chapter by reviewing the development and limitations of conventional macroeconomic models, in order to contrast the current prevalent approach to ABMs.

5.2 Conventional Macroeconomic Modelling

This section briefly discusses the prevalent approach to macroeconomic modelling, *Dynamic Stochastic General Equilibrium* (DSGE) models. Before the financial crisis of 2007-2008, DSGE models were the dominant framework for monetary policy analysis (Clarida et al., 1999; Woodford, 2011) and for studying business cycles within a general equilibrium context (Christiano et al., 2018; Eisenbach and Sannikov, 2013). However, neither the crisis nor the subsequent *Great Recession* was predicted by these models at the time (Christiano et al., 2018; Lindé et al., 2016), which may be attributed to certain problematic assumptions, such as perfect agent rationality and a lack of heterogeneity among agents.

DSGE models are constructed from macroeconomic foundations. They typically consist of optimising agents[7] forming rational expectations and maximising their

[2]For the analysis of the spread of epidemics (Situngkir, 2004), population dynamics (Caplat et al., 2008), or plant-animal interactions (Fedriani et al., 2018).

[3]Including social network effects (Franchi, 2014) or supply chain optimisation (Okada et al., 2016).

[4]For an early economic ABM, see the model by Orcutt (1957); for a later encouragement to use ABMs in economics, see Farmer and Foley (2009).

[5]For a potential solution, see Section 5.4.

[6]Referred to as "the wilderness of bounded rationality" by Sims (1980).

[7]Not all models feature optimising agents. For example, in Coenen and Straub (2005), a fraction of households are non-optimising and deplete all their earnings in each period.

objective functions subject to their respective constraints. Specifically, DSGE models usually represent households and firms. Households consume goods, supply labour, and accumulate capital to maximise a specified utility function, which may be subject to various constraints (e.g., the budget constraint, labour demand constraint, or Calvo constraints on the frequency of wage adjustments). Firms produce goods and hire labour, seeking to maximise their stream of profits, subject to their respective constraints (e.g., demand constraint, Calvo constraint on the frequency of price adjustments). Additionally, these models often include a central bank, which utilises a Taylor rule.

A non-exhaustive list of DSGE models includes DSGE models with fully flexible prices (the subclass of RBC models) or nominal rigidities (the NK models), both of which are prevalent in central banks for monetary policy analysis. There are also DSGE models with a representative household or with heterogeneous households and firms (Khan and Thomas, 2007; Kaplan et al., 2018), as well as DSGE models with either infinitely lived agents or finitely lived ones (Nishiyama and Smetters, 2014). Additionally, there are DSGE models with complete financial markets or with financial frictions and incomplete markets (Fernández-Villaverde et al., 2023), with rational expectations or with learning (Primiceri, 2006), and featuring either full rationality or behavioural biases (Gabaix, 2020).

We provide a brief review of the development of these models (Section 5.2.1) and highlight their strengths and limitations (Section 5.2.2).

5.2.1 Review of DSGE Modelling

This section provides a brief review of the development of DSGE modelling, starting with *Neo-Keynesian* models in the 1950s (Section 5.2.1.1) to *Real Business Cycle* models (Section 5.2.1.2), to *New-Keynesian* models (Section 5.2.1.3), until modern *Heterogeneous Agent New Keynesian* models (Section 5.2.1.4).[8]

5.2.1.1 Neo-Keynesian Models

Keynes's *General Theory* (Keynes, 1936) sparked the construction of large-scale macroeconomic models. The Keynesian revolution emphasised the importance of aggregate demand in determining output and employment, marking a departure from classical economics. Keynes argued that government intervention was necessary to stabilise the economy during times of insufficient demand. *Keynesian-style* models are systems of equations with ad-hoc decision rules inspired by Keynesian macroeconomic theory. These models gained popularity in the 1950s[9]; however,

[8]This classification mostly follows the review by Slanicay (2014).
[9]See, for example, Prais (1956).

in the 1970s, their popularity began to decline due to their failure to explain the simultaneous increase in inflation and unemployment. This stemmed from their dependency on the *Phillips curve*, implying a negative correlation between inflation and unemployment. Several empirical studies highlighted the weak forecasting performance of such models, showing that simple statistical extrapolation of time series provided better predictive accuracy (see, e.g. Nelson (1972)).

Economists became increasingly dissatisfied with the theoretical foundations of these models, particularly criticising their reliance on disequilibrium conditions and the absence of microfoundations for assumptions regarding price and wage rigidities. This criticism led to the emergence of the *New Keynesian* research agenda (see Section 5.2.1.3), which sought to provide solid microfoundations for these rigidities.[10] Furthermore, following the influential work of Muth (1961), which introduced the concept of rational expectations, economists also began to challenge the ad-hoc use of adaptive expectations in these models.

The most significant criticism of Neo-Keynesian models came from the "Lucas critique" (Lucas Jr, 1976), which challenged the reliability of large-scale macroeconomic models built in an ad-hoc manner for forecasting the impact of changes in economic policy. Lucas argued that because the parameters of Neo-Keynesian models are not structural (i.e., they depend on policy), they would shift when economic policy changes. Consequently, predictions could be inaccurate as they are solely based on relationships observed in historical data. This critique led economists to prioritise models with microeconomic foundations, where macroeconomic outcomes are driven by the decisions of individual agents, such as households and firms.

The creation of DSGE models stems from the effort to design models less vulnerable to the Lucas critique, i.e., models whose parameters remain stable regardless of changes in economic policy. These "deep parameters" include factors like the discount rate, the elasticity of substitution between goods, and the elasticity of labour supply — elements considered unaffected by monetary policy. This development paved the way for the rise of *Real Business Cycle* (Section 5.2.1.2) models, which offered a fresh framework for analysing macroeconomic fluctuations.

5.2.1.2 Real Business Cycle Models

Kydland and Prescott introduced the first DSGE model (Kydland and Prescott, 1982), which laid the foundation for the *Real Business Cycle* (RBC) theory. Following their work, numerous RBC models were developed, including those by Long Jr and Plosser (1983) and Prescott (1986). The central insight of these models is that economic fluctuations are primarily driven by real, rather than

[10] For a collection of influential papers, see Mankiw and Romer (1991).

nominal, shocks, with productivity shocks playing a key role. RBC models view business cycles as optimal responses to real shocks, with markets continually clearing. The calibrated RBC models were able to replicate certain unconditional second moments of key macroeconomic time series, such as relative standard deviations and correlations (see Long Jr and Plosser (1983); Prescott (1986)).

The assumption that real forces, especially productivity shocks, are the main causes of business cycles represented a substantial departure from Keynesian models, which emphasised demand-side shocks and the role of monetary policy. However, two key predictions of RBC models — positive comovement of output, labour inputs, and productivity in response to technology shocks, as well as the dominance of technology shocks as drivers of business cycle fluctuations — have been challenged (Galí and Rabanal, 2004). Positive technology shocks frequently lead to negative short-term comovement of output and labour inputs, and the impact of technology shocks on business cycle fluctuations is relatively minor. Additionally, RBC models assume that monetary policy has no influence on real variables, which contradicts the evidence presented by Friedman and Schwartz (2008) and Christiano et al. (1998). Summers (1986) offers a detailed critique, highlighting that several RBC model parameters[11] are calibrated to implausible values.

Despite these shortcomings, RBC models were instrumental in shaping the development of DSGE models. They introduced key elements such as dynamic optimisation, rational expectations, and general equilibrium, which were subsequently refined and incorporated into New Keynesian DSGE models. The RBC framework provided a foundation for future progress, particularly in how macroeconomists understand the behaviour of economic agents over time.

5.2.1.3 New Keynesian Models

The shortcomings of RBC models prompted the development of *New Keynesian* (NK) models. These models retained the methodological framework and core structure of RBC models, particularly the principle of optimising agents. However, unlike RBC models, NK models introduced Keynesian elements, such as monopolistic competition in goods or labour markets and price or wage rigidities. A key innovation was the incorporation of nominal rigidities through mechanisms like Calvo pricing, where only a fraction of firms adjust their prices in any given period, leading to price stickiness. New Keynesian DSGE models also incorporated rational expectations, addressing criticisms raised by Lucas. By integrating microfoundations with sticky prices, these models preserved important Keynesian insights while resolving many

[11] For instance, the proportion of household time spent on market activities, the average real interest rate, and the elasticity of labour supply.

of the theoretical limitations of earlier models. As a result, New Keynesian DSGE models became central to modern macroeconomic policy analysis and are now widely used by central banks to assess the impact of monetary policy.

5.2.1.4 Heterogeneous Agent New Keynesian Models

Heterogeneous Agent New Keynesian (HANK) models evolved from the New Keynesian framework (Kaplan et al., 2018). They can be summarised by three key aggregate equations: the Phillips curve, which defines the relationship between inflation and output dynamics; the Taylor rule, which describes how the monetary authority sets the nominal interest rate; and the Fisher equation, which connects the real interest rate, policy rate, and expected inflation (Slanicay, 2014). The primary innovation in HANK models is the replacement of the representative consumer with a more modern theory of consumption and saving. In HANK models, households face idiosyncratic risks, such as unemployment or health shocks, which they mitigate through saving or borrowing. These shocks introduce heterogeneity across households in terms of income, consumption, and wealth, contrasting with earlier DSGE models that assumed a representative agent, overlooking the distribution of income and wealth. As a result, HANK models can produce a realistic cross-sectional distribution of income, consumption, and wealth.

HANK models represent a major step forward in understanding the effects of monetary and fiscal policies on different household groups. By incorporating heterogeneity in income and wealth, these models offer deeper insights into the distributional impacts of policy interventions, which is crucial for addressing inequality. Furthermore, the increasing availability of microdata has strengthened the empirical relevance of HANK models, allowing policymakers to gauge household-level responses to macroeconomic shocks better.

5.2.2 Limitations of DSGE Models

This section briefly summarises the limitations and shortcomings of DSGE models. First, DSGEs often include numerous *ad-hoc* economic mechanisms designed primarily to fit the persistence properties of data. Second, they require many unexplained shocks, which are frequently assumed to be highly persistent. Third, there is substantial evidence that the pure rational expectations assumption is flawed, as agents encounter aggregate, endogenous uncertainty (Gatti et al., 2010). Fourth, research on HANK models (see Section 5.2.1.4) has demonstrated that household and firm heterogeneity significantly impacts macroeconomic aggregates. However, HANK models still incorporate only a limited degree of heterogeneity (Fagiolo and Roventini, 2016). Finally, in cases where no microdata is utilised (which

is common), the information available from macro time series and prior distributions is limited compared to the number of calibrated variables. For instance, the well-known DSGE model by Smets and Wouters (2007) is estimated using Bayesian methods and only incorporates seven quarterly macro time series. As a result, they struggle to model large-scale phenomena such as regime shifts or significant fluctuations within this framework.[12]

None of these criticisms imply that DSGE models are not useful for analysing and estimating aggregate economic behaviour in a meaningful way.[13] Unlike time series models (such as VARs), DSGE models impose budget constraints, provide a consistent narrative for agent behaviour, and offer a coherent framework for handling expectations.

5.3 Agent-based Modelling in Macroeconomics

We now turn to agent-based modelling as an alternative approach to conventional macroeconomic modelling. This section provides an overview of the field (Section 5.3.1), outlining its key principles and applications, and compares several popular models and frameworks (Section 5.3.2), which have been instrumental in shaping the modelling approaches adopted in Chapter 6.

5.3.1 Overview

The aim of this section is to summarise why agent-based models could complement the DSGE approach in the future. Farmer and Axtell (2022) recently provided a direct comparison by listing essential features of economic agent-based models and corresponding standard neoclassical abstractions, highlighting which of these idealisations have been relaxed in the economics literature. However, unlike conventional economic models, the agent-based framework allows us to realistically model particular interactions between agents.[14]

We briefly explore the development of agent-based models in macroeconomics (Section 5.3.1.1), highlight a recent success story (Section 5.3.1.2), and specifically review the latest progress in constructing *data-driven* agent-based models (Section 5.3.1.3).

[12] For more comprehensive critiques of DSGE models, see Blanchard (2018); Canova and Sala (2009); Colander et al. (2009); Kirman (2010); Krugman (2011); Mankiw (2006); Romer (2016); Solow (2010); Stiglitz (2011). For an overview of theoretical, empirical, and political-economy issues associated with DSGE modelling, see Fagiolo and Roventini (2016).

[13] See Christiano et al. (2018) for a recent defence of DSGEs.

[14] For example, price formation should be a consequence of transactions, not a precondition for them (Kirman, 2006; Gaffeo et al., 2007, 2008).

5.3.1.1 Development

Early work on macroeconomic agent-based models includes the study by Allen and Carroll (2001), who introduce an agent-based model of individuals attempting to learn effective consumption-saving behaviour without performing the type of optimisation found in DSGE models. Their trial-and-error approach takes far too long for a single person's lifetime, pointing to the potential role of social learning, where individuals learn from others' experiences. Similarly, Bullard and Duffy (2001) modelled how agents learn about macro volatility. All models at the time faced constraints due to limited computing power. However, around fifteen years ago, the idea of using high-performance computing to simulate large-scale agent models in macroeconomics emerged. The EURACE model (Deissenberg et al. (2008); Teglio et al. (2010), Section 5.3.2.2) was the first major example of this endeavour, yielding a complex representation of the European economy, incorporating thousands of agents (households, individuals, banks, central banks, and governments). Simultaneously, in the aftermath of the Great Financial Crisis, the *CRISIS* (Complexity Research Initiative for Systemic Instabilities) project[15] aimed to develop agent-based models for understanding the events surrounding the crisis. Specifically, it produced several interesting ABMs for understanding system-wide risks, feedback loops within the financial system, and studying different regulatory policies (Poledna and Thurner, 2016).

Over the past decade, agent-based models have made considerable advances, both in terms of the economic and financial processes they simulate, and their empirical foundations (Farmer and Axtell, 2022). Dawid and Gatti (2018) provides a comparison of the scope, structure, and modelling assumptions of 7 "families" of macroeconomic ABMs. We build upon their work and review a larger set of modern macroeconomic agent-based models in terms of scope and relation to data in Section 5.3.2.

5.3.1.2 A Recent Success Story

Recently, Pichler et al. (2020) analysed several scenarios for the phased reopening of the British economy after COVID-19 by explicitly modelling production network effects. At one extreme, they found that reopening only a minimal number of industries could create supply chain miscoordination issues that might reduce aggregate output. Conversely, reopening all industries without implementing extra precautionary measures would increase the infection rate above one and could lead to a second wave of the epidemic. Their results suggested that a reasonable compromise exists, which would produce a relatively small rise in epidemic spread

[15] See https://cendef.uva.nl/projects/crisis-eu-fp7/crisis.html.

while substantially boosting economic output. Crucially, their model uniquely addressed the specific features of the crisis, simultaneously handling both supply and demand shocks, along with sluggish consumption recovery, and included sectoral inventory buffers. The core of the model was an industrial production network that distinguished between critical and non-critical inputs.[16] The model was calibrated to the UK economy and initialised with first-order supply and demand shocks derived by del Rio-Chanona et al. (2020). Validation exercises demonstrated that the model aligned with recent estimates by the Office for National Statistics and was consistent with US sectoral unemployment data. Their predicted conditional forecast for a GDP loss of 21.5% in the second quarter of 2020 was remarkably close to the observed loss of 22.1%.

5.3.1.3 Data-Driven Modelling

In general, ABMs have become increasingly data-driven in recent years, aiming to reflect a specific economy more accurately at a particular point in time (Pangallo, 2024). To achieve this, it is insufficient to merely calibrate a few parameters; it is also essential to initialise the state variables of individual agents in a manner that aligns with real-world data. The main advantage of this approach is that data-driven ABMs can track and forecast empirical time series, thereby explicitly demonstrating their ability to faithfully represent the economy and increasing confidence in counterfactual analyses and policy recommendations. Most data-driven ABM applications have focused on specific sectors of the economy, such as housing (Geanakoplos et al., 2012) and labour markets (del Rio-Chanona et al., 2021), as well as on assessing the economic impact of disasters (Hallegatte, 2008; Pichler et al., 2022). Nevertheless, data-driven macroeconomic ABMs have also been developed (Papadopoulos, 2019; Kaszowska-Mojsa and Pipień, 2020).

We will explicitly distinguish between model *initialisation* (estimating initial conditions for agent-level attributes) and model *calibration* (estimating fixed model parameters). In recent years, the initialisation of agent-based models has shifted from randomly setting variables at the start of a simulation to more methodological approaches utilising microdata and macrodata.[17] Calibration, on the other hand, involves estimating model parameters, which can often be directly obtained from data or expert knowledge. However, when this is not feasible, all existing calibration methods require sampling from the parameter space (in the simplest case, uniformly at random) and selecting the parameters based on a chosen criterion and summary statistic (such as moments of aggregate quantities). The framework we introduce

[16] To determine which inputs were critical, they conducted a survey with IHS Markit sector analysts, asking, "Can production continue in sector X if input Y is unavailable for two months?".

[17] E.g. *Iterative Proportional Fitting* (Kruithof, 1937).

in Chapter 6 first employs initialisation directly from microdata and then applies Bayesian calibration to a subset of parameters (see Section 5.4 below).[18]

Additionally, to ensure that the agent-based model output tracks empirical time series over time, it may be necessary to continuously estimate the *latent* (not directly observable) variables of the model. For an example of data assimilation applied to a macroeconomic agent-based model, see Papadopoulos (2019). However, we do not discuss data assimilation further here as it is not utilised in the macroeconomic model presented in Chapter 6.

5.3.2 Notable Examples of Macroeconomic ABMs

We continue by summarising some of the most notable large macroeconomic agent-based models, ranging from simple toy models to complex models that attempt to incorporate many real-world features. Our selection is based on models that had the most significant impact on the literature and were developed relatively recently or feature unusual behavioural choices. We primarily follow Dawid and Gatti (2018) in selecting the models we compare and the dimensions we consider. Specifically, we examine the types of agents (banks, firms, households, etc.) and the markets in which they interact (goods, labour, credit, and housing). Additionally, we introduce another particularly relevant dimension: how models relate to data (Pangallo, 2024). We first consider which parts of the ABM are derived from data, including the calibration of parameters and the initialisation of individual agents' state variables. We also assess what aspect of the data is targeted, that is, whether the models aim at reproducing summary statistics or stylised facts, or whether they attempt to forecast empirical time series. Our comparison is summarised in Table 5.1.

We consider eleven models or families of models. First, we examine the so-called CATS model, whose baseline version is described in the book by Delli Gatti et al. (2011). Additionally, we include the extended version of the CATS model that incorporates capital goods (Assenza et al., 2015), as it has been calibrated (Delli Gatti and Grazzini, 2020; Glielmo et al., 2023) and used as the foundation for Poledna et al. (2023a). We also consider a simplified version of the CATS model, called Mark-0 (Gualdi et al., 2015b), which has been the subject of rich theoretical studies. These studies demonstrate that contrary to the common critique of ABMs, a deep mathematical understanding of their behaviour is possible (Gualdi et al., 2015a), and that advanced numerical techniques can facilitate efficient exploration of their parameter space (Naumann-Woleske et al., 2021). Next, we include the two descendants of the EURACE project, developed at the University of Bielefeld (UNIBI) and initially described by Deissenberg et al. (2008), and at the University of

[18] Platt (2020) finds that Bayesian models outperform when applied to simple models.

		Agents							Markets				Rel. to Data			
Model	Year	Banks	Central Bank	Firms	Central Gov.	Gov. Entities	Households	Individuals	Goods	Labour	Credit	Housing	Calibration	Initialisation	Stylised Facts	Forecasting
CATS	2005	x		x			x		x	x	x				x	
K+S	2006	x	x	x	x		x		x	x	x		x		x	
EURACE UNIBI	2008	x	x	x	x	x	x	x	x	x	x	x			x	
EURACE UNIGE	2008	x	x	x	x	x	x	x	x	x	x	x			x	
Lagom GeneriC	2010		x	x	x		x		x	x				x	x	
Lengnick Model	2013	x		x			x		x	x	x				x	
CC-MABM	2015	x	x	x			x		x	x	x		x		x	x
Mark-0	2015	x		x			x		x	x	x					
Ashraf Model	2017	x	x	x	x	x	x		x	x	x		x		x	
CANVAS Model	2022	x	x	x	x	x	x		x	x	x			x		x
IIASA Model	2023	x	x	x	x	x	x		x	x	x			x		x
Our Model	2024	x	x	x	x	x	x	x	x	x	x	x	x	x		x

Table 5.1: Comparison of 11 macroeconomic agent-based models against the one we are building Chapter 6 in terms of agents, markets, and relation to data (calibration, initialisation, stylised facts, and forecasting). Our model features all types of agents and markets, and is validated by forecasting economic time series instead of satisfying stylised facts.

Genoa (UNIGE), first introduced by Cincotti et al. (2012a). Furthermore, we include the family of Keynes meets Schumpeter (K+S) ABMs developed at Sant'Anna Pisa (see, for instance, Dosi et al. (2010, 2015)), the framework developed by Ashraf et al. (2016, 2017), and the baseline macroeconomic ABM by Lengnick (2013). Finally, we consider the main baseline for the framework developed in Chapter 6, namely the IIASA model (Poledna et al., 2023a),[19] along with its extension by the Bank of Canada, CANVAS (Hommes et al., 2022). These models are compared with the large agent-based model we are introducing in Chapter 6 and employ in Chapter 7.

In terms of types of agents, all models feature one or more banks, firms, and households, at least in some versions.[20] Among the models we consider, only the EURACE models (Deissenberg et al., 2008; Cincotti et al., 2012a) and our model explicitly distinguish between households and individuals. Every model features markets for goods, labour, and credit; however, only the EURACE models and ours include a housing market.

[19] For a more detailed comparison, see Section 6.3.

[20] For instance, Dosi et al. (2010) do not include banks, but Dosi et al. (2015) do.

In terms of their relation to real-world data, most models aim to replicate stylised facts by calibrating parameters. The majority of papers calibrate parameters heuristically, either by setting reasonable values or by adopting values from previous studies. A formal parameter calibration procedure is only followed in the CC-MABM (Delli Gatti and Grazzini, 2020; Glielmo et al., 2023), the K+S model (Martinoli et al., 2022), Ashraf et al. (2017), and in our model. Even without a formal calibration procedure, most models succeed in replicating stylised facts.[21]. As discussed in the introduction, only the IIASA model, the CANVAS model, and our model initialise the variables of individual agents to be compatible with real-world data, which enables them to perform out-of-sample forecasting.[22]

As highlighted in Table 5.1, our model distinguishes itself by its comprehensive scope. It covers more economic sectors than other models, making it broader in scope, and introduces robust Bayesian calibration for empirical validation, enhancing the benchmarks set by Poledna et al. (2023a) and Hommes et al. (2022). Furthermore, we are applying this model to a large set of countries — 38 — increasing the statistical validity of our approach. These innovations position our model as a significant advancement in the field of macroeconomic ABMs.

For the remainder of this section, we discuss the mentioned models in more detail.

5.3.2.1 The CATS Framework

The CATS Framework (*Complex Adaptive Trivial Systems*), first introduced by Gatti et al. (2005) and further developed by Russo et al. (2007), employs an agent-based model of firms and banks. Firms raise funds in the credit market based on their investment expenditures, which are linked to banks' interest rates. The model produces a right-skewed firm size distribution that follows a Zipf power law probability density function (Axtell, 2001; Gaffeo et al., 2003).

5.3.2.2 The EURACE Models

The EURACE project was funded by the European Commission to develop an agent-based simulation platform suitable for macroeconomic analysis and for studying the effects of different policy measures. The initial model is a complex

[21] Such stylised facts are typically related to comovements between aggregates (e.g. gross output is positively correlated with firm intermediate consumption, gross fixed capital formation, household consumption, government consumption, and exports; inflation is negatively correlated with the unemployment rate), moments (e.g. gross fixed capital formation is more volatile than gross output), and microdistributions (e.g. fat-tailed distributions in aggregate growth rates (Castaldi and Dosi, 2009; Fagiolo et al., 2008); only a relatively small subset of firms are exporters, and those firms tend to be larger and more productive than non-exporters (Dosi et al., 2010, 2019)).

[22] See also Delli Gatti and Grazzini (2020) for an approach to forecasting with ABMs that relies on a kind of Bayesian surrogate model.

representation of the European economy (Deissenberg et al., 2008), incorporating households, individuals, banks, central banks, and governments. Each iteration in the model corresponds to a business day, although not all economic activities occur daily. Agents are linked by networks representing supply chains, employment, and the supply of loans, which form based on matching processes in five markets: consumption goods, investment goods, labour, credit, and financial assets. Except for the investment goods market and the market for financial assets, these markets are cleared separately within each NUTS-2 region. The project's developed agenda and model versions can be found in Holcombe et al. (2013); Cincotti et al. (2012a). After the project's conclusion, the EURACE framework was further developed at the University of Bielefeld and the University of Genoa. An extensive description of the Bielefeld version is provided in Dawid et al. (2016). For an application to firm dynamics, see Dawid and Harting (2012). For an analysis of fiscal policies, see Harting (2015), and for the integration and analysis of social networks, see Dawid et al. (2014). The continuation at the University of Genoa incorporates different specifications and framework scales. For analyses of monetary policy, see Raberto et al. (2008); for studies on the effects of financial regulation, see Cincotti et al. (2010); Teglio et al. (2010, 2012); Cincotti et al. (2012b); Raberto et al. (2012, 2017); and for fiscal policy and sovereign debt analysis, see Raberto et al. (2014); Teglio et al. (2019). For applications to the housing market, see Erlingsson et al. (2013, 2014); Ozel et al. (2019). A multi-country version of the model is presented in Petrovic et al. (2017).

5.3.2.3 The Lagom GeneriC Model

In the Lagom framework (Mandel et al., 2010), households and firms apply genetic evolution to setting consumption technology, production, wages, and prices. See also Wolf et al. (2013) or the simpler model in Gintis (2007), which do not include capital accumulation, inter-sectoral trade, or life-saving behaviour.

The core economic cycle of the original model by Mandel et al. (2010) consists of goods production, consumption, and trading. Each firm produces the maximum possible quantity, given its stock of intermediate, capital, and labour inputs, according to a Leontief production function. Intermediate inputs are used up, and capital inputs depreciate. Household demand is based on total wealth and observed prices, assuming unitary elasticity. Households receive income from wages, unemployment benefits, and firm dividends. The government taxes income at the rate necessary to pay precisely for total unemployment benefits. Households have reservation wages and actively look for higher-paying jobs. The target employment of each firm is determined by the workforce it can put to work, given its current

stock of capital inputs. Interest rates are set based on a Taylor rule, which defines a natural interest rate, target inflation, and a target unemployment rate.

5.3.2.4 The K+S Model

The Keynes-meeting-Schumpeter (K+S) model (Dosi et al., 2010) is part of the "Keynes meeting Schumpeter" framework, which also includes the earlier models by Dosi et al. (2006) and Dosi et al. (2009). For an analysis of labour market issues, see Napoletano et al. (2012), and for an application to the study of climate change, see Lamperti et al. (2018).

The basic model by Dosi et al. (2010) consists of firms producing consumption goods, firms producing capital goods, households acting as consumers and workers, and a central government paying unemployment benefits and collecting taxes on firm profits. Firms producing capital goods invest in research and development and produce heterogeneous capital inputs using labour. Consumption-good firms produce goods for household consumption using capital inputs and labour. Capital-good firms set prices with a fixed markup on production costs (labour), and consumption-good firms set prices based on a dynamic markup on costs, which reflects the previous change in market share. A new firm replaces firms with either a close-to-zero market share or negative net assets, and its stocks of intermediate and capital inputs are set to an average across incumbent firms. The authors validate the model against various stylised facts, including the lead/lag correlation structure in output, growth rate distributions, and firm size distributions.

5.3.2.5 The Model by Lengnick (2013)

The model by Lengnick (2013) is a simple model of households and firms interacting in goods and labour markets. Households are characterised by two properties: their wealth and their reservation wage. In each iteration, their wealth increases through labour income and decreases through spending on consumption goods. The rate of consumption expenditure increases with wealth but at a decreasing rate. Households are employed by firms, which set wage rates based on their success in attracting workers during the previous period. Firms similarly set prices based on their success in selling goods. Production in firms follows a linear production function that depends solely on labour. Firms maintain trading relationships with other firms. These trading networks are fixed in the short term but can evolve over time. A key aspect of the model is that actions can occur at different time intervals (e.g. monthly or daily), allowing coordination problems to emerge. The model replicates several stylised facts, such as an endogenous business cycle, a Phillips curve, and a Beveridge curve.

The model shares similarities with those in Dosi et al. (2009) and Gaffeo et al. (2008), particularly in its use of adaptive rules for pricing and wage-setting. However, the models mentioned focus more on growth driven by research and development (R&D). For further developments of the original Lengnick model, see Lengnick and Wohltmann (2013) and Lengnick and Wohltmann (2016).

5.3.2.6 The CC-MABM Model

The CC-MABM model (Assenza et al., 2015) focuses on the systemic risk and crises arising from the financial sector. The model describes an economy consisting of households, firms, and banks. In this framework, the household sector is divided into workers and firm investors ("capitalists", with one investor per firm). Workers supply labour and purchase consumption goods, while firm investors receive dividends. Both groups accumulate wealth in the form of bank deposits. The corporate sector comprises firms producing both consumption and capital goods. Consumption goods firms demand labour and capital goods to produce and sell goods to households, while capital goods firms require only labour to produce and sell capital goods to the consumption goods firms. On the other hand, banks serve as intermediaries by receiving deposits from households and extending credit to firms. Each bank determines both the interest rate and the amount of credit to provide to a firm.

All agents — households, firms, and banks — follow simple adaptive behavioural rules. Economic transactions occur within goods, labour, and credit markets, all of which are cleared through random search-and-matching processes. The model generates time series for key economic indicators such as GDP, consumption, investment, and unemployment, which are comparable to real-world data, particularly in terms of standard deviation and first-lag autocorrelation.

5.3.2.7 The Mark-0 Model

The Mark-0 model (Gualdi et al., 2015b) is a simplified version of the Mark-1 model, developed by Gaffeo et al. (2008); Gatti et al. (2008, 2011). Both Mark-0 and Mark-1 include firms, households, firm owners, and banks. Firms produce a single, non-durable good in proportion to the number of employees. When firms are unable to meet wage payments, they apply for loans, and the bank sets interest rates based on the firms' financial fragility. Firm owners supply the workforce in exchange for a salary and allocate a fixed fraction of their savings to consumption. Firms follow simple heuristic rules to update production and prices in response to observed demand.

The Mark-0 model is amenable to analytical treatment and extends Mark-1 by introducing dynamic wage-setting and variable spending propensities. Both

models display a *tipping point* — a discontinuous shift — between high and low employment levels, which is influenced by the baseline interest rate.

5.3.2.8 The Model by Ashraf et al. (2017)

The model introduced by Ashraf et al. (2017) consists of households, firms, a central government, banks, and a central bank. The goods market is cleared by randomly matching households with at most one firm, while firms also trade with each other. The government is responsible for fiscal policy, monetary policy, and bank regulation: it collects taxes, sets interest rates, lends to banks, enforces capital adequacy requirements, and insures bank deposits. Additionally, the government issues money and bonds.

Banks finance new and existing firms, invest in government bonds, and provide fully collateralised commercial loans. The central bank follows a Taylor rule for setting interest rates. The authors' main finding is that while the model typically tracks full capacity utilisation, it may become unstable due to firm turnover. During such periods of instability, the economy performs better when banks are lightly regulated (i.e., higher loan-to-value ratios or lower capital requirements). The model successfully replicates key macroeconomic relationships, including the Phillips curve, the Beveridge curve, and Okun's law. For an extension of this model to monetary and macro-prudential analysis, see Popoyan et al. (2017).

5.3.2.9 The IIASA Model

Poledna et al. (2023a) introduce a macroeconomic agent-based model (which we refer to as the IIASA model) of an open European country that uses (mostly[23]) public macroeconomic data from national sector accounts, input-output tables, government statistics, and census data for initialisation. The model is based on Assenza et al. (2015) and incorporates households, firms, a central government, government entities, a bank, and a central bank. Gill et al. (2021) present a case study of implementing the model on a high-performance computing cluster, while Magnuszewski et al. (2023) and Poledna et al. (2023b) apply it to migration policy. Additionally, Bachner et al. (2023) use the model to study flood risk in Austria. The version of Hommes and Poledna (2023) extends the original model to the whole Euro area and applies it directly to the Financial crisis of 2007/2008, the subsequent Great Recession, the European sovereign debt crisis, and the COVID-19 recession. The modelling rules are mostly unchanged, except for a minor adjustment for existing inventory when planning future production and a rule for accounting for liquidity shortages.

[23] except data on household investment

Firms produce goods using labour, intermediate inputs, and capital inputs. They operate in 62 industries, corresponding to NACE-level 2. Unlike some other models, there is no explicit distinction between households and individuals. Households are either employed, unemployed, or not economically active. Employed households supply labour and receive wages, while unemployed or not economically active households receive social benefits. They operate on the goods market; their consumption is homogeneous and depends solely on household income. A single bank grants loans to firms, and the interest rate is set as a markup on the policy rate, which is determined by a simple Taylor rule set by the central bank. To close the model, countries not explicitly modelled are represented by a "rest of the world" agent, interacting via imports and exports. There is an international market for goods and local markets for labour and credit. The goods and housing markets create competition for prices, the credit market for interest rates, and the labour market for wages.

The model's strengths lie in its careful consideration of stock flows and its ability to be initialised using Eurostat macro data alone. However, some of the behavioural assumptions in the model appear unconventional. First, inflation and future growth are estimated using an AR(1)-process. It seems unusual that every firm employs the same forecast, and this growth estimate should ideally be sector-specific. Second, even though there are 62 distinct types of goods/sectors, they are all assumed to be perfect substitutes in production, and the current stock of finalised goods is not considered, which artificially inflates prices and demand. Third, firms set prices by adding a markup to their per-unit production costs without considering supply or demand factors. Fourth, the goods market is cleared by a random search-and-matching process, where the probability of selecting a particular seller is based on the seller's size and price. The model prescribes a Boltzmann probability for prices with an inverse temperature of 2, and the probability for size is based on the firm's share of total output. The seller's probability is the average of the probabilities corresponding to size and price. This setup allows a large firm to set excessively high prices yet still be chosen with a probability of $\frac{1}{2}$, which is unrealistic. Finally, every household spends the same fixed fraction of income on consumption and investment, which is unrealistic. We compare the IIASA model and our model (see Chapter 6) in Section 6.3.

5.3.2.10 The CANVAS Model

The "Canadian Behavioral Agent-Based Model" (CANVAS) (Hommes et al., 2022) is being developed at the Bank of Canada, with a focus on evaluating the effects of multiple macroeconomic policies. It shares much of the structural foundation

with the model developed at IIASA (Section 5.3.2.9) but incorporates Canadian microdata to create heterogeneous households. Additionally, CANVAS introduces depreciation of finalised goods, firm-specific growth rates based on previously observed demand, and more nuanced price-setting mechanisms driven by *cost-push* and *demand-pull* inflation. The rule for price setting is examined in more detail by Grazzini et al. (2023), particularly in the context of post-pandemic inflation dynamics in Canada. The Bank of Canada runs CANVAS at a scale of 1:100 to inform their understanding of inflation targeting, making it one of the first agent-based models actively used within a central bank (Hommes et al., 2024).

When it comes to calibration, of the surveyed models, parameters of the K+S model (Section 5.3.2.4) and the CC-MABM model (Section 5.3.2.6) are chosen ad-hoc, and parameters of the model by Ashraf et al. (2017) can be directly taken from public US data.[24] In Chapter 6, we will introduce a large macroeconomic agent-based model, whose parameters are mostly calibrated directly from data; however, we also perform *Bayesian* calibration on a small subset of parameters. The following section reviews the literature on calibrating agent-based models, including the methods we will employ in Chapter 7.

5.4 Calibrating Agent-Based Models

Calibrating agent-based models involves estimating the fixed parameters that govern agent behaviour. Often, these parameters cannot be directly observed or derived from empirical data, necessitating the use of calibration techniques. Calibration aims to identify parameter sets that minimise the model's error, ensuring that simulated data aligns with real-world observations. This is critical for making the model both predictive and analytically useful. In Chapter 7, we will use Bayesian calibration to find parameters that ensure the simulated data matches observed real-world data.

Calibrating agent-based models is challenging for the following reasons: First, running agent-based models can be computationally expensive, and the calibration procedure may require a large number of model runs. This is particularly true in our case — in Chapter 7, we calibrate a large, stochastic representation of entire countries, which is expensive and may require millions of runs, depending on the calibrated parameters. Second, agent-based models often generate complex, high-dimensional data, which leads to the need to simplify assumptions in the

[24] It is noteworthy, that the model by Ashraf et al. (2017) was the first macroeconomic agent-based models which had all of its parameters fit to real-world data.

calibration procedure. Finally, agent-based models are only simplifications of real-world systems, so there will inevitably be a gap between the model's predictions and the actual data (McCulloch et al., 2022).

In the following sections, $\mathcal{M}$ denotes the model containing free parameters $\boldsymbol{\theta}$. Real-world data is denoted by $\boldsymbol{y}$ and simulation output by $\boldsymbol{x}$. The inference aims to obtain a posterior density $\pi(\boldsymbol{\theta} \mid \boldsymbol{y}, \mathcal{M})$ for $\boldsymbol{\theta}$ given y.

5.4.1 Likelihood-based Approaches

Exact Bayesian inference proceeds by specifying a prior distribution $\pi(\boldsymbol{\theta} \mid \mathcal{M})$ for $\boldsymbol{\theta}$, which is updated with the observed data $\boldsymbol{y}$ using *Bayes' theorem* to form a posterior distribution:

$$\pi(\boldsymbol{\theta} \mid \boldsymbol{y}, \mathcal{M}) = \frac{p(\boldsymbol{y} \mid \boldsymbol{\theta}, \mathcal{M})}{m(\boldsymbol{y} \mid \mathcal{M})} \pi(\boldsymbol{\theta} \mid \mathcal{M}), \tag{5.1}$$

where

$$m(\boldsymbol{y} \mid \mathcal{M}) = \int p(\boldsymbol{y} \mid \boldsymbol{\theta}, \mathcal{M}) \, \pi(\boldsymbol{\theta} \mid \mathcal{M}) \, \mathrm{d}\theta \tag{5.2}$$

is the marginal likelihood of the model. This allows the posterior distribution to reflect the uncertainty in the parameters $\boldsymbol{\theta}$ given the observed data $\boldsymbol{y}$ (Robert et al., 2007). However, agent-based models typically have an intractable likelihood function, making it impossible to evaluate the posterior density directly via Bayes's theorem. In contrast, *simulation-based Bayesian inference* procedures approximate the posterior density using only our ability to simulate the model.

5.4.2 Simulation-based Inference

Simulation-based inference (SBI) generates simulated data under various parameter values and compares it to the observed data (Rubin, 1984). More specifically, the process typically involves sampling a parameter set $\boldsymbol{\theta_1}$ from the prior, running the agent-based model to generate simulation outputs $p(\boldsymbol{x} \mid \boldsymbol{\theta_1}, \mathcal{M})$, and using the simulations to estimate the likelihood function $\hat{p}(\boldsymbol{y} \mid \boldsymbol{\theta_1}, \mathcal{M})$ at the observed data $\boldsymbol{y}$. These steps are repeated for many parameter values to build up an approximation of the posterior distribution.

However, SBI is particularly challenging for ABMs because these models can be computationally expensive and often generate complex, high-dimensional data. Many approaches to SBI require a large number of simulations to achieve a small Monte Carlo error, which increases computational costs. Moreover, modelling the data distribution is difficult when dealing with high-dimensional data, often leading to the need for simplifying assumptions.

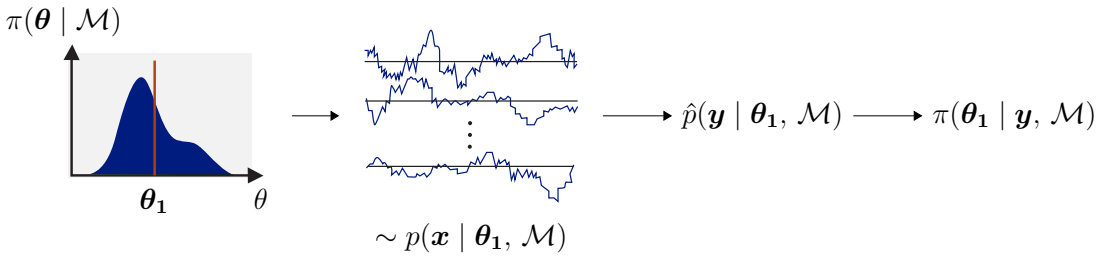


Figure 5.1: Estimating the posterior using simulation-based inference. The process involves sampling a parameter set $\boldsymbol{\theta_1}$ from the prior, running the agent-based model to generate simulation outputs $p(\boldsymbol{x} \mid \boldsymbol{\theta_1}, \mathcal{M})$, and using the simulations to estimate the likelihood function $\hat{p}(\boldsymbol{y} \mid \boldsymbol{\theta_1}, \mathcal{M})$ at the observed data $\boldsymbol{y}$. The illustration follows Dyer (2022).

In this thesis, we will employ two modern methods (Dyer et al., 2024) that address these challenges: *Neural Posterior Estimation* (Section 5.4.2.1) and *Neural Ratio Estimation* (Section 5.4.2.2). Both methods use neural networks to approximate the posterior distribution efficiently, allowing for faster calibration compared to traditional approaches such as *Approximate Bayesian Computation* (Tavaré et al., 1997; Pritchard et al., 1999; Beaumont et al., 2002).

5.4.2.1 Neural Posterior Estimation

Neural Posterior Estimation (NPE) models the posterior density $\pi(\boldsymbol{\theta} \mid \boldsymbol{y}, \mathcal{M})$ directly using a neural network $q_{\boldsymbol{\phi}}$, where $\boldsymbol{\phi}$ are the weights of the network (Papamakarios and Murray, 2016; Lueckmann et al., 2017; Greenberg et al., 2019). The trained weights $\hat{\boldsymbol{\phi}}$ are obtained by minimising the Kullback-Leibler divergence between the joint density over parameters and data, i.e., by minimising

$$\begin{aligned} &\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{\theta},\mathcal{M})\,\pi(\boldsymbol{\theta}|\mathcal{M})}\left[\log \frac{p(\boldsymbol{x} \mid \boldsymbol{\theta}, \mathcal{M})\,\pi(\boldsymbol{\theta} \mid \mathcal{M})}{q_{\boldsymbol{\phi}}(\boldsymbol{\theta} \mid \boldsymbol{x}, \mathcal{M})\,m(\boldsymbol{x} \mid \mathcal{M})}\right] \\ &\qquad = \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{\theta},\mathcal{M})\,\pi(\boldsymbol{\theta}|\mathcal{M})}\left[-\log q_{\boldsymbol{\phi}}(\boldsymbol{\theta} \mid \boldsymbol{x}, \mathcal{M})\right] + \text{constants in } \boldsymbol{\phi}. \end{aligned} \tag{5.3}$$

This motivates the loss function

$$\mathcal{L}(\boldsymbol{\phi}) = \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{\theta},\mathcal{M})\,\pi(\boldsymbol{\theta}|\mathcal{M})}\left[-\log q_{\boldsymbol{\phi}}(\boldsymbol{\theta} \mid \boldsymbol{x}, \mathcal{M})\right], \tag{5.4}$$

which is, in practice, estimated using a finite Monte Carlo sample of N data-parameter pairs $(\boldsymbol{x}^{(i)}, \boldsymbol{\theta}^{(i)})_{i=1}^{N} \sim p(\boldsymbol{x} \mid \boldsymbol{\theta}, \mathcal{M})\,\pi(\boldsymbol{\theta} \mid \mathcal{M})$ simulated from the prior and the ABM. The posterior $\pi(\boldsymbol{\theta} \mid \boldsymbol{y}, \mathcal{M})$ is then approximated as $q_{\hat{\boldsymbol{\phi}}}(\boldsymbol{\theta} \mid \boldsymbol{y}, \mathcal{M})$. For our calibrating exercise in Chapter 7, we take $q_{\boldsymbol{\phi}}$ to be a normalising flow (Tabak and Vanden-Eijnden, 2010); further details on the neural network architecture and training procedure are provided in Section 7.2.2.

5.4.2.2 Neural Ratio Estimation

Neural Ratio Estimation (NRE) is an alternative approach that also uses neural networks (Hermans et al., 2020; Thomas et al., 2022; Dyer et al., 2022). Here, a neural network $r_{\boldsymbol{\varphi}}$ with parameters $\boldsymbol{\varphi}$ is trained to approximate the log ratio of the likelihood to the marginal likelihood for the model. While this can be achieved in various ways, we do so here by finding neural network weights $\hat{\varphi}$ that minimise the loss function

$$\mathcal{L}(\boldsymbol{\varphi}) = -\frac{1}{B}\sum_{b=1}^{B} \log\left(\frac{\exp\left(r_{\boldsymbol{\varphi}}(\boldsymbol{x}^{(b)}, \boldsymbol{\theta}^{(b)})\right)}{\exp\left(r_{\boldsymbol{\varphi}}(\boldsymbol{x}^{(b)}, \boldsymbol{\theta}^{(b)})\right) + \sum_{k=1}^{K} \exp\left(r_{\boldsymbol{\varphi}}(\boldsymbol{x}^{(b)}, \boldsymbol{\theta}^{(b,k)})\right)}\right), \tag{5.5}$$

where a batch of $B > 1$ data-parameter pairs are drawn from the true joint density, $\boldsymbol{x}^{(b)}, \boldsymbol{\theta}^{(b)} \sim p(\boldsymbol{x} \mid \boldsymbol{\theta}, \mathcal{M})\,\pi(\boldsymbol{\theta} \mid \mathcal{M})$, and each of the $\boldsymbol{\theta}^{(b,k)} \sim \pi(\boldsymbol{\theta} \mid \mathcal{M})$, $K > 1$. Durkan et al. (2020) showed that this induces the neural network to approximate, up to some additive constant $c(\boldsymbol{x})$, the log-ratio of the likelihood and marginal likelihood at $\boldsymbol{x}$, i.e.,

$$r_{\hat{\boldsymbol{\varphi}}}(\boldsymbol{x}, \boldsymbol{\theta}) \approx \log \frac{p(\boldsymbol{x} \mid \boldsymbol{\theta}, \mathcal{M})}{m(\boldsymbol{x} \mid \mathcal{M})} + c(\boldsymbol{x}). \tag{5.6}$$

This can then be used to recover approximate posterior samples for y, for example, by Markov chain Monte Carlo procedures. Further details on the training procedure are provided in Section 7.2.2.

5.5 Conclusion

This chapter extensively reviews a range of influential macroeconomic agent-based models in terms of scope and relation to data. We contrast them to conventional Dynamic Stochastic General Equilibrium models (DSGEs), for which we discuss their development and limitations. Agent-based models offer a promising alternative, allowing for great flexibility in modelling agent heterogeneity, interaction, and non-equilibrium dynamics. However, they also pose new challenges, particularly regarding calibration and validation. We introduced modern calibration techniques that are particularly suitable for agent-based models, which we will apply to the model detailed in Chapter 6. We will detail the calibration and validation of the model in Chapter 7.

6

A Framework for Macroeconomic Agent-based Modelling

This chapter presents a novel framework for constructing large, data-driven, agent-based macroeconomic models. The framework is modular, allowing for the flexible replacement and comparison of behavioural rules and datasets. Section 6.2 provides an overview of the agents, their interactions, the data used for initialization, and the sequence of events. We significantly diverge from the benchmark IIASA model regarding implementation, scope, and mechanisms (Section 6.3), with key innovations such as modelling 38 countries, a detailed credit market, a housing market, and more realistic behavioural rules. Section 6.4 outlines the modelling choices regarding the agents, which include firms, households, individuals, a central government, government entities, banks, and a central bank. Section 6.5 describes the goods, labour, credit, and housing markets. The modelling specifications are based on standard theory, common choices from the literature on macroeconomic agent-based modelling, and discussions with experts. Section 6.6 concludes the chapter, particularly highlighting future development on the framework.

Contents

6.1 Introduction **122**
6.2 Model Overview **123**
6.3 Key Differences to the IIASA Model **138**
6.4 Agents **140**
6.5 Markets **191**
6.6 Conclusion **195**

6.1 Introduction

Although macroeconomic ABMs have been around for twenty years (Dawid and Gatti, 2018), they have predominantly been successful at generating qualitative insights, particularly regarding policy (Fagiolo and Roventini, 2016). Typical macroeconomic ABMs are validated by demonstrating their ability to reproduce stylised facts or summary statistics from real-world data. In recent years, ABMs have become more data-driven, striving to faithfully represent a specific economy at a specific point in time (Pangallo, 2024).[1] A breakthrough in using data-driven ABMs for macroeconomics was achieved by Poledna et al. (2023a), and we will use their IIASA model[2] as a benchmark.

This chapter introduces a novel framework for building data-driven agent-based models. Our framework allows for modelling a set of countries, each comprising firms, individuals, households, a central government, government entities, banks, and a central bank. Agents interact in markets for goods, labour, credit, and housing. We depart from the IIASA model in several key specifications designed to enhance realism and scope. The modelling specifications are based on standard theory, common choices from the literature on macroeconomic agent-based modelling, and discussions with experts. Our main innovations include:

- **Housing Market**: Households buy, sell, let, and rent properties. Because houses represent a substantial fraction of household wealth, slowdowns in the housing market can have ripple effects throughout the economy.
- **Credit Market**: Banks make loans to firms and households, considering realistic supply-side loan requirements and setting interest rates using a statistical model employed by private sector banks. Our more detailed modelling of the credit market improves our model's suitability for studying macroprudential policy.
- **Using Microdata**: Individuals are grouped into households and differ by several socio-economic characteristics, including age, income, wealth, and occupational status. In general, the additional heterogeneity that accurately matches the demographic and economic characteristics of the simulated countries allows us to better capture the economic conditions and the diverse behaviours of the agents within the model.
- **Enhanced Behavioural Rules and Market Mechanisms**: We have integrated more realistic behavioural rules and mechanisms that are absent

[1] See Section 5.3.1.3.
[2] See Section 5.3.2.9.

in the IIASA model. For instance, in the IIASA model, firms set production targets limited by past resource constraints, which significantly restricts their expansion capacity and constrains aggregate growth. In contrast, we allow firms to set production targets that only partly consider past resource constraints, enabling them to establish more ambitious targets.

An important property of our framework is its *modularity*, allowing for easy replacement of agents, behavioural rules, and datasets. This modularity enables us to reproduce the rules laid out in other models, such as the IIASA or the CANVAS[3] models, as well as new ones like ours. The specification of agents, rules, and inputs that we present below is only one possible realisation intended to highlight the breadth of the model and showcase its performance in forecasting economic aggregates in Chapter 7. However, we also summarise alternative approaches for each agent and behavioural rule.

Notation We use the following notation for this chapter: By overlines $\overline{x}$, we denote values predicted by agents, and with hats $\hat{x}$, we denote target values. If not otherwise specified, values are in nominal terms. We write $[x]^+ = \max(0, x)$, $[x]^- = -\min(0, x)$, and $\Delta x(t) = x(t) - x(t-1)$.

We begin by providing a thorough overview of the model, describing the agents and their interactions, the data used, and the sequence of events.

6.2 Model Overview

Our model enables the simulation of multiple countries simultaneously, providing a closed framework in which the entire world economy is portrayed. The modeller has the option to model specific countries explicitly, while all others are aggregated into a collective entity we call the "rest of the world" (ROW). For example, one could simulate Germany and the United States. The simulation would then be organised into three main blocks: one for Germany, one for the United States, and one for the rest of the world. In this work, however, we restrict ourselves to a simple configuration where we run only one country at a time, paired with the ROW. This systematic approach will subsequently be applied to 38 countries.

Our model is calibrated to represent the economic structure of countries at a given point in time. Within each country, different economic agents — firms, individuals, households, banks, a central bank, and the government — play distinct roles and interact with each other. Each country is initialised with 18 sectors[4]. To manage

[3] See Section 5.3.2.10.
[4] See Table 6.1.

computational complexity, we apply a scale factor: one agent (firms, households, etc.) in the simulation represents 1000 real agents.[5,6] Every agent is calibrated using microdata, allowing for heterogeneity. As in the IIASA model, we ensure full stock-flow consistency. Figure 6.1 schematises the basic structure of the agent-based model.

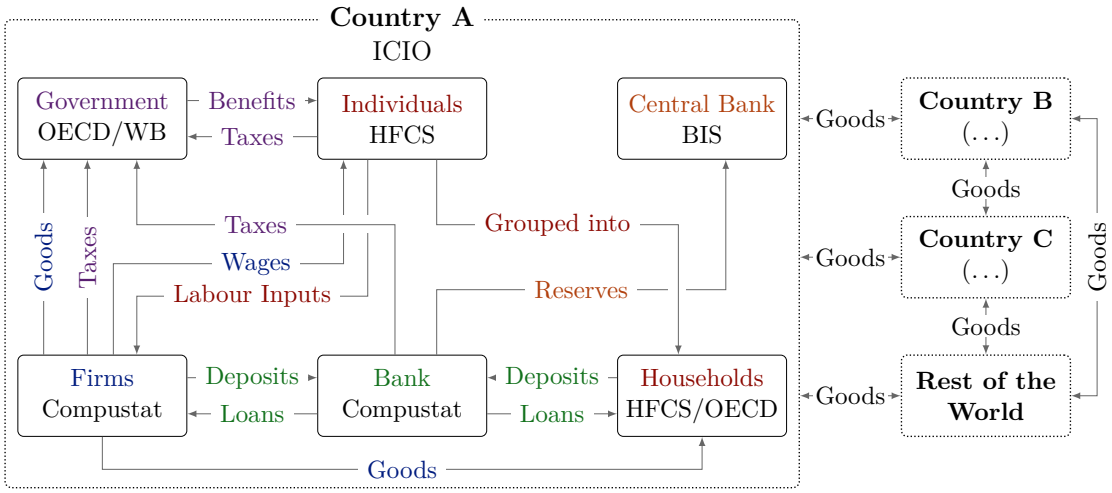


Figure 6.1: The basic structure of the agent-based model with the most important datasets used to calibrate each agent. These include the inter-country input-output tables from the OECD, World Bank data, Compustat microdata for firms and banks, Household Finance and Consumption Survey data for households and individuals, and policy rates provided by the Bank for International Settlements.

The firm sector comprises 18 industry sectors according to the NACE-2 classification, as shown in Table 6.1.

Section 6.2.1 briefly describes the agents in the model along with their key interactions, and Section 6.2.2 summarises the datasets used. Section 6.2.3 provides the sequence of events in a single iteration of the model, while Section 6.2.4 summarises economic aggregates, and Section 6.2.5 summarises the expectation setting.

[5]There are no general results on the optimal number of agents in (macroeconomic) agent-based models. Too few agents may prohibit desired emergent behaviour, while running simulations becomes too expensive if the number of agents is too large. Many empirical quantities in economic data are distributed with heavy tails, such as Pareto-distributed wealth, firm sizes, and city sizes. These distributions are so skewed that it is difficult to assess the characteristics that arise from small numbers of agents. Di Domenico and Riccetti (2023) discuss the impact of scaling the IIASA model and find no qualitative or quantitative differences between models of different scales. Hosszú et al. (2024) study the effect of scaling down the 1-1 model described in Mérő et al. (2023). They find that with fewer agents, the variance of key variables increases, making calibration less efficient, and forecasting performance (goodness of fit of key variables) deteriorates.

[6]The increasing availability of large-scale computing power represents an opportunity for agent-based modellers — for instance, Gill et al. (2021) describes running the model in (Poledna et al., 2023a) on a large-scale computing cluster, and the authors of Gleiser et al. (2024) describe running the model in Kaszowska-Mojsa and Pipień (2020) on Amazon cloud infrastructure.

Description	NACE Rev. 2
Agriculture, forestry and fishing	A
Mining and quarrying	B
Manufacturing	C
Electricity, gas, steam and air conditioning supply	D
Water supply; sewerage; waste management, and remediation	E
Construction	F
Wholesale and retail trade; repair of motor vehicles	G
Transporting and storage	H
Accommodation and food service activities	I
Information and communication	J
Financial and insurance activities	K
Real estate activities	L
Professional, scientific and technical activities	M
Administrative and support service activities	N
Public administration and defence; compulsory social security	O
Education	P
Human health and social work activities	Q
Arts, entertainment and recreation + Other services activities	R + S

Table 6.1: Classification of economic activities NACE Rev. 2 Level 1.

6.2.1 Description of Agents

This section briefly describes the agents in the model and their key interactions, as illustrated in Figure 6.1.

6.2.1.1 Individuals and Households

Individuals in the model represent the inhabitants of a country, with demographic characteristics such as age reflecting the real population distribution. These individuals are grouped into households. Each household's income and wealth are derived from the economic activities of its members, shaping their consumption habits and decisions. Households participate in the housing market, where they can rent, let, or buy homes; in the credit market, where they can apply for loans; and in the goods market, where they purchase what they consume. Individuals can also be hired and fired in the labour market.

6.2.1.2 Firms

Firms are categorised by sector and serve as the primary production units in the economy. They interact by purchasing inputs from local and international

markets and transforming them into goods and services. Firms hire workers and apply for loans.

6.2.1.3 Banks and the Central Bank

Banks play a crucial role by providing credit to agents and holding their deposits. They operate within the credit market, where loans are issued with varying maturities and interest rates that are contingent on the type of loan and on the agent to whom it is given — for example, consumption loans or mortgages for households, or investment loans for firms. The borrower's financial health and the broader economic context influence these terms, particularly the interest rate. The interest rate charged by banks is typically a markup over the policy rate, which is set by the central bank.

6.2.1.4 The Government

The government is responsible for a wide range of activities. It collects taxes from individuals and firms, which constitutes its primary source of revenue, and has the authority to set fiscal policy. These funds are used to purchase goods in the goods market, fuelling public investment. The government also provides individuals with benefits based on specific criteria, such as income, employment status, or other demographic characteristics.

6.2.1.5 The Rest of the World

This agent, external to all countries, models the global economy's non-explicit components. It interacts with domestic economies exclusively via the goods market, encapsulating import/export activities. This allows for a mathematically closed model.

6.2.2 Data

At its core, our simulation is structured around the detailed description of the agents' balance sheets and income statements described above. These balance sheets and their relationships are initialised using a comprehensive array of data sources that are crucial for accurately representing a country's characteristics. Most importantly, the *Inter-Country Input-Output Tables* (ICIO) from the OECD[7] clearly describe the input dependencies of sectors within and across national economies. We perform a few simple steps to allow the raw ICIO tables to be used in our model:

[7] See https://data-explorer.oecd.org.

Handling Input-Output Tables The *Inter-Country Input-Output Tables* of the OECD[8] consist of the intermediate-use matrix **Z**, a taxes less subsidies on intermediate and final products row **TLS**, a value-added row **VA**, the output row **X**, the household final consumption expenditure column **HFCE**, the non-profit institutions serving households column **NPISH**, the general government final consumption column **GGFC**, the gross fixed capital formation column **GFCF**, the changes in inventories and valuables column **INVNT**, and direct purchases abroad by residents **DPABR**. The row-sum of a country/sector pair i satisfies

$$\sum_c \mathrm{Z}_{ic} + \mathrm{HFCE}_i + \mathrm{NPISH}_i + \mathrm{GGFC}_i + \mathrm{GFCF}_i + \mathrm{INVNT}_i + \mathrm{DPABR}_i = \mathrm{X}_i \quad (6.1)$$

and the column-sum of a country/sector pair i satisfies

$$\sum_r \mathrm{Z}_{ri} + \mathrm{TLS}_i + \mathrm{VA}_i = \mathrm{X}_i. \quad (6.2)$$

We include consumption by non-profit institutions serving households in final consumption by households, denoted as **HFCE**. For simplicity, we set the changes in the inventories and valuables column **INVNT** to 0 and include that (small) offset in the gross fixed capital formation column **GFCF**. We also set the direct purchases abroad by residents column **DPABR** to 0 and include that offset in household final consumption **HFCE**. We split the value-added row **VA** into an employee compensation row **LAB**[9] and a capital compensation row **CAP**, so that

$$\mathrm{VA}_i = \mathrm{LAB}_i + \mathrm{CAP}_i \quad (6.3)$$

for all country/sector pairs i. We also split[10] the gross fixed capital formation column **GFCF** by sector into investment by firms **FGFCF**, by households **HGFCF**, and by the general government **GGFCF**, so that

$$\mathrm{GFCF}_i = \mathrm{FGFCF}_i + \mathrm{HGFCF}_i + \mathrm{GGFCF}_i \quad (6.4)$$

for all country/sector pairs i.

Additional data for initialising the model comes from the IMF[11], the World Bank[12], Compustat[13], the European Central Bank (ECB)[14], and the Bank for International Settlements (BIS)[15]. Table 6.2 summarises these data sources.

[8] OECD: Inter-country input-output tables (Code: *ICIO*).

[9] OECD Data Archive (Code: *F_DP_LIVE*). This is consistent with the socio-economic accounts of the World Input-Output Database (Timmer et al., 2015).

[10] OECD: Annual investment by asset and institutional sector (Code: *DSD_NASEC10 @ DF_TABLE14_GFCF*).

[11] See https://www.imf.org/en/Data.

[12] See https://data.worldbank.org/.

[13] See https://www.lseg.com/en/data-analytics/financial-data/company-data/fundamentals-data/standardized-fundamentals/sp-compustat-database.

[14] See https://www.ecb.europa.eu/stats/ecb_surveys/hfcs/html/index.en.html.

[15] See https://data.bis.org/.

Dataset	Usage
OECD	Input-Output tables, balance sheets, business demography
IMF	Quarterly time series of economic aggregates
World Bank	Taxes, government spending, government revenue
BIS	Policy rates
ECB	Microdata on households and individuals
Compustat	Microdata on firms and banks

Table 6.2: Data sources used for model initialisation and calibration: from the OECD, the IMF, the World Bank, the BIS, the ECB, and Compustat (S&P).

Running the agent-based model consists of updating agent-level information sequentially. Each iteration represents a discrete step in the evolution of our economic system. The following section describes the various steps involved in each of these iterations, outlining how agent decisions and interactions drive the system forward over time.

6.2.3 Sequence of Events

The timeframe corresponding to each iteration is one quarter, reflecting the frequency of the macroeconomic data utilised in the model. During these periods, agents make decisions based on partial information and use heuristic "rules of thumb". The type of heuristics includes, for example, those used by firms to determine their target production levels by considering factors such as inflation, historical demand or input costs. These rules are formulated into functions that translate market conditions and agent-specific information into decisions, such as production targets or hiring strategies. Similarly, market mechanisms, like those governing the labour or goods markets, are modelled through matching functions that aim to align supply and demand but typically do not result in a perfect match. These functions are flexible enough to be adapted or replaced, depending on the context. We outline the different intermediate steps of each iteration below. Overall, the steps described are similar to those in Poledna et al. (2023a).[16] The full model description is provided in Section 6.4 and Section 6.5.

6.2.3.1 Estimation Phase

This step lays the foundation for each iteration of our simulations. In this phase, agents generate forecasts based on historical data and predictive models, which inform their decisions in the upcoming quarter. By the end of this phase, all agents

[16]We outline the key differences between our model and the IIASA model in Section 6.3.

have expectations about future economic conditions, which guide their choices in the *Target Setting Phase* (Section 6.2.3.2).

Central Economic Forecasts This step constitutes the core of this phase, as it includes forecasts for various economic conditions, such as inflation across different markets (consumer prices, producer prices, and housing prices) and overall economic growth rates. These forecasts are based on historical data and are designed to provide a common set of expectations about the future state of the economy. This centralised forecast provides a shared baseline that agents can build upon or deviate from when forming their own individual projections. For detailed explanations, see Section 6.2.5.

Firm-Specific Forecasts Following the establishment of central economic forecasts, individual firms create their own expectations about the economy's future. Each firm takes into account the history of supply availability, demand for its products, and price movements to project its own future growth (Equation (6.60)), forecast future demand for its products (Equation (6.61)), and forecast future profits (Equation (6.62)).

6.2.3.2 Target Setting Phase

During this phase, agents make strategic decisions based on earlier forecasts. This phase lays the groundwork for future market interactions. It determines a considerable portion of the demand that will influence transactions in the Goods Market and sets the stage for the subsequent Labour Market Phase.

Firm Decisions Firms begin by setting their production level objective (Equation (6.63)). They consider several key factors, including the future market conditions forecasted in the *Estimation Phase* and, more specifically, their projections of future demand and their current finished goods inventory. Each firm also evaluates its financial health, including debt levels and equity, which helps determine its capacity for expansion. After establishing the production target, firms set objectives for labour market activities (Equation (6.69)), defining their hiring strategies and wage offers. Firms also assess their procurement needs based on their current stock of intermediate inputs.

Individual Decisions Individuals in the economy set their reservation wages — the minimum wage they are willing to accept from a potential employer. This decision is informed by their personal wage history and employment conditions.

6.2.3.3 Labour Market Phase

This phase serves as a clearinghouse where the supply of labour from individuals is matched with the demand for labour from firms, considering both offered wages and individual reservation wages. At the same time, firms may lay off employees if necessary to meet their objectives, and individuals may also leave their jobs voluntarily, based on a certain probability. Once these changes occur, the states of both firms and individuals are updated to reflect the new employment conditions, with firms' employee rosters and individuals' employment statuses and incomes adjusting accordingly. This process is vital in determining firms' production capacities, which are linked to their workforce size, and the incomes available to households. The detailed functioning of the labour market is explained in Section 6.5.2.

6.2.3.4 Planning Metrics Phase

In this phase, agents revise their expectations and set further targets in response to changes in unemployment. This section is, therefore, crucial for preparing the housing, goods, and credit markets.

Government Decisions The government updates its social welfare expenditures in response to changes in unemployment and inflation expectations (Equation (6.97) and Equation (6.98)). The government also sets its target level of consumption for goods and services depending on economic conditions, forecasts, and its historical spending trajectory.[17]

Central Bank Decisions The Central Bank sets its policy rate based on current inflation trends (Equation (6.46)).

Household Decisions Households update their information about their current income based on changes in their employment status and changes in social benefits. With updated income estimates, they can determine their spending capacity and adjust their demand for housing and goods and services (Equation (6.107) and Equation (6.136)).

[17]Note, however, that in our implementation, this last step is done in the Goods Market Phase without impacting the ordering of the steps described here.

Firm Production Knowing their output capacity, which is determined by their workforce, their stock of intermediate inputs, and their capital stock, and guided by the target level described above, firms carry out their production (Equation (6.73)). They set their selling price (Equation (6.74)) depending on inflation forecasts and previous market conditions. This sets the available supply of each good.

6.2.3.5 Housing Market Phase

This phase involves matching sellers and buyers of housing and matching renters with landlords. The housing market is detailed in Section 6.5.4. It is essentially a reimplementation of the model described by Geanakoplos et al. (2012) and Carro et al. (2023), and it consists of two distinct phases:

Preparation Phase First, the valuation of each property is adjusted depending on market conditions. Households also update their priors on the amount they are willing to pay to buy or rent property and decide whether they want to own or rent. Landlords also decide whether to update the rent on their property depending on inflation.

Clearing Phase With the volume and prices set for both supply and demand in the housing market, the market is cleared by matching properties up for sale with prospective buyers and properties up for rent with renters. The wealth and housing expenses of households are updated accordingly. Note that a household attempting to buy a property may need to obtain a mortgage in the next phase for the transaction to happen.

6.2.3.6 Credit Market Phase

The credit market is detailed in Section 6.5.4. This phase also consists of two steps:

Preparation Phase Firms and households set their credit targets by studying their financial health and economic forecasts. Banks set the interest rates they offer on different loan types (mortgages and consumption loans to households, loans for firms) as a markup on the current Central Bank rate. This markup may also depend on the borrower's financial health.

Clearing Phase The demand for loans from households and firms is matched with the loans offered by banks. The transactions of households that apply for and obtain mortgages are executed, and the corresponding real estate is updated to be in their name. If mortgages are not approved, no transaction is registered.

6.2.3.7 Goods Market Phase

In contrast to the phases described above, which all happen within a country, this phase involves the global economy and matches the global supply and demand of goods. The goods market is described in Section 6.5.1; the phase again involves two steps:

Preparation Phase The supply and demand of agents is aggregated into an order book in USD, requiring demand in local currencies to be converted to demand for goods in USD using the current exchange rates.

Clearing Phase The order book defined above is cleared by matching buyers and sellers. Note that the degree to which agents prefer to buy goods domestically rather than in the international market is a parameter that can be adjusted. We use international trade data to match flows of goods and services in global markets realistically.

6.2.3.8 Realised Metrics Phase

Here, agents update the actual outcomes of their operations and financial activities over the past quarter. This phase is crucial for comparing predicted and realised outcomes and making the necessary updates for future planning.

Firm Decisions After subtracting the amounts sold, firms update their records to reflect their inventory levels (Equation (6.86)). They also register the realised demand (Equation (6.89)), which will be used to predict demand for the next quarter. Firms also review their financial health (Equation (6.90)) and analyse their solvency after updating their financial data with the proceeds from their sales. The financial records of firms and their bank accounts are also updated. Firms under stress may file for bankruptcy.

Bank Decisions Banks compute their profits (Equation (6.41)), liabilities (Equation (6.43)), and reserves (Equation (6.44)) after all the operations described above, including the consequences of possible borrowers' defaults. Banks under stress may also go bankrupt.

Household Decisions Households update their income (Equation (6.108)), wealth (Equation (6.129)) and debt (Equation (6.130)) after operations in the goods and credit market.

Government Decisions The government updates its revenue (Equation (6.99)), deficit (Equation (6.100)), and debt (Equation (6.101)).

6.2.4 Economic Aggregates

In this section, we provide an overview of the economic aggregates in the model, all resulting from the endogenous behaviour of agents. Table 6.3 provides a summary; the aggregates related to production (Section 6.2.4.1), prices (Section 6.2.4.2), housing (Section 6.2.4.3), and credit (Section 6.2.4.4) are summarised below. Finally, we provide a full GDP identity in Section 6.2.4.5.

6.2.4.1 Production

The total/sectoral real gross output of domestically producing firms is set to

$$Y(t) = \sum_{f \in \mathcal{F}} Y_f(t) \tag{6.5}$$

$$Y_s(t) = \sum_{f \in \mathcal{F}_s} Y_f(t) \tag{6.6}$$

where $\mathcal{F}$ denotes the set of firms, $\mathcal{F}_s$ the set of firms operating in sector s, and $Y_f(t)$ real production of firm f.

6.2.4.2 Price

The *producer price index* $P^{\text{PPI}}(t)$ and the *consumer price index* $P^{\text{CPI}}(t)$ are updated as[18]

$$P^{\text{PPI}}(t) = \frac{\sum_{f \in \mathcal{F}} P_f(t)(Y_f(t) + S_f(t-1)) + \text{IMP}^{\text{N}}(t)}{\sum_{f \in \mathcal{F}} (Y_f(t) + S_f(t-1)) + \text{IMP}^{\text{R}}(t)} \tag{6.7}$$

$$P^{\text{CPI}}(t) = \sum_s b_s^{\text{CPI}} P_s(t) \tag{6.8}$$

where $\mathcal{F}$ is the set of firms, $P_f(t)$ is the price set by firm f, $Y_f(t)$ is real output of firm f, $S_f(t-1)$ are real finished goods inventories of firm f, $\text{IMP}^{\text{N}}(t)$ / $\text{IMP}^{\text{R}}(t)$ are total imports in nominal / real terms, b_s^{CPI} are aggregate household consumption weights, and $P_s(t)$ denotes the price index for goods produced in sector s,

$$P_s(t) = \frac{\sum_{f \in \mathcal{F}_s} P_f(t)(Y_f(t) + S_f(t-1)) + \text{IMP}_s^{\text{N}}(t)}{\sum_{f \in \mathcal{F}_s} (Y_f(t) + S_f(t-1)) + \text{IMP}_s^{\text{R}}(t)} \tag{6.9}$$

where $\text{IMP}_s^{\text{R}}(t)$ are imports of sector s in real terms and $\text{IMP}_s^{\text{N}}(t)$ are imports of sector s in nominal terms.

[18] From looking at Equations (6.7)–(6.8), $P^{\text{PPI}}(t)$ and $P^{\text{CPI}}(t)$ look more like weighted price averages than price indexes. However, as explained in Section 6.4.3.2, we set all initial prices to 1, so $P^{\text{PPI}}(t)$ and $P^{\text{CPI}}(t)$ can be correctly interpreted as price indexes.

Category	Description	Notation
Sets	Sectors	$\mathcal{S}$
Production	Total real production	$Y(t)$
	Predicted total real production	$\overline{Y}(t)$
	Sectoral real production	$Y_s(t)$
	Predicted sectoral real production	$\overline{Y}_s(t)$
Price	Sectoral price index	$P_s(t)$
	Producer price index	$P^{\mathrm{PPI}}(t)$
	Predicted producer price index	$\overline{P}^{\mathrm{PPI}}(t)$
	Consumer price index	$P^{\mathrm{CPI}}(t)$
	Predicted consumer price	$\overline{P}^{\mathrm{CPI}}(t)$
	House price index	$P^{\mathrm{HPI}}(t)$
	Predicted house price index	$\overline{P}^{\mathrm{HPI}}(t)$
	Rental price index	$P^{\mathrm{RPI}}(t)$
	Predicted rental price index	$\overline{P}^{\mathrm{RPI}}(t)$
	Producer price index inflation	$\pi^{\mathrm{PPI}}(t)$
	Predicted producer price index inflation	$\overline{\pi}^{\mathrm{PPI}}(t)$
	Consumer price index inflation	$\pi^{\mathrm{CPI}}(t)$
	Predicted consumer price index inflation	$\overline{\pi}^{\mathrm{CPI}}(t)$
	House price index inflation	$\pi^{\mathrm{HPI}}(t)$
	Predicted house price index inflation	$\overline{\pi}^{\mathrm{HPI}}(t)$
	Rental price index inflation	$\pi^{\mathrm{RPI}}(t)$
	Predicted rental price index inflation	$\overline{\pi}^{\mathrm{RPI}}(t)$
Credit	Total debt	$L(t)$
	Total debt of firms operating in sector s	$L_s^{\mathrm{L}}(t)$
	Total household debt in consumption loans	$L^{\mathrm{C}}(t)$
	Total mortgage debt	$L^{\mathrm{M}}(t)$
	Ratio of nonperforming firm loans in sector s	$\nu_s^{\mathrm{F}}(t)$
	Ratio of nonperforming household consumption loans	$\nu^{\mathrm{C}}(t)$
	Ratio of nonperforming mortgages	$\nu^{\mathrm{M}}(t)$

Table 6.3: Aggregate variables in the model, in particular related to production, prices, and credit.

6.2.4.3 Housing

The *house price index* is updated according to

$$P^{\mathrm{HPI}}(t) = \frac{\sum_{p\in\mathcal{P}} V_p(t)}{\sum_{p\in\mathcal{P}} V_p(0)} \tag{6.10}$$

where $\mathcal{P}$ is the set of properties and $V_p(t)$ is the value of property p at time t. The *rental price index* is updated according to

$$P^{\mathrm{RPI}}(t) = \frac{\sum_{p\in\mathcal{P}} r_p(t)}{\sum_{p\in\mathcal{P}} r_p(0)} \tag{6.11}$$

where $\mathcal{P}$ is the set of properties and $r_p(t)$ is the (imputed) rent of property p at time t.

6.2.4.4 Credit

The total amount of loans granted is computed as

$$L(t) = \sum_{b\in\mathcal{B}} L_b(t) \tag{6.12}$$

where $\mathcal{B}$ is the set of banks, and $L_b(t)$ is the total amount of loans granted by bank b. The total amount of loans granted to firms of sector s is given by

$$L_s^{\mathrm{F}}(t) = \sum_{f\in\mathcal{F}_s} L_f(t) \tag{6.13}$$

where $\mathcal{F}_s$ is the set of firms operating in sector s, and $L_f(t)$ the debt of firm f. The total amount of household consumption loans is given by

$$L^{\mathrm{C}}(t) = \sum_{h\in\mathcal{H}} L_h^{\mathrm{C}}(t) \tag{6.14}$$

where $\mathcal{H}$ is the set of households and $L_h^{\mathrm{C}}(t)$ the debt of household h in consumption loans. The total amount of mortgage debt is given by

$$L^{\mathrm{M}}(t) = \sum_{h\in\mathcal{H}} L_h^{\mathrm{M}}(t) \tag{6.15}$$

where $L_h^{\mathrm{M}}(t)$ the mortgage debt of household h.

The ratio of Non-Performing Loans (NPL) to firms in sector s is given by

$$\nu_s^{\mathrm{F}}(t) = \frac{\sum_{f\in\mathcal{F}'_s(t)} L_f(t)}{L_s^{\mathrm{F}}(t)} \tag{6.16}$$

where $\mathcal{F}'_s(t)$ is the set of firms operating in sector s that become insolvent at time t. The ratio of non-performing household consumption loans is given by

$$\nu^{\mathrm{C}}(t) = \frac{\sum_{h\in\mathcal{H}'(t)} L_h^{\mathrm{C}}(t)}{L^{\mathrm{C}}(t)} \tag{6.17}$$

where $\mathcal{H}'(t)$ is the set of insolvent households. Similarly, the ratio of non-performing mortgages is given by

$$\nu^{\mathrm{M}}(t) = \frac{\sum_{h\in\mathcal{H}'(t)} L_h^{\mathrm{M}}(t)}{L^{\mathrm{M}}(t)}. \tag{6.18}$$

6.2.4.5 GDP Identity

Nominal gross domestic product in the model can be calculated via the output approach, the expenditure approach, and the income approach:

$$
\begin{aligned}
&\underbrace{\sum_{s\in\mathcal{S}} \tau_s^{\text{PROD}} \sum_{f\in\mathcal{F}_s} P_f(t)Y_f(t) + \tau^{\text{VAT}} \sum_{h\in\mathcal{H}} C_h(t) + \tau^{\text{CF}} \sum_{h\in\mathcal{H}} K_h(t) + \tau^{\text{EXP}}\text{EXP}(t)}_{\text{Taxes on Products}} \\
&+ \underbrace{\sum_{s\in\mathcal{S}} \left(1-\tau_s^{\text{PROD}}\right) \sum_{f\in\mathcal{F}_s} P_f(t)Y_f(t)}_{\text{Production}} - \underbrace{\sum_{s\in\mathcal{S}}\sum_{s'\in\mathcal{S}} \frac{1}{m_{s's}} \sum_{f\in\mathcal{F}_s} P_s(t)Y_f(t)}_{\text{Intermediate inputs}} \\
&= \underbrace{\sum_{h\in\mathcal{H}} \left(1+\tau^{\text{VAT}}\right) C_h(t)}_{\text{Household consumption}} + \underbrace{\sum_{g\in\mathcal{G}} C_g(t)}_{\text{Government consumption}} + \underbrace{\left(1+\tau^{\text{EXP}}\right)\text{EXP}(t)}_{\text{Exports}} \\
&- \underbrace{\text{IMP}(t)}_{\text{Imports}} + \underbrace{\left(1+\tau^{\text{CF}}\right) \sum_{h\in\mathcal{H}} K_h(t) + \sum_{s\in\mathcal{S}}\sum_{f\in\mathcal{F}_s} P_s(t)K_f(t)}_{\text{Gross fixed capital formation}} \\
&+ \underbrace{\sum_{s\in\mathcal{S}} \left[\sum_{f\in\mathcal{F}_s} P_s(t)\Delta S_f(t) + \sum_{s'\in\mathcal{S}} P_{s'}(t)\left(\Delta M_{fs'}(t) - \frac{1}{m_{s's}}Y_f(t)\right)\right]}_{\text{Changes in stocks and inventories}} \\
&= \underbrace{\sum_{s\in\mathcal{S}} \tau_s^{\text{PROD}} \sum_{f\in\mathcal{F}_s} P_f(t)Y_f(t) + \tau^{\text{VAT}} \sum_{h\in\mathcal{H}} C_h(t) + \tau^{\text{CF}} \sum_{h\in\mathcal{H}} K_h(t) + \tau^{\text{EXP}}\text{EXP}(t)}_{\text{Taxes on Products}} \\
&+ \underbrace{\sum_{s\in\mathcal{S}} \left(1-\tau_s^{\text{PROD}}\right) \sum_{f\in\mathcal{F}_s} P_f(t)Y_f(t) - \sum_{i\in\mathcal{I}^{\text{E}}(t)} w_i(t) - \sum_{s\in\mathcal{S}}\sum_{f\in\mathcal{F}_s}\sum_{s'\in\mathcal{S}} P_{s'}(t)\frac{1}{m_{s's}}Y_f(t))}_{\text{Gross operating surplus and mixed income}} \\
&+ \underbrace{\sum_{i\in\mathcal{I}^{\text{E}}(t)} w_i(t)}_{\text{Compensation of employees}}
\end{aligned}
\tag{6.19}
$$

where $\mathcal{F}_s$ is the set of firms operating in sector s, $\mathcal{H}$ is the set of households, $\mathcal{G}$ is the set of government entities, and $\mathcal{I}^{\text{E}}(t)$ is the set of employed individuals at time t. In addition, τ_s^{PROD} are tax rates on production in sector s, τ^{VAT} is the value-added tax rate on household consumption $C_h(t)$, τ^{CF} is the tax rate on capital formation $K_h(t)$ of households, and τ^{EXP} is the tax rate on exports $\text{EXP}(t)$. Finally, $P_f(t)$ is the price set by firm f, $Y_f(t)$ is the real production of firm f, $K_f(t)$ is the real capital investment of firm f, $\text{IMP}(t)$ are total imports, $P_s(t)$ are average sectoral prices, $\Delta S_f(t)$ changes in real inventories of firm f, $\Delta M_{fs}(t)$ changes in real intermediate inputs of sector s, $m_{s's}$ denotes the real amount of intermediate inputs of sector s' necessary to produce one real unit of output of sector s, and $w_i(t)$ is the wage paid to individual i.

6.2.5 Expectations

We conclude the overview of our model by explaining the role of expectations. Section 6.2.5.1 summarizes the general approach, which is applied to obtain forecasts for growth expectations (Section 6.2.5.2) and inflation (Section 6.2.5.3).

6.2.5.1 Modeling Expectations

The standard approach to modelling expectations in DSGE models is to use rational, or model-consistent, expectations, where the expectation operator coincides with the realised future value of the variables, leading to a fixed point dynamical equation that, in most cases, must be log-linearised around the steady state to be solved. In ABMs, we cannot analytically aggregate the expectations of heterogeneous interacting agents.[19]

When deciding which expectation rules we should consider, we are faced with a choice between simple and more complicated decision rules. Simple decision rules that are less likely to lead to overfitting are often optimal (Artinger et al., 2022), and we employ a deterministic autoregressive process with lag one for expectation formation. Hence, output stochasticity for our model results from the random search-and-matching processes on the goods, labour, credit, and housing markets. This distinguishes our model from the IIASA model (Poledna et al., 2023a), which assumes an autoregressive model with lag one (AR(1)) process for expectation formation.[20] The autoregressive coefficient is recalibrated at every timestep using a mix of real-world data (from 2000-Q1 up to the initialisation time) and simulation output (from the initialisation time until the previous timestep) for the relevant variable. Therefore, the length of the time series used for calibration increases with every timestep. The following details the data we use to inform expectation setting.

6.2.5.2 Expectations on Growth

Forecasts for total real gross output $\bar{Y}(t)$ and total real sectoral gross output $\bar{Y}_s(t)$ are used by every agent in updating their decisions. Realised real gross output and total real sectoral gross output enter the autoregressive process in log levels.[21] Predicted

[19] Rational expectations are also challenging for DSGE models that incorporate heterogeneity, such as Heterogeneous Agent New Keynesian (HANK) models. See, for instance, Moll (2024).

[20] Most ABMs in the literature employ linear forecasts. Other approaches are discussed in Brayton et al. (1997).

[21] OECD: Quarterly GDP and Components - Output Approach (Code: *DSD_NAMAIN1 @ DF_QNA_BY_ACTIVITY_OUTPUT*).

growth in real gross output and predicted sectoral real gross output growth are set as

$$\bar{\gamma}(t) = \log \frac{\bar{Y}(t)}{Y(t-1)} \tag{6.20}$$

$$\bar{\gamma}_s(t) = \log \frac{\bar{Y}_s(t)}{Y_s(t-1)} \tag{6.21}$$

where $Y(t-1)$ is the previous realised total gross output and $Y_s(t-1)$ is the previous realised total gross output of sector s.

6.2.5.3 Expectations on Inflation

Similarly to growth, every agent uses global PPI / CPI / HPI / RPI forecasts to update their decisions.[22] These forecasts are again obtained using autoregressive models given previous PPI[23], CPI[24], HPI/RPI[25]. Then, predicted inflation is set as

$$\overline{\pi}^{\text{PPI}}(t) = \log \frac{\overline{P}^{\text{PPI}}(t)}{P^{\text{PPI}}(t-1)} \tag{6.22}$$

$$\overline{\pi}^{\text{CPI}}(t) = \log \frac{\overline{P}^{\text{CPI}}(t)}{P^{\text{CPI}}(t-1)} \tag{6.23}$$

$$\overline{\pi}^{\text{HPI}}(t) = \log \frac{\overline{P}^{\text{HPI}}(t)}{P^{\text{HPI}}(t-1)} \tag{6.24}$$

$$\overline{\pi}^{\text{RPI}}(t) = \log \frac{\overline{P}^{\text{RPI}}(t)}{P^{\text{RPI}}(t-1)} \tag{6.25}$$

where $\overline{P}^{\text{PPI}}(t)$ is predicted PPI, $P^{\text{PPI}}(t-1)$ is previous PPI, $\overline{P}^{\text{CPI}}(t)$ is predicted CPI, $P^{\text{CPI}}(t-1)$ is previous CPI, $\overline{P}^{\text{HPI}}(t)$ is predicted HPI, $P^{\text{HPI}}(t-1)$ is previous HPI, $\overline{P}^{\text{RPI}}(t)$ is predicted RPI, $P^{\text{RPI}}(t-1)$ is previous RPI.

6.3 Key Differences to the IIASA Model

The key differences between the IIASA model (Poledna et al., 2023a) and our model are summarized in Table 6.4. While there are many similarities in terms of agents and structure, the main differences are that our model can be calibrated to many countries, features a housing market[26], and explicitly distinguishes between

[22] One reason for using global forecasts is that it is unclear whether forecasting inflation improves when considering individual components (Roma et al., 2004).

[23] IMF: International Financial Statistics (Code: *PPPI_IX*).

[24] IMF: International Financial Statistics (Code: *PCPI_IX*).

[25] OECD: Analytical House Prices Indicators (Code: *DSD_AN_HOUSE_PRICES @ DF_HOUSE_PRICES*)

[26] See Section 6.5.4.

households and individuals. The detailed breakdown of household income, wealth, and debt leads to heterogeneous consumption.

Both models employ a Leontief production function in labour, intermediate, and capital inputs, but we do not assume perfect substitutability between intermediate and capital inputs. Instead, we have a Leontief production function for all inputs. Firms take liquidity shortages into account when setting production and purchasing targets. They dynamically adjust wages based on labour market tightness.

Instead of a single bank, our model features multiple, heterogeneous banks with different risk and credit supply profiles. Banks employ realistic credit conditions and grant four types of loans to firms and households instead of just one type of loan to firms. Banks may fall into bankruptcy and will be bailed in.

	Category	IIASA Model	Our Model
Scope	Coverage	One country	Multiple countries
	Markets	Goods, Labour, Credit	Goods, Labour, Credit, Housing
Ind. & Hhs.	Matching	One-to-one	Multiple individuals living in the same household
	Income	Wages, benefits, dividends	Wages, benefits, rental income, investment profits
	Wealth	Deposits, properties	Deposits, properties, other real assets, other financial assets
	Debt	Overdrafts	Consumption loans, mortgages, overdrafts
	Consumption	Fixed fraction of income	Heterogeneous fraction of income, consumption smoothing, minimal consumption
	Reservation Wage	No	Yes
Firms	Production	Nested Leontief that allows substitution between intermediate goods and capital goods	Nested Leontief
	Target setting	Only buffer	Active management
	Wages	Static	Dynamic

Banks	Number of banks	1	Multiple
	Loans recipients	Firms, households (only overdrafts)	Firms, households (overdrafts, consumption loans, mortgages)
	Bankruptcy	Cannot fail	Bail-in mechanism
	Supply-side loan requirements	Capital requirement	Capital requirement and risk diversification
	Demand-side loan requirements	Loan-to-value	Loan-to-value, return on equity/assets (firms), loan-to-income/debt-service-to-income (households)
	Interest rates	Fixed markup on the central bank policy rate	Statistical model that considers pass-through from the central bank, inflation and non-performing loans
Gov.	Benefits	Unemployment benefits grow with the economy	Countercyclical unemployment benefits

Table 6.4: Key differences of our model to the IIASA Model in terms of scope and agent behaviour.

6.4 Agents

We now turn to an in-depth description of the agents in the model, providing details on initial conditions, parameters, behavioural rules, and data. Specifically, of a given country, we describe banks (Section 6.4.1), the central bank (Section 6.4.2), firms (Section 6.4.3), the government and government entities (Section 6.4.4), households (Section 6.4.5), individuals (Section 6.4.6), and the "rest of the world" (Section 6.4.7). For each agent, we also state the alternative modelling choices in the literature, including the models reviewed in Chapter 5.

6.4.1 Banks

In our model, banks set interest rates, hold central bank reserves, household and firm deposits, and grant loans to firms (long-term and short-term) and households

(consumption loans and mortgages).[27] When firms or households they lent to go bankrupt, they take losses. For simplicity, we do not explicitly model people employed by banks. We also assume that banks do not have investment strategies in derivatives and real assets. Banks may fail and will be bailed in by their creditors.[28]

Table 6.5 lists variables and parameters related to banks, and Table 6.6 lists variables and parameters related to loans granted by banks.

6.4.1.1 Initial Conditions

In this section, we briefly summarize the initial conditions for banks.

Drawing Banks from Compustat Data Banks are sampled with replacement from Compustat data[29] so that the total number of banks matches IMF aggregates[30]. This includes total deposits (Code: *dptcq*) and total liabilities (Code: *ltq*). The maximum number of banks recorded in Compustat from 2013-Q1 to 2017-Q4 is 3 for Australia, 7 for Canada, 3 for Chile, 2 for Columbia, 1 for Denmark, 2 for France, 1 for Germany, 1 for Greece, 2 for Ireland, 1 for Italy, 3 for Japan, 3 for Korea, 2 for Spain, 1 for Sweden, 7 for the UK, and 467 for the United States. For countries that are not included in the Compustat data for banks, we follow the approach in Poledna et al. (2023a) and create a single bank which is calibrated to aggregates.

Initial Profits Initial profits of bank b (see Equation (6.41)) are set according to

$$\begin{aligned}\Pi_b(0) = & \underbrace{\sum_{l\in\mathcal{L}_b(0)} r_l V_l(0)}_{\text{Interest received on loans}} + \underbrace{r_b^{\text{F-O}}(0)\sum_{f\in\mathcal{F}_b(0)}[D_f(0)]^- + r_b^{\text{H-O}}(0)\sum_{h\in\mathcal{H}_b(0)}[D_h(0)]^-}_{\text{Interest received on overdrafts}} \\ & + \underbrace{r(0)\,[R_b(0)]^+}_{\text{Interest received on reserves}} - \underbrace{r(0)\left(\sum_{f\in\mathcal{F}_b}[D_f(0)]^+ + \sum_{h\in\mathcal{H}_b}[D_h(0)]^+\right)}_{\text{Interest paid on deposits}} \\ & - \underbrace{r(0)\,[R_b(0)]^-}_{\text{Interest paid on reserves}}\end{aligned} \tag{6.26}$$

which is the interest on loans granted to firms or households, plus the interest received on overdrafts, plus the interest received on reserves (if positive), minus

[27] Different models mainly vary in terms of lending conditions. In some models, all credit demands are satisfied (Seppecher, 2012; Mandel et al., 2010), others introduce probabilities of loan approvals (Ashraf et al., 2016; Assenza et al., 2015), or assume (regulatory) upper bounds on the volume of credit a bank can grant (Dawid et al., 2014; Dosi et al., 2010; Poledna et al., 2023a).

[28] In Assenza et al. (2015); Ashraf et al. (2016), banks will always be bailed out by the central government, in Dawid et al. (2014) by the central bank. In Seppecher (2012), the single bank's bankruptcy ends the simulation.

[29] CRSP/Compustat Merged Database - Bank Quarterly (Code: *crsp_a_ccm*).

[30] IMF: Financial Access Survey (Code: *FAS*).

Category	Description	Notation
Sets	Set of banks	$\mathcal{B}$
	Set of insolvent banks	$\mathcal{B}'(t)$
	Loans provided	$\mathcal{L}_b(t)$
	Firms with deposits at the bank	$\mathcal{F}_b(t)$
	Households with deposits at the bank	$\mathcal{H}_b(t)$
Assets	Reserves	$R_b(t)$
	Total loans granted to firms of sector s	$V_{bs}^{\mathrm{F}}(t)$
	Total consumption loans granted to households	$V_b^{\mathrm{C}}(t)$
	Total mortgages granted to households	$V_b^{\mathrm{M}}(t)$
Liabilities	Liability	$L_b(t)$
	Equity	$E_b(t)$
P&L Account	Profits	$\Pi_b(t)$
Credit Supply	Maximum credit willing to grant	$V_b^{\max}(t)$
	Supply of credit to firms of sector s	$\hat{V}_{bs}^{\mathrm{F}}(t)$
	Supply of credit to households for consumption	$\hat{V}_b^{\mathrm{C}}(t)$
	Supply of mortgage credit to households	$\hat{V}_b^{\mathrm{M}}(t)$
	Influence of the NPL ratio when allocating credit	ϕ^{CS}
Interest Rates	Interest rate on firm deposit overdrafts	$r_b^{\text{F-O}}(t)$
	Interest rate on household deposit overdrafts	$r_b^{\text{H-O}}(t)$
Regulatory Requirements	Bank capital adequacy ratio	ρ^{CAR}
	Solvency ratio between equity and assets	ρ^{SR}
	Firm loan debt-to-equity ratio	ρ^{DtE}
	Firm loan return-on-equity ratio	ρ^{RoE}
	Firm loan return-on-assets ratio	ρ^{RoA}
	Household consumption loan-to-income ratio	$\rho^{\text{LTI-C}}$
	Mortgage loan-to-value ratio	ρ^{LTV}
	Mortgage loan-to-income ratio	$\rho^{\text{LTI-M}}$
	Mortgage debt-service-to-income ratio	ρ^{DSTI}
	Equity injections for other insolvent banks	$I_b(t)$

Table 6.5: Variables and parameters in the model related to banks, in particular on balance sheets, the supply of credit, interest rate setting, and regularity requirements.

the interest paid on deposits or reserves (if negative). In the equation above, r_l is the interest rate on loan l, V_l its amount, $r_b^{\text{F-O}}(0)$ / $r_b^{\text{H-O}}(0)$ are the initial overdraft rates on firm/household deposits, and $r(0)$ is the initial central bank policy rate.

Agent	Category	Description	Notation
Loan l	Attributes	Principal amount	$V_l(t)$
		Interest rate	r_l
		Maturity	m_l

Table 6.6: Variables and parameters in the model related to loans, in particular principal amount, interest rate, and maturity.

Initial Reserves Initial central bank reserves (see Equation (6.44)) of bank b are set according to

$$R_b(0) = \underbrace{\sum_{f \in \mathcal{F}_b(0)} D_f(0)}_{\text{Firm deposits}} + \underbrace{\sum_{h \in \mathcal{H}_b(0)} D_h(0)}_{\text{Household deposits}} + \underbrace{E_b(0)}_{\text{Bank equity}} - \underbrace{\sum_{l \in \mathcal{L}_b(0)} V_l(0)}_{\text{Loans granted}} \tag{6.27}$$

which is the total amount of firm and household deposits held at the bank, plus the bank's equity, minus the total value of loans granted to firms or households. Here, $\mathcal{F}_b(0)/\mathcal{H}_b(0)$ denote the initial set of firms/households with deposits at bank b, and $\mathcal{L}_b(0)$ represents the total initial set of loans granted by the bank.

Initial Equity Initial bank equity (see Equation 6.42) is set proportionally to the total amount of loans granted and then rescaled so that the sum of the equity of all banks matches the OECD aggregates.[31] The rescaling ensures that the overall equity in the model aligns with the real-world financial data reported by the OECD.

Initial Matching with Firms Banks are initially matched with firms by solving a linear sum assignment problem. This minimises the sum of the differences between the total firm deposits and debt, and the total bank deposits and debt in short-term and long-term firm loans. In other words, at each model initialisation, we construct a bipartite graph between banks and firms to ensure that their respective initial financial states are as closely matched as possible.

Initial Matching with Households Similarly, banks are initially matched with households by solving a linear sum assignment problem. This minimises the sum of the differences between the total household deposits and debt, and the total bank deposits and debt in household consumption loans and mortgages. This ensures that household financial states are aligned with the banks' initial conditions at the start of the model.

[31] OECD: Quarterly Financial Balance Sheets (Code: DSD_NASEC20 @ DF_T720R_Q), non-consolidated equity (Code: *F5*) and financial net worth (Code: *BF90*) of monetary financial institutions other than the central bank (Code: *S12T*).

6.4.1.2 Parameters

Bank parameters include regulatory requirements, loan maturities, and credit supply parameters.

Regulatory Requirements We assume each bank's capital adequacy ratio to be $\rho^{\text{CAR}} = 0.08$, corresponding to Basel III, and a solvency ratio between equity and assets to be $\rho^{\text{SR}} = 0.1$.[32] For firm loans, we assume the firm loan debt-to-equity ratio to be $\rho^{\text{DtE}} = 1.0$, the firm loan return-on-equity ratio to be $\rho^{\text{RoE}} = 0.15$, and the firm loan return-on-assets ratio to be $\rho^{\text{RoA}} = 0.05$. For household consumption loans, we assume the household consumption loan to income ratio to be $\rho^{\text{LTI-C}} = 0.36$. For mortgages ($\rho^{\text{LTV}}$, $\rho^{\text{LTI-M}}$, ρ^{DSTI}), we follow the *Overview of national macroprudential measures*[33] compiled by the *European Systemic Risk Board* (ESRB). Alternatively, Muellbauer and Kaszowska-Mojsa (2023) compiled regulatory requirements over time for the 28 EU countries.

Loan Maturities Firm short-term loans are assumed to mature in one quarter, long-term loans in two years, household consumption loans in one quarter, and mortgages in 25 years. This follows standard maturity limits set by central banks and compiled by the ESRB.

Credit Supply The influence of the ratio of nonperforming loans (NPL) when allocating credit supply ϕ^{CS} is set to be 2.0, see Equation (6.37). This leads banks to optimise their credit supply away from sectors that are in crisis and towards sectors with less risk.

6.4.1.3 Rules

In this section, we discuss the behavioural rules for each bank, most notably the supply of credit, setting interest rates, and lending requirements.

Interest Rates Setting Bank interest rates are set using a single equation error correction model derived from an Autoregressive Distributed Lags (ARDL)

[32] Compustat does not provide microdata on regulatory requirements of individual banks; see Section 6.4.1.1. Therefore, we assume the same capital adequacy ratio for each bank. Future work could use Moody's *BankFocus*.

[33] https://www.esrb.europa.eu/national_policy/html/index.en.html

model.[34,35] We estimate a time-series model specified as

$$\begin{aligned}\Delta r_l(t) = &\underbrace{\phi^{\mathrm{EC}}\left(r_l(t-1) - \phi^{\mathrm{LR}} r(t)\right)}_{\text{Error correction mechanism}} + \underbrace{\sum_{j=1}^{p-1} \alpha_j \Delta r_l(t-j)}_{\text{Lagged interest rate}} + \underbrace{\sum_{j=0}^{q-1} \beta_j \Delta r(t-j)}_{\text{Lagged policy rate}} \\ &+ \underbrace{\sum_{j=0}^{r-1} \gamma_j \Delta \pi^{\mathrm{PPI}}(t-j)}_{\text{Lagged inflation}} + \underbrace{\sum_{j=0}^{s-1} \delta_j \Delta \nu(t-j)}_{\text{Lagged NPL ratio}} + \mu + \varepsilon_t\end{aligned} \tag{6.28}$$

for each type of loan l. The first term is the error correction mechanism, ϕ^{LR} is the long-run pass-through between the central bank policy rate $r(t)$[36] and the interest rate $r_l(t)$[37], ϕ^{EC} is the error-correction term. We additionally consider PPI inflation[38] and the ratio of non-performing loans[39] ($\nu_s(t)$ for loans to firms in sector s, $\nu^{\mathrm{C}}(t)$ for household consumption loans, $\nu^{\mathrm{M}}(t)$ for mortgages), α_j, β_j, γ_j, δ_j are coefficients, and the choice of lag structure p, q, r, s is based on Akaike's Information Criterion. These parameters are estimated using historic real data at model initialisation for each country.

We further assume that the interest rate on firm or household deposits is equal to the policy rate, that the household deposits overdraft rate $r_b^{\text{H-O}}(t)$ is equal to the interest rate on household consumption loans, and that the firm deposits overdraft rate $r_b^{\text{F-O}}(t)$ is equal to the interest rate on short-term firm loans.

Firm Loans Lending Requirements For short- or long-term firm loans, banks consider three ratios limiting the maximum amount of a new loan:

1. **Debt to Equity**: The credit provided to a firm f cannot exceed

$$\begin{aligned}V_l(t) \leq &\rho^{\mathrm{DtE}} \underbrace{\sum_{s\in\mathcal{S}} P_s(t) K_{fs}(t)}_{\text{Capital stock value}} - \underbrace{L_f(t-1)}_{\text{Loans}} + \underbrace{[D_f(t-1)]^-}_{\text{Overdrafts}} \\ &\underbrace{+ r_b^{\text{F-O}}(t)\,[D_f(t-1)]^- - \sum_{l\in\mathcal{L}_f(t)} r_l V_l(t)}_{\text{Interest payments}}\end{aligned} \tag{6.29}$$

[34]For more extensive models, see Leroy and Lucotte (2016).

[35]The models in Poledna et al. (2023a) and Hommes et al. (2022) assume fixed markups on the policy rate. In Assenza et al. (2015), the interest rate increases with the borrower's leverage; in Seppecher (2012), the interest rate is fixed. In Olmez et al. (2024), the mortgage rate is set by a reinforcement learner attempting to keep the housing market "healthy" with a low median house price to income ratio and stable house prices.

[36]Obtained from the BIS central bank policy rates dataset.

[37]ECB: Corporates' total loans (Code: *MIR.M.U2.B.A2A.A.R.A.2240.EUR.N*), Household loans for consumption (Code: *MIR.M.U2.B.A2B.A.R.A.2250.EUR.N*), Household loans for house purchase (Code: *MIR.M.U2.B.A2C.A.R.A.2250.EUR.N*).

[38]IMF: International Financial Statistics (Code: *PPPI_IX*).

[39]World Bank Data: bank non-performing loans to gross loans (Code: *GFDD.SI.02*).

where ρ^{DtE} is a parameter, $\sum_{s\in\mathcal{S}} P_s(t)K_{fs}(t)$ is the current value of the capital stock of firm f, $L_f(t)$ is the current debt of the firm, $D_f(t)$ are firm deposits at its bank, $r_b^{\text{F-O}}(t)$ are firm overdraft rates, $\mathcal{L}_f(t)$ is the set of loans of firm f, r_l is the interest rate of loan l, and $V_l(t)$ is the amount of loan l.

2. **Return on Equity**: The credit provided to a firm f cannot exceed

$$V_l(t) \leq \underbrace{\sum_{s\in\mathcal{S}} P_s(t)K_{fs}(t)}_{\text{Capital stock value}} + \underbrace{D_f(t-1)}_{\text{Firm deposits}} - \underbrace{L_f(t-1)}_{\text{Firm debt}} - \frac{\overline{\Pi}_f(t)}{\rho^{\text{RoE}}} \tag{6.30}$$

where ρ^{RoE} is a parameter and $\overline{\Pi}_f(t)$ are predicted quarterly profits of firm f.

3. **Return on Assets**: Credit is only provided to a firm f if

$$\frac{\overline{\Pi}_f(t)}{L_f(t-1) + E_f(t-1)} \geq \rho^{\text{RoA}} \tag{6.31}$$

where ρ^{RoA} is a parameter.

Household Consumption Loans Lending Requirement Banks perform a household risk assessment based on their average income over the last six months (corresponding to the previous two timesteps). The credit provided to a household h cannot exceed

$$V_l(t) \leq \rho^{\text{LTI-C}} \underbrace{\frac{1}{2}\left(Y_h(t-2) + Y_h(t-1)\right)}_{\text{Average income over the last 6 months}} - \underbrace{L_h(t-1)}_{\text{Household debt}} \tag{6.32}$$

where $\rho^{\text{LTI-C}}$ is a parameter.

Mortgage Lending Requirements Each bank has its own policies for extending mortgages to households. Generally, banks consider three ratios:

1. **Loan-to-Value**: Bank b sets a maximum value for the principal the household can borrow. Since the down payment is the household's full financial wealth, the maximum loan value is restricted by

$$V_l(t) \leq \frac{\rho^{\text{LTV}}}{1-\rho^{\text{LTV}}} \underbrace{W_h^{\text{FA}}(t)}_{\text{Wealth in financial assets}} \tag{6.33}$$

where ρ^{LTV} is a parameter.

2. **Loan-to-Income**: The maximum loan value is restricted by

$$V_l(t) \leq \rho^{\text{LTI-M}} \underbrace{\frac{1}{2}\left(Y_h(t-2) + Y_h(t-1)\right)}_{\text{Average income over the last 6 months}} - \underbrace{L_h(t-1)}_{\text{Household debt}} \tag{6.34}$$

where $\rho^{\text{LTI-M}}$ is a parameter.

3. **Debt-Service-to-Income**: The maximum loan value is restricted by

$$V_l(t) \leq \rho^{\text{DSTI}} \underbrace{\frac{1}{2}\left(Y_h(t-2) + Y_h(t-1)\right)}_{\text{Average income over the last 6 months}} \frac{1-(1+r_l)^{-m_l}}{r_l} \tag{6.35}$$

where ρ^{DSTI} is a parameter, r_l is the mortgage rate offered by the bank, and m_l is the maturity of the mortgage.

Supply-Side Credit Constraints Each bank b also has supply-side constraints. The bank cannot grant loans above

$$V_b^{\max}(t) = \frac{E_b(t-1)}{\rho^{\text{CAR}}} - \underbrace{\sum_{l \in \mathcal{L}_b(t)} V_l(t)}_{\text{Loans granted}} \tag{6.36}$$

where $E_b(t-1)$ is previous bank equity, and ρ^{CAR} is a capital adequacy ratio. Each bank distributes the total amount of loans they're willing to grant among firm loans, household consumption loans, and mortgages based on the corresponding ratios of previous non-performing loans. This means that there are no specialised banks that only grant a specific type of loan, but instead, banks optimise their credit supply profile to shift away from sectors hit by crises.[40] Specifically,

$$\begin{aligned}
\hat{V}_{bs}^{\text{F}}(t) &\propto V_{bs}^{\text{F}}(0) \exp\left(-\phi^{\text{CS}} \nu_s^{\text{F}}(t-1)\right) \\
\hat{V}_b^{\text{C}}(t) &\propto V_b^{\text{C}}(0) \exp\left(-\phi^{\text{CS}} \nu^{\text{C}}(t-1)\right) \\
\hat{V}_b^{\text{M}}(t) &\propto V_b^{\text{M}}(0) \exp\left(-\phi^{\text{CS}} \nu^{\text{M}}(t-1)\right) \\
V_b^{\max}(t) &= \sum_{s \in \mathcal{S}} \hat{V}_{bs}^{\text{F}}(t) + \hat{V}_b^{\text{C}}(t) + \hat{V}_b^{\text{M}}(t)
\end{aligned} \tag{6.37}$$

where $V_{bs}^{\text{F}}(0)$, $V_b^{\text{C}}(0)$, $V_b^{\text{M}}(0)$ are the initial amounts of credit granted by bank b to firms of sector s,

$$V_{bs}^{\text{F}}(0) = \sum_{f \in \mathcal{F}_b(0) \cap \mathcal{F}_s} L_{fb}(0), \tag{6.38}$$

[40] Compustat does not provide microdata on credit supply profiles of individual banks; see Section 6.4.1.1.

where $L_{fb}(0)$ are initial loans from bank b to firm f; households for consumption,

$$V_b^{\mathrm{C}}(0) = \sum_{h \in \mathcal{H}_b(0)} L_{hb}^{\mathrm{C}}(0), \tag{6.39}$$

where $L_{hb}^{\mathrm{C}}(0)$ are initial consumption loans from bank b to household h; and mortgages

$$V_b^{\mathrm{M}}(0) = \sum_{h \in \mathcal{H}_b(0)} L_{hb}^{\mathrm{M}}(0). \tag{6.40}$$

where $L_{hb}^{\mathrm{M}}(0)$ are initial mortgage loans from bank b to household h. The parameter ϕ^{CS} governs the influence of the ratio of non-performing ratios in allocating the total supply of credit.

Profits Profits of bank b are updated according to

$$\begin{aligned}\Pi_b(t) = & \underbrace{\sum_{l \in \mathcal{L}_b(t)} r_l V_l}_{\text{Interest received on loans}} + \underbrace{r_b^{\text{F-O}}(t) \sum_{f \in \mathcal{F}_b(t)} [D_f(t)]^- + r_b^{\text{H-O}}(t) \sum_{h \in \mathcal{H}_b(t)} [D_h(t)]^-}_{\text{Interest received on overdrafts}} \\ & + \underbrace{r(t)\,[R_b(t)]^+}_{\text{Interest received on reserves}} - \underbrace{r(t) \left(\sum_{f \in \mathcal{F}_b} [D_f(t)]^+ + \sum_{h \in \mathcal{H}_b} [D_h(t)]^+ \right)}_{\text{Interest paid on deposits}} \\ & - \underbrace{r(t)\,[R_b(t)]^-}_{\text{Interest paid on reserves}} - \underbrace{I_b(t)}_{\text{Equity injection}}\end{aligned} \tag{6.41}$$

which is the interest on granted loans to firms or households plus the interest received on overdrafts plus the interest received on reserves (if positive) minus the interest paid on deposits or reserves (if negative). In the equation above, r_l is the interest rate on loan l, V_l its amount, $r_b^{\text{F-0}}(t)/r_b^{\text{H-0}}(t)$ are the overdraft rates on firm/household deposits, and $r(t)$ is the central bank policy rate. The equity injection $I_b(t)$ is for bailing in other insolvent banks.

Equity The equity of bank b is updated according to

$$\begin{aligned}E_b(t) = & \underbrace{E_b(t-1)}_{\text{Previous bank equity}} + \underbrace{\Pi_b(t)}_{\text{Bank profits}} - \underbrace{\tau^{\text{CORP}}\,[\Pi_b(t)]^+}_{\text{Corporate Income Taxes}} \\ & + \underbrace{\sum_{f \in \mathcal{F}_b(t) \cap \mathcal{F}'(t)} D_f(t) - \sum_{f \in \mathcal{F}'(t)} L_{fb}(t) + \sum_{h \in \mathcal{H}_b(t) \cap \mathcal{H}'(t)} W_h^{\mathrm{D}}(t) - \sum_{h \in \mathcal{H}'(t)} L_{hb}(t)}_{\text{Write-off of bad debt}}\end{aligned} \tag{6.42}$$

where $\mathcal{F}'(t)$ / $\mathcal{H}'(t)$ is the set of insolvent firms / households, $L_{fb}(t)$ / $L_{hb}(t)$ are loans granted by bank b to firm f / household h, $D_f(t)$ are deposits of firm f, and $W_h^{\mathrm{D}}(t)$ are current household deposits. Whenever a firm or a household goes

bankrupt, its bank appropriates what is left as a deposit but loses all the loan (i.e., it cannot resell the firm assets).

Liabilities The liabilities of bank b are updated as

$$L_b(t) = \underbrace{E_b(t)}_{\text{Bank equity}} + \underbrace{\sum_{f \in \mathcal{F}_b(t)} [D_f(t)]^+ + \sum_{h \in \mathcal{H}_b(t)} [D_h(t)]^+}_{\text{Deposits}} - \underbrace{[R_b(t-1)]^-}_{\text{Reserves}} \tag{6.43}$$

where $\mathcal{F}_b(t)/\mathcal{H}_b(t)$ is the set of firms/households that have their deposits at bank b.

Reserves The reserves of bank b at the central bank are obtained by

$$R_b(t) = \underbrace{\sum_{f \in \mathcal{F}_b(t)} D_f(t) + \sum_{h \in \mathcal{H}_b(t)} D_h(t)}_{\text{Deposits}} + \underbrace{E_b(t)}_{\text{Bank equity}} - \underbrace{\sum_{l \in \mathcal{L}_b(t)} V_l(t)}_{\text{Loans granted}} \tag{6.44}$$

where $\mathcal{F}_b(t)/\mathcal{H}_b(t)$ is the set of firms/households that have their deposits at bank b.

Insolvency If a bank has a low solvency ratio, it goes bankrupt. Specifically, if its ratio of equity to assets falls below a certain threshold

$$\frac{E_b(t)}{L_b(t) + [R_b(t)]^+} < \rho^{\text{SR}}, \tag{6.45}$$

then the bank is *bailed-in* by all other non-insolvent banks cancelling a fixed fraction of their debt until the equity of the insolvent bank is equal to the average equity of non-insolvent banks.

6.4.2 Central Bank

The central bank provides liquidity to banks and sets the central bank interest rate[41].[42] Table 6.7 lists variables and parameters related to the central bank.

6.4.2.1 Parameters

We assume target CPI inflation to be $\pi^\star = 0.02$.

[41] This is typically done using different variations of the Taylor rule, as originally proposed in Taylor (1993). For instance, in the models by Mandel et al. (2010) and Ashraf et al. (2016), the central bank sets the interest rate using a Taylor rule based on inflation and unemployment gaps. Meanwhile, in Poledna et al. (2023a), the central bank employs a version of the Taylor rule that follows Blattner and Margaritov (2010), which does not consider the output gap.

[42] In contrast, in the model by Dawid et al. (2014), the central bank may engage in quantitative easing.

Category	Description	Notation
Policy Rate	Policy Rate	$r(t)$
	Target CPI inflation	$\pi^\star$
	Autoregressive parameter	ρ
	Real equilibrium interest rate	$r^\star$
	Weight on inflation targeting	ξ^π
	Weight on economic growth	ξ^γ

Table 6.7: Variables and parameters in the model related to the central bank, in particular for setting the policy rate.

6.4.2.2 Rules

The central bank sets the policy rate.

Policy Rate Following Poledna et al. (2023a), a Taylor rule models the policy rate according to

$$\bar{r}(t) = \left[\rho\bar{r}(t-1) + (1-\rho)\left(r^\star + \pi^\star + \xi^\pi(\pi^{\text{CPI}}(t) - \pi^\star) + \xi^\gamma\gamma(t)\right)\right]^+ \tag{6.46}$$

where ρ is a measure for gradual adjustment of the policy rate, $r^\star$ is the real equilibrium interest rate, $\pi^\star$ is the inflation target by the central bank, ξ^π is the weight on inflation targeting, and ξ^γ the weight on economic growth. The parameters used in the Taylor rule are estimated at model initialisation using the central bank policy rate $r(t)$[43], CPI inflation $\pi^{\text{CPI}}(t)$[44], and total growth rates $\gamma(t)$[45]. The OLS regression is given by

$$r(t) = \alpha + \rho r(t-1) + \beta_\pi\left(\pi^{\text{CPI}}(t) - \pi^\star\right) + \beta_\gamma\gamma(t) + \varepsilon_t \tag{6.47}$$

where α is the intercept, ρ, β_π and β_γ are coefficients to be estimated and ε_t is an error term. From the Taylor rule equation, we obtain

$$r^\star = \frac{\alpha}{1-\rho} - \pi^\star \tag{6.48}$$

$$\xi^\pi = \frac{\beta_\pi}{1-\rho} \tag{6.49}$$

$$\xi^\gamma = \frac{\beta_\gamma}{1-\rho} \tag{6.50}$$

[43]Obtained from the BIS central bank policy rates dataset.
[44]IMF: International Financial Statistics (Code: *PCPI_IX*).
[45]IMF: International Financial Statistics (Code: *NGDP_SA_XDC*).

6.4.3 Firms

Macroeconomic agent-based models typically distinguish between consumption good firms, which use labour and capital to sell to households, and capital good firms, which use labour and sell to consumption good firms.[46] In our model, we do not impose such a formal distinction. Firms set the quantity and the price of the goods they produce by utilising labour, intermediate inputs, and capital inputs. Production quantities are set using standard inventory planning, and prices are determined as a dynamic markup on production costs. Firms also set the wages they pay to employees. Wages depend on whether a firm is able to fill all its vacancies or not.[47] Firms may face a financing gap and will seek short-term or long-term loans from banks to cover this deficit. Firms can go bankrupt and exit the market; for simplicity, a new firm immediately replaces the bankrupt one.[48]

Table 6.8 lists variables and parameters related to firms.

Category	Description	Notation
Sets	Set of firms	$\mathcal{F}$
	Set of insolvent firms	$\mathcal{F}'(t)$
	Set of firms operating in sector s	$\mathcal{F}_s$
	Set of insolvent firms operating in sector s	$\mathcal{F}'_s(t)$
	Individuals employed by the firm	$\mathcal{I}_f(t)$
	Loans taken out by the firm	$\mathcal{L}_f(t)$
Production	Real production	$Y_f(t)$
	Target real production	$\hat{Y}_f(t)$
	Real production limited by labour inputs	$H_f(t)$
	Real production limited by intermediate inputs	$M_f(t)$
	Real production limited by capital inputs	$K_f(t)$
	Predicted growth	$\overline{\gamma}_f(t)$
	Target inventory to production fraction	ϕ^{StY}
	Influence of labour inputs in limiting target production	χ^{H}
	Influence of interm. inputs in limiting target production	χ^{M}
	Influence of capital inputs in limiting target production	χ^{K}
Demand	Real demand	$Q_f(t)$
	Real demand	$\overline{Q}_f(t)$

[46] Exceptions include the models in Ashraf et al. (2016); Poledna et al. (2023a), where firms produce different types of goods but do not distinguish between consumption goods and capital, and the models in Seppecher (2012); Lengnick (2013), where there is only one type of (consumption good) firm that uses labour only.

[47] In Lengnick (2013), wages follow a stochastic process, while in several other models (Poledna et al., 2023a; Dawid et al., 2014; Dosi et al., 2010; Mandel et al., 2010), wages are also adjusted according to the evolution of labour productivity.

[48] The model in Ashraf et al. (2016) explicitly describes market entry and associated decisions a firm needs to make.

	Predicted real demand	$\tilde{Q}_f(t)$
	Demand adjustment speed on firm growth	ϕ_F^Q
Inventory and stocks	Real inventory	$S_f(t)$
	Sectoral depreciation rates of inventory	δ_s
	Real stock of intermediate inputs	$M_{fs}(t)$
	Intermediate inputs utilisation rate	ω^{M}
	Target real intermediate inputs before financial frictions	$\hat{M}_{fs}^{\mathrm{FF}}(t)$
	Influence of the fin. situation on target interm. inputs	ϕ^{FM}
	Target real intermediate inputs	$\hat{M}_{fs}(t)$
	Realised real intermediate inputs	$\tilde{M}_{fs}(t)$
	Influence of existing stock on target interm. inputs	ϕ^M
	Real stock of capital inputs	$K_{fs}(t)$
	Capital inputs utilisation rate	ω^{K}
	Target real capital inputs before financial frictions	$\hat{K}_{fs}^{\mathrm{FF}}(t)$
	Influence of the fin. situation on target capital inputs	ϕ^{FK}
	Target real capital inputs	$\hat{K}_{fs}(t)$
	Realised real capital inputs	$\tilde{K}_{fs}(t)$
	Influence of existing stock in setting target capital inputs	ϕ^K
	Delay in capital acquisition	T_s^{KD}
Labour	Labour inputs	$H_f(t)$
	Target labour inputs	$\hat{H}_f(t)$
	Labour productivity factor	$h_f(t)$
	Maximum increase in work effort	$h^{\max}$
	Total wages paid	$w_f(t)$
	Markup factor on new employee wages	μ_f^{WN}
	Adjustment of the average wage for new employees	ϕ^{WN}
	Time window for wage markups	T^{WN}
Price	Price	$P_f(t)$
	Unit costs	$U_f(t)$
	Demand-pull inflation	$\overline{\pi}_f^{\mathrm{DP}}(t)$
	Cost-push inflation	$\overline{\pi}_f^{\mathrm{CP}}(t)$
	Influence of demand-pull inflation when setting prices	ϕ^{DP}
	Influence of cost-push inflation when setting prices	ϕ^{CP}
Financials	Deposits	$D_f(t)$
	Predicted change in deposits	$\overline{D}_f^{\Delta}(t)$
	Profits	$\Pi_f(t)$
	Predicted profits	$\overline{\Pi}_f(t)$
	Production costs	$C_f(t)$
	Equity	$E_f(t)$
	Debt	$L_f(t)$
	Target short-term loans	$\hat{L}_f^{\mathrm{S}}(t)$
	Target long-term loans	$\hat{L}_f^{\mathrm{L}}(t)$
	Acquired short-term loans	$L_f^{\mathrm{S}}(t)$

	Acquired long-term loans	$L_f^{\mathrm{L}}(t)$
	Total loans from bank b	$L_{fb}(t)$
Weights	Interm. inputs of s' required to produce one unit of s	$m_{s's}$
	Capital inputs of s' required to produce one unit of s	$k_{s's}$
	Depreciation of capital goods of s' for production of s	$d_{s's}$
	Labour productivity individuals employed in sector s	h_s

Table 6.8: Variables and parameters in the model related to firms, in particular production, demand, inventory and stocks, labour, prices, and balance sheets.

6.4.3.1 Parameters

This section summarizes firm parameters.

Sector-specific Weights We are assuming that initial sectoral prices are set to 1. The matrix $m_{ss'}$ denotes the real amount of intermediate inputs of sector s required to produce one real unit of sector s'. It is computed by dividing each column of the intermediate-use matrix of the input-output table[49] by the corresponding sectoral gross output.

To obtain a matrix for capital inputs depreciation, we assume that at initialisation, only depreciated capital inputs are replaced. The matrix $d_{ss'}$, which denotes the real amount of capital inputs of sector s that depreciate for one real unit of produced output of sector s', is calculated by dividing each column of a matrix for sectoral capital compensation by corresponding gross sectoral production. The matrix for sectoral capital compensation is obtained from the capital compensation row of the input-output table so that the column-sums of the matrix are equal to the total capital compensation row and so that the row-sums are proportional to the gross fixed capital formation column attributed to firms.

We make a similar argument to find a matrix $k_{ss'}$, which denotes the necessary real capital inputs of sector s to produce one real unit of sector s'. We create a matrix of sectoral net fixed assets from aggregate sectoral net fixed assets[50] so that the column-sums of the matrix are equal to the total amount of net fixed assets and so that the row sums are proportional to the gross fixed capital formation column attributed to firms. Then, $k_{ss'}$ is obtained by dividing the columns of the matrix of net fixed assets by the corresponding sectoral gross output.

The average sectoral labour productivity (h_s) is calculated by dividing total sectoral gross output by the number of people employed in that sector.

[49] OECD: Inter-country input-output tables (Code: *ICIO*).

[50] OECD: Annual fixed assets by economic activity and by asset (Code: *DSD_NAMAIN10 @ DF_TABLE9A*).

Inventories (Finished Goods) In our modelling framework, we assume that firms decide on production also considering a target level of finished goods inventories (see 6.61). Since we are running the model over short time horizons, we assume inventory depreciation $\delta_s = 0$ for all sectors s (see Equation (6.86)).

Stocks (Materials and Supplies) After a sensitivity analysis concluded that the initial intermediate inputs and capital goods utilisation rates do not affect simulation output in a range of $\omega^{\mathrm{M}}, \omega^{\mathrm{K}} \in [0.5, 0.9]$, we follow Poledna et al. (2023a) and set $\omega^{\mathrm{M}} = \omega^{\mathrm{K}} = 0.85$. For simplicity, we set the sector-specific delay between the purchase of capital and being able to use that capital in production to be $T_s^{\mathrm{KD}} = 1$ for all sectors s. We also assume that firms attempt to keep their fraction of intermediate inputs and capital goods to production constant, assuming $\phi^M = \phi^K = 1$ (see Equation 6.79 and Equation 6.80).

Wage Adjustments The discussion of heterogeneity in individual labour productivity is out of the scope of this thesis. Therefore, we assume no wage mark-ups when firms fail to meet the labour targets, $\phi^{\mathrm{WN}} = 0$ (see Equation 6.71).[51]

Maximum Increase in Work Effort Following Poledna et al. (2023a), the maximum increase in work effort is assumed to be $h^{\max} = 1.5$, see Equation 6.67.

Target Setting The study of the effect of firm liquidity shortages on target production is left for future research $\phi^{\mathrm{FM}} = \phi^{\mathrm{FK}} = 0$, see Equation 6.84 and Equation 6.85).

6.4.3.2 Initial Conditions

This section discusses setting the initial conditions for firms. Figure 6.2 shows an example for initial distributions of firm sectors, number of employees, deposits, and debt for Australia, Austria, and Canada in 2013 at a scale of 1:1000.

Drawing Firms from Compustat Data Firms are sampled with replacement from Compustat data[52] so that the number of firms by sector (Code: *gsector*) matches OECD aggregates[53]. This includes the total number of employees (Code:

[51] If the model were to be used to study business cycles or unemployment dynamics, this assumption should be revisited.

[52] Compustat Global - Fundamentals Annual/Quarterly (Code: *comp_global_daily*).

[53] OECD: Structural Business Statistics by Size Class and Economic Activity (Code: *DSD_SDBSBSC_ISIC4 @ DF_SDBS_ISIC4*).

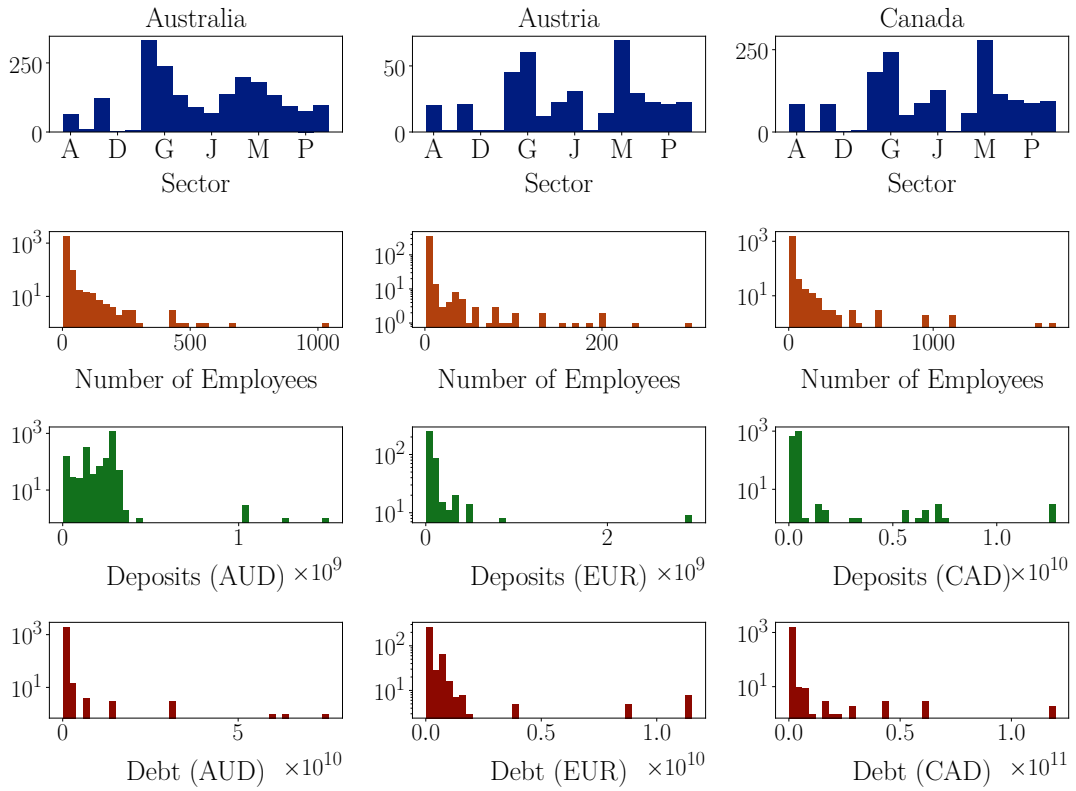


Figure 6.2: Initial distributions of firm sectors, number of employees, deposits, and debt for Australia, Austria, and Canada in 2013 at a scale of 1:1000, as obtained from OECD and Compustat data.

emp) which is rescaled to match OECD aggregates[54], total debt (Code: *dlttq*) which is rescaled to match OECD aggregates[55], total deposits (Code: *dptbq*) which is rescaled to match OECD aggregates[56]. Table 6.9 summarized the Compustat codes used to initialize firms.

The maximum number of firms recorded in Compustat from 2013-Q1 to 2017-Q4 is shown in Figure 6.3. For countries that are not included in the Compustat data for firms, we follow the approach in Poledna et al. (2023a) and create homogeneous firms.

[54]OECD: Quarterly Employment by Economic activity (Code: *DSD_NAMAIN1 @ DF_QNA_BY_ACTIVITY_EMPDC*).

[55]OECD: Quarterly Financial Balance Sheets (Code: *DSD_NASEC20 @ DF_T710R_Q*), loans (Code: *F4*) of non-financial corporations (Code: *S11*).

[56]OECD: Quarterly Financial Balance Sheets (Code: *DSD_NASEC20 @ DF_T710R_Q*), non-consolidated deposits (Code: *F2*) held by non-financial corporations (Code: *S11*).

Description	Compustat Code
Number of firms by sector	*gsector*
Number of employees	*emp*
Total debt	*dlttq*
Total deposits	*dptbq*

Table 6.9: Compustat codes for firms, in particular on the number of firms by sector, the number of employees of each firm, and each firm's debt and deposits.

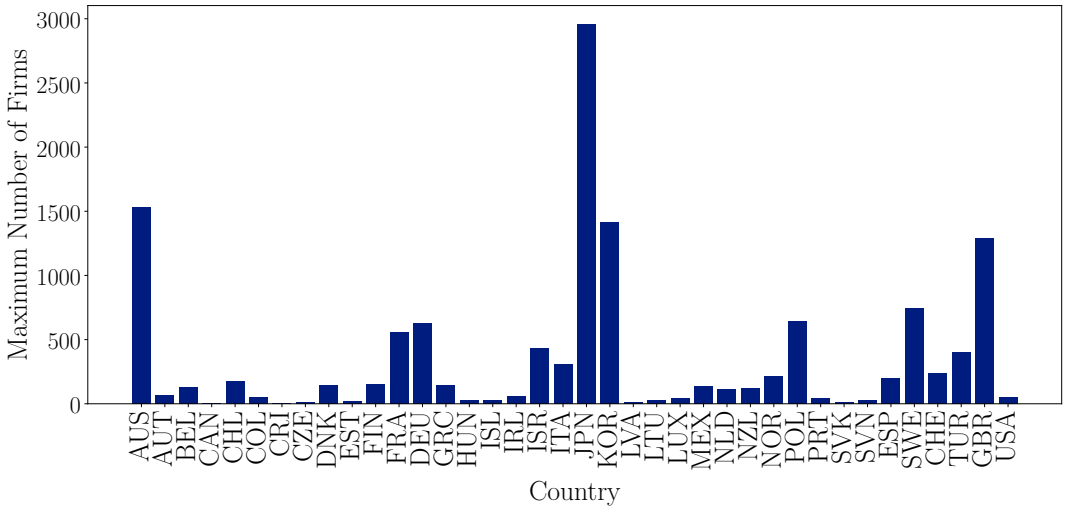


Figure 6.3: The maximum number of firms recorded in Compustat from 2013-Q1 to 2017-Q4 by country.

Initial Total Wages The initial total wages paid by a firm f are proportional to its number of employees $|\mathcal{I}_f(0)|$,

$$w_f(0) = \frac{|\mathcal{I}_f(0)|}{|\mathcal{I}_s(0)|} w_s(0), \tag{6.51}$$

where $w_s(0)$ is the initial total labour compensation of sector s obtained from socio-economic accounts and $\mathcal{I}_s(0)$ is the set of individuals employed in sector s.

Initial Production For each sector s, the initial production of the firms operating in the sector is set proportional to each firm's initial number of employees. For each sector s and each firm $f \in \mathcal{F}_s$,

$$Y_f(0) = \frac{|\mathcal{I}_f(0)|}{|\mathcal{I}_s(0)|} Y_s(0) \tag{6.52}$$

where $Y_s(0)$ is the total sectoral output and $\mathcal{I}_s(0)$ is the set of individuals employed in sector s.

Initial Prices Initial prices are set at 1. For each firm f,

$$P_f(0) = 1. \tag{6.53}$$

Initial Demand The initial demand each firm faces is set to be equal to initial production. For each firm f,

$$Q_f(0) = Y_f(0). \tag{6.54}$$

Initial Inventory The initial inventory of firm f is set according to

$$S_f(0) = \phi^{\mathrm{StY}} Y_f(0) \tag{6.55}$$

where $Y_f(0)$ is initial production and ϕ^{StY} is the target inventory to production fraction.

Initial Stock of Intermediate Inputs The initial stock of intermediate inputs of firm f operating in sector s is given by

$$M_{fs'}(0) = \frac{1}{\omega^{\mathrm{M}}} \frac{1}{m_{s's}} Y_f(0) \tag{6.56}$$

where $m_{s's}$ denotes the real amount of intermediate inputs of sector s' required to produce one real unit of sector s and ω^{M} is the initial utilisation rate of intermediate inputs.

Initial Stock of Capital Inputs The initial stock of capital inputs of firm f operating in sector s is given by

$$K_{fs'}(0) = \frac{1}{\omega^{\mathrm{K}}} \frac{1}{k_{s's}} Y_f(0) \tag{6.57}$$

where $k_{s's}$ denotes the real amount of capital inputs of sector s' required to produce one real unit of sector s and ω^{K} is the initial utilisation rate of capital inputs.

Initial Matching with Employees Firms are initially matched with employees by solving a linear sum assignment problem so that the sum of the differences between total firm wages and the sum of received wages before taxes by employees is minimal.

Initial Matching with Banks Firms are initially matched with banks by solving a linear sum assignment problem so that the sum of the differences between total firm deposits plus debt to total bank deposits and debt is minimal.

Initial Profits Total initial production costs of firm f operating in sector s with deposits at bank b are computed as

$$C_f(0) = \underbrace{w_f(t)}_{\text{Labour costs}} + \underbrace{\sum_{s' \in \mathcal{S}} m_{s's} P_{s'}(0) Y_f(0)}_{\text{Intermediate inputs bought}} + \underbrace{\sum_{s' \in \mathcal{S}} d_{s's} P_{s'}(0) Y_f(0)}_{\text{Capital inputs bought}} + \underbrace{\tau_s^{\text{PROD}} P_f(0) Y_f(0)}_{\text{Taxes on production}} + \underbrace{r_b^{\text{F-O}}(0) \left[D_f(0)\right]^- - r(0) \left[D_f(0)\right]^+}_{\text{Interest paid on deposits}} + \underbrace{\sum_{l \in \mathcal{L}_f(0)} r_l V_l(0)}_{\text{Interest on loans}}, \tag{6.58}$$

where $w_f(0)$ are total initial wages paid by firm f, $P_s(0) = 1$ are average initial sectoral prices, $r_b^{\text{F-O}}(0)$ is the initial overdraft rate on firm deposits, $r(0)$ is the initial central bank policy rate, and $\mathcal{L}_f(0)$ is the initial set of loans that firm took out. The matrix $d_{s's}$ denotes depreciation rates of capital goods of sector s' for firms producing goods of sector s. We assume that at initialisation, the firm buys as many real units of intermediate inputs as used for production and that the depreciated capital is exactly replaced.

Initial firm profits are then computed as

$$\Pi_f(0) = \underbrace{P_f(0) Y_f(0)}_{\text{Output}} - \underbrace{C_f(0)}_{\text{Production costs}}. \tag{6.59}$$

6.4.3.3 Rules

Firms estimate demand and set production targets. They also decide on the wages they pay employees, target credit when they face liquidity shortages, and may go bankrupt.

Predicted Idiosyncratic Growth Firms also make idiosyncratic estimates for future growth. Predicted growth of firm f is set as

$$\overline{\gamma}_f(t) = \frac{Q_f(t-1)}{Y_f(t-1) + S_f(t-2)} \tag{6.60}$$

if there is either (real) excess demand ($Y_f(t-1) + S_f(t-2) \leq Q_f(t-1)$) and the firms' price is above the market average ($P_f(t-1) \geq P_s(t-1)$) or if there is (real) excess supply ($Y_f(t-1) + S_f(t-2) \geq Q_f(t-1)$) and the firms' price is below the market average ($P_f(t-1) \geq P_s(t-1)$); otherwise $\overline{\gamma}_f = 0$. In the equation, $Q_f(t-1)$ is previous real demand, $Y_f(t-1)$ is previous real production, and $S_f(t-2)$ is real inventory previously offered on the goods market alongside newly produced goods.

Predicted Demand The real demand for firm f's goods is set based on predicted sectoral growth and predicted idiosyncratic firm growth,

$$\overline{Q}_f(t) = \underbrace{(1+\overline{\gamma}_s(t))}_{\text{Predicted sectoral growth}} \times \underbrace{\left(1+\phi_F^Q\overline{\gamma}_f(t)\right)}_{\text{Predicted idiosyncratic growth}} \times \underbrace{Q_f(t-1)}_{\text{Previous demand}} \tag{6.61}$$

where $\phi_F^Q \in [0,1]$ is a parameter that controls the extent to which predicted firm-specific growth determines predicted demand.

Predicted Profits Predicted profits of firm f operating in sector s are set according to

$$\overline{\Pi}_f(t) = \underbrace{\left(1+\overline{\pi}^{\text{PPI}}(t)\right)}_{\text{Predicted PPI inflation}} \times \underbrace{\left(1+\overline{\gamma}_f(t)\right)}_{\text{Predicted idiosyncratic growth}} \times \underbrace{\Pi_f(t-1)}_{\text{Previous profits}} \tag{6.62}$$

where $\Pi_f(t-1)$ are realised previous profits.

Target Production A firm's real target production is set based on predicted real demand $\overline{Q}_f(t)$, current real inventory, and the current financial situation of the firm. It may also be limited by the firm's current workforce, stock of intermediate inputs, and stock of capital inputs. Specifically,

$$\hat{Y}_f(t) = \min\Bigg(\underbrace{\overline{Q}_f(t)+\phi^{\text{StY}}Y_f(t-1)-S_f(t-1)}_{\text{Predicted demand given current inventory}}, \underbrace{\overline{Q}_f(t)+\chi^H\left(H_f(t)-\overline{Q}_f(t)\right)}_{\text{Labour inputs}},$$
$$\underbrace{\overline{Q}_f(t)+\chi^M\left(M_f(t)-\overline{Q}_f(t)\right)}_{\text{Intermediate inputs}}, \underbrace{\overline{Q}_f(t)+\chi^K\left(K_f(t)-\overline{Q}_f(t)\right)}_{\text{Capital inputs}}\Bigg) \tag{6.63}$$

where ϕ^{StY} is the target inventory to production fraction and $S_f(t-1)$ are real inventories kept from the previous timestep. Target production (as in production, see Equation (6.73)) is additionally constrained by the firm's labour inputs $H_f(t)$, intermediate inputs $M_f(t)$ and capital inputs $K_f(t)$ (assets used in production). Parameters $\chi^{\text{H}}, \chi^{\text{M}}, \chi^{\text{K}} \in [0,1]$ determine the influence of these limiting factors on target production; for instance, if $\chi^{\text{H}} = 0$, firms would ignore their level of labour inputs in determining target production and if $\chi^{\text{H}} = 1$, they would fully consider labour input constraints. Target production determines the firm's demand for labour, intermediate inputs, and capital.

The maximum production allowed given the stock of intermediate inputs $M_f(t)$ and the maximum permitted production given the stock of capital inputs

$K_f(t)$ is set as

$$M_f(t) = \min_{s' \in \mathcal{S}} \left(\frac{M_{fs'}(t-1)}{m_{s's}} \right) \tag{6.64}$$

$$K_f(t) = \min_{s' \in \mathcal{S}} \left(\frac{K_{fs'}(t-1)}{k_{s's}} \right). \tag{6.65}$$

These functional forms are substantially different from those of the IIASA model, which aggregates all intermediate inputs into a single intermediate composite and all capital inputs into a single capital composite, essentially assuming a linear production function in each intermediate and capital input. Here, we take a stricter approach regarding substitution possibilities in the economy and assume a Leontief production function for each input instead. However, our framework also allows for incorporating less restrictive production functions. One example is the *partially binding* Leontief production function (Pichler et al., 2022), which is a Leontief production function that, for each sector, only considers certain inputs to be critical for production, as evaluated in a survey of sector analysts.

Labour Inputs The labour inputs of firm f are given by

$$H_f(t) = h_f(t) \underbrace{\sum_{i \in \mathcal{I}_f(t)} H_i(t)}_{\text{Labour supply from employees}}, \tag{6.66}$$

where $\mathcal{I}_f(t)$ is the set of individuals employed by firm f, $H_i(t)$ are labour inputs from individual i, and $h_f(t)$ is a firm-specific factor denoting work effort, set as

$$\phi^{\text{WE}} = \min \left(\underbrace{h^{\max}}_{\text{Maximum factor}}, \underbrace{\frac{\min\left(M_f(t), K_f(t)\right)}{h_f(0) \sum_{i \in \mathcal{I}_f(t)} H_i(t)}}_{\text{Additional production factor}} \right) \tag{6.67}$$

$$h_f(t) = \phi^{\text{WE}} h_f(0). \tag{6.68}$$

This follows the approach in the IIASA model, where $\phi^{\text{WE}} > 1.0$ corresponds to overtime work allowed by the availability of intermediate inputs $M_f(t)$ and capital inputs $K_f(t)$. The parameter $h^{\max}$ is the maximum increase in work effort. Similarly, $\phi^{\text{WE}} < 1.0$ corresponds to part-time employment. The term $\min\left(M_f(t), K_f(t)\right)$ denotes the maximum possible production given the level of intermediate inputs and capital inputs; the term $h_f(0) \sum_{i \in \mathcal{I}_f(t)} H_i(t)$ denotes the maximum possible production given the firm's labour inputs.

The target labour inputs of firms are set based on target production,

$$\hat{H}_f(t) = \hat{Y}_f(t). \tag{6.69}$$

Wages The wages paid by firm f to its currently employed individuals $i \in \mathcal{I}_f(t)$ are set according to

$$w_i(t) = \underbrace{\left(1 + \overline{\pi}^{\text{PPI}}(t)\right)}_{\text{Predicted PPI inflation}} \times \underbrace{\left(1 + \mu_f^{\text{WN}}\right)}_{\text{Labour market tightness}} \times \underbrace{\phi^{\text{WE}}}_{\text{Work effort}} \times \underbrace{w_i(t-1)}_{\text{Previous wage}} \tag{6.70}$$

where ϕ^{WE} denotes the level of work effort defined in Equation (6.67) above, $w_i(t-1)$ is the previous wage paid to employee i, and μ_f^{WN} is a firm-specific markup on the average salary paid to improve the chances of keeping the current or hiring new employees.[57] It is set to be

$$\mu_f^{\text{WN}} = \phi^{\text{WN}} \frac{1}{T^{\text{WN}}} \underbrace{\sum_{t'=1}^{T^{\text{WN}}} \left[\frac{\hat{H}_f(t-t') - H_f(t-t')}{\hat{H}_f(t-t')} \right]^+}_{\text{Relative failure in hiring labour}} \tag{6.71}$$

where ϕ^{WN} is a parameter, and T^{WN} is the timeframe firms consider when raising prices after failing to meet their labour targets $\hat{H}_f(t-t') - H_f(t-t')$. The offered wage to a prospective employee i that contributes labour inputs $H_i(t)$ is given by

$$w_i(t) = \underbrace{\left(1 + \mu_f^{\text{WN}}\right)}_{\text{Labour market tightness}} \times \underbrace{\phi^{\text{WE}}}_{\text{Work effort}} \times \underbrace{\frac{\sum_{i' \in \mathcal{I}_f(t-1)} w_{i'}(t-1)}{\sum_{i' \in \mathcal{I}_f(t-1)} H_{i'}(t)}}_{\text{Average wage by labour inputs}} \underbrace{H_i(t)}_{\text{Prospective labour inputs}} . \tag{6.72}$$

Production Firms produce goods of the sector they operate in according to a Leontief production function using labour inputs $H_f(t)$, intermediate inputs $M_f(t)$, and capital inputs $K_f(t)$.[58] Firm f's production is given by

$$Y_f(t) = \min \left(\underbrace{\hat{Y}_f(t)}_{\text{Target production}}, \underbrace{H_f(t)}_{\text{Labour supply}}, \underbrace{M_f(t-1)}_{\text{Intermediate inputs}}, \underbrace{K_f(t-1)}_{\text{Capital inputs}} \right). \tag{6.73}$$

Production by firm f may not be equal to target production $\hat{Y}_f(t)$, and it has to scale down activity if the available labour supply, intermediate inputs, or capital inputs limit it. This approach is consistent with the data and in line with similar large models (Poledna et al., 2023a; Hommes et al., 2022).

[57] The functional form (6.70) does not include wage stickiness. This should be improved in further iterations, as this is an important driver for inflation.

[58] In most models, there are two types of firms: consumption goods firms that require capital and labour, and capital goods firms that only require labour. There is a vast variety of used production technologies in different models, for instance, a Leontief production function in Assenza et al. (2015); Hommes et al. (2022), Cobb-Douglas in Dawid and Gatti (2018), or a CRS technology in Gaffeo et al. (2008).

Prices Firms set prices based on their expectations for inflation. The price set by firm f operating in sector s is given by

$$P_f(t) = \underbrace{\left(1+\overline{\pi}^{\text{PPI}}(t)\right)}_{\text{Predicted PPI inflation}} \times \underbrace{\left(1+\phi^{\text{DP}}\overline{\pi}_f^{\text{DP}}(t)\right)}_{\text{Demand-pull inflation}} \times \underbrace{\left(1+\phi^{\text{CP}}\overline{\pi}_f^{\text{CP}}(t)\right)}_{\text{Cost-push inflation}} \times \underbrace{P_f(t-1)}_{\text{Previous price}} \tag{6.74}$$

where $\phi^{\text{DP}} \in [0,1]$ and $\phi^{\text{CP}} \in [0,1]$ are parameters, demand-pull inflation is defined as

$$\overline{\pi}_f^{\text{DP}}(t) = \frac{Q_f(t-1)}{Y_f(t-1)+S_f(t-2)} - 1 \tag{6.75}$$

if the total good real supply offered by the firm $Y_f(t-1)+S_f(t-2)$ exceeds previous real demand $Q_f(t-1)$ and the firm had a price $P_f(t-1)$ that was higher than the average sector price $P_s(t-1)$, or if total good supply is lower than previous demand, and firm prices are lower than the average sector price, otherwise

$$\overline{\pi_f}^{\text{DP}}(t) = 0. \tag{6.76}$$

On the other hand, cost-push inflation is defined as

$$\overline{\pi}_f^{\text{CP}}(t) = \frac{U_f(t-1)}{P_f(t-1)} - 1 \tag{6.77}$$

where $U_f(t-1)$ are previous unit costs, computed as

$$U_f(t) = \underbrace{\frac{w_f(t)}{Y_f(t)}}_{\text{Labour costs}} + \underbrace{\sum_{s'\in\mathcal{S}} m_{s's}P_{s'}(t-1)}_{\text{Intermediate inputs used}} + \underbrace{\sum_{s'\in\mathcal{S}} d_{s's}P_{s'}(t-1)}_{\text{Depreciation}} + \underbrace{\tau_s^{\text{PROD}}P_f(t-1)}_{\text{Taxes on production}} \tag{6.78}$$

where the first term $\frac{w_f(t)}{Y_f(t)}$ denotes total wages paid by real units of output. For the predicted PPI inflation see Section 6.2.4.2.

Demand for Goods before Financial Frictions Each firm needs intermediate inputs for production. Each firm holds a real stock of intermediate inputs $M_{fs}(t)$ that is depleted in each production step. Target intermediate inputs before financial frictions of sector s' of firm f operating in sector s are based on the chosen target production $\hat{Y}_f(t)$ and current stocks and given by

$$\hat{M}_{fs'}^{\text{FF}}(t) = \left[\underbrace{m_{s's}\hat{Y}_f(t)}_{\text{Targetted amount}} - \phi^M \underbrace{\left(M_{fs'}(t-1) - M_{fs'}(0)\frac{Y_f(t)}{Y_f(0)}\right)}_{\text{Existing stock}}\right]^+ \tag{6.79}$$

where the parameter ϕ^M governs how much the current stock of intermediate inputs is considered.

Similarly, each firm needs capital inputs for production. Each firm holds a real stock of capital inputs $K_{fs}(t)$ that depreciates when used for production. Target capital inputs before financial frictions are set as

$$\hat{K}^{\mathrm{FF}}_{fs'}(t) = \left[\underbrace{d_{s's}\hat{Y}_f(t)}_{\text{Targetted amount}} - \phi^K \underbrace{\left(K_{fs'}(t-1) - K_{fs'}(0)\frac{Y_f(t)}{Y_f(0)} \right)}_{\text{Existing stock}} \right]^+ \tag{6.80}$$

where the parameter ϕ^K governs how much the current stock of capital inputs is considered.

Demand for Loans The predicted change in deposits without new loans and without taking the purchase of new inputs into account is given by

$$\begin{aligned} \overline{D}^{\Delta}_f(t) = & \underbrace{P_f(t)Y_f(t)}_{\text{Production}} - \underbrace{w_f(t)}_{\text{Total wages paid}} - \underbrace{\tau^{\mathrm{CORP}} \left[\overline{\Pi}_f(t)\right]^+}_{\text{Predicted corporate taxes}} \\ & - \underbrace{r^{\text{F-O}}_b(t) \left[D_f(t-1)\right]^- - r(t) \left[D_f(t-1)\right]^+}_{\text{Interest paid on deposits}} - \underbrace{\sum_{l \in \mathcal{L}_f(t)} r_l V_l(t)}_{\text{Interest on loans}} \\ & - \underbrace{\tau^{\mathrm{PROD}}_s P_f(t)Y_f(t)}_{\text{Taxes on production}} - \underbrace{\sum_{l \in \mathcal{L}_f(t)} \frac{V_l(t)}{m_l}}_{\text{Debt installment}} \end{aligned} \tag{6.81}$$

where $w_f(t)$ are total wages paid by firm f, $\overline{\Pi}_f(t)$ are predicted firm profits, $r^{\text{F-O}}_b(t)$ is the overdraft rate on firm deposits, $r(t)$ is the central bank policy rate, and $\mathcal{L}_f(t)$ is the set of loans that firm took out.

The firms' financing needs determine the demand for credit. If the operating costs exceed the internal funds, the firm has that difference as a financing gap and will apply for a bank loan. Firms apply for short-term and long-term loans. Each firm applies for short-term loans to cover any financing gap due to wages or intermediate inputs purchasing costs by applying for loans to the value of

$$\hat{L}^{\mathrm{S}}_f(t) = \left[\underbrace{D_f(t-1) + \overline{D}^{\Delta}_f(t)}_{\text{Predicted deposits}} - \underbrace{\sum_{s' \in \mathcal{S}} P_{s'}(t-1)\hat{M}^{\mathrm{FF}}_{fs'}(t)}_{\text{Costs of intermediate inputs}} \right]^- \tag{6.82}$$

where $P_{s'}(t-1)$ is the previous average price of goods from sector s. Each firm may also apply for long-term loans to finance the purchase of capital inputs,

$$\hat{L}^{\mathrm{L}}_f(t) = \left[\underbrace{D_f(t-1) + \Delta\overline{D}_f(t)}_{\text{Predicted deposits}} - \underbrace{\sum_{s' \in \mathcal{S}} P_{s'}(t-1)\hat{K}^{\mathrm{FF}}_{fs'}(t)}_{\text{Costs of capital inputs}} \right]^- . \tag{6.83}$$

Constrained Demand for Goods Firms purchase fewer intermediate inputs and capital goods if they fail to acquire the loans they applied for. Specifically, target intermediate inputs are set as

$$\hat{M}_{fs'}(t) = \underbrace{\hat{M}^{\mathrm{FF}}_{fs'}(t)}_{\text{Demand before frictions}} - \phi^{\mathrm{FM}} \underbrace{\left(\frac{\hat{L}^{\mathrm{S}}_f(t) - L^{\mathrm{S}}_f(t)}{(1+\overline{\pi}^{\mathrm{PPI}}(t))\, P_s(t-1)} \right)}_{\text{Credit gap}}, \tag{6.84}$$

where $\hat{L}^{\mathrm{S}}_f(t) - L^{\mathrm{S}}_f(t)$ is the difference between requested and granted short-term loans, and ϕ^{FM} is a parameter. Target capital inputs are set as

$$\hat{K}_{fs'}(t) = \underbrace{\hat{K}^{\mathrm{FF}}_{fs'}(t)}_{\text{Demand before frictions}} - \phi^{\mathrm{FK}} \underbrace{\left(\frac{\hat{L}^{\mathrm{L}}_f(t) - L^{\mathrm{L}}_f(t)}{(1+\overline{\pi}^{\mathrm{PPI}}(t))\, P_s(t-1)} \right)}_{\text{Credit gap}}, \tag{6.85}$$

where $\hat{L}^{\mathrm{L}}_f(t) - L^{\mathrm{L}}_f(t)$ is the difference between requested and granted long-term loans, and ϕ^{FK} is a parameter.

Inventory and Stocks The inventory of finalised goods of firm f operating in sector s is updated according to

$$S_f(t) = \max\left(0, \underbrace{(1-\delta_s)\, S_f(t-1)}_{\text{Previous inventory}} + \underbrace{Y_f(t)}_{\text{Production}} - \underbrace{\tilde{Q}_f(t)}_{\text{Units sold}}\right) \tag{6.86}$$

where δ_s is sector-specific depreciation of finalised goods $S_f(t)$ in the inventory of firm f. There is no direct distinction between goods bought as intermediate inputs and goods bought as capital inputs in the goods market. Firms prioritize bought goods to satisfy their demand for intermediate inputs first; the remainder goes into their capital stock. The updated stock of intermediate inputs is

$$M_{fs'}(t) = \left[\underbrace{M_{fs'}(t-1)}_{\text{Previous stock}} - \underbrace{m_{s's} Y_f(t)}_{\text{Used stock}} + \underbrace{\tilde{M}_{fs'}(t)}_{\text{Newly bought stock}} \right]^+ \tag{6.87}$$

and the new stock of capital inputs is

$$K_{fs'}(t) = \left[\underbrace{K_{fs'}(t-1)}_{\text{Previous stock}} - \underbrace{d_{s's} Y_f(t)}_{\text{Depreciated stock}} + \underbrace{\tilde{K}_{fs'}(t - T^{\mathrm{D}}_{s'})}_{\text{Newly bought stock}} \right]^+ \tag{6.88}$$

where $T^{\mathrm{D}}_{s'}$ is a sector-specific delay between purchasing capital and using that capital in production.

Demand Firms face demand for goods from other firms, households, government entities, and the rest of the world. The real demand a firm faces is how much that firm managed to sell to buyers plus what they could have sold additionally. Specifically, firm f's real demand is given by

$$Q_f(t) = \underbrace{\tilde{Q}_f(t)}_{\text{Amount sold}} + \underbrace{\tilde{Y}_f^{\text{E}}(t)}_{\text{Excess demand}} \tag{6.89}$$

where $\tilde{Q}_f(t)$ is how much firm f sold on the goods market and $\tilde{Y}_f^{\text{E}}(t)$ is how much firm f could have sold additionally if it had more supply. To calculate excess demand, firms get allocated a share of the total real amount potential buyers would have purchased at the offered price if the random search-and-matching process of the goods market had continued after firms ran out of supply. The share the firm receives is based on the firms' characteristics; see Section 6.5.1.

Profits Total production costs of firm f operating in sector s with deposits at bank b are computed as

$$\begin{aligned} C_f(t) = & \underbrace{w_f(t)}_{\text{Total wages paid}} + \underbrace{\sum_{s' \in \mathcal{S}} P_{s'}(t) \left(M_{fs'}(t) - M_{fs'}(t-1) + m_{s's} Y_f(t)\right)}_{\text{Intermediate inputs bought}} \\ & + \underbrace{\sum_{s' \in \mathcal{S}} P_{s'}(t) \left(K_{fs'}(t) - K_{fs'}(t-1) + d_{s's} Y_f(t)\right)}_{\text{Capital inputs bought}} \\ & + \underbrace{r_b^{\text{F-O}}(t) \left[D_f(t-1)\right]^- - r(t) \left[D_f(t-1)\right]^+}_{\text{Interest paid on deposits}} + \underbrace{\sum_{l \in \mathcal{L}_f(t)} r_l V_l(t)}_{\text{Interest on loans}} \\ & + \underbrace{\tau_s^{\text{PROD}} P_f(t) Y_f(t)}_{\text{Taxes on production}}, \end{aligned} \tag{6.90}$$

where $w_f(t)$ are total wages paid by firm f, $P_s(t)$ is the avetage sectoral price of good s, $r_b^{\text{F-O}}(t)$ is the overdraft rate on firm deposits, $r(t)$ is the central bank policy rate, and $\mathcal{L}_f(t)$ is the set of loans that firm took out. Firm profits are then computed as

$$\Pi_f(t) = \underbrace{P_f(t)\tilde{Q}_f(t)}_{\text{Sales}} + \underbrace{P_f(t)\Delta S_f(t)}_{\text{Inventory change}} - \underbrace{C_f(t)}_{\text{Production costs}}. \tag{6.91}$$

Deposits Deposits of firm f that has its deposits with bank b and is operating in sector s are updated as

$$\begin{aligned} D_f(t) = & \underbrace{D_f(t-1)}_{\text{Previous deposits}} + \underbrace{P_f(t)\tilde{Q}_f(t)}_{\text{Units sold}} - \underbrace{C_f(t)}_{\text{Production costs}} - \underbrace{\tau^{\text{CORP}} \left[\Pi_f(t)\right]^+}_{\text{Corporate taxes}} \\ & - \underbrace{\sum_{l \in \mathcal{L}_f(t)} \frac{V_l(t)}{m_l}}_{\text{Debt installment}} + \underbrace{\Delta L_f(t)}_{\text{New credit}}. \end{aligned} \tag{6.92}$$

Debt The debt of firm f is updated as

$$L_f(t) = \underbrace{L_f(t-1)}_{\text{Previous debt}} - \underbrace{\sum_{l \in \mathcal{L}_f(t)} \frac{V_l(t)}{m_l}}_{\text{Debt installment}} + \underbrace{\Delta L_f(t).}_{\text{New credit}} \tag{6.93}$$

Equity Equity of firm f evolves according to

$$E_f(t) = \underbrace{D_f(t)}_{\text{Deposits}} + \underbrace{P_f(t) S_f(t)}_{\text{Inventory}} + \underbrace{\sum_{s' \in \mathcal{S}} P_{s'}(t) \left(M_{fs'}(t) + K_{fs'}(t) \right)}_{\text{Stock}} - \underbrace{L_f(t).}_{\text{Debt}} \tag{6.94}$$

Bankruptcy If a firm f is cash-flow insolvent ($D_f(t) < 0$) and balance-sheet insolvent ($E_f(t) < 0$), it goes bankrupt. A bankrupt firm is replaced by a new firm that enters the same sector. That new firm keeps the same stock and inventory as the bankrupt firm and is initialised with

$$\underbrace{D_f(t+1)}_{\text{Deposits}} = \underbrace{L_f(t+1)}_{\text{Debt}} = 0. \tag{6.95}$$

As described in Equation (6.42), the losses are borne by the bank b that gave credit to firm f.

6.4.4 Government

The government consumes according to predefined fractions, collects taxes, and pays social benefits to individuals and households. Variables and parameters related to the government are listed in Table 6.10.

6.4.4.1 Initial Conditions

This section summarizes the government's initial conditions on consumption, social benefits, and initial debt.

Number of Entities Following Poledna et al. (2023a), the number of government entities in each country (agents operating on the goods market) is set to be 25% of the number of domestically producing firms.

Initial Consumption Initial total consumption $C^{\text{CG}}(0) = \hat{C}^{\text{CG}}(0)$ and government consumption weights c_s^{CG} are matched to input-output tables[59], see Equation (6.96). Consumption is distributed evenly among all government entities.

[59] OECD: Inter-country input-output tables (Code: *ICIO*).

Category	Description	Notation
Sets	Set of government entities	$\mathcal{G}$
Revenue	Revenue	$Y^{\mathrm{CG}}(t)$
	Sectoral taxes on production	τ_s^{PROD}
	Income tax rate	τ^{INC}
	Corporate tax rate	τ^{CORP}
	Value-added tax rate	τ^{VAT}
	Capital formation tax rate	τ^{CF}
	Social insurance tax rate (employers)	τ^{SIF}
	Social insurance tax rate (employees)	τ^{SIW}
	Exports tax rate	τ^{EXP}
Expenditures	Consumption	$C^{\mathrm{CG}}(t)$
	Target consumption	$\hat{C}^{\mathrm{CG}}(t)$
	Government consumption weights	c_s^{CG}
	Per-capita real unemployment benefits	$w^{\mathrm{U}}(t)$
	Total real other social benefits	$\mathrm{sb}^{\mathrm{O}}(t)$
Deficit	Deficit	$\Pi^{\mathrm{CG}}(t)$
Debt	Debt	$L^{\mathrm{CG}}(t)$

Table 6.10: Variables and parameters in the model related to the government, in particular, tax rates, expenditures, deficit, and debt.

Initial Social Benefits Initial social benefits are matched with OECD aggregates[60] and distributed according to *Household Finance and Consumption Survey* microdata, see Equation (6.97) and Equation (6.98).

Initial Debt Initial government debt $L^{\mathrm{CG}}(0)$ is initialised using World Bank data[61], see Equation (6.101).

6.4.4.2 Parameters

Sectoral taxes on production τ_s^{PROD} are taken from input-output tables[62]. The income tax rate, corporate tax rate, export taxes, value-added tax rate, and social insurance rates are taken directly from the OECD database.[63]

[60] OECD: Social Expenditure Aggregates (Code: *DSD_SOCX_AGG @ DF_SOCX_AGG*).
[61] World Bank Data: Government debt (Code: *GC.DOD.TOTL.GD.ZS*)
[62] OECD: Inter-country input-output tables (Code: *ICIO*).
[63] OECD: Annual Government Taxes and Social Contributions Receipts (Code: *DSD_NASEC10 @ DF_TABLE10*) with income taxes (Code: *D51A*), corporate taxes (Code: *D51B*), export taxes (Code: *D214K*), value-added taxes (Code: *D211*), employees' contributions to social insurance (Code: *D613CE*), employers' contribution to social insurance (Code: *D611*).

6.4.4.3 Rules

The government consumes and pays social benefits, based on which we update revenue, deficits, and debt.

Consumption Government entities update their total real target consumption $\hat{C}^{\mathrm{CG}}(t)$ based on an AR(1) model on historical data. Nominal sectoral target consumption is then given by

$$\hat{C}_s^{\mathrm{CG}}(t) = \underbrace{c_s^{\mathrm{CG}}}_{\text{Weights}} \times \underbrace{\left(1+\overline{\pi}^{\mathrm{PPI}}(t)\right)}_{\text{Predicted PPI inflation}} \times \underbrace{P_s(t-1)}_{\text{Previous average price}} \times \underbrace{\hat{C}^{\mathrm{CG}}(t)}_{\text{Total target consumption}} \tag{6.96}$$

which is then distributed evenly among each government entity $g \in \mathcal{G}$.

Social Benefits The government pays unemployment benefits to unemployed individuals and other social benefits to households based on household composition and individual characteristics (income, wealth, and debt). Per-capita real unemployment benefits are increased when agents predict that the economy will enter a recession (see Section 6.2.5.2),

$$w^{\mathrm{U}}(t) = \underbrace{\max\left(1, \frac{1}{1+\overline{\gamma}(t)}\right)}_{\text{Inverse predicted growth}} w^{\mathrm{U}}(t-1) \tag{6.97}$$

Total real other benefits grow according to

$$\mathrm{sb}^{\mathrm{O}}(t) = \underbrace{(1+\overline{\gamma}(t))}_{\text{Predicted growth}} \mathrm{sb}^{\mathrm{O}}(t-1). \tag{6.98}$$

Revenue The government collects social contributions (τ^{SIF} employers', τ^{SIW} employees') and taxes according to fixed rates on labour income (τ^{INC}), corporate income (τ^{CORP}), value-added (τ^{VAT}), household capital formation (τ^{CF}), net taxes on production (τ_s^{PROD}), and exports (τ^{EXP}):

$$\begin{aligned} Y^{\mathrm{G}}(t) = {} & \underbrace{P^{\mathrm{CPI}}(t)\left(\tau^{\mathrm{SIW}} + \tau^{\mathrm{INC}}\left(1-\tau^{\mathrm{SIW}}\right)\right)\sum_{i\in\mathcal{I}^{\mathrm{E}}(t)} w_i(t)}_{\text{Social security contributions and labour income taxes}} + \underbrace{\tau^{\mathrm{INC}}\sum_{h\in\mathcal{H}}\sum_{p\in\mathcal{P}_h^{\mathrm{R}}(t)} r_p(t)}_{\text{Rental income taxes}} \\ & + \underbrace{\tau^{\mathrm{CORP}}\left(\sum_{f\in\mathcal{F}}[\Pi_f(t)]^+ + \sum_{b\in\mathcal{B}}[\Pi_b(t)]^+\right)}_{\text{Corporate income taxes}} + \underbrace{\tau^{\mathrm{VAT}}\sum_{h\in\mathcal{H}} C_h(t)}_{\text{Value added taxes}} \\ & + \underbrace{\tau^{\mathrm{CF}}\sum_{h\in\mathcal{H}} K_h(t)}_{\text{Taxes on capital formation}} + \underbrace{\sum_{s\in\mathcal{S}}\tau_s^{\mathrm{PROD}}\sum_{f\in\mathcal{F}_s} P_f(t)Y_f(t)}_{\text{Taxes on Production}} + \underbrace{\tau^{\mathrm{EXP}}\mathrm{EXP}(t)}_{\text{Export taxes}} \end{aligned} \tag{6.99}$$

where we assume that taxes are indexed to consumer price index inflation $P^{\mathrm{CPI}}(t)$. Furthermore, $\mathcal{I}^{\mathrm{E}}(t)$ is the set of employed individuals, $w_i(t)$ is the wage paid by the employer of individual i, $\mathcal{P}_h^{\mathrm{R}}(t)$ is the set of properties owned by household h that are rented-out to other households, $r_p(t)$ is the rent paid for living in property p, $\Pi_f(t)$ / $\Pi_b(t)$ are profits of firm f / bank b, $C_h(t)$ is consumption of household h, $K_h(t)$ is investment of household h, and $\mathrm{EXP}(t)$ are total nominal exports.

Deficit/Surplus Government deficits $\Pi^{\mathrm{CG}}(t)$ are updated according to

$$\Pi^{\mathrm{CG}}(t) = \underbrace{P^{\mathrm{CPI}}(t)\left(\mathrm{sb}^{\mathrm{O}}(t) + \left|\mathcal{I}^{\mathrm{U}}(t)\right| w^{\mathrm{U}}(t)\right)}_{\text{Social benefits}} + \underbrace{C^{\mathrm{CG}}(t)}_{\text{Spending}} + \underbrace{r(t)L^{\mathrm{CG}}(t-1)}_{\text{Interest payments}} - \underbrace{Y^{\mathrm{CG}}(t)}_{\text{Revenue}}. \tag{6.100}$$

where $\mathrm{sb}^{\mathrm{O}}(t)$ are total real other social benefits and $\left|\mathcal{I}^{\mathrm{U}}(t)\right| w^{\mathrm{U}}(t)$ are total real unemployment benefits. We are assuming that social housing revenue is equal to costs.

Debt Government debt is updated according to

$$L^{\mathrm{CG}}(t) = \underbrace{L^{\mathrm{CG}}(t-1)}_{\text{Previous debt}} + \underbrace{\Pi^{\mathrm{CG}}(t)}_{\text{Deficit}}. \tag{6.101}$$

6.4.5 Households

Households are groups of individuals. As economic units, households are the relevant entities in the goods, credit, and housing markets. They may own housing, live in rented properties, or reside in social housing, and can apply for consumption loans or mortgages. Household income is composed of the sum of income from their individuals, along with additional social benefits, income from financial assets, and rental income. They hold wealth in the form of both real and financial assets.

Table 6.11 lists variables and parameters related to households.

Category	Description	Notation
Sets	Set of households	$\mathcal{H}$
	Set of insolvent households	$\mathcal{H}'(t)$
	Set of corresponding individuals	$\mathcal{I}_h(t)$
	Set of loans	$\mathcal{L}_h(t)$
	Set of owned properties	$\mathcal{P}_h(t)$
	Set of rented-out properties	$\mathcal{P}_h^{\mathrm{R}}(t)$
Income	Income	$Y_h(t)$
	Predicted income	$\overline{Y}_h(t)$
	Rent	$r_h(t)$
	Other social benefits	$\mathrm{sb}_h^{\mathrm{O}}(t)$
	Income from financial assets coefficient	ϕ^{FA}

	Variance of the Gaussian noise impacting income	σ^{FA}
Consumption	Consumption	$C_h(t)$
	Target consumption	$\hat{C}_{hs}(t)$
	Aggregate consumption weights	$c_s^{\mathrm{CPI}}(t)$
	Saving rate	$\phi_h^{\mathrm{SR}}(t)$
	Consumption smoothing fraction	ϕ^{CO}
	Time frame for consumption smoothing	T^{CO}
Investment	Investment	$K_h(t)$
	Target investment	$\hat{K}_{hs}(t)$
	Investment weights	k_s
	Investment rate	ϕ^{IR}
Wealth	Wealth	$W_h(t)$
	Wealth in properties	$W_h^{\mathrm{P}}(t)$
	Wealth in other real assets	$W_h^{\mathrm{ORA}}(t)$
	Wealth in real assets	$W_h^{\mathrm{RA}}(t)$
	Wealth in deposits	$W_h^{\mathrm{D}}(t)$
	Wealth in other financial assets	$W_h^{\mathrm{OFA}}(t)$
	Wealth in financial assets	$W_h^{\mathrm{FA}}(t)$
	Net Wealth	$W_h^{\mathrm{N}}(t)$
	Rate of depreciation of real assets	d^{RA}
Debt	Total debt	$L_h(t)$
	Debt to bank b	$L_{hb}(t)$
	Debt in consumptions loans	$L_h^{\mathrm{C}}(t)$
	Debt in consumptions loans to bank b	$L_{hb}^{\mathrm{C}}(t)$
	Debt in mortgages	$L_h^{\mathrm{M}}(t)$
	Debt in mortgages to bank b	$L_{hb}^{\mathrm{M}}(t)$
	Demand for consumption loans	$\hat{L}_h^{\mathrm{C}}(t)$
	Demand for mortgages	$\hat{L}_h^{\mathrm{M}}(t)$
	Newly granted consumption loans	$\hat{L}_h^{\Delta\mathrm{C}}(t)$
	Newly granted mortgages	$\hat{L}_h^{\Delta\mathrm{M}}(t)$
Housing as a tenant	Probability of not changing housing when renting	p^{RS}
	Maximum price for buying a property	$P_h(t)$
	Factor for setting the maximum price	ϕ^{HP}
	Exponent for setting the maximum price	β^{HP}
	Mean of the Gaussian noise impacting the price	μ^{HP}
	Var. of the Gaussian noise impacting the price	σ^{HP}
	Psychological pressure of renting	μ^{PS}
	Influence of costs in deciding whether to buy/rent	ϕ^{B}
	Factor for setting the maximum rent	ϕ^{HR}
	Exponent for setting the maximum rent	β^{HR}
	Probability of buying over renting	$p_h^{\mathrm{B}}(t)$
Housing as an owner-occupier	Probability of not changing housing when owning	p^{OS}

	Probability of price reduction	p^{PM}
	Mean of the Gaussian noise reducing the price	μ^{PM}
	Variance of the Gaussian noise reducing the price	σ^{PM}
Housing as an investor	Probability of rent reduction	p^{RM}
	Mean of the Gaussian noise reducing the rent	μ^{RM}
	Variance of the Gaussian noise reducing the rent	σ^{RM}
	Partial indexation for setting rent	ϕ^{PIR}
	Lag for setting rent	T^{PIR}

Table 6.11: Variables and parameters in the model related to households, in particular income, consumption, investment, wealth, debt and housing.

6.4.5.1 Calculating Household Consumption Weights

This section explains our approach for obtaining household consumption weights by income quintile. Our treatment of household saving rates and household consumption weights (Section 6.4.5.2) allows for consumption heterogeneity that remains as close as possible to real data. Our method implicitly creates the three types of households typically found in standard macroeconomic models — hand-to-mouth households (low liquidity, high propensity to consume, deriving income from wages and government transfers), middle class (strong precautionary savings, income from wages), and high net-worth individuals (low propensity to consume, income from capital).

We match consumption weights by income quintile to agent behaviour using the OECD's dataset on distributional information on household income, consumption, and saving. The dataset presents experimental distributional results for a range of OECD countries and results from the work of the *Expert Group on Distributional Information on Household Income, Consumption, and Saving* (EG DNA), which defined an internationally comparable methodology to produce measures of disparities consistent with national accounts concepts and totals using existing microdata sources. These results have been provided by countries participating in the EG DNA and Eurostat (for a selection of items) for EU countries not participating in the exercise (or did not provide results for a specific reference year). Here, we focus on consumption data in the *Classification of Individual Consumption by Purpose* (COICOP).

Our approach starts from OECD EGDNA data, which are vectors of consumption of 12 COICOP products by the five income quintiles. These vectors are given in purchaser prices and must be mapped to OECD ICIO domestic final consumption, which is given in NACE and basic prices. The first step maps the consumption vectors to ISIC notation; the second step converts them from purchaser prices to basic prices.

Conversion from COICOP to ISIC Consumption data in COICOP must be converted into ISIC using contingency or bridge matrices. The European Union national statistical offices build these bridge tables to compile national accounts, but they are not always published due to various considerations. Only eight EU countries make them available, so Cazcarro et al. (2022) suggest a systematic procedure to estimate bridge matrices for the other 20 countries. This procedure combines information from official contingency matrices with the consumption vector in the *Classification of Products by Activity* (CPA) and in the COICOP classification from Eurostat. Note that CPA refers to products, and ISIC refers to sectors. However, each product is described by the sector that produces it, so the categories in the two classification systems are the same. In principle, one should transform these matrices when working with sector-to-sector input-output tables, but here, for simplicity, we consider CPA and ISIC as interchangeable. The contingency tables are matrices with dimensions of 64 (CPA products) $\times$ 47 (COICOP categories). Each matrix element represents the total quantity of a product used. We normalize the concordance tables so that each column sums to one and then multiply each row of the matrix by each of the five consumption vectors. The result is the set of consumption vectors in CPA by the five income quintiles, which is still given in purchaser prices.

At this step, we verify that the OECD EGDNA data are consistent with national accounts. In principle, the OECD started the EGDNA project precisely to make consumption surveys consistent with national accounts. Still, it was not clear to us if, for instance, imputed rents had been added. To check this, we verify if the OECD EGDNA data used as input are consistent with the data used in Cazcarro et al. (2022) to build the concordance tables. Cazcarro et al. (2022) start from the European Union Household Budget Survey (EU HBS), which is a harmonized survey that collects data on household expenditures and income across EU member states and then adjusts the data, for instance, adding imputed rents. We find that the two datasets are highly correlated, and in particular, consumption of housing services (top right in the panel), which includes imputed rents, is very similar in the two datasets. This suggests that, indeed, OECD EGDNA data are compatible with national accounts and represent a reasonable starting point for distinguishing consumption habits by income groups.

Conversion from Purchaser Prices to Basic Prices COICOP data are given in purchaser prices, while data in input-output tables are in basic prices. Basic and purchaser prices are two different ways of measuring the value of goods and services. Basic prices refer to the price received by the producer, excluding any taxes on the product, while purchaser prices include any taxes paid by the

consumer, such as value-added tax (VAT). Moreover, basic prices exclude trade and transport margins, whereas purchaser prices include them. For instance, when measuring the consumption of manufacturing products in COICOP, purchases are only allocated to the manufacturing sector, without allocating part of the cost to the transport and trade sectors.

We still need to attribute the trade and transport margins and consider VAT. To do so, we follow the procedure in Cazcarro et al. (2022), which was itself based on Mongelli et al. (2010). For each country and product, the authors compute the positive or negative percentage of trade and transport margins **TTM**, as well as taxes less subsidies **TLS**, and then compute basic prices **BP** from purchaser prices **PP** as

$$\mathrm{BP}_i = \mathrm{PP}_i - \mathrm{TTM}_i - \mathrm{TLS}_i \tag{6.102}$$

for all country/product pairs i. This data is available for 65 CPA products, as the authors break down the real estate sector (L) into two sub-sectors, distinguishing imputed rents and the rest of the sector. We aggregate these back into one sector (as this is the level of aggregation of the concordance tables) and aggregate into the desired level of aggregation. We find a high correlation between the CPA-converted OECD EGDNA data and the Household Final Consumption Expenditures **HFCE** column from the OECD ICIO data. The only exceptions are Trade and Manufacturing, respectively, which were underestimated and overestimated in the EGDNA data. This is precisely because the trade margins of manufacturing consumption have not been allocated yet. After the allocation, the quantities in EGDNA and ICIO match more closely.

Ideally, we would like to match private households' total final domestic consumption (the same as **HFCE**) exactly to the quantities in the ICIO tables. After adding TTM and TLS, the two quantities come closer but are not identical. This is due, in part, to different concepts used in the EGDNA data compared to the HBS data that Cazcarro et al. (2022) use. Thus, the final step is to increase or decrease the consumption of each product by multiplying the consumption of each income quintile by aggregate multiplying factors that ensure that the total consumption of the five income groups matches that in input-output tables. Finally, we normalize the consumption vectors so that they sum to one. This makes it easier to adapt them to new vintages of input-output tables simply by multiplying them by the total household consumption.

6.4.5.2 Initial Conditions

This section describes setting the initial conditions of households. Figure 6.4 shows an example of initial distributions of household income, wealth, and debt for Australia, Austria, and Canada in 2013 at a scale of 1:1000.

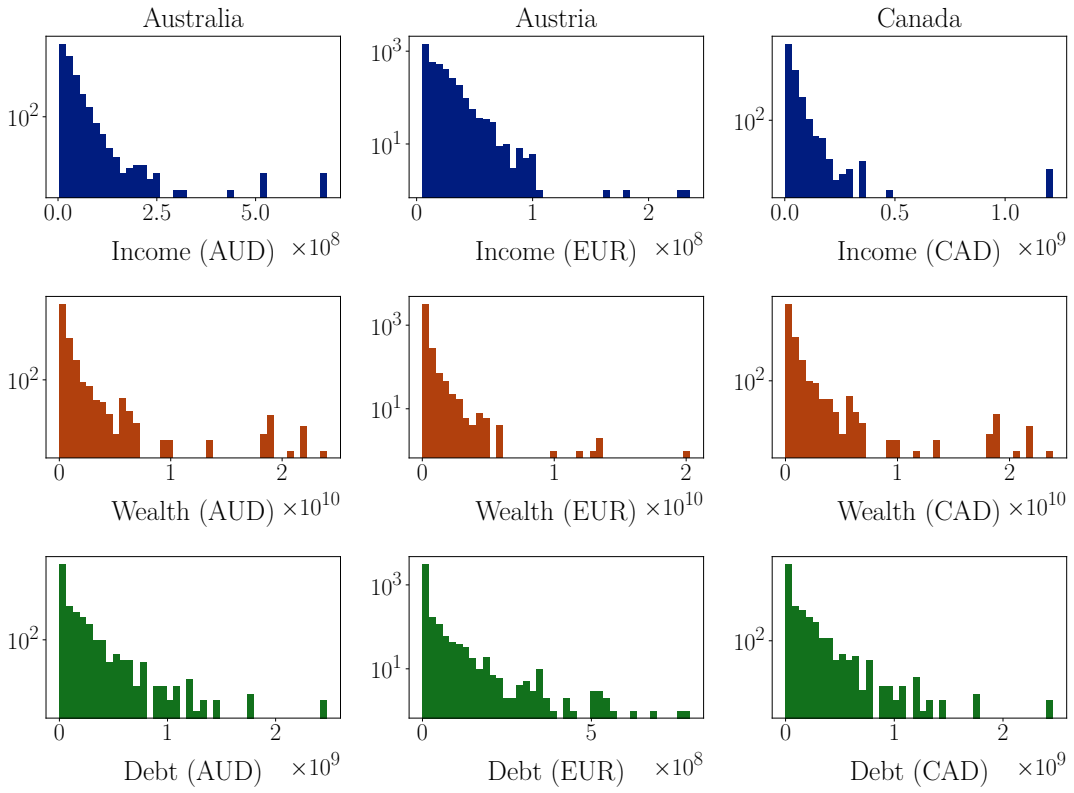


Figure 6.4: Initial distributions of household income, wealth, and debt for Australia, Austria, and Canada in 2013 at a scale of 1:1000, as obtained from the Household Finance and Consumption Survey.

Drawing Households from HFCS Data Households are sampled[64] from *Household Finance and Consumption Survey* (HFCS) data to match OECD aggregates[65] according to the given weights (Code: *HW0010*).[66] The data includes microdata for the fields described in Table 6.12.

Initial Income Initial household income (see Equation (6.108)) from real estate and initial income from financial assets is matched to HFCS data. Other social benefits paid to households are set proportional to regular social transfers and rescaled to match OECD aggregates[67]. Income after rent $Y_h^{\text{-r}}(0)$ is set as the difference between total income (Code: *DI2000*) and rent paid (Code: *HB2300*).

[64] See Network (2023) for details on the sampling procedure and Harland et al. (2012) for an overview of alternative approaches.

[65] OECD: Infra-annual Labour Statistics (Code: *DSD_NASEC20 @ DF_T720R_Q*).

[66] The HFCS is published in four waves: 2010, 2014, 2017, and 2021. For each time of initialisation, we use the most recent HFCS wave. The survey covers the 19 euro area countries, as well as the Czech Republic, Croatia and Hungary. Households of OECD-countries not included are generated homogeneously as in Poledna et al. (2023a).

[67] OECD: Social Expenditure Aggregates (Code: *DSD_SOCX_AGG @ DF_SOCX_AGG*).

Category	Description	Code
Attributes	Type	*DHHTYPE*
	Country	*SA0100*
Income	Rental income from real estate	*DI1300*
	Income from financial assets	*DI1400*
	Income from regular social transfers	*DI1620*
	Income	*DI2000*
Assets	Value of the MHR	*DA1110*
	Value of other properties	*DA1120*
	Value of vehicles	*DA1130*
	Value of valuables	*DA1131*
	Wealth in deposits	*DA2101*
	Mutual funds	*DA2102*
	Bonds	*DA2103*
	Value of private businesses	*DA2104*
	Shares	*DA2105*
	Managed accounts	*DA2106*
	Money owed	*DA2107*
	Other assets	*DA2108*
	Voluntary Pension	*DA2109*
Liabilities	Outstanding balance of mortgages on the MHR	*DL1110*
	Outstanding balance of mortgages on others	*DL1120*
	Outstanding balance of credit line	*DL1210*
	Outstanding balance of credit card debt	*DL1220*
	Outstanding balance of other loans	*DL1230*
Housing	Tenure status of the MHR	*HB0300*
	Rent paid	*HB2300*
	Number of properties other than MHR	*HB2410*
Consumption	Income share spent on consumption	DOCOGOODP

Table 6.12: Household Finance and Consumption Survey codes related to households, in particular income, assets, liabilities, household, and consumption.

Initial Wealth Total household wealth (see Equation (6.129)) $W_h(0)$, wealth in properties $W_h^{\mathrm{P}}(0)$, and wealth in deposits $W_h^{\mathrm{D}}(0)$ is matched to HFCS data. The initial wealth in other real assets $W_h^{\mathrm{ORA}}(0)$ is set to be the sum of the values of vehicles (Code: *DA1130*) and of other valuables (Code: *DA1131*). Wealth in other financial assets W_h^{OFA}, since not explicitly modelled, are assumed to be the sum of mutual funds (Code: *DA2102*), bonds (Code: *2103*), shares (Code: *2105*), managed accounts (Code: *2106*), money owed (Code: *DA2107*), other assets (Code: *2108*), and voluntary pensions (Code: *DA2109*).

Initial Debt Total household mortgage debt $L_h^{\mathrm{M}}(0)$ is set to be the sum of outstanding balances of mortgages on household main residences (Code: *DL1110*) and outstanding balances of mortgages on other properties (Code: *DL1120*). Total household debt in consumption loans $L_h^{\mathrm{C}}(0)$ is set to be the sum of outstanding balances of lines of credit (Code: *DL1210*), credit card debt (Code: *DL1220*), and other non-mortgage loans (Code: *DL1230*).

Initial Consumption The initial household saving rate $\phi_h^{\mathrm{SR}}(0)$ is set to be proportional to 1 − the household income share spent on consumer goods and services (Code: *DOCOGOODP*), and rescaled so that initial household consumption

$$C_{hs}(0) = \frac{1}{\tau^{\mathrm{VAT}}} \left(1 - \phi_h^{\mathrm{SR}}(0)\right) c_{hs}^{\mathrm{CPI}} Y_h(t) \tag{6.103}$$

matches the aggregate input-output tables[68] household final consumption expenditure column. The value-added tax rate is denoted by τ^{VAT}. The consumption weights c_{hs}^{CPI} of household h are set to match input-output tables and shared among households in the same income quintile.[69] Letting $\tilde{c}_{1s}^{\mathrm{CPI}}, \ldots, \tilde{c}_{5s}^{\mathrm{CPI}}$ denote the raw OECD consumption weights by income quintile, we find the model consumption weights by income quintile $c_{1s}^{\mathrm{CPI}}, \ldots, c_{5s}^{\mathrm{CPI}}$ by finding the roots of

$$f(\sigma) = \begin{cases} \displaystyle\sum_{q \in \{1,\ldots,5\}} \sum_{h \in \mathcal{H}_q(0)} \frac{1}{\tau^{\mathrm{VAT}}} \left(1 - \phi_h^{\mathrm{SR}}(0)\right) c_{qs}^{\mathrm{CPI}} Y_h(0) - C_s(0) \text{ for } s \in \mathcal{S} \\ \displaystyle\sum_{s \in \mathcal{S}} \sigma_s - |\mathcal{S}| \\ \displaystyle\sum_{s \in \mathcal{S}} c_{qs}^{\mathrm{CPI}} - 1 \text{ for } q \in \{1, \ldots, 5\} \end{cases} \tag{6.104}$$

where

$$c_{qs}^{\mathrm{CPI}} = c_s^{\mathrm{CPI}} + \sigma_s \left(\tilde{c}_{qs}^{\mathrm{CPI}} - c_s^{\mathrm{CPI}}\right) \tag{6.105}$$

for all quintiles $q \in \{1, \ldots, 5\}$ and for all sectors $s \in \mathcal{S}$. We interpolate between the aggregate consumption weights across all quintiles and the consumption weights by income quintile using Equation (6.105). Our starting estimate is $\sigma_s = 1$ for all $s \in \mathcal{S}$. In Equation (6.104), $C_s(0)$ is the initial total household consumption in sector s obtained from input-output tables. The first term in Equation (6.104) guarantees that the total household consumption at initialisation time matches the aggregate in the input-output table. The second term guarantees that we stay close to the consumption weights by income quintiles and not just revert to the aggregate c_s^{CPI}. The third term guarantees that the consumption weights are normalized.

[68] OECD: Inter-country input-output tables (Code: *ICIO*).

[69] Household consumption weights by income are taken from OECD data (Code: *EGDNA_PUBLIC*, only available in the old data portal: https://stats.oecd.org/).

Initial Investment The investment rate ϕ^{IR} is assumed to be homogenous among households[70] and set so that initial household investment (residential fixed investment)

$$K_{hs}(0) = \frac{1}{\tau^{\mathrm{CF}}}\phi^{\mathrm{IR}} k_s Y_h(t) \tag{6.106}$$

matches the input-output table. The investment weights k_s are set to match input-output tables, and τ^{CF} are taxes on capital formation.

Initial Matching with Banks Households are initially matched with banks by solving a linear sum assignment problem so that the sum of the differences between total household deposits plus debt to total bank deposits and debt from households is minimal.

Initial Matching with Properties Households are initially matched with properties by solving a linear sum assignment problem so that the sum of the differences between property values and the households' total wealth in real assets is minimal. Properties themselves are generated based on the tenure status of the main residence (Codes: *HB0300*, *DA1110*) of each household and their additional number of properties owned (Code: *HB2410*, *DA1120*).

6.4.5.3 Parameters

This section summarizes setting household initial conditions.

Coefficient for Income from Financial Assets The coefficient for household income from financial assets ϕ^{FA} is obtained as the slope of the linear regression of household income from financial assets on household wealth in other financial assets. We assume $\sigma^{\mathrm{FA}} = 0$ for simplicity; see Equation (6.108).[71]

Consumption Smoothing Timeframe The timeframe for consumption smoothing is assumed to be $T^{\mathrm{CO}} = 12$, see Equation (6.109).

Depreciation of Other Real Assets We are assuming a depreciation rate of $d^{\mathrm{RA}} = 5\%$, see Equation (6.126).

[70] This simplifying assumption is due to data availability. In reality, poor households do not invest and rich households invest in large fractions. Further work focussing on inequality should highlight this difference.

[71] This assumption should be revisited in the future. The variations of the return on financial investments are a principal driver of recessions.

Housing as a Tenant We are taking the estimated values for household parameters related to housing as a tenant from Carro et al. (2023). Specifically, when setting the maximum purchasing price in Equation (6.111), we assume $\phi^{\mathrm{HP}} = 42.90$, $\beta^{\mathrm{HP}} = 0.79$, $\mu^{\mathrm{HP}} = -0.018$, and $\sigma^{\mathrm{HP}} = 0.17$. The psychological pressure of renting is set as $\mu^{\mathrm{PS}} = 0.4$, and the influence of costs in deciding whether to buy or rent is $\phi^{\mathrm{B}} = 0.001$. When setting the desired rental price in Equation (6.115), the factor is $\phi^{\mathrm{HR}} = 17.22$ and the exponent is $\beta^{\mathrm{HR}} = 0.35$. We are assuming that the average period for which tenants hold their property is 2 years so that the quarterly probability of not changing housing is given by $p^{\mathrm{RS}} = 7/8$.

Housing as an Owner-Occupier We are taking the values for household parameters related to housing as an owner-occupier from Carro et al. (2023). Specifically, the probability of quarterly price reduction for properties not yet sold is given by $p^{\mathrm{PM}} = 0.1964$, so that $\mu^{\mathrm{PM}} = 1.4531$ and $\sigma^{\mathrm{PM}} = 0.4889$ in Equation (6.117). We are assuming that the average period for which owner-occupiers hold their houses is 20 years so that the quarterly probability of not changing housing is given by $p^{\mathrm{OS}} = 79/80$.

Housing as a Buy-to-Let Investor We are taking the values for household parameters related to housing as a buy-to-let investor from Carro et al. (2023). Specifically, the probability of quarterly rent reduction for properties not yet rented out is given by $p^{\mathrm{RM}} = 0.2848$, so that $\mu^{\mathrm{RM}} = 1.6559$ and $\sigma^{\mathrm{RM}} = 0.7855$ in Equation (6.119). For simplicity, we assume that rents are fully indexed with CPI, $\phi^{\mathrm{PIR}} = 1$ with a lag of one quarter, $T^{\mathrm{PIR}} = 1$.

6.4.5.4 Rules

Households receive income, and based on their income, they choose their target consumption. They operate on the housing market, apply for consumption loans and mortgages, and update their wealth. Households can go bankrupt.

Predicted Income The predicted income of household h is the sum of the predicted incomes of its individuals, predicted other social transfers to the household that depend on household properties (household type and wealth), rental income, and predicted income from financial assets,

$$\overline{Y}_h(t) = \underbrace{\sum_{i \in \mathcal{I}_h(t)} \overline{Y}_i(t)}_{\text{Predicted individual incomes}} + \underbrace{\overline{P}^{\mathrm{CPI}}(t)\mathrm{sb}_h^{\mathrm{O}}(t)}_{\text{Predicted social transfers}} + \underbrace{\sum_{p \in \mathcal{P}_h^{\mathrm{R}}(t)} r_p(t)}_{\text{Rental income}} + \underbrace{\phi^{\mathrm{FA}} W_h^{\mathrm{OFA}}(t-1)}_{\text{Predicted income from financial assets}} \tag{6.107}$$

where $\mathcal{I}_h(t)$ is the set of individuals corresponding to household h, $\overline{P}^{\mathrm{CPI}}(t)$ is the predicted consumer price index, $\mathrm{sb}_h^{\mathrm{O}}(t)$ are real other social transfers to the household, $\mathcal{P}_h^{\mathrm{R}}(t)$ is the set of properties the household has rented out, ϕ^{FA} is a parameter[72], and $W_h^{\mathrm{OFA}}(t)$ are other financial assets of the household.

Income The income of household h is the sum of the incomes of its individuals, other social transfers to the household that depend on household properties (household type and wealth), rental income, and income from financial assets,

$$Y_h(t) = \underbrace{\sum_{i \in \mathcal{I}_h(t)} Y_i(t)}_{\text{Individual incomes}} + \underbrace{P^{\mathrm{CPI}}(t)\mathrm{sb}_h^{\mathrm{O}}(t)}_{\text{Social transfers}} + \underbrace{\sum_{p \in \mathcal{P}_h^{\mathrm{R}}(t)} r_p(t)}_{\text{Rental income}} + \underbrace{(1+\varepsilon)\,\phi^{\mathrm{FA}} W_h^{\mathrm{OFA}}(t-1)}_{\text{Income from financial assets}} \tag{6.108}$$

where $\mathcal{I}_h(t)$ is the set of individuals corresponding to household h, $P^{\mathrm{CPI}}(t)$ is the consumer price index, $\mathrm{sb}_h^{\mathrm{O}}(t)$ are other social transfers to the household, $\mathcal{P}_h^{\mathrm{R}}(t)$ is the set of properties the household has rented out, $\varepsilon \sim N(0, \sigma^{\mathrm{FA}})$ is noise, ϕ^{FA} is a parameter, and $W_h^{\mathrm{ORA}}(t)$ are other financial assets of the household.

Consumption Households participate in the goods market. Target consumption of household h is determined given a minimum consumption amount, an average over previous consumption levels, and a fraction of predicted income[73],

$$\hat{C}_{hs}(t) = \frac{c_s^{\mathrm{CPI}}(t)}{1+\tau^{\mathrm{VAT}}} \max\left(\underbrace{\left(1-\phi_h^{\mathrm{SR}}(t)\right)\overline{P}^{\mathrm{CPI}}(t) w^{\mathrm{U}}(t)}_{\text{Minimal consumption}}, \underbrace{\left(1-\phi_h^{\mathrm{SR}}(t)\right)\overline{Y}_h(t)}_{\text{Fraction of predicted income}}, \underbrace{\frac{\phi^{\mathrm{CO}}}{T^{\mathrm{CO}}} \sum_{t'=1}^{T^{\mathrm{CO}}} C_h(t-t')}_{\text{Consumption smoothing}} \right), \tag{6.109}$$

where $c_s^{\mathrm{CPI}}(t)$ are consumption weights, $\phi_h^{\mathrm{SR}}(t)$ are saving rates of household h, $\overline{P}^{\mathrm{CPI}}(t)$ is the predicted consumer price index, $w^{\mathrm{U}}(t)$ are unemployment benefits, ϕ^{CO} and T^{CO} determine consumption smoothing, $C_h(t)$ is total consumption, and $\overline{Y}_h(t)$ is predicted income. Saving rates $\phi_h^{\mathrm{SR}}(t)$ are updated using a linear model

[72] There is a minimalistic stock market in Dawid et al. (2014), where only an index containing all firms in the economy is traded.

[73] Most generally, the consumption budget is a linear combination of human and financial wealth, e.g. Dawid and Gatti (2018). In Dosi et al. (2010), the household's consumption budget is precisely its income; in Mandel et al. (2010), the consumption budget is a linear function of real money balances, in Lengnick (2013) it is an increasing concave function of real money balances, and in Assenza et al. (2015); Ashraf et al. (2016), it is a function of past incomes and financial wealth. A more complicated implementation in Seppecher (2012) incorporates consumer sentiment and opinion dynamics.

estimated at model initialisation using the *Household Finance and Consumption Survey* (HFCS) data on household income, wealth, and debt.

Investment Households also purchase investment goods in the goods market. The target investment of household h is set as

$$\hat{K}_{hs}(t) = \frac{k_s}{1+\tau^{\text{CF}}} \underbrace{\phi_h^{\text{IR}} \overline{Y}_h(t)}_{\text{Fraction of predicted income}} \tag{6.110}$$

where k_s are fixed sectoral investment weights, τ^{CF} are taxes on capital formation, ϕ_h^{IR} is the investment rate, and $\overline{Y}_h(t)$ is predicted income of household h.[74]

Housing The household may live in social housing, rent, or own its main residence. It may also own additional properties rented out to other households. The household's decisions depend on its current financial situation.

- *Households in Social Housing*: Households in social housing first decide on a desired purchase price based on their predicted income, then predict the maximum housing value they can afford with the purchase price, and then compare the costs of buying with the cost of renting a house of the same value. The household is willing to pay up to

$$P_h(t) = \phi^{\text{HP}} \overline{Y}_h(t)^{\beta^{\text{HP}}} \exp(\varepsilon) \tag{6.111}$$

where ϕ^{HP} and β^{HP} are parameters, and $\varepsilon \sim N(\mu^{\text{HP}}, \sigma^{\text{HP}})$. Households then estimate the value $V^\star$ of housing they can afford by regressing the value of previously sold properties on corresponding prices. They then compare the annual costs of renting a property of value $V^\star$ to the cost of purchasing a property of that value. The perceived annual cost of renting is given by

$$C_{V^\star}^{\text{R}}(t) = 4\left(1+\mu^{\text{PS}}\right)\overline{r}_{V^\star}(t) \tag{6.112}$$

where μ^{PS} represents the psychological pressure of having to rent and $\overline{r}_{V^\star}(t)$ is the predicted rent of a property of value $V^\star$, also regressed on previously previous sales. The annual cost of purchasing the house is given by

$$\begin{aligned} C_{V^\star}^{\text{B}}(t) = & \underbrace{4\frac{P_h(t) - W_h^{\text{FA}}(t)}{m_l} + 4\frac{r^\star\left(P_h(t) - W_h^{\text{FA}}(t)\right)}{1-(1+r^\star)^{m_l}}}_{\text{Annual mortgage repayment}} \\ & - \underbrace{\left(\left(1+\overline{\gamma}^{\text{HPI}}(t)\right)^4 - 1\right)V^\star}_{\text{Predicted house appreciation/depreciation}} \end{aligned} \tag{6.113}$$

[74] Future research should improve the functional form for household investment in Equation 6.110. For instance, it should be pro-cyclical.

where $W_h^{\mathrm{FA}}(t)$ is the financial wealth of household h, m_l is the maturity of the potential new mortgage, $r^\star$ is the average mortgage interest rate observed in the previous timestep, and $\overline{\gamma}^{\mathrm{HPI}}(t)$ is predicted house price index growth. The first term of the annual mortgage repayment above corresponds to paying back the principal, and the second term corresponds to interest. We assume that the mortgage is fixed-rate and that there is no remortgaging.

Then, the probability of buying over renting is given by

$$p_h^{\mathrm{B}}(t) = \frac{1}{1 + \exp\left(\phi^{\mathrm{B}}\left(C_{V^\star}^{\mathrm{R}}(t) - C_{V^\star}^{\mathrm{B}}(t)\right)\right)} \tag{6.114}$$

where ϕ^{B} determines the influence of the difference in predicted costs. If a household buys the property and its financial wealth $W_h^{\mathrm{FA}}(t)$ is greater than the price of the property $P_h(t)$, it will pay for the property using its existing financial wealth (other financial assets first, then deposits). Otherwise, it will apply for a mortgage. If deciding to rent, the rent the household is willing to pay depends on its income and is given by

$$r_h(t) \leq \phi^{\mathrm{HR}} Y_h(t)^{\beta^{\mathrm{HR}}} \tag{6.115}$$

where ϕ^{HR} and β^{HR} are parameters.

- *Households Currently Renting*: Households renting attempt to move with a probability $1 - p^{\mathrm{RS}}$. Otherwise, they follow the same decision process as households in social housing above.

- *Households Currently Owning*: Households currently owning do not attempt to move with a probability p^{OS}. Otherwise, based on their income, households currently decide whether to move into a rented property or buy a different property in the same way as presently renting households. If a household finds a property they would like to rent, they will put their current property p on the market at a price of

$$P_p(t) = \underbrace{\left(1 + \overline{\pi}^{\mathrm{HPI}}(t)\right)}_{\text{Markup}} \underbrace{V_p(t),}_{\text{Property value}} \tag{6.116}$$

where $V_p(t)$ is the value of the property and $\overline{\pi}^{\mathrm{HPI}}(t)$ is the predicted house price index inflation. Each timestep they do not manage to sell the property, the household will reduce the price with a probability of p^{PM} by some random fraction

$$P_p(t) = \underbrace{(1 - \exp\varepsilon)}_{\text{Random reduction}} \; P_p(t-1) \tag{6.117}$$

where $\varepsilon \sim N(\mu^{\mathrm{PM}}, \sigma^{\mathrm{PM}})$, until they manage to sell it.

If the household decides to buy a new home for themselves, they will put their current home on the market and take out a mortgage to buy a new home with the price as before. They will choose their full wealth in financial assets as a down payment.

- *Households as Buy-to-Let Investors*: Buy-to-let investors rent-out their vacant properties. A property p of value $V^\star$ is put on the rental market at a rent

$$r_p(t) = \underbrace{\left(1 + \overline{\pi}^{\mathrm{RPI}}(t)\right)}_{\text{Markup}} \underbrace{r_{V^\star}(t)}_{\text{Average property rent}} \tag{6.118}$$

where $\overline{\pi}^{\mathrm{RPI}}(t)$ is the predicted RPI inflation, $r_{V^\star}(t)$ is the average rent of a property of value $V^\star$, obtained by regressing the value of previously newly rented-out properties on corresponding rents.

Each quarter, a property remains on the rental market, the offered rent is reduced according to

$$r_p(t) = \underbrace{(1 - \exp\varepsilon)}_{\text{Random reduction}} \; r_p(t-1) \tag{6.119}$$

with probability p^{RM}, where $\varepsilon \sim N(\mu^{\mathrm{RM}}, \sigma^{\mathrm{RM}})$.

If a property p is rented out, the rent set by the owner of the house is partially indexed with lagged CPI inflation,

$$r_p(t) = \underbrace{\left(1 + \phi^{\mathrm{PIR}} \pi^{\mathrm{CPI}}(t - T^{\mathrm{PIR}})\right)}_{\text{Lagged partially indexed inflation}} r_p(t-1) \tag{6.120}$$

where ϕ^{PIR} is the partial indexation and T^{PIR} is a lag.[75]

Demand for Credit Households apply for consumption loans to cover liquidity shortages and mortgages for purchasing properties.

If target consumption $\hat{C}_h(t)$ is above disposable income $Y_h^{\text{-r}}(t)$, the household needs to use a fraction of its wealth to make up for the difference. Households use their financial assets (first other financial assets, then deposits) to accommodate additional consumption. If that is not sufficient, the household applies for a consumption loan to cover the gap

$$\hat{L}_h^{\mathrm{C}}(t) = \left[\underbrace{\hat{C}_h(t)}_{\text{Target consumption}} - \underbrace{Y_h^{\text{-r}}(t)}_{\text{Disposable income}} - \underbrace{W_h^{\mathrm{FA}}(t-1)}_{\text{Financial wealth}} \right]^+ . \tag{6.121}$$

[75] In reality, rents are sticky. This assumption should be revisited when studying inflation.

Households may additionally have decided to buy a property for themselves or for renting it out. If the price of the property $P_p(t)$ is above the household financial wealth minus additional necessary spending for consumption, the household will apply for a mortgage of

$$\hat{L}_h^{\mathrm{M}}(t) = \left[\underbrace{P_p(t)}_{\text{Property price}} - \underbrace{\left[W_h^{\mathrm{FA}}(t-1) - \left(\hat{C}_h(t) - \overline{Y}_h^{\mathrm{r}}(t) \right) \right]^+}_{\text{Down-payment}} \right]^+ \tag{6.122}$$

where we assume that the household's desired down payment is their full wealth and financial assets.[76]

Wealth Households[77] hold wealth in real assets,

$$W_h^{\mathrm{RA}}(t) = \underbrace{W_h^{\mathrm{P}}(t)}_{\text{Wealth in properties}} + \underbrace{W_h^{\mathrm{ORA}}(t)}_{\text{Wealth in other real assets}} \tag{6.123}$$

and in financial assets,

$$W_h^{\mathrm{FA}}(t) = \underbrace{W_h^{\mathrm{D}}(t)}_{\text{Wealth in deposits}} + \underbrace{W_h^{\mathrm{OFA}}(t).}_{\text{Wealth in other financial assets}} \tag{6.124}$$

If the total growth of wealth in the current period $Y_h^{\mathrm{r}}(t) - C_h(t)$ is positive, that difference is shared among deposits and other financial assets in fixed fractions depending on current household income, wealth, and debt. If $Y_h^{\mathrm{r}}(t) - C_h(t)$ is negative, a portion of financial household wealth is used up, first other financial assets $W_h^{\mathrm{OFA}}(t)$ and then deposits $W_h^{\mathrm{D}}(t)$. If that is not sufficient, the household applies for a consumption loan to cover the gap $\hat{L}_h^{\mathrm{C}}(t)$.

The wealth in properties of the household is updated as

$$W_h^{\mathrm{P}}(t) = \sum_{p \in \mathcal{P}_h(t)} \underbrace{V_p(t)}_{\text{Value of the property}} \tag{6.125}$$

where $\mathcal{P}_h(t)$ is the set of properties owned by the household and $V_p(t)$ is the value of property p. The value of other real assets depreciates

$$W_h^{\mathrm{ORA}}(t) = \underbrace{\left(1 - d^{\mathrm{RA}}\right) W_h^{\mathrm{ORA}}(t-1)}_{\text{Depreciated current wealth}} + \underbrace{K_h(t)}_{\text{New other real assets}} \tag{6.126}$$

[76] When using the model to study macroprudential policy, this assumption should be revisited.

[77] In Dosi et al. (2010); Assenza et al. (2015); Seppecher (2012); Mandel et al. (2010) households' wealth consists only of deposits, whereas in Lengnick (2013), households' wealth consists of only cash. In Dawid et al. (2014), households hold deposits at banks, a portfolio of shares, and government bonds.

where d^{RA} is the rate of depreciation. The deposits of the household are updated as

$$\begin{aligned} W_h^{\mathrm{D}}(t) = & \underbrace{W_h^{\mathrm{D}}(t-1)}_{\text{Previous deposits}} + \underbrace{\Delta W_h^{\mathrm{D}}(t)}_{\text{New deposits}} - \underbrace{\sum_{l \in \mathcal{L}_h(t)} \frac{V_l(t)}{m_l}}_{\text{Debt installment}} + \underbrace{L_h^{\Delta \mathrm{C}}(t) + L_h^{\Delta \mathrm{M}}(t)}_{\text{Granted loans}} \\ & + \underbrace{r(t)\left[W_h^{\mathrm{D}}(t-1)\right]^{+} - r_b^{\text{H-O}}(t)\left[W_h^{\mathrm{D}}(t-1)\right]^{-}}_{\text{Interest on deposits}} - \underbrace{\sum_{l \in \mathcal{L}_h(t)} r_l V_l(t)}_{\text{Interest paid on debt}} \\ & - \underbrace{\tau^{\mathrm{CF}} \sum_{h \in \mathcal{H}} K_h(t)}_{\text{Taxes on capital formation}} \end{aligned} \tag{6.127}$$

where $\mathcal{L}_h(t)$ is the set of household loans, $L_h^{\Delta \mathrm{C}}(t)$ are newly granted consumption loans, $L_h^{\Delta \mathrm{M}}(t)$ are newly granted mortgages, $r(t)$ is the policy rate, $r_b^{\text{H-O}}(t)$ is the overdraft rate on deposits, $V_l(t)$ is the amount of loan l, m_l is the maturity of loan l, r_l is the interest rate of loan l, τ^{CF} is the tax rate on capital formation, and $K_h(t)$ is new household investment.

The value of other financial assets is updated as

$$W_h^{\mathrm{OFA}}(t) = \underbrace{W_h^{\mathrm{OFA}}(t-1)}_{\text{Wealth in other financial assets}} + \underbrace{\Delta W_h^{\mathrm{OFA}}(t)}_{\text{Change in other financial assets}} \tag{6.128}$$

and total wealth as

$$W_h(t) = \underbrace{W_h^{\mathrm{P}}(t) + W_h^{\mathrm{ORA}}(t)}_{\text{Real assets}} + \underbrace{W_h^{\mathrm{D}}(t) + W_h^{\mathrm{OFA}}(t)}_{\text{Financial assets}}. \tag{6.129}$$

Debt Household debt is set according to

$$L_h(t) = \underbrace{L_h(t-1)}_{\text{Previous debt}} - \underbrace{\sum_{l \in \mathcal{L}_h(t)} \frac{V_l(t)}{m_l}}_{\text{Debt installment}} + \underbrace{L_h^{\Delta \mathrm{C}}(t) + L_h^{\Delta \mathrm{M}}(t)}_{\text{Granted loans}} \tag{6.130}$$

where $\mathcal{L}_h(t)$ is the set of current loans the household pays instalments on, $L_h^{\Delta \mathrm{C}}(t)$ are newly granted household consumption loans, and $L_h^{\Delta \mathrm{M}}(t)$ are newly granted mortgages.

Net Wealth Household net wealth is set to be

$$W_h^{\mathrm{N}}(t) = \underbrace{W_h(t)}_{\text{Wealth}} - \underbrace{L_h(t)}_{\text{Debt}}. \tag{6.131}$$

Bankruptcy If a household h is insolvent, $W_h^{\mathrm{N}}(t) < 0$ (negative net wealth) and $W_h^{\mathrm{D}}(t) < 0$ (negative deposits), it goes bankrupt. The banks receive all of the households' financial wealth, as well as its other owned properties $\mathcal{P}_h^{\mathrm{R}}(t)$. The remaining debt and deposit overdrafts are written off. If the household is in the process of paying off a mortgage for its current residence, the bank takes the residence, and the household will seek to rent or buy in the next iteration.

6.4.6 Individuals

Individuals are part of households. They may be employed, unemployed, or not economically active. If employed, they supply firms with labour and receive a wage. If they are unemployed, they receive social benefits from the government.

Table 6.13 shows variables and parameters related to individuals.

Category	Description	Notation
Sets	Set of individuals	$\mathcal{I}$
	Set of employed individuals	$\mathcal{I}^{\mathrm{E}}(t)$
	Set of unemployed individuals	$\mathcal{I}^{\mathrm{U}}(t)$
	Set of not-economically-active individuals	$\mathcal{I}^{\mathrm{N}}$
Income	Income	$Y_i(t)$
	Predicted income	$\overline{Y}_i(t)$
	Wage	$w_i(t)$
	Reservation wage	$w_i^{\mathrm{RW}}(t)$
	Time period for reservation wages	T^{RW}
Labour	Labour inputs contributed per timestep	$H_i(t)$
	Increase in the quality of labour for employed individuals	h^{E}
	Decrease in the quality of labour for unemployed ind.	h^{U}

Table 6.13: Variables and parameters in the model related to individuals, in particular income and labour.

6.4.6.1 Parameters

We are running the model on relatively short time scales so that we will assume no change in the quality of labour provided by individuals, $h^{\mathrm{U}} = h^{\mathrm{E}} = 0$. The period for reservation wages is set to be $T^{\mathrm{RW}} = 8$, see Equation (6.135).

6.4.6.2 Initial Conditions

This section describes setting initial conditions for individuals.

Drawing Individuals from HFCS Data Each of the randomly drawn households in Section 6.4.5.2 is linked to corresponding individuals (Code: *iid*). The data on individuals includes microdata for the fields described in Table 6.14.

Category	Description	Code
Attributes	Gender	*RA0200*
	Age	*RA0300*
	Labour status	*PE0100*
	Education level	*PA0200*
	Employment sector	*PE0400*
Income	Employment income	*PG0110*
	Self-employment income	*PG0210*
	Income from unemployment benefits	*PG0510*

Table 6.14: Household Finance and Consumption Survey codes related to individuals, in particular personal characteristics and income.

Initial Labour Status The initial labour status of an individual is matched to HFCS data using the mapping described in Table 6.15.

Description	HFCS Code	Activity Status in the Model
Regular work	1	Employed
On leave	2	Unemployed
Unemployed	3	Unemployed
Student	4	Not economically active
Retiree	5	Not economically active
Disabled	6	Not economically active
Military social service	7	Employed
Domestic tasks	8	Unemployed
Other not for pay	9	Unemployed

Table 6.15: Matching between the Household Finance and Consumption Survey codes for the labour status of an individual to the activity status in the model.

The labour status is then adjusted to match the unemployment rate[78] and the vacancy rate[79].

[78]World Bank Data: Unemployment rate (Code: *API_SL.UEM.TOTL.ZS_DS2*).
[79]OECD: Infra-annual Registered Unemployment and Job Vacancies (Code: *DSD_OLAB @ DF_OIALAB_INDIC*).

Initial Employment Sector The initial employment sector of an individual is matched to HFCS data. It is adjusted to match sectoral aggregates[80] if necessary, by changing the employment sector of randomly chosen individuals.

Initial Labour Inputs Each individual's initial labour inputs are assumed to be one: for each $i \in \mathcal{H}$,

$$H_i(0) = 1. \tag{6.132}$$

Initial Wages Initial wages $w_i(0)$ of employed individuals are assumed to be the sum of HFCS employment income (Code: *PG0110*) plus HFCS self-employment income (Code: *PG0210*), since self-employment is not explicitly modelled. Initial wages of unemployed- or non-economically-active- individuals are assumed to be zero.

Initial Income The initial labour income of employed individuals is assumed to be just wages. Initial income of unemployed individuals is proportional to their income from unemployment benefits taken from HFCS data (Code: *PG0510*) and rescaled to match aggregate social benefits[81]. Not-economically active individuals do not directly receive income, but their corresponding household receives social benefits (e.g., for retirement).

Initial Matching with Firms Employed individuals are initially matched with firms by solving a linear sum assignment problem so that the sum of the differences between total firm wages and the sum of received wages before taxes by employees is minimal.

6.4.6.3 Rules

Individuals update their labour supply based on their employment status, choose their reservation wage, and update their predicted income and realized income.

Labour Supply Every individual supplies labour inputs $H_i(t)$ to the firm it employs. Not economically active individuals can not supply labour. The (hypothetical) labour inputs of currently unemployed individuals decrease according to

$$H_i(t) = \underbrace{\frac{1}{1+h^{\mathrm{U}}}}_{\text{Markup}} H_i(t-1) \tag{6.133}$$

[80] OECD: Quarterly Employment by Economic Activity (Code: *DSD_NAMAIN1 @ DF_QNA_BY_ACTIVITY_EMPDC*).

[81] OECD: Social Expenditure Aggregates (Code: *DSD_SOCX_AGG @ DF_SOCX_AGG*).

and the labour inputs of currently employed individuals increase according to

$$H_i(t) = \underbrace{\left(1 + h^{\mathrm{E}}\right)}_{\text{Markup}} H_i(t-1) \tag{6.134}$$

where h^{U} is a parameter that determines the fall in the quality of labour of unemployed individuals and h^{E} is a parameter that determines the rise in the quality of labour of employed individuals.

Reservation Wages The reservation wage of an unemployed individual i is set as the mean of previously received wages and at least what they get from unemployment benefits,

$$w_i^{\mathrm{R}}(t) = \max\left(\underbrace{\overline{P}^{\mathrm{CPI}}(t) w^{\mathrm{U}}(t)}_{\text{Unemployment benefits}}, \underbrace{\frac{1}{T^{\mathrm{RW}}} \sum_{t'=1}^{T^{\mathrm{RW}}} w_i(t-t')}_{\text{Average previous wages}}\right) \tag{6.135}$$

where $\overline{P}^{\mathrm{CPI}}(t)$ is the predicted consumer price index, $w^{\mathrm{U}}(t)$ are real per-capita unemployment benefits, and T^{RW} is the time-span over which an individual is considering its average wage.

Predicted Income The predicted income of an employed individual is

$$\overline{Y}_i(t) = \overline{P}^{\mathrm{CPI}}(t) w_i(t) \underbrace{\left(1 - \tau^{\mathrm{SIW}} - \tau^{\mathrm{INC}}\left(1 - \tau^{\mathrm{SIW}}\right)\right)}_{\text{Taxes}} \tag{6.136}$$

where $\overline{P}^{\mathrm{CPI}}(t)$ is the predicted consumer price index, τ^{SIW} are employees' social contributions and τ^{INC} is the income tax rate. Each unemployed individual receives unemployment benefits, $\overline{Y}_i(t) = \overline{P}^{\mathrm{CPI}}(t) w^{\mathrm{U}}(t)$, not-economically-active individuals estimate their income as part of the other social benefits received by their household, which depends on the household type (see Section 6.4.5.4).

Income The income of an employed individual is

$$Y_i(t) = P^{\mathrm{CPI}}(t) w_i(t) \underbrace{\left(1 - \tau^{\mathrm{SIW}} - \tau^{\mathrm{INC}}\left(1 - \tau^{\mathrm{SIW}}\right)\right)}_{\text{Taxes}} \tag{6.137}$$

where $P^{\mathrm{CPI}}(t)$ is the consumer price index, τ^{SIW} are employees' social contributions and τ^{INC} is the income tax rate. Each unemployed individual receives unemployment benefits, $Y_i(t) = P^{\mathrm{CPI}}(t) w^{\mathrm{U}}(t)$, not-economically-active individuals receive income as part of the other social benefits received by their household, which depends on the household type (see Section 6.4.5.4).

6.4.7 Rest of the World

The "rest of the world" (ROW) models every country we do not specifically simulate. It operates on the global goods market through imports and exports. Corresponding variables are listed in Table 6.16.

Category	Description	Notation
Trade	Total real exports	$Y^{\text{ROW}}(t)$
	Total real target exports	$\hat{Y}^{\text{ROW}}(t)$
	Export weights	c_s^{ROW}
	Total imports	$C^{\text{ROW}}(t)$
	Total target imports	$\hat{C}^{\text{ROW}}(t)$
	Import weights	y_s^{ROW}
	Adjustment speed[82]	ϕ^{ROW}
	Net exports	$\text{NX}^{\text{ROW}}(t)$
Prices	Sectoral price	$P_s^{\text{ROW}}(t)$

Table 6.16: Variables in the model related to the rest of the world, in particular trade and prices.

6.4.7.1 Parameters

The export weights c_s^{ROW} and the import weights y_s^{ROW} are matched to input-output tables[83]. Since we are only simulating a single country interacting with the rest of the world, the adjustment speed is assumed to be $\phi^{\text{ROW}} = 1.0$.

6.4.7.2 Initial Conditions

Initial conditions of the rest of the world are the initial exports $Y^{\text{ROW}}(0)$ and imports $C^{\text{ROW}}(0)$ of every country not explicitly modelled. These are obtained from input-output tables. Initial prices set by the rest of the world are assumed to be $P_s^{\text{ROW}}(0) = 1$ for all sectors s.

6.4.7.3 Rules

The rest of the world determines target imports and exports to the modelled countries and prices.

[83] OECD: Inter-country input-output tables (Code: *ICIO*).

Prices The aggregate price index of all modelled countries is given by

$$\tilde{P}(t) = \frac{\sum_{f\in\tilde{\mathcal{F}}} P_f(t)(Y_f(t) + S_f(t-1))}{\sum_{f\in\tilde{\mathcal{F}}} (Y_f(t) + S_f(t-1))} \tag{6.138}$$

where $\tilde{\mathcal{F}}$ is the set of all firms of all countries, $Y_f(t)$ is output of firm f, $S_f(t)$ is the level of inventories of firm f, and $P_f(t)$ is price of goods produced by firm f. The prices for goods bought from the rest of the world are indexed with the aggregate price index of each modelled country,

$$P_s^{\text{ROW}}(t) = \left[\underbrace{\left(1 + \phi^{\text{ROW}} \left(\tilde{P}(t) - 1\right)\right)}_{\text{Price index}} \times \underbrace{P_s^{\text{ROW}}(0)}_{\text{Initial prices}} \right]^+ \tag{6.139}$$

where $\phi^{\text{ROW}} \in [0, 1]$ is a parameter denoting the effect that price changes in the modelled country have on prices in the rest of the world.

Exports Target exports of the rest of the world are indexed by an aggregate production index.

$$\tilde{Y}(t) = \frac{\sum_{f\in\tilde{\mathcal{F}}} Y_f(t)}{\sum_{f\in\tilde{\mathcal{F}}} Y_f(0)} \tag{6.140}$$

where $\tilde{\mathcal{F}}$ is the set of all firms of all countries. Total real target exports are then set as

$$\hat{Y}^{\text{ROW}}(t) = \left[\underbrace{\left(1 + \phi^{\text{ROW}} \left(\tilde{Y}(t) - 1\right)\right)}_{\text{Production index}} \times \underbrace{\hat{Y}^{\text{ROW}}(0)}_{\text{Initial exports}} \right]^+ \tag{6.141}$$

with weights y_s^{ROW}.

Imports Target imports into the rest of the world are indexed with both the aggregate domestic price index and the aggregate production index,

$$\hat{C}^{\text{ROW}}(t) = \left[\underbrace{\left(1 + \phi^{\text{ROW}} \left(\tilde{P}(t) - 1\right)\right)}_{\text{Price index}} \times \underbrace{\left(1 + \phi^{\text{ROW}} \left(\tilde{Y}(t) - 1\right)\right)}_{\text{Production index}} \times \underbrace{\hat{C}^{\text{ROW}}(0)}_{\text{Initial imports}} \right]^+ \tag{6.142}$$

with weights c_s^{ROW}.

Accounting The cumulative net exports of the rest of the world evolve according to

$$\text{NX}^{\text{ROW}}(t) = \underbrace{\text{NX}^{\text{ROW}}(t-1)}_{\text{Previous cumulative NX}} + \underbrace{\sum_{s\in\mathcal{S}} P_s^{\text{ROW}}(t) Y_s^{\text{ROW}}(t)}_{\text{Exports}} - \underbrace{C^{\text{ROW}}(t)}_{\text{Imports}} \tag{6.143}$$

where $\mathcal{S}$ is the set of sectors.

6.5 Markets

We model four markets[84]: a market for goods (Section 6.5.1) matching buyers (firms, households, government entities, the rest of the world) with sellers (firms, the rest of the world), a market for labour (Section 6.5.2) matching firms with individuals, a market for credit (Section 6.5.3) matching firms and households with banks, and a market for housing (Section 6.5.4) matching households with properties. Below, we provide a detailed description of behavioural rules and parameter calibration.

6.5.1 Goods Market

Firms of all modelled countries sell goods and set prices. Firms also buy goods for use as intermediate and capital inputs. Households and government entities buy goods. The rest of the world sells goods, sets prices, and buys goods. Our random search-and-matching process[85] preserves realistic export- and import ratios while also allowing us to prioritise agents, for instance, when they offer better prices.[86]

6.5.1.1 Parameters

We assume $\phi^{\mathrm{GM}} = 2$, which controls the likelihood of choosing a seller based on the prices they offer; see Equation (6.144). This follows Poledna et al. (2023a).

6.5.1.2 Rules

This section summarizes the modelling rules for clearing the goods market.

Exchange Rates For simplicity, we are assuming that exchange rates between countries are constant.

[84] We do not explicitly model a stock market. For instance, the agents in EURACE (Deissenberg et al., 2008) can purchase an index bond.

[85] All models we reviewed choose buyers and sellers one at a time based on offered prices. In Poledna et al. (2023a), households consider the size of a firm when deciding which firm to buy from.

[86] Some models, such as CATS (Assenza et al., 2015) or EURACE (Deissenberg et al., 2008), differentiate between consumption goods markets and capital goods markets, while others, like Poledna et al. (2023a), do not. In Ashraf et al. (2017), CATS (Assenza et al., 2015), EURACE (Deissenberg et al., 2008), JAMEL (Seppecher, 2012), and LAGOM (Wolf et al., 2013), buyers in the consumption goods market observe a random subset of sellers and choose based on the best-offered prices. The model described in Lengnick and Wohltmann (2013) explicitly models a supply chain, whereas the KS model (Dosi et al., 2010) does not model interactions explicitly. There is no market for capital goods in Ashraf et al. (2017) and Lengnick (2013).

Clearing Supply and demand are allocated among firms, households, government entities, and the rest of the world proportional to their total supply or demand. Individual agents are matched using a random search-and-matching process. Firms are always prioritised as buyers and as sellers if they are larger or offer a better price ranked by

$$\underbrace{\frac{\exp(-\phi^{\mathrm{GM}} P_f(t))}{\sum_{f' \in \mathcal{F}_s(t)} \exp(-\phi^{\mathrm{GM}} P_{f'}(t))}}_{\text{Relative prices}} \times \underbrace{\frac{Y_f(t)}{\sum_{f' \in \mathcal{F}_s(t)} Y_{f'}(t)}}_{\text{Relative production}} \tag{6.144}$$

where ϕ^{GM} is a parameter that governs the influence of the better price in prioritizing sellers. Equation (6.144) specifies the supply chain formation between agents and countries.

Excess Demand If demand is left after matching buyers and sellers, it is distributed among sellers as if the allocation process continues. This can be interpreted as if producers were to send out queries to buyers, asking how much they would additionally purchase at a given price if more supply was available.

6.5.2 Labour Market

The labour market randomly matches firms with potential employees in each country separately. Each firm compares its target labour inputs $\hat{H}^{\mathrm{t}}_f(t)$ to its current labour inputs $H_f(t-1)$. Firms with higher target production attempt to hire more individuals, and firms with current labour inputs higher than the target fire individuals. All firms fire first, and then the remaining firms hire.[87]

Table 6.17 lists variables and parameters related to the labour market.

Category	Description	Notation
Clearing	Number of fired individuals	$N^{\mathrm{F}}_f(t)$
	Number of hired individuals	$N^{\mathrm{H}}_f(t)$
	Firing speed of firms	γ^{F}
	Hiring speed of firms	γ^{H}

Table 6.17: Variables and parameters in the model related to labour market clearing.

[87] The model in Ashraf et al. (2017) considers wages in the search-and-matching processes, in CATS (Assenza et al., 2015) unemployed workers visit a subset of firms and choose the first open vacancy, EURACE (Deissenberg et al., 2008) considers the skill-level of workers, and workers in JAMEL (Seppecher, 2012), LAGOM (Wolf et al., 2013), and the model in Lengnick and Wohltmann (2013) consider reservation wages.

6.5.2.1 Parameters

We assume $\gamma^{\mathrm{F}} = \gamma^{\mathrm{H}} = 1$ to simplify the calibration exercise; see Equation (6.145) and Equation (6.146).

6.5.2.2 Rules

This section summarizes the rules for clearing the labour market.

Firing Each firm fires employees at random until firing any other employee brings its current labour inputs below its target labour inputs or until the labour inputs lost due to firing employees exceeds

$$\underbrace{\gamma^{\mathrm{F}}}_{\text{Firing speed}} \times \underbrace{\left(H_f(t-1) - \hat{H}_f(t)\right)}_{\text{Excess labour supply}} \tag{6.145}$$

where $\gamma^{\mathrm{F}} \in [0, 1]$ is the speed at which workers can be fired, and $H_f(t-1) - \hat{H}_f(t)$ is the difference between current and target labour inputs.

Hiring We iterate over each firm with hiring needs at random. Firms hire individuals seeking jobs in random order until their target labour inputs fall below their given labour inputs or until their labour inputs gained due to new hires exceeds

$$\underbrace{\gamma^{\mathrm{H}}}_{\text{Hiring speed}} \times \underbrace{\left(\hat{Y}_f(t) - H_f(t-1)\right)}_{\text{Missing labour supply}} \tag{6.146}$$

where $\gamma^{\mathrm{H}} \in [0, 1]$ is the speed at which workers can be hired and $\hat{H}_f(t) - H_f(t-1)$ is the difference target labour inputs and current labour inputs. Individuals only accept a job offer if the offered wage is at least at the level of their reservation wage.

6.5.3 Credit Market

The credit market randomly matches firms and households looking for loans to banks willing to provide loans. The market for short-term firm loans is cleared first, then long-term firm loans, then household consumption loans, and finally mortgages. In each case, firms/households are drawn at random, select a random subset of n^{LF} / n^{LH} banks, visit them in order of their offered interest rates, and succeed or fail

to obtain loans based on the credit conditions of the corresponding bank.[88]

Table 6.18 lists parameters related to the credit market.

Category	Description	Notation
Clearing	The maximum number of banks each firm visits	n^{LF}
	The maximum number of banks each household visits	n^{LH}

Table 6.18: Parameters in the model related to credit market clearing.

6.5.4 Housing Market

The housing market module randomly matches buyers and sellers of properties and households looking to rent with households offering rental properties. There are separate markets for buying and renting properties. The market for buying properties is cleared first.[89] We follow the double-auction process proposed in Carro et al. (2023).

Table 6.19 lists variables and parameters related to properties.

Category	Description	Notation
Price	Price	$P_p(t)$
	Predicted annual price of renting	$\overline{P}_p^{\mathrm{R}}(t)$
	Predicted annual price of buying	$\overline{P}_p^{\mathrm{B}}(t)$
	Property value	$V_p(t)$
	Rent	$r_p(t)$
	Predicted rental yield	$\overline{Y}_p(t)$

Table 6.19: Variables in the model related to properties, in particular prices and rent.

[88] There are various implementations of a credit market in the literature. In Ashraf et al. (2017), there is only one bank per sector, and firms can only approach banks in their sector. In EURACE (Deissenberg et al., 2008), firms contact a random set of banks and choose the bank with the lowest offered interest rate. In CATS (Assenza et al., 2015), JAMEL (Seppecher, 2012), and KS (Dosi et al., 2010), there is only a single bank. The models LAGOM (Wolf et al., 2013) and the model in Lengnick and Wohltmann (2013) do not model banks explicitly. See Seppecher (2012); Mandel et al. (2010) for models without any restrictions on credit demands, Dawid et al. (2014); Dosi et al. (2010); Poledna et al. (2023a) for models with certain regulatory requirements, or Ashraf et al. (2016) for a model that allows lenders to seize a collateral.

[89] The most prominent examples of agent-based models for the housing market include Carro et al. (2023); Bardoscia et al. (2024); Gamal et al. (2024) for the United Kingdom, Carro (2023) for Spain, and Mérő et al. (2023) for Hungary. In all cases, the market for sales is cleared before the market for rentals by matching bids and offers in several rounds until no more matches are possible.

Clearing Households looking to buy or rent are drawn at random. They visit the property whose price or rent is closest to what they hope to spend. The purchase of a new property does not go ahead if the household fails to obtain a mortgage if required.

6.6 Conclusion

This chapter introduced a modular framework for constructing large, data-driven, agent-based macroeconomic models that simulate interactions across multiple countries. The framework models interactions between various agents — firms, households, individuals, governments, banks, and central banks — across multiple markets, including goods, labour, credit, and housing. Its modular structure allows for the easy replacement of agents and behavioural rules, enabling comparison with other ABM approaches and providing flexibility for various simulations. Our model is calibrated using real-world data from 38 countries and captures sectoral diversity through 18 NACE-2 economic sectors. Integrating richer market mechanisms and more sophisticated behavioural rules enhances the model's suitability for forecasting and policy analysis. A key feature of this framework is the use of microdata to create heterogeneous agents. The model maintains stock-flow consistency and employs a scale factor to manage computational complexity.

The improvements highlighted here only represent a small exploration of the capabilities of our model. In future work, we intend to explore our model's capabilities further, making the individual behavioural rules more realistic and using more granular and extensive datasets. In the long term, we will run all 38 countries *in parallel*, capturing all the relevant feedback loops of international trade.

At the moment, our framework does not support any form of spatiality, seasonality, an explicit stock market, household demography (and, e.g. bequests), entrepreneurship, mergers & acquisitions, an explicit construction sector, floating rate mortgages, or labour unions. When it comes to specific potential improvements, future work should revisit the pure Leontief production function, which does not permit any substitutions and support of marginal tax rates. The labour market is currently underdeveloped: most critical is that wages do not affect firm labour demand and that there is no nominal wage rigidity. Additionally, the assumption that workers can freely move between sectors should be revisited; the fact that the labour supply decisions of workers do not depend on wealth, on the interest rate, or on income taxes is problematic; and firms should consider their own market power when setting wages. Table 6.20 summarises these points and provides references to (stylised) models in the literature that in some form incorporate what we are potentially missing.

The next chapter will focus on calibrating the agent-based models developed within this framework, and comparing its forecasting performance against a benchmark time series model and another agent-based model.

Aspect	Description
Spatiality	We have no spatiality at the moment, which is crucial for realistically modelling supply chains and the housing market. *June* (Aylett-Bullock et al., 2021) is a spatial macro-epidemic model of the UK. For an overview, see Heppenstall and Crooks (2019).
Seasonality	We do not explicitly consider seasonality. Kimmich et al. (2024) presents a case study using the IIASA model on how prolonged heat affects labour productivity
Firms	The current pure Leontief production function does not permit any substitution. A more realistic function is discussed in Pichler et al. (2020).
Households	We currently do not model household demography (e.g. the *SimPaths* framework (Richiardi et al., 2023)). We also do not explicitly model the level of education of individuals, Kaszowska-Mojsa and Pipień (2020) does.
Government	We are not aware of a macroeconomic agent-based model that features realistic marginal tax rates.
Labour Market	Wages do not affect firm labour demand and there is no nominal wage rigidity. Additionally, the assumption that workers can freely move between sectors should be revisited (see e.g. del Rio-Chanona et al. (2021)); the fact that the labour supply decisions of workers do not depend on wealth, on the interest rate, or on income taxes is problematic; and firms should consider their own market power when setting wages. Our model does not include labour unions. The model by Dosi et al. (2021) does.
Credit Market	Our framework does not currently support floating-rate mortgages. The model for the Hungarian housing market (Mérő et al., 2023) does. We do not include specialised banks that only grant a single type of loan, risk-weighting of assets, and capital buffers (see Bardoscia et al. (2024)).
Housing Market	Residential fixed investment is not pro-cyclical. The model for the Hungarian housing market (Mérő et al., 2023) features an explicit construction sector.
Stock Market	The agents in EURACE (Deissenberg et al., 2008) can purchase firm equity and an index bond. The model by Paulin et al. (2019) explicitly models banks, funds, assets, and traders.
Market Entry	There is no entrepreneurship (e.g. *AMOEBA* (Pablo-Martí et al., 2013)), mergers, acquisitions, or other corporate actions (e.g. Kaszowska-Mojsa and Pipień (2020) features all).

Table 6.20: Aspects of the agent-based modelling framework that could be improved in the future, with references to other macroeconomic agent-based models/frameworks that incorporate them.

7 Economic Forecasting

This chapter presents both the calibration and forecasting experiments of a large-scale data-driven agent-based model. We outline the selection of priors and the training process (Section 7.2). Our model's performance is compared to the IIASA model, which we calibrate for 38 countries (Section 7.3). Section 7.4 highlights the out-of-sample forecasting results for key economic aggregates. To further evaluate our model's predictive capabilities, we estimate a Bayes factor in comparison to the benchmark model (Section 7.5). Finally, we assess the statistical significance of our improvements (Section 7.6), concluding that our model shows statistically significant improvements at the 0.005 *significance level under the specified hypothesis test. Section 7.7 concludes the chapter.*

Contents

7.1 Introduction 199
7.2 Calibration 200
7.2.1 Prior and Posterior Estimation 201
7.2.2 Estimation Details 204
7.3 Running the IIASA Model as a Benchmark 205
7.4 Forecasting Economic Aggregates 206
7.4.1 Average Forecasting Performance 207
7.4.2 Country-Level Forecasting Performance 207
7.5 Bayesian Model Selection 208
7.5.1 Estimating the Bayes Factor 210
7.5.2 Classification Results 211
7.6 Statistical Significance 214
7.7 Conclusion 216

7.1 Introduction

The release of the *IIASA model* (Poledna et al., 2023a) marked a breakthrough in using data-driven agent-based models for macroeconomics: the authors demonstrated that agent-based models could match or exceed the performance of traditional models[1] in both analysing and predicting economic dynamics. The IIASA model was initially calibrated only to Austria, and later expanded to Canada (Hommes et al., 2022) and Italy (Di Domenico and Riccetti, 2023).[2]

This chapter seeks to convince the reader that more realistic modelling rules, the use of comprehensive microdata, and careful calibration lead to substantially better forecasting performance.[3] Specifically, we focus on measuring the improvement in root mean squared error (RMSE) compared to the IIASA model and a baseline AR(1) as a simple time series model. Compared to the literature (Hommes et al., 2022; Poledna et al., 2023a), our validation is more robust by performing Bayes factor estimation and encompassing multiple (38) countries with individual Monte Carlo trajectories. Section 7.4 discusses the forecasting performance of different components of GDP and inflation, and Section 7.5 discusses the methodology of Bayes factor estimation.

Figure 7.1 shows a forecast spanning from 2015-Q1 to 2018-Q1 for real GDP, real consumption, and real investment for Austria. The green line represents true data, and the yellow line corresponds to an iterated AR(1) multiperiod forecast, fitted on quarterly first differences of the log nominal gross domestic product between 1990-Q1 and 2015-Q1. The blue lines indicate the average of 1000 simulation runs of our model, sampled from the posterior predictive distribution associated with the estimated posterior. The solid line corresponds to NPE, and the dashed line corresponds to NRE, each targeting moments of economic aggregates between 1990-Q1 and 2010-Q1. The orange line corresponds to the IIASA model and shows the average of 1000 simulation runs.[4] For this particular country and period of time, our model with NPE performs exceptionally well in predicting consumption; however, all models struggle to predict investment accurately. This may be due to not modelling residential fixed investment (see Table 6.20).

[1] See Section 5.2.

[2] We calibrate the IIASA model to the 38 OECD countries, see Section 7.3.

[3] We use the OECD Quarterly National Accounts (Code: *DSD_NAMAIN1@DF_QNA*) as the ground truth dataset. This includes the Gross Domestic Product (Code: *B1_GE*), Government Consumption (Code: *P3S13*), Household Consumption (Code: *P31S14_S15*), and Gross Fixed Capital Formation (Code: *P51G*).

[4] We run the IIASA model for all 38 OECD member countries; see Section 7.3.

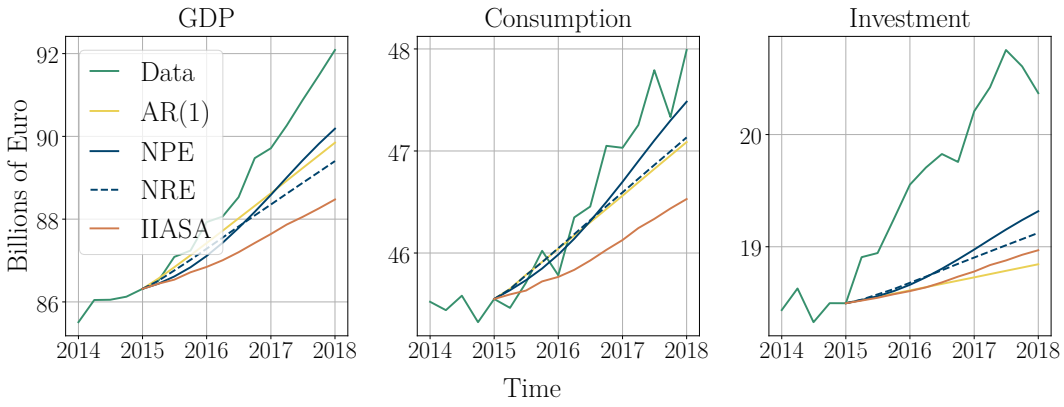


Figure 7.1: Forecasts for real GDP, consumption, and investment for Austria from 2015-Q1 to 2018-Q1. The green line represents true data, the yellow line corresponds to an iterated AR(1) multiperiod forecast, the blue lines indicate the average of 1000 simulation runs of our model, and the orange line corresponds to the IIASA model. For this particular country and period of time, our model with NPE performs exceptionally well in predicting consumption; however, all models struggle to predict investment accurately.

7.2 Calibration

Beyond initialisation, the comprehensive data sources[5] we integrate into our model help establish specific fixed parameters. For instance, the tax rates applied within the model are directly derived from World Bank data. However, it is essential to distinguish these from the "free" parameters of the model, which are not directly observable and therefore require careful calibration. Free parameters typically include behavioural factors that govern how agents extrapolate past data to predict future conditions and make decisions. Since these parameters fundamentally influence the model's dynamics, they must be inferred by systematically comparing the model's time series output against actual economic data. This process forms the core of the calibration task. We discuss the subset of free parameters subject to systematic calibration in Section 7.2.1.

In this work, we employ two advanced Bayesian simulation-based inference procedures: *Neural Posterior Estimation* (NPE) and *Neural Ratio Estimation* (NRE), which have been shown (Dyer et al., 2024) to be both faster and more accurate than other methods commonly employed in the agent-based modelling literature.[6] These methods were introduced in Section 5.4.2.1 and Section 5.4.2.2; we detail the estimation process itself in Section 7.2.2.

[5] See Section 6.2.2 for an overview of the data sources.
[6] For example, the approaches discussed in Grazzini et al. (2017).

7.2.1 Prior and Posterior Estimation

We choose seven parameters among the full set of free model parameters[7] for estimation. This choice was made using observational experiments — the chosen parameters appear to have the strongest positive influence on forecasting performance[8]. Table 7.1 summarises the priors and estimated posteriors using NPE for Austria, which serves as a representative example, as it is the country the IIASA model is initially calibrated to. We calibrate these parameters to match the first moment of growth rates of five time series between 1990-Q1 and 2013-Q1: real gross domestic product, inflation, household consumption, government consumption, and investment. In Table 7.1, $\mathrm{U}(\{0,1\})$ refers to binary parameters chosen with equal likelihood, while $\mathrm{U}([0,1])$ refers to a continuous uniform distribution on the closed interval $[0,1]$.

The first three rows in Table 7.1 are not calibrated using approximate Bayesian inference methods. Instead, for each of the 8 possible combinations, corresponding to $\left(\phi_F^Q, \phi^{\mathrm{DP}}, \phi^{\mathrm{CP}}\right) = (\pm 1, \pm 1, \pm 1)$, we ran the calibration procedure using 10^6 samples and a uniform prior for the remaining four rows of the table. For these rows, we show the first two moments of the estimated posterior corresponding to $\phi_F^Q = \phi^{\mathrm{DP}} = \phi^{\mathrm{CP}} = 0$, which is the configuration that minimises the forecasting error[9] to the true nominal gross domestic product growth rates between 1990-Q1 and 2013-Q1.

Table 7.1 shows summary statistics describing the estimated posterior for Austria. In the following discussion, we explain the meaning and significance of the parameters chosen for estimation. We provided a thorough explanation of the behavioural rules of all agents in Chapter 6, but for the benefit of the reader, we repeat the explanation of the behavioural rules related to firm demand, prices, and target production below.

7.2.1.1 Predicted Firm Demand

In our model, growth is driven by expectations. Specifically, firms set target production based on their expectation of future demand. The predicted future demand of a firm f operating in sector s is set as a markup on previously observed demand $Q_f(t-1)$,

$$\overline{Q}_f(t) = \underbrace{(1+\overline{\gamma}_s(t))}_{\text{Predicted sectoral growth}} \times \underbrace{\left(1+\phi_F^Q \overline{\gamma}_f(t)\right)}_{\text{Predicted idiosyncratic growth}} \times \underbrace{Q_f(t-1).}_{\text{Previous demand}} \tag{7.1}$$

[7] For the full model description, including parameters, see Chapter 6.

[8] Our chosen metric is described in Section 7.4. While we could employ various alternative model selection rules, such as selecting the maximum a posteriori model through a procedure mirroring the Bayes factor computations presented below, we opt for model selection based on this RMSE metric for simplicity.

[9] See the previous footnote 8.

Description	Not.	Prior	Posterior Mean	Posterior Var.
Demand adjustment speed on firm growth	ϕ_F^Q	$\mathrm{U}(\{0,1\})$	0	-
Influence of demand-pull inflation on prices	ϕ^{DP}	$\mathrm{U}(\{0,1\})$	0	-
Influence of cost-push inflation on prices	ϕ^{CP}	$\mathrm{U}(\{0,1\})$	0	-
Target inventory to production fraction	ϕ^{StY}	$\mathrm{U}([0,1])$	0.10	0.00
Influence of labour inputs on target production	χ^{H}	$\mathrm{U}([0,1])$	0.53	0.02
Influence of interm. inputs on target production	χ^{M}	$\mathrm{U}([0,1])$	0.03	0.00
Influence of capital inputs on target production	χ^{K}	$\mathrm{U}([0,1])$	0.18	0.03

Table 7.1: The prior and the corresponding posterior moments estimated by NPE for Austria between 1990-Q1 and 2013-Q1. $\mathrm{U}(\{0,1\})$ refers to binary parameters chosen with equal likelihood, $\mathrm{U}([0,1])$ to the continuous uniform distribution on the closed interval $[0,1]$. The parameters in the prior are related to firm demand predictions, price setting, and target production.

The first term in Equation (7.1) represents the global sectoral growth forecast $\overline{\gamma}_s(t)$, while the second term reflects the firm-specific idiosyncratic growth forecast $\overline{\gamma}_f(t)$[10]. A parameter ϕ_F^Q governs the extent to which the idiosyncratic forecast is considered in predicting future demand. For both NPE and NRE methods, we set $\phi_F^Q = 0$ for all 38 countries, thereby excluding firm-specific growth forecasts.

7.2.1.2 Firm Prices

Firms set prices $P_f(t)$ as a markup on previous prices $P_f(t-1)$,

$$P_f(t) = \underbrace{\left(1+\overline{\pi}^{\mathrm{PPI}}(t)\right)}_{\text{Predicted PPI inflation}} \times \underbrace{\left(1+\phi^{\mathrm{DP}}\overline{\pi}_f^{\mathrm{DP}}(t)\right)}_{\text{Demand-pull inflation}} \times \underbrace{\left(1+\phi^{\mathrm{CP}}\overline{\pi}_f^{\mathrm{CP}}(t)\right)}_{\text{Cost-push inflation}} \times \underbrace{P_f(t-1)}_{\text{Previous price}}. \tag{7.2}$$

The first term in Equation (7.2) represents a global forecast for producer price index (PPI) inflation $\overline{\pi}^{\mathrm{PPI}}(t)$. The second term captures firm-specific demand-pull inflation $\overline{\pi}_f^{\mathrm{DP}}(t)$, which is driven by observed demand, with its effect controlled by the parameter ϕ^{DP}. The third term accounts for firm-specific cost-push inflation $\overline{\pi}_f^{\mathrm{CP}}(t)$, determined by observed unit costs, and governed by the parameter ϕ^{CP}. These firm-specific inflation terms enable the model to account for variations in pricing behaviour across different firms.[11] For NPE and NRE, $\phi^{\mathrm{DP}} = \phi^{\mathrm{CP}} = 0$ for all 38 countries,

[10] The firm-specific growth forecast was introduced in the CANVAS model (Hommes et al., 2022). It is influenced by demand-pull and cost-push inflation, as described in Equation (6.60).

[11] The firm-specific inflation drivers were introduced in the CANVAS model (Hommes et al., 2022), see Equation (6.75) and Equation (6.77).

effectively removing the influence of firm-specific inflation in price-setting decisions.

7.2.1.3 Firm Target Production

Firm target production $\hat{Y}_f(t)$ is set based on predicted demand $\overline{Q}_f(t)$, current inventory, and considers the current financial situation of the firm. It may also be limited by the firm's current workforce, stock of intermediate inputs, and stock of capital inputs. Specifically,

$$\hat{Y}_f(t) = \min \Bigg(\underbrace{\overline{Q}_f(t) + \phi^{\mathrm{StY}} Y_f(t-1) - S_f(t-1)}_{\text{Predicted demand given current inventory}}, \underbrace{\overline{Q}_f(t) + \chi^H \left(H_f(t) - \overline{Q}_f(t) \right)}_{\text{Labour inputs}}, \underbrace{\overline{Q}_f(t) + \chi^M \left(M_f(t) - \overline{Q}_f(t) \right)}_{\text{Intermediate inputs}}, \underbrace{\overline{Q}_f(t) + \chi^K \left(K_f(t) - \overline{Q}_f(t) \right)}_{\text{Capital inputs}} \Bigg), \tag{7.3}$$

where $Y_f(t-1)$ is real output of firm f in the previous time step, ϕ^{StY} is the target inventory to production fraction and $S_f(t-1)$ are inventories kept from the previous time step. For NPE, we obtain a mean of $\phi^{\mathrm{StY}} = 0.06$ across all 38 countries; for NRE, we find $\phi^{\mathrm{StY}} = 0.04$, implying that the role of inventories in setting target production is small. Moreover, $H_f(t)$, $M_f(t)$ and $K_f(t)$ denote maximum production given the firm's labour, intermediate, and capital inputs, respectively.

The firms' labour inputs $H_f(t)$ are set as

$$H_f(t) = \underbrace{h_f(t)}_{\text{Work effort}} \underbrace{\sum_{i \in \mathcal{I}_f(t)} H_i(t)}_{\text{Labour supply}} \tag{7.4}$$

where $\mathcal{I}_f(t)$ is the set of individuals employed by firm f, $H_i(t)$ are labour inputs from individual i, and $h_f(t)$ is a firm-specific factor denoting work effort. This factor follows Poledna et al. (2023a) and corresponds to increased or decreased work effort due to overtime or part-time employment (see Equation (6.67)).

The firm's production is also limited by its level of intermediate and capital inputs, set as

$$M_f(t) = \min_{s' \in \mathcal{S}} \left(\frac{M_{fs'}(t-1)}{m_{s's}} \right) \tag{7.5}$$

$$K_f(t) = \min_{s' \in \mathcal{S}} \left(\frac{K_{fs'}(t-1)}{k_{s's}} \right) \tag{7.6}$$

where $\mathcal{S}$ is the set of sectors, $M_{fs}(t-1)$ is the firms' previous stock of intermediate inputs, and $K_{fs}(t-1)$ is the firms' previous stock of capital inputs of sector s. The matrix $m_{s's}$ denotes the real amount of intermediate inputs of sector s' necessary to

produce one real unit of output of sector s. The matrix $k_{s's}$ denotes the real amount of capital inputs of sector s' necessary to produce one real unit of output of sector s. These functional forms differ significantly from Poledna et al. (2023a). Indeed, Poledna et al. (2023a) aggregates all intermediate inputs into a single intermediate composite and all capital inputs into a single capital composite, essentially assuming a linear production function in each intermediate and capital input. Here, we are stricter about substitution possibilities in the economy and assume a Leontief production function for each input instead.[12]

The parameters $\chi^{\mathrm{H}}, \chi^{\mathrm{M}}, \chi^{\mathrm{K}} \in [0, 1]$, determining the influence of resource constraints on target production, are another key innovation in our framework. We think it unreasonable to assume that firms limit target production from the constraints on labour, intermediate and capital inputs. Instead, it is likely that firms have imperfect knowledge of these constraints and wish to produce more than they could — although they would naturally fail to produce as much as planned if such constraints bind.

Thus, for instance, if $\chi^{\mathrm{H}} = 0$, firms would ignore labour inputs in determining target production, while if $\chi^{\mathrm{H}} = 1$, they would fully consider labour input constraints. Target production directly determines the firm's demand for labour, intermediate inputs, and capital, and so it is clear that limiting it ex-ante creates substantial problems for macroeconomic dynamics. Across all 38 countries, and for both NPE and NRE, the role of intermediate inputs (determined by χ^{M}) is small, suggesting that firms would substantially expand production if the constraints were not binding. Labour inputs and capital inputs are partially binding (averaged across all 38 countries, $\chi^{\mathrm{H}} = 0.54$ and $\chi^{\mathrm{K}} = 0.56$ for NPE, $\chi^{\mathrm{H}} = 0.45$ and $\chi^{\mathrm{K}} = 0.55$ for NRE), and taken into account by firms when setting their production targets.

7.2.2 Estimation Details

We estimate the posterior on the parameters described above using two different methods: Neural Posterior Estimation (NPE, see Section 5.4.2.1) and Neural Ratio Estimation (NRE, see Section 5.4.2.2). Here we provide details on the training procedure.

To estimate the posterior using NPE, we employ a Masked Autoregressive Flow (MAF, Papamakarios et al. (2017)) as a density estimator. The MAF comprises five stacked Masked Autoencoders for Distribution Estimation (MADE, Germain et al. (2015)), each containing two blocks and fifty hidden features. The density estimator is trained using the Adam optimiser (Kingma and Ba, 2014) with a

[12] However, for reasons of data availability, we only run our model with 18 sectors (NACE Rev. 2 Level 1). The IIASA model (Poledna et al., 2023a) and the CANVAS model (Hommes et al., 2022) are ran with 62 sectors (NACE Rev. 2 Level 1).

learning rate of 5×10^{-4}. We allocate 90% of the data for training, reserving the remaining 10% as a hold-out validation set. Training ceases after twenty epochs have elapsed without improvement in performance on the validation set, thereby mitigating the risk of overfitting.

To estimate the posterior using NRE, we utilise a residual neural network (ResNet, He et al. (2016)) with two layers, each containing fifty hidden features, as a classifier to learn the density ratio. Once again, we employ 90% of the data for training, with the remaining 10% set aside as a validation set. We optimise using the Adam algorithm (Kingma and Ba, 2014) with a learning rate of 5×10^{-4}. Training is halted after twenty successive epochs show no improvement in validation set performance, to prevent overfitting.

7.3 Running the IIASA Model as a Benchmark

We compare our calibrated model to the IIASA model, which is the state-of-the-art macroeconomic agent-based model used for economic forecasting. While it is possible to implement the IIASA model in our flexible framework described in Chapter 6, in order to be absolutely certain we represent their model accurately, we simply run their code[13]. The IIASA model was initially calibrated only to Austria (Poledna et al., 2023a), but recently expanded to Canada (Hommes et al., 2022) and Italy (Di Domenico and Riccetti, 2023). We calibrate the IIASA model to the 38 OECD countries by adjusting their routine for setting initial conditions and parameters. This is possible by using OECD data[14] instead of *Eurostat* data[15]. The OECD datasets used are summarised in Table 7.2 below. This table presents the Eurostat codes from Table 1 in Poledna et al. (2023a) along with the corresponding OECD codes.

As in Poledna et al. (2023a); Hommes et al. (2022); Di Domenico and Riccetti (2023), we are running their model on a scale of 1:10000, but, for reasons of data availability, only run their model for 18 sectors (as our framework, see Table 6.1). In the IIASA model, the policy rate is set using a Taylor rule, which is calibrated to Euro area economic growth and inflation. For running non-Euro area countries, we instead calibrate the Taylor rule to the domestic level of economic growth and inflation.

[13] See `https://github.com/iiasa/abm`.
[14] See `https://data-explorer.oecd.org`.
[15] See `https://ec.europa.eu/eurostat/data/database`.

Description	Eurostat Code	OECD Code
Population by current activity status	cens_11an_r2	DSD_LFS @DF_IALFS_INDIC
Business demography by legal form	bd_9ac_l_form_r2	DSD_SDBSBSC_ISIC4 @DF_SDBS_ISIC4
Symmetric input-output table	naio_10_cp1700	ICIO Tables
Cross-classification of fixed assets	nama_10_nfa_st	DSD_NAMAIN10 @DF_TABLE9A
Government revenue and main aggregates	gov_10a_main	DSD_NASEC10 @DF_TABLE12_REV
Government expenditure by function	gov_10a_exp	DSD_NASEC10 @DF_TABLE11
Government non-financial accounts	gov_10q_ggnfa	DSD_NASEC1 @DF_QSA_TRANSACTIONS_C/D
Government debt	gov_10q_ggdebt	DSD_NASEC20 @DF_T7PSD_Q
Financial balance sheets	nasq_10_f_bs	DSD_NASEC20 @DF_T710R_Q
Non-financial transactions	nasq_10_nf_tr	DSD_NASEC1 @DF_QSA_TRANSACTIONS_C
GDP and main components	namq_10_gdp	DSD_NAMAIN1 @DF_QNA
Money market interest rates	irt_st_q	DSD_STES @DF_FINMARK

Table 7.2: Datasets used for calibration of the IIASA model (Poledna et al., 2023a) with corresponding Eurostat and OECD codes.

7.4 Forecasting Economic Aggregates

We can now use our calibrated agent-based model for forecasting economic aggregates and compare its performance against the IIASA model across 38 countries. We measure the forecasting performance for different economic aggregates of a chosen

country and a selected time of initialisation T_1 by calculating the median of the root mean squared error (RMSE) across 1000 Monte Carlo trajectories:

$$\text{RMSE} = \underset{i\in\{1,\ldots,1000\}}{\text{median}} \sqrt{\frac{1}{h}\sum_{t=T_1+1}^{T_1+1+h}\left(X_i(t) - \hat{X}(t)\right)^2} \tag{7.7}$$

where h is the length of the forecasting window, $X_i(t)$ is the simulation output of the i-th Monte Carlo run, and $\hat{X}(t)$ is the corresponding true data at time t. For each time T_1 and each country for which we can initialise the model, we calculate the median of the quantity in Equation (7.7). Specifically, we compute the median across 38 OECD member countries and 20 times of initialisation between 2013-Q1 and 2017-Q4. We have chosen the median in the previous steps because it is a robust metric that is not swayed by outliers. We compare the performance of our model to a benchmark AR(1) iterated multiperiod forecast, fitted on the first differences of the log, and to the IIASA model by examining its relative improvement in median RMSE.

For the remainder of this section, we present results on the average forecasting performance across the full set of 38 countries in Section 7.4.1, and on the results by country in Section 7.4.2.

7.4.1 Average Forecasting Performance

The first super-column of Table 7.3 ("AR1 RMSE") shows the median RMSE of an AR(1) model predicting five different economic aggregates for six different forecasting horizons. As described above, we take the median over the complete set of 38 countries, 20 initialisation times, and 1000 trajectories. The remaining rows of Table 7.3 compare the performance of the agent-based models (the IIASA model and our calibrated models) to the AR(1) by showing the relative improvement[16] in forecasting error (measured as a median across countries and initialisation times) compared to the AR(1). Overall, our model performs significantly better than both the AR(1) and the IIASA model. At the country level (see 7.4.2), our model outperforms the IIASA model across all initialisation times and economic aggregates for 33 out of 38 countries.

7.4.2 Country-Level Forecasting Performance

Table 7.4 shows the median relative improvement in forecasting performance of nominal gross domestic product on a country-level of our model over the IIASA model (Poledna et al., 2023a). The values on the left of the '/' correspond to

[16] Measured by subtracting the median RMSE of the agent-based model from the median RMSE of the AR(1), and then dividing that difference by the median RMSE of the AR(1).

	Aggregate	Horizon 1Q	2Q	3Q	1Y	2Y	3Y
AR1 RMSE	GDP	$7 \cdot 10^{-3}$	$9 \cdot 10^{-3}$	$10 \cdot 10^{-3}$	$10 \cdot 10^{-3}$	$10 \cdot 10^{-3}$	$11 \cdot 10^{-3}$
	Inflation	$8 \cdot 10^{-3}$	$10 \cdot 10^{-3}$	$12 \cdot 10^{-3}$	$12 \cdot 10^{-3}$	$12 \cdot 10^{-3}$	$12 \cdot 10^{-3}$
	Hh. Cons.	$7 \cdot 10^{-3}$	$10 \cdot 10^{-3}$	$10 \cdot 10^{-3}$	$11 \cdot 10^{-3}$	$11 \cdot 10^{-3}$	$12 \cdot 10^{-3}$
	Gov. Cons.	$8 \cdot 10^{-3}$	$10 \cdot 10^{-3}$	$11 \cdot 10^{-3}$	$12 \cdot 10^{-3}$	$12 \cdot 10^{-3}$	$13 \cdot 10^{-3}$
	Inv.	$28 \cdot 10^{-3}$	$39 \cdot 10^{-3}$	$41 \cdot 10^{-3}$	$43 \cdot 10^{-3}$	$45 \cdot 10^{-3}$	$49 \cdot 10^{-3}$
IIASA	GDP	-3%	-5%	-8%	-9%	-12%	-10%
	Inflation	-2%	0%	0%	2%	4%	2%
	Hh. Cons.	-12%	-15%	-16%	-19%	-19%	-22%
	Gov. Cons.	1%	0%	-1%	-1%	-2%	0%
	Inv.	30%	30%	28%	28%	31%	31%
NPE	GDP	41%	43%	38%	43%	63%	62%
	Inflation	44%	55%	49%	50%	59%	63%
	Hh. Cons.	60%	55%	35%	38%	47%	44%
	Gov. Cons.	73%	71%	67%	59%	48%	45%
	Inv.	62%	53%	51%	57%	61%	58%
NRE	GDP	49%	48%	38%	44%	50%	50%
	Inflation	44%	55%	52%	53%	61%	66%
	Hh. Cons.	63%	62%	47%	43%	44%	40%
	Gov. Cons.	66%	75%	69%	59%	48%	37%
	Inv.	56%	58%	53%	59%	62%	56%

Table 7.3: The first super-column ("AR1 RMSE") shows the median RMSE of an AR(1) model predicting five different economic aggregates for six different forecasting horizons. The median is taken over the full set of 38 countries, 20 times of initialisation, and 1000 trajectories. The remaining rows of Table 7.3 compare the performance of the agent-based models (the IIASA model and our calibrated models) to the AR(1) by showing the relative improvement in forecasting error (measured as a median across countries and initialisation times) compared to the AR(1). Overall, our model performs significantly better than both the AR(1) and the IIASA model.

model calibration using NPE, and those on the right correspond to NRE. We highlight country/horizon combinations for which the IIASA model performs better than our model in red.

7.5 Bayesian Model Selection

To further test the forecasting capabilities of our model relative to the benchmark ABM (Poledna et al., 2023a), we estimate a *Bayes factor* for each country and initialisation time. Section 7.5.1 introduces the approach and details the training procedure, and Section 7.5.2 summarises the results.

	Horizon					
	1Q	2Q	3Q	1Y	2Y	3Y
AUS	154 / 113	103 / 57	80 / 76	41 / 92	30 / 88	39 / 111
AUT	30 / 38	37 / 51	31 / 40	36 / 44	46 / 46	56 / 44
BEL	25 / 60	28 / 35	34 / 42	41 / 50	37 / 39	38 / 47
CAN	93 / 107	38 / 38	55 / 51	43 / 50	66 / 70	77 / 101
CHE	82 / 89	79 / 85	62 / 60	96 / 93	102 / 102	82 / 83
CHL	59 / 38	94 / 79	76 / 71	87 / 70	88 / 74	84 / 79
COL	101 / 104	130 / 101	103 / 93	90 / 91	75 / 83	91 / 83
CRI	108 / 211	150 / 206	179 / 166	174 / 163	141 / 143	152 / 132
CZE	1 / -22	-14 / -8	4 / 0	8 / 3	19 / 8	32 / 6
DEU	29 / 56	-21 / -5	-7 / 2	-10 / -8	21 / -16	20 / -7
DNK	107 / 109	104 / 109	106 / 103	103 / 103	85 / 97	73 / 77
ESP	40 / 39	-29 / 20	-36 / 30	-32 / 47	0 / 77	15 / 92
EST	77 / 78	56 / 56	73 / 74	78 / 80	89 / 90	99 / 98
FIN	98 / 111	49 / 50	51 / 47	72 / 70	94 / 97	92 / 92
FRA	-3 / -19	-38 / 2	-33 / -14	-35 / -28	-34 / -48	-27 / -44
GBR	200 / 174	135 / 70	99 / 62	44 / 66	72 / 55	73 / 66
GRC	0 / 0	36 / 53	48 / 77	58 / 60	71 / 97	65 / 83
HUN	53 / 25	37 / 29	33 / 20	47 / 33	47 / 35	66 / 43
IRL	2 / 3	12 / 13	29 / 28	39 / 40	25 / 25	25 / 26
ISL	82 / 85	42 / 39	32 / 40	56 / 57	62 / 60	67 / 66
ISR	117 / 112	94 / 91	84 / 82	88 / 87	98 / 98	87 / 87
ITA	13 / 15	-87 / -78	-85 / -74	-84 / -72	-77 / -58	-73 / -51
JPN	66 / 87	32 / -1	18 / -7	0 / -26	-14 / -90	-88 / -92
KOR	88 / 110	83 / 143	74 / 161	56 / 162	99 / 230	149 / 225
LTU	188 / 188	217 / 198	216 / 208	176 / 204	149 / 146	152 / 130
LUX	164 / 149	197 / 181	163 / 172	155 / 150	198 / 197	206 / 213
LVA	184 / 185	205 / 259	215 / 229	235 / 259	269 / 249	277 / 263
MEX	24 / 24	81 / 79	85 / 78	106 / 95	112 / 103	111 / 109
NLD	4 / 3	13 / 13	32 / 32	35 / 35	38 / 38	36 / 36
NOR	47 / 48	41 / 41	46 / 46	52 / 53	37 / 39	27 / 28
NZL	2 / 7	21 / 22	32 / 37	54 / 58	91 / 95	78 / 79
POL	189 / 35	182 / 0	157 / 9	163 / 24	135 / 50	110 / 72
PRT	-9 / -18	-14 / -9	-8 / -5	10 / 13	28 / 22	29 / 21
SVK	116 / 149	136 / 173	134 / 149	132 / 162	140 / 141	164 / 188
SVN	32 / 2	47 / 6	30 / 0	63 / 36	65 / 59	52 / 21
SWE	14 / 68	13 / 67	27 / 57	29 / 44	24 / 35	32 / 44
TUR	88 / 87	75 / 75	39 / 42	-17 / 21	3 / 13	9 / 19
USA	75 / 84	146 / 159	136 / 172	121 / 129	106 / 122	103 / 109

Table 7.4: Median relative improvements in RMSE of our model's growth rates of nominal GDP over the IIASA model (Poledna et al., 2023a). The values on the left of the '/' correspond to model calibration using NPE, and those on the right correspond to NRE. The model was run for 20 initialisation times between 2013-Q1 and 2017-Q4 in both cases.

7.5.1 Estimating the Bayes Factor

The Bayes factor $B(\mathcal{M}, \mathcal{M}', \boldsymbol{y})$ (Jeffreys, 1998) for our model $\mathcal{M}$ against the IIASA model $\mathcal{M}'$ and for data $\boldsymbol{y}$ is the ratio of the two models' marginal likelihood functions:

$$B(\mathcal{M}, \mathcal{M}', \boldsymbol{y}) = \frac{m(\boldsymbol{y} \mid \mathcal{M})}{m(\boldsymbol{y} \mid \mathcal{M}')}. \tag{7.8}$$

For our model, the marginal likelihood function is the likelihood function at data $\boldsymbol{y}$ integrated over the prior distribution over the model's free parameters, $\boldsymbol{\theta}$:

$$m(\boldsymbol{y} \mid \mathcal{M}) = \int p(\boldsymbol{y} \mid \boldsymbol{\theta}, \mathcal{M})\, \pi(\boldsymbol{\theta} \mid \mathcal{M})\, \mathrm{d}\theta. \tag{7.9}$$

In contrast, the IIASA benchmark $\mathcal{M}'$ has no free parameters; $m(y \mid \mathcal{M}')$ is, therefore, simply the distribution over model output resulting from the model's use of pseudorandom numbers internally.

Given the use of the prior (rather than a posterior) density in Equation 7.9, Bayes factors assess the relative predictive performance of the two models $\mathcal{M}$ and $\mathcal{M}'$ treating all data as out-of-sample. That is, Bayes factors can be interpreted as the predictive probability (density) for the data before it becomes available for use in parameter estimation tasks, and a summary of the evidence provided by unseen data in favour of one probabilistic model over another (Kass and Raftery, 1995). Further, Bayes factors naturally penalise a model for any increased complexity introduced by the appearance of additional model parameters, due to the fact that they integrate the model's likelihood function over the prior. Finally, in contrast to alternative approaches to model comparison (e.g., non-Bayesian likelihood ratio tests), Bayes factors do not require the models being compared to be nested[17]. Therefore, Bayes factors provide a suitable vehicle for comparing the predictive performance of our more highly parameterised model against the IIASA baseline.

In our Bayes factor computations, we choose a uniform parameter prior $\pi(\boldsymbol{\theta} \mid \mathcal{M})$ over the (bounded) support of the parameters for our model[18]. Given that the models' likelihood functions (and therefore marginal likelihood functions) are difficult to derive, given the complexity of $\mathcal{M}$ and $\mathcal{M}'$, we estimate Bayes factors using density ratio estimation: we train a real-valued neural network $b_{\boldsymbol{\phi}}$ to minimise (Monte Carlo estimates of) the loss function

$$\mathcal{L}(\boldsymbol{\phi}) = -\mathbb{E}_{\boldsymbol{x} \sim m(\cdot \mid \mathcal{M})} \left[\log \frac{\exp(b_{\boldsymbol{\phi}}(\boldsymbol{x}))}{1 + \exp(b_{\boldsymbol{\phi}}(\boldsymbol{x}))} \right] - \mathbb{E}_{\boldsymbol{x}' \sim m(\cdot \mid \mathcal{M}')} \left[\log \frac{1}{1 + \exp(b_{\boldsymbol{\phi}}(\boldsymbol{x}'))} \right], \tag{7.10}$$

[17] Model A is nested within model B with parameterisation $\xi = (\alpha, \beta)$ (with α and β scalars or vectors) if model A is attained from model B when $\beta = \beta_0$ for some β_0.

[18] We discuss in Section 7.6 that this prior is not entirely suitable, and is likely to lead to Bayes factors that are less favourable for our model than other more sensible priors.

which induces the network to learn the log-ratio of the marginal likelihoods (Gutmann and Hyvärinen, 2012), i.e.,

$$b_{\hat{\boldsymbol{\phi}}}(\boldsymbol{x}) \approx \log \frac{m(\boldsymbol{x} \mid \mathcal{M})}{m(\boldsymbol{x} \mid \mathcal{M}')} = \log B(\mathcal{M}, \mathcal{M}', \boldsymbol{x}) \tag{7.11}$$

for any given data $\boldsymbol{x}$, where $\hat{\boldsymbol{\phi}}$ are the trained neural network weights. We then estimate $B(\mathcal{M}, \mathcal{M}', \boldsymbol{y})$ as $\exp(b_{\hat{\boldsymbol{\phi}}}(\boldsymbol{y}))$.

We estimate the Bayes factors using neural networks. We take this neural network to be a 2-layer gated recurrent unit (GRU, Cho et al. (2014)) followed by a feedforward neural network with hidden sizes 32, 32, 32, 16, which produces as output a single real value estimating the log of the Bayes factor. Each hidden layer uses ReLU activations. The GRU's input layer size corresponds to the number of observed time series[19], and the size of the hidden layer is 32. The FF neural network takes 32 input dimensions and has two 32-dimensional hidden layers, with each layer using ReLU activations. As input features, we use the mean of real GDP, GDP deflator, real household consumption, real government consumption, and real gross fixed capital formation, resulting in 5 input features. We train using AdamW (Loshchilov and Hutter, 2019) and a learning rate of 10^{-3}, using 80% of the data for training and the remaining 20% as validation. To avoid overfitting, we train for a maximum of 500 epochs, with early stopping if the validation set performance does not improve for 50 consecutive epochs.

7.5.2 Classification Results

We estimate Bayes factors for the 38 OECD countries at 20 times of initialisation between 2013-Q1 and 2017-Q4. Each ratio estimator is trained using 1000 Monte Carlo samples for three years. In particular, the ratio estimators are trained on summary statistics derived from five macroeconomic aggregates[20] simulated over these three-year periods. Of all estimated Bayes factors, approximately 35% support the IIASA model, and the remaining Bayes factors support our model to a varying degree.

Tables 7.5 and 7.6 show the classification of the Bayes factors are estimated by country and time of initialisation, using the classification scheme from Jeffreys (1998), summarised in Table 7.7 below. Overall Bayes factors are computed for each initialisation period using the fact that each country is simulated separately and shares no parameters; thus, the overall Bayes factors are then products of the country-wise Bayes factors within each initialisation period.

[19] We use six time series: Real GDP, GDP deflator, real gross output, real household consumption, real government consumption, real gross fixed capital formation.

[20] See Table 7.3.

	2013				2014				2015			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
AUS	D	D	D	D	D	D	VSt	VSt	D	D	Su	Su
AUT	N	N	Su	D	VSt	D	D	D	D	D	D	D
BEL	N	N	N	N	N	St	D	N	Su	B	D	D
CAN	N	N	D	St	N	N	D	D	B	VSt	St	St
CHL	N	N	N	N	N	N	N	N	N	N	N	N
COL	B	Su	N	N	N	Su	N	N	N	N	N	N
CRI	D	D	St	VSt	D	D	D	D	D	D	D	D
CZE	D	St	N	St	N	N	N	Su	VSt	Su	D	D
DNK	D	D	N	D	N	N	N	D	D	D	N	N
EST	N	VSt	VSt	St	D	D	N	N	N	N	N	D
FIN	B	N	D	St	D	VSt	D	D	D	D	D	D
FRA	VSt	VSt	Su	Su	N	Su	Su	D	D	D	D	D
DEU	St	Su	D	VSt	N	St	B	Su	VSt	D	VSt	B
GRC	N	N	N	N	D	D	D	D	D	D	D	N
HUN	N	N	St	N	N	N	N	N	Su	D	N	N
ISL	N	N	D	B	N	N	N	N	N	N	N	N
IRL	D	D	D	D	D	D	D	D	D	D	D	D
ISR	N	St	Su	N	B	N	St	VSt	D	D	D	N
ITA	Su	N	St	D	N	St	D	D	D	VSt	Su	D
JPN	N	N	D	N	N	N	N	N	N	N	N	N
KOR	Su	D	D	Su	D	D	D	D	St	D	St	St
LVA	N	D	D	VSt	N	D	Su	D	Su	N	D	D
LTU	D	N	D	N	N	D	D	D	D	D	D	D
LUX	N	N	N	N	N	VSt	N	N	N	N	N	N
MEX	N	N	St	VSt	N	N	N	N	N	N	N	N
NLD	D	D	N	D	D	D	D	D	D	D	D	D
NZL	D	D	St	D	N	D	N	D	D	D	D	D
NOR	N	N	N	N	N	N	N	N	N	N	N	N
POL	N	D	N	N	D	D	D	D	D	D	D	D
PRT	N	N	N	N	D	D	N	N	D	D	D	D
SVK	D	D	D	D	D	D	D	D	VSt	D	N	N
SVN	N	N	N	N	D	D	D	N	D	D	St	D
ESP	Su	D	D	D	D	Su	D	D	St	VSt	St	VSt
SWE	D	D	D	D	D	St	D	D	D	D	D	B
CHE	D	D	D	D	B	D	D	D	D	D	D	Su
TUR	N	N	N	N	N	N	N	B	N	N	N	D
GBR	D	D	D	D	D	VSt	D	D	D	B	Su	St
USA	N	B	D	N	D	D	D	D	St	D	D	St
Overall	N	D	N	N	D	D	D	D	D	D	D	D

Table 7.5: Estimated Bayes factors by the time of initialisation and country, from 2013 to 2015.

	2016				2017			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
AUS	B	St	N	VSt	VSt	VSt	D	D
AUT	D	D	D	D	D	N	N	N
BEL	D	D	N	D	VSt	D	D	D
CAN	St	St	VSt	Su	B	D	N	D
CHL	N	Su	D	B	D	VSt	D	D
COL	N	N	N	D	N	N	D	N
CRI	St	N	N	N	N	N	N	N
CZE	D	D	D	D	D	D	D	D
DNK	N	N	N	N	N	N	N	N
EST	N	D	D	N	St	N	N	D
FIN	D	St	N	N	D	N	N	N
FRA	D	D	D	D	St	D	D	D
DEU	N	N	N	N	N	D	N	N
GRC	D	D	D	D	D	D	D	Su
HUN	D	D	D	D	D	N	N	B
ISL	N	N	N	N	N	N	N	N
IRL	D	D	D	D	D	N	D	D
ISR	St	St	N	VSt	Su	VSt	N	D
ITA	B	D	N	B	N	B	N	D
JPN	N	N	N	N	N	B	N	N
KOR	B	St	St	D	D	B	VSt	St
LVA	D	D	D	D	D	N	N	N
LTU	D	N	D	D	D	St	D	D
LUX	N	N	VSt	N	N	N	D	N
MEX	N	N	N	N	N	N	D	D
NLD	Su	St	N	VSt	D	N	N	VSt
NZL	D	D	D	D	D	D	N	D
NOR	N	N	N	N	N	N	N	N
POL	VSt	D	D	D	D	D	D	St
PRT	D	D	D	D	D	D	D	N
SVK	VSt	VSt	D	St	N	N	N	N
SVN	D	D	D	D	D	N	D	N
ESP	D	D	N	D	N	D	D	D
SWE	N	N	N	D	N	N	N	N
CHE	D	D	D	D	VSt	N	VSt	N
TUR	St	N	VSt	N	N	D	N	D
GBR	D	St	St	Su	B	D	D	D
USA	N	D	Su	N	D	D	D	VSt
Overall	D	D	D	D	D	D	D	D

Table 7.6: Estimated Bayes factors by the time of initialisation and country, from 2016 to 2017.

Range	Interpretation	Count
$[0, 1)$	Negative (N)	282
$[1, 10^{1/2})$	Barely Worth Mentioning (B)	24
$[10^{1/2}, 10)$	Substantial (Su)	33
$[10, 10^{3/2})$	Strong (St)	47
$[10^{3/2}, 10^2)$	Very Strong (VSt)	42
$> 10^2$	Decisive (D)	332

Table 7.7: Classification of Bayes factor estimates as reported in Jeffreys (1998), and counts of each Bayes factor class in Tables 7.5 and 7.6.

7.6 Statistical Significance

In this section, we detail our calculations for assessing the statistical significance of the improvements in model performance over the baseline IIASA model.

Let n be the number of countries on which we test the performance of our model and the performance of the IIASA model. Let X_i, $i = 1, \ldots, n$ be an indicator variable which takes value 1 if our model outperforms the IIASA model according to some metric (e.g., Bayes factor, RMSE forecasting error, etc.) on country i, and 0 otherwise. We assume that the X_i are independent and identically distributed Bernoulli random variables, each with a common success probability p.[21] Then, the number of successes

$$T = \sum_{i=1}^{n} X_i \sim \text{Bin}(\cdot \mid n, p) =: g_{n,p}(\cdot), \tag{7.12}$$

where $\text{Bin}(\cdot \mid n, p)$ denotes the Binomial distribution with n trials and success probability p.

We consider the following composite hypothesis test, in which we test

$$H_0 : p \in \Theta_0 = [0, 1/2] \quad \text{vs.} \quad H_1 : p \in \Theta_1 = (1/2, 1], \tag{7.13}$$

with $\Theta := \Theta_0 \cup \Theta_1$. The null hypothesis H_0 corresponds to the case in which the IIASA model outperforms our model, while the alternative corresponds to our own model outperforming the IIASA baseline.

A common test statistic for composite null-composite alternative hypothesis tests is the generalised likelihood ratio,

$$\lambda_{\text{GLR}}(T) = -2 \log \frac{\sup_{p \in \Theta_0} g_{n,p}(T)}{\sup_{p \in \Theta} g_{n,p}(T)}. \tag{7.14}$$

[21] This is a common assumption. If the random variables are independent and identically distributed Bernoulli random variables, then if all the T binary outcomes are divided into $\frac{T}{N}$ groups of N subsets of binary outcomes, then the total number of successes within each group should resemble a Binomial distribution. This is the case.

The maximum likelihood estimate $\hat{p}_{\text{MLE}}$ for p is

$$\hat{p}_{\text{MLE}} = \arg\sup_{p\in\Theta} g_{n,p}(T) = \frac{T}{n}. \tag{7.15}$$

Thus, we have that

$$\lambda_{\text{GLR}}(T) = -2\log\frac{\sup_{p\in\Theta_0} g_{n,p}(T)}{g_{n,\frac{T}{n}}(T)}. \tag{7.16}$$

From the formula for the Binomial distribution, we have that

$$\arg\sup_{p\in\Theta_0} g_{n,p}(T) = \begin{cases} T/n & \text{if} \quad T \le n/2 \\ 1/2 & \text{if} \quad T > n/2, \end{cases} \tag{7.17}$$

since $p \mapsto g_{n,p}(T)$ is strictly concave in p on $[0,1]$, which gives

$$\lambda_{\text{GLR}}(T) = \begin{cases} 0 & \text{if} \quad T \le n/2 \\ 2\left(n\log 2 + T\log\frac{T}{n} + (n-T)\log\left(1-\frac{T}{n}\right)\right) & \text{if} \quad T > n/2. \end{cases} \tag{7.18}$$

Further, for a given p, we have that

$$\lambda_{\text{GLR}}(T) \mid p = \begin{cases} 0 & \text{with prob.} \quad \sum_{k=1}^{n/2} g_{n,p}(k) \\ 2\left(n\log 2 + k\log\frac{k}{n} + (n-k)\log\left(1-\frac{k}{n}\right)\right) & \text{with prob.} \quad g_{n,p}(k),\ k > n/2. \end{cases} \tag{7.19}$$

For this test to have a significance level of α, we seek a $\tilde{\lambda}$ such that

$$\sup_{p\in\Theta_0} \mathbb{P}\left(\lambda_{\text{GLR}}(T) \ge \tilde{\lambda} \mid p\right) \le \alpha. \tag{7.20}$$

That is, we seek the value $\tilde{\lambda} \in (0,\infty)$ such that values for the test statistic that are more extreme than $\tilde{\lambda}$ occur no more than $(100\alpha)\%$ of the time under the null hypothesis. Since $\lambda_{\text{GLR}}(k)$ is non-decreasing in k, the probability with which $\lambda_{\text{GLR}}(k) \ge \tilde{\lambda}$ is the probability with which $k \ge \tilde{k}$ for some $\tilde{k}$ corresponding to $\tilde{\lambda}$. It is therefore sufficient to find $\tilde{k}$ such that

$$\sup_{p\in\Theta_0} \mathbb{P}(k \ge \tilde{k} \mid p) \le \alpha \tag{7.21}$$

for our chosen significance level α. We also have that $\mathbb{P}(k \ge \tilde{k} \mid \frac{1}{2}) > \mathbb{P}(k \ge \tilde{k} \mid p)$ for any $0 \le p < 1/2$, so in particular it is sufficient to find a $\tilde{k}$ such that

$$\mathbb{P}\left(k \ge \tilde{k} \,\middle|\, \frac{1}{2}\right) \le \alpha. \tag{7.22}$$

For $n = 38$ countries, we can easily obtain numerically the fact that the lowest such $\tilde{k}$ is $\tilde{k} = 25$, giving

$$\mathbb{P}\left(k \geq 25 \;\middle|\; \frac{1}{2}\right) \approx 0.036 \leq 0.05. \tag{7.23}$$

Thus, a significance level of $\alpha = 0.05$ corresponds to succeeding in outperforming a baseline model in 25 of the 38 countries, and a significance level of $\alpha = 0.005$ corresponds to 28 out of 38 countries. Similarly, treating all 760 entries of Tables 7.5 & 7.6 and all 228 entries of Table 7.4 as independent, we require that our model outperforms the baseline on, respectively, 402 and 126 entries for a significance level of 0.05, or 415 and 133 entries for a significance level of 0.005.

From the results presented in Table 7.4 of Section 7.4.2, as well as Tables 7.5 & 7.6 of Section 7.5.2, we see that our improvement in predictive performance is statistically significant at the 0.005 significance level under this hypothesis test across all prediction horizons when the model is calibrated with either NPE or NRE. In particular, for the Bayes factor computations, it is noteworthy that our chosen uniform prior is not entirely suitable, given that a uniform prior places equal weight on regions of the parameter space that produce economic collapse in our model, which we reasonably believe to be implausible behaviour. It is, therefore, likely that adjusting the prior place greater weight on more sensible economic outcomes than economic collapse will result in even more favourable Bayes factors for our model.

7.7 Conclusion

This chapter demonstrates the effectiveness of calibrating a large-scale data-driven agent-based model to multiple countries, highlighting its utility in economic forecasting. By employing advanced Bayesian simulation-based inference techniques, we significantly improved the model's predictive performance compared to the benchmark model. The inclusion of more realistic behavioural rules, calibrated across the 38 OECD countries, underscores the robustness of our model in forecasting key economic aggregates such as GDP, inflation, and consumption. The detailed comparison with the IIASA model, conducted via out-of-sample forecasts and Bayes factor estimation, provides strong evidence that our model outperforms the benchmark in both forecasting accuracy and predictive robustness. The results suggest a statistically significant improvement at the 0.005 level. Moreover, our methodological advances, such as incorporating microdata and applying more flexible behavioural rules, have proven pivotal in achieving these performance gains. The ability of our model to outperform simpler time series models and

established agent-based models underscores the importance of careful calibration and comprehensive data integration.

Since our focus for this chapter was on the benefits that come from careful calibration, additional data, and improved behavioural rules, we compared our agent-based model against another agent-based model. Future work could compare the performance of a calibrated agent-based model against conventional economic models (such as HANK, see Section 5.2.1.4), or against more advanced time series models (e.g. Bayesian VARs). Additionally, future research can explore further refinements in the model's structural assumptions, including expanding the scope of behavioural rules by removing some of the simplifying assumptions made to lower computational complexity and improving the prior selection process for Bayesian inference. These refinements could further enhance the model's predictive accuracy, solidifying its role as a valuable tool for macroeconomic analysis and policy evaluation alongside conventional models.

Conclusion

8 Conclusion

In this thesis, we have studied economic dynamics and forecasting using two complementary approaches — (i) a microeconomic model of agents repeatedly playing normal-form games and (ii) a large-scale multi-country macroeconomic model. This chapter summarises our key contributions and highlights the limitations and potential future directions of this work.

Part I: Random Games as a Null Model At the microeconomic scale on the level of individual agents, we investigated the dynamics of various learning processes on randomly generated normal-form games. Our baseline is the best-response dynamic, under which players myopically choose the best responses to their current environment.

Results on the Convergence Frequency and Likelihood of Games Our key result is that the sequence of players selecting their optimal strategies is crucial for convergence. Specifically, we showed that the clockwork best-response dynamic does not converge to a pure Nash equilibrium as the number of players or actions increases (Theorem 1), which contrasts sharply with the convergent asymptotic behaviour of the random sequence best-response dynamics (Proposition 1). Additionally, we provided new expressions for the frequency of convergence in the 2-player case (Theorem 3) and for the convergence frequency of network games (Theorem 5). We also studied the likelihood of generating games with a specific structure at random and found explicit expressions for games with certain attractor structures (Theorem 8). Using computer simulations, we compared the convergence properties of best-response dynamics to more complex learning rules

(such as reinforcement learning, fictitious play, and replicator dynamics). We showed that the frequency of convergence of these learning rules, which are defined under simultaneous updating, is most closely tracked by clockwork best-response dynamics, which have not received significant attention in the literature (Section 3.3.2). Additionally, we generalised the basic model to games with correlated payoffs, and found that games with a large number of players or a large number of actions are unlikely to converge to a Nash equilibrium (Section 3.3.1.1).

Future Work on Dynamics on Random Games For our theoretical results, we had to impose the assumption that payoffs are drawn independently. Recently, Mimun et al. (2024) show that the best-response dynamic converges asymptotically with the number of actions in 2-player games with correlated payoffs if a pure Nash equilibrium exists. Future research could extend this result to the multi-player case and to games on networks. Other potential theoretical extensions could aim at finding explicit expressions for the frequency of convergence in multi-player games (improving Theorem 2) or generalising the likelihood of certain attractor structures beyond games with a unique pure Nash equilibrium (extending Theorem 8). Finally, we did not analytically investigate why specifically the *clockwork* best-response dynamic predicts the convergence of complex learning rules so well (e.g. Figure 3.15). To do so, future research could study the convergence frequency of the complex learning rules by calculating the sizes of the basins of attraction of those learning rules on simple Boolean games (such as those presented in Figure 3.16).

Part II: Large-Scale Macroeconomic Agent-based Modelling At the macroeconomic scale, our novel agent-based model continues the recently established stream of macroeconomic ABMs capable of out-of-sample forecasting (see Section 5.3.2). Their ability to track empirical time series, rather than just reproducing summary statistics, allows agent-based models to be directly compared to conventional economic models by their forecasting performance.

Main Contributions to the ABM Literature We improved the benchmark model (Poledna et al., 2023a) in several key specifications, most notably by modelling the credit market (Section 6.5.3) and the housing market (Section 6.5.4) in greater detail, and by adding enhanced behavioural rules and market mechanisms. Importantly, our model can be calibrated to any set of countries, in particular to the 38 OECD member states, for which we described the calibration procedure (see Section 6.2.2 for an overview of the data sources used). This represents a significant improvement over the current state-of-the-art in macroeconomic agent-based models for forecasting, which have only been calibrated to a single country each (Austria

(Poledna et al., 2023a), Canada (Hommes et al., 2022), Italy (Di Domenico and Riccetti, 2023)). Additionally, unlike those ABMs, every agent in our model is initialised using agent-level microdata, which provides heterogeneity and defines behavioural rules that are not crucial to the model, thereby reducing the number of arbitrary assumptions. We demonstrated its calibration using approximate Bayesian estimation methods, which allowed us to question some of the assumptions typically made in the literature (Section 7.2). We performed forecasting exercises across 38 countries and 20 initialisation times, showing that our calibrated model outperforms both a benchmark time series model and the benchmark agent-based model (Poledna et al., 2023a) (Section 7.4). Specifically, our model provides a lower forecasting error for all considered economic aggregates in 33 out of 38 countries (Table 7.4). Additionally, we estimated a Bayes factor for each time of initialisation and country. Of all estimated Bayes factors, 18% support the IIASA model over our model; the remaining Bayes factors support our model to varying degrees. We show that both the improvement in forecasting performance of the calibrated model and the fraction of Bayes factors supporting our model are statistically significant (Section 7.6). We argue that this improvement comes from more realistic modelling assumptions, the use of microdata, and careful calibration.

For the remainder of this chapter, we discuss the limitations of our approach and highlight two streams of future work — one on the use of agent-based models for policymaking and one on the task of building a *global* macroeconomic agent-based model.[1]

Limitations of the Modelling Approach We have suggested a list of additions and improvements to the model in Section 6.6. These include regionalising the model, modelling seasonality, improving the production function of firms (Section 6.4.3.3), as well as the household consumption and investment functions (Section 6.4.5.4), and modelling the labour market more realistically (Section 6.5.2). Additionally, we are currently using S&P Compustat microdata to initialise firms and banks, which has poor coverage (see Sections 6.4.1.1 and 6.4.3.2) for non-US countries. This issue could be addressed by switching to Moody's Orbis data for firms and BankFocus data for banks.[2] Finally, we currently utilise similar modelling rules across all OECD countries, even though there are significant institutional

[1]There are many additional applications of the agent-based model. For instance, future work could use the calibrated agent-based model to study the green transition. By disaggregating certain sectors, e.g. breaking down *B: Mining and quarrying, energy-producing sectors* into petroleum extraction, coal mining, natural gas, etc., one could study the macroeconomic consequences of shifting from unsustainable practices to more sustainable ones.

[2]Which, however, are not nearly as widely available to universities and other institutions as Compustat — Orbis *Academics* costs about 60k GBP for all countries, BankFocus 17k GBP.

differences. To address those issues, we have built the model to be entirely modular, allowing for easy replacement of entire agents, behavioural rules or datasets. In this sense, we have not built a single model because our framework allows the implementation and comparison of a multitude of models. This allows institutional detail to be included for different countries by working with subject-matter experts.

Use of our Agent-based Model for Policymaking When it comes to the use of agent-based modelling for policymaking, the authors of Hommes et al. (2024) discuss the use of CANVAS (Hommes et al. (2022), Section 5.3.2.10) for policymaking at the Bank of Canada, in particular regarding inflation expectations.[3] At the current stage, the model presented here is not sufficiently developed for studying macroprudential policy. For instance, in order to implement Basel III (Basel Committee on Banking Supervision, 2010), we would have to model the risk-weighting of assets and include capital buffers (as e.g. in Bardoscia et al. (2024)). As mentioned above, Moody's BankFocus data would provide detailed balance sheets, income statements, and key regulatory ratios along with regulatory metrics such as CAR and Basel III compliance.[4] Additionally, the model is currently built as a short-term model and is meant to produce forecasts for up to 3 years. For longer time horizons that are often interesting to central banks, the model would need to be adapted to include structural change and equipped with the ability to produce business cycles and credit cycles. However, the agent-based model we have presented in Chapter 6 is discussed in a report (Muellbauer and Kaszowska-Mojsa, 2023) for *Finansinspektionen*, the financial supervisory authority of Sweden, and by Muellbauer and Aron (2024) in a review on analytical frameworks used for policy analysis for the Czech National Bank.

Towards a Global Macroeconomic Agent-based Model Instead of building an agent-based model for a specific country in great detail (like Kaszowska-Mojsa and Pipień (2020) for Poland), we focussed on calibrating a large number of countries. This represents an important step towards the long-term goal of developing a global macroeconomic agent-based model, as we provide a foundation for researchers in each country to incorporate institutional detail. The work presented in this thesis only included the 38 OECD member states, covering about 40% of global GDP in 2023. We hope to convince groups at other institutions to adopt our approach and improve future model-building efforts.

[3]The report by central bankers at the Banks of Hungary, Spain, Italy, England, and Poland (Borsos et al., 2024) summarises the research on agent-based models conducted at 24 different central banks and 7 related institutions.

[4]Our approach presents a starting point by modelling realistic regulatory requirements following data by the *European Systemic Risk Board* (Section 6.4.1.2).

Bibliography

Aktipis, C. A. (2004). Know when to walk away: contingent movement and the evolution of cooperation. *Journal of theoretical biology*, 231(2):249–260.

Aktipis, C. A. (2006). Recognition memory and the evolution of cooperation: How simple strategies succeed in an agent-based world. *Adaptive Behavior*, 14(3):239–247.

Allen, T. W. and Carroll, C. D. (2001). Individual learning about consumption. *Macroeconomic Dynamics*, 5(02):255–271.

Alon, N., Rudov, K., and Yariv, L. (2021). Dominance solvability in random games. *Available at SSRN 3850992*.

Amiet, B., Collevecchio, A., Scarsini, M., and Zhong, Z. (2021). Pure nash equilibria and best-response dynamics in random games. *Mathematics of Operations Research*. Publisher: INFORMS.

Arratia, R., Goldstein, L., Gordon, L., and others (1989). Two moments suffice for Poisson approximations: the Chen-Stein method. *The Annals of Probability*, 17(1):9–25. Publisher: Institute of Mathematical Statistics.

Arthur, W. B. (1991). Designing economic agents that act like human agents: A behavioral approach to bounded rationality. *The American economic review*, 81(2):353–359. Publisher: JSTOR.

Artinger, F. M., Gigerenzer, G., and Jacobs, P. (2022). Satisficing: Integrating two traditions. *Journal of Economic Literature*, 60(2):598–635.

Ashraf, Q., Gershman, B., and Howitt, P. (2016). How inflation affects macroeconomic performance: An agent-based computational investigation. *Macroeconomic dynamics*, 20(2):558–581. Type: Journal article.

Ashraf, Q., Gershman, B., and Howitt, P. (2017). Banks, market organization, and macroeconomic performance: an agent-based computational analysis. *Journal of Economic Behavior & Organization*, 135:143–180. Type: Journal article.

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

Austin, T. L. (1960). The enumeration of point labelled chromatic graphs and trees. *Canadian Journal of Mathematics*, 12:535–545. Publisher: Cambridge University Press.

Axtell, R., Epstein, J., and Young, H. (2000). The emergence of classes in a multi-agent bargaining model. *Generative Social Science: Studies in Agent-Based Computational Modeling*.

Axtell, R. L. (2001). Zipf distribution of us firm sizes. *science*, 293(5536):1818–1820.

Axtell, R. L. and Farmer, J. D. (2024). Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature (Forthcoming)*, pages 1–101.

Aylett-Bullock, J., Cuesta-Lazaro, C., Quera-Bofarull, A., Icaza-Lizaola, M., Sedgewick, A., Truong, H., Curran, A., Elliott, E., Caulfield, T., Fong, K., et al. (2021). June: open-source individual-based epidemiology simulation. *Royal Society open science*, 8(7):210506.

Babichenko, Y. (2013). Best-reply dynamics in large binary-choice anonymous games. *Games and Economic Behavior*, 81:130–144. Publisher: Elsevier.

Babu, S. and Mohan, U. (2018). An integrated approach to evaluating sustainability in supply chains using evolutionary game theory. *Computers & operations research*, 89:269–283. Type: Journal article.

Bachner, G., Knittel, N., Poledna, S., Hochrainer-Stigler, S., and Reiter, K. (2023). Revealing indirect risks in complex socioeconomic systems: A highly detailed multi-model analysis of flood events in austria. *Risk Analysis*.

Bankes, S. C. (2002). Agent-based modeling: A revolution? *Proceedings of the National Academy of Sciences*, 99(suppl 3):7199–7200. Type: Journal article.

Bardoscia, M., Carro, A., Hinterschweiger, M., Napoletano, M., Popoyan, L., Roventini, A., and Uluc, A. (2024). The impact of prudential regulations on the uk housing market and economy: insights from an agent-based model. Technical report, Bank of England.

Basel Committee on Banking Supervision (2010). Basel iii: International framework for liquidity risk measurement, standards and monitoring. Technical report, Bank for International Settlements Basel.

Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035.

Benes, M. J., Kumhof, M. M., and Laxton, M. D. (2014). *Financial crises in DSGE models: A prototype model*. International Monetary Fund.

Berg, J. and Weigt, M. (1999). Entropy and typical properties of Nash equilibria in two-player games. *EPL (Europhysics Letters)*, 48(2):129–135. Publisher: IOP Publishing.

Blanchard, O. (2018). On the future of macroeconomic models. *Oxford Review of Economic Policy*, 34(1-2):43–54. Type: Journal article.

Blattner, T. S. and Margaritov, E. (2010). Towards a robust monetary policy rule for the euro area. Working paper series, European Central Bank. Issue: 1210.

Blume, L. E. and others (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5(3):387–424.

Borsos, A., Carro, A., Glielmo, A., Hinterschweiger, M., Kaszowska-Mojsa, J., and Uluc, A. (2024). Agent-based modeling at central banks: Recent developments and new challenges. unpublished.

Boucher, V. (2017). Selecting equilibria using best-response dynamics. *Economics Bulletin*, 37(4):2728–2734.

Brayton, F., Mauskopf, E., Reifschneider, D., and Tinsley, P. (1997). The role of expectations in the frb/us macroeconomic model. *Fed. Res. Bull.*, 83:227.

Brown, G. W. (1951). *Iterative solutions of games by fictitious play*. Activity analysis of production and allocation. Wiley, New York. Pages: 374–376.

Bullard, J. and Duffy, J. (2001). Learning and excess volatility. *Macroeconomic Dynamics*, 5(02):272–302.

Bush, R. R. and Mosteller, F. (1953). A stochastic model with applications to learning. *The Annals of Mathematical Statistics*, 24(4):559–585.

Börgers, T. and Sarin, R. (1997). Learning through reinforcement and replicator dynamics. *Journal of Economic Theory*, 77(1):1–14.

Candogan, O., Ozdaglar, A., and Parrilo, P. A. (2013). Dynamics in near-potential games. *Games and Economic Behavior*, 82:66–90. Publisher: Elsevier.

Canova, F. and Sala, L. (2009). Back to square one: Identification issues in DSGE models. *Journal of Monetary Economics*, 56(4):431–449. Type: Journal article.

Caplat, P., Anand, M., and Bauch, C. (2008). Symmetric competition causes population oscillations in an individual-based model of forest dynamics. *Ecological Modelling*, 211(3):491–500. Type: Journal article.

Carro, A. (2023). Taming the housing roller coaster: The impact of macroprudential policy on the house price cycle. *Journal of Economic Dynamics and Control*, 156:104753.

Carro, A., Hinterschweiger, M., Uluc, A., and Farmer, J. D. (2023). Heterogeneous effects and spillovers of macroprudential policy in an agent-based model of the uk housing market. *Industrial and Corporate Change*, 32(2):386–432.

Castaldi, C. and Dosi, G. (2009). The patterns of output growth of firms and countries: Scale invariances and scale specificities. *Empirical Economics*, 37:475–495.

Cazcarro, I., Amores, A. F., Arto, I., and Kratena, K. (2022). Linking multisectoral economic models and consumption surveys for the european union. *Economic Systems Research*, 34(1):22–40.

Cederman, L.-E. (2005). Computational models of social forms: Advancing generative process theory. *American journal of sociology*, 110(4):864–893. Type: Journal article.

Challet, D., Marsili, M., and Zhang, Y.-C. (2005). *Minority games: Interacting agents in financial markets (oxford finance series)*. Oxford University Press, Inc. Type: Book.

Cheng, L., Liu, G., Huang, H., Wang, X., Chen, Y., Zhang, J., Meng, A., Yang, R., and Yu, T. (2020). Equilibrium analysis of general N-population multi-strategy games for generation-side long-term bidding: An evolutionary game perspective. *Journal of Cleaner Production*, 276:124123. Type: Journal article.

Cheung, Y.-W. and Friedman, D. (1997). Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior*, 19(1):46–76. Publisher: Elsevier.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. (2014). Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

Christiano, L., Eichenbaum, M. S., and Evans, C. (1998). Modeling money.

Christiano, L. J., Eichenbaum, M. S., and Trabandt, M. (2018). On DSGE models. *Journal of Economic Perspectives*, 32(3):113–40. Type: Journal article.

Christodoulou, G., Mirrokni, V. S., and Sidiropoulos, A. (2012). Convergence and approximation in potential games. *Theoretical Computer Science*, 438:13–27. Publisher: Elsevier.

Cincotti, S., Raberto, M., and Teglio, A. (2010). Credit money and macroeconomic instability in the agent-based model and simulator Eurace. *Economics: The Open-Access, Open-Assessment E-Journal*, 4. Type: Journal article.

Cincotti, S., Raberto, M., and Teglio, A. (2012a). The eurace macroeconomic model and simulator. In *Agent-based Dynamics, Norms, and Corporate Governance. The proceedings of the 16-th World Congress of the International Economic Association, Palgrave*, volume 2, pages 1–26.

Cincotti, S., Raberto, M., and Teglio, A. (2012b). Macroprudential policies in an agent-based artificial economy. *Revue de l'OFCE*, 1(5):205–234.

Clarida, R., Gali, J., and Gertler, M. (1999). The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature*, 37(4):1661–1707. Type: Journal article.

Coenen, G. and Straub, R. (2005). Does government spending crowd in private consumption? theory and empirical evidence for the euro area. *International finance*, 8(3):435–470.

Colander, D., Goldberg, M., Haas, A., Juselius, K., Kirman, A., Lux, T., and Sloth, B. (2009). The financial crisis and the systemic failure of the economics profession. *Critical Review*, 21(2-3):249–267. Type: Journal article.

Collevecchio, A., Mimun, H. A., Quattropani, M., and Scarsini, M. (2024). Basins of attraction in two-player random ordinal potential games. *arXiv preprint arXiv:2407.05460*.

Coucheney, P., Durand, S., Gaujal, B., and Touati, C. (2014). General revision protocols in best response algorithms for potential games. In *2014 7th international conference on NETwork games, COntrol and OPtimization (NetGCoop)*, pages 239–246. IEEE.

Dawid, H. and Gatti, D. D. (2018). Agent-based macroeconomics. In *Handbook of computational economics*, volume 4, pages 63–156. Elsevier. Type: Book section.

Dawid, H. and Harting, P. (2012). Capturing firm behavior in agent-based models of industry evolution and macroeconomic dynamics. *Applied Evolutionary Economics, Behavior and Organizations, ed. by G. Bünsdorf, Edward-Elgar*, pages 103–130. Type: Journal article.

Dawid, H., Harting, P., and Neugart, M. (2014). Economic convergence: policy implications from a heterogeneous agent model. *Journal of Economic Dynamics and Control*, 44:54–80. Type: Journal article.

Dawid, H., Harting, P., van der Hoog, S., and Neugart, M. (2016). A heterogeneous agent macroeconomic model for policy evaluation: Improving transparency and reproducibility. Technical report, Bielefeld University.

De Vany, A. (1996). The emergence and evolution of self-organized coalitions. In *Computational Economic Systems: Models, Methods & Econometrics*, pages 25–50. Springer.

Deissenberg, C., Van Der Hoog, S., and Dawid, H. (2008). EURACE: A massively parallel agent-based model of the European economy. *Applied Mathematics and Computation*, 204(2):541–552. Type: Journal article.

del Rio-Chanona, R. M., Mealy, P., Beguerisse-Díaz, M., Lafond, F., and Farmer, J. D. (2021). Occupational mobility and automation: a data-driven network model. *Journal of the Royal Society Interface*, 18(174):20200898.

del Rio-Chanona, R. M., Mealy, P., Pichler, A., Lafond, F., and Farmer, J. D. (2020). Supply and demand shocks in the covid-19 pandemic: An industry and occupation perspective. *Oxford Review of Economic Policy*, 36(Supplement_1):S94–S137.

Delli Gatti, D., Desiderio, S., Gaffeo, E., Cirillo, P., and Gallegati, M. (2011). *Macroeconomics from the bottom-up.* Springer.

Delli Gatti, D. and Grazzini, J. (2020). Rising to the challenge: Bayesian estimation and forecasting techniques for macroeconomic agent based models. *Journal of Economic Behavior & Organization*, 178:875–902.

Di Domenico, J. and Riccetti, L. (2023). The relevance for modeling market exchanges of local interaction, heterogeneity, and number of agents. *Available at SSRN 4669871*.

Dosi, G., Fagiolo, G., Napoletano, M., Roventini, A., and Treibich, T. (2015). Fiscal and monetary policies in complex evolving economies. *Journal of Economic Dynamics and Control*, 52:166–189.

Dosi, G., Fagiolo, G., and Roventini, A. (2006). An evolutionary model of endogenous business cycles. *Computational Economics*, 27(1):3–34. Type: Journal article.

Dosi, G., Fagiolo, G., and Roventini, A. (2009). The microfoundations of business cycles: an evolutionary, multi-agent model. In *Schumpeterian perspectives on innovation, competition and growth*, pages 161–180. Springer. Type: Book section.

Dosi, G., Fagiolo, G., and 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. Type: Journal article.

Dosi, G., Freeman, R. B., Pereira, M. C., Roventini, A., and Virgillito, M. E. (2021). The impact of deunionization on the growth and dispersion of productivity and pay. *Industrial and Corporate Change*, 30(2):377–408.

Dosi, G., Roventini, A., and Russo, E. (2019). Endogenous growth and global divergence in a multi-country agent-based model. *Journal of Economic Dynamics and Control*, 101:101–129. Type: Journal article.

Dresher, M. (1970). Probability of a pure equilibrium point in n-person games. *Journal of Combinatorial Theory*, 8(1):134–145. Type: Journal article.

Duong, M. H. (2016). On the expected number of equilibria in a multi-player multi-strategy evolutionary game. *Dynamic Games and Applications*, 6(3):324–346. Type: Journal article.

Durand, S., Garin, F., and Gaujal, B. (2019). Distributed best response dynamics with high playing rates in potential games. *Performance Evaluation*, 129:40–59. Publisher: Elsevier.

Durand, S. and Gaujal, B. (2016). Complexity and optimality of the best response algorithm in random potential games. In *International symposium on algorithmic game theory*, pages 40–51. Springer.

Durkan, C., Murray, I., and Papamakarios, G. (2020). On contrastive learning for likelihood-free inference. In *International conference on machine learning*, pages 2771–2781. PMLR.

Dyer, J. (2022). Black-box bayesian inference for economic agent-based models.

Dyer, J., Cannon, P., Farmer, J. D., and Schmon, S. M. (2024). Black-box Bayesian inference for agent-based models. *Journal of Economic Dynamics and Control*, 161:104827.

Dyer, J., Cannon, P. W., and Schmon, S. M. (2022). Amortised likelihood-free inference for expensive time-series simulators with signatured ratio estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 11131–11144. PMLR.

Dyer, J., Quera-Bofarull, A., Chopra, A., Farmer, J. D., Calinescu, A., and Wooldridge, M. (2023). Gradient-assisted calibration for financial agent-based models. In *Proceedings of the Fourth ACM International Conference on AI in Finance*, pages 288–296.

Eisenbach, T. and Sannikov, Y. (2013). Macroeconomics with financial frictions: A survey. *Advances in Economics and Econometrics*. Type: Journal article.

Erev, I. and Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, pages 848–881. Publisher: JSTOR.

Erlingsson, E. J., Raberto, M., Stefánsson, H., and Sturluson, J. T. (2013). Integrating the housing market into an agent-based economic model. *Managing Market Complexity: The Approach of Artificial Economics*, pages 65–76.

Erlingsson, E. J., Teglio, A., Cincotti, S., Stefansson, H., Sturluson, J. T., and Raberto, M. (2014). Housing market bubbles and business cycles in an agent-based credit economy. *Economics*, 8(1):20140008.

Fabrikant, A., Jaggard, A. D., and Schapira, M. (2013). On the structure of weakly acyclic games. *Theory of Computing Systems*, 53(1):107–122. Publisher: Springer.

Fagiolo, G., Napoletano, M., and Roventini, A. (2008). Are output growth-rate distributions fat-tailed? some evidence from oecd countries. *Journal of Applied Econometrics*, 23(5):639–669.

Fagiolo, G. and Roventini, A. (2016). Macroeconomic policy in DSGE and agent-based models redux: New developments and challenges ahead. *Available at SSRN 2763735*. Type: Journal article.

Farmer, J. D. and Axtell, R. L. (2022). Agent-Based Modeling in Economics and Finance: Past, Present, and Future. INET Oxford Working Papers 2022-10, Institute for New Economic Thinking at the Oxford Martin School, University of Oxford.

Farmer, J. D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256):685–686. Type: Journal article.

Fedriani, J. M., Wiegand, T., Ayllón, D., Palomares, F., Suárez-Esteban, A., and Grimm, V. (2018). Assisting seed dispersers to restore oldfields: An individual-based model of the interactions among badgers, foxes and Iberian pear trees. *Journal of Applied Ecology*, 55(2):600–611. Type: Journal article.

Feldman, M. and Tamir, T. (2012). Convergence of best-response dynamics in games with conflicting congestion effects. In *International workshop on internet and network economics*, pages 496–503. Springer.

Fernández-Villaverde, J., Hurtado, S., and Nuno, G. (2023). Financial frictions and the wealth distribution. *Econometrica*, 91(3):869–901.

Flood, M. and Dresher, M. (1950). The basis of the game theory: Prisoner's dilemma model of cooperation and conflict. *California: Rand Corporation*.

Foster, D. P. and Young, H. P. (1998). On the nonconvergence of fictitious play in coordination games. *Games and Economic Behavior*, 25(1):79–96.

Foster, D. P. and Young, H. P. (2006). Regret testing: Learning to play Nash equilibrium without knowing you have an opponent. *Theoretical Economics*, 1(3):341–367. Publisher: New York, NY: The Econometric Society.

Franchi, E. (2014). An agent-based modeling framework for social network simulation. In *State of the art applications of social network analysis*, pages 75–96. Springer. Type: Book section.

Friedman, D. (1996). Equilibrium in evolutionary games: Some experimental results. *Economic Journal*, pages 1–25. Publisher: JSTOR.

Friedman, J. W. and Mezzetti, C. (2001). Learning in games by random sampling. *Journal of Economic Theory*, 98(1):55–84. Publisher: Elsevier.

Friedman, M. and Schwartz, A. J. (2008). *A monetary history of the United States, 1867-1960*, volume 14. Princeton university press.

Fudenberg, D. and Levine, D. K. (1998). *The theory of learning in games*. MIT Press.

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review*, 110(8):2271–2327.

Gaffeo, E., Catalano, M., Clementi, F., Gatti, D. D., Gallegati, M., and Russo, A. (2007). Reflections on modern macroeconomics: Can we travel along a safer road? *Physica A: Statistical Mechanics and its Applications*, 382(1):89–97. Type: Journal article.

Gaffeo, E., Gallegati, M., and Palestrini, A. (2003). On the size distribution of firms: additional evidence from the g7 countries. *Physica A: Statistical Mechanics and its Applications*, 324(1-2):117–123.

Gaffeo, E., Gatti, D. D., Desiderio, S., and Gallegati, M. (2008). Adaptive microfoundations for emergent macroeconomics. *Eastern Economic Journal*, 34(4):441–463. Type: Journal article.

Galí, J. and Rabanal, P. (2004). Technology shocks and aggregate fluctuations: How well does the real business cycle model fit postwar us data? *NBER macroeconomics annual*, 19:225–288.

Galla, T. and Farmer, J. D. (2013). Complex dynamics in learning complicated games. *Proceedings of the National Academy of Sciences*, 110(4):1232–1236. Publisher: National Academy of Sciences.

Gamal, Y., Elsenbroich, C. J., Gilbert, N., Heppenstall, A., and Zia, K. (2024). A behavioural agent-based model for housing markets: Impact of financial shocks in the uk. *Journal of artificial societies and social simulation*, 27(4).

Gatti, D., Gaffeo, E., Gallegati, M., Giulioni, G., and Palestrini, A. (2008). *Emergent macroeconomics: an agent-based approach to business fluctuations*. Springer Science & Business Media.

Gatti, D. D., Desiderio, S., Gaffeo, E., Cirillo, P., and Gallegati, M. (2011). *Macroeconomics from the Bottom-up*, volume 1. Springer Science & Business Media.

Gatti, D. D., Di Guilmi, C., Gaffeo, E., Giulioni, G., Gallegati, M., and Palestrini, A. (2005). A new approach to business fluctuations: heterogeneous interacting agents, scaling laws and financial fragility. *Journal of Economic Behavior & Organization*, 56(4):489–512. Type: Journal article.

Gatti, D. D., Gaffeo, E., and Gallegati, M. (2010). Complex agent-based macroeconomics: a manifesto for a new paradigm. *Journal of Economic Interaction and Coordination*, 5(2):111–135. Type: Journal article.

Geanakoplos, J., Axtell, R., Farmer, J. D., Howitt, P., Conlee, B., Goldstein, J., Hendrey, M., Palmer, N. M., and Yang, C.-Y. (2012). Getting at systemic risk via an agent-based model of the housing market. *American Economic Review*, 102(3):53–58. Type: Journal article.

Germain, M., Gregor, K., Murray, I., and Larochelle, H. (2015). MADE: Masked autoencoder for distribution estimation. In *International conference on machine learning*, pages 881–889. PMLR.

Germano, F. and Lugosi, G. (2007). Global nash convergence of foster and young's regret testing. *Games and Economic Behavior*, 60(1):135–154. Publisher: Elsevier.

Gill, A., Lalith, M., Poledna, S., Hori, M., Fujita, K., and Ichimura, T. (2021). High-performance computing implementations of agent-based economic models for realizing 1: 1 scale simulations of large economies. *IEEE Transactions on Parallel and Distributed Systems*, 32(8):2101–2114.

Gintis, H. (2007). The dynamics of general equilibrium. *The Economic Journal*, 117(523):1280–1309. Type: Journal article.

Gleiser, I., Farmer, D., Kaszowska-Mojsa, J., and Bydlon, S. (2024). How agent-based models powered by hpc are enabling large scale economic simulations. Blog Article.

Glielmo, A., Favorito, M., Chanda, D., and Delli Gatti, D. (2023). Reinforcement learning for combining search methods in the calibration of economic abms. In *Proceedings of the Fourth ACM International Conference on AI in Finance*, pages 305–313.

Goemans, M., Mirrokni, V., and Vetta, A. (2005). Sink equilibria and convergence. In *46th annual IEEE symposium on foundations of computer science (FOCS'05)*, pages 142–151. IEEE.

Gokhale, C. S. and Traulsen, A. (2010). Evolutionary games in the multiverse. *Proceedings of the National Academy of Sciences*, 107(12):5500–5504. Publisher: National Acad Sciences.

Goldberg, K., Goldman, A., and Newman, M. (1968). The probability of an equilibrium point. *Journal of Research of the National Bureau of Standards*, 72(2):93–101. Type: Journal article.

Goldman, A. J. (1957). The probability of a saddlepoint. *American Mathematical Monthly*, 64:729. Type: Journal article.

Grazzini, J., Hommes, C. H., Poledna, S., and Zhang, Y. (2023). Understanding post-pandemic inflation dynamics with a behavioral macroeconomic model of the canadian economy. *Available at SSRN 4381235*.

Grazzini, J., Richiardi, M. G., and Tsionas, M. (2017). Bayesian estimation of agent-based models. *Journal of Economic Dynamics and Control*, 77:26–47.

Greenberg, D., Nonnenmacher, M., and Macke, J. (2019). Automatic posterior transformation for likelihood-free inference. In *International Conference on Machine Learning*, pages 2404–2414. PMLR.

Gualdi, S., Bouchaud, J.-P., Cencetti, G., Tarzia, M., and Zamponi, F. (2015a). Endogenous crisis waves: Stochastic model with synchronized collective behavior. *Physical Review Letters*, 114(8).

Gualdi, S., Tarzia, M., Zamponi, F., and Bouchaud, J.-P. (2015b). Tipping points in macroeconomic agent-based models. *Journal of Economic Dynamics and Control*, 50:29–61.

Guerrieri, V., Lorenzoni, G., Straub, L., and Werning, I. (2022). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? *American Economic Review*, 112(5):1437–1474.

Gutmann, M. U. and Hyvärinen, A. (2012). Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of machine learning research*, 13(2).

Haldane, A. G. and Turrell, A. E. (2018). An interdisciplinary model for macroeconomics. *Oxford Review of Economic Policy*, 34(1-2):219–251. Type: Journal article.

Hallegatte, S. (2008). An adaptive regional input-output model and its application to the assessment of the economic cost of katrina. *Risk Analysis*, 28(3):779–799.

Harland, K., Heppenstall, A., Smith, D., and Birkin, M. H. (2012). Creating realistic synthetic populations at varying spatial scales: A comparative critique of population synthesis techniques. *Journal of Artificial Societies and Social Simulation*, 15(1).

Harting, P. (2015). Stabilization policies and long term growth: Policy implications from an agent-based macroeconomic model. *Bielefeld Working Papers in Economics and Management*.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

Heinrich, T., Jang, Y., Mungo, L., Pangallo, M., Scott, A., Tarbush, B., and Wiese, S. (2023). Best-response dynamics, playing sequences, and convergence to equilibrium in random games. *International Journal of Game Theory*, 52(3):703–735.

Heppenstall, A. and Crooks, A. (2019). Guest editorial for spatial agent-based models: Current practices and future trends. *GeoInformatica*, 23:163–167.

Hermans, J., Begy, V., and Louppe, G. (2020). Likelihood-free mcmc with amortized approximate ratio estimators. In *International conference on machine learning*, pages 4239–4248. PMLR.

Hofbauer, J. and Sigmund, K. (1998). *Evolutionary games and population dynamics*. Cambridge university press.

Holcombe, M., Chin, S., Cincotti, S., Raberto, M., Teglio, A., Coakley, S., Deissenberg, C., van der Hoog, S., Greenough, C., Dawid, H., et al. (2013). Large-scale modelling of economic systems. *Complex Systems*, 22(2).

Holland, J. H. and Miller, J. H. (1991). Artificial adaptive agents in economic theory. *The American Economic Review*, 81(2):365–370. Type: Journal article.

Hommes, C., He, M., Poledna, S., Siqueira, M., and Zhang, Y. (2022). CANVAS: A Canadian Behavioral Agent-Based Model. Staff Working Papers 22-51, Bank of Canada.

Hommes, C., Kozicki, S., Poledna, S., and Zhang, Y. (2024). How an agent-based model can support monetary policy in a complex evolving economy. unpublished.

Hommes, C. H. and Poledna, S. (2023). Analyzing and forecasting economic crises with an agent-based model of the euro area. *Available at SSRN 4381261*.

Hosszú, Z., Borsos, A., Mérő, B., and Vágó, N. (2024). The more the merrier?-the optimal choice of scaling in economic agent-based models. *Available at SSRN 4751602*.

Huberman, B. A. and Glance, N. S. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences*, 90(16):7716–7718.

Jackson, M. O. and Zenou, Y. (2015). Games on networks. In *Handbook of game theory with economic applications*, volume 4, pages 95–163. Elsevier.

Jeffreys, H. (1998). *The theory of probability*. OuP Oxford.

Johnston, T., Savery, M., Scott, A., and Tarbush, B. (2023). Game connectivity and adaptive dynamics. *arXiv preprint arXiv:2309.10609*.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743. Type: Journal article.

Kash, I. A., an, E. J., and Halpern, J. Y. (2011). Multiagent learning in large anonymous games. *Journal of Artificial Intelligence Research*, 40:571–598.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the american statistical association*, 90(430):773–795.

Kaszowska-Mojsa, J. and Pipień, M. (2020). Macroprudential policy in a heterogeneous environment–an application of agent-based approach in systemic risk modelling. *Entropy*, 22(2):129.

Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. Macmillan. 14th edition, 1973.

Khan, A. and Thomas, J. K. (2007). Inventories and the business cycle: An equilibrium analysis of (s, s) policies. *American Economic Review*, 97(4):1165–1188.

Kimmich, C., Weyerstraß, K., Czypionka, T., Fauster, N. F., Kinner, M., Laa, E., Mateeva, L., Plank, K., Ulrici, L., Zenz, H., et al. (2024). Economic impact of labor productivity losses induced by heat stress: An agent-based macroeconomic approach. unpublished.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kirman, A. (2006). Heterogeneity in economics. *Journal of Economic Interaction and Coordination*, 1(1):89–117. Type: Journal article.

Kirman, A. (2010). The economic crisis is a crisis for economic theory. *CESifo Economic Studies*, 56(4):498–535. Type: Journal article.

Krugman, P. (2011). The profession and the crisis. *Eastern Economic Journal*, 37(3):307–312. Type: Journal article.

Kruithof, J. (1937). Telefoonverkeersrekening. *De Ingenieur*, 52:15–25.

Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica: Journal of the Econometric Society*, pages 1345–1370.

Lamperti, F., Roventini, A., and Sani, A. (2018). Agent-based model calibration using machine learning surrogates. *Journal of Economic Dynamics and Control*, 90:366–389. Type: Journal article.

Lengnick, M. (2013). Agent-based macroeconomics: A baseline model. *Journal of Economic Behavior & Organization*, 86:102–120. Type: Journal article.

Lengnick, M. and Wohltmann, H.-W. (2013). Agent-based financial markets and New Keynesian macroeconomics: a synthesis. *Journal of Economic Interaction and Coordination*, 8(1):1–32. Type: Journal article.

Lengnick, M. and Wohltmann, H.-W. (2016). Optimal monetary policy in a new Keynesian model with animal spirits and financial markets. *Journal of Economic Dynamics and Control*, 64:148–165. Type: Journal article.

Leroy, A. and Lucotte, Y. (2016). Structural and cyclical determinants of bank interest-rate pass-through in the eurozone. *Comparative Economic Studies*, 58:196–225.

Lindé, J., Smets, F., and Wouters, R. (2016). Challenges for central banks' macro models. In *Handbook of macroeconomics*, volume 2, pages 2185–2262. Elsevier. Type: Book section.

Long Jr, J. B. and Plosser, C. I. (1983). Real business cycles. *Journal of political Economy*, 91(1):39–69.

Lorenz, H.-W. (1993). *Nonlinear dynamical economics and chaotic motion*, volume 334. Springer. Type: Book.

Loshchilov, I. and Hutter, F. (2019). Decoupled weight decay regularization. In *International Conference on Learning Representations*. PMLR.

Lucas Jr, R. E. (1976). Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46. North-Holland.

Lueckmann, J.-M., Goncalves, P. J., Bassetto, G., Öcal, K., Nonnenmacher, M., and Macke, J. H. (2017). Flexible statistical inference for mechanistic models of neural dynamics. *Advances in neural information processing systems*, 30.

Macy, M. W. and Flache, A. (2002). Learning dynamics in social dilemmas. *Proceedings of the National Academy of Sciences of the United States of America*, 99:7229–7236.

Magnuszewski, P., Campo, P., Strelkovskii, N., Fresolone-Caparrós, A., Linnerooth-Bayer, J., Poledna, S., Rovenskaya, E., Pajak, M., Goujon, A., Conte, A., et al. (2023). The migration policy simulation: Engaging stakeholders in austria's migration future by linking an agent-based model with a policy exercise. Technical report, IIASA.

Mandel, A., Jaeger, C., Fürst, S., Lass, W., Lincke, D., Meissner, F., Pablo-Marti, F., and Wolf, S. (2010). Agent-based dynamics in disaggregated growth models. Documents de travail du centre d'Economie de la sorbonne, Université Panthéon-Sorbonne (Paris 1), Centre d'Economie de la Sorbonne. Issue: 10077.

Mankiw, N. G. (2006). The macroeconomist as scientist and engineer. *Journal of Economic Perspectives*, 20(4):29–46. Type: Journal article.

Mankiw, N. G. and Romer, D. (1991). *Imperfect competition and sticky prices*, volume 1. MIT Press.

Marchi, S. d. and Page, S. E. (2014). Agent-based models. *Annual Review of Political Science*, 17(1):1–20. Type: Journal article.

Martinoli, M., Moneta, A., and Pallante, G. (2022). Calibration and validation of macroeconomic simulation models by statistical causal search. *Available at SSRN 4593650*.

Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press.

McCulloch, J., Ge, J., Ward, J. A., Heppenstall, A., Polhill, J. G., and Malleson, N. (2022). Calibrating agent-based models using uncertainty quantification methods. *Journal of Artificial Societies and Social Simulation*, 25(2).

Mérő, B., Borsos, A., Hosszú, Z., Oláh, Z., and Vágó, N. (2023). A high-resolution, data-driven agent-based model of the housing market. *Journal of Economic Dynamics and Control*, 155:104738.

Miller, J. H. and Page, S. E. (2009). *Complex adaptive systems: An introduction to computational models of social life*, volume 17. Princeton university press. Type: Book.

Mimun, H. A., Quattropani, M., and Scarsini, M. (2024). Best-response dynamics in two-person random games with correlated payoffs. *Games and Economic Behavior*, 145:239–262.

Minsky, H. P. (1977). The financial instability hypothesis: An interpretation of keynes and an alternative to "standard" theory. *Challenge*, 20(1):20–27.

Mirrokni, V. S. and Skopalik, A. (2009). On the complexity of Nash dynamics and sink equilibria. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 1–10.

Moll, B. (2024). Heterogeneous agent macroeconomics: Eight lessons and a challenge. Economic Journal Lecture, Royal Economic Society. https://www.youtube.com/watch?v=H6ycBM_64tk (Consulted June 05, 2024).

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and economic behavior*, 14(1):124–143. Publisher: Elsevier.

Mongelli, I., Neuwahl, F., and Rueda-Cantuche, J. M. (2010). Integrating a household demand system in the input–output framework. methodological aspects and modelling implications. *Economic Systems Research*, 22(3):201–222.

Muellbauer, J. and Aron, J. (2024). External Monetary Policy Review of the Czech National Bank: A Review of the Analytical Framework used for Policy Analysis and Forecasting for Monetary Policy Making. INET Oxford Working Papers 2024-11, Institute for New Economic Thinking at the Oxford Martin School, University of Oxford.

Muellbauer, J. and Kaszowska-Mojsa, J. (2023). The effects of bbms. Technical report, Finansinspektionen.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica: journal of the Econometric Society*, pages 315–335.

Napoletano, M., Dosi, G., Fagiolo, G., and Roventini, A. (2012). Wage formation, investment behavior and growth regimes: An agent-based analysis. *Revue de l'OFCE*, 1(5):235–261. Type: Journal article.

Naumann-Woleske, K., Knicker, M. S., Benzaquen, M., and Bouchaud, J.-P. (2021). Exploration of the parameter space in macroeconomic agent-based models.

Nelson, C. R. (1972). The prediction performance of the frb-mit-penn model of the us economy. *The American Economic Review*, 62(5):902–917.

Network, C. (2023). Household finance and consumption survey: Results from the 2021 wave. Technical report, ECB Statistics Paper.

Nishiyama, S. and Smetters, K. (2014). Analyzing fiscal policies in a heterogeneous-agent overlapping-generations economy. In *Handbook of Computational Economics*, volume 3, pages 117–160. Elsevier.

Nowak, M. A. and May, R. M. (1992). Evolutionary games and spatial chaos. *nature*, 359(6398):826–829.

Okada, T., Namatame, A., and Sato, H. (2016). An agent-based model of smart supply chain networks. In *Intelligent and evolutionary systems*, pages 373–384. Springer. Type: Book section.

Olmez, S., Heppenstall, A., Ge, J., Elsenbroich, C., and Birks, D. (2024). Mitigating housing market shocks: an agent-based reinforcement learning approach with implications for real-time decision support. *Journal of Simulation*, 0(0):1–19.

Opper, M. and Diederich, S. (1992). Phase transition and 1/f noise in a game dynamical model. *Physical review letters*, 69(10):1616.

Orcutt, G. H. (1957). A new type of socio-economic system. *The review of economics and statistics*, 39(2):116–123. Type: Journal article.

Ozel, B., Nathanael, R. C., Raberto, M., Teglio, A., and Cincotti, S. (2019). Macroeconomic implications of mortgage loan requirements: an agent-based approach. *Journal of Economic Interaction and Coordination*, 14:7–46.

Pablo-Martí, F., García-Tabuenca, A., and Mancha, T. (2013). Amoeba: An agent-based model of entrepreneurship and business activities. In *53rd ERSA Congress, Palermo, Italy*, pages 27–31.

Page, S. E. (2008). Uncertainty, difficulty, and complexity. *Journal of Theoretical Politics*, 20(2):115–149. Type: Journal article.

Pangallo, M. (2024). Data-driven economic agent-based models. In *The economy as an evolving complex system IV*, page ? SFI Press, Santa Fe, N.M.

Pangallo, M. (2025). Synchronization of endogenous business cycles. *Journal of Economic Behavior & Organization*, 229:106827.

Pangallo, M., Heinrich, T., and Farmer, J. D. (2019). Best reply structure and equilibrium convergence in generic games. *Science advances*, 5(2):eaat1328. Type: Journal article.

Papadopoulos, G. (2019). Income inequality, consumption, credit and credit risk in a data-driven agent-based model. *Journal of Economic Dynamics and Control*, 104:39–73.

Papamakarios, G. and Murray, I. (2016). Fast ε-free inference of simulation models with bayesian conditional density estimation. *Advances in neural information processing systems*, 29.

Papamakarios, G., Pavlakou, T., and Murray, I. (2017). Masked autoregressive flow for density estimation. *Advances in neural information processing systems*, 30.

Paulin, J., Calinescu, A., and Wooldridge, M. (2019). Understanding flash crash contagion and systemic risk: A micro–macro agent-based approach. *Journal of Economic Dynamics and Control*, 100:200–229.

Pei, T. and Takahashi, S. (2019). Rationalizable strategies in random games. *Games and Economic Behavior*, 118:110–125. Type: Journal article.

Petrovic, M., Ozel, B., Teglio, A., Raberto, M., and Cincotti, S. (2017). Eurace open: An agent-based multi-country model. Working Papers 2017/09, Economics Department, Universitat Jaume I, Castellón (Spain).

Pichler, A., Pangallo, M., del Rio-Chanona, R. M., Lafond, F., and Farmer, J. D. (2020). Production networks and epidemic spreading: How to restart the uk economy? *arXiv preprint arXiv:2005.10585*.

Pichler, A., Pangallo, M., del Rio-Chanona, R. M., Lafond, F., and Farmer, J. D. (2022). Forecasting the propagation of pandemic shocks with a dynamic input-output model. *Journal of Economic Dynamics and Control*, page 104527.

Platt, D. (2020). A comparison of economic agent-based model calibration methods. *Journal of Economic Dynamics and Control*, 113:103859.

Poledna, S., Miess, M. G., Hommes, C., and Rabitsch, K. (2023a). Economic forecasting with an agent-based model. *European Economic Review*, 151:104306.

Poledna, S., Strelkovskii, N., Conte, A., Goujon, A., Linnerooth-Bayer, J., Catalano, M., and Rovenskaya, E. (2023b). Economic and labour market impacts of migration in austria: An agent-based modelling approach. *Available at SSRN 4414968*.

Poledna, S. and Thurner, S. (2016). Elimination of systemic risk in financial networks by means of a systemic risk transaction tax. *Quantitative finance*, 16(10):1599–1613.

Popoyan, L., Napoletano, M., and Roventini, A. (2017). Taming macroeconomic instability: Monetary and macro-prudential policy interactions in an agent-based model. *Journal of Economic Behavior & Organization*, 134:117–140. Type: Journal article.

Powers, I. Y. (1990). Limiting distributions of the number of pure strategy Nash equilibria in n-person games. *International Journal of Game Theory*, 19(3):277–286. Type: Journal article.

Prais, S. J. (1956). An Econometric Model of the United States, 1929-1952. *The Economic Journal*, 66(262):325–327.

Prescott, E. C. (1986). Theory ahead of business-cycle measurement. *Carnegie-Rochester Conference Series on Public Policy*, 25:11–44.

Primiceri, G. E. (2006). Why inflation rose and fell: policy-makers' beliefs and us postwar stabilization policy. *The Quarterly Journal of Economics*, 121(3):867–901.

Pritchard, J. K., Seielstad, M. T., Perez-Lezaun, A., and Feldman, M. W. (1999). Population growth of human y chromosomes: a study of y chromosome microsatellites. *Molecular biology and evolution*, 16(12):1791–1798.

Quint, T., Shubik, M., and Yan, D. (1997). Dumb bugs vs. bright noncooperative players: A comparison. In Albers, W., Güth, W., Hammerstein, P., Moldvanu, B., and van Damme, E., editors, *Understanding strategic interaction*, pages 185–197. Springer, Berlin, Heidelberg.

Raberto, M., Nathanael, R. C., Ozel, B., Teglio, A., and Cincotti, S. (2017). Credit-driven business cycles in an agent-based macro model. In *Theory and method of evolutionary political economy: A cyprus symposium*, page 182. Taylor & Francis. Type: Conference proceedings.

Raberto, M., Teglio, A., and Cincotti, S. (2008). Integrating real and financial markets in an agent-based economic model: an application to monetary policy design. *Computational Economics*, 32:147–162.

Raberto, M., Teglio, A., and Cincotti, S. (2012). Debt, deleveraging and business cycles: An agent-based perspective. *Economics*, 6(1):20120027.

Raberto, M., Teglio, A., and Cincotti, S. (2014). Fiscal consolidation and sovereign debt risk in balance-sheet recessions. *Economic Policy and the Financial Crisis*, page 163.

Reis, R. (2022). Losing the inflation anchor. *Brookings Papers on Economic Activity*, 2021(2):307–379.

Richiardi, M., Bronka, P., van de Ven, J., Kopasker, D., and Katikireddi, S. V. (2023). Simpaths: an open-source microsimulation model for life course analysis. *Available at SSRN 4808042*.

Rinott, Y. and Scarsini, M. (2000). On the number of pure strategy Nash equilibria in random games. *Games and Economic Behavior*, 33(2):274–293. Publisher: Elsevier.

Robert, C. P. et al. (2007). *The Bayesian choice: from decision-theoretic foundations to computational implementation*, volume 2. Springer.

Robinson, J. (1951). An iterative method of solving a game. *The Annals of Mathematics*, 54(2):296.

Roma, M., Skudelny, F., Benalal, N., Diaz del Hoyo, J. L., and Landau, B. (2004). To aggregate or not to aggregate? Euro area inflation forecasting. Working paper series, European Central Bank. Issue: 374.

Romer, P. (2016). The trouble with macroeconomics. unpublished.

Rubin, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *The Annals of Statistics*, pages 1151–1172.

Russo, A., Catalano, M., Gaffeo, E., Gallegati, M., and Napoletano, M. (2007). Industrial dynamics, fiscal policy and R&D: Evidence from a computational experiment. *Journal of Economic Behavior & Organization*, 64(3-4):426–447. Type: Journal article.

Sanders, J. B., Farmer, J. D., and Galla, T. (2018). The prevalence of chaotic dynamics in games with many players. *Scientific reports*, 8(1):4902. Publisher: Nature Publishing Group.

Sarin, R. and Vahid, F. (2001). Predicting how people play games: a simple dynamic model of choice. *Games and Economic Behavior*, 34(1):104–122. Publisher: Elsevier.

Schelling, T. C. (1969). Models of segregation. *The American Economic Review*, 59(2):488–493. Type: Journal article.

Seppecher, P. (2012). Flexibility of wages and macroeconomic instability in an agent-based computational model with endogenous money. *Macroeconomic Dynamics*, 16(S2):284–297. Type: Journal article.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica: journal of the Econometric Society*, pages 1–48. Type: Journal article.

Situngkir, H. (2004). Epidemiology through cellular automata: case of study avian influenza in Indonesia. *arXiv preprint nlin/0403035*. Type: Journal article.

Slanicay, M. (2014). Some notes on historical, theoretical, and empirical background of dsge models. *Review of Economic Perspectives*, 14(2):145–164.

Smets, F. and Wouters, R. (2007). Shocks and frictions in US business cycles: A bayesian DSGE approach. *American Economic Review*, 97(3):586–606.

Smith, V. L. et al. (1992). Game theory and experimental economics: beginnings and early influences. *History of Political Economy*, 24(5):241–282.

Solow, R. (2010). Building a science of economics for the real world. *House Committee on Science and Technology Subcommittee on Investigations and Oversight*, 20. Type: Journal article.

Stanford, W. (1995). A note on the probability of k pure Nash equilibria in matrix games. *Games and Economic Behavior*, 9(2):238–246. Type: Journal article.

Stanford, W. (1999). On the number of pure strategy Nash equilibria in finite common payoffs games. *Economics Letters*, 62(1):29–34. Publisher: Elsevier.

Stiglitz, J. E. (2011). Rethinking macroeconomics: What failed, and how to repair it. *Journal of the European Economic Association*, 9(4):591–645. Type: Journal article.

Stiglitz, J. E. (2018). Where modern macroeconomics went wrong. *Oxford Review of Economic Policy*, 34(1-2):70–106.

Summers, L. H. (1986). Does the stock market rationally reflect fundamental values? *The Journal of Finance*, 41(3):591–601.

Swenson, B., Murray, R., and Kar, S. (2018). On best-response dynamics in potential games. *SIAM Journal on Control and Optimization*, 56(4):2734–2767. Publisher: SIAM.

Tabak, E. G. and Vanden-Eijnden, E. (2010). Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1):217–233.

Takahashi, S. and Yamamori, T. (2002). The pure Nash equilibrium property and the quasi-acyclic condition. *Economics Bulletin*, 3(22):1–6.

Tavaré, S., Balding, D. J., Griffiths, R. C., and Donnelly, P. (1997). Inferring coalescence times from dna sequence data. *Genetics*, 145(2):505–518.

Taylor, C. and Nowak, M. A. (2006). Evolutionary game dynamics with non-uniform interaction rates. *Theoretical Population Biology*, 69(3):243–252.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, volume 39, pages 195–214. Elsevier. Type: Conference proceedings.

Teglio, A., Mazzocchetti, A., Ponta, L., Raberto, M., and Cincotti, S. (2019). Budgetary rigour with stimulus in lean times: Policy advices from an agent-based model. *Journal of Economic Behavior & Organization*, 157:59–83.

Teglio, A., Raberto, M., and Cincotti, S. (2010). Endogenous credit dynamics as source of business cycles in the EURACE model. In *Progress in artificial economics*, pages 203–214. Springer. Type: Book section.

Teglio, A., Raberto, M., and Cincotti, S. (2012). The impact of banks' cappital adequacy regulation on the economic system: An agent-based approach. *Advances in Complex Systems*, 15(supp02):1250040.

Thomas, O., Dutta, R., Corander, J., Kaski, S., and Gutmann, M. U. (2022). Likelihood-free inference by ratio estimation. *Bayesian Analysis*, 17(1):1–31.

Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., and De Vries, G. J. (2015). An illustrated user guide to the world input–output database: the case of global automotive production. *Review of International Economics*, 23(3):575–605.

Van Huyck, J., Battalio, R., Mathur, S., Van Huyck, P., and Ortmann, A. (1995). On the origin of convention: Evidence from symmetric bargaining games. *International Journal of Game Theory*, 24(2):187–212. Publisher: Springer.

Wiese, S. (2020). Convergence of learning dynamics in generic games. Master's thesis, Leipzig University.

Wiese, S., Kaszowska-Mojsa, J., Dyer, J., Moran, J., Pangallo, M., Lafond, F., Muellbauer, J., Calinescu, A., and Farmer, J. D. (2024). Forecasting macroeconomic dynamics using a calibrated data-driven agent-based model. *arXiv preprint arXiv:2409.18760*.

Wiese, S. C. and Heinrich, T. (2022). The frequency of convergent games under best-response dynamics. *Dynamic Games and Applications*, pages 1–12.

Wolf, S., Fürst, S., Mandel, A., Lass, W., Lincke, D., Pablo-Marti, F., and Jaeger, C. (2013). A multi-agent model of several economic regions. *Environmental modelling & software*, 44:25–43. Type: Journal article.

Woodford, M. (2011). *Interest and prices: Foundations of a theory of monetary policy*. princeton university press. Type: Book.

Wooldridge, M. (2009). *An introduction to multiagent systems*. John wiley & sons.

Wooldridge, M. (2012). Does game theory work? *IEEE Intelligent Systems*, 27(6):76–80.