Bounded Rationality and Macroeconomic Policy

James Moberly

University College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Economics
Trinity Term 2024

Bounded Rationality and Macroeconomic Policy

James Moberly

University College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Economics
Trinity Term 2024

Abstract

This thesis comprises three self-contained chapters which explore a variety of topics in new Keynesian macroeconomics. Chapter 1 considers an extension of Gabaix's (2020) behavioural new Keynesian model in which the degree of cognitive discounting is endogenous, so that the degree of forward-looking behaviour varies with the volatility and persistence of macroeconomic variables. I show that (i) the determinacy condition, (ii) optimal policy, and (iii) identification and estimation results change qualitatively compared to a model with exogenous discounting.

Chapter 2, co-authored with Alistair Macaulay, uses novel survey data to explore heterogeneity how households form inflation expectations. Specifically, we elicit (i) uncertainty in perceptions of current inflation, and (ii) how persistent households perceive inflation to be. Combining these with standard survey questions on inflation, we infer laws of motion for expectations at the individual level. Based on averages alone, a standard model calibrated to our data predicts inflation shocks generate small and transitory responses in expectations and consumption. The considerable heterogeneity we observe in expectation formation, however, amplifies the transmission to aggregate consumption by an order of magnitude, and substantially increases its persistence.

In Chapter 3, I study the implications of non-homothetic preferences for macroeconomic dynamics and optimal monetary policy in a new Keynesian model, in both a representative agent and two agent setting. For a broad class of preferences, I show that the textbook IS and Phillips curve hold, but for a differently weighted inflation index: the inflation rate of the marginal consumption basket. If one rewrites the model in terms of consumption-weighted inflation, it features a new variable: the relative price of the marginal consumption basket. I derive a law of motion for this relative price to give

an analytically tractable four-equation model. The first-order role of relative prices in this model (i) qualitatively changes optimal policy, (ii) biases the Phillips curve slope obtained when estimating the textbook model, and (iii) in a TANK extension, generates fluctuations in inequality.

Supervisor This thesis was supervised by Professor Martin Ellison and Professor Guido Ascari.

Statement of Authorship This thesis is based on my work at the University of Oxford between 2019 and 2023. Chapter 1 is single-authored, and an earlier version was submitted in partial fulfillment of the requirements for the degree of Master of Philosophy in Economics at the Department of Economics, University of Oxford in 2020. Chapter 2 is co-authored with Alistair Macaulay. Chapter 3 is single-authored.

Acknowledgement of Funding I gratefully acknowledge funding from the Oxford-Chellgren Graduate Scholarship at University College and the David Walton Distinguished Doctoral Scholarship.

Acknowledgements

I am very grateful to my supervisors, Martin Ellison and Guido Ascari, for their invaluable advice and encouragement throughout my time as a graduate student. They have taught me a great deal, and I have very much enjoyed working with both of them. More broadly, I have greatly enjoyed interacting with my peers and the faculty at the Department of Economics, and I am thankful for their helpful comments and suggestions on my work. Special thanks goes to my co-author Alistair Macaulay; working together on our paper, which is the second chapter of this thesis, was the most enjoyable part of my doctoral studies.

I am very grateful for the generous financial support provided by the Oxford-Chellgren Graduate Scholarhip and the David Walton Distinguished Doctoral Scholarship.

Finally, I would like to thank my parents and my sisters Emma and Charlotte for their constant support and encouragement.

James Moberly
April 2024

Contents

1 Endogenous Cognitive Discounting **12**
1.1 Motivation . . . 13
1.1.1 Related Literature . . . 16
1.1.2 Outline . . . 18
1.2 Model . . . 18
1.2.1 The Model with Exogenous Attention . . . 18
1.2.2 Optimal Choices of Attention . . . 22
1.2.3 Attention Equilibrium . . . 26
1.3 Determinacy . . . 27
1.3.1 A New Statement of the Taylor Principle . . . 27
1.3.2 Path Dependence . . . 29
1.3.3 Shock Dependence . . . 32
1.4 Stability at the ELB . . . 34
1.5 Optimal Policy . . . 37
1.5.1 The Loss Function . . . 37
1.5.2 Optimal Policy under Discretion . . . 38
1.5.3 Optimal Policy under Commitment . . . 41
1.6 Identification . . . 45
1.7 Estimation . . . 48
1.8 External Validation . . . 58
1.9 Conclusions . . . 60

Appendices to Chapter 1 **63**
1.A Proofs and Derivations . . . 63
1.A.1 Exogenous Attention . . . 63
1.A.2 Consumer Attention . . . 66
1.A.3 Effects of Attention on Macroeconomic Volatility . . . 67
1.A.4 Firm Losses under Cognitive Discounting . . . 68
1.B Numerical Solution Method . . . 70
1.B.1 Smoothed Maximum . . . 71

1.C Policy at the ELB . . . 71
1.C.1 A Regime Switching Model . . . 71
1.C.2 Average Inflation Targeting . . . 72
1.C.3 Model Solution under Regime Switching . . . 74
1.D Alternative Loss Functions . . . 78
1.E Further Exercises in Identification . . . 80
1.E.1 Identification without Interest Rate Smoothing . . . 80
1.E.2 Potential Biases from Neglecting Discounting . . . 80
1.F Estimation Methodology . . . 83
1.G Additional Empirical Results . . . 83

2 Heterogeneity in Imperfect Inflation Expectations: Theory and Evidence from a Novel Survey 88
2.1 Introduction . . . 89
2.2 Expectations Framework . . . 93
2.2.1 The Agent . . . 93
2.2.2 Expectations Impulse Responses . . . 95
2.2.3 The Role of Heterogeneity . . . 95
2.3 Data . . . 96
2.4 Empirical Results . . . 98
2.4.1 Relationships between Expectation Components . . . 100
2.4.2 Correlations with Household Characteristics . . . 101
2.5 Implications for Aggregate Consumption . . . 103
2.5.1 Consumption-Saving Model . . . 103
2.5.2 Aggregate Shock Transmission . . . 105
2.6 Conclusion . . . 107

Appendices to Chapter 2 108
2.A Proofs . . . 108
2.A.1 Baseline Model . . . 108
2.A.2 Extended Subjective Laws of Motion . . . 111
2.B Survey details . . . 113
2.B.1 Summary statistics . . . 113
2.B.2 Survey questions (English and German) . . . 114
2.B.3 Novel questions for November 2021 (in German) . . . 114
2.C Variable construction . . . 117
2.C.1 Main variables . . . 117
2.C.2 Alternative measurement for V_i^f . . . 120

2.D Additional empirical results . . . 123
2.D.1 Additional parameter distributions . . . 123
2.D.2 Relationship between components of expectation formation and point estimates . . . 124
2.D.3 Relationship between components of expectation formation: further details . . . 125
2.D.4 Correlations of expectation components with household characteristics: further details . . . 129
2.E Further impulse response exercises . . . 130

3 A Four Equation New Keynesian Model with Many Sectors and Non-Homothetic Preferences 133
3.1 Motivation . . . 134
3.1.1 Related Literature . . . 138
3.1.2 Outline . . . 139
3.2 Baseline Model . . . 140
3.2.1 Households . . . 140
3.2.2 A More General Result . . . 143
3.2.3 Final Goods Firms . . . 144
3.2.4 Intermediate Goods Firms . . . 144
3.2.5 Fiscal Policy . . . 146
3.2.6 Monetary Policy . . . 146
3.3 An Alternative IS and Phillips Curve . . . 147
3.4 A Four Equation New Keynesian Model . . . 148
3.4.1 Demand Side . . . 148
3.4.2 Supply Side . . . 149
3.4.3 Relative Price Inflation . . . 150
3.4.4 The Four Equation Model . . . 152
3.5 Implications of Relative Price Shocks in the RANK Model . . . 152
3.5.1 An Apparently Flat Phillips Curve . . . 153
3.5.2 Qualitatively Different Optimal Co-Movements . . . 157
3.6 A Two Agent Extension . . . 161
3.6.1 Inequality in the TANK Model . . . 164
3.7 Implications of Relative Price Shocks in the TANK Model . . . 165
3.7.1 Inequality with a Taylor Rule . . . 165
3.7.2 Optimal Policy in TANK . . . 169
3.8 Conclusion . . . 170

Appendices to Chapter 3 **172**
3.A Proofs and Derivations 172
3.A.1 A Broader Class of Non-Homothetic Preferences 172
3.A.2 Law of Motion for Relative Prices 176
3.A.3 Analytical Solution with a Taylor Rule 177
3.A.4 Welfare Function 181
3.A.5 Optimal Policy under Discretion 186
3.A.6 Optimal Policy under Commitment 189
3.B How Large are Shocks to the Relative Price of the Marginal Basket? . . . 193
3.B.1 Data and Calculation 193
3.B.2 The Marginal and Average Baskets 194
3.C Additional Figures 196

List of Figures

1.1 Attention equilibria 30
1.2 Attention equilibria, assuming determinacy in the previous period 31
1.3 Attention equilibria with fundamental shocks 34
1.4 Attention equilibria 36
1.5 Losses under Different $\tilde{\lambda}$: Rationality and Exogenous Attention 40
1.6 Losses under Different $\tilde{\lambda}$: Endogenous Attention 41
1.7 Attention and Response of Nominal Rates to Cost-Push Shocks under Discretion in a Calibration with Larger Shocks 41
1.8 Example of Optimal Response of Price Level to a 1pp Cost-Push Shock with $\rho_\eta = 0.8$ 42
1.9 Optimal Responses to 1ppt Cost-Push Shock under Commitment with $\rho_\eta = 0.8$ 44
1.10 Optimal Responses to 1ppt Cost-Push Shock under Commitment with $\rho_\eta = 0$ 45
1.11 Strength of Identification under Exogenous and Endogenous Attention. . 48
1.12 Counterfactual Determinacy Regions: Baseline Specification 55
1.13 Counterfactual Determinacy Regions: Low Attention Cost Specification . 55
1.14 Counterfactual Inflation Volatility 56
1.15 Counterfactual Output Gap Volatility 56
1.16 Counterfactual Central Bank Losses 56
1.17 Implied Lower Bound on β 60
1.18 Optimal choice of $m_{f,E}$ 75
1.19 Optimal choice of $m_{c,E}$ 75
1.20 Attention Equilibria. Note that points marked in red are equilibria in $m_{c,E}$, and blue denotes equilibria in $m_{f,E}$. Intersections of the blue and red curves give attention equilibria. The black line denotes the boundary of the determinacy region. 75

2.1 CDFs of key parameters. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 99

2.2 Implied IRFs of one-period ahead inflation expectations and consumption. Source: Bundesbank-Online-Panel-Households, November 2021 wave. . . 106
2.3 CDF of $St.Dev_i(\pi_{t+1})$ under the narrower and broader alternative calculation measures for V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 121
2.4 IRF of $\tilde{E}_t\pi_{t+1}$ under the narrower and broader alternative calculation measures for V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 122
2.5 IRF of $\hat{c}_t$ under the narrower and broader alternative calculation measures for V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 122
2.6 CDF of raw responses to question 1, for both the group shown a $\pm 1\%$ range and the group show a $\pm 2\%$ range, and the implied CDFs of V_i^p and V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 123
2.7 CDFs of upper and lower bounds for inferred parameters. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 124
2.8 CDF of perceived persistence, only respondents whose response to initial inflation expectations question does not end in .0 or .5. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 124
2.9 Implied IRF of aggregate expectations and aggregate consumption. The homogeneity and heterogeneity cases are as described in Section 2.5.2. The remaining cases set χ_i and $\tilde{\rho}_i$ respectively to their average values for all households. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 130
2.10 Implied IRFs of one-period ahead inflation expectations and consumption by shock. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 131
2.11 Implied $\hat{c}_0$ under different assumptions on $\tilde{\phi}_i$. In panel (a) all households form interest rate expectations using the same perceived Taylor Rule parameter $\bar{\phi}$. In panel (b) this parameter varies across households according to $\tilde{\phi}_i = \bar{\phi} + \nu(\tilde{\rho}_i - E[\tilde{\rho}_i])$. Source: Bundesbank-Online-Panel-Households, November 2021 wave. 132

3.1 Budget Shares for Selected Consumption Categories by Equivalized Disposable Income Decile, Based on 2019 UK Data 134
3.2 Impulse responses to a one percentage point relative price shock ($\tilde{\varepsilon}_t^M$) in the RANK model with $\rho = 0$. 156

3.3 Impulse responses to a one percentage point relative price shock ($\tilde{\varepsilon}_t^M$) in the RANK model with $\rho = 0.6$. 156
3.4 Estimated Phillips Curve slope if the econometrician neglects the first-order role of relative price shocks. The x-axis varies the standard deviation of the innovation in the relative price shock process, $\tilde{\varepsilon}_t^M$. The different coloured dots vary the persistence of the relative price shock ρ. 158
3.5 Impulse responses to a relative price shock in the RANK and TANK models with $\rho = 0$ under the baseline calibration. 166
3.6 Impulse responses to a cost-push shock accompanied by positive or negative relative price shock in the TANK model with $\rho = 0$ under the baseline calibration. Formally, I assume $\eta_0 = 1$ in both cases. The blue line shows the responses if $\tilde{\eta}_t^M = 2$, while the red line shows the responses if $\tilde{\eta}_t^M = -2$. 168
3.7 Optimal policy responses to a relative price shock under discretion in the $\rho = 0$ case, RANK vs. TANK model. 170
3.8 Composition of Average and Marginal Baskets, Based on 2019 UK Data . 194
3.9 Differences Between Inflation On Average and Marginal Basket (Year-on-Year), and Contributions from Different COICOP Categories 195
3.10 Impulse responses to a relative price shock in the RANK model with $\rho = 0$ for a variety of calibrations. 197
3.11 Impulse responses to a relative price shock in the RANK model with $\rho = 0.6$ for a variety of calibrations. 198
3.12 Impulse responses to a relative price shock in the RANK model with $\rho = 0$ for a variety of calibrations, with policy set optimally under discretion. . 199
3.13 Impulse responses to a relative price shock in the RANK model with $\rho = 0.6$ for a variety of calibrations, with policy set optimally under discretion. 200
3.14 Impulse responses to a relative price shock in the RANK model with $\rho = 0$ for a variety of calibrations, with policy set optimally under commitment. 201
3.15 Impulse responses to a relative price shock in the RANK model with $\rho = 0.6$ for a variety of calibrations, with policy set optimally under commitment. 202

Chapter 1

Endogenous Cognitive Discounting

Abstract Gabaix (2020) shows that cognitive discounting resolves a number of puzzles in the new Keynesian model and qualitatively alters the model's policy implications. I analyze a model in which cognitive discounting is endogenous, so that the degree of forward-looking behaviour varies with the volatility and persistence of macroeconomic variables. I show that (i) the determinacy condition, (ii) optimal policy, and (iii) identification and estimation results change qualitatively compared to a model with exogenous discounting.

1.1 Motivation

In an influential paper, Gabaix (2020) analyzes a new Keynesian model in which agents pay limited attention to future inflation, output, and interest rates when making decisions today. Gabaix shows that this simple modification, known as "cognitive discounting", alters the new Keynesian model's theoretical predictions in seven important ways. Four of these changes bring the model's predictions closer to the data: Del Negro et al.'s (2012) forward guidance puzzle is resolved; the model predicts stability at the effective lower bound (ELB); the costs of the ELB are reduced; and neo-Fisherian paradoxes disappear. The remaining three changes considerably affect the model's policy implications: the Taylor principle is relaxed; price level targeting is no longer optimal under commitment; and fiscal policy is more powerful.

Gabaix shows how one could calculate the optimal choice of attention to the future for a given set of macroeconomic dynamics, using the "sparsity" microfoundation developed in earlier papers.[1] When analyzing the model's implications, however, he assumes that the degree of cognitive discounting is an exogenous constant.

Studying the implications of a model with constant attention makes sense in the context of a study that already makes an enormous contribution to the new Keynesian literature, and focuses on deriving analytical results. Nonetheless, existing empirical work indicates that attention to macroeconomic variables is regime-dependent.[2] This suggests that it is important to take into account the endogeneity of deviations from rationality when analyzing optimal policy or determinacy, since such analysis necessarily involves considering large changes in the policy regime and hence the macroeconomic environment.[3] This raises the question: do the implications of Gabaix's (2020) behavioural new Keynesian model change when cognitive discounting is endogenous?

In this paper, I analyze the implications of a model with endogenous cognitive dis-

[1] See the appendices to Gabaix (2020) for details on these calculations. The sparsity microfoundation is developed in Gabaix (2014) and Gabaix (2016).

[2] For example, Coibion and Gorodnichenko (2015a) show that the degree of informational rigidities in professional forecasts displays a negative co-movement with macroeconomic volatility. This finding matters because Gabaix (2020) shows that his cognitive discounting parameter places a lower bound on Coibion and Gorodnichenko's (2015a) informational rigidity parameter. There is also evidence that attention to macroeconomic variables is regime-dependent for consumers. For example, Cavallo et al. (2017) find that consumers have stronger priors about the inflation rate in high inflation environments.

[3] As Gabaix (2020) notes, the sparsity microfoundation implies that attention remains at anchored at a minimum, "default" level until a threshold level of macroeconomic volatility is reached, and only then starts to rise with volatility. This might be taken to indicate that at low levels of volatility, small changes in policy do not affect attention, leaving the exogenous attention model robust to the Lucas critique. However, my analysis of the determinacy condition shows that even if fundamental macroeconomic shocks are small, there exist policy rules for which the threshold level of volatility is exceeded. Moreover, given that we are currently in a period of high macroeconomic volatility, uncovering the policy implications of a model in which attention is above the default level is relevant to the current policy debate.

counting, and show that the answer is definitively "yes". I demonstrate that the determinacy condition, optimal monetary policy, and identification and estimation results all change qualitatively compared to both the model with exogenous discounting and the rational expectations benchmark.

To construct the model, I first extend Gabaix's (2020) derivation for the optimal choice of attention given a set of macroeconomic dynamics to a more general class of models, allowing me to accommodate multiple shocks, lagged terms in the model equations, and indeterminacy. Then, whilst Gabaix's derivations show the optimal choice of attention for a particular set of macroeconomic dynamics, I acknowledge that the choice of attention in turn influences those dynamics. I define an "attention equilibrium" as a set of attention choices and macroeconomic dynamics that are mutually consistent. Defining, solving for, and studying the properties of this equilibrium in the context of the new Keynesian model goes beyond Gabaix's (2020) analysis.[4]

I then study the theoretical predictions of the model with endogenous cognitive discounting. This is a novel contribution relative to Gabaix (2020), who derives the seven theoretical predictions listed above under the assumption of constant attention. In particular, I show that four of the seven main predictions of the model are significantly altered when attention is endogenous.

- **The Taylor Principle.** Gabaix shows that the determinacy condition is relaxed in the fixed attention model; with a moderate degree of cognitive discounting, even an interest rate peg results in determinacy. I show that the determinacy condition changes qualitatively under endogenous attention; if the rational expectations Taylor principle is violated, then there always exists an attention equilibrium that allows for multiple stable solutions. There may also exist another attention equilibrium that gives a unique stable solution, depending on the size and persistence of the fundamental shocks. I also study conditions under which the economy can "jump" from one equilibrium to the other.

- **Stability at the ELB.** In contrast to Gabaix's findings in the fixed attention model, I show that while economic stability at the ELB is possible in the endogenous attention model, it is not assured. This follows from the finding that under an interest rate peg, there always exists an attention equilibrium that results in indeterminacy.

- **Costs of the ELB.** In the exogenous attention model, the costs of the ELB are

[4]While attention parameters are determined in equilibrium some examples in other settings in Gabaix's earlier papers, this equilibrium is not defined for the behavioural new Keynesian model in Gabaix (2020).

significantly reduced. I show that while this may be the case under endogenous attention, there often also exists an attention equilibrium in which attention is high and the ELB is highly costly. This equilibrium exists because high attention results in a large recession at the ELB, which in turn justifies a high level of attention. I then show how average inflation targeting can rule out the possibility of costly ELB episodes, making this policy potentially much more powerful under endogenous than exogenous discounting.

- **Optimal Policy.** I show that optimal policy changes qualitatively in the model with endogenous attention relative to the exogenous attention model.
  - Under **discretion**, it is well known that welfare losses can be reduced by appointing a "conservative" central banker who has a stronger preference for inflation stabilization than society does; this mitigates the "stabilization bias" that stems from taking the expectation term in the Phillips curve as exogenous. Under exogenous attention, the optimal choice of central banker is *less* conservative than under rationality; if expectations are discounted, stabilization bias is reduced and the benefit of appointing a conservative central banker is therefore lower. Under endogenous attention, by contrast, I find that the optimal choice of central banker may be *more* conservative than under rationality. This is because appointing a more conservative central banker can endogenously reduce the discount factor in the Phillips curve and consequently improve the inflation-output tradeoff.
  - Under **commitment**, Gabaix shows that under exogenous discounting the central banker optimally returns the price level at a *higher* level after a cost-push shock. Under endogenous discounting, I show that when the shocks are transitory, the central banker may instead choose to return the price level to a *lower* level. This finding results from a novel tradeoff facing the central bank. On the one hand, higher attention amplifies the effect of persistent shocks. That said, it also increases the ability of the central bank to affect inflation today by committing to a future policy tightening. If shocks are transitory, the latter effect dominates and it is beneficial for the central bank to induce higher attention by committing to reduce inflation further below target in the future. Put differently, agents pay little attention to a commitment to a small future policy tightening, and so such commitments have little effect on expectations. Agents pay more attention to a commitment to a larger future policy tightening, and so such a commitment is more powerful.

I then turn to an empirical assessment of the model. As a precursor to estimating the

model, I first check the strength of identification under endogenous cognitive discounting. I show that the exogenous discounting model suffers from very weak identification when the autocorrelations of the shocks are similar. My results suggest that endogenizing the cognitive discount factors mitigates this issue by pinning down the cognitive discount factors in terms of the remaining model parameters. That said, this identification strategy relies on having an estimate of attention costs. I argue that one can validate the choice of attention cost parameters by comparing the implied informational rigidities with those estimated using expectations microdata.

Finally, I take the endogenous attention model to the data. I estimate the model using US data for the Great Inflation and Great Moderation periods. Unlike previous analyses that used the exogenous discounting model (Ilabaca et al., 2020), I find strong evidence against indeterminacy as a cause of the Great Inflation.[5] The endogenous discounting model also allows me to conduct counterfactuals robust to Lucas Jr's (1976) critque. I show that because the shocks hitting the economy during the Great Inflation were large, a stronger policy response to inflation would have substantially reduced inflation volatility. These findings demonstrate the importance of endogenising the discount factor when estimating the model, and also contribute to the "good luck" versus "good policy" debate regarding the Great Inflation.[6]

1.1.1 Related Literature

My paper relates to a growing literature that applies cognitive discounting to a range of policy questions such as the effects of average inflation targeting (Budianto et al., 2023), and in a variety of theoretical settings such as medium-scale models (Meggiorini, 2021) or heterogeneous-agent models (Pfäuti and Seyrich, 2022). There are also a number of papers estimating behavioural new Keynesian models, such as (Ilabaca et al., 2020). The prevailing approach in this literature is to follow Gabaix (2020) in assuming that the degree of cognitive discounting is an exogenous constant. My contribution to this literature is to highlight the importance of endogenizing the discount factor, which I show leads to novel policy conclusions. Future research could consider medium-scale or heterogeneous agent models with an endogenous discount factor.

I also contribute to the wider literature incorporating bounded rationality into DSGE models. Influential approaches include the sticky information approach studied by Mankiw

[5] Another estimation of a model with exogenous discounting model is conducted by Meggiorini (2021). In that paper, however, the consumer and firm discount factors are assumed to be equal. One important conclusion from the empirical sections of this paper is that firm and consumer incentives to pay attention may differ very greatly.

[6] See Clarida et al. (2000) or Lubik and Schorfheide (2004) for the "good policy" view, and Fernández-Villaverde et al. (2010) or Justiniano and Primiceri (2008) for the "good luck" hypothesis.

and Reis (2002), and microfounded in Reis (2006a) and Reis (2006b), learning (see Evans and Honkapohja (2012)), heuristic switching (see Hommes (2021)), and level-k thinking (see Farhi and Werning (2019)). Many other examples are surveyed by Woodford (2013). Perhaps the most closely related approach is that of Woodford (2019), who assumes that agents have finite planning horizons, and derives an IS and Phillips curve of the same form as Gabaix. Woodford's approach provides an alternative microfoundation to derive a model with fixed discount factors.

Also closely related is the rational inattention approach proposed by Sims (2003) and developed in a DSGE setting by Mackowiak and Wiederholt (2009) and Maćkowiak and Wiederholt (2015).[7] This approach also features agents trading off a cost of being more rational against the losses associated with deviating from rationality.

Nonetheless, my model differs considerably from rational inattention models. When forecasting a macroeconomic variable, an agent needs to know (i) the current value of that variable, (ii) its persistence, and (iii) news about future shocks. Whilst rational inattention models typically focus on attention to the current value of the variable, my model captures endogenous changes in perceived persistence and attention to news about future shocks. This leads to different policy implications. For example, the novel trade-offs that emerge under commitment and the changes in the determinacy condition stem specifically from the fact that agents must choose how much attention to pay to future rather than current realizations of the endogenous variables.

My model is distinguished in a similar fashion Châu's (2019) recent study, which uses Gabaix's sparsity based framework to endogenize attention to contemporaneous macroeconomic variables to generate time variation in the of the Phillips curve. This form of inattention differs from cognitive discounting; Chau assumes that agents discount macroeconomic variables at all horizons equally. In my model, it is the dynamic nature of discounting that leads to the key conclusions on determinacy and optimal policy under commitment.

In independent work, Lubik and Marzo (2021) investigate fiscal policy in a model with cognitive discounting. They include a brief section in which they calculate the optimal choice of attention for a particular calibration. However, they only calculate this optimal choice of attention for one calibration, whereas this paper investigates how the equilibrium level of attention varies with changes in policy rules, shock processes, and so on. They do not look at different monetary policy rules under endogenous attention.

This paper also relates to a broader literature studying different means of microfounding discounting in the Euler equation. Examples include McKay et al.'s (2016) model with uninsurable income risk, and Bilbiie's (2021) model which features discounting when

[7]See Maćkowiak et al. (2021) for a review of the rational inattention literature.

inequality is procyclical. My results show that how the discounting is microfounded matters; as such, adjudicating between different sources of discounting remains a critical objective for future research.

1.1.2 Outline

Section 1.2 outlines the model. In section 1.3, I prove a number of novel results regarding the determinacy condition. Section 1.4 extends this analysis to study the stability of the economy at the ELB. Section 1.5 demonstrates how optimal policy under discretion and commitment changes when discounting is endogenous. I then turn to an empirical assessment of the model. Section 1.6 assesses the strength of identification under exogenous and endogenous discounting. Section 1.7 estimates the model on the Great Inflation and Great Moderation periods, and illustrates how endogenizing attention leads to a new interpretation of those periods. In section 1.8, I then present a method for externally validating the calibration of attention costs used in the estimation using expectations data. Section 1.9 concludes.

1.2 Model

I now derive the model with endogenous cognitive discounting. I begin by reviewing Gabaix's (2020) model with exogenous attention in 1.2.1. In 1.2.2, I take Gabaix's derivations to solve for the choice of attention for a given set of macroeconomic dynamics and generalize to a broader class of models. This allows me to incorporate multiple shocks, indeterminacy and lagged terms in the model equations. In 1.2.3, I define attention equilibrium: a set of attention choices and macroeconomic dynamics that are mutually consistent.

1.2.1 The Model with Exogenous Attention

Cognitive Discounting

Gabaix (2020) assumes that households and firms do not pay full attention to future fluctuations of macroeconomic variables around their steady states when making decisions today. Formally, suppose that $\boldsymbol{s_t}$ is a vector of de-meaned endogenous variables. An agent's subjective expectation $\tilde{E}_t$ is a discounted version of the rational expectation E_t:

$$\tilde{E}_t \boldsymbol{s_{t+h}} = m^h E_t \boldsymbol{s_{t+h}} \tag{1.1}$$

Here, $m \in [0, 1]$ is a cognitive discount factor. Where $m = 1$, agents pay full attention to the future and so have rational expectations. The lower is m, the more myopic agents are. I denote the cognitive discount factors of individual firms and consumers m_f and m_c respectively.

Gabaix's Behavioural New Keynesian Model

Under this assumption, Gabaix shows that one can derive the following Phillips and IS curves:

$$\pi_t = \beta M_f E_t \pi_{t+1} + \kappa y_t + \eta_t \tag{1.2}$$

$$y_t = M_c E_t y_{t+1} - \sigma(i_t - E_t \pi_{t+1}) - \sigma z_t \tag{1.3}$$

These equations are standard, except for the aggregate cognitive discount factors M_f and M_c.[8] These parameters are functions of the individual cognitive discount factors:[9]

$$M_f = m_f \left(\theta + (1 - \theta) \frac{1 - \beta\theta}{1 - \beta m_f \theta} \right) \tag{1.4}$$

$$M_c = m_c \tag{1.5}$$

Here, θ is the Calvo pricing parameter. For the derivation of these parameters, the reader is referred to Gabaix (2020).

The Phillips and IS curve slopes are standard:[10]

$$\kappa = (\gamma + \phi) \frac{(1 - \beta\theta)(1 - \theta)}{\theta}$$

$$\sigma = \gamma^{-1}$$

[8] Note that the inflation expectation term in the IS curve is not discounted in Gabaix's baseline model, so that agents do not suffer from nominal illusion. Gabaix does discuss the possibility, however, that consumers could also discount this term. Future work could consider this extension.

[9] To derive the aggregate discount factor in the Phillips curve, one solves for the firm's optimal choice of relative price in terms of the macroeconomic variables, and discounts the expectations of those variables using the individual firm's discount factor m_f. One then solves for the implied Phillips curve, which features the discount factor M_f in front of the expectation term. This is not equivalent to simply adding the individual firm's discount factor in front of the expectation terms in the log-linearized new Keynesian Phillips curve, which is why the aggregate cognitive discount factor in the Phillips curve M_f does not equal the individual firms' cognitive discount factor m_f. A similar process is followed to derive the aggregate discount factor in the IS curve, although in this case the individual and aggregate discount factors do coincide.

[10] I work with a log-linearized model, rather than the linearized model used by Gabaix. This makes a very fractional difference to the IS curve slope, which is γ^{-1} is my model, rather than $\beta\gamma^{-1}$ in Gabaix's paper. I also follow Gabaix in using a constant returns to scale production function, resulting in this Phillips curve slope formulation.

Here, γ is the coefficient of relative risk aversion, and ϕ is the inverse Frisch elasticity.

The cost-push shock η_t and demand shock z_t follow independent exogenous AR(1) processes:

$$\eta_t = \rho_\eta \eta_{t-1} + \varepsilon_t^\eta \tag{1.6}$$
$$z_t = \rho_z z_{t-1} + \varepsilon_t^z \tag{1.7}$$
$$\varepsilon_t^\eta \sim^{i.i.d.} N(0, \sigma^2_{\varepsilon,\eta}) \tag{1.8}$$
$$\varepsilon_t^z \sim^{i.i.d.} N(0, \sigma^2_{\varepsilon,z}) \tag{1.9}$$

Differently from Gabaix, I motivate the demand shock z_t as a discount factor shock.

$$\beta_t = \beta \exp z_t$$

I use discount rather than technology shocks because Justiniano and Primiceri (2008) find discount factor shocks account for a large fraction of consumption variance.[11] Note that I write the model in terms of the log deviation of output from its steady-state, denoted y_t. Absent technology shocks, y_t also equals the output gap.

The model is then closed using an interest rate rule. For the rest of this section, and for the section on determinacy, I use a simple Taylor rule of the form:

$$\begin{aligned} i_t &= \phi_\pi \pi_t + \phi_y y_t + v_t \\ v_t &= \rho_v v_{t-1} + \varepsilon_t^v \\ \varepsilon_t^v &\sim^{i.i.d.} N(0, \sigma^2_{\varepsilon,v}) \end{aligned} \tag{1.10}$$

In the sections on identification and estimation, I instead use a Taylor rule with interest rate smoothing, as is standard in estimated new Keynesian models:

$$i_t = \rho_r i_{t-1} + (1-\rho_r)(\phi_\pi \pi_t + \phi_y y_t) + \varepsilon_t^v \tag{1.11}$$

The solution method outlined below accommodates lagged terms in the model equations, and so can handle either form of the policy rule. The determination of interest rates in the optimal policy exercises is discussed in section 1.5.

[11] This assumption is not entirely innocuous from the perspective of consumer attention. As I argue below, the fact that demand shocks stem from the discount factor affects consumers' attention choices.

It will prove useful to define the following parameter vectors:

$$\begin{aligned}
\boldsymbol{M} &= \begin{bmatrix} M_f & M_c \end{bmatrix} \\
\boldsymbol{m} &= \begin{bmatrix} m_f & m_c \end{bmatrix} \\
\boldsymbol{\chi} &= \begin{bmatrix} \beta & \theta & \gamma & \phi_\pi & \phi_y & \rho_\eta & \rho_z & \rho_v & \sigma_{\varepsilon,\eta} & \sigma_{\varepsilon,z} & \sigma_{\varepsilon,v} \end{bmatrix}
\end{aligned}$$

For the model closed by the interest rate rule in (1.10), $\boldsymbol{\chi}$ summarizes the model parameters other than the cognitive discounting parameters.

The Determinacy Condition

Assuming interest rates follow (1.10), Gabaix shows that the determinacy condition is relaxed under cognitive discounting:

$$\phi_\pi + \frac{1-\beta M_f}{\kappa}\phi_y + \frac{(1-\beta M_f)(1-M_c)}{\kappa\sigma} > 1 \tag{1.12}$$

Gabaix notes that the system is determinate even with a constant interest rate if M_f and M_c are sufficiently low.

Note that the left-hand side of this equation is a function $\boldsymbol{M}$ and $\boldsymbol{\chi}$. Given that $\boldsymbol{M}$ is itself a function of $\boldsymbol{m}$ and $\boldsymbol{\chi}$ by (1.4) and (1.5), one can then write the left-hand side of (1.12) as a function of a vector of $\boldsymbol{m}$ and $\boldsymbol{\chi}$.

$$h(\boldsymbol{m}, \boldsymbol{\chi}) = \phi_\pi + \frac{1-\beta M_f}{\kappa}\phi_y + \frac{(1-\beta M_f)(1-M_c)}{\kappa\sigma} \tag{1.13}$$

I define the set of possible attention levels as the set $\mathcal{A}$. I then define $\mathcal{D}$ and $\mathcal{I}$ as the subsets of $\mathcal{A}$ that lead to determinacy and indeterminacy respectively.

$$\begin{aligned}
\mathcal{A} &= [m_{f,d}, 1] \times [m_{c,d}, 1] \\
\mathcal{D} &= \{\boldsymbol{m} : \boldsymbol{m} \in \mathcal{A} \wedge h(\boldsymbol{m}, \boldsymbol{\chi}) > 1\} \\
\mathcal{I} &= \{\boldsymbol{m} : \boldsymbol{m} \in \mathcal{A} \wedge h(\boldsymbol{m}, \boldsymbol{\chi}) \leq 1\}
\end{aligned}$$

The terms $m_{f,d}$ and $m_{c,d}$ refer to "default", minimum levels of attention, and are explained below.

Model Solution

I begin by solving for the model's dynamics, taking cognitive discounting parameters as given. Following Lubik and Schorfheide (2004), I write the system as:

$$\begin{aligned} \boldsymbol{\Gamma_0 s_t} &= \boldsymbol{\Gamma_1 s_{t-1}} + \boldsymbol{\Psi \varepsilon_t} + \boldsymbol{\Pi \omega_t} \\ \boldsymbol{s_t} &\coloneqq \begin{bmatrix} \pi_t & y_t & i_t & E_t\pi_{t+1} & E_t y_{t+1} & \eta_t & z_t & v_t \end{bmatrix}' \\ \boldsymbol{\varepsilon_t} &\coloneqq \begin{bmatrix} \varepsilon_t^{\eta} & \varepsilon_t^{z} & \varepsilon_t^{v} \end{bmatrix}' \\ \boldsymbol{\omega_t} &\coloneqq \begin{bmatrix} \pi_t - E_{t-1}\pi_t & y_t - E_{t-1} y_t \end{bmatrix}' \end{aligned} \tag{1.14}$$

Note that $\boldsymbol{\Gamma_0}$ depends on $\boldsymbol{m}$ as well as $\boldsymbol{\chi}$. Following Lubik and Schorfheide's (2004) methodology, one can find the set of stable solutions to the model, which can be written as follows:

$$\boldsymbol{s_t} = \boldsymbol{Z_1 s_{t-1}} + \boldsymbol{Z_2 \varepsilon_t} + \boldsymbol{Z_3} \zeta_t$$

Here, the eigenvalues of $\boldsymbol{Z_1}$ all lie within the unit circle.[12] If $\boldsymbol{m} \in \mathcal{D}$, then $\boldsymbol{Z_2}$ is uniquely determined by $\boldsymbol{\chi}$ and $\boldsymbol{m}$. If $\boldsymbol{m} \in \mathcal{I}$, conversely, then $\boldsymbol{Z_2}$ is not uniquely determined. In order to calculate an equilibrium level of attention, I must specify a particular solution for the response to the fundamental shocks in the indeterminate region. I use Lubik and Schorfheide's (2004) "continuity solution", which ensures that the response to the fundamental shocks is continuous at the boundary of the indeterminacy region.[13] If $\boldsymbol{m} \in \mathcal{D}$, then $\boldsymbol{Z_3}$ is a vector of zeros, because the sunspot shock does not enter the solution, and $\boldsymbol{Z_2}$ is uniquely determined by $\boldsymbol{\chi}$ and $\boldsymbol{m}$.

With this equation in hand, one can calculate the variance-covariance matrix of $\boldsymbol{s_t}$, which I denote $\boldsymbol{\Sigma_s}$.

1.2.2 Optimal Choices of Attention

Summary of Sparsity-Based Approach

I begin by summarizing Gabaix's sparse dynamic programming methodology, which is used throughout. For further details, see Gabaix (2014) and Gabaix (2016). Gabaix supposes that an agent faces a constrained optimisation problem in which they choose

[12] The read might wonder why this is the case, given that the eigenvalues of $\boldsymbol{\Gamma_0^{-1}\Gamma_1}$ lie outside the unit circle. The reason is that we impose the requirement that solutions are stable. As such, the expectation errors are such that the unstable components of the system are always set to zero, and so are not relevant for the dynamics of the system.

[13] There is a slight complication in that to get to the edge of the determinacy region one could adjust m_f or m_c instead of ϕ_π, but to be consistent with Lubik and Schorfheide (2004) I use ϕ_π.

an action a, and that the variable b is relevant to the agent's decision. Denote by $\widehat{b}$ the deviation of b from its mean. Gabaix supposes that it is effortful for agents to pay attention to the value of $\widehat{b}$. By default, the agent pays some level of attention $m_d \in [0, 1]$. They then perceive $m_d\widehat{b}$ and make their decision over a as if that were the true value of $\widehat{b}$. They can increase their level of attention, but only at the expense of a cognition cost $\mathcal{K}w(m - m_d)$. Paying more attention allows the agent to take an action closer to the optimal action that the agent would take under full attention ($m = 1$). Gabaix suggests agents approximate average utility losses from not paying attention using a second-order expansion of the value function V around the default levelof attention.[14] This yields the approximated loss L:[15]

$$L \approx -\frac{1}{2}(1 - m)^2\Lambda$$
$$\Lambda \coloneqq V_{aa}E[(\partial a/\partial m)^2]$$

Note that both terms in Λ are evaluated at m_d.[16] The agents' optimisation problem with respect to attention is then as follows.[17]

$$\max_m -\frac{1}{2}(1 - m)^2\Lambda - \mathcal{K}w(m - m_d)$$

The optimal choice of attention (denoted m^{opt}) is then:[18]

$$m^{opt} \coloneqq \arg\max_m -\frac{1}{2}(1 - m)^2\lambda - w(m - m_d)$$
$$\lambda \coloneqq \frac{V_{aa}E[(\partial a/\partial m)^2]}{\mathcal{K}}$$

Throughout this paper, I assume a linear form for the attention cost.

$$w(m - m_d) = |m - m_d|$$

[14] See Gabaix (2020) Appendix VII B for an explanation of the value function.

[15] This is the first term in equation (8) of Gabaix (2016).

[16] Note that this is an approximation, as these terms should strictly be evaluated at the chosen level of m. Gabaix assumes that agents simplify the problem with a further Taylor expanison, and so evaluate these at m_d. See the online appendix to Gabaix (2020) for further details. This means that the default level of attention affects the optimal choice of attention even when attention is above its default level.

[17] This is equation (70) in Gabaix (2020).

[18] See equation (72) in Gabaix (2020).

Gabaix shows that the optimal choice of attention is then given as follows.[19]

$$m^{opt} = \max\left(1 - \frac{1}{\lambda}, m_d\right)$$

As Gabaix notes, a linear cost means attention is continuous in λ,[20] which helps to establish the existence of equilibrium. Additionally, it induces "sparsity", in the sense that attention is fixed at m_d for low levels of λ, and then after a certain threshold begins to rise.

Throughout, I use the scale-free attention cost discussed in Gabaix (2016), which is invariant to transformations of the utility function. Formally, the cost of cognition $\mathcal{K}$ is:[21]

$$\mathcal{K} = k^2(a_d)^2|V_{aa}|$$

Here, k is the scale-free attention cost, a_d is the action taken under the default level of attention, and V_{aa} is the second derivative of the value function with respect to the action, evaluated at m_d. The intuition for k is that if $k = 0.01$, for example, then agents pay attention to variables that on average make more than a 1% difference to their optimal action at the typical scale. I denote the cost of attention for firms and consumers as k_c and k_f respectively. To make the derivations easier, I replace the default action a_d with the steady-state action $\bar{a}$. This makes little difference quantitatively, since the optimal price and consumption tend to differ from their steady-states by at most a few percentage points under a standard parameterisation. Steady-state consumption is 1 because Gabaix (2020) assumes a corrective production subsidy that offsets the distortions from monopolistic competition. Equally, the zero-inflation steady-state choice of relative price for firm i, $P_{i,t}/P_t$, is equal to 1.[22] Consequently, λ simplifies to:

$$\lambda = \frac{E[(\partial a/\partial m)^2]}{k^2}$$

What happens when the agent multiple choice variables? The consumer, for example, decides on both consumption and labour supply. Fortunately, the intratemporal optimality condition imposes a relationship between these two actions, so the problem can be formulated solely in terms of consumption.[23] I calibrate the scale-free attention cost for

[19]This result is given in Gabaix (2016), equation (15).

[20]See Gabaix (2016) p.12 for further details on the continuity and sparsity of the attention function.

[21]See Gabaix (2016), equation (75).

[22]However, throughout this paper, deviations are expressed in percentage point terms, and so the steady state action becomes 100, and a 1% scale-free cost of attention is given by $k = 1.0$ rather than $k = 0.01$.

[23]This is shown in equation (221) in the online appendix to Gabaix (2020).

the consumer problem based on the percetange difference that paying attention makes to consumption.

Consumer Attention

I now solve for the optimal choice of cognitive discounting parameters, taking as given the model's dynamics. As discussed above, the key quantity in the consumer case will be $E[(\partial c_t/\partial m_c)^2]$. Gabaix (2020) shows that consumption can be written as:

$$c_t = E_t \sum_{h\geq 0} (\beta m_c)^h \left((1-\beta) y_{t+h} - \beta\gamma^{-1} r_{t+h}\right) \tag{1.15}$$

In our case, the equation is slightly altered because of the presence of the discount factor shock.

$$c_t = E_t \sum_{h\geq 0} (\beta m_c)^h \left((1-\beta) y_{t+h} - \beta\gamma^{-1} r_{t+h} - \beta\gamma^{-1} z_{t+h}\right) \tag{1.16}$$

This expression is derived in appendix 1.A. Gabaix derives an analytical expression for $E[(\partial c_t/\partial m_c)^2]$ under the assumption of determinacy with only technology shocks.[24] In appendix 1.A, I show how to solve for my more general case, which accommodates multiple shocks, lagged terms in the model equations, and indeterminacy. I show that one arrives at:

$$\begin{aligned} E[(\partial c_t/\partial m_c)^2] &= \boldsymbol{P_c \Sigma_s P_c'} \\ \boldsymbol{P_c} &= \beta \boldsymbol{k_c} (\boldsymbol{I} - \beta m_{c,d} \boldsymbol{Z_1})^{-2} \boldsymbol{Z_1} \end{aligned} \tag{1.17}$$

Here, $\boldsymbol{k_c}$ is a selection matrix defined in appendix 1.A. Note that the dynamics of the model are taken as given, so that $\boldsymbol{Z_1}$, and indeed $\boldsymbol{\Sigma_s}$, are unaffected by evaluating at the default.

Firm Attention

The firm cognitive discount factor can be derived in a similar fashion. Gabaix shows that the firms' choice of relative price is described by:

$$q_t = (1-\beta\theta) \sum_{h\geq 0} (\beta\theta m_f)^h E_t(\pi_{t+1} + ... + \pi_{t+h} + mc_{t+h}) \tag{1.18}$$

[24] Specifically, see Gabaix (2020) equations (175) and (227).

The critical term for the firm choice of attention is then $E[(\partial q_t/\partial m_f)^2]$. In appendix 1.A, I show that this expression can be written as:

$$E[(\partial q_t/\partial m_f)^2] = \boldsymbol{P_f \Sigma_s P_f'} \tag{1.19}$$
$$\boldsymbol{P_f} = (1-\beta\theta)\beta\theta(\boldsymbol{k_\pi Z_1 (I - Z_1)^{-1}}((\boldsymbol{I} - \beta\theta m_{f,d}\boldsymbol{I})^{-2} - (\boldsymbol{I} - \beta\theta m_{f,d}\boldsymbol{Z_1})^{-2}\boldsymbol{Z_1})$$
$$+\boldsymbol{k_{mc}}(\boldsymbol{I} - \beta\theta m_{f,d}\boldsymbol{Z_1})^{-2}\boldsymbol{Z_1})$$

Here, $\boldsymbol{k_\pi}$ and $\boldsymbol{k_{mc}}$ are selection matrices defined in appendix 1.A.

1.2.3 Attention Equilibrium

The optimal choices of attention for consumers and firms are then functions of the model's parameters $\boldsymbol{\chi}$, the parameters of the attention problem (which I will denote using the vector $\boldsymbol{\xi}$), and the model dynamics as summarized by $\boldsymbol{Z_1}$ and $\boldsymbol{\Sigma_S}$. Recall that $\boldsymbol{Z_1}$ and $\boldsymbol{\Sigma_s}$ are in turn functions $\boldsymbol{m}$ and $\boldsymbol{\chi}$. The optimal choices of attention (denoted $\boldsymbol{m}^{opt}$) can consequently be written as a function of $\boldsymbol{\chi}$, $\boldsymbol{\xi}$, and $\boldsymbol{m}$.

$$\boldsymbol{m}^{opt} = \boldsymbol{g}(\boldsymbol{m}, \boldsymbol{\chi}, \boldsymbol{\xi}) \tag{1.20}$$

I define an attention equilibrium as a fixed-point of the $\boldsymbol{g(\cdot)}$ mapping.

Definition 1 *An attention equilibrium is a vector* $\boldsymbol{m(\chi, \xi)} \in \mathcal{A}$ *such that*

$$\boldsymbol{m(\chi, \xi) = g(m(\chi, \xi), \chi, \xi)} \tag{1.21}$$

In equilibrium, the level of attention and the dynamics of the system are mutually consistent. That is, the levels of attention implies a set of dynamics which in turn justifies those levels of attention as optimal choices.

The Phillips and IS curves of the behavioural new Keynesian model then become:

$$\pi_t = \beta M_f(\boldsymbol{\chi}, \boldsymbol{\xi}) E_t \pi_{t+1} + \kappa y_t + \eta_t \tag{1.22}$$
$$y_t = M_c(\boldsymbol{\chi}, \boldsymbol{\xi}) E_t y_{t+1} - \sigma(i_t - E_t \pi_{t+1}) - \sigma z_t \tag{1.23}$$
$$M_f(\boldsymbol{\chi}, \boldsymbol{\xi}) = m_f(\boldsymbol{\chi}, \boldsymbol{\xi}) \left(\theta + (1-\theta)\frac{(1-\beta\theta)}{1 - \beta m_f(\boldsymbol{\chi}, \boldsymbol{\xi})\theta}\right)$$
$$M_c(\boldsymbol{\chi}, \boldsymbol{\xi}) = m_c(\boldsymbol{\chi}, \boldsymbol{\xi})$$

Note that the equilibrium choices of attention are a function of the parameters $\boldsymbol{\chi}$ and $\boldsymbol{\xi}$, but do not vary over time. That is, they are regime-dependent rather than state-dependent. As such, for a given set of parameters, the model can be solved using a two-

step procedure. First, one solves for the equilibrium attention levels $\boldsymbol{m}(\boldsymbol{\chi}, \boldsymbol{\xi})$. Second, one can solve the model treating those choices of attention as fixed values.

1.3 Determinacy

In this section I analyze the determinacy condition of the model. To simplify the analysis, I set $M_f = 0$ throughout, giving a static Phillips curve, while allowing M_c to be determined endogenously.[25] To begin with, I also drop the fundamental shocks from the model. These are not necessary to obtain the main result, presented in 1.3.1, that an indeterminate attention equilibrium exists whenever the rational expectations Taylor principle is violated, although it may exist alongside a determinate attention equilibrium. Section 1.3.2 then shows that the determinacy condition is *path dependent*. If one starts in the determinate attention equilibrium, it is not always possible to switch into the indeterminate state. Section 1.3.3 then reintroduces fundamental shocks, and argues that whether a determinate attention depends on the variance of the fundamental shocks.

1.3.1 A New Statement of the Taylor Principle

As noted above, I begin with a simple version of the model in which I exogenously set $M_f = 0$, while allowing M_c to be determined endogenously as described in section 1.2. I initially drop the fundamental shocks from the model, to focus on the role of sunspot shocks, and assume a simplified version of the policy rule with $\rho_r = \phi_y = 0$. These assumptions give the following model:

$$\pi_t = \kappa y_t \tag{1.24}$$

$$y_t = M_c E_t y_{t+1} - \sigma(i_t - E_t \pi_{t+1}) \tag{1.25}$$

$$i_t = \phi_\pi \pi_t \tag{1.26}$$

Taking as given the level of attention, the model can be re-written as:

$$y_t = \delta E_t y_{t+1} \tag{1.27}$$

$$\delta = \frac{M_c + \kappa\sigma}{1 + \kappa\sigma\phi_\pi}$$

If $\delta < 1$ then the model has a unique stable solution: $y_t = 0$ for all t. This requires:

$$\phi_\pi > 1 - \frac{1 - M_c}{\kappa\sigma} \tag{1.28}$$

[25] One can equally set $M_c = 0$ and allow M_f to vary. This does not change the results outlined below.

If $M_c = 1$, so that agents pay full attention, then expectations are rational, and so one obtains the familiar condition that $\phi_\pi > 1$ ensures determinacy. As M_c decreases, so agents discount the future more, the determinacy condition relaxes. As Gabaix notes, if κ and σ are sufficiently small, then even moderate discounting ensures determinacy for any $\phi_\pi \geq 0$.

If $\delta > 1$, then the model is indeterminate. Specifically, y_t follows an autoregressive process:

$$\begin{aligned} y_t &= \delta^{-1} y_{t-1} + \omega_t \\ \omega_t &= y_t - E_{t-1} y_t \end{aligned} \tag{1.29}$$

Absent fundamental shocks, the expectation error is just given by a sunspot shock, which I denote ζ_t. The sunspot shock must be mean zero, but the model places no restrictions on higher moments of ζ_t. I assume that it has constant variance σ_ζ^2.[26]

If $\omega_t = \zeta_t$, then the stationary distribution of y_t then has variance:

$$\sigma_y^2 = \frac{\sigma_\zeta^2}{1 - \delta^{-2}} \tag{1.30}$$

Notice that at the very edge of the indeterminacy region, as δ approaches unity from above, the stationary variance of output rises becomes unboundedly large. Why? In the indeterminacy region, sunspot shocks are mean-reverting, which is why we do not rule them out when finding stable solutions to the model. In the determinacy region they are explosive, which is why they are ruled out. At the boundary, they have a unit root. As one approaches the boundary, then, the sunspot shocks approach a unit root and so generate unboundedly large variance in output over a sufficiently long time horizon.

If M_c is sufficiently low that $\delta < 1$, then there is no incentive to pay attention: $c_t = 0$ in all t, so agents always choose to pay the default level of attention. Provided that $\delta < 1$ when $M_c = m_{c,d}$, then there is an attention equilibrium at the default level.

Suppose instead that $\delta > 1$. What are the incentives to pay attention? We have:

$$c_t = \sum_{h \geq 0} (\beta m_c)^h ((1 - \beta) E_t y_{t+h} - \beta \sigma E_t i_{t+h} + \beta \sigma E_t \pi_{t+h+1}) \tag{1.31}$$

After some algebra, one finds that the expectation of the square of the derivative (eval-

[26]The assumption that the variance does not change over time is innocuous. Less innocuous is the assumption that it does not depend on δ. That is, there is no tendency for sunspot shocks of different persistence to have different variance. In particular, for the derivation below, one would have to assume that σ_ζ^2 does not tend to zero as δ approaches unity from above.

uated at the default) is:

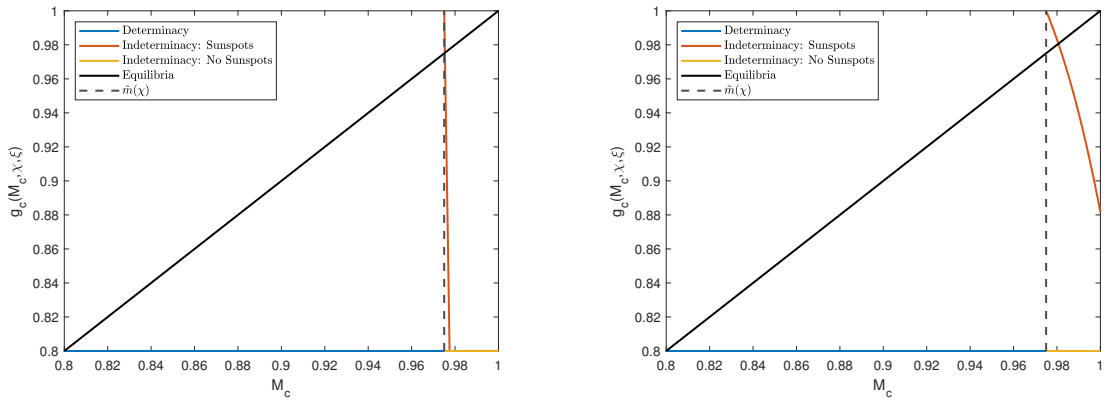
$$E[(c_{m,t})^2] = \frac{\beta^2\delta^{-2}}{(1-\beta m_{c,d}\delta^{-1})^4}(1-\beta-\beta\sigma\kappa(\phi_\pi-\delta^{-1}))^2\sigma_y^2 \tag{1.32}$$

$$= \frac{\beta^2\delta^{-2}}{(1-\beta m_{c,d}\delta^{-1})^4}(1-\beta-\beta\sigma\kappa(\phi_\pi-\delta^{-1}))^2\frac{\sigma_\zeta^2}{1-\delta^{-2}} \tag{1.33}$$

Denote the function for the optimal consumer choice of attention as g_c (corresponding to the second row of the $\boldsymbol{g(\cdot)}$ mapping described above). Figure 1.1 plots the g_c mapping in two calibrated examples. The calibration used has: $\beta = 0.99$, $\kappa = 0.10$, and $\sigma = 0.50$. I assume that $\phi_\pi = 0.50$, so that the rational expectations Taylor Principle is violated. I assume that $m_{c,d} = 0.80$, and $k_c = 1.0$. Attention equilibria are on the 45 degree line, where a given level of aggregate attention maps into the same choice of individual (and hence aggregate) attention. The edge of the determinacy region is shown by the vertical dashed line, with the corresponding level of attention (that gives $\delta = 1$) denoted as $\tilde{m}(\boldsymbol{\chi})$. In the determinate region, the g_c mapping is shown by the blue line. Absent fundamental shocks, output volatility and hence incentives to pay attention are zero. So, the blue line remains at the default level of attention, whatever is aggregate attention. That gives an attention equilibrium at the default level of 0.80. In the indeterminate region, the yellow line shows the case where the sunspot shock variance is equal to zero. Here, the no sunspot shock case implies zero output volatility, no incentives to pay attention, and so an optimal choice of attention at the default. The red line shows the g_c mapping. In panel (a), I set $\sigma_\zeta = 0.2$, while in panel (b), I use $\sigma_\zeta = 1.0$. Notice that as the economy approaches the very edge of the indeterminacy region, δ approaches unity from above, and the variance of output rises without limit. Incentives to pay attention become unboundedly large, and so the optimal choice of attention (shown by the red line) approaches unity. The continuity of $g_c(M_c, \boldsymbol{\chi}, \boldsymbol{\xi_c})$ in M_c within the indeterminacy region then ensures that an indeterminate equilibrium always exists. This situation holds whenever the rational expectations Taylor principle fails. Our specific example always allows for a determinate equilibrium to exist. As such, the model features two kinds of multiplicity. First, there are multiple equilibrium choices of attention. Second, one of those equilibrium choices of attention allows for multiple stable solutions.

1.3.2 Path Dependence

The indeterminacy result above relied on a quite specific assumption; I showed that there always exists an indeterminate equilibrium, provided that the economy has been in the indeterminate state for an arbitrarily long time, so that the variance of output was given



(a) Small sunspot shocks **(b)** Large sunspot shocks

Figure 1.1: Attention equilibria

by:

$$\sigma_y^2 = \frac{\sigma_\zeta^2}{1-\delta^{-2}}$$

From now on, I will refer to this equilibrium as a *long-run* equilibrium, to denote that it is an equilibrium that exists if the economy has been in the indeterminate region already for a long time.

Suppose instead that the economy was in the determinate state in time $t-1$. Then $y_{t-1}=0$. I now ask: could the economy make an unexpected jump into an indeterminate equilibrium in time t? In time t, output is given by:

$$y_t = \delta^{-1} y_{t-1} + \zeta_t \tag{1.34}$$

The variance of output is:

$$Var(y_t) = \delta^{-2} Var(y_{t-1}) + \sigma_\zeta^2 \tag{1.35}$$

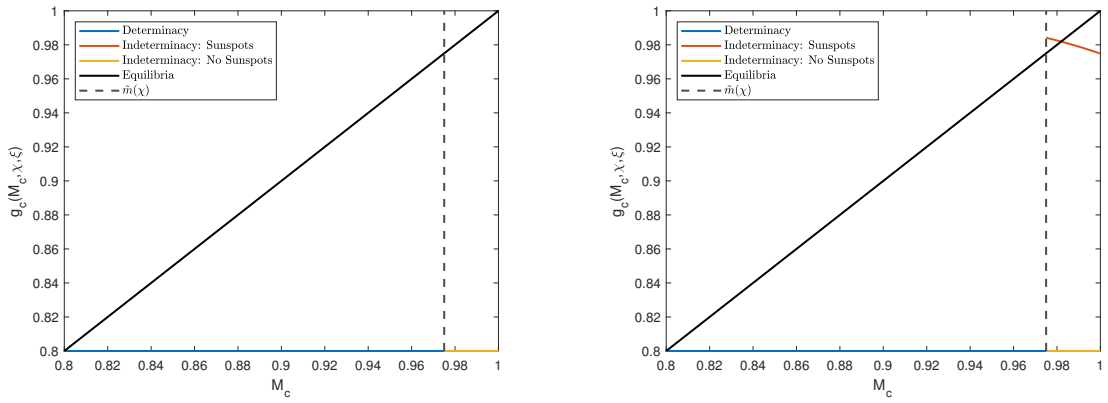
$$= \sigma_\zeta^2 \tag{1.36}$$

Then the critical term in calculating the g_c mapping is:[27]

$$E[(c_{m,t})^2] = \frac{\beta^2\delta^{-2}}{(1-\beta m_{c,d}\delta^{-1})^4}(1-\beta-\beta\sigma\kappa(\phi_\pi-\delta^{-1}))^2\sigma_y^2 \tag{1.37}$$

$$= \frac{\beta^2\delta^{-2}}{(1-\beta m_{c,d}\delta^{-1})^4}(1-\beta-\beta\sigma\kappa(\phi_\pi-\delta^{-1}))^2\sigma_\zeta^2 \tag{1.38}$$

[27]Strictly, this simplifies a little by supposing that agents assume that if a particular M_c prevails in time t, then it will always prevail. Otherwise, the δ would be expected to change over time, which would make the calculation more complex. This simplification makes the derivations that follow more tractable.



(a) Large sunspot shocks **(b)** Very, very large sunspot shocks

Figure 1.2: Attention equilibria, assuming determinacy in the previous period

Note that this term does *not* become unboundedly large as δ approaches unity from above. Hence although there is a discontinuity in the g_c mapping at $\tilde{m}(\boldsymbol{\chi})$, it is a small discontinuity, rather than taking the g_c mapping up to unity. Indeed, if the variance of the sunspot shocks is small enough, then attention may remain at its default level as one crosses the edge of the determinacy region. This situation is shown in figure 1.2a, which uses the same calibration as figure 1.1b. Notice that under either level of attention cost, the output of the g_c mapping remains at its default level at the boundary. As such, there is no indeterminate equilibrium. So, if the economy starts in the determinate equilibrium, it will only be able to reach the indeterminate equilibrium if the variance of the sunspot shocks is sufficiently large. Consequently, the calibration considered in 1.1b and 1.2a allows an indeterminate equilibrium to be sustained if the model has already been in that equilibrium for an arbitrarily long time, but does not allow the economy ever to switch into the indeterminate state. In this sense, the long-run indeterminate equilibrium shown in 1.2b is *unattainable*.

The determinacy condition of the model is consequently dynamic, and path dependent. Starting off in the determinate regime means that the model is likely to remain there. Looking at long-run equilibria restored a version of the rational expectations Taylor principle; whenever the rational expectations Taylor principle fails, a long-run indeterminate equilibrium exists. If one looks only at the equilibria that can be reached, however, then policy rules which break the rational expectations Taylor principle will often only allow for a determinate equilibrium to exist.

Of course, the model does not place restrictions on the size of the sunspot shocks. One can always find a sufficiently large σ_ζ such that an indeterminate equilibrium exists. This situation is shown in figure 1.2b, in which I set the sunspot shock variance to be ten times larger, so that $\sigma_\zeta = 10$. This generates enormous volatility in consumption, which allows an indeterminate equilibrium will exist even in the period after the switch; notice

the intersection of the red line in figure 1.2b with the 45 degree line. Nonetheless, one might be skeptical about the possibility of such large sunspot shocks; these would require an extraordinary degree of coordination between agents.

What happens if an indeterminate equilibrium does exist in the period after the switch? I now introduce some additional notation. First, denote $M_{c,t+\tau}$ as the aggregate level of attention chosen in period $t+\tau$, and $\delta_{t+\tau}$ as δ implied by this level of attention. Now in period $t+\tau$, supposing that the economy remains within the indeterminate region:

$$Var(y_{t+\tau}) = \delta_{t+\tau-1}^{-2} Var(y_{t+\tau-1}) + \sigma_\zeta^2 \tag{1.39}$$

It must be that the variance of $y_{t+\tau}$ is at least as great as the variance of y_t, which means that the resulting g_c will be at least as great for any M_c. This means that if an indeterminate equilibrium exists in the first period after the switch, it will exist in all subsequent periods. In the simple example considered so far, then, if an indeterminate equilibrium exists in the first period after the switch, then it will exist in all subsequent periods, with attention drifting up towards the long-run equilibrium over time. To formally define the notion of equilibrium attainability, I first define a new mapping $g_c^{att}(M_c, \boldsymbol{\chi}, \boldsymbol{\xi_c})$, which is identical to the original g_c mapping but using the derivative of consumption with respect to attention in (1.38). I then define an indeterminate equilibrium given by the long-run g_c mapping as attainable only if there also exists an indeterminate equilibrium of the short-run g_c^{att} mapping.

Definition: An equilibrium choice of attention $M_c(\boldsymbol{\chi}, \boldsymbol{\xi_c})$ that lies in the indeterminate region is *attainable* if an only if there exists some $M_c^* \in [\tilde{m}(\boldsymbol{\chi}), 1]$ such that:

$$M_c^* = g_c^{att}(M_c^*, \boldsymbol{\chi}, \boldsymbol{\xi_c}) \tag{1.40}$$

1.3.3 Shock Dependence

So far I abstracted entirely from fundamental shocks. That meant that in the determinate state, output and inflation were always zero, and hence there were no incentives to pay attention, guaranteeing the existence of a determinate attention equilibrium at the default level of attention whenever $m_{c,d} < \tilde{m}(\boldsymbol{\chi})$.

Let's now introduce a fundamental shock. Suppose that there is a policy shock, so that the nominal interest rate is:

$$i_t = \phi_\pi \pi_t + v_t \tag{1.41}$$

$$v_t = \rho v_{t-1} + \varepsilon_t^v \tag{1.42}$$

$$\varepsilon_t^v \sim N(0, \sigma_{\varepsilon,v}^2)$$

The model now reduces to:

$$y_t = \delta E_t y_{t+1} - \tilde{v}_t \tag{1.43}$$

$$\tilde{v}_t = \frac{\sigma}{1 + \kappa\sigma\phi_\pi} v_t \tag{1.44}$$

I denote the variance of $\tilde{v}_t$ as $\tilde{\sigma}_v^2$.

Under determinacy, the model solution is:

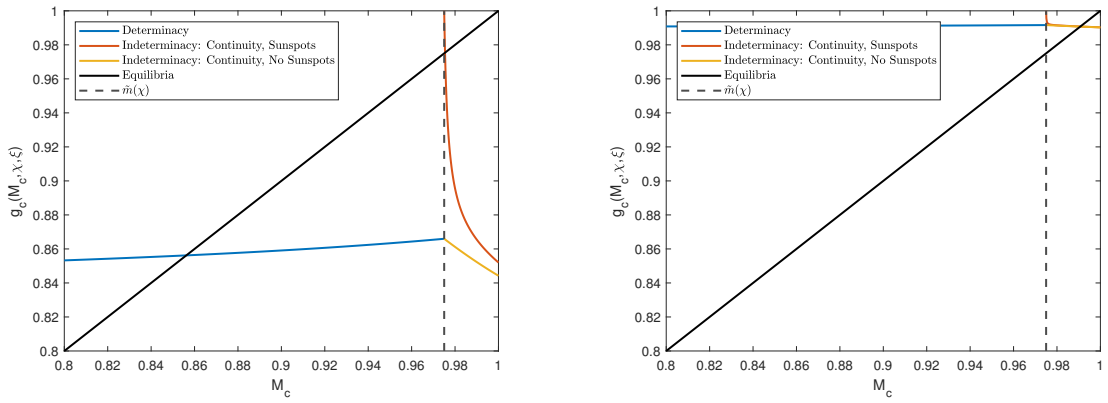
$$y_t = -\frac{1}{1 - \delta\rho}\tilde{v}_t \tag{1.45}$$

The consumption choice is linear in $\tilde{v}_t$, because all of the endogenous variables and their expectations are just linear functions of $\tilde{v}_t$. Hence the derivative of consumption with respect to m_c is also linear in $\tilde{v}_t$. The expectation of the square of this quantity is then just linear in $\tilde{\sigma}_v^2$ and hence in $\sigma_{\varepsilon,v}^2$.

Now suppose that $\tilde{m}(\boldsymbol{\chi}) < 1$, so that the economy is indeterminate under rationality. For any aggregate M_c within the determinate region, there exists some sufficiently large $\sigma_{\varepsilon,v}^2$ such that the implied individual choice of m_c approaches unity. That will ensure there is no determinate equilibrium. As such, the existence of determinate equilibrium under any policy rule with $\phi_\pi < 1$ depends critically on the size of the shocks hitting the economy. This situation is shown in figures 1.3a and 1.3b. These figures assume that $\rho = 0.8$. 1.3a assumes that $\sigma_{\varepsilon,v} = 0.5$, while 1.3b assumes that $\sigma_{\varepsilon,v} = 2.0$. Notice the difference in the blue line, that shows the optimal choice of attention in the determinate region. In 1.3a, with the smaller shocks, incentives to pay attention are sufficient to move the g_c mapping above the default level. However, a determinate equilibrium still exists at the intersection of the blue line with the 45 degree line. In 1.3b, however, the fundamental shocks are larger. The g_c mapping in the determinate region is shifted upwards, to the extent that no equilibrium exists in the determinate region.

What does the long-run equilibrium look like in the indeterminate region? To pin this down, one has to not only specify the sunspot shock variance, but also the response of the endogenous variables to the fundamental shocks. The latter is not pinned down in the indeterminacy region, but one has to specify a solution in order to calculate the attention level. I will throughout use Lubik and Schorfheide's (2004) continuity solution. In this simple example, that means that at the very edge of the indeterminacy region, expectation errors are given by:

$$\omega_t = \frac{1}{1 - \rho}\varepsilon_t^v + \zeta_t \tag{1.46}$$



(a) Small fundamental shocks, small sunspot shocks**(b)** Large fundamental shocks, small sunspot shocks

Figure 1.3: Attention equilibria with fundamental shocks

The reason is that this is the response at the boundary under determinacy (plus a sunspot shock), so the responses to fundamental shocks are continuous as one crosses the boundary. This assumption gives a finite variance for output and the boundary of the indeterminacy region if and only if $\sigma_\zeta^2 = 0$. This case is shown by the yellow lines in figures 1.3a and 1.3b. Because the responses to the fundamental shocks are continuous as one crosses the boundary of the determinacy region, and there are no sunspot shocks, then the g_c mapping is continuous at the boundary, and there is no jump from the blue to the yellow line. However, if one deviates even slightly from the continuity solution, or if the sunspot shock variance is non-zero, then one must have an infinite output variance at the boundary, meaning that a long-run indeterminate equilibrium exists. The red line shows the g_c mapping when the sunspot shock variance is non-zero. Once again, the cumulation of sunspot shocks over time ensures that the g_c mapping approaches unity as the boundary of the indeterminate region. This effect in turn ensures that a long-run indeterminate equilibrium is sure to exist because of the continuity of the g_c mapping in M_c.

1.4 Stability at the ELB

I now explore a corollary of the findings in section 1.3; if the economy is at the ELB and is expected to remain there permanently, then the economy *may* be stable but there will always exist an attention equilibrium that implies indeterminacy. This contrasts with Gabaix's finding that exogenous cognitive discounting implies stability at the ELB.

I once again assume a static Phillips curve:

$$\pi_t = \kappa y_t$$

However, I now assume that $i_t = 0$ for all t, and that the neutral real rate of interest is negative. The economy's IS curve is then:

$$y_t = M_c E_t y_{t+1} + \sigma E_t \pi_{t+1} + \sigma \bar{r}$$

Where $\bar{r} < 0$. In steady-state, this implies:

$$\bar{y} = \frac{1}{1 - M_c - \kappa\sigma} \bar{r}$$

Now define $y_t^D \coloneqq y_t - \bar{y}$. It follows that:

$$y_t^D = (M_c + \kappa\sigma) y_{t+1}^D$$

If $M_c < 1 - \kappa\sigma$ then the model is determinate, and $y_t = \bar{y}$ for all t. y_t^D then has zero variance. If instead $M_c > 1 - \kappa\sigma$, the model is indeterminate, and so:

$$y_t^D = \frac{1}{M_c + \kappa\sigma} y_{t-1}^D + \zeta_t$$

Where ζ_t is a mean zero sunspot shock as before. Solving for the variance of y_t^D in the indeterminate case:

$$\sigma_{y^D}^2 = \frac{1}{1 - \left(\frac{1}{M_c + \kappa\sigma}\right)^2} \sigma_\zeta^2$$

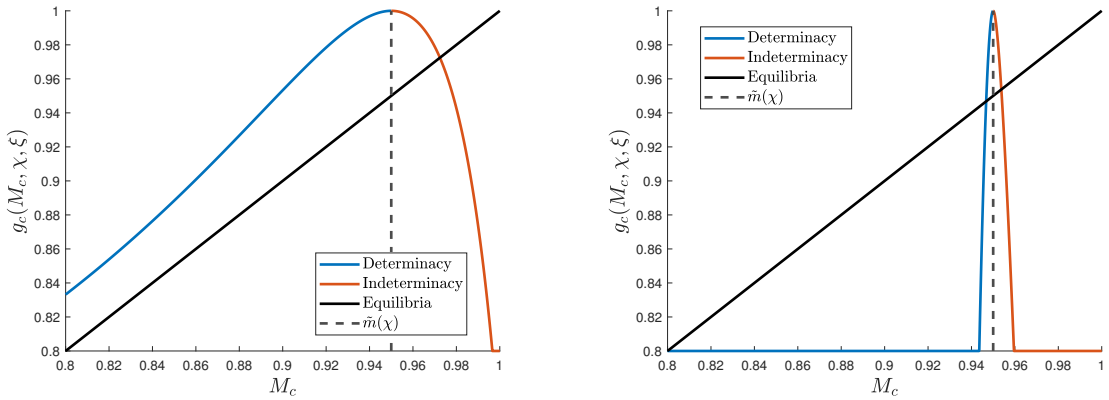
Notice again that this expression becomes unboundedly large as the economy approaches the edge of the indeterminacy region. I now solve for the optimal choice of attention. Consumption can be written as:

$$\begin{aligned} c_t &= E_t \sum_{h \geq 0} (\beta m_c)^h \left((1 - \beta) y_{t+h} - \beta\sigma (i_{t+h} - \pi_{t+h+1} - \bar{r}) \right) \\ &= E_t \sum_{h \geq 0} (\beta m_c)^h \left((1 - \beta) y_{t+h} + \beta\sigma \pi_{t+h+1} + \beta\sigma \bar{r} \right) \end{aligned}$$

Substituting in the expressions above, one can obtain:

$$\begin{aligned} c_t = \frac{1}{1 - \beta m_c} &\left((1 - \beta + \beta\kappa\sigma) \frac{1}{1 - M_c - \kappa\sigma} \bar{r} + \beta\sigma\bar{r} \right) \\ &+ \left(1 - \beta + \frac{\beta\kappa\sigma}{M_c + \kappa\sigma} \right) \frac{M_c + \kappa\sigma}{M_c + \kappa\sigma - \beta m_c} y_t^D \end{aligned}$$

Taking the derivative with respect to m_c, evaluated at $m_{c,d}$, and then taking the expec-



(a) Very large $\bar{r}$ **(b)** Very small $\bar{r}$

Figure 1.4: Attention equilibria

tation of the square yields the key expression for the attention problem:

$$E\left[\left(\frac{\partial c_t}{\partial m_c}\right)^2\right] = \frac{\beta^2}{(1-\beta m_{c,d})^4}\left((1-\beta+\beta\kappa\sigma)\frac{1}{1-M_c-\kappa\sigma}\bar{r}+\beta\sigma\bar{r}\right)^2 + \left(1-\beta+\frac{\beta\kappa\sigma}{M_c+\kappa\sigma}\right)^2\frac{\beta^2(M_c+\kappa\sigma)^2}{(M_c+\kappa\sigma-\beta m_{c,d})^4}\sigma^2_{y^D}$$

For the purposes of this example, I will assume that $m_{c,d} < 1-\kappa\sigma$, such that the economy is determinate if attention is at its default level, but is indeterminate under rationality. Let us first analyze the choice of attention within the determinacy region, that is if $M_c < 1-\kappa\sigma$. The second term is then zero. However, note that the first term becomes unboundedly large as the economy approaches the boundary of the determinacy region. That is because $\bar{y}$ becomes unboundedly large (that is, there is an unboundedly large recession) at the edge of the determinacy region. Whether an attention equilibrium exists in the determinacy region then depends on the parameterization. If $\bar{r}$ is very large, for example, then any $M_c \in [M_{c,d}, 1]$ will result in a sufficiently large recession to mean that the optimal choice of attention exceeds $1-\kappa\sigma$, and hence no determinate equilibrium will exist. This case is shown in figure 1.4a.

An alternative case is shown in figure 1.4b. Here, $\bar{r}$ is very small. As such, when M_c is low the optimal choice of m_c will lie at the default, giving an equilibrium at $m_{c,d}$. At the economy approaches the edge of the determinacy region, however, the recession induced by the ELB must become unboundedly large and the optimal choice of attention consequently approaches unity. This results in a second, higher attention and more costly equilibrium close to the edge of the determinacy region.

Now let us consider the indeterminacy region. Both the first and the second term are unboundedly large at the edge of the region, meaning that the optimal choice of attention is always 1. Both terms are continues in M_c for $M_c \in (1-\kappa\sigma, 1]$, which means that an

equilibrium must exist within the indeterminacy region.

This analysis points to two main conclusions. First, while the economy *may* be stable at the ELB if the neutral real rate is not too negative, there always also exists an attention equilibirum in the indeterminate region. As such, stability is not a given. Second, if the economy is stable at the ELB, then there will often be multiple determinate equilibria. While lower attention equilibria lie further from the boundary of the determinacy region and imply a smaller recession, there may also be a higher attention equilibrium closer to the boundary that results in a much higher output cost. This implies that the ELB *may* be much less costly under endogenous discounting than rationality, but this is not a given as it is in the exogenous discounting model.

In appendix 1.C, I consider more complex setting in which the economy switches between a regime in which the ELB binds and a regime in which it does not. In this regime switching setting, I show that targeting average inflation when the ELB does not bind can in numerical examples rule out the possibility of these costly determinate equilibria in the ELB regime. As such, average inflation targeting is potentially much more powerful under endogenous than exogenous discounting.

An extension of my regime switching model to a somewhat richer setting has been considered by Ascari et al. (2023).

1.5 Optimal Policy

In this section, I develop a method to solve for optimal policy. Section 1.5.1 discusses the assumed central bank loss function. Section 1.5.2 illustrates how optimal policy under discretion differs under endogenous attention compared to the exogenous attention or rational expectations models. Section 1.5.3 solves for the optimal policy under commitment, identifying a new tradeoff for the central bank when cognitive discounting is endogenous.

1.5.1 The Loss Function

I assume that the central bank minimizes a standard loss function:

$$L_t = E_t \sum_{h \geq 0} \beta^h \left(\pi_{t+h}^2 + \vartheta y_{t+h}^2 \right) \tag{1.47}$$

Minimizing this objective function is equivalent to minimizing the loss function given in equation 40 in Gabaix (2020), which allows for direct comparison with the results that Gabaix obtains under the assumption of exogenous discounting. As Gabaix notes, this

loss function corresponds to a welfare loss under two assumptions. First, subjective rather than objective expectations appear in the welfare function. Second, cognitive costs are not included in welfare. These assumptions could be debated, and other papers in the literature take different approaches.[28] As such, I will refer to the loss function as "central bank losses" rather than "welfare". The policies that I discuss below are optimal in the sense of minimizing this objective function. I briefly discuss how the analysis might change under alternative assumptions for the loss function in appendix 1.D, and future work could consider this question in more depth.

1.5.2 Optimal Policy under Discretion

Under discretion, the central bank cannot commit to a particular policy regime. As such, the central bank must take the equilibrium choices of attention as given. Minimizing 1.47 subject to the Phillips curve, taking as given the expectations terms and the cognitive discount factor, yields the familiar choice of inflation-output tradeoff:

$$y_t = -\frac{\kappa}{\vartheta}\pi_t$$

Given this choice of inflation-output tradeoff, and taking as given the cognitive discount factor in the Phillips curve, one arrives at the following paths for output, inflation, and real interest rates:

$$\begin{aligned}
\pi_t &= \frac{\vartheta}{\kappa^2 + \vartheta(1 - \beta M_f \rho)}\eta_t \\
y_t &= -\frac{\kappa}{\kappa^2 + \vartheta(1 - \beta M_f \rho)}\eta_t \\
i_t &= \frac{\kappa\sigma^{-1}(1 - M_c\rho_\eta) + \vartheta\rho_\eta}{\kappa^2 + \vartheta(1 - \beta M_f \rho_\eta)}\eta_t
\end{aligned}$$

These solutions for the endogenous variables in terms of the attention levels are the same as in Gabaix's exogenous discounting model. The difference versus Gabaix's results will be that the attention choices are determined endogenously. It is worth noting that if $\rho > 0$, then higher firm attention M_f strictly increases inflation and output volatility, and consequently leads to higher central bank losses. This is because higher firm attention amplifies the effects of persistent cost-push shocks on the inflation rate. This generates a worse "static" tradeoff between π_t and y_t in the Phillips curve.

The level of consumer attention does not matter for inflation and output volatility, but does matter for implementation; the optimal path for output can be achieved with a lower level of real interest rate variability if consumers are more attentive.

[28] See footnote 41 in Gabaix (2020) for a review.

Before endogenizing attention, there is one point that is worth making about the exogenous attention model. Under discretion, one has the classic "stabilization bias" problem. In response to a cost-push shock, a policymaker operating under discretion stabilizes output too much and inflation too little. The reason is that they treat the expectation term in the Phillips curve as exogenous, and so do not internalize the fact that responding more aggressively to shocks reduces inflation indirectly through inflation expectations as well as directly through the output term in the Phillips curve. Under cognitive discounting, expected inflation has a smaller effect on current inflation, and so the degree of stabilization bias is reduced.

To see this formally, suppose that the central bank's remit is to minimize a loss function given by 1.47. Suppose that the government appoints a central bank whose preferences are described by the same loss function, but with a weight on output deviations of $\tilde{\vartheta}$. By appointing a "conservative" central banker à la Rogoff (1985) with preferences $\tilde{\vartheta} < \vartheta$, losses can be reduced. In particular, the optimal $\tilde{\vartheta}$ to choose is:

$$\tilde{\vartheta}^* = \vartheta(1 - \beta M_f \rho_\eta)$$

As such, when firms pay more attention to the future, it becomes advantageous to appoint more conservative central bankers to offset the stabilization bias.

Figure 1.5 illustrates this points graphically. The y-axis shows the expected value of the per-period loss in 1.47, for different values of $\tilde{\vartheta}$. The vertical dashed line shows the value of ϑ, which is calibrated to be 0.0625. The blue line shows the rational expectations case ($M_f = 1$), while the red line shows an example with exogenous discounting ($M_f = 0.6$). The loss-minimizing value of $\tilde{\vartheta}$, shown by the dots, is lower in the rational expectations case than under exogenous discounting.[29]

I now consider the case where the cognitive discount factor is endogenous. A given set of attention choices M_c and M_f implies a set of macroeconomic dynamics under discretion, which in turn imply particular choices of attention as optimal. The equilibrium choices of attention will be those that imply a set of macroeconomic dynamics under discretion, which in turn implies those same particular choices of attention as optimal. For different values of $\tilde{\vartheta}$, I numerically find the equilibrium M_c and M_f, and then the welfare losses.

Figure 1.6 shows the losses under endogenous discounting, compared to the exogenous discounting and rationality cases. Figure 1.6a shows a calibration with smaller shocks, while figure 1.6b shows a calibration with larger shocks. If the shocks are smaller, then under endogenous discounting firms choose the default level of attention unless the central banker's concern for output losses ($\tilde{\vartheta}$) is very large. That is, it is only when the central

[29]Note also that the difference between the loss at the optimal $\tilde{\vartheta}$ and the loss where $\tilde{\vartheta} = \vartheta$ is proportionally larger under rationality.

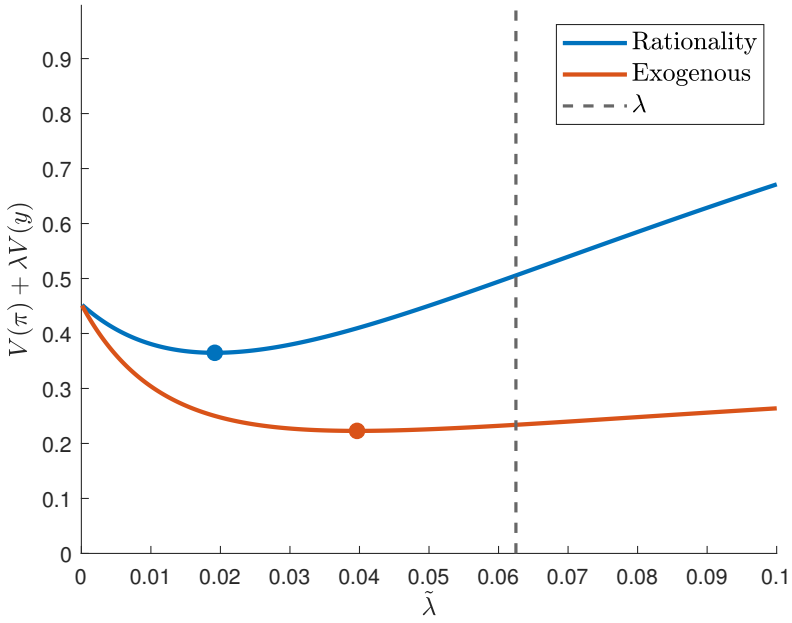
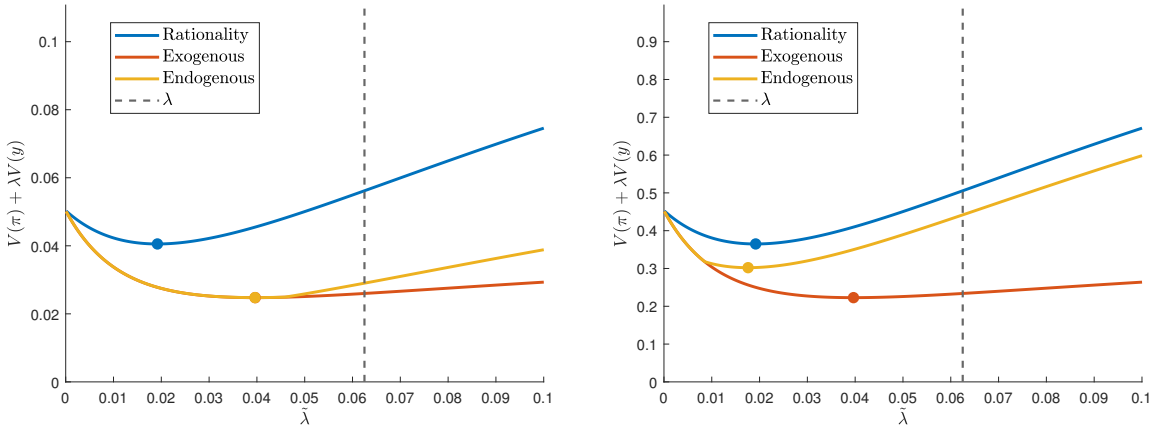


Figure 1.5: Losses under Different $\tilde{\lambda}$: Rationality and Exogenous Attention

banker places a relatively high weight on output versus inflation stabilization, and consequently allows a considerable amount of inflation volatility, that firms pay attention. As such, the level of $\tilde{\vartheta}$ that minimizes welfare losses is the same under endogenous and exogenous discounting.

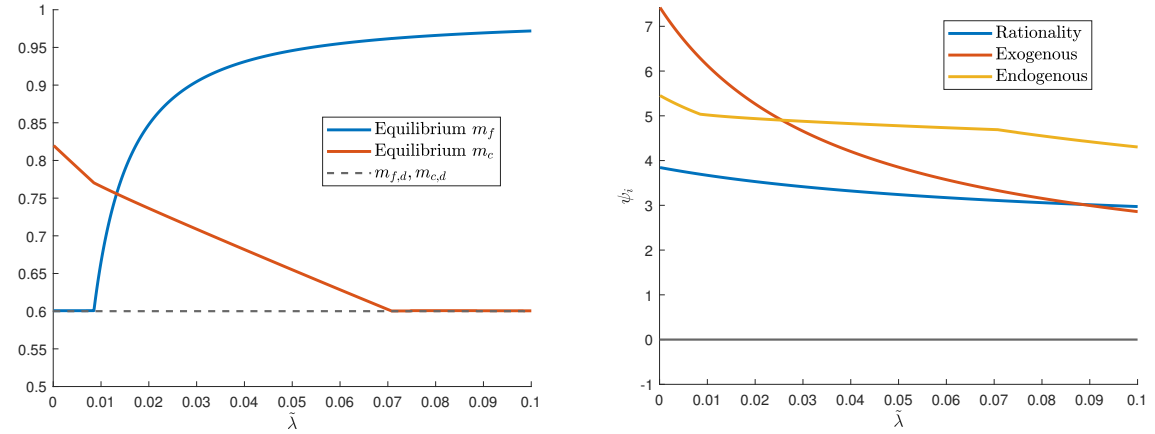
If, on the hand, shocks are large, then it is not necessarily the case that that the loss-minimizing central banker will be less conservative under endogenous discounting than under rationality. Figure 1.6b shows a calibration in which $\tilde{\vartheta}$ may be *lower* than in the rationality case, a qualitative difference compared to the exogenous discounting model. Why does endogenous discounting call the appointment of such a conservative central banker when shocks a large? The reason is illustrated in figure 1.7a. More conservative central bankers focus more on inflation stabilization, which reduces inflation volatility and hence firms' incentives to pay attention. A fall in firm attention then improves the tradeoff between output and inflation, allowing for smaller welfare losses.

As in the exogenous discounting model, consumer attention does not affect the optimal output-inflation tradeoff and so does not affect welfare losses, but it does affect implementation of the optimal policy. In particular, a more conservative central banker responds more aggressively to cost-push shocks, which increases real rate volatility and consequently raises consumer attention. In turn, this increases the effect of persistent changes in real interest rates on output. This effect then attenuates the required increase in real real variability needed to achieve greater stabilization of inflation. As such, a more conservative central banker will still respond more vigorously to cost-push shocks with the nominal rate, but this effect is less pronounced than under exogenous discounting.



(a) Small Shocks **(b)** Large Shocks

Figure 1.6: Losses under Different $\tilde{\lambda}$: Endogenous Attention



(a) Firm Attention **(b)** Response of i_t to η_t

Figure 1.7: Attention and Response of Nominal Rates to Cost-Push Shocks under Discretion in a Calibration with Larger Shocks

This effect is illustrated in figure 1.7b.

1.5.3 Optimal Policy under Commitment

When the central bank can commit to a future path of interest rates, it is well known that the optimal policy response to a cost-push shock under rationality involves bringing prices back to their pre-shock level. By committing to generate an undershoot in inflation relative to the target in the future, the central bank reduces the overshoot in inflation when the shock hits. Under exogenous discounting, Gabaix shows that the optimal policy response instead involves stabilizing prices above the level that they had stood before the shock. When firms discount the future, committing to ensure that inflation falls below target in the future has a smaller effect on the level of inflation today, and so is less worthwhile. An example is shown in figure 1.8.

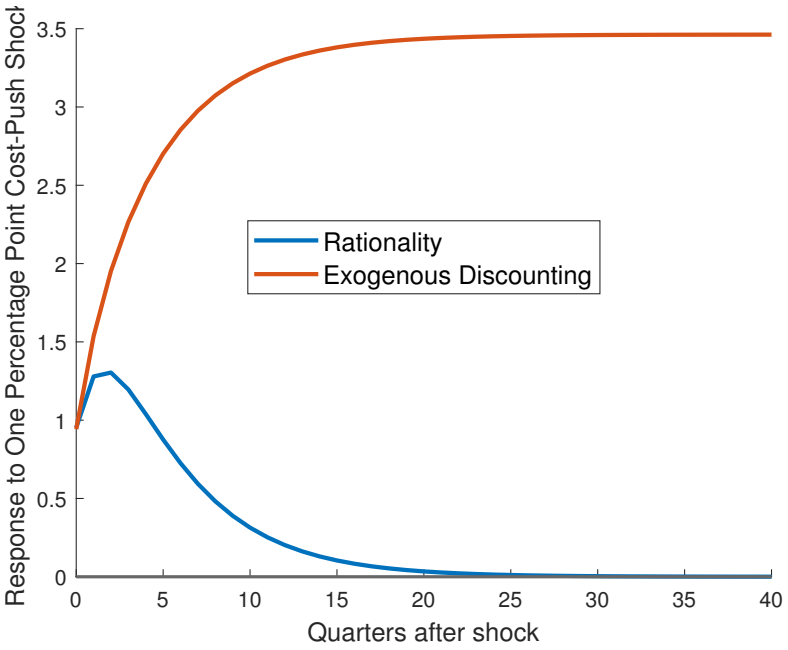


Figure 1.8: Example of Optimal Response of Price Level to a 1pp Cost-Push Shock with $\rho_\eta = 0.8$

The central bank faces a more interesting tradeoff when attention is endogenous. If the central bank is able to commit to a particular policy regime, then it is able to influence the attention level. However, it is now not obvious how the central bank would ideally like to influence the level of attention. As in the discretion case, higher firm attention (M_f) amplifies the effect of persistent cost-push shocks on inflation, which gives a worse static tradeoff between π_t and y_t in the Phillips curve, and so increases losses. On the other hand, the power of committing to a future disinflation after a cost-push shock increases when attention to the future is high. This gives a better *dynamic* tradeoff between inflation today (π_t) and future inflation (π_{t+h}). If shocks are highly persistent, then the amplification channel is very powerful, meaning that higher attention increases losses. In this case, the central bank is incentivised to commit to a policy that reduces firm attention. On the other hand, if shocks are entirely transitory, then the amplification channel disappears, and so higher attention reduces losses. As such, when shocks are transitory the central bank can reduce losses by commitment to a policy regime that increases the attention level.

Persistent Shocks

Figure 1.9 shows an example with a persistent shock, with $\rho_\eta = 0.8$. The path for the price level that is optimal under exogenous discounting may prove to be highly suboptimal when discounting is endogenous. The reason is that this path for the price level implies that inflation remains well above target for an extended period of time; the effects of the

cost-push shock are consequently very persistent. As a result of this high persistence, firms are strongly incentivized to pay attention to the future. If the variability of cost-push shocks is large enough, this may result in a very high attention level. If the attention level is high, then the amplification channel implies a poor tradeoff between inflation and output, which worsens welfare losses. This point emphasizes that policy prescriptions which are optimal when discounting is endogenous may be highly suboptimal once the potential endogeneity of discounting is taken into account.

Likewise, the price level path that is optimal under rationality also may result in high attention and a worse static tradeoff between output and inflation. The reason for this is that this path for the price level also implies inflation remaining away from target for an extended period of time, albeit in this case inflation is *below* rather than *above* target. If shocks are sufficiently large, then this may also result in an endogenous increase in attention.

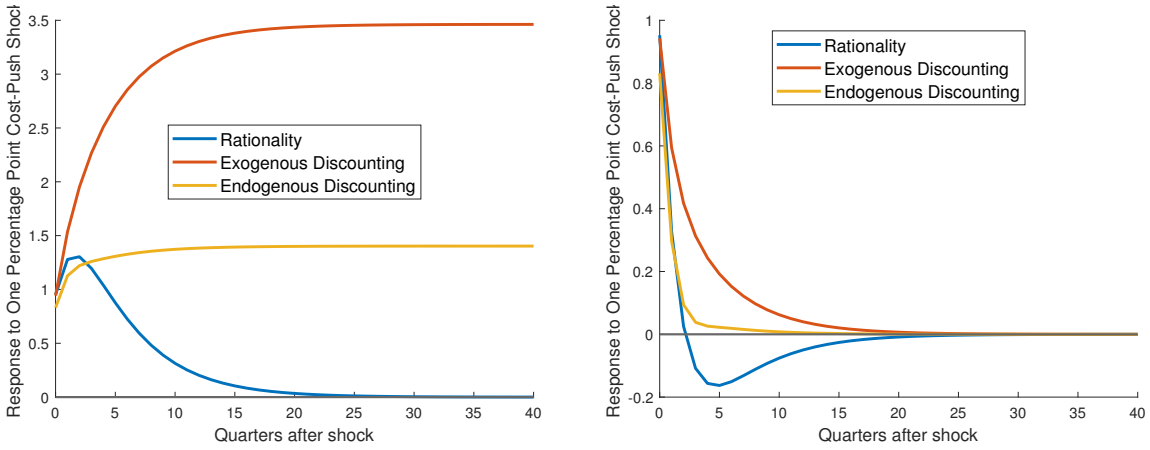
In this example, the optimal path under endogenous attention is one which returns inflation close to the target very quickly after the shock hits. By ensuring that shocks do not have persistent effects on the inflation rate, the central bank ensures that incentives to pay attention to the future are low. This reduces firm attention, and improves the static tradeoff between inflation and output, resulting in smaller welfare losses.

Transitory Shocks

Figure 1.10 shows an example with a transitory shock, with $\rho_\eta = 0$. Because the shock is not persistent, the amplification channel is shut down entirely, and so necessarily the forward-guidance channel will dominate. As such, welfare losses are reduced when M_f is higher, as the dynamic tradeoff between inflation today and inflation tomorrow is improved.

In this case, the optimal path for the price level changes qualitatively from both the rationality and exogenous discounting cases. Under rationality, it is optimal to return the price level to wherever it stood before the cost-push shock hit. Under exogenous discounting, the effect of future inflation on inflation today is smaller. As such, it is less worthwhile for the central bank to commit to ensuring that inflation undershoots the target in the future. The central bank consequently optimally generates a smaller undershoot of inflation relative to the target in the periods after the shock hits. The price level is consequently always stabilized at a higher level than where it stood before the shock hit.

Under endogenous discounting, this conclusion changes qualitatively. In particular, as figure 1.10 shows, for some calibrations it may be optimal for the central bank to stabilize the price level at a *lower* level than where it stood before the shock hit. Why



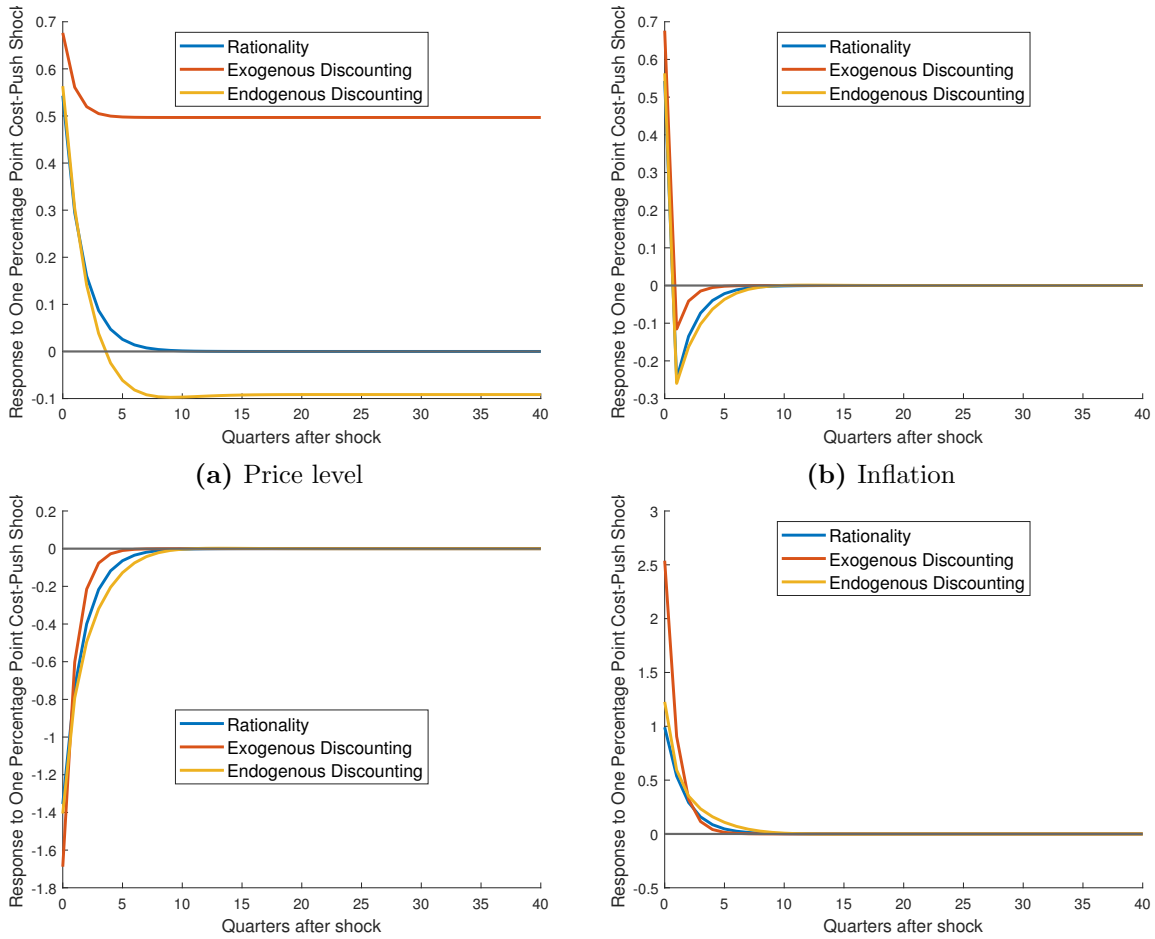
(a) Price level

(b) Inflation

(c) Output

(d) Nominal rate

Figure 1.9: Optimal Responses to 1ppt Cost-Push Shock under Commitment with $\rho_\eta = 0.8$



(a) Price level **(b)** Inflation

(c) Output **(d)** Nominal rate

Figure 1.10: Optimal Responses to 1ppt Cost-Push Shock under Commitment with $\rho_\eta = 0$

does this result emerge? By generating a more aggressive undershoot of inflation in the periods after the shock, the central bank increases firm attention. This effect results in an improved dynamic tradeoff; the undershoot in inflation in the periods after the shock hits ($t \geq 1$) is more effective in bringing down inflation in the period that the shock hits ($t = 0$).

Put differently, if a central bank were to try to pursue the commitment policy and drive down inflation today through the expectations channel, the commitment would have to be one that was large enough to get firms' attention and cause them to respond. Promising that inflation will undershoot the target in the future is unproductive if the undershoot is small enough that firms neglect that information today.

1.6 Identification

I now turn to an empirical assessment of the model. As a precursor to estimating the model, I first check the strength of identification under endogenous cognitive discounting.

Identification is typically a concern in models with cognitive discount parameters; Gabaix (2020) shows that in a single equation example with an unobserved forcing variable attention parameters may not be identified, because they have the same effect on the observed moments as the volatility of the forcing variable.[30] Andrade et al. (2019) also note the possibility of weak identification in a systems of equations approach. My analysis of identification differs from theirs in three respects: (i) I consider a setting with three observables rather than two;[31] (ii) I assume firms and consumers have different cognitive discount factors;[32] and (iii) I also consider endogenous attention.

I assess the strength of identification using a method due to Andrews and Mikusheva (2014).[33] They note that under strong identification both the quadratic variation of the score vector and the negative of the Hessian of the log likelihood converge to the theoretical Fisher information matrix. Under weak identification, only the former does so. One can then assess the strength of identification using the scaled difference between the two:

$$M = J_T^{-\frac{1}{2}}(I_T - J_T)J_T^{-\frac{1}{2}}$$

Here J_T is the quadratic variation of the score vector, and I_T is the negative Hessian. If identification is weak, the largest eigenvalue of M in absolute value, denoted $||M||$, becomes large.

Under endogenous attention, the econometrician estimates:

$$\boldsymbol{\chi} = \begin{bmatrix} \phi_\pi & \phi_y & \theta & \gamma & \rho_\eta & \rho_z & \rho_r & \sigma_\eta^e & \sigma_z^e & \sigma_v^e \end{bmatrix}$$

The econometrician is assumed to know the structure of the attention problem, and so can calculate the implied M_f and M_c for any parameter vector. Under exogenous attention,

[30] See Gabaix (2020) p.15.

[31] In their section 3 where they analyse identification, Andrade et al. (2019) assume only the output gap and inflation are observable, although their estimated model features three observables.

[32] In the maximum likelihood estimation of the behavioural DSGE model in their paper Andrade et al. (2019) assume that the same discount factor is shared by firms and consumers, although they do also conduct separate single equation estimations of the IS and Phillips curves. While imposing the restriction that firms and consumers have the same cognitive discount factor when estimating a behavioral DSGE model may strengthen identification, the microfoundation shows that these cognitive discount factors are driven by rather different factors; for the consumer problem, real interest rate volatility is the key driver, whereas for firms it is marginal cost and inflation volatility. Indeed, my results in the next section suggest that the two may differ substantially.

[33] The estimation procedure for the endogenous attention model does not use Dynare, because for each possible parameter vector one must solve for the equilibrium choices of attention. That said, one can estimate the exogenous attention model in Dynare. In that case, one could instead use the identification tests available in that platform to check the strength of identification. Given the risks of weak identification outlined in this section, using these tests would be worthwhile for econometricians estimating such models.

the econometrician must estimate M_f and M_c.

I set ϕ_π set to 1.2 and ϕ_y to 0.2 to ensure determinacy.[34] ρ_r is set to 0.7.[35] I simulate samples of 100 quarters for a range of values of ρ_z and ρ_η ranging from 0.05 to 0.95. The reason for assessing different values of the autocorrelation parameters is that, as noted in Andrews and Mikusheva (2014), points of identification failure in new Keynesian models are often associated with equality of the autocorrelation parameters. Future research could examine how the strength of identification varies with other model parameters.

The median $||M||$ under exogenous and endogenous attention is shown in figure 1.11 (with a logarithmic scale on the y-axis). Under exogenous attention, the median $||M||$ is exceptionally high where $\rho_z = \rho_\eta$ or $\rho_z = \rho_r$, suggesting point identification failure. Around these points, the measure remains large, indicating that identification is very weak when these autocorrelation parameters are close even if they are not exactly equal.[36] Under endogenous attention, the median $||M||$ does not become large where $\rho_z = \rho_\eta$ or $\rho_z = \rho_r$, suggesting that endogenizing the discount factors mitigates this identification issue. Even away from the points where $\rho_z = \rho_\eta$ or $\rho_z = \rho_r$, identification is often notably stronger under endogenous rather than exogenous cognitive discounting, even if identification does become somewhat less strong in the endogenous attention model as attention starts to move away from the default.[37]

The application to the Great Inflation and Great Moderation uses a Bayesian approach. If the likelihood is weakly identified, then the posterior is determined largely by the priors. Under exogenous attention it is difficult to ascertain whether, for example, a high variance of inflation results from high attention or volatility in the cost-push shock; there is a range of combinations of M_f and σ_η^ε that match the data. Which combination is more likely per the posterior depends largely on the priors. Under endogenous attention, a combination of M_f and σ_η^ε is selected which is consistent with optimising behaviour.[38]

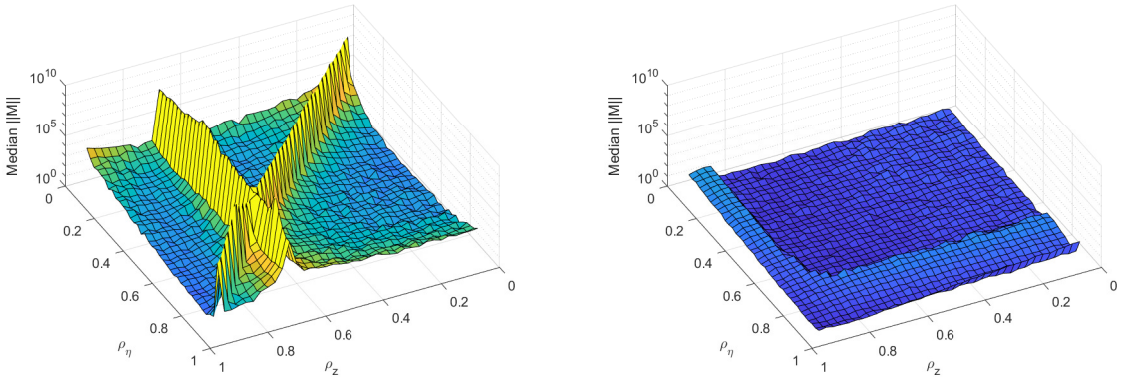
[34] I choose standard parameterizations for γ and θ, which are set to 2.0 and 0.8 respectively. For the shock variances, I set the variance of demand shock to be larger than the cost-push or monetary policy shocks, taking into account the results I the next section. I set the shock variances sufficiently large so that m_f rises above the default level of attention when the persistence of cost-push shocks is high, given that the results shown in the next section have m_f rising above the default level in the Great Inflation period.

[35] This is roughly between the values estimated for the Great Inflation and Great Moderation samples in the next section.

[36] This possibility is noted by Andrews and Mikusheva (2014). In their example, weak identification arises even if the difference between autocorrelation parameters is as large as 0.3.

[37] As figure 1.11 shows, identification weakens more notably under endogenous attention for some very high values of ρ_z and ρ_η. However, these combinations of ρ_z and ρ_η also typically see weaker identification under exogenous cognitive discounting.

[38] For procedures for assessing identification specific to a Bayesian context, see Koop et al. (2013). Simply comparing priors and posteriors is not sufficient to check identification. As Koop et al. (2013) argue, requiring the parameters to be consistent with determinacy or indeterminacy imposes a joint restriction that may lead posteriors to differ from priors even if they are unidentified. This matters because I estimate the model imposing either determinacy or indeterminacy.



(a) Exogenous Attention **(b)** Endogenous Attention

Figure 1.11: Strength of Identification under Exogenous and Endogenous Attention.

An alternative approach is to use methods robust to weak identification. Andrade et al. (2019) conduct both a single equation GMM estimation of the IS and Phillips curves in a behavioural model, and a system of equations maximum likelihood estimation. They use methods proposed by Andrews and Mikusheva (2015) and Andrews (2018) to generate robust confidence sets for the attention parameters. However, their robust confidence intervals are very wide. In their GMM estimation, for example, the interval for M_f ranges from 0.14 to 0.95. Values at the lower and upper ends of this interval have vastly different implications for indeterminacy. Moreover, large regions within this intervals are inconsistent with the model's microfoundations. Endogenising attention avoids this issue by finding a combination of attention and shock volatilities consistent with both the data and the microfoundations.

Identification under endogenous attention relies on correct specification of the attention problem. One must make assumptions on the form of the attention cost, the default level of attention and the attention cost. If one estimates these parameters, attention can vary independently of the shock variances, which is the issue that affects the exogenous case. In section 1.8, I present a method to validate the choice of attention costs, using the time variation in informational rigidities in expectations data uncovered by Coibion and Gorodnichenko (2015a).

1.7 Estimation

I now estimate the model on the Great Inflation and Great Moderation periods. This allows me to show empirical relevance of endogenous discounting for identification, determinacy, and the effects of policy.

Model

I estimate the model given by (1.22), (1.23), and (1.11). I assume that β and ϕ are known, and take values of 0.99 and 1.0 respectively. Attention is considered endogenous, and I once again assume values for default attention and the attention costs. Under determinacy, I estimate the parameters:

$$\boldsymbol{\chi} = \begin{bmatrix} \phi_\pi & \phi_y & \theta & \gamma & \rho_\eta & \rho_z & \rho_r & \sigma_\eta^\varepsilon & \sigma_z^\varepsilon & \sigma_v^\varepsilon \end{bmatrix} \tag{1.48}$$

Under indeterminacy, I impose the the continuity solution, to avoid issues of weak identification.[39] I then estimate the parameter vector above, augmented with σ_ζ.

As above, I specify default levels of attention and attention costs. The attention parameters have a considerable impact on the results. By specifying an arbitrarily high cost of attention, and a low default level, one could ensure attention is sufficiently low such that the model is always determinate, resulting in a arbitrarily small marginal likelihood for the indeterminate model. Equally, if one sets the cost of attention arbitrarily low, the model reduces to the rational expectations model, making indeterminacy very likely in the Great Inflation sub-sample because the interest rate response to inflation is weak.

I calibrate default attention to $m_{f,d} = m_{c,d} = 0.85$. This corresponds to the attention level M_c suggested by Gabaix (2020) for calibrating the behavioural new Keynesian model. As Gabaix notes, this implies discounting of c.50% at a one-year horizon. Gabaix notes that this value is broadly in line with studies estimating the weight on forward looking terms in the Phillips and IS curves, such as Galı and Gertler (1999) or Fuhrer and Rudebusch (2004).

For consumers, I assume an attention cost of $k_c = 4.5$, which is the value estimated by Ganong and Noel (2017).[40] To my knowledge, this is the only estimate for scale-free attention costs; no estimates exist for firms. One would expect costs for firms to be lower. My baseline specification assumes $k_f = 1.5$. I calculate that if paying attention to a variable changes pricing decisions by 1.5% on average, then doing so increases profits by c.0.1% of sales. This accords with Coibion et al. (2020) finding that firms lose profits of roughly c.$0.1 - 0.3$% of sales as a result of inaccurate inflation perceptions. For details of my calculation, see the Appendix. To show the robustness of the results to this assumption, I also show an alternative specification in which k_f is half the baseline level at 0.75.

[39] Nonetheless, I note that this may disadvantage the indeterminate model in the model comparison conducted below. Future research could assess the sensitivity of the results to allowing response to fundamental shocks to vary.

[40] See p.24 of Ganong and Noel (2017). In their model, they interpret this parameter as "the largest possible income shock for which the agent would not cut spending in advance at all".

I impose the requirement that the equilibrium found must be attainable. When calculating the unconditional covariance matrix for the variables, which I use to initialize the Kalman filter, I suppose that the first period of the sample is the first period that the economy was in the indeterminate state. In theory, one should allow attention to drift upwards over time. In practice, it is then technically challenging to calculate the likelihood. I instead hold the level of attention fixed over time in the indeterminate state. Moreover, I assume that if an indeterminate equilibrium exists in the first period, then it exists in all subsequent periods.

One has to make an assumption about the source of IS curve shocks to calculate consumer attention. Justiniano and Primiceri (2008) find that discount factor shocks account for a large part of consumption variance. For simplicity, I therefore assume all IS curve shocks are driven by the discount factor.

Data

All data are from Federal Reserve Economic Data. The baseline specification follows Coibion and Gorodnichenko (2015b) in using unemployment (UNRATE) as the forcing variable in the Phillips curve. I filter the unemployment rate using the Hodrick-Prescott (HP) filter to obtain the unemployment gap, and use the negative of this as an output gap measure.[41] As a robustness check, I follow Lubik and Schorfheide (2004) in using the HP filter to de-trend the log of real GDP per capita (A939RX0Q048SBEA). Deviations from trend are multiplied by 100 to obtain a percentage point output gap measure. I include this specification because one might expect the choice of forcing variable to affect the Phillips curve slope and hence the estimated θ, which in turn affects the determinacy region.

For inflation, I use the quarter-on-quarter percentage change in the GDP deflator (GDPDEF). For the nominal interest rate I use the effective Federal Funds rate (FEDFUNDS), which I convert to a quarterly rate. I de-mean the nominal variables in each sub-sample.[42] I de-mean nominal variables, rather than HP filtering. Canova and Ferroni (2011) note that "while real variables typically show long run drifts, nominal variables just display low frequency fluctuations". Removing such fluctuations could be a concern when estimating the model under indeterminacy. As highlighted in above, sunspot shocks can have very persistent effects on inflation; de-trending could remove this low frequency variation.

For the Great Inflation period, I use data from 1960:I to 1979:II, and for the Great

[41] For the HP filter, I set $\lambda = 1600$, as standard in the literature. To avoid the end point problems associated with the HP filter, I de-trend a longer time series and take sub-samples of the longer series. As such, the first and last four observations of the filtered data are not used in any sub-sample.

[42] Lubik and Schorfheide (2004) instead include additional parameters for their average levels.

Moderation period I use 1984:I to 2007:IV. I also report results from 1990:I to 2007:IV, motivated by Fernández-Villaverde et al.'s (2010) finding that the policy response to inflation in this period was weak.

Priors

Table 1.1: Priors

Parameter	Distribution	Mean	St. Dev.	90 pct. interval
ϕ_π	Γ	1.10	0.50	[0.43,2.03]
ϕ_y	Γ	0.25	0.15	[0.06,0.54]
θ	B	0.70	0.075	[0.57,0.82]
γ	Γ	2.00	0.50	[1.25,2.89]
ρ_η	B	0.70	0.10	[0.52,0.85]
ρ_z	B	0.70	0.10	[0.52,0.85]
ρ_r	B	0.50	0.20	[0.17,0.83]
σ_η^ε	Γ^{-1}	0.30	0.50	[0.08,0.80]
σ_z^ε	Γ^{-1}	1.00	1.00	[0.32,2.43]
σ_v^ε	Γ^{-1}	0.30	0.50	[0.08,0.80]
σ_ζ	Γ^{-1}	0.30	0.50	[0.08,0.80]

For the policy parameters, shock autocorrelations, and the coefficient of relative risk aversion, I set the mean and standard deviations for each prior in line with Lubik and Schorfheide (2004).[43] Lubik and Schorfheide (2004) estimate the Phillips curve slope directly, so I cannot follow their prior for θ. I use a beta distribution to constrain θ to the interval (0,1). I centre the distribution at 0.7, giving an average price duration of 10 months, close to the average duration found in Nakamura and Steinsson (2008). The 90% credibility interval implies price duration between 7 months and 17 months.[44] For shock volatilities, I use similar means to Lubik and Schorfheide (2004), but larger standard deviations; these parameters matter for attention and hence determinacy, so I prefer to give more weight to the data.

Methodology

My process for estimating the model broadly follows techniques outlined in Lubik and Schorfheide (2004) and Herbst and Schorfheide (2016), adapting these methods to allow for endogenous attention. The key difference is that for each set of parameters drawn, the equilibrium level of attention must be calculated in order to evaluate the log-likelihood. A full description of the algorithm is given in the Appendix.

[43] My prior for γ corresponds to Lubik and Schorfheide's (2004) prior for the inverse of the IS curve slope.

[44] I calculate credibility intervals using the 5th and 95th percentiles of the distribution. Note that this differs from Lubik and Schorfheide (2004), who calculate the shortest such interval.

Results

Table 1.2 compares the log marginal likelihood of determinacy and indeterminacy for each sub-sample in each specification, and shows the evidence against indeterminacy according to Kass and Raftery's (1995) descriptive scale. Table 1.3 shows full results for the specification. Full results for the alternative specifications are shown in the Appendix.

Comparison of the marginal likelihood suggests that, for the Great Inflation period, the evidence in favour of determinacy is "very strong" in the baseline specification, despite the mean of ϕ_π lying well below unity. This result also holds in the specification using the output gap. Even in the low cost of attention specification, the evidence against indeterminacy is "strong" on the Kass and Raftery (1995) scale. As such, the findings represent more conclusive evidence against the good policy hypothesis than is found in the exogenous attention estimation of Ilabaca et al. (2020).

Why does this result emerge? First, even under endogenous attention, the estimated Phillips curve is rather low. This generally pushes the economy towards indeterminacy. Second, for consumers, the Ganong and Noel (2017) estimate of k_c implies that in the present specification consumer attention always remains at its default level. In this model consumers care about real interest rates. Fluctuations in the real rate are smaller than nominal fluctuations in this period, because ϕ_π is close to the autocorrelation parameter for the cost-push and demand shocks. With a fairly high consumer cost of attention, the volatility in the real rate is insufficient to move attention above the default level. It might be that a richer model with medium-scale features might cause consumers to pay more attention. Third, ϕ_y matters more for determinacy under endogenous attention than under rational expectations. In the unemployment gap specification, ϕ_y tends to be rather high.

As such, one would require an exceptionally high m_f to generate indeterminacy. The results do suggest that m_f was well above its default level during the Great Inflation period in response to the high variance and persistence of cost-push shocks. However, firms remain some way from rational expectations. As I discuss below, this pattern is externally validated by empirical expectations data from Coibion and Gorodnichenko (2015a).

The reader might wonder why the marginal likelihoods of the determinate and indeterminate models are so different, despite the estimates for many parameters seeming to be close. However, the estimates imply rather different inflation volatility and persistence. This can be seen from the high m_f obtained under indeterminacy. Indeed, the model is pushed towards generating high inflation volatility and persistence because m_f must be very high for system to be indeterminate. The reason is the difference in ρ_η. When ϕ_π is low, small differences in ρ_η dramatically change the response of inflation to cost-push

shocks, and this more than offsets the lower estimate for σ_η^ε under indeterminacy.

For the Great Moderation, indeterminacy is strongly rejected in all specifications. For the 1984:I to 2007:IV sub-sample, I obtain a mean estimate for ϕ_π above unity in all specifications, in line with the literature. Like Fernández-Villaverde et al. (2010), however, I find that the point estimate for ϕ_π falls below unity for the period after 1990, albeit the 90% credibility interval is very wide. This result holds across all specifications. Despite this, the probability of indeterminacy is negligible even in the lowest attention cost specification. Indeterminacy is particularly unlikely for this period because (i) the cost-push shock is less volatile and persistent than in the Great Inflation period, lowering firm attention, and (ii) the estimate for θ is even greater. Notice how extreme the estimate must be (ϕ_π of 0.20) to generate indeterminacy; this leads to the very low marginal likelihoood.

Note that this model assumes that persistence in inflation arises from the persistence of cost-push shocks, and the results outlined above show that this matters for determinacy. Other approaches assume that this arises from indexation.[45] Benati (2008) finds that the extent of indexation has changed substantially over time. Future research might model the degree of backward looking behaviour endogenously as a behavioural phenomenon, perhaps using similar techniques to those outlined in this paper.

Table 1.2: Log Marginal Likelihood Comparison

	Determinacy	Indeterminacy	Evidence
Baseline			
Great Inflation	-53.0	-67.8	Very Strong
Great Moderation (a)	84.4	23.6	Very Strong
Great Moderation (b)	64.0	20.2	Very Strong
Low Attn. Cost			
Great Inflation	-53.9	-57.2	Strong
Great Moderation (a)	82.4	43.7	Very Strong
Great Moderation (b)	62.3	35.9	Very Strong
Output Gap			
Great Inflation	-128.4	-141.7	Very Strong
Great Moderation (a)	15.6	-42.3	Very Strong
Great Moderation (b)	17.3	-25.4	Very Strong

"Evidence" refers to evidence against indeterminacy given by the Bayes factor, assessed using the scale given by Kass and Raftery (1995). Great Moderation (a) refers to the sample beginning in 1984:I, and (b) to the sample beginning 1990:I.

Counterfactuals

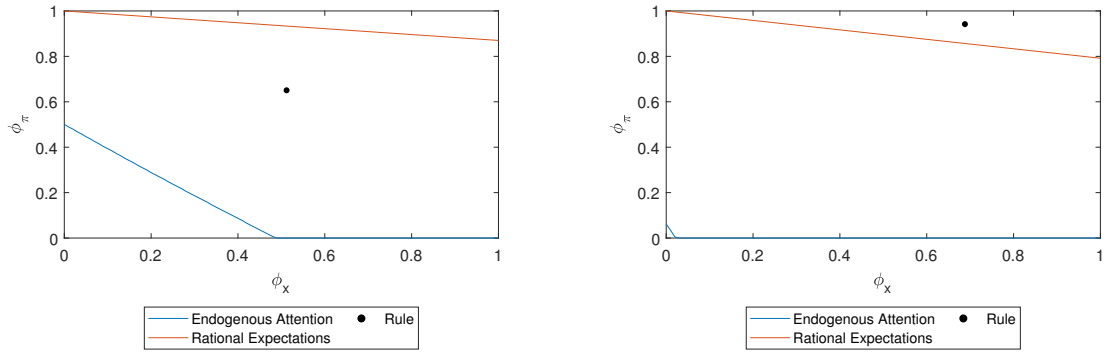
I now consider how the economy would have behaved under counterfactual policy rules. I begin by considering determinacy regions. Figure 1.12 shows the determinacy regions

[45] See Christiano et al. (2005), for example.

Table 1.3: Posterior Distribution - Baseline Specification ($k_f = 1.5$)

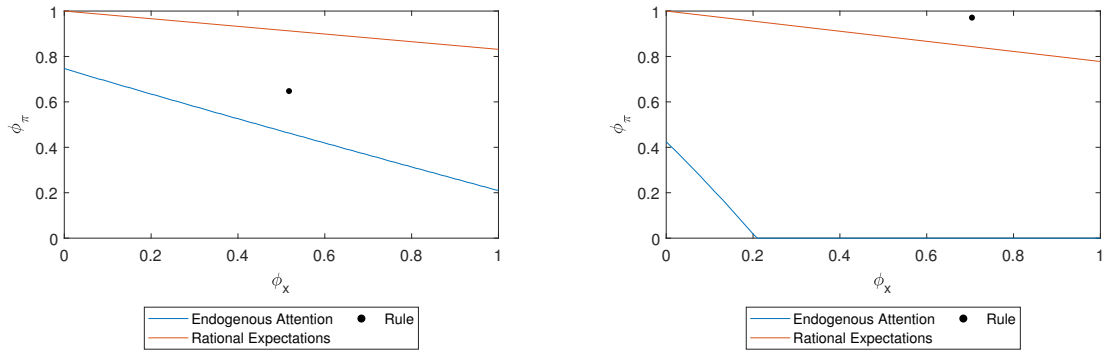
	Great Inflation (1960:I to 1979:II)				Great Moderation (a) (1984:I to 2007:IV)				Great Moderation (b) (1990:I to 2007:IV)			
	Determinacy		Indeterminacy		Determinacy		Indeterminacy		Determinacy		Indeterminacy	
Parameter	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval
ϕ_π	0.65	[0.55,0.76]	0.53	[0.52,0.62]	1.55	[1.00,2.17]	0.19	[0.08,0.32]	0.94	[0.49,1.47]	0.20	[0.08,0.33]
ϕ_y	0.51	[0.39,0.65]	0.48	[0.37,0.60]	0.73	[0.52,0.95]	0.64	[0.47,0.85]	0.69	[0.51,0.88]	0.49	[0.33,0.66]
θ	0.87	[0.83,0.92]	0.86	[0.82,0.89]	0.90	[0.87,0.93]	0.88	[0.86,0.90]	0.91	[0.88,0.93]	0.87	[0.79,0.90]
γ	2.90	[2.13,3.78]	3.42	[2.57,4.39]	4.02	[3.10,5.04]	4.91	[3.78,6.23]	3.49	[2.63,4.46]	4.22	[2.75,5.44]
ρ_η	0.74	[0.67,0.81]	0.86	[0.81,0.90]	0.66	[0.57,0.76]	0.86	[0.76,0.92]	0.63	[0.53,0.73]	0.84	[0.71,0.90]
ρ_z	0.80	[0.73,0.86]	0.76	[0.69,0.83]	0.91	[0.88,0.94]	0.88	[0.82,0.94]	0.90	[0.85,0.94]	0.86	[0.79,0.92]
ρ_r	0.56	[0.45,0.66]	0.56	[0.46,0.65]	0.84	[0.80,0.87]	0.84	[0.80,0.87]	0.83	[0.78,0.87]	0.81	[0.76,0.85]
σ_η^ε	0.12	[0.08,0.17]	0.08	[0.05,0.11]	0.08	[0.06,0.090]	0.04	[0.03,0.07]	0.08	[0.06,0.10]	0.06	[0.03,0.09]
σ_z^ε	0.48	[0.36,0.63]	0.62	[0.47,0.81]	0.28	[0.22,0.35]	0.39	[0.25,0.53]	0.28	[0.22,0.36]	0.38	[0.26,0.51]
σ_v^ε	0.16	[0.14,0.19]	0.16	[0.14,0.19]	0.12	[0.10,0.14]	0.12	[0.11,0.14]	0.10	[0.09,0.12]	0.10	[0.09,0.12]
σ_ζ			0.12	[0.06,0.20]			0.08	[0.05,0.13]			0.09	[0.05,0.15]
m_f	0.93	[0.85,0.98]	0.98	[0.96,0.99]	0.85	[0.85,0.89]	0.97	[0.95,0.98]	0.85	[0.85,0.88]	0.95	[0.85,0.98]
m_c	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]
Log ML		-53.0		-67.8		84.4		23.6		64.0		20.2

generated by the structural and shock process parameters for the Great Inflation and post-1990 Great Moderation periods under the determinate estimate in the baseline specification.[46] Figure 1.13 shows results for the low cost of attention specification. I use the mean as a point estimate.[47]



(a) Great Inflation **(b)** Great Moderation

Figure 1.12: Counterfactual Determinacy Regions: Baseline Specification



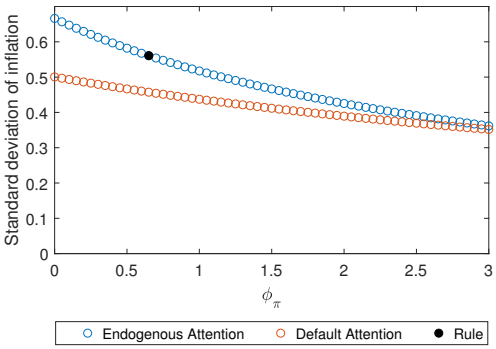
(a) Great Inflation **(b)** Great Moderation

Figure 1.13: Counterfactual Determinacy Regions: Low Attention Cost Specification

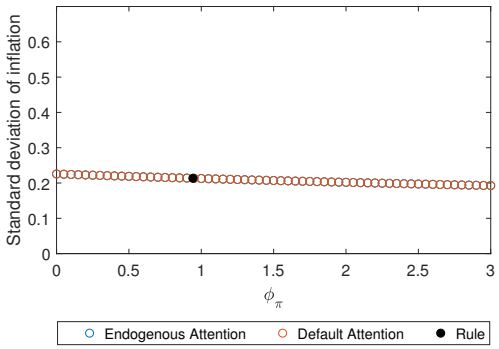
Although the estimated rule for the Great Inflation leads to indeterminacy under rational expectations, it lies well within the determinate region in the baseline specification. Figure 1.12a illustrates how unlikely is indeterminacy in the endogenous attention model; even for the Great Inflation period, the indeterminacy region is small. Under the low attention cost, the indeterminacy region expands, but the estimated rule still lies well within the determinate region. For the post-1990 sub-sample, the point estimate for ϕ_π lies below unity across specifications. The point estimate for the rule still lies within the rational expectations determinacy region because of the high ϕ_y, but 39% of draws result in indeterminacy under rational expectations. However, under endogenous attention, the indeterminacy region is very small; volatility and hence attention were so low during this period that one would have to set ϕ_π and ϕ_y very close to 0 to allow indeterminacy.

[46] Note that I hold ρ_r constant in either case and just vary ϕ_π and ϕ_y.

[47] Using the posterior mean as a point estimate minimises a quadratic loss function; see Herbst and Schorfheide (2016).

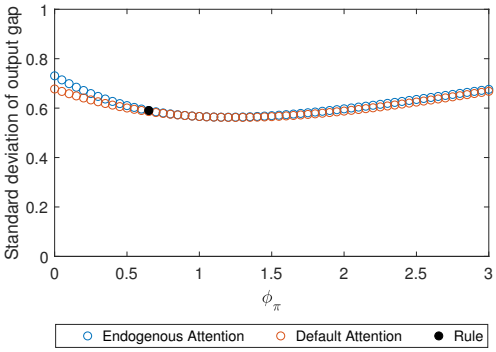


(a) Great Inflation

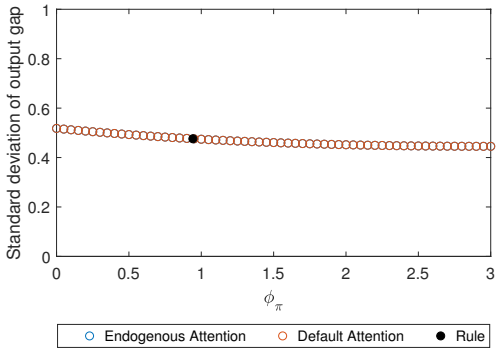


(b) Great Moderation

Figure 1.14: Counterfactual Inflation Volatility

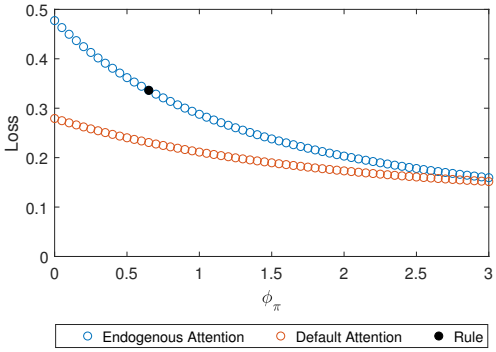


(a) Great Inflation

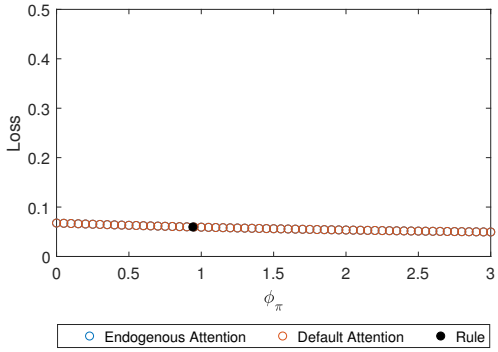


(b) Great Moderation

Figure 1.15: Counterfactual Output Gap Volatility



(a) Great Inflation



(b) Great Moderation

Figure 1.16: Counterfactual Central Bank Losses

I now conduct a counterfactual experiment by varying ϕ_π, holding the other parameters fixed. I show how the standard deviation of inflation and the output gap would have changed under different policy rules. I also show central bank losses, which I calculate as:

$$L = V(\pi) + \vartheta V(x)$$

I use $\vartheta = 0.0625$.[48] In each figure, the black dot shows the volatility implied by the estimated rule. The blue series shows how volatility changes with ϕ_π under endogenous attention. The red series shows, for comparison, how volatility changes if attention is fixed at the default. For the Great Moderation attention is always at its default, so only the red series is visible. Table 1.4 compares volatility between the estimated ϕ_π for each period and $\phi_\pi = 2.0$.[49] To make a fair comparison across periods, as estimated policy rules differ, I also compare $\phi_\pi = 2.0$ and $\phi_\pi = 0.5$.

For the Great Inflation, inflation volatility and losses are high, and fall steeply as ϕ_π increases. For the Great Moderation, inflation volatility and losses are lower, and change proportionally less with ϕ_π. It makes sense that the absolute change in inflation volatility is greater for the Great Inflation, but why is it proportionally much larger? First, the estimated parameters imply a higher κ, allowing interest rate changes to have stronger effects on inflation. Second, ρ_η is higher during the Great Inflation, so raising ϕ_π tends to have a larger impact on inflation volatility. Consider the change in inflation volatility as ϕ_π moves from 2.0 to 0.5. In the Great Moderation, where attention is always at its default level, inflation volatility is just 8% higher when ϕ_π is set to 0.5. For the Great Inflation, even fixing attention at its default, the difference is 20%. However, a further effect results from shocks being large enough during the Great Inflation to push m_f above its default. Under endogenous attention, inflation volatility is 37% higher when ϕ_π is set to 0.5 than when it is set to 2.0. That is shown by the divergence of the red and the blue series in figure 1.14a. As noted above, this divergence arises both because attention rises as ϕ_π decreases, and because when attention is higher the effects of ϕ_π on inflation volatility are more powerful. The results suggest these behavioural channels, which are only active when the economy is sufficiently volatile to push attention above the default, are quantitatively large.

One might conclude from these results that there is no role for policy in the Great

[48] As noted in Debortoli et al. (2019), former Janet Yellen has suggested that the Fed places equal weight on annualized inflation and the unemployment gap. Here, π_t is quarter-on-quarter inflation, or one-quarter of annualized inflation. This translates into a weight of 1/16, or 0.0625, on the unemployment gap relative to inflation

[49] A response to inflation of 2.0 is roughly in line with the level estimated by Fernández-Villaverde et al. (2010) for the majority of the 1980s.

Table 1.4: Counterfactual Macroeconomic Volatility and Central Bank Losses

	$\phi_\pi = 0.5$	Est. ϕ_π	$\phi_\pi = 2.0$	$\phi_\pi = 0.5$ vs. $\phi_\pi = 2.0$	Est. ϕ_π vs. $\phi_\pi = 2.0$
Standard Deviation π					
Great Moderation	0.219	0.213	0.202	8%	6%
Great Inflation (default)	0.467	0.457	0.389	20%	17%
Great Inflation	0.582	0.561	0.425	37%	32%
Standard Deviation y					
Great Moderation	0.493	0.476	0.452	9%	5%
Great Inflation (default)	0.601	0.586	0.588	2%	0%
Great Inflation	0.610	0.590	0.597	2%	-1%
Central Bank Loss					
Great Moderation	0.063	0.060	0.054	18%	12%
Great Inflation (default)	0.240	0.231	0.173	38%	33%
Great Inflation	0.362	0.336	0.203	78%	66%

The post-1990 sub-sample is used for the Great Moderation. The estimated ϕ_π is 0.65 for the Great Inflation, and 0.94 for the post-1990 period.

Inflation and Great Moderation. After all, the point estimate for ϕ_π changes little between the Great Inflation and the post-1990 sub-sample. The results in table 1.4 reject this view. Weak stabilization of inflation during the Great Inflation resulted in losses that were 66% higher than they would have been if ϕ_π were set to 2.0. After 1990, the difference is just 12%. The policy rule implemented during the Great Inflation was a poor choice for that time; this was precisely the moment when a strong response to inflation would have been most beneficial. After 1990, the response to inflation may have been nearly as weak as during the Great Inflation. However, this mattered less for inflation volatility.

1.8 External Validation

The paucity of empirical estimates of the costs of attention in the microdata presents a challenge for estimating the endogenous attention model. My analysis shows that endogenizing the attention level is important empirically, and so calls for further empirical work to pin down attention costs. Until there is greater confidence in estimates of attention costs in the microdata, researchers may desire an alternative method for externally validating their results when estimating the endogenous attention model. This section demonstrates such a method, using the Great Inflation and Great Moderation estimation exercise as an example. Specifically, I verify whether the dynamics of firm cognitive discounting that drive my results match empirical data on expectation formation. Coibion and Gorodnichenko (2015a) study Survey of Professional Forecasters (SPF) data. They run regressions of the form:

$$x_{t+h} - F_t x_{t+h} = c + \beta(F_t x_{t+h} - F_{t-1} x_{t+h}) + \varepsilon_{t,h}$$

Here, x_t is some macroeconomic variable, and $F_t x_{t+h}$ denotes the time t forecast at the h quarter horizon. This regression asks whether the forecast revision at time t predicts the ex-post error of the time t forecast. Under rational expectations, there is no relationship between forecast revisions and ex-post errors, and hence β is zero. An estimated β significantly different from zero then constitutes evidence against rational expectations. Coibion and Gorodnichenko refer to β as a measure of "informational rigidity".

Gabaix shows that under cognitive discounting, β should be different from zero. Specifically, at horizon h, β is bounded below by:

$$\underline{\beta}(m,h) = \frac{1-m^h}{m^h}$$

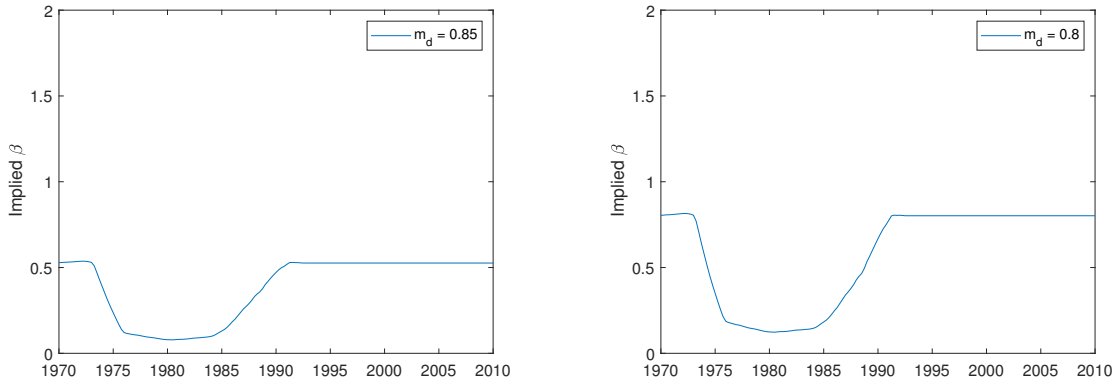
As such, the endogenous cognitive discounting model predicts that in periods of higher macroeconomic volatility, when m is higher, β should be lower.

Coibion and Gorodnichenko do indeed find evidence of this. Pooling forecast revisions and errors across a number of variables in the SPF, and across forecast horizons, they estimate β for each quarter from the beginning of the SPF sample in 1968, and take a smoothed average, shown in their figure 3. For comparison, they show the rolling 5 year standard deviation of US real GDP growth as a proxy for macroeconomic volatility. Informational rigidities are macroeconomic volatility are clearly inversely related; in the late 1970s, when macroeconomic volatility was high, informational rigidities declined. The estimated β then rose again as macroeconomic volatility declined in the late 1980s. While the professional forecasts in the SPF are not entirely comparable to firm forecasts, they should provide a guide to the direction and magnitude of changes in informational rigidities at least for large firms.

To compare their results with the predictions of my model, I estimate the model on rolling 40 quarter periods of data, and find the implied m_f at the posterior mode. I then calculate the implied lower bound on β. I note that Coibion and Gorodnichenko pool across horizons up to four periods ahead. To make things comparable, since $\underline{\beta}$ depends on h, I calculate an approximate lower bound for β in a pooled approach using an average of the lower bound for one to four quarter horizons:

$$\begin{aligned}\underline{\beta}_{pooled}(m) &= \sum_{h=1}^{4} \underline{\beta}(m,h) \\ &= \sum_{h=1}^{4} \frac{1-m^h}{m^h}\end{aligned}$$

Coibion and Gorodnichenko smooth their results, and I do the same by averaging across periods.



(a) $m_{f,d} = 0.85$ **(b)** $m_{f,d} = 0.8$

Figure 1.17: Implied Lower Bound on β

The results are shown in figure 1.17, using two different values of m_f. 1.17a uses $m_{f,d} = 0.85$, as in the baseline sepecification, and 1.17b uses a slightly lower value of 0.8. The degree of informational rigidity predicted by the endogenous cognitive discounting model decreases during the late 1970s, and subsequently increases after the mid-1980s, much as in the Coibion and Gorodnichenko findings. In the model, the lower bound on m_f is given by $m_{f,d}$, which is why the implied $\underline{\beta}_{pooled}$ reaches a maximum in the late 1980s. Nonetheless, the magnitude of the changes in $\underline{\beta}_{pooled}$ are in line with Coibion and Gorodnichenko's estimates of β, particularly for $m_{f,d} = 0.8$. As such, empirical data on expectations provides external validation for the direction and magnitudes of the changes in firms' cognitive discounting over this period, which are the key driver of my results.

1.9 Conclusions

In his review of Gabaix's (2020) behavioural new Keynesian model, Cochrane (2016) argues that future work should take seriously Gabaix's behavioural microfoundation: "because the paper is so important," Cochrane writes, "its foundations matter". The microfoundation implies that attention to future inflation, output, and interest rates rises above its usual levels when the volatility and persistence of macroeconomic variables is high. This prediction appears to be consistent with the co-movement of informational rigidities and macroeconomic volatility uncovered by Coibion and Gorodnichenko (2015a).

Although Gabaix shows how one could calculate the optimal choice of attention to the future for a given set of macroeconomic dynamics, when analyzing the model's implications he assumes that the degree of cognitive discounting is fixed. The seven main predictions of the behavioural new Keynesian model presented in Gabaix (2020) are therefore derived under exogenous discounting.

Differently from Gabaix, I analyzed the implications of a behavioural new Keynesian

model in which the degree of discounting varies endogenously as the microfoundation implies. I demonstrated that four of the key implications of the model with fixed attention change qualitatively when the discounting is instead endogenous:

1. Endogenizing the cognitive discount factor qualitatively changes the model's determinacy condition.

2. Endogenous discounting implies different stability properties at the ELB from the model with exogenous discounting.

3. Under endogenous discounting there may exist equilibria in which the cost of the ELB is very high as well as those in which the cost is low, and AIT can help to rule out the former.

4. Optimal policy changes qualitatively relative to the model with exogenous cognitive discounting.

Future work could investigate whether the remaining three key implications of the model with constant attention also change when discounting is endogenous.

Aside from the novel theoretical implications of my model, I showed that estimating the model with endogenous discounting qualitatively changes the interpretation of the Great Inflation and Great Moderation periods.

These findings highlight the importance of considering the potential endogeneity of behavioural parameters in new Keynesian models. When assessing optimal policy or determinacy, one must consider large changes in the policy regime and hence the macroeconomic environment, meaning that endogenous changes in these parameters are likely to be important.

The model could be extended in two main directions. First, one could incorporate medium-scale features such as sticky wages, investment, consumption habits, and so on. Such features would likely affect firms' and consumers' incentives to pay attention to the future, potentially leading to different strategic interactions and novel policy implications. Second, one could allow for heterogeneity in asset market participation or in costs of paying attention, to understand how heterogeneity affects agents' choices of cognitive discount factor in equilibrium.

My study also highlights the need for further empirical work in this area. The extent to which agents are myopic could be tested by collecting survey data on how agents respond to shocks of different durations. Tracking the degree of myopia over time would then allow one to trace out how the degree of forward-looking behavior varies across different macroeconomic regimes, and compare the results with the predictions of the sparsity-based microfoundation. Given the importance of the behavioural microfoundation in a

theoretical model, testing the microfoundation using microdata is an important avenue for future research.

Appendices to Chapter 1

1.A Proofs and Derivations

1.A.1 Exogenous Attention

I begin by solving the system under exogenous attention, and then use those dynamics to formulate the g mapping. To solve the system, one can simply follow Lubik and Schorfheide (2004). This allows one to solve the system under both determinacy and indeterminacy. Denote by $\boldsymbol{s_t}$ a vector of endogenous variables, augmented by the rational expectations terms $E_t\pi_{t+1}$ and $E_t x_{t+1}$. Then the system can be written as:

$$\boldsymbol{\Gamma_0(M,\chi)s_t = \Gamma_1(\chi)s_{t-1} + \Psi\varepsilon_t + \Pi\omega_t}$$

Here, $\boldsymbol{\omega_t}$ is a vector of expectation errors, and $\boldsymbol{\varepsilon_t}$ is a vector comprising the innovations in the fundamental shock processes. Following Lubik and Schorfheide (2004), one can write the solution to the model as:

$$\boldsymbol{s_t = Z_1(M,\chi)s_{t-1} + Z_2(M,\chi,Q)\varepsilon_t + Z_3(M,\chi)\zeta_t}$$

Under determinacy, the vector $\boldsymbol{Z_3(M)}$ is restricted to be a vector of zeros; sunspot shocks are ruled out by assumption because they have greater than unit root persistence. Under indeterminacy, however, sunspot shocks may affect the endogenous variables. Under determinacy, the matrix $\boldsymbol{Z_2(M,\chi,Q)}$ is uniquely determined by $\boldsymbol{M}$ and $\boldsymbol{\chi}$. Under indeterminacy, the response of the endogenous variables to fundamental shocks may vary. Throughout, I use Lubik and Schorfheide (2004) continuity solution as a baseline.[50] The vector $\boldsymbol{Q}$ parameterises how responses differ from this continuity solution. Thus, under the continuity solution it is just a vector of zeros, and one can just write $\boldsymbol{Z_2}$ as a function of $\boldsymbol{M}$ and $\boldsymbol{\chi}$.

[50] There is a slight complication in that one has to calculate the "nearest" point at the edge of the the determinacy region to calculate the continutiy solution. To get to the edge of the determinacy region one could adjust m_f or m_c instead of ϕ_π, but to be consistent with Lubik and Schorfheide (2004) I use ϕ_π.

The dynamics above imply a particular variance-covariance matrix of the endogenous variables, which I will denote by $\boldsymbol{\Sigma_s(M, \chi)}$.

Consumption with a Time-Varying Discount Factor

The consumer's problem changes slightly when demand shocks are considered to be discount factor shocks (which stem from within the consumer problem) rather than technology shocks (which stem from outside the firm problem). Here, I derive an expression for $\hat{c}_t$ in the presence of discount factor shocks, which will be needed to calculate consumer attention. I use the same strategy as Gabaix (2020), modified to include a discount factor shock. I consider the deterministic version of the problem. Taking income (and hence labour supply) as given, the consumer's budget constraint can be solved forwards as follows:

$$C_t + \frac{C_{t+1}}{R_t} + \frac{C_{t+2}}{R_t R_{t+1}} + ... = Y_t + \frac{Y_{t+1}}{R_t} + \frac{Y_{t+2}}{R_t R_{t+2}} + ...$$

Using the consumption Euler equation then gives:

$$C_t = \left(1 + \beta_t^{\frac{1}{\gamma}} R_t^{\frac{1}{\gamma}-1} + ...\right)^{-1} \left(Y_t + \frac{Y_{t+1}}{R_t} + \frac{Y_{t+2}}{R_t R_{t+1}} + ...\right)$$

Linearising then yields the following relationship.

$$\widehat{c}_t = \sum_{\tau \geq 0} \beta^\tau \left((1-\beta)\widehat{y}_{t+\tau} - \frac{\beta^2}{\gamma}\widehat{r}_{t+\tau} - \frac{1}{\gamma} z_{t+\tau}\right)$$

Consumers' Cognitive Discount Factor

The calculation of consumer attention is now slightly more complex because under indeterminacy and with interest rate smoothing the expectations of the endogenous variables cannot just be written as linear functions of the contemporaneous shocks. I instead propose the following method. First, define the selection vectors $\boldsymbol{k_y}$, $\boldsymbol{k_r}$, and $\boldsymbol{k_z}$ such that:

$$\boldsymbol{k_y s_t} = \widehat{y}_t$$
$$\boldsymbol{k_r s_t} = \widehat{r}_t$$
$$\boldsymbol{k_z s_t} = z_t$$

The above expression can then be rewritten as:

$$\begin{aligned}\partial c_t/\partial m_c &= \frac{1}{m_c} E_t \sum_{\tau\geq 0} \tau(\beta m_c)^\tau \left((1-\beta)\boldsymbol{k_y} - \frac{\beta^2}{\gamma}\boldsymbol{k_r} - \frac{1}{\gamma}\boldsymbol{k_z}\right) \boldsymbol{s_{t+\tau}} \\ &= \frac{1}{m_c} E_t \sum_{\tau\geq 0} \tau(\beta m_c)^\tau \left((1-\beta)\boldsymbol{k_y} - \frac{\beta^2}{\gamma}\boldsymbol{k_r} - \frac{1}{\gamma}\boldsymbol{k_z}\right) \boldsymbol{Z_1(M,\chi)}^\tau \boldsymbol{s_t}\end{aligned}$$

For compactness of notation, define:

$$\boldsymbol{k_c} = (1-\beta)\boldsymbol{k_y} - \frac{\beta^2}{\gamma}\boldsymbol{k_r} - \frac{1}{\gamma}\boldsymbol{k_z}$$

Noting the presence of the τ in the summation, and recalling that the eigenvalues of $\boldsymbol{Z_1(M,\chi)}$ lie within the unit circle, one can write this as:

$$\partial c_t/\partial m_c = \boldsymbol{k_c}(\boldsymbol{I} - \beta m_c \boldsymbol{Z_1(M,\chi)})^{-2}\beta \boldsymbol{Z_1(M,\chi)}\boldsymbol{s_t}$$

I take the expectation of the square, and evaluate at the default as explained above. I obtain the key quantity in the attention function:

$$\begin{aligned}&E[(\partial c_t/\partial m_c)^2] = \\ &\quad \beta^2 \boldsymbol{k_c}(\boldsymbol{I} - \beta m_{c,d}\boldsymbol{Z_1(M,\chi)})^{-2}\boldsymbol{Z_1(M,\chi)}\boldsymbol{\Sigma_s(M,\chi)}\ \boldsymbol{Z_1(M,\chi)}'\left((\boldsymbol{I} - \beta m_{c,d}\boldsymbol{Z_1(M,\chi)})^{-2}\right)'\boldsymbol{k_c'}\end{aligned}$$

Note that the dynamics of the model are taken as given, so that $\boldsymbol{Z_1(M,\chi)}$, and indeed $\boldsymbol{\Sigma_s(M,\chi)}$, are unaffected by evaluating at the default.

Firms' Cognitive Discount Factor

I now show the full derviation for equation the firms' choice of cognitive discount factor. I begin with the following equation from Gabaix (2020), which gives the optimal choice of relative price q for firms resetting prices in time t:[51]

$$\begin{aligned}q_t &= (1-\beta\theta)\sum_{\tau\geq 0}(\beta\theta m_f)^\tau E_t(\pi_{t+1} + ... + \pi_{t+\tau} + \widehat{mc}_{t+\tau}) \\ \partial q_t/\partial m_f &= (1-\beta\theta)\frac{1}{m_f}\sum_{\tau\geq 0}\tau(\beta\theta m_f)^\tau E_t(\pi_{t+1} + ... + \pi_{t+\tau} + \widehat{mc}_{t+\tau})\end{aligned}$$

Note that no adjustment to the firm's attention problem is required because: (i) linearisation means that fluctuations in the discount factor do not enter into the expression for the optimal relative price, and (ii) the cost-push shock η_t originates outside the firm's

[51] See Gabaix (2020) equation (173).

problem.

I define selection matrices $\boldsymbol{k_\pi}$ and $\boldsymbol{k_{mc}}$ such that:

$$\boldsymbol{k_\pi s_t} = \pi_t$$
$$\boldsymbol{k_{mc} s_t} = \widehat{mc}_t$$

I then follow the same procedure as before. To lighten the notation, I denote $\boldsymbol{Z_1(m)}$ simply as $\boldsymbol{Z_1}$, though of course this still depends on the choices of attention.

$$\begin{aligned}
\partial q_t/\partial m_f &= (1-\beta\theta)\frac{1}{m_f}\sum_{\tau\geq 0}\tau(\beta\theta m_f)^\tau E_t(\boldsymbol{k_\pi s_{t+1}} + ... + \boldsymbol{k_\pi s_{t+\tau}} + \boldsymbol{k_{mc} s_{t+\tau}}) \\
&= (1-\beta\theta)\frac{1}{m_f}\sum_{\tau\geq 0}\tau(\beta\theta m_f)^\tau(\boldsymbol{k_\pi Z_1}(\boldsymbol{I} + ... + \boldsymbol{Z_1^{\tau-1}}) + \boldsymbol{k_{mc} Z_1})\boldsymbol{s_t} \\
&= (1-\beta\theta)\frac{1}{m_f}\sum_{\tau\geq 0}\tau(\beta\theta m_f)^\tau(\boldsymbol{k_\pi Z_1}(\boldsymbol{I} - \boldsymbol{Z_1})^{-1}(\boldsymbol{I} - \boldsymbol{Z_1}^\tau) + \boldsymbol{k_{mc} Z_1})\boldsymbol{s_t} \\
&= (1-\beta\theta)\frac{1}{m_f}\sum_{\tau\geq 0}\tau(\beta\theta m_f)^\tau(\boldsymbol{k_\pi Z_1}(\boldsymbol{I} - \boldsymbol{Z_1})^{-1} - \boldsymbol{k_\pi Z_1}(\boldsymbol{I} - \boldsymbol{Z_1})^{-1}\boldsymbol{Z_1}^\tau) + \boldsymbol{k_{mc} Z_1})\boldsymbol{s_t}
\end{aligned}$$

One then arrives at the following:

$$\partial q_t/\partial m_f = (1-\beta\theta)\beta\theta\boldsymbol{P s_t}$$
$$E[(\partial q_t/\partial m_f)^2] = (1-\beta\theta)^2(\beta\theta)^2\boldsymbol{P\Sigma_s P'}$$

The matrix $\boldsymbol{P}$ is defined as follows (noting that the termabove is evaluated at $m_{f,d}$):

$$\begin{aligned}
\boldsymbol{P} \coloneqq \boldsymbol{k_\pi Z_1}(\boldsymbol{I - Z_1})^{-1}((\boldsymbol{I} - \beta\theta m_{f,d}\boldsymbol{I})^{-2} - (\boldsymbol{I} - \beta\theta m_{f,d}\boldsymbol{Z_1})^{-2}\boldsymbol{Z_1}) \\
+\boldsymbol{k_{mc}}(\boldsymbol{I} - \beta\theta m_{f,d}\boldsymbol{Z_1})^{-2}\boldsymbol{Z_1}
\end{aligned}$$

1.A.2 Consumer Attention

I now discuss the determinants of consumer attention. In the model, consumer attention is largely driven by real interest rates rather than output. The coefficient on output in the selection matrix $\boldsymbol{k_c}$ is $(1-\beta)$, which is the marginal propensity to consume today out of a change in the discounted value of lifetime income. In an infinite horizon optimization problem, persistent deviations in real interest rates may cause large changes in consumption. Small changes in income, even if persistent, tend to have smaller effects.

For cost-push or technology shocks, the volatility of the real interest rate rises as ϕ_π increases. In the case of cost-push shocks, inflation stabilization is achieved by causing

the real rate to deviate from its natural level, causing changes in the output gap and real wages that partially offset the exogenous change in marginal costs. For technology shocks, stabilisation of inflation is achieved by stabilising the real rate around its natural level. That involves causing larger deviations from steady state, which is what matters for the consumer's problem. For some combinations of shock variances, greater stabilization of inflation may therefore raise consumer attention, potentially increasing output gap volatility. For discount factor shocks, a higher ϕ_π increases real interest rate volatility by stabilising the real rate around a fluctuating natural level. In this case, however, that lowers consumer attention, because the shock comes from within the consumer's problem. The consumer cares about fluctuations in $\beta_t R_t$, and this term is less volatile if ϕ_π is higher. Under endogenous attention, then, the source of shocks matters for policy analysis. For policy shocks, greater stabilisation of inflation partially reverses the initial shock to the real interest rate, and so lowers consumer attention. The effect of ϕ_y is ambiguous. For cost-push, discount factor, or policy shocks, greater output gap stabilisation lowers consumer attention. For technology shocks, though, it increases attention; stabilisation occurs by pushing the real interest rate towards its natural level, and hence further from its steady state.

1.A.3 Effects of Attention on Macroeconomic Volatility

How far changes in attention affect inflation volatility depends on the model's other parameters. To see this, note that under determinacy inflation can be written as: $\pi_t = \psi_{\pi\eta}\eta_t$. $\psi_{\pi\eta}$ can be solved for using the method of undetermined coefficients, to give:

$$\psi_{\pi\eta} = \frac{(1 - M_c\rho_\eta + \sigma\phi_y)}{(1 - M_c\rho_\eta + \sigma\phi_y)(1 - \beta M_f\rho_\eta) + \kappa\sigma(\phi_\pi - \rho_\eta)}$$

The derivative with respect to M_f is:

$$\partial\psi_{\pi\eta}/\partial M_f = \beta\rho_\eta\psi_{\pi\eta}^2$$

The effect of changes in M_f on inflation volatility consequently tend to be larger when shocks are more persistent, since an increase in ρ_η also increases $\psi_{\pi\eta}$.

1.A.4 Firm Losses under Cognitive Discounting

Under Calvo pricing, the expected real-terms that firms make when they reset prices (until the next price reset) is:

$$D = \sum_{\tau \geq 0} \theta^{\tau} Q_{t,t+\tau} \left(\left(\frac{P_i}{P_{t+\tau}} \right)^{1-\epsilon} - \left(\frac{P_i}{P_{t+\tau}} \right)^{-\epsilon} MC^r_{t+\tau} \right)$$

Denote D^* the profits if P_i is set to the optimum P^*. Denote $\tilde{D}$ the profits under some alternative $\tilde{P}$. Now the expected loss in profits as a percentage of (optimum) sales is:

$$\begin{aligned} \frac{D^* - \tilde{D}}{S^*} &\approx \frac{\frac{\partial D}{\partial P_i}(\tilde{P} - P^*) + \frac{1}{2}\frac{\partial^2 D}{\partial P_i^2}(\tilde{P} - P^*)^2}{S^*} \\ &= \frac{\frac{1}{2}\frac{\partial^2 D}{\partial P_i^2}(\tilde{P} - P^*)^2}{S^*} \end{aligned}$$

Derivatives are evaluated at the optimum, which is why the first order term disappears. Expected sales are, in real terms:

$$S = \sum_{\tau \geq 0} \theta^{\tau} Q_{t,t+\tau} Y_{t+\tau} \left(\frac{P_i}{P_{t+\tau}} \right)^{1-\epsilon}$$

Next, I find the optimal price. The first-order condition is:

$$\frac{\partial D}{\partial P_i} = \sum_{\tau \geq 0} \theta^{\tau} Q_{t,t+\tau} Y_{t+\tau} \left((1-\epsilon) \frac{P_i^{-\epsilon}}{P_{t+\tau}^{1-\epsilon}} + \epsilon \frac{P_i^{-\epsilon-1}}{P_{t+\tau}^{-\epsilon}} MC^r_{t+\tau} \right)$$

The optimum price P^* equates this first-order condition to zero. Hence:

$$0 = \sum_{\tau \geq 0} \theta^{\tau} Q_{t,t+\tau} Y_{t+\tau} \left((1-\epsilon) \frac{(P^*)^{-\epsilon}}{P_{t+\tau}^{1-\epsilon}} + \epsilon \frac{(P^*)^{-\epsilon-1}}{P_{t+\tau}^{-\epsilon}} MC^r_{t+\tau} \right)$$

Rearranging gives:

$$(\epsilon - 1) \sum_{\tau \geq 0} \theta^{\tau} Q_{t,t+\tau} Y_{t+\tau} \left(\frac{(P^*)^{1-\epsilon}}{P_{t+\tau}^{1-\epsilon}} \right) = \epsilon \sum_{\tau \geq 0} \theta^{\tau} Q_{t,t+\tau} Y_{t+\tau} \left(\frac{(P^*)^{-\epsilon}}{P_{t+\tau}^{-\epsilon}} MC^r_{t+\tau} \right)$$

Denote these quantities as:

$$(\epsilon - 1)\zeta^* = \epsilon \upsilon^*$$

Now note that:

$$\begin{aligned} D^* &= \zeta^* - \upsilon^* \\ &= \frac{\epsilon}{\epsilon - 1}\upsilon^* - \upsilon^* \\ &= \frac{1}{\epsilon - 1}\upsilon^* \end{aligned}$$

Sales are:

$$S^* = \zeta^* = \frac{\epsilon}{\epsilon - 1}\upsilon^*$$

Next, I find the second derivative:

$$\frac{\partial^2 D}{\partial P_i^2} = \sum_{\tau \geq 0} \theta^\tau Q_{t,t+\tau} Y_{t+\tau} \left(\epsilon(\epsilon - 1) \frac{P_i^{-\epsilon-1}}{P_{t+\tau}^{1-\epsilon}} - \epsilon(\epsilon + 1) \frac{P_i^{-\epsilon-2}}{P_{t+\tau}^{-\epsilon}} MC_{t+\tau}^r \right)$$

Evaluated at the optimum:

$$\frac{\partial^2 D}{\partial P_i^2} = \sum_{\tau \geq 0} \theta^\tau Q_{t,t+\tau} Y_{t+\tau} \left(\epsilon(\epsilon - 1) \frac{(P^*)^{-\epsilon-1}}{P_{t+\tau}^{1-\epsilon}} - \epsilon(\epsilon + 1) \frac{(P^*)^{-\epsilon-2}}{P_{t+\tau}^{-\epsilon}} MC_{t+\tau}^r \right)$$

Mutiply both sides by $(P^*)^2$:

$$\begin{aligned} (P^*)^2 \frac{\partial^2 D}{\partial P_i^2} &= \sum_{\tau \geq 0} \theta^\tau Q_{t,t+\tau} Y_{t+\tau} \left(\epsilon(\epsilon - 1) \frac{(P^*)^{1-\epsilon}}{P_{t+\tau}^{1-\epsilon}} - \epsilon(\epsilon + 1) \frac{(P^*)^{-\epsilon}}{P_{t+\tau}^{-\epsilon}} MC_{t+\tau}^r \right) \\ &= \epsilon(\epsilon - 1)\zeta^* - \epsilon(\epsilon + 1)\upsilon^* \\ &\quad - \epsilon^2 \upsilon^* - \epsilon^2 \upsilon^* - \epsilon e \upsilon^* \\ &= -\epsilon \upsilon^* \end{aligned}$$

As a consequence:

$$\begin{aligned} \frac{D^* - \tilde{D}}{S^*} &= \frac{1}{2} \frac{(P^*)^2 \frac{\partial^2 D}{\partial P_i^2}}{S^*} \left(\frac{\tilde{P} - P^*}{P^*} \right)^2 \\ &= \frac{1}{2} \frac{-\epsilon \upsilon^*}{\frac{\epsilon}{\epsilon - 1}\upsilon^*} \left(\frac{\tilde{P} - P^*}{P^*} \right)^2 \\ &= -\frac{1}{2}(\epsilon - 1) \left(\frac{\tilde{P} - P^*}{P^*} \right)^2 \end{aligned}$$

In expectation, the loss is then a constant times the expectation of the square of the percentage deviation in the pricing decision from the optimum. Suppose that the firm is not paying full attention to the future, and doing so changes the square of its pricing decision by $1.5\%^2$ on average. Then, using the above equation, and specifying $\epsilon = 10$, one obtains an expected loss of $c.0.1\%$ of sales.

1.B Numerical Solution Method

As illustrated by the charts in the main text, if rates are not constrained by the ELB, then the $\boldsymbol{g(\cdot)}$ mapping tend to have a unique fixed point within the determinacy region (where such a fixed point exists at all) and a unique fixed point within the indeterminacy region. It is also worth noting that, if the fixed points within each region are indeed unique, where they exist, then both exist where the function crosses from above the 45 degree line to below it.

These features are helpful when finding fixed points numerically. I use a very simple fixed point iteration algorithm. For the single dimensional case, assume some starting value for m vector, denoted m_0, I perform the following iteration until convergence.

$$m_{n+1} = (1 - c) \times m_n + c \times g(m_n, \boldsymbol{\chi}, \boldsymbol{\xi})$$

Here, $c \in (0, 1]$ is an adjustment factor. To find a determinate equilibrium, I found that it often works best to set $m_0 = m_d$. It may be that this point is a fixed point. If there is no fixed point at m_d, then $g(m_n, \boldsymbol{\chi}, \boldsymbol{\xi}) > m_d$, and so m_n increases. Because a fixed point will involve $g(m_n, \boldsymbol{\chi}, \boldsymbol{\xi})$ crossing the 45 degree line from above, if m_n is below the equilibrium level, it will increase, while if it is above the equilibrium level, it will decrease. As such, if such a fixed point does exist, I have found that m_n typically converges to it, provided c is sufficiently small. If there is no determinate equilibrium, then m_n will eventually exceed $\tilde{m}(\boldsymbol{\chi})$, at which point the algorithm stops. For the indeterminate case, the same features of the $g(m_n, \boldsymbol{\chi}, \boldsymbol{\xi})$ mapping help to find the fixed point. Often it works best to start at $m_0 = 1$, and the algorithm decreases m_n until the fixed point is reached.

For the two-dimensional case, I use the same solution method as for the one-dimensional case, except that now the g function is bivariate rather than univariate. To assess attainability of the indeterminate equilibrium, I search for an equilibrium in the first period of indeterminacy. That means τ is set equal to 1.

I found that the performance of the algorithm improved when I replaced the maximum operator in the attention function with the "smoothed maximum", described below.

1.B.1 Smoothed Maximum

Here I explain the "smoothed max" function which I found allowed for better performance of the fixed point iteration for finding attention equilibrium. One replaces the maximum in the attention function with the following differentiable approximation:[52]

$$\frac{m_d exp\{\alpha m_d\} + \Lambda exp\{\alpha \Lambda\}}{exp\{\alpha m_d\} + exp\{\alpha \Lambda\}}$$

Here, $\Lambda \coloneqq 1 - \frac{1}{\lambda}$. As $\alpha \to \infty$ this converges to the maximum.

1.C Policy at the ELB

In this section, I develop a version of the model with regime switching and analyze the effects of average inflation targeting.

1.C.1 A Regime Switching Model

In the section on average inflation targeting, I use a regime switching model. I assume there are two states: state N (for Normal times), and state E (for ELB). I allow attention to vary with the regime; it seems plausible that agents would choose a different level of attention in an ELB episode if the ELB leads to very different macroeconomic dynamics. To simplify the derviations, I assume state N is absorbing. Because it is an absorbing state, the choice of attention is not affected by switching probabilities, and so can be calculated as in the previous subsection. The state N model is as in the previous section, except for a different Taylor rule. I consider two monetary policy rules. The first is simply a Taylor rule without interest rate smoothing or monetary policy shocks:

$$i_t = \phi_\pi \pi_t + \phi_y y_t$$

The alternative is a rule in which the central bank responds to the average inflation rate over N quarters:

$$i_t = \phi_\pi \frac{1}{N} \sum_{s=0}^{N-1} \pi_{t-s} + \phi_y y_t$$

In state E, I assume that the economy faces an exogenous probability p of remaining in state E, but switches to state N with probability $1 - p$. In state E, the Phillips and

[52] This particular formulation of the smoothed maximum is from Lange et al. (2014).

IS curve are given by:

$$\pi_t = \beta M_{f,E}(\boldsymbol{\chi}, \boldsymbol{\xi}) E_t \pi_{t+1} + \kappa y_t + \eta_t$$
$$y_t = M_{c,E}(\boldsymbol{\chi}, \boldsymbol{\xi}) E_t x_{t+1} - \sigma(-E_t \pi_{t+1}) + \tilde{z}_t + \sigma r_E^n$$

Here, $r_E^n < 0$ is the natural real rate of interest in the ELB state. Note that I assume that this arises from a discount factor shock. In this case, the source of the shock does matter because it is an internal shock that affects consumer attention directly. In state E, the output gap, inflation, and real interest rate steady states differ from the steady state in normal times. In keeping with Gabaix's methodology, I assume that consumers discount future deviations from the normal times steady state.

The equations for calculating consumer and firm attention are now somewhat more complex, because the expectation at time t of the endogenous variables in time $t + h$ depends not just on whether the economy switched into state N but also *when* it did so; a switch in $t + 1$ implies different $t + h$ state variables from a switch in $t + h$. I leave a description of the solution methodology to the Appendix.

1.C.2 Average Inflation Targeting

Average inflation targeting reduces the costs of the ELB under rational expectations. However, as Budianto et al. (2023) show, AIT may have small effects under exogenous cognitive discounting. This is because AIT operates through altering expectations while the economy is in the ELB state; this effect is smaller when agents discount the future more heavily.

I now show that the effects are potentially far larger under endogenous cognitive discounting. The calibration considered is: $\beta = 0.99$ $\theta = 0.85$, $\gamma = 3$, $\phi = 1$, $\phi_\pi = 2.50$, $\phi_y = 0.25$, $\rho_\eta = 0.60$, $\rho_z = 0.80$, $\sigma_{e,\eta} = 0.10$, and $\sigma_{e,z} = 0.20$. The strong policy response to inflation gives AIT the greatest chance to affect macroeconomic dynamics. I assume $p = 0.85$, giving an expected duration of the ELB episode of 6 to 7 quarters. I assume that the natural real rate is -1.5% during the period of the ELB. I use a slightly lower default level of attention than in the previous section, $m_{c,d} = m_{f,d} = 0.70$, since Budianto et al. (2023) emphasise that the power of AIT is muted particularly when the discount factor is far from rationality. I continue to assume $\kappa_c = 4.5$ and $\kappa_f = 1.5$. I then consider two versions of the policy rule. The first is that the central bank follows a simple Taylor rule in state N. The second assumes that the central bank responds to the average inflation rate over an extended period. I choose a long window of 16 quarters, again giving AIT the best chance of mitigating the costs of the ELB.

To solve the model, one has to find the equilibrium level of m_c and m_f in state N. In

the calibration chosen, the equilibrium in all cases involves $m_{c,N} = m_{c,d}$ and $m_{f,N} = m_{f,d}$. One can then solve for the state N dynamics, which in turn allows one to solve for the state E dynamics for any combination of $m_{c,E}$ and $m_{f,E}$. Then, one can use standard numerical routines to find fixed points of the attention vector $\boldsymbol{m_E}$.

Figures 1.18 and 1.19 show the implied choices of $m_{f,E}$ and $m_{c,E}$ for different combinations of $m_{c,E}$ and $m_{f,E}$ under AIT and under the Taylor rule. In this calibration, high attention pushes the economy into the indeterminate region when monetary policy is set using a Taylor rule. At the boundary of the determinacy region, the recession induced by the ELB becomes unboundedly large. That generates the dramatic increase in the choice of attention close to the boundary; the choices of $m_{c,E}$ and $m_{f,E}$ approach unity at the boundary. I do not consider the possibility of indeterminate equilibria here, but future research could do so.

Figures 1.20 show equilibria in $m_{f,E}$ and $m_{c,E}$ under the two policy regimes. The red lines show which combinations of $m_{f,E}$ and $m_{c,E}$ imply that the same $m_{c,E}$ is the optimal choice of attention for consumers. The blue lines show the same for $m_{f,E}$. Attention equilibria occur where the two lines intersect. Under AIT, the unique attention equilibrium is with $m_{c,E} = m_{c,d}$ and $m_{f,E} = m_{f,d}$. The analysis is simplified considerably in this case because even full attention does not push the economy into the indeterminate region in this calibration.

Under the Taylor rule, the situation is more complex. The optimal choice of firm attention is much higher. Under AIT, the expected rate of inflation is a combination of (i) the low rate expected if the economy remains in state E, and (ii) the high rate that will be justified over an extended period if the economy moves into state N. Expectations are consequently closer to zero, and incentives to pay attention are lower. Under the Taylor rule, effect (ii) is absent, and so expectations are further from zero, and incentives to pay attention are much higher.

Moreover, for any given $m_{f,E}$ there may be two or even three equilibria in $m_{c,E}$. If consumers pay little attention, then output and inflation are not too far from zero. The real interest rate consequently does not rise too much, and consumers have low incentives to pay attention. However, if consumers pay sufficient attention, then the economy approaches the boundary of the determinacy region, an arbitrarily large recession and deflation ensues, the real interest rate rises dramatically, and incentives to pay attention are high.

This effect results in two attention equilibria existing under the Taylor rule. These are summarized in table 1.5, which shows the equilibrium levels of attention under AIT and the two Taylor rule equilibria, the resulting average deviation in inflation and the output gap in state E, and the variability of inflation and the output gap in the two states. The

Table 1.5: Equilibria under Taylor Rule vs. Average Inflation Targeting

	Avg. Inflation Targeting			Taylor Rule		
	Equilibrium	Counterfactual 1	Counterfactual 1	Equilibrium 1	Equilibrium 2	Counterfactual
$m_{c,E}$	0.70	0.70	0.87	0.70	0.87	0.70
$m_{f,E}$	0.70	0.96	1.00	0.96	1.00	0.70
$\bar{\pi}_E$	-0.31	-0.54	-0.91	-1.18	-5.80	-0.42
$\bar{y}_E$	-1.28	-1.31	-1.95	-2.06	-8.32	-1.53
$St.Dev.\pi_N$	0.64	0.64	0.64	0.59	0.59	0.59
$St.Dev.y_N$	0.64	0.64	0.64	0.64	0.64	0.64
$St.Dev.\pi_E$	0.73	1.01	1.01	1.28	1.47	0.78
$St.Dev.y_E$	0.29	0.37	0.46	0.50	0.72	0.29

average deviation of inflation from the target is dramatically higher under the Taylor rule, and the Taylor rule also induces a considerably larger recession. To show the source of this difference, I analyze three counterfactuals. First, I show the dynamics if one keeps AIT but uses the equilibrium attention level under the two Taylor rule equilibria. Note the size of the deflation and recession induced by the ELB is much smaller than under the Taylor rule; this is because AIT is particularly strongly effective when agents are close to rationality and hence very forward-looking. I then analyze a counterfactual where the policy rule used is the Taylor rule but the attention levels are the AIT attention levels. Here, the size of the deflation and recession induced by the ELB is quite similar to the AIT case.

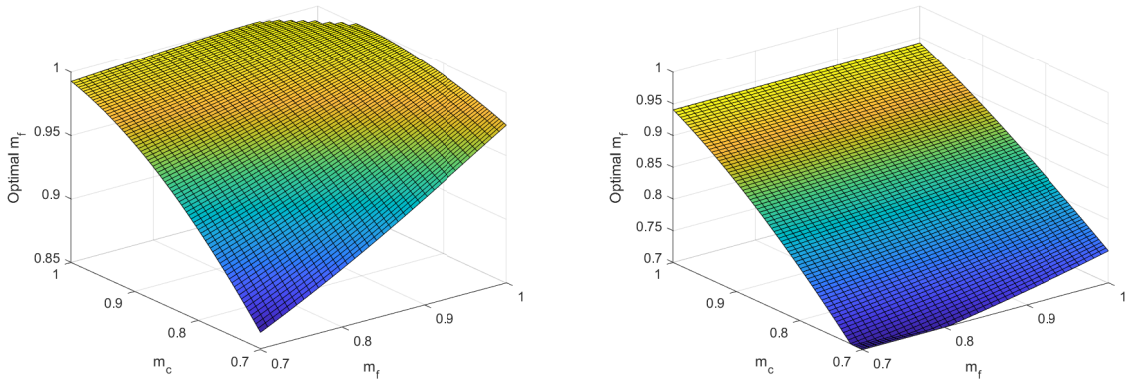
To summarize, AIT is effective at reducing the costs of the ELB when agents are very forward-looking, and less so when agents discount the future heavily. However, AIT induces agents to discount the future far more than does the Taylor Rule. As such, the benefits of AIT may be higher under endogenous than exogenous discounting. The possibility of multiplicity does complicate the analysis, and the extent of the differences in equilibrium attention between the two regimes likely depends on the calibration. Future work could explore this further.

1.C.3 Model Solution under Regime Switching

State N

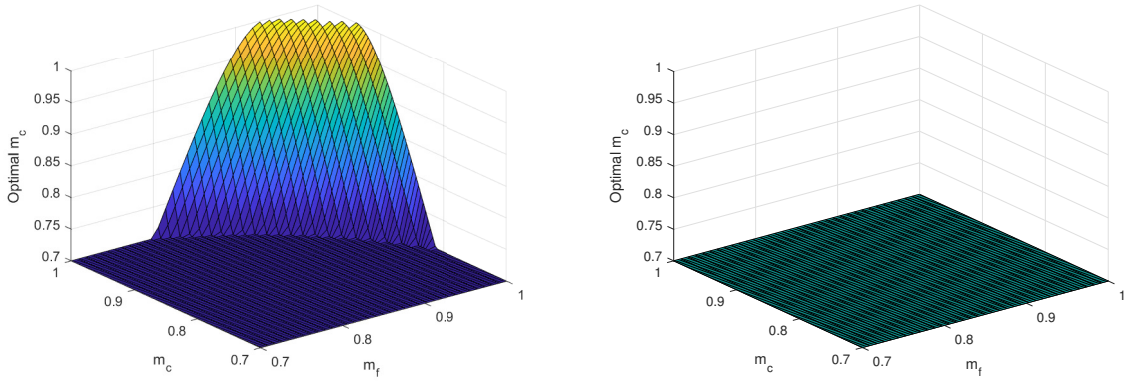
In the regime switching model, I focus on determinate equilibria. This allows me to adopt a somewhat simpler method of solving the model. To solve the model with exogenous attention, one can iterate to find the minimum state variable solution. Taking as given the level of attention, the state N model can be written as:

$$\boldsymbol{A_N s_t} = \boldsymbol{B_N s_{t-1}} + \boldsymbol{C_N} E_t \boldsymbol{s_{t+1}} + \boldsymbol{D_N \varepsilon_t} + \boldsymbol{E_N}$$



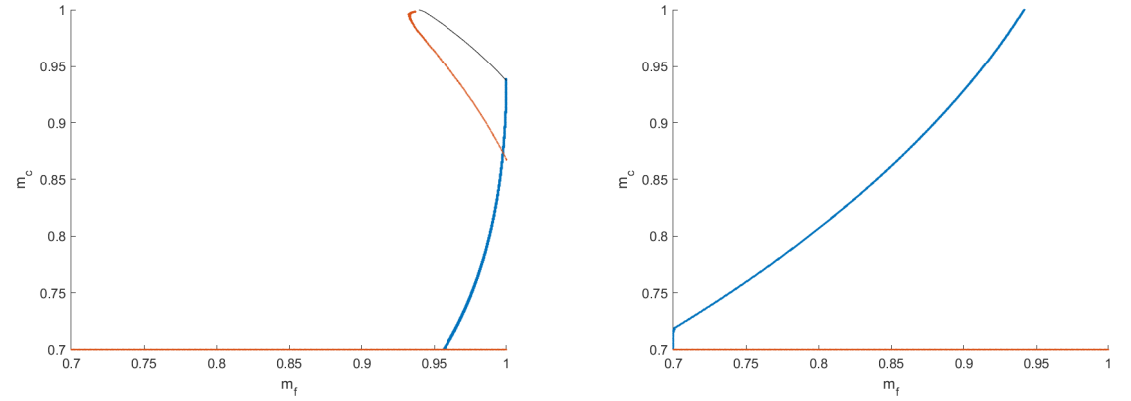
(a) Taylor Rule **(b)** Average Inflation Targeting

Figure 1.18: Optimal choice of $m_{f,E}$



(a) Taylor Rule **(b)** Average Inflation Targeting

Figure 1.19: Optimal choice of $m_{c,E}$



(a) Taylor Rule **(b)** Average Inflation Targeting

Figure 1.20: Attention Equilibria. Note that points marked in red are equilibria in $m_{c,E}$, and blue denotes equilibria in $m_{f,E}$. Intersections of the blue and red curves give attention equilibria. The black line denotes the boundary of the determinacy region.

Here, note that $\boldsymbol{s_t}$ does not contain expectation terms. In state N, the vector $\boldsymbol{E_N}$ is just a vector of zeros, giving a steady state that is just a vector of zeros. One can use a straightforward iteration to solve for the minimum state variable solution to the model:

$$\boldsymbol{s_t} = \boldsymbol{F_N s_t} + \boldsymbol{G_N \varepsilon_t}$$

With the solution in hand, one can then straightforwardly obtain the covariance matrix $\boldsymbol{\Sigma_{s.N}}$.

To calculate the implied level of attention, one can use the same method as above, with one slight difference. Here, $\boldsymbol{s_t}$ does not contain the expectation terms. For consumers, one has to calculate:

$$\hat{c}_t = -\sigma \boldsymbol{k_i} \sum_{h \geq 0} m_c^h E_t \boldsymbol{s_{t+h}} + \sigma \boldsymbol{k_\pi} \sum_{h \geq 0} m_c^h E_t \boldsymbol{s_{t+h}} + \beta^{-2} \sigma \boldsymbol{k_z} \sum_{h \geq 0} m_c^h E_t \boldsymbol{s_{t+h}}$$

Calculating the expectation in state N is straightforward, because it is an absorbing state. One obtains:

$$\hat{c}_t = -\sigma(\boldsymbol{k_i} - \beta^{-2}\boldsymbol{k_z})(\boldsymbol{I} - m_c \boldsymbol{F_N})^{-1} \boldsymbol{s_t} + \sigma \boldsymbol{k_\pi} (\boldsymbol{I} - m_c \boldsymbol{F_N})^{-1} \boldsymbol{F_N s_t}$$

One could differentiate this expression with respect to m_c analytically. I will actually proceed numerically, because the analytical solution becomes too complicated in state E. So, I write this expression as:

$$\hat{c}_t = \boldsymbol{\Omega_N}(m_c) \boldsymbol{s_t}$$

To calculate the expectation of the square of the derivative with respect to m_c, evaluated at the default:

$$\hat{c}_t = \left(\frac{\boldsymbol{\Omega_N}(m_{c,d} + \epsilon) - \boldsymbol{\Omega_N}(m_{c,d})}{\epsilon} \right) \boldsymbol{\Sigma_{s,N}} \left(\frac{\boldsymbol{\Omega_N}(m_{c,d} + \epsilon) - \boldsymbol{\Omega_N}(m_{c,d})}{\epsilon} \right)'$$

The calculation proceeds in the same fashion for firms. To calculate the equilibrium level of attention, I proceed numerically using a minimisation routine in MATLAB. Multiplicity is generally not a concern in state N, but if desired one can use a routine such as *fmincon* to search for equilibria in different intervals of m_f and m_c.

State E

In state E, matters are complicated by the fact that this is not an absorbing state. To solve the model for a given set of attention parameters, write the model as:

$$\boldsymbol{A_E s_t} = \boldsymbol{B_E s_{t-1}} + \boldsymbol{C_E} E_t \boldsymbol{s_{t+1}} + \boldsymbol{D_E \varepsilon_t} + \boldsymbol{E_E}$$

Then note that $E_t \boldsymbol{s_{t+1}}$ is given by:

$$\begin{aligned} E_t \boldsymbol{s_{t+1}} &= p E_t \boldsymbol{s^E_{t+1}} + (1-p) E_t \boldsymbol{s^N_{t+1}} \\ &= p E_t \boldsymbol{s^E_{t+1}} + (1-p) \boldsymbol{F_N s^E_t} \end{aligned}$$

Here, $E_t \boldsymbol{s^E_{t+1}}$ denotes the expectation condition on being in state E. Then one can rewrite the model as:

$$\boldsymbol{A_E s^E_t} = \boldsymbol{B_E s^E_{t-1}} + p \boldsymbol{C_E} E_t \boldsymbol{s^E_{t+1}} + (1-p) \boldsymbol{C_E F_N s^E_t} + \boldsymbol{D_E \varepsilon_t} + \boldsymbol{E_E}$$

One can then again proceed using an iteration to solve for:

$$\boldsymbol{s^E_t} = \boldsymbol{F_E s^E_t} + \boldsymbol{G_E \varepsilon_t} + \boldsymbol{H_E}$$

I then use this equation to solve for the steady state $\boldsymbol{\bar{s}_E}$. One can then write:

$$\boldsymbol{s^E_t} - \boldsymbol{\bar{s}_E} = \boldsymbol{F_E}(\boldsymbol{s^E_t} - -\boldsymbol{\bar{s}_E}) + \boldsymbol{G_E \varepsilon_t}$$

One can use this to solve for the variance-covariance matrix $\boldsymbol{\Sigma_{s,E}}$.

Solving for the optimal choice of attention is now much more complex, because now the expectation $\boldsymbol{s_{t+h}}$ depends on whether the economy switched into state N, and if so at what time. Formally:

$$E_t \boldsymbol{s_{t+h}} = p^h \boldsymbol{F^h_E}(\boldsymbol{s_t} - \boldsymbol{\bar{s}_E}) + \boldsymbol{\bar{s}_E} + (1-p)(p^{h-1} \boldsymbol{F_N}(\boldsymbol{F^{h-1}_E}(\boldsymbol{s_t} - \boldsymbol{\bar{s}_E}) + \boldsymbol{\bar{s}_E}) + p^{h-2} \boldsymbol{F^2_N}(\boldsymbol{F^{h-2}_E}(\boldsymbol{s_t} - \boldsymbol{\bar{s}_E}) +$$

Combining this with the expression for $\hat{c}_t$ given above, one obtains:

$$\hat{c}_t = \boldsymbol{\Omega_{1,E}}(m_c) \boldsymbol{\bar{s}_E} + \boldsymbol{\Omega_{2,E}}(m_c)(\boldsymbol{s_t} - \boldsymbol{\bar{s}_E})$$

Where:

$$\boldsymbol{\Omega_{1,E}}(m_c) = -\sigma(\boldsymbol{k_i} - \beta^{-2}\boldsymbol{k_z}) \sum_h \geq 0 m_c^h \left(p^h \boldsymbol{F_E^h} + (1-p) \sum_{s=1}^{h} p^{h-s} \boldsymbol{F_N^s F_E^{h-s}} \right) +$$
$$\sigma \boldsymbol{k_\pi} \sum_{h \geq 0} m_c^h \left(p^{h+1} \boldsymbol{F_E^{h+1}} + (1-p) \sum_{s=1}^{h+1} p^{h+1-s} \boldsymbol{F_N^s F_E^{h+1-s}} \right)$$

And:

$$\boldsymbol{\Omega_{2,E}}(m_c) = -\sigma(\boldsymbol{k_i} - \beta^{-2}\boldsymbol{k_z} - p\boldsymbol{k_\pi}\boldsymbol{F_E})(\boldsymbol{I} - m_c p \boldsymbol{F_E})^{-1} -$$
$$\sigma(1-p)(\boldsymbol{k_i} - \beta^{-2}\boldsymbol{k_z}) \sum_{h \geq 0} m_c^h \sum_{s=1}^{h} p^{h-s} \boldsymbol{F_N^s F_E^{h-s}} + \sigma(1-p)\boldsymbol{k_\pi} \sum_{h \geq 0} m_c^h \sum_{s=1}^{h+1} p^{h+1-s} \boldsymbol{F_N^s F_E^{h+1-s}}$$

These matrices can be calculated numerically, summing from $h = 0$ to $h = H$ for some very large H. One can then solve for the expectation of the square of the derivative, and hence the optimal choice of attention, as above. The derivation for firms proceeds in the same fashion. To find attention equilibria, I once again use a MATLAB minimisation routine. Here, multiplicity is much more of a concern. To deal with this, one can either: (i) visualize the set of equilibria, as I do in the main body of the paper, to check for multiplicity; or (ii) use a minimisation routine to check for equilibria in a set of specific intervals of m_f and m_c.

1.D Alternative Loss Functions

I now briefly discuss how optimal policy might change under alternative assumptions on the loss function.

The central bank loss function considered in the main paper made two assumptions. First, subjective rather than objective expectations appear in the loss function. One could instead assume that the central bank too discounts the future some rate $m_{cb} \in [0, 1]$ and hence use the loss function:

$$L_t = E_t \sum_{h \geq 0} (\beta m_{cb})^h \left(\vartheta y_{t+h}^2 + \pi_{t+h}^2 \right) \tag{1.49}$$

Second, I assumed that consumers' and firms' thinking costs do not appear in the central bank loss function. One could instead assume that the central bank cares about consumers' and firms' thinking costs. For example, one could suppose that the central bank

loss function features additively separable terms in m_c and m_f:

$$L_t = E_t \sum_{h \geq 0} \beta^h \left(\vartheta y_{t+h}^2 + \pi_{t+h}^2 \right) + z_c(m_c) + z_f(m_f) \tag{1.50}$$

I allow arbitrary functions $z_c(\cdot)$ and $z_f(\cdot)$ in the discussion below, allowing the central bank to place any particular weight on thinking costs, rather than necessarily reflecting the agents' welfare loss. However, these functions could be chosen to mirror the attention costs in the agents' problems described in section 1.2. At a minimum, one would want these functions to be continuous and strictly increasing in attention for $m \in [m_d, 1]$.

Under discretion, neither modification of the central bank loss function would affect optimal policy. Changing the discount factor in the central bank's loss function does not affect the first order conditions of the central bank's loss minimization problem, which still yields the solution:

$$y_t = -\frac{\kappa}{\vartheta} \pi_t$$

Equally, supposing that the central bank loss function features an additively separable term in the consumers' or firms' attention costs does not affect optimal policy under discretion because the central bank cannot commit to a particular policy regime and hence consumer and firm attention is invariant to policy.

Under commitment, however, both modifications of the central bank loss function would affect the optimal policy. Indeed, changing the central bank's discount factor would affect optimal policy under commitment even under exogenous cognitive discounting. As Gabaix highlights, the fact that price level targeting is optimal in the baseline new Keynesian model with rational expectations arises from the fact that the discount factor in the central bank loss function is equal to the discount factor in the Phillips curve. As such, if the central bank's discount factor is equal to the discount factor in the Phillips curve then this result would be restored. Exploration of how the tradeoffs faced by the central bank would change in this case under endogenous cognitive discounting would be an interesting avenue for future research. Changing m_{cb} would change the value of future commitments, and so this might affect the new tradeoff described in the main text.

Equally, if there are additively separable terms in consumer and firm discounting in the loss function, that would further quantitatively affect optimal policy under commitment if discounting is endogenous. Under the original loss function, the central bank faced the tradeoff that higher firm attention on the one hand led to greater amplification of persistent shocks but at the same time made commitments to a future policy path more powerful at affecting output and inflation today. Under the alternative loss function in

1.50, the central bank would have the additional consideration that higher firm attention is inherently costly. Moreover, under the original loss function, consumers' attention does not affect the optimal paths for output and inflation, but only affects the policy path required to achieve those output and inflation outcomes. Under the loss function in 1.50, the central bank would need to factor in the implications for consumer attention when choosing a path for output and inflation. I leave more formal exploration of this topic for future work.

1.E Further Exercises in Identification

1.E.1 Identification without Interest Rate Smoothing

For the model without interest rate smoothing, applying the procedure outlined in Iskrev (2010) procedure yields the following results regarding point identification:

- If $\rho_\eta = \rho_z = \rho_v$, point identification fails even if M^f and M^c are known.
- If $\rho_\eta = \rho_z$, the model is identified if either M^f or M^c is known, but not if both are estimated.
- If $\rho_v = \rho_\eta$ or $\rho_v = \rho_z$ then point identification fails if either M^f or M^c has to be estimated.

1.E.2 Potential Biases from Neglecting Discounting

In this section, I show that if households and firms discount the future, but one assumes that they are rational, then one will underestimate the slopes of the Phillips and IS curves. In the rational expectations model, the expectation terms generate amplification of persistence shocks. For example, the expectation term in the Phillips curve amplifies the effect of persistent changes in the output gap on inflation. If these expectations are dampened due to cognitive discounting, then the effect of the output gap on inflation is attenuated. As such, an econmetrician assuming rational expectations would mistakenly estimate a flatter Phillips curve.

Periods of macroeconomic stability should, according to the endogenous cognitive discounting model, lead households and firms to discount the future more heavily. This leads to the appearance of flatter Phillips and IS curves, if the econometrician is using a rational expectations model. This effect could have contributed to the apparent flattening of the Phillips curve often discussed in the literature.[53] This finding contributes to the

[53]See McLeay and Tenreyro (2020) for an overview of the empirical evidence.

growing literature that suggests that changes in expectation formation could account for the apparent flatness of the Phillips curve during the Great Moderation. See, for example, Jørgensen and Lansing (2019), who find a similar result holds using a signal-extraction model.

To illustrate this point, I consider three model specifications. The first has a default attention level of 0.85, as in the estimation section, and large enough shocks to push the firm attention level well above the default. The second specification has the same attention level, but smaller shocks, so that both firm and consumer attention are at the default. The third specification considers a lower default level of attention, of 0.60, and again uses smaller shocks so that attention remains at the default for both firms and consumers. For each specification, I simulate 100 datasets of 200 periods. I then estimate an endogenous cognitive discounting model and a rational expectations model on the resulting data using maximum likelihood. For comparison, I also estimate a model that assumes exogenous cognitive discounting.[54] I then take the range of parameter estimates obtained, and take the (winsorized) mean and standard deviation.[55]

The results are shown in tables 1.6, 1.7, and 1.8. The rational expectations model gives a fairly accurate picture of the Phillips curve slope in the high shock specification. This is unsurprising, since the true value of m_f is close to unity, and so the rational expectations model is a fairly good approximation to the true model. In the lower shock specification, however, m_f is at its default, and so the rational expectations model is a poor approximation. In this case, the tendency is for the rational expectations model to underestimate the Phillips and IS curve slopes.

The results also demonstrate the identification problem with the exogenous attention model. Note how large is the standard deviation of the point estimates for m_f and m_c. The imprecise estimates of the discounting parameters spills over into inaccurate estimates of the Phillips and IS curve slopes; note how much larger are the standard deviation of the point estimates of κ under exogenous than endogenous attention. The problem of identifying the Phillips and IS curve slopes consequently cannot be easily resolved by simply adding exogenous discounting parameters into the model.

[54]The exogenous attention model is estimated using Dynare. Often Dynare finds for the exogenous attention model that the hessian at the mode is not positive definite. I discount any samples where this is the case, and take the first 100 samples where the hessian at the mode is positive definite.

[55]I trim the sample at the 5th and 95th percentiles for each of the parameters, to try to give an accurate representation of the typical results obtained.

Table 1.6: Distribution of Points Estimates Obtained Under MLE: High Default, Large Shocks

		Endogenous Attention			Exogenous Attention			Rational Expectations		
Parameter	True Value	Median	Mean	St.Dev.	Median	Mean	St.Dev.	Median	Mean	St.Dev.
ϕ_π	1.50	1.49	1.50	0.05	1.50	1.50	0.05	1.50	1.50	0.05
ϕ_x	0.25	0.25	0.25	0.03	0.24	0.25	0.04	0.25	0.25	0.03
θ	0.80	0.80	0.79	0.05	0.73	0.74	0.07	0.80	0.79	0.05
γ	2.00	1.90	2.01	0.51	1.60	1.73	0.50	2.19	2.31	0.56
ρ_η	0.75	0.75	0.74	0.05	0.75	0.75	0.05	0.75	0.74	0.05
ρ_z	0.90	0.90	0.89	0.04	0.90	0.90	0.03	0.89	0.89	0.03
ρ_i	0.60	0.60	0.60	0.02	0.59	0.59	0.02	0.60	0.60	0.02
σ_η^ε	0.25	0.26	0.28	0.06	0.36	0.40	0.15	0.23	0.24	0.04
σ_z^ε	0.50	0.51	0.52	0.11	0.49	0.54	0.18	0.45	0.46	0.08
σ_v^ε	0.25	0.25	0.25	0.01	0.25	0.25	0.01	0.25	0.25	0.01
m_f	0.96	0.95	0.94	0.04	0.82	0.70	0.32	1.00	1.00	0.00
m_c	0.85	0.85	0.85	0.00	0.79	0.73	0.23	1.00	1.00	0.00
κ	0.16	0.16	0.18	0.08	0.26	0.27	0.14	0.18	0.19	0.07
σ	0.50	0.52	0.52	0.13	0.62	0.62	0.16	0.45	0.45	0.11

Table 1.7: Distribution of Points Estimates Obtained Under MLE: High Default, Small Shocks

		Endogenous Attention			Exogenous Attention			Rational Expectations		
Parameter	True Value	Median	Mean	St.Dev.	Median	Mean	St.Dev.	Median	Mean	St.Dev.
ϕ_π	1.50	1.49	1.50	0.06	1.50	1.49	0.06	1.49	1.49	0.06
ϕ_x	0.25	0.25	0.25	0.03	0.24	0.25	0.04	0.25	0.25	0.03
θ	0.80	0.80	0.80	0.03	0.76	0.77	0.07	0.85	0.84	0.04
γ	2.00	1.95	2.05	0.56	1.65	1.76	0.53	2.45	2.56	0.67
ρ_η	0.75	0.74	0.73	0.05	0.73	0.73	0.05	0.73	0.73	0.05
ρ_z	0.90	0.89	0.89	0.04	0.90	0.90	0.03	0.88	0.88	0.03
ρ_i	0.60	0.60	0.60	0.02	0.59	0.59	0.02	0.60	0.60	0.02
σ_η^ε	0.10	0.10	0.10	0.01	0.13	0.14	0.05	0.07	0.07	0.01
σ_z^ε	0.20	0.20	0.21	0.05	0.21	0.24	0.09	0.18	0.19	0.04
σ_v^ε	0.10	0.10	0.10	0.00	0.10	0.10	0.01	0.10	0.10	0.00
m_f	0.85	0.85	0.85	0.01	0.72	0.59	0.37	1.00	1.00	0.00
m_c	0.85	0.85	0.85	0.00	0.75	0.67	0.28	1.00	1.00	0.00
κ	0.16	0.16	0.16	0.03	0.20	0.21	0.13	0.10	0.11	0.04
σ	0.50	0.51	0.52	0.13	0.60	0.61	0.17	0.40	0.41	0.10

Table 1.8: Distribution of Points Estimates Obtained Under MLE: Low Default, Small Shocks

		Endogenous Attention			Exogenous Attention			Rational Expectations		
Parameter	True Value	Median	Mean	St.Dev.	Median	Mean	St.Dev.	Median	Mean	St.Dev.
ϕ_π	1.50	1.50	1.50	0.09	1.50	1.50	0.08	1.49	1.49	0.09
ϕ_x	0.25	0.25	0.25	0.04	0.24	0.24	0.05	0.25	0.25	0.05
θ	0.80	0.80	0.80	0.03	0.79	0.79	0.07	0.86	0.85	0.05
γ	2.00	1.94	2.02	0.41	1.82	1.91	0.48	2.80	2.93	0.70
ρ_η	0.75	0.74	0.74	0.04	0.73	0.73	0.05	0.75	0.75	0.06
ρ_z	0.90	0.89	0.89	0.03	0.90	0.89	0.03	0.88	0.88	0.03
ρ_i	0.60	0.60	0.60	0.03	0.59	0.59	0.03	0.60	0.60	0.03
σ_η^ε	0.10	0.10	0.10	0.01	0.12	0.12	0.04	0.05	0.05	0.01
σ_z^ε	0.20	0.20	0.21	0.04	0.19	0.20	0.07	0.14	0.14	0.02
σ_v^ε	0.10	0.10	0.10	0.01	0.10	0.10	0.01	0.10	0.10	0.01
m_f	0.60	0.60	0.60	0.00	0.48	0.44	0.41	1.00	1.00	0.00
m_c	0.60	0.60	0.60	0.00	0.62	0.54	0.33	1.00	1.00	0.00
κ	0.16	0.16	0.16	0.04	0.18	0.19	0.13	0.10	0.12	0.06
σ	0.50	0.51	0.51	0.10	0.54	0.55	0.13	0.35	0.36	0.08

1.F Estimation Methodology

My process for estimating the model broadly follows techniques outlined in Lubik and Schorfheide (2004) and Herbst and Schorfheide (2016), adapting these methods to allow for endogenous attention.

1. I write a fixed point iteration procedure to find the equilibrium level of attention for a given set of parameters.

2. I write a function which calculates the log-likelihood of a given set of parameters under determinacy or indeterminacy, using the Kalman filter. This log-likelihood takes into account the equilibrium level of attention implied by that set of parameters.

3. I combine the likelihood with the prior to compute the posterior.

4. I use a numerical procedure to find the mode of the posterior distribution within the determinate region and the indeterminate region. In order to force the function to remain within the determinate or indeterminate region, I use a penalty function.

5. I calculate the inverse Hessian of the posterior distribution at the mode. If the mode is close to the boundary of the region, and so affected by the penalty function, I find the Hessian at the mode under an alternative specification for attention, and/or an alternative sub-sample. I scale the Hessian by a constant c^2. I chose c to give a c.30-40% acceptance rate.[56]

6. I use the mode and scaled inverse Hessian as inputs for a random walk Metropolis Hastings algorithm to find the unconditional posterior distribution of each parameter. I take 110,000 draws, discarding the first 10,000 as burn-in. I find the mean and 90 percent credibility interval. I use Geweke's (1999) modified harmonic mean to estimate the marginal likelihood.

1.G Additional Empirical Results

I now present the results from the specification using de-trended log real GDP per capita as an output gap measure. The results are rather similar to the baseline specification. Once again, the mean estimate for ϕ_π is well below unity for the Great Inflation period, above unity for the longer Great Moderation sub-sample, and just below unity for the

[56]This acceptance rate is achieved in most sub-samples for most specifications, but the acceptance rate did turn out slightly below this range in some indeterminate sub-samples.

post-1990 sub-sample. Again, the evidence against indeterminacy is "very strong" in all sub-samples. Two differences worth noting are: (i) the estimate for θ tends to be slightly higher than in the baseline specification, which is perhaps unsurprising given that the output gap series is more volatile than in the baseline specification, and (ii) the estimate for ϕ_y is somewhat lower, and more in line with estimates typically obtained in the literature.

I also present a robustness check showing that the evidence in favour of determinacy is strong even halving the firm cost of attention k_f to 0.75 to demonstrate that the results are robust to different specifications of the attention cost.

I then show results over the entire sample from 1960:I to 2007:IV under the assumption of determinacy, using the alternative output gap specification. I first estimate the model under endogenous discounting. I the show results imposing that $m_f = m_c = 1$, which corresponds to the rationality case. Finally, I present results imposing that $m_f = m_{f,d} = 0.85$ and $m_c = m_{c,d} = 0.85$ by assuming a very high cost of attention for both consumers and firms. This corresponds to the exogenous attention case, albeit with the attention parameters calibrated rather than estimated given the concerns regarding weak identification in the eoxgenous attention model discussed in the main text.

Table 1.9: Posterior Distribution - Alternative Output Gap Specification

	Great Inflation (1960:I to 1979:II)				Great Moderation (a) (1984:I to 2007:IV)				Great Moderation (b) (1990:I to 2007:IV)			
	Determinacy		Indeterminacy		Determinacy		Indeterminacy		Determinacy		Indeterminacy	
Parameter	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval
ϕ_π	0.68	[0.52,0.84]	0.47	[0.32,0.59]	1.49	[0.90,2.14]	0.18	[0.08,0.30]	0.93	[0.43,1.51]	0.18	[0.08,0.31]
ϕ_y	0.30	[0.21,0.42]	0.27	[0.19,0.37]	0.47	[0.33,0.63]	0.28	[0.05,0.45]	0.49	[0.35,0.67]	0.37	[0.26,0.48]
θ	0.92	[0.88,0.95]	0.89	[0.87,0.92]	0.92	[0.90,0.94]	0.90	[0.87,0.92]	0.92	[0.90,0.94]	0.89	[0.87,0.91]
γ	2.04	[1.41,2.79]	2.16	[1.53,2.97]	2.71	[2.00,3.53]	3.02	[2.18,3.84]	2.46	[1.77,3.28]	2.80	[2.08,3.60]
ρ_η	0.75	[0.67,0.82]	0.85	[0.80,0.90]	0.63	[0.54,0.72]	0.83	[0.59,0.91]	0.61	[0.50,0.70]	0.87	[0.82,0.91]
ρ_z	0.77	[0.70,0.85]	0.72	[0.64,0.80]	0.89	[0.84,0.92]	0.84	[0.76,0.91]	0.87	[0.81,0.92]	0.81	[0.74,0.88]
ρ_r	0.71	[0.62,0.79]	0.70	[0.61,0.78]	0.86	[0.83,0.90]	0.87	[0.82,0.93]	0.87	[0.83,0.91]	0.84	[0.80,0.88]
σ_η^ε	0.11	[0.07,0.15]	0.08	[0.05,0.12]	0.08	[0.06,0.10]	0.05	[0.03,0.10]	0.08	[0.06,0.10]	0.04	[0.03,0.06]
σ_z^ε	0.84	[0.59,1.17]	1.04	[0.71,1.49]	0.43	[0.32,0.56]	0.57	[0.36,0.78]	0.44	[0.31,0.59]	0.62	[0.43,0.88]
σ_v^ε	0.16	[0.14,0.19]	0.16	[0.14,0.18]	0.12	[0.10,0.13]	0.13	[0.11,0.15]	0.10	[0.08,0.11]	0.10	[0.09,0.11]
σ_ζ			0.13	[0.06,0.23]			0.10	[0.05,0.17]			0.10	[0.05,0.17]
m_f	0.97	[0.94,0.99]	0.98	[0.98,0.99]	0.86	[0.85,0.91]	0.97	[0.95,0.99]	0.86	[0.85,0.89]	0.97	[0.96,0.98]
m_c	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]
Log ML		-128.4		-141.7		15.6		-42.3		17.3		-25.4

Table 1.10: Posterior Distribution - Low Attention Cost Specification ($k^f = 0.75$)

	Great Inflation (1960:I to 1979:II)				Great Moderation (a) (1984:I to 2007:IV)				Great Moderation (b) (1990:I to 2007:IV)			
	Determinacy		Indeterminacy		Determinacy		Indeterminacy		Determinacy		Indeterminacy	
Parameter	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval
ϕ_π	0.65	[0.54,0.75]	0.57	[0.46,0.66]	1.61	[1.04,2.28]	0.26	[0.12,0.41]	0.97	[0.51,1.53]	0.24	[0.12,0.38]
ϕ_y	0.52	[0.40,0.66]	0.49	[0.38,0.62]	0.74	[0.53,0.97]	0.66	[0.47,0.85]	0.70	[0.53,0.90]	0.58	[0.42,0.73]
θ	0.89	[0.84,0.93]	0.87	[0.83,0.90]	0.90	[0.86,0.94]	0.90	[0.87,0.92]	0.91	[0.87,0.94]	0.89	[0.86,0.91]
γ	2.93	[2.15,3.81]	3.10	[2.29,4.06]	3.84	[2.94,4.88]	4.61	[3.66,5.70]	3.33	[2.49,4.28]	3.82	[2.95,4.77]
ρ_η	0.75	[0.67,0.81]	0.77	[0.69,0.84]	0.64	[0.54,0.73]	0.80	[0.72,0.86]	0.60	[0.50,0.70]	0.78	[0.69,0.85]
ρ_z	0.80	[0.73,0.86]	0.78	[0.70,0.85]	0.90	[0.87,0.93]	0.89	[0.84,0.92]	0.88	[0.84,0.92]	0.87	[0.82,0.92]
ρ_r	0.56	[0.46,0.66]	0.55	[0.45,0.65]	0.84	[0.80,0.87]	0.84	[0.80,0.87]	0.83	[0.78,0.87]	0.81	[0.76,0.85]
σ_η^ε	0.10	[0.07,0.14]	0.10	[0.07,0.014]	0.07	[0.05,0.09]	0.05	[0.03,0.07]	0.08	[0.06,0.10]	0.05	[0.04,0.08]
σ_z^ε	0.49	[0.37,0.64]	0.53	[0.39,0.71]	0.28	[0.22,0.35]	0.35	[0.27,0.45]	0.29	[0.22,0.37]	0.34	[0.25,0.45]
σ_v^ε	0.16	[0.14,0.19]	0.16	[0.14,0.19]	0.12	[0.11,0.14]	0.12	[0.11,0.14]	0.10	[0.09,0.12]	0.10	[0.09,0.12]
σ_ζ			0.12	[0.06,0.20]			0.08	[0.05,0.12]			0.08	[0.05,0.13]
m^f	0.98	[0.97,1.00]	0.98	[0.97,0.99]	0.92	[0.85,0.97]	0.98	[0.97,0.99]	0.91	[0.85,0.96]	0.98	[0.97,0.99]
m^c	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]	0.85	[0.85,0.85]
Log ML		-53.9		-57.2		82.4		43.7		62.3		35.9

Table 1.11: Posterior Distribution - Full Sample from 1960:I to 2007:IV - Alternative Output Gap

	Endogenous Attention		$m_f = m_c = 1$		$m_f = m_c = 0.85$	
Parameter	Mean	90-pct. interval	Mean	90-pct. interval	Mean	90-pct. interval
ϕ_π	1.09	[0.83,1.37]	1.12	[0.87,1.38]	1.08	[0.82,1.35]
ϕ_y	0.53	[0.38,0.71]	0.52	[0.38,0.68]	0.52	[0.37,0.69]
θ	0.95	[0.93,0.97]	0.95	[0.93,0.97]	0.95	[0.93,0.96]
γ	3.13	[2.39,3.98]	3.38	[2.58,4.28]	3.24	[2.48,4.11]
ρ_η	0.83	[0.79,0.87]	0.83	[0.79,0.87]	0.85	[0.80,0.89]
ρ_z	0.83	[0.78,0.87]	0.83	[0.79,0.88]	0.84	[0.79,0.89]
ρ_r	0.85	[0.82,0.89]	0.84	[0.81,0.87]	0.85	[0.82,0.88]
σ_η^ε	0.05	[0.04,0.07]	0.05	[0.04,0.06]	0.08	[0.07,0.10]
σ_z^ε	0.94	[0.73,1.19]	0.73	[0.56,0.94]	0.93	[0.72,1.19]
σ_v^ε	0.20	[0.18,0.22]	0.20	[0.18,0.22]	0.20	[0.18,0.22]
m^f	0.99	[0.98,1.00]	1.00	[1.00,1.00]	0.85	[0.85,0.85]
m^c	0.85	[0.85,0.85]	1.00	[1.00,1.00]	0.85	[0.85,0.85]
Log ML		-247.3		-251.1		-242.8

Chapter 2

Heterogeneity in Imperfect Inflation Expectations: Theory and Evidence from a Novel Survey

Abstract Using novel survey data from Germany, we study heterogeneity in how households form inflation expectations. We elicit (i) uncertainty in perceptions of current inflation, and (ii) how persistent households perceive inflation to be. Combining these with standard survey questions on inflation, we infer laws of motion for expectations at the individual level. Based on averages alone, a standard model calibrated to our data predicts inflation shocks generate small and transitory responses in expectations and consumption. The considerable heterogeneity we observe in expectation formation, however, amplifies the transmission to aggregate consumption by an order of magnitude, and substantially increases its persistence.

2.1 Introduction

Households' information and subjective models of inflation shape how their inflation expectations respond to macroeconomic shocks. However, commonly used survey data on expectations is consistent with various combinations of information and subjective models, with contrasting implications. The well-known regressions in Coibion and Gorodnichenko (2015a), for example, are consistent with models in which information is noisy but households know the true law of motion for inflation (Coibion and Gorodnichenko, 2015a), or models with full information but misspecified forecasting rules (Gabaix, 2020; Hajdini, 2020). How then do households form their expectations?

This paper uses novel survey data to answer that question. We add new questions to the Bundesbank's Survey on Consumer Expectations to elicit the uncertainty in household perceptions of current inflation, and how persistent households perceive inflation to be. When combined with responses to other standard survey questions, this allows us to infer subjective laws of motion and details of information processing at the individual level.

We find that, on average, uncertainty about current inflation is low, and the perceived persistence of inflation is close to that of realized inflation. However, these averages mask considerable heterogeneity, which is important for understanding aggregate behavior. Calibrating a standard consumption-saving model to our data, heterogeneity in the expectations process amplifies the aggregate consumption response to inflation shocks by an order of magnitude, relative to the representative-agent case. Expectations may therefore play a substantially larger role in business cycle fluctuations than implied by models based on aggregate expectations alone.[1]

How do these questions pin down households' subjective models of inflation? Suppose that each household i believes inflation follows the AR(1) process:[2]

$$\begin{aligned} \pi_{t+1} &= \tilde{\rho}_i \pi_t + \varepsilon_{t+1} \\ \varepsilon_{t+1} &\sim N(0, \tilde{\sigma}^2_{\varepsilon,i}) \end{aligned} \tag{2.1}$$

Suppose further that in each period, the household receives a noisy signal $s_{i,t}$ about current inflation:

$$\begin{aligned} s_{i,t} &= \pi_t + q_{i,t} \\ q_{i,t} &\sim N(0, \sigma^2_q) \end{aligned} \tag{2.2}$$

[1] e.g. Fuster et al. (2010), Bhandari et al. (2019), Angeletos et al. (2020), and many others.

[2] Numerous studies suggest household forecasts are well-characterized by such simple forecasting rules (e.g. Adam, 2007). We relax this assumption in Appendix 2.A.

We assume that households perceive the variance of $q_{i,t}$ to be $\tilde{\sigma}^2_{q,i}$, which is not necessarily equal to σ^2_q (as in e.g. Broer and Kohlhas, 2019).[3] To the extent that households perceive the signals received to be noisy, they are uncertain about the current inflation rate.

Uncertainty in the inflation forecast can then stem from two sources. First, the household can be uncertain about future inflation because they are uncertain about current inflation, and they believe inflation to be persistent. Second, the household can be uncertain about future inflation because they expect there to be shocks in the future. Formally, we can decompose uncertainty about future inflation as follows:[4]

$$\tilde{Var}_{i,t}(\pi_{t+1}) = \tilde{\rho}_i^2 \tilde{Var}_{i,t}(\pi_t) + \tilde{\sigma}^2_{\varepsilon,i} \tag{2.3}$$

Of the terms in this equation, existing survey questions only measure subjective uncertainty in the forecast $\tilde{Var}_{i,t}(\pi_{t+1})$. However, different combinations of the two components of this uncertainty imply very different aggregate dynamics, for the same observed forecast uncertainty. To illustrate, consider three ways of decomposing a given $\tilde{Var}_{i,t}(\pi_{t+1})$, assuming that the true law of motion for inflation is persistent but stationary:

1. $\tilde{\rho}_i = 1$, $\tilde{\sigma}^2_{\varepsilon,i} = 0$, $\tilde{Var}_{i,t}(\pi_t) = \tilde{Var}_{i,t}(\pi_{t+1})$. Agents are uncertain in their perceptions, suggesting imperfect information about current inflation. This implies their inflation perceptions respond sluggishly to shocks. Agents overextrapolate from their perceptions to expectations, generating 'delayed overshooting' (Angeletos et al., 2020).

2. $\tilde{\rho}_i = 1$, $\tilde{\sigma}^2_{\varepsilon,i} = \tilde{Var}_{i,t}(\pi_{t+1})$, $\tilde{Var}_{i,t}(\pi_t) = 0$. Agents overextrapolate, but are certain about perceived current inflation, so there is no delay in the response of perceptions and expectations to realized inflation.

3. $\tilde{\rho}_i = 0$, $\tilde{\sigma}^2_{\varepsilon,i} = \tilde{Var}_{i,t}(\pi_{t+1})$, $\tilde{Var}_{i,t}(\pi_t)$ undefined. Agents underextrapolate. Forecasts do not respond to current inflation.

A given level of uncertainty in inflation forecasts could be consistent with a continuum of models between these extremes. Our questions pin down the correct model at the household level by eliciting the variance of the inflation perception $\tilde{Var}_{i,t}(\pi_t)$, and the perceived persistence $\tilde{\rho}_i$. Combined with existing questions eliciting $\tilde{Var}_{i,t}(\pi_{t+1})$, we then calculate the perceived variance of the innovations $\tilde{\sigma}^2_{\varepsilon,i}$.

To measure uncertainty in perceptions, we start with an existing survey question eliciting a point estimate of the respondent's inflation perception. We then add a new

[3] This allows the model to be consistent with any survey respondents who are simultaneously incorrect, but very certain, about their inflation perception.

[4] Note this abstracts from uncertainty about $\tilde{\rho}_i$, as is common in models with parameter learning (see e.g. Bullard and Suda, 2016).

question, asking for the probability that the inflation rate lies within a specified range around that point estimate. Fitting a triangular distribution to these responses yields an estimate of $\tilde{Var}_{i,t}(\pi_t)$. To measure $\tilde{\rho}_i$, we present respondents with hypothetical scenarios of macroeconomic shocks, as in Andre et al. (2022). We specify that the shock has increased current inflation by one percentage point, and ask how the respondent would update their inflation expectations as a result. These novel questions are critical to answer our question: even with panel data on $\tilde{Var}_{i,t}(\pi_{t+1})$, as in the US Survey of Consumer Expectations, it is not possible to jointly identify $\tilde{Var}_{i,t}(\pi_t)$, $\tilde{Var}_{i,t}(\pi_{t+1})$, and $\tilde{\rho}_i$ at the individual level without restrictions on other aspects of the expectations process.

Our finding that, on average, respondents are less uncertain about current inflation than future inflation suggests that much of the uncertainty in expectations comes from perceived noise in the inflation process, not a lack of information about current inflation. However, a minority of households are very uncertain in their inflation perceptions. Perceived persistence is similarly heterogeneous: while two-thirds of households perceive no inflation persistence at the one-year horizon, those who do update their expectations after the hypothetical shock often do so by a large amount.

This heterogeneity has a large impact on the dynamics of aggregate consumption. In a standard consumption-saving model, the individual consumption response to inflation is convex in perceived inflation persistence. The same path of aggregate inflation expectations is therefore associated with much larger fluctuations in aggregate consumption if the individual-level subjective laws of motion are heterogeneous. The response of aggregate consumption to an inflation shock is consequently an order of magnitude larger when we calibrate the model using the heterogeneity observed in our data, relative to a representative-agent calibration based on average parameters. The persistence of aggregate consumption also rises by nearly one-half.

Moreover, elements of household expectation formation correlate systematically with each other, and with household characteristics, further distorting aggregate consumption away from the representative-agent case. In particular, hand-to-mouth households with little liquid wealth are more uncertain about future inflation relative to current inflation, and believe inflation is substantially more persistent. Since these households are likely to be less responsive to changes in their expectations, this dimension of heterogeneity further underlines the importance of investigating beyond average parameters of expectation formation for understanding macroeconomic behavior.

Related Literature. Angeletos et al. (2020) also study imperfect inflation expectations, directly estimating impulse responses of aggregate expectations to shocks. Our approach is complementary. While they can capture richer subjective models than the linear approximations we identify, our approach reveals heterogeneity in expectation for-

mation, which cannot be observed with average forecasts.

One key difference relative to Angeletos et al. (2020) is that our study assumes that agents follow a misspecified forecasting rule. In particular, we assume that households in the survey assume an AR(1) process for inflation, which we do not suppose to be the true law of motion for inflation. In that respect, our study is related to the heuristic switching literature (see Hommes (2021) for a survey). Models in this literature often assume that households either forecast inflation returning immediately to target or fully extrapolate the latest realization (e.g. De Grauwe, 2012; Hommes and Lustenhouwer, 2019). This would be similar to a special case of our expectations model in which $\tilde{\rho}_i \in \{0, 1\}$.[5]

In another related study, Ryngaert (2018) estimates perceived persistence and signal precision among professional forecasters, but cannot uncover heterogeneity in these parameters.

There are large literatures measuring noisy information and perceived persistence separately, using data from surveys (e.g. Branch, 2004; Coibion and Gorodnichenko, 2012, 2015a; Boneva et al., 2020; Laudenbach et al., 2021; Link et al., 2023) and lab experiments (e.g. Adam, 2007; Beshears et al., 2013; Caplin and Dean, 2015; Afrouzi et al., 2023). Consistent with our results, these papers frequently find substantial heterogeneity in the aspect of expectation formation they study.[6] We contribute to these literatures by directly measuring the perceived persistence of inflation simultaneously with variance in inflation perceptions and expectations, all at the individual level. While previous studies typically infer at least one of these objects from assumptions on information or subjective laws of motion, we show that identifying them separately is important to understand the dynamics of aggregate expectations and consumption. In particular, to the best of our knowledge there are no existing quantitative measures of the variance in inflation perceptions.[7]

Finally, we contribute to the literature on the role of expectations in business cycles.

[5]Note though that the way in which we model consumption decisions differs from the approach often used in the heuristic switching literature. Commonly in this literature the expectation terms in the aggregate IS and Phillips curves are simply replaced with the average inflation expectation implied by the fraction of agents who extrapolate the latest realization and the fraction who forecast output and inflation returning to target. As is explained in the appendix to Hommes and Lustenhouwer (2019), this approach implicitly assumes that all agents believe that their probability of following a given heuristic in the next period is independent of the heuristic they have used in the past. As such, supposing that nominal rates are at their steady state, that approach does not imply the same level of consumption for a given level of lagged inflation as our partial equilibrium approach below.

[6]Coibion and Gorodnichenko (2012, 2015a) find no evidence of heterogeneous signal-to-noise ratios. However, their tests are derived under the assumption that forecasters believe inflation follows a random walk, contrary to the low average perceived persistence in our sample, and in other surveys (Jain, 2017; Ryngaert, 2018).

[7]Armona et al. (2019) ask households to rate their uncertainty over past house-price growth on a 1-5 scale, but do not relate this quantitatively to the variance in expectations. The uncertainty over *future* inflation elicited in existing surveys is qualitatively different, as that inflation is not yet realized.

Like us, Branch and Evans (2006), Hommes and Lustenhouwer (2019), Macaulay (2022), Pedemonte et al. (2023) (among others) find that heterogeneity in the expectation formation process can substantially alter macroeconomic outcomes. We directly measure the relevant heterogeneity, and show that the resulting distribution of expectation processes amplifies the aggregate consumption response to inflation by an order of magnitude.

Of course, our approach is not the only possible way to capture the effects of heterogeneity in expectation formation we identify. While data on market-based uncertainty (Bauer et al., 2022) or disagreement between agents (Mankiw et al., 2004; Andrade et al., 2016) cannot identify individual-level properties of expectation formation, with further structural assumptions it may be possible to use questions in existing individual-level surveys to estimate our parameters. We are not aware of any studies currently taking this approach for both information and subjective laws of motion simultaneously. For example, Jain (2017) estimates perceived inflation persistence using the panel dimension of the US Survey of Professional Forecasters, but does not disentangle the uncertainty over current and future inflation. That survey, however, contains density forecasts of inflation over very short horizons, which could be used to approximate the uncertainty over current inflation we measure for households.

We pursue our approach over such a method for two reasons. First, short-horizon density forecasts are typically only available for professional forecasters, who are substantially better-informed about current economic events than households (Link et al., 2023). Second, our direct measurements allow us to document some of the key properties of expectation heterogeneity with less restrictive structural assumptions. For example, methods inferring perceived persistence by comparing expectations at different horizons frequently assume households hold identical long-run expectations (Reis, 2020; Andre et al., 2021), and that there is no 'forward information' about future inflation innovations (Goldstein and Gorodnichenko, 2022). Our method is robust to relaxing both of these assumptions (Appendix 2.A.2), and with no panel dimension we also avoid survey tenure effects (Crossley et al., 2017), which can be large for questions on inflation (Kim and Binder, 2022).

2.2 Expectations Framework

2.2.1 The Agent

Each agent i believes inflation follows the AR(1) process in equation 2.1. This subjective law of motion for inflation may or may not coincide with the true data generating process. The agent receives noisy signals about the current inflation rate as in equation 2.2, and

as above their perception of the noise in the signal may differ from the actual noise in the signal.

After observing the signal, agents update their perception of inflation using the steady-state Kalman filter.[8] The posterior one-period ahead inflation forecast is:

$$\tilde{E}_{i,t}\pi_{t+1} = (1-\chi_i)\tilde{\rho}_i\tilde{E}_{i,t-1}\pi_t + \tilde{\rho}_i\chi_i(\pi_t + q_{i,t}) \tag{2.4}$$

$$\chi_i = 1 - \frac{V_i^p}{V_i^f} \tag{2.5}$$

where χ_i is agent i's Kalman gain. V_i^p and V_i^f denote respectively the steady-state subjective variances of perceived inflation ($\tilde{Var}_{i,t}(\pi_t)$) and one-period ahead inflation ($\tilde{Var}_{i,t}(\pi_{t+1})$). These are such that $V_i^f \geq V_i^p$. All derivations for results here, and in Section 2.5, are in Appendix 2.A.

The formula for χ_i is intuitive: the lower is V_i^p relative to V_i^f, the more informative is the signal received, and hence the larger the update the agent optimally makes to their inflation perceptions and forecasts. Note that the forecast variance in a given period, as measured in several existing surveys, does not place any restrictions on the Kalman gain. This is why we require a novel question, not present in existing surveys, to identify V_i^p alongside V_i^f.

Using equation 2.1, $\tilde{\sigma}^2_{\varepsilon,i}$ is given by:

$$\tilde{\sigma}^2_{\varepsilon,i} = V_i^f - \tilde{\rho}_i^2 V_i^p \tag{2.6}$$

Finally, the perceived signal noise is:

$$\tilde{\sigma}^2_{q,i} = \frac{V_i^f V_i^p}{V_i^f - V_i^p} \tag{2.7}$$

Our novel questions measuring V_i^p and $\tilde{\rho}_i$ therefore allow us to infer all parameters of the law of motion for inflation expectations (equation 2.4), and the variances of both fundamental shocks and signal noise.

Of course, interpreting our data in this way assumes that an AR(1) process is a good description of household beliefs about the law of motion for inflation, as documented in e.g. Adam (2007); Goldstein and Gorodnichenko (2022). In Appendix 2.A.2 we extend the model to a richer set of subjective laws of motion, which may include other variables (e.g. output, interest rates), longer lags, and heterogeneous long-run expectations. We

[8]The assumption of steady-state filtering is required to identify the Kalman gain in the absence of panel data. The same assumption is commonly made in the rational inattention literature for tractability (Mackowiak and Wiederholt, 2009).

show that Equations 2.5 and 2.7 still capture household i's Kalman gain and perceived noise variance respectively, and all the transmission channels discussed below continue to operate in this extended version of the model. The main effect of the extension of the law of motion that we consider is that equation 2.6 no longer captures the perceived variance of inflation shocks, but rather gives a composite of all sources of uncertainty that are not related to current inflation.

2.2.2 Expectations Impulse Responses

For simplicity, assume inflation is indeed an exogenous AR(1) process, with true autocorrelation ρ:

$$\pi_t = \rho\pi_{t-1} + \varepsilon_t \tag{2.8}$$

We consider the impulse response of expectations to a one percentage-point shock to inflation at $t = 0$, with inflation and inflation expectations at steady-state (zero) before the shock. Abstracting from the effect of realized signal noise $q_{i,t}$,[9] the one-period ahead inflation forecast of agent i, t periods after the shock, is:

$$\tilde{E}_{i,t}\pi_{t+1} = \tilde{\rho}_i\chi_i\frac{\rho^{t+1} - (1-\chi_i)^{t+1}\tilde{\rho}_i^{t+1}}{\rho - (1-\chi_i)\tilde{\rho}_i} \tag{2.9}$$

Different combinations of χ_i and $\tilde{\rho}_i$ therefore imply very different impulse responses of expectations, even for the same V_i^f. On impact, the response of expectations is increasing in χ_i and $\tilde{\rho}_i$. The persistence of the expectation response increases in $\tilde{\rho}_i$, but decreases in χ_i. If $\tilde{\rho}_i$ is sufficiently large, and χ_i sufficiently small, then expectations display hump-shaped impulse responses (as observed in e.g. Angeletos et al., 2020). Equally, as in Angeletos et al. (2020), if $\tilde{\rho}_i > \rho$ then expectations overshoot, rising above realized inflation some periods after the shock.

2.2.3 The Role of Heterogeneity

If agents were homogeneous, equation 2.9 would also describe the impulse response of aggregate inflation expectations to the shock. With heterogeneity in $\tilde{\rho}_i$ and χ_i, however, the initial response of aggregate inflation expectations to the shock becomes:

$$\tilde{E}_0\pi_1 = E[\tilde{\rho}_i]E[\chi_i] + Cov(\tilde{\rho}_i, \chi_i) \tag{2.10}$$

[9]This captures the average expectation across many agents who share the same $\tilde{\rho}_i$ and χ_i.

A positive correlation between $\tilde{\rho}_i$ and χ_i therefore amplifies the initial effect of the shock on expectations, because those who extrapolate the most from perceived to expected inflation also update their perceptions the most in the period of the shock.

Heterogeneity continues to affect aggregate expectations in the periods after the shock. In Appendix 2.A, we show that heterogeneity in $\tilde{\rho}_i$ increases the persistence of the response of aggregate expectations to the shock.

2.3 Data

We use the November 2021 wave of the Bundesbank-Online-Panel-Households survey, which is administered online to a representative sample of the German population. 4110 households were asked our questions.

In the main survey, households give a point estimate of the inflation rate over the past 12 months, and give both point and density forecasts of inflation over the next 12 months. The density forecast involves households filling out the probabilities of inflation falling within various intervals, as in other recent household surveys (Armantier et al., 2017; Coibion et al., 2021). Additionally, a range of household characteristics are collected. We report summary statistics in Appendix 2.B.

We add two novel questions for the November 2021 wave, reproduced in Table 2.1 (see Appendix 2.B for the German translations seen by respondents, and the point and density forecast questions).

Question 1 elicits the uncertainty in the household's perceptions of current inflation.[10] The high and low inflation values seen by the respondent are their point estimate of current inflation, ± 1 percentage-point. If the respondent's point inflation estimate is $\geq 5\%$, this range is widened to ± 2 percentage-points, as uncertainty in inflation expectations is known to rise with point estimates (Ben-David et al., 2018). Answers are in percent, and must be within $[0, 100]$. Respondents also see a note giving further explanation of the question (see Appendix 2.B).

To calculate the variance of perceived inflation V_i^p, we fit a symmetric triangular distribution using the respondent's answer and their point estimate. This is similar to the approach in Coibion et al. (2021), and in the Survey of Consumer Expectations when respondents only report positive probabilities in two bins of a density forecast question (Armantier et al., 2017). Note that this is different from the method used to measure V_i^f, the uncertainty over future inflation, which uses all the information in the density forecast for one-year ahead inflation.

[10] We do not use a multiple-bin density forecast, as for inflation expectations, as these questions are cognitively demanding. Too many can result in households dropping out of the survey.

Table 2.1: Novel questions added to the BOP-HH survey in November 2021

Question	Text	Sample
1	Now we would like to know how certain you are about your information on the inflation rate or deflation rate over the past 12 months ([Value of point estimate])%. In your opinion, how likely is it that the inflation rate has been between [Low inflation level]% and [High inflation level]% over the past twelve months?	All respondents
Respondents randomly shown one of three scenarios before Question 2		
General	Imagine the following hypothetical situation: Due to an unexpected economic event, the inflation rate increased by one percentage point in the past year.	Group A
Supply	Imagine the following hypothetical situation: Due to unexpected problems with local production technology in the Middle East, the price of crude oil rose in the past year, causing the inflation rate to rise by one percentage point.	Group B
Demand	Imagine the following hypothetical situation: Due to increased defense spending, government spending rose unexpectedly more than usual in the past year, causing the inflation rate to rise by one percentage point. The change is temporary and occurs even though the government's assessment of national security or economic conditions has not changed. In addition, taxes do not change in response to the spending program.	Group C
2	In this situation, would you adjust your inflation expectations for the next 12 months as stated in the first part of the questionnaire? If so, to what extent?	All respondents

In computing the Kalman gain for each respondent, we take the ratio of these two variances (equation 2.5). To confirm that the difference in measurement approaches does not bias our results, in Appendix 2.C.2 we construct an alternative measure of V_i^f, which utilizes less of the available information from the density forecasts, but which corresponds closely to the measurement of V_i^p. All results below are robust to this alternative. This is consistent with the hypothesis in Kumar et al. (2022) that the discrepancy they find between their triangular and density variance measures is driven by a difference in the treatment of the end-points of each distribution, which is not present here.[11]

[11]Specifically, Kumar et al. (2022) fit a triangular distribution using questions about firms' most optimistic, and most pessimistic, growth expectations. As they use these as the end-points of the distribution, they are assigned zero probability mass. In contrast, their density forecast question allows firms to place positive probability mass on these highest and lowest forecasts. Our measure of V_i^p does

Question 2 elicits perceived inflation persistence $\tilde{\rho}_i$. Following Andre et al. (2022), respondents are given a hypothetical scenario describing an exogenous shock, and asked how they would expect that to affect future inflation. Unlike Andre et al. (2022), in each scenario we tell the respondents that the shock caused current inflation to increase by 1 percentage-point. Their answers on how that would change their inflation expectations consequently reflect their estimates of inflation persistence, not their predictions of the immediate impact of the shock, which Andre et al. (2022) find to be heterogeneous across households.

In Section 2.2 we did not distinguish between different types of shocks, and our main empirical analysis will do the same. However, households may associate different shocks with different levels of persistence. To investigate this, we randomly split respondents into three groups. The first group are not told the nature of the shock, the second see a hypothetical supply shock (oil price), and the third see a demand shock (government spending). The specific scenarios are adapted from Andre et al. (2022).

To answer, respondents see:

1) Yes, from [Value of point estimate]% to __%
2) No

and either input a number in the first line or select 'No'.

Using these questions we obtain $\tilde{\rho}_i$ and V_i^p for each respondent. We then infer the implied χ_i, $\tilde{\sigma}^2_{\varepsilon,i}$, and $\tilde{\sigma}^2_{q,i}$ using equations 2.5-2.7. Full details of the variable construction are in Appendix 2.C.1.

2.4 Empirical Results

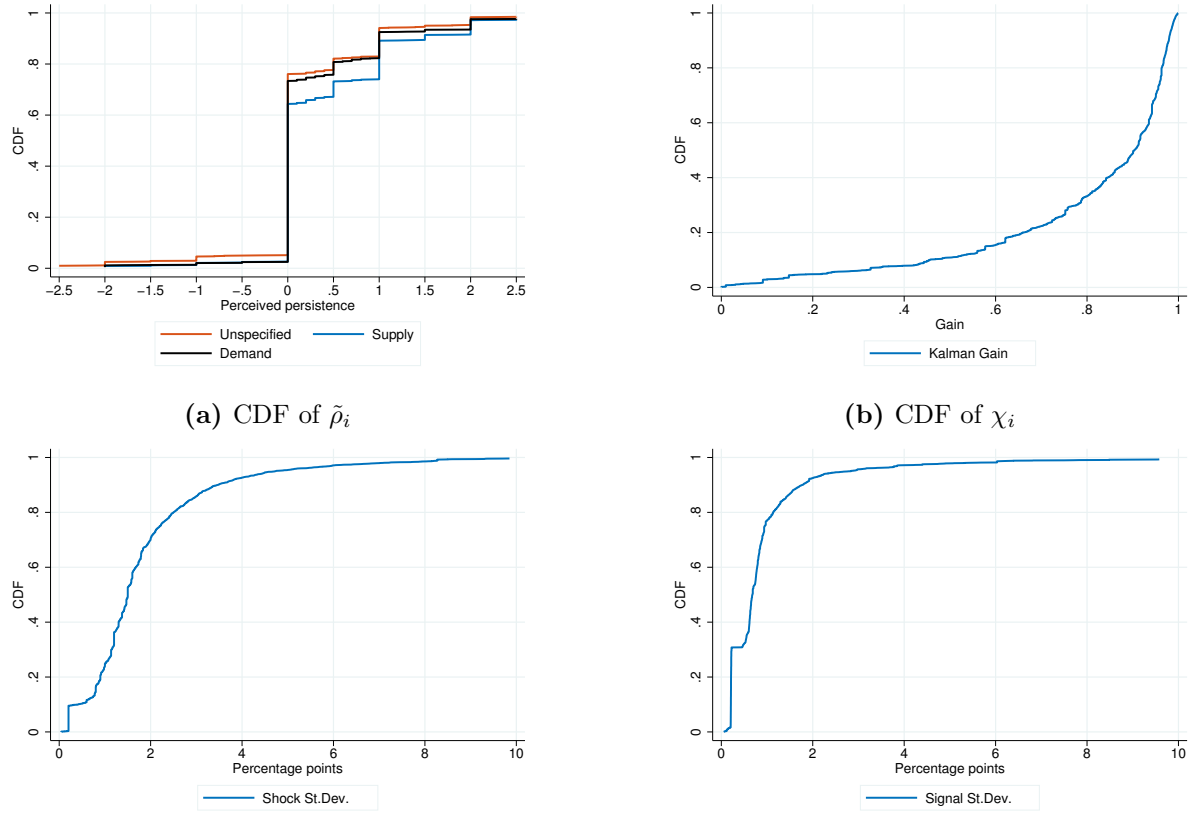
Figure 2.1a plots the CDF of $\tilde{\rho}_i$, truncated to remove the approximately 1% of responses outside $[-5, 5]$. Of the remaining responses, 89% report $\tilde{\rho}_i \in [0, 1]$, and 68% do not revise their expectations at all.

Result 1 *Conditional on perceiving persistence in* $[0, 1]$*, the average perceived persistence is broadly consistent with the data. However, the cross-sectional heterogeneity is large.*

Of those with $\tilde{\rho}_i \in [0, 1]$, the mean $\tilde{\rho}_i$ is 0.18, close to the 'correct' answer of 0.21 based on recent German data.[12] Including all responses in $[-5, 5]$, the mean $\tilde{\rho}_i$ is 0.29. The heterogeneity, however, is large. Restricting to responses in [0,1], the maximum possible standard deviation is 0.5: the standard deviation in our sample is 0.36.

not involve elicited end-points, so we avoid this.

[12]This is the coefficient from a linear projection of annual CPI inflation on its lagged value (data from *www.destatis.de*, 2002-2021).



(a) CDF of $\tilde{\rho}_i$

(b) CDF of χ_i

(c) CDF of $\tilde{\sigma}_{\varepsilon,i}$

(d) CDF of $\tilde{\sigma}_{q,i}$

Figure 2.1: CDFs of key parameters. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

In Appendix 2.D we show that this heterogeneity is not driven by households rounding to the nearest percentage-point when reporting expectations (as studied in Binder, 2017). We continue with the full sample here, as rounded expectations may still matter for consumption decisions. The large fraction of households choosing not to update expectations is consistent with other information treatments in previous waves of this survey (Dräger et al., 2022), and with other hypothetical scenario surveys (Christelis et al., 2021; Fuster et al., 2021).

Splitting by the shock type presented to respondents, the mean $\tilde{\rho}_i$ measurements for those within the $[0, 1]$ interval are 0.16 (unspecified shock), 0.22 (supply shock), and 0.16 (demand shock). Supply shocks are therefore perceived to be slightly more persistent than demand shocks, consistent with evidence that supply shocks are particularly important for household inflation expectations (Coibion and Gorodnichenko, 2015b). While some of this difference could stem from the stronger reference to the temporary nature of the demand shock in our scenarios, note that the unspecified scenario is very similar to the supply shock in this respect; neither state explicitly that the shock is temporary. The fact that average $\tilde{\rho}_i$ is higher for supply shocks than *both* alternatives supports the

interpretation that supply shocks are perceived to be more persistent.

Figure 2.1b shows the CDF of χ_i.[13]

Result 2 *The average Kalman gain is high, at 0.8. There is considerable cross-sectional heterogeneity.*

The high average Kalman gain stems from most consumers being considerably more certain about their inflation perceptions than their expectations. There is, however, a long tail of very uncertain households, with low Kalman gains.

The average χ_i exceeds values obtained from regressions on average forecast errors, which are typically close to 0.5 (e.g. Coibion and Gorodnichenko, 2015a). The discrepancy is unsurprising, since such regressions yield biased estimates if agents hold inaccurate beliefs about inflation persistence (Ryngaert, 2018). Moreover, at the time of the survey, CPI inflation in Germany exceeded 5%. Consumers were plausibly better-informed about inflation than in other periods because of media coverage.

Figure 2.1c shows the CDF of $\tilde{\sigma}_{\varepsilon,i}$. There is considerable heterogeneity; a tail of households believe future inflation is extremely volatile. This reinforces that much of the uncertainty in inflation expectations relates to future shocks, rather than uncertainty about current inflation. Finally, Figure 2.1d shows the CDF of $\tilde{\sigma}_{q,i}$. Reflecting the high average χ_i, most households have little noise in their signals, though a minority have very imprecise information. Further distributions, including those of V_i^p and V_i^f used to calculate χ_i, and the relationships of these variables with point inflation forecasts, are presented in Appendix 2.D.

2.4.1 Relationships between Expectation Components

Table 2.2 shows our next main result:

Result 3 *Households who are more uncertain about current inflation are also more uncertain about future inflation, believe inflation shocks are more volatile, and have lower Kalman gains.*

This is consistent with noisy information models, in which greater uncertainty arises when households process less information. Importantly, this result is not imposed by our structural assumptions. Our model allows more uncertainty about current inflation to be

[13]For households who are completely certain that π_t is within the interval shown in our question, V_i^p cannot be point-identified, so we obtain a range of possible values for χ_i (details in Appendix 2.C.1). Figure 2.1b uses the mid-point in these cases, dropping respondents with a range of width >0.2. The CDFs using the upper and lower bounds on χ_i are shown in Appendix 2.D. Figures 2.1c and 2.1d similarly use midpoints of implied ranges in these cases.

Table 2.2: Cross-sectional correlations of subjective law of motion elements.

	(1)				
	$SD_i(\pi_{t+1})$	$SD_i(\pi_t)$	$\tilde{\rho}_i$	$SD_i(\varepsilon_{t+1})$	χ_i
$SD_i(\pi_{t+1})$	1.000				
$SD_i(\pi_t)$	0.476***	1.000			
$\tilde{\rho}_i$	0.036*	-0.015	1.000		
$SD_i(\varepsilon_{t+1})$	0.988***	0.443***	-0.034	1.000	
χ_i	0.316***	-0.393***	0.038*	0.337***	1.000

Note: Bundesbank-Online-Panel-Households, November 2021 wave. For cases where χ_i is set-identified, respondents are excluded if the parameters are estimated very imprecisely (range$>$ 0.2). For all remaining set-identified parameters, the mid-point of the range is used. Observations of $SD_i(\pi_{t+1})$, and $SD_i(\pi_t)$ below the 1st or above the 99th percentile of that variable's distribution are also excluded as outliers, as are observations of $\tilde{\rho}_i$ outside $[-5, 5]$ (c.1% of observations). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

associated with higher or lower Kalman gains, depending on the relationships between V_i^p and V_i^f (equation 2.5). Indeed, noisy information can only partly explain the distribution of uncertainty in the data. More uncertain households also believe that inflation shocks are more volatile (higher $\tilde{\sigma}_{\varepsilon,i}$). This further highlights the importance of measuring the uncertainty in perceptions and expectations separately.

There are small positive correlations of $\tilde{\rho}_i$ with uncertainty about future inflation, and with Kalman gains. However, in Appendix 2.D we break this down, and find that households who believe inflation is non-stationary ($|\tilde{\rho}_i| \geq 1$) are qualitatively different from households with $|\tilde{\rho}_i| < 1$. Within households who believe inflation is persistent and stationary, greater perceived persistence is associated with less uncertainty about current and future inflation, less perceived noise in the inflation process, and greater Kalman gains. That is consistent with models of endogenous information acquisition; if inflation is more persistent, information about current inflation is more valuable. We also split the sample by shock scenario, but find little difference across shock types.

2.4.2 Correlations with Household Characteristics

As our application in Section 2.5.2 is to consumption, Table 2.3 shows results from regressing each component of expectation formation on household characteristics known to relate to Marginal Propensities to Consume (MPCs). The key variables are liquid wealth (bank deposits plus securities), illiquid wealth (property plus firm ownership), other wealth, debt, and household income. There is also an indicator for if the household is hand-to-mouth, defined here as having liquid wealth of less than €1250.

Since households with $\tilde{\rho}_i = 0$ may differ qualitatively from those with $\tilde{\rho}_i \neq 0$, the final column restricts the sample to those with $\tilde{\rho}_i \neq 0$. This gives the estimated associations conditional on the household revising expectations in light of the hypothetical shock.[14] In Appendix 2.D we show that selection into $\tilde{\rho}_i \neq 0$ is not significantly related to wealth or income.

The first row of coefficients shows our next main result:

Result 4 *Hand-to-mouth households are more uncertain about future inflation, but no more uncertain about current inflation, than other households. They believe inflation is noisier and more persistent, and have 8.6% higher Kalman gains on average.*

Table 2.3: Regressions of components of subjective laws of motion on household characteristics.

	(1) $\log(SD_i(\pi_{t+1}))$	(2) $\log(SD_i(\pi_t))$	(3) $\log(SD_i(\varepsilon_{t+1}))$	(4) $\log(\chi_i)$	(5) $\tilde{\rho}_i$
Hand-to-mouth	0.1407***	0.0296	0.1518***	0.0859**	0.2716**
	(0.0451)	(0.0384)	(0.0484)	(0.0432)	(0.1326)
Liquid wealth	0.0001	-0.0001	0.0001*	0.0002***	0.0001
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0002)
Illiquid wealth	-0.0000	-0.0000	-0.0000	0.0000	-0.0001
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)
Other wealth	0.0001	-0.0003*	0.0001	-0.0000	0.0001
	(0.0002)	(0.0001)	(0.0002)	(0.0003)	(0.0003)
Debt	0.0001	-0.0000	0.0001	0.0002*	-0.0001
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0003)
log(income)	-0.0972***	-0.0826***	-0.1130***	-0.0120	-0.0035
	(0.0293)	(0.0261)	(0.0313)	(0.0364)	(0.0856)
HH Controls	Yes	Yes	Yes	Yes	Yes
Observations	1900	1900	1900	1900	567
R^2	0.0661	0.0603	0.0591	0.0237	0.0703

Note: Bundesbank-Online-panel-Households, November 2021 wave. The units of the wealth and debt variables are €1000s. The household controls are age (in years up to a top bin of ≥ 80, coded as 80), age^2, gender, region (north/south/east/west), education, occupation category, and employment status (all categorical, for details see the full questionnaire at https://www.bundesbank.de/en/bundesbank/research/survey-on-consumer-expectations/questionnaires-850746). All controls except age and age^2 are treated as categorical. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[14] This is equivalent to the linear component of a hurdle model in which, unlike Cragg (1971), we do not impose that $\tilde{\rho}_i$ is truncated to be ≥ 0.

Angeletos et al. (2020) find that aggregate inflation expectations display delayed overshooting, suggesting households perceive inflation to be more persistent than implied by the true data-generating process. Our results suggest that aggregate over-persistence may partly reflect the expectations of hand-to-mouth households, who are less able to respond to expected inflation by adjusting consumption and saving. Their expectations may not therefore have much impact on aggregate dynamics. While high Kalman gains for hand-to-mouth households is inconsistent with simple models of rational inattention, this result could be driven by those who are close to leaving their borrowing constraints, who have highly non-linear policy functions and value information as a result (Broer et al., 2021).

Above the hand-to-mouth threshold, higher liquid wealth has little relationship with expectation formation. Higher liquid wealth is associated with a statistically significant increase in the Kalman gain, but this is quantitatively small: each further €1000 is associated with a 0.02% rise in χ_i.

The final row of coefficients gives our next main result:

Result 5 *Higher income is associated with less uncertainty about current and future inflation, and less perceived noise in the inflation process, but is not associated with differences in perceived inflation persistence or Kalman gains.*

As documented in other contexts (e.g. Ben-David et al., 2018), higher-income households are less uncertain about future inflation: a 10% rise in household income is associated with a 1% fall in $SD_i(\pi_{t+1})$. There is however no evidence that this comes from high-income households acquiring more precise information, as they are also less uncertain about current inflation. Since the differences across current and future uncertainty are small, there is no significant relationship between income and χ_i. Rather, the bulk of the lower uncertainty for high-income households is explained by them believing the inflation process is less volatile. A 10% rise in income is associated with a 1.1% reduction in $SD_i(\varepsilon_{t+1})$. There is no significant correlation between income and $\tilde{\rho}_i$.

2.5 Implications for Aggregate Consumption

We now analyse how our results affect aggregate dynamics through a simple partial-equilibrium model with both unconstrained and hand-to-mouth households.

2.5.1 Consumption-Saving Model

Unconstrained households have an infinite horizon and no borrowing constraint. They choose consumption $\hat{c}_{i,t}$ to maximize the expected discounted sum of CRRA utility over

consumption, and invest any unspent exogenous income $y_{i,t}$ in risk-free one-period bonds with gross nominal interest rate i_t. The log-linearized consumption function is:[15]

$$\hat{c}_{i,t} = \sum_{h=0}^{\infty} \beta^h \left((1-\beta)\tilde{E}_{i,t} y_{i,t+h} - \beta\gamma^{-1}\tilde{E}_{i,t} i_{t+h} + \beta\gamma^{-1}\tilde{E}_{i,t}\pi_{t+h+1} \right) \tag{2.11}$$

where β is the discount factor and γ is the coefficient of relative risk aversion (see Gabaix (2020) Proposition 29). To isolate the effect of a shock to expected inflation, we hold expected $y_{i,t+h}$ and i_{t+h} constant in all exercises (relaxed in Appendix 2.E). Using equation 2.1, consumption is:[16]

$$\hat{c}_{i,t} = \frac{\beta\gamma^{-1}}{1-\beta\tilde{\rho}_i}\tilde{E}_{i,t}\pi_{t+1} \tag{2.12}$$

A higher $\tilde{\rho}_i$ therefore increases the responsiveness of consumption to expected inflation, as it implies larger changes in longer-horizon expectations.

Using equation 2.9 and aggregating across households, the aggregate consumption response to a one percentage-point inflation shock in $t = 0$ is:

$$\hat{c}_0 = \beta\gamma^{-1}\left(E[\chi_i]E\left[\frac{\tilde{\rho}_i}{1-\beta\tilde{\rho}_i}\right] + Cov\left(\chi_i, \frac{\tilde{\rho}_i}{1-\beta\tilde{\rho}_i}\right)\right) \tag{2.13}$$

Heterogeneity in expectation formation therefore affects aggregate consumption in two ways. First, heterogeneity in $\tilde{\rho}_i$ amplifies the aggregate consumption responses to inflation, because $\tilde{\rho}_i/(1-\beta\tilde{\rho}_i)$ is convex in $\tilde{\rho}_i$. If even a few households believe that inflation is close to a unit root, they respond very strongly to current inflation, generating large aggregate consumption responses. Note that heterogeneity in $\tilde{\rho}_i$ amplifies the consumption response *relative* to the response of inflation expectations. Similar aggregate impulse responses in inflation expectations may therefore correspond to very different impulse responses in consumption.

Second, any correlation between $\tilde{\rho}_i$ and χ_i will further distort the aggregate consumption response away from the representative-agent case. Intuitively, the response of aggregate consumption is amplified if the households who obtain precise information about the shock are also the ones who respond most strongly to that information. This is an example of the 'narrative heterogeneity channel' discussed in Macaulay (2022).

Constrained (hand-to-mouth) households, in contrast, do not respond to expectations.

[15]Note that by studying a log-linearized model, we abstract away from any direct effects of uncertainty on consumption. The uncertainties measured in the survey are still important, however, because they determine the responsiveness of expectations to shocks through the Kalman gain χ_i (equation 2.5).

[16]This assumes $|\beta\tilde{\rho}_i| < 1$. When calibrating to the data we drop the minority of households for whom this is not true.

Since we abstract from indirect effects of nominal shocks through incomes, they have $\hat{c}_{i,t} = 0$.

2.5.2 Aggregate Shock Transmission

We now generate impulse responses of aggregate one-year ahead inflation expectations and consumption in three cases. First, we consider full information rational expectations (FIRE): all households know that $\rho = 0.21$, and observe π_t precisely. Second, we maintain homogeneity, and calibrate the model using the population averages for χ_i and $\tilde{\rho}_i$ in the survey. Finally, we allow for heterogeneity, calibrating to the observed joint distribution of χ_i and $\tilde{\rho}_i$.[17]

Figure 2.2a plots the IRFs for $\tilde{E}_t \pi_{t+1}$. Expectations respond less on impact in both cases calibrated to the survey data than with FIRE, because $E[\chi_i] < 1$ and $E[\tilde{\rho}_i] \approx \rho$. While on impact expectations in these two cases are similar, they are somewhat more persistent under heterogeneity: a year after the shock, average $\tilde{E}_{i,t} \pi_{t+1}$ is c.35% greater than under homogeneity.

Figure 2.2b plots the IRFs for aggregate consumption. They differ considerably between cases:

Result 6 *The model-implied consumption response under heterogeneity is* $15.5\times$ *greater on impact than under homogeneity. The persistence of the consumption response under heterogeneity* ($\frac{\hat{c}_1}{\hat{c}_0}$) *is c.45% greater than under homogeneity.*

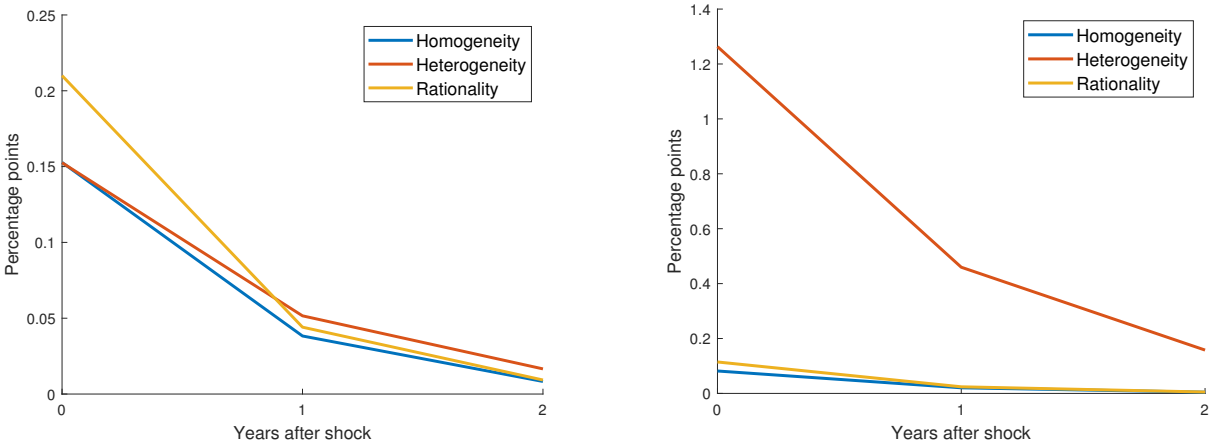
On impact, the FIRE response is small (0.11%). It has persistence of 0.21, so the deviation from steady-state in $t = 1$ is negligible. The homogeneous case has an even smaller initial consumption response (0.08%), and marginally greater persistence.

The heterogeneous case, however, has a vastly larger initial consumption response of 1.26%. The response is also more persistent, with a persistence of 0.36 between $t = 0$ and $t = 1$. Aggregate consumption therefore remains substantially above steady-state in the two years following the shock. In Appendix 2.E, we show that these effects principally reflect the heterogeneity in $\tilde{\rho}_i$.

These figures demonstrate the challenges involved in inferring how expectations affect macroeconomic dynamics using only aggregate or consensus expectations data. The homogeneity and heterogeneity cases yield similar IRFs in aggregate inflation expectations, but entirely different IRFs in aggregate consumption.

Discussion. The model studied here is simple, which serves two purposes: it highlights how heterogeneity in expectation formation can have large effects even in textbook

[17] To exclude outliers, in all cases we exclude observations with $\tilde{\rho}_i \notin [0, 1]$. The excluded households are disproportionately the hand-to-mouth, who are least able to respond to expectations (Result 4).



(a) Expectations **(b)** Consumption

Figure 2.2: Implied IRFs of one-period ahead inflation expectations and consumption. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

models, and it allows us to study the mechanisms behind those effects analytically. However, to achieve this we have made a number of simplifying assumptions. In particular, the results rely on the properties of the consumption function (equation 2.11), which assumes households follow an Euler equation. While we cannot test this directly in our data, Dräger and Nghiem (2021) document such behavior for German households. Hanspal et al. (2021) also find that perceived persistence is important for consumption decisions in the context of Covid-19 expectations.

The amplification result survives a number of extensions and robustness checks. Appendix 2.A.2 demonstrates that the key channels continue to operate if households use richer subjective laws of motion. Additionally, Appendix 2.E shows that amplification from heterogeneity remains large when the model is calibrated separately to results for each of the shock scenarios in survey Question 2, when we exclude rounded expectations, and when we allow expectations of future nominal interest rates to react to expected inflation.[18]

In general equilibrium, these effects will be amplified if the consumption increase leads to rising real incomes for hand-to-mouth households, and rising income expectations for unconstrained households. A further round of general equilibrium effects may then also occur through the Phillips curve. We leave exploration of these effects to future research.

[18]Figures 2.2a and 2.2b can be interpreted as the results if households believe nominal interest rates are unresponsive to inflation. This is a reasonable baseline, since at the time of the survey, CPI inflation in Germany had risen sharply over the preceding year (from -0.3% to 5.2%), but the ECB was still a long way from raising nominal interest rates (Lagarde, 2021).

2.6 Conclusion

Inflation expectations are important in many theories of the business cycle. However, the quantities measured by existing expectation surveys are consistent with many different laws of motion for expectations, with contrasting aggregate implications. To distinguish between these models, we use novel survey data to elicit (i) households' uncertainty over current inflation, and (ii) how persistent they perceive inflation to be, at the individual level.

We find that, on average, consumers are relatively confident about current inflation, and perceive little persistence in inflation. However, these averages mask considerable heterogeneity, which increases the aggregate consumption response to an inflation shock by an order of magnitude in an otherwise standard consumption-savings model. The persistence of consumption responses to shocks also increases substantially.

This effect occurs because individual consumption functions are highly non-linear in the components of expectation formation. Heterogeneity in those parameters, and correlations between them, can therefore have large effects on aggregate consumption, even for a given path of aggregate expectations. The components of expectation formation are also correlated with household wealth and income, both of which correlate with consumption behavior (Kaplan et al., 2014; Kueng, 2018). Exploring the distribution of these components for expectations of other variables, and how the distributions change over time and states of the world, could be a fruitful avenue for future research.

Appendices to Chapter 2

2.A Proofs

2.A.1 Baseline Model

Steady-State Kalman Filter

Before receiving the signal at time t, the agent's subjective distribution for π_t and the signal is:

$$\begin{pmatrix} \pi_t \\ \pi_t + q_{i,t} \end{pmatrix} \sim N\left(\begin{pmatrix} \tilde{\rho}_i \tilde{E}_{i,t-1}\pi_{t-1} \\ \tilde{\rho}_i \tilde{E}_{i,t-1}\pi_{t-1} \end{pmatrix}, \begin{pmatrix} V^f_{i,t-1} & V^f_{i,t-1} \\ V^f_{i,t-1} & V^f_{i,t-1} + \tilde{\sigma}^2_{q,i} \end{pmatrix} \right) \tag{2.14}$$

The conditional mean of π_t given the signal is then:

$$\tilde{E}_{i,t}\pi_t = (1-\chi_{i,t})\tilde{\rho}_i \tilde{E}_{i,t-1}\pi_{t-1} + \chi_{i,t}(\pi_t + q_{i,t}), \text{ where } \chi_{i,t} = \frac{V^f_{i,t-1}}{V^f_{i,t-1} + \tilde{\sigma}^2_{q,i}} \tag{2.15}$$

The conditional variance of π_t:

$$V^p_{i,t} = V^f_{i,t-1}\left(1 - \frac{(V^f_{i,t-1})^2}{V^f_{i,t-1}(V^f_{i,t-1} + \tilde{\sigma}^2_{q,i})}\right) = \frac{V^f_{i,t-1}\tilde{\sigma}^2_{q,i}}{V^f_{i,t-1} + \tilde{\sigma}^2_{q,i}}$$

In steady state, this variance is:

$$V^p_i = \frac{V^f_i \tilde{\sigma}^2_{q,i}}{V^f_i + \tilde{\sigma}^2_{q,i}} \tag{2.16}$$

The steady state Kalman gain is then:

$$\chi_i = \frac{V^f_i}{V^f_i + \tilde{\sigma}^2_{q,i}} = 1 - \frac{V^p_i}{V^f_i} \tag{2.17}$$

Response of Inflation Expectations to Shocks

Throughout, we assume that both inflation and the agent's inflation perception start in steady state in $t = -1$. That is, $\pi_{-1} = \tilde{E}_{i,-1}\pi_{-1} = 0$. The individual inflation perception

is given by equation 2.15. Iterating backwards to time 0, we obtain:

$$\tilde{E}_{i,t}\pi_t = \chi_i \sum_{s=0}^{t}((1-\chi_i)\tilde{\rho}_i)^s(\pi_{t-s} + q_{i,t-s}) \tag{2.18}$$

Abstracting from $q_{i,t}$, the h period ahead forecast is then:

$$\tilde{E}_{i,t}\pi_{t+h} = \chi_i\tilde{\rho}_i^h \sum_{s=0}^{t}((1-\chi_i)\tilde{\rho}_i)^s\pi_{t-s} = \chi_i\tilde{\rho}_i^h\rho^t \sum_{s=0}^{t}((1-\chi_i)\tilde{\rho}_i\rho^{-1})^s\varepsilon_0 \tag{2.19}$$

where the second equality uses that $\pi_{t-s} = \rho^{t-s}\varepsilon_0$.

Provided that $(1-\chi_i)\tilde{\rho}_i\rho^{-1} \neq 1$, then evaluating the summation and rearranging yields the result:

$$\tilde{E}_{i,t}\pi_{t+h} = \chi_i\tilde{\rho}_i^h \frac{\rho^{t+1} - ((1-\chi_i)\tilde{\rho}_i)^{t+1}}{\rho - (1-\chi_i)\tilde{\rho}_i}\varepsilon_0 \tag{2.20}$$

Setting $h = 1, \varepsilon_0 = 1$ yields equation 2.9.

Persistence of Expectations

From equation 2.20 we have:

$$\frac{\tilde{E}_{i,1}\pi_2}{\tilde{E}_{i,0}\pi_1} = \rho + (1-\chi_i)\tilde{\rho}_i \tag{2.21}$$

The persistence of the response of expectations to the shock is increasing in ρ and $\tilde{\rho}_i$, and decreasing in χ_i (assuming $\tilde{\rho}_i > 0$). If $\rho + (1-\chi_i)\tilde{\rho}_i > 1$, then the expectation rises between the period the shock hits and the period after, giving a hump-shaped response. This condition is both necessary and sufficient for a hump-shaped impulse response.

Under heterogeneity, the impact response of aggregate inflation expectations to the shock is given by equation 2.10. To see the role of heterogeneity in future periods, we consider the special case of $\rho = 0$, in which case:

$$\tilde{E}_t\pi_{t+1} = E[\tilde{\rho}_i^{t+1}]E[\chi_i(1-\chi_i)^t] + Cov(\tilde{\rho}_i^{t+1}, \chi_i(1-\chi_i)^t) \tag{2.22}$$

$\tilde{\rho}_i^{t+1}$ is strictly convex in $\tilde{\rho}_i$ for all $t \geq 1$, so heterogeneity in $\tilde{\rho}_i$ increases the time t response for $t \geq 1$. This convexity increases with t, so the persistence of the response of expectations also increases with heterogeneity in $\tilde{\rho}_i$. In addition, note that $\chi_i(1-\chi_i)^t$ is linear in χ_i for $t = 0$, concave in χ_i if $t = 1$, but becomes convex in χ_i as t becomes large. Heterogeneity in χ_i consequently decreases persistence for small t, but may increase persistence for large enough t.

To understand the role of the covariance term, consider the simple case where $\tilde{\rho}_i$ is monotonically increasing in χ_i. The covariance term is then positive for small t, but

negative for large t. As such, a positive correlation between $\tilde{\rho}_i$ and χ_i tends to result in lower persistence of the response in inflation expectations, despite a larger initial response.

Formally, the persistence of aggregate expectations between $t = 0$ and $t = 1$ is:

$$\frac{\tilde{E}_1\pi_2}{\tilde{E}_0\pi_1} = \rho + \left(E[\tilde{\rho}_i] + \frac{Var(\tilde{\rho}_i)}{E[\tilde{\rho}_i]}\right)\left(1 - E[\chi_i] - \frac{Var[\chi_i]}{E(\chi_i)}\right) + \text{ covariance terms} \tag{2.23}$$

This highlights that heterogeneity in $\tilde{\rho}_i$ helps to generate hump-shaped IRFs in expectations, while heterogeneity in χ_i makes hump-shaped responses less likely to emerge.

Consumption responses

Consider an unconstrained agent facing an infinite-horizon consumption savings problem. As in Gabaix (2020), take income as given. The consumption function is:

$$\hat{c}_{i,t} = \sum_{h\geq 0} \beta^h((1-\beta)\tilde{E}_{i,t}\hat{y}_{i,t+h} - \beta\gamma^{-1}\tilde{E}_{i,t}i_{t+h} + \beta\gamma^{-1}\tilde{E}_{i,t}\pi_{t+h+1}) \tag{2.24}$$

Since we hold expected income and nominal interest rates at steady state, we have $\tilde{E}_{i,t}\hat{y}_{i,t+h} = 0$ and $\tilde{E}_{i,t}i_{t+h} = 0$ for all t and h. The consumption function then reduces to:

$$\hat{c}_{i,t} = \beta\gamma^{-1}\sum_{h\geq 0}\beta^h\tilde{E}_{i,t}\pi_{t+h+1} = \beta\gamma^{-1}\frac{1}{1-\beta\tilde{\rho}_i}\tilde{E}_{i,t}\pi_{t+1} \tag{2.25}$$

To proceed, substitute in for the one period ahead expectation in time t using equation 2.20:

$$\hat{c}_{i,t} = \beta\gamma^{-1}\frac{1}{1-\beta\tilde{\rho}_i}\tilde{\rho}_i\chi_i\frac{\rho^{t+1} - (1-\chi_i)^{t+1}\tilde{\rho}_i^{t+1}}{\rho - (1-\chi_i)\tilde{\rho}_i} + d_{i,t} \tag{2.26}$$

Here, $d_{i,t}$ is an idiosyncratic noise term, which is a linear function of $q_{i,t}$, $q_{i,t-1}$, ...,$q_{i,0}$, and so has mean zero. Averaging across agents, one obtains:

$$\hat{c}_t = \beta\gamma^{-1}E\left[\frac{1}{1-\beta\tilde{\rho}_i}\tilde{\rho}_i\chi_i\frac{\rho^{t+1} - (1-\chi_i)^{t+1}\tilde{\rho}_i^{t+1}}{\rho - (1-\chi_i)\tilde{\rho}_i}\right] \tag{2.27}$$

Which in $t = 0$ becomes:

$$\hat{c}_0 = \beta\gamma^{-1}E\left[\frac{\tilde{\rho}_i}{1-\beta\tilde{\rho}_i}\chi_i\right] \tag{2.28}$$

Applying the definition of a covariance then leads to equation 2.13.

2.A.2 Extended Subjective Laws of Motion

Suppose households hold a subjective law of motion for inflation of the form:

$$\pi_{t+1} = \tilde{\rho}_i \pi_t + \xi_t + \varepsilon_{t+1} \tag{2.29}$$

$$\varepsilon_{t+1} \sim N(0, \tilde{\sigma}^2_{\varepsilon i}) \tag{2.30}$$

where ξ_t is (perceived to be) a normally distributed random variable. Inflation and this new term are therefore jointly normal:

$$\begin{pmatrix} \pi_t \\ \xi_t \\ \varepsilon_{t+1} \end{pmatrix} \sim N\left(\begin{pmatrix} \tilde{\mu}_{\pi i} \\ \tilde{\mu}_{\xi i} \\ 0 \end{pmatrix}, \begin{pmatrix} \tilde{\sigma}^2_{\pi i} & \tilde{r}_i & 0 \\ \tilde{r}_i & \tilde{\sigma}^2_{\xi i} & \tilde{r}_{\varepsilon i} \\ 0 & \tilde{r}_{\varepsilon i} & \tilde{\sigma}^2_{\varepsilon i} \end{pmatrix} \right) \tag{2.31}$$

where $\tilde{r}_i$ is household i's perception of the covariance between π_t and ξ_t, and $\tilde{r}_{\varepsilon i}$ is the perceived covariance between ξ_t and ε_{t+1}. $\tilde{\mu}_{\pi i}, \tilde{\mu}_{\xi i}$ are the perceived means of each variable. From equation 2.29, $\tilde{\mu}_{\xi i} = (1 - \tilde{\rho}_i)\tilde{\mu}_{\pi i}$.

ξ_t could include, for example, longer lags of inflation, long-run inflation expectations, or realizations of other variables such as output or interest rates. The key restriction is that ξ_t is included additively in the law of motion, and is normally distributed. This is necessary to ensure that the Kalman filter remains the appropriate way for households to interpret signals on inflation.

With this extended law of motion, Bayesian updating implies that household i uses $s_{i,t}$ to form perceptions according to:[19]

$$\tilde{E}_{i,t}\pi_t = (1 - \chi_i)(\tilde{\rho}_i \tilde{E}_{i,t-1}\pi_{t-1} + \tilde{E}_{i,t-1}\xi_t) + \chi_i(\pi_t + q_{it}) \tag{2.32}$$

$$\tilde{E}_{i,t}\xi_t = \tilde{E}_{i,t-1}\xi_t + \frac{\tilde{r}_i}{V_i^f}\chi_i(\pi_t + q_{i,t} - \tilde{E}_{i,t-1}\pi_t) \tag{2.33}$$

where:

$$\chi_i = \frac{V_i^f}{V_i^f + \tilde{\sigma}^2_{q,i}} \tag{2.34}$$

The Kalman gain, reflecting the sensitivity of inflation perceptions to realized inflation, therefore has exactly the same form as in our baseline model (equation 2.15). The only difference is that V_i^f now also includes uncertainty from ξ_t, and any related covariance terms. Since uncertainty about π_t is still Gaussian, V_i^p is calculated exactly as in equation 2.16. We can therefore measure the Kalman gain using equation 2.5, which implies the

[19] As in the baseline model, we restrict attention to the steady state Kalman filter here.

perceived signal noise is given by equation 2.7.

The inclusion of ξ_t in the subjective law of motion does however affect other aspects of perceptions. The prior belief over inflation in equation 2.32 now includes the prior perception of ξ_t, and if $\tilde{r}_i \neq 0$ household i uses the inflation signal to update beliefs about ξ_t as well as π_t. In addition, the variance calculated in equation 2.6 will now measure:

$$V_i^f - \tilde{\rho}_i^2 V_i^p = \tilde{\sigma}_{\varepsilon i}^2 + \tilde{\sigma}_{\xi i}^2 + 2\tilde{\rho}_i \tilde{r}_i + 2\tilde{r}_{\varepsilon i} \tag{2.35}$$

That is, this measurement combines uncertainty from inflation shocks (as in the baseline model) with uncertainty arising from the extra term in the law of motion, including from its covariance with either inflation or shocks.

We now turn to the response of expectations to inflation shocks in this extended model. As in Section 2, suppose that in period -1, household i's expectations of all variables are at steady state ($E_{i,-1}\pi_{-1} = \tilde{\mu}_{\pi i}, E_{i,-1}\xi_{-1} = \tilde{\mu}_{\xi i}$). In that case, abstracting from signal noise as in equation 2.9, we have:

$$\tilde{E}_{i,0}\pi_1 = \left(\tilde{\rho}_i(1-\chi_i) - \frac{\tilde{r}_i}{V_i^f}\chi_i\right)\tilde{\mu}_{\pi i} + \tilde{\mu}_{\xi i} + \left(\tilde{\rho}_i\chi_i + \frac{\tilde{r}_i}{V_i^f}\chi_i\right)\pi_0 \tag{2.36}$$

The change in the aggregate inflation expectation when π_0 rises by one unit is then:[20]

$$\Delta\tilde{E}_0\pi_1 = E[\tilde{\rho}_i]E[\chi_i] + Cov(\tilde{\rho}_i, \chi_i) + E[\tilde{r}_i/V_i^f]E[\chi_i] + Cov(\tilde{r}_i/V_i^f, \chi_i) \tag{2.37}$$

The first two terms are exactly as in equation 2.10, so our key channels are still present. However, they are now supplemented with two additional channels, reflecting the same mechanisms for expectations of ξ_t whenever $\tilde{r}_i \neq 0$ for at least some households.

Specific examples of extended subjective laws of motion

Suppose ξ_t reflects the long-run mean of inflation. For simplicity, assume that while long-run expectations may be heterogeneous across households (heterogeneous $\tilde{\mu}_{\xi i}$), they are constant at the household level ($\tilde{\sigma}_{\xi i}^2 = 0$). In that case, $\tilde{r}_i = 0$ for all households, as ξ_t has no variance. From equations 2.36, expectations are therefore exactly the same as in the baseline case, plus an individual-specific constant term. This constant does not affect the response of expectations to shocks (equation 2.37), so all results on shock transmission are identical to those in Sections 2.2 and 2.5. The same will apply if ξ_t reflects the forward information studied in Goldstein and Gorodnichenko (2022).

However in other cases the new terms in equation 2.37 will be non-zero. For example,

[20] We use this rather than the raw aggregate inflation expectation because steady state priors are now no longer necessarily zero or equal across households in this extended model. If $\tilde{\mu}_{\pi i} = \tilde{\mu}_{\xi i} = 0$ for all i then this exercise is identical to that in equation 2.10.

households who believe in a wage-price spiral may believe that current incomes affect future inflation ($\xi_t = \kappa_i y_t$), and that current income comoves positively with current inflation ($\tilde{r}_i \geq 0$). Through equation 2.37, that implies that the channels studied in Section 2.2 understate the response of aggregate inflation expectations to shocks, especially if the households who believe in a very strong spiral (high $\tilde{r}_i$) are systematically those obtaining more precise signals about inflation (high χ_i).

2.B Survey details

2.B.1 Summary statistics

Table 2.4 shows summary statistics for the key variables used in our analysis, and several other household characteristics. The construction of $SD_i(\pi_t), SD_i(\pi_{t+1}), SD_i(\varepsilon_{t+1}), \chi_i, \tilde{\rho}_i$ is described in Appendix 2.C.1.

Income and wealth variables are reported in bins. We take the mid-point of each bin. We code the lowest bin for income as if the lower bound is zero, and again take the midpoint (all wealth variables have a separate bin for zero). The top bin is coded as if it had the same width as the second-highest bin. Liquid wealth is (bank deposits + securities). Illiquid wealth is (property + firm ownership). Debt is (secured + unsecured debt). A respondent is classified as hand-to-mouth if their liquid wealth is $< €1250$.

Table 2.4: Summary statistics for expectations (point estimates and components), respondent characteristics, and income/wealth.

	Mean	Std Dev.	Min	Max
Panel A: Expectations				
$\tilde{E}_{i,t}\pi_t$	4.10	2.61	0	30
$\tilde{E}_{i,t}\pi_{t+1}$	4.90	4.99	-3	60
$SD_i(\pi_t)$	1.75	4.70	0.41	40.41
$SD_i(\pi_{t+1})$	1.72	1.35	0.30	8.80
$SD_i(\varepsilon_{t+1})$	1.87	1.52	0.04	12.12
χ_i	0.80	0.23	0	1
$\tilde{\rho}_i$	0.29	0.84	-5	5
Panel B: Demographics				
Age	56.87	14.66	16	80
Female	0.37	0.48	0	1
Higher Education	0.39	0.49	0	1
Is Working	0.55	0.50	0	1
Panel C: Income and Wealth				
Income	3.95	1.97	0.25	11
Liquid Wealth	90.49	154.90	0	1250
Illiquid Wealth	315.38	383.64	0	2375
Other Wealth	12.41	48.75	0	625
Debt	47.89	109.34	0	955
Owns Securities	0.62	0.48	0	1
Hand-to-mouth	0.14	0.34	0	1

Note: Bundesbank-Online-Panel-Households, November 2021 wave. For cases where χ_i is set-identified, respondents are excluded if the parameters are estimated very imprecisely (range>0.2). For all remaining set-identified parameters, the mid-point of the range is used. Observations of $\tilde{E}_{i,t}\pi_t, \tilde{E}_{i,t}\pi_{t+1}, SD_i(\pi_{t+1})$, and $SD_i(\pi_t)$ below the 1st or above the 99th percentile of that variable's distribution are also excluded as outliers, as are observations of $\tilde{\rho}_i$ outside $[-5, 5]$ (c.1% of observations). All income and wealth variables are in €1000s, and income refers to monthly net income of the household. Higher Education is an indicator for if the respondent has a bachelor's degree or higher, not including vocational training.

2.B.2 Survey questions (English and German)

Table 2.5 contains the existing questions in the Bundesbank survey eliciting point estimates of current and future inflation, and the density forecast of future inflation.

2.B.3 Novel questions for November 2021 (in German)

Table 2.6 contains the German text of our novel survey questions.

Table 2.5: Existing questions in the BOP-HH survey

Label	Text
Inflation Development	
Question	What do you think the rate of inflation or deflation in Germany was over the past twelve months?
Note	If you assume there was deflation, please enter a negative value. Values may have one decimal place.
Question	Wass denken Sie, wie hoch war die Inflationsrate oder Deflationsrate in den letzten zwölf Monaten in Deutschland?
Note	Im Falle einer angenommenen Deflationsrate tippen Sie bitte einen negativen Wert ein. Die Eingabe maximal einer Nachkommastelle ist möglich. Bitte geben Sie einen Wert hier ein.
Input Field	Percent
Inflation Expectations Qualitative	
Question	Do you think inflation or deflation is more likely over the next twelve months?
Note	Inflation is the percentage increase in the general price level. It is mostly measured using the consumer price index. A decrease in the price level is generally described as "deflation".
Question	Was denken Sie, ist in den kommenden zwölf Monaten eher mit einer Inflation oder einer Deflation zu rechnen?
Note	Inflation ist der prozentuale Anstieg des allgemeinen Preisniveaus. Sie wird meist über den Verbraucher-preisindex gemessen. Ein Rückgang des Preisniveaus wird gemeinhin als "Deflation" bezeichnet.
Input Field	Select one answer
Inflation Expectations Quantitative	
Question	What do you think the rate of inflation/deflation will roughly be over the next twelve months? (select based on answer to *Inflation Expectations Qualitative*)
Note	Inflation is the percentage increase in the general price level. It is mostly measured using the consumer price index. A decrease in the price level is generally described as "deflation".
Question	Was denken Sie, wie hoch wird die Inflationsrate/Deflationsrate in den kommenden zwölf Monaten in etwa sein? (select based on answer to *Inflation Expectations Qualitative*)

Note	Inflation ist der prozentuale Anstieg des allgemeinen Preisniveaus. Sie wird meist über den Verbraucherpreisindex gemessen. Ein Rückgang des Preisniveaus wird gemeinhin als "Deflation" bezeichnet. Bitte tippen Sie einen Wert in das Zahlenfeld ein (eine Nachkommastelle möglich).
Input Field	Percent

Inflation Expectations Probabilistic

Question	In your opinion, how likely is it that the rate of inflation will change as follows over the next twelve months?
Note	The aim of this question is to determine how likely you think it is that something specific will happen in the future. You can rate the likelihood on a scale from 0 to 100, with 0 meaning that an event is completely unlikely and 100 meaning that you are absolutely certain it will happen. Use values between the two extremes to moderate the strength of your opinion. Please note that your answers to the categories have to add up to 100.
Input Field	The rate of deflation (opposite of inflation) will be 12% or higher. __ The rate of deflation (opposite of inflation) will be between 8% and less than 12%. __ The rate of deflation (opposite of inflation) will be between 4% and less than 8%. __ The rate of deflation (opposite of inflation) will be between 2% and less than 4%. __ The rate of deflation (opposite of inflation) will be between 0% and less than 2%. __ The rate of inflation will be between 0% and less than 2%. __ The rate of inflation will be between 2% and less than 4%. __ The rate of inflation will be between 4% and less than 8%. __ The rate of inflation will be between 8% and less than 12%. __ The rate of inflation will be 12% or higher. __

Question	Für wie wahrscheinlich halten Sie es, dass sich die Inflationsrate in den kommenden zwölf Monaten wie folgt entwickelt?
Note	Bei dieser Frage geht es darum, wie Sie die Wahrscheinlichkeit einschätzen, dass ein bestimmter Sachverhalt in der Zukunft eintritt. Ihre Antworten können in einer Spanne zwischen 0 und 100 liegen, wobei 0 absolut unwahrscheinlich bedeutet und 100 absolut sicher. Mit Werten dazwischen können Sie Ihre Einschätzung abstufen. Bitte beachten Sie, dass sich die Angaben über alle Kategorien auf 100 summieren müssen.
Input Field	die Deflationsrate (Gegenteil von Inflation) wird 12% oder höher sein. __ die Deflationsrate (Gegenteil von Inflation) wird zwischen 8% und 12% liegen. __ die Deflationsrate (Gegenteil von Inflation) wird zwischen 4% und 8% liegen. __ die Deflationsrate (Gegenteil von Inflation) wird zwischen 2% und 4% liegen. __ die Deflationsrate (Gegenteil von Inflation) wird zwischen 0% und 2% liegen. __ die Inflationsrate wird zwischen 0% und 2% liegen. __ die Inflationsrate wird zwischen 2% und 4% liegen. __ die Inflationsrate wird zwischen 4% und 8% liegen. __ die Inflationsrate wird zwischen 8% und 12% liegen. __ die Inflationsrate wird 12% oder höher sein. __

Table 2.6: Novel questions added to the BOP-HH survey in November 2021

Question	Text
1	Nun möchten wir wissen, wie sicher Sie sich über Ihre Angabe zur Inflationsrate oder Deflationsrate in den letzten 12 Monaten sind ([Value of point estimate])%. Wie wahrscheinlich ist es Ihrer Meinung nach, dass die Inflationsrate in den letzten zwölf Monaten zwischen [Low inflation level]% und [High inflation level]% lag?
Hinweis	Bei dieser Frage geht es darum, wie Sie die Wahrscheinlichkeit einschätzen, dass die von Ihnen angegebene Inflationsrate oder Deflationsrate in den letzten 12 Monaten tatsächlich ungefähr diesen Wert angenommen hat. Ihre Antworten können zwischen 0 und 100 liegen, wobei 100 bedeutet, dass Sie absolut sicher sind. Kleinere Zahlen bedeuten, dass Sie sich weniger sicher sind.
Input Field	Prozent
Respondents randomly shown one of three scenarios before Question 2	
Group A	Stellen Sie sich die folgende hypothetische Situation vor: Aufgrund eines unerwarteten wirtschaftlichen Ereignisses hat sich die Inflationsrate im vergangenen Jahr um einen Prozentpunkt erhöht.
Group B	Stellen Sie sich die folgende hypothetische Situation vor: Aufgrund von uner-warteten Problemen mit der lokalen Produktionstechnologie im Nahen Osten ist der Rohölpreis im vergangenen Jahr gestiegen, was zu einem Anstieg der Inflationsrate um einen Prozentpunkt geführt hat.
Group C	Stellen Sie sich die folgende hypothetische Situation vor: Aufgrund gestiegener Verteidigungsausgaben sind die Staatsausgaben im vergangenen Jahr unerwartet stärker als üblich gestiegen, was zu einem Anstieg der Inflationsrate um einen Prozentpunkt geführt hat. Die Änderung ist vorübergehend und tritt ein, obwohl sich die Einschätzung der Regierung zur nationalen Sicherheit oder den wirtschaftlichen Bedingungen nicht geändert hat. Darüber hinaus ändern sich die Steuern nicht als Reaktion auf das Ausgabenprogramm.
2	Würden Sie in dieser Situation Ihre im vorderen Teil des Fragebogens genannten Inflationserwartungen für die nächsten 12 Monate anpassen? Wenn ja, inwiefern?
Input Field	1) Ja, von [Value of point estimate]Prozent auf __Prozent 2) Nein

2.C Variable construction

2.C.1 Main variables

To obtain V_i^p, we fit a symmetric triangular distribution to household i's answers:

$$V_i^p = \begin{cases} \frac{1}{6}\left(1 - \sqrt{1 - \frac{x_{1i}}{100}}\right)^{-2} & \text{if } \tilde{E}_i(\pi_t) \in (-5, 5) \\ \frac{2}{3}\left(1 - \sqrt{1 - \frac{x_{1i}}{100}}\right)^{-2} & \text{if } \tilde{E}_i(\pi_t) \notin (-5, 5) \end{cases} \tag{2.38}$$

where x_{1i} is respondent i's response to Question 1. Note that for households who report $x_i = 0$, this method provides an upper bound on their $Var_i(\pi_t)$.

To obtain $\tilde{\rho}_i$, we set $\tilde{\rho}_i = 0$ for households who select 'No' in answer to Question 2. For all others, we set:

$$\tilde{\rho}_i = x_{2i} - \tilde{E}_i(\pi_{t+1}) \tag{2.39}$$

where x_{2i} is respondent i's response to Question 2.

We then calculate V_i^f. For agents who are certain future inflation will lie within one specific bin, we calculate an upper bound on the variance using the symmetric triangular distribution, just as for the perception. The lower bound on V_i^f is given by zero.

For the remaining agents, we calculate V_i^f by taking the midpoints of each of the bins in the probability distribution. Denote these midpoints as z_j for the bins $j = 1, ..., n$. Denote the probability assigned to each bin as p_j. We then calculate the mean:

$$\bar{z}_i = \sum_{j=1}^{n} p_{i,j} z_j \tag{2.40}$$

The variance is then:

$$V_i^f = \sum_{j=1}^{n} p_{i,j}(z_j - \bar{z}_i) \tag{2.41}$$

The calculation of the Kalman gain is complicated by the fact that that for some respondents we have ranges of possible V_i^p or V_i^f, in which case we can only find ranges for χ_i and the other key parameters. We now describe how we calculate these parameters for each such case.

Case (i): V_i^p and V_i^f both point-identified

Calculate χ_i using:

$$\chi_i = 1 - \frac{V_i^p}{V_f^i} \tag{2.42}$$

Back out $\tilde{\sigma}^2_{\varepsilon,i}$ and $\tilde{\sigma}^2_{q,i}$ using:

$$\tilde{\sigma}^2_{\varepsilon,i} = V_i^f - \tilde{\rho}_i^2 V_i^p \tag{2.43}$$

$$\tilde{\sigma}^2_{q,i} = \frac{V_i^f V_i^p}{V_i^f - V_i^p} \tag{2.44}$$

Datapoints are inconsistent with Kalman filtering (and so are dropped) if $\chi_i < 0$ or $\tilde{\sigma}^2_{\varepsilon,i} < 0$.

Case (ii): V_i^f point-identified, V_i^p set-identified

This occurs if the respondent is certain that π_t lies within the specified interval, but places strictly positive probability in multiple intervals in the expectation question. V_i^p is then bounded below by zero, and the upper bound is calculated using the symmetric

triangular distribution as above.

Denote the upper bound on V_i^p by a, so that $V_i^p \in [0, a]$. Under steady state Kalman filtering, it must be that $V_i^p \leq V_i^f$ and $V_i^p \leq \tilde{\rho}_i^{-2} V_i^f$. The latter is more restrictive if $|\tilde{\rho}_i| > 1$. This may shrink the upper bound on V_i^p, and hence raise the lower bound on the Kalman filter. As such, $V_i^p \in [0, \tilde{a}]$, where $\tilde{a}$ is given by:

$$\tilde{a} = \min(V_i^f, \tilde{\rho}_i^{-2} V_i^f, a) \tag{2.45}$$

Then we have the following ranges for the key parameters:

$$\chi_i \in \left[1 - \frac{\tilde{a}}{V_i^f}, 1\right], \; \tilde{\sigma}_{q,i}^2 \in \left[0, \frac{V_i^f \tilde{a}}{V_i^f - \tilde{a}}\right], \; \tilde{\sigma}_{\varepsilon,i}^2 \in \left[V_i^f - \tilde{\rho}_i^2 \tilde{a}, V_i^f\right] \tag{2.46}$$

Case (iii): V_i^f set-identified, V_i^p point-identified

In this case, the consumer is not certain that current inflation lies within the specified interval, but is certain that future inflation lies within one specific interval. As such, V_i^p is known, but $V_i^f \in [0, b]$, where b is given by the symmetric triangular distribution.

Under steady state Kalman filtering, it must be the case that $V_i^f \geq V_i^p$ and $V_i^f \geq \tilde{\rho}_i^2 V_i^p$. Hence $V_i^f \in [\tilde{b}, b]$, where:

$$\tilde{b} = \max(V_i^p, \tilde{\rho}_i^2 V_i^p) \tag{2.47}$$

Note that if $\tilde{b} > b$, then the observations must be dropped as they are inconsistent with steady state Kalman filtering. Using the equation for the Kalman gain, we then have:

$$\chi_i \in \left[1 - \frac{V_i^p}{\tilde{b}}, 1 - \frac{V_i^p}{b}\right] \tag{2.48}$$

The variance of the signal then lies in the interval:

$$\tilde{\sigma}_{q,i}^2 \in \left[\frac{b V_i^p}{b - V_i^p}, \frac{\tilde{b} V_i^p}{\tilde{b} - V_i^p}\right] \tag{2.49}$$

Note that if $\tilde{b} = V_i^p$, then the upper end of this interval is infinite, implying the signal may be infinitely noisy (i.e. contains no information).

Finally, the perceived variance of the shock lies in the range:

$$\tilde{\sigma}_{\varepsilon,i}^2 \in [\tilde{b} - \tilde{\rho}_i^2 V_i^p, b - \tilde{\rho}_i^2 V_i^p] \tag{2.50}$$

Case (iv): V_i^f and V_i^p both set-identified

In this case, the consumer is certain that current inflation lies within the specified

interval, and certain that future inflation will lie within one specific interval. Hence, we have $V_i^p \in [0, a]$ and $V_i^f \in [0, b]$. If $|\tilde{\rho}_i| \leq 1$, then χ_i is unrestricted within the interval $[0, 1]$. If $|\tilde{\rho}_i| > 1$, then χ_i is bounded below as described above. Hence $\chi_i \in [0, 1]$ if $|\tilde{\rho}_i| \leq 1$, and $\chi_i \in [1 - \tilde{\rho}_i^{-2}, 1]$ if $|\tilde{\rho}_i| > 1$.

We then know that:

$$\tilde{\sigma}_{q,i}^2 = \frac{V_i^f V_i^p}{V_i^f - V_i^p} \tag{2.51}$$

If $|\tilde{\rho}_i| < 1$, this can take any value. It could be infinite large if $V_i^p = V_i^f$, and could be zero if $V_i^p = 0$ but $V_i^f > 0$. If $|\tilde{\rho}_i| > 1$, then $V_i^f \geq \tilde{\rho}_i^2 V_i^p$. In that case, $\tilde{\sigma}_{q,i}^2$ could still be zero, but the maximum value it can now take is:

$$\tilde{\sigma}_{q,i}^2 = \frac{V_i^f V_i^p}{V_i^f - V_i^p} \leq \frac{V_i^f V_i^p}{\tilde{\rho}_i^2 V_i^p - V_i^p} \tag{2.52}$$

$$= \frac{V_i^f}{\tilde{\rho}_i^2 - 1} \leq \frac{b}{\tilde{\rho}_i^2 - 1} \tag{2.53}$$

To summarize, then, $\tilde{\sigma}_{q,i}^2 \in [0, \infty)$ if $|\tilde{\rho}_i| \leq 1$, and $\tilde{\sigma}_{q,i}^2 \in [0, \frac{b}{\tilde{\rho}_i^2 - 1}]$ if $|\tilde{\rho}_i| > 1$.

Turning to $\tilde{\sigma}_{\varepsilon,i}^2$, this could always be zero in this case. The maximum it could be is b if $V_i^f = b$ and $V_i^p = 0$. Hence $\tilde{\sigma}_{\varepsilon,i}^2 \in [0, b]$.

2.C.2 Alternative measurement for V_i^f

We consider two alternative measures for V_i^f. Like the measurement of V_i^p, both make use of just two pieces of information per respondent: their point estimate for inflation in the following year, and the probability that inflation will be within a particular range around that point estimate. We use these two pieces of information to fit a symmetric triangular distribution to beliefs about future inflation, and infer the variance from that.

For both of these measures, we take the point estimate for future inflation from the existing question in the survey (see Table 2.5). For the first (broader) measure, we then consider the density forecast question, and focus just on the inflation rate bin containing the respondent's point estimate. The probability assigned to this inflation range gives us the second piece of information. That is, we observe:

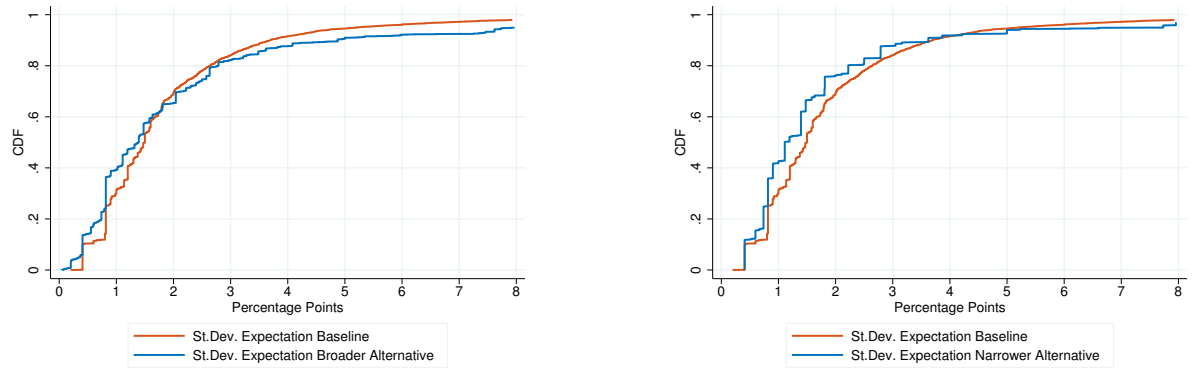
1. $\tilde{E}_{i,t}\pi_{t+1}$
2. $\Pr(lb < \pi_{t+1} \leq ub)$

for lb, ub defined by the edges of the relevant bin in the density forecast. We then fit the symmetric triangular distribution as described in Appendix 2.C.1. If the point estimate is on the boundary between two bins in the density forecast, we combine the bins to form

one wider inflation range, and take the sum of the probabilities given. This disregards some information contained within the future inflation density forecasts, and indeed requires dropping a small number of observations where the point estimate is completely inconsistent with the density forecast (i.e. the density forecast assigns 0 probability to the bin containing the point estimate). The sample size therefore shrinks somewhat, to c.93% of the original sample size. It is however much closer to the measurement of V_i^p: the only differences are that the respondent has been simultaneously asked about the probabilities of inflation being in several ranges rather than just one, and that the bin we use is not necessarily symmetric about their point estimate.

In the second (narrower) alternative measure for V_i^f, we go further and remove the second of these points of difference. That is, we restrict the sample to only respondents whose point estimate is at the mid-point of one of the bins in the density forecast question, then apply the same method described above. This substantially reduces the number of observations, but does leave us with a measure of V_i^f computed in the same way as V_i^p. The only assumption required to make them exactly comparable is an independence of irrelevant alternatives: the fact that respondents are also asked about the probability of inflation being in other ranges far away from their point estimate does not affect their answer for the range around their point estimate.

Figure 2.3 shows the distributions of V_i^f computed using our baseline measure using all of the information in density forecasts, and our two alternative triangular measures. They are all extremely similar. Moreover, the ranking of individuals within these distributions is strongly correlated. The Spearman's rank correlations of the first and second alternative measures with the baseline measure across individuals is 0.76 and 0.84 respectively.

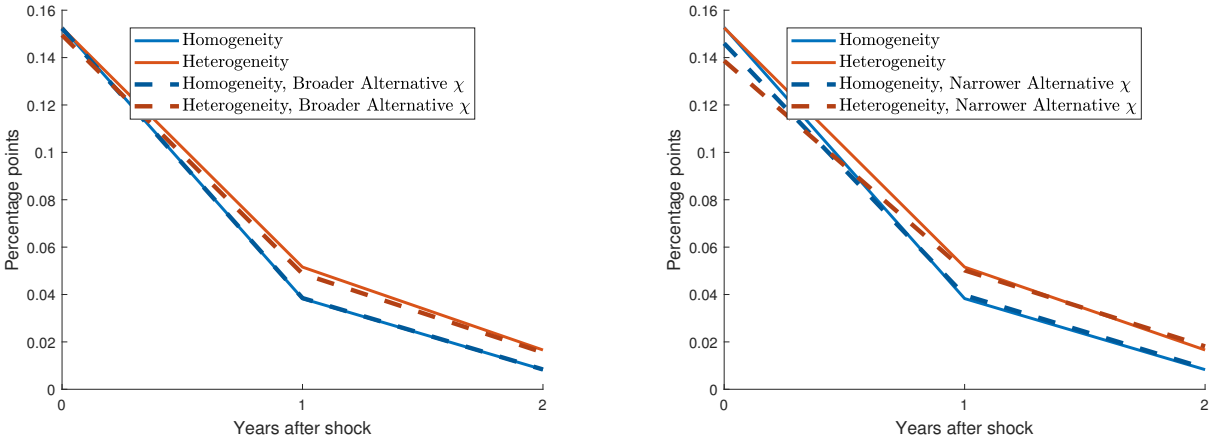


(a) Broader Alternative Measure **(b)** Narrower Alternative Measure

Figure 2.3: CDF of $St.Dev_i(\pi_{t+1})$ under the narrower and broader alternative calculation measures for V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

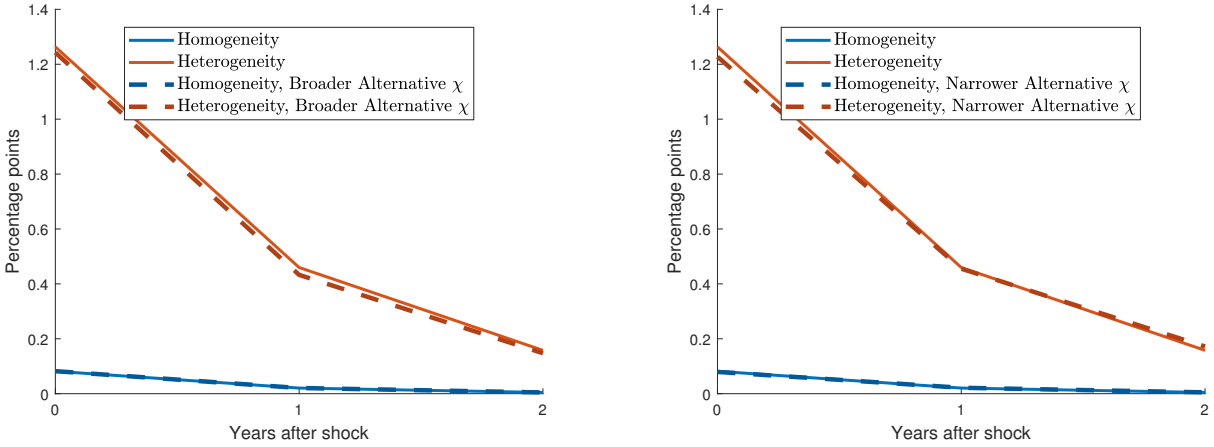
Unsurprisingly, the key results are therefore robust to these alternative variance measures. The mean Kalman gains in the two cases are 0.79 and 0.72, similar to to the 0.80

we find using our baseline measure. The different ways of calculating the Kalman gain correlate strongly across individuals, giving a Spearman's rank correlation of our baseline measure of χ_i with the first (broader) alternative of 0.69, and with the second (narrower) measure of 0.75. The impulse responses to an inflation shock in the model calibrated using the alternative measures are extremely close to those using the baseline measure (figures 2.4 and 2.5).



(a) Broader Alternative Measure **(b)** Narrower Alternative Measure

Figure 2.4: IRF of $\tilde{E}_t\pi_{t+1}$ under the narrower and broader alternative calculation measures for V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave.



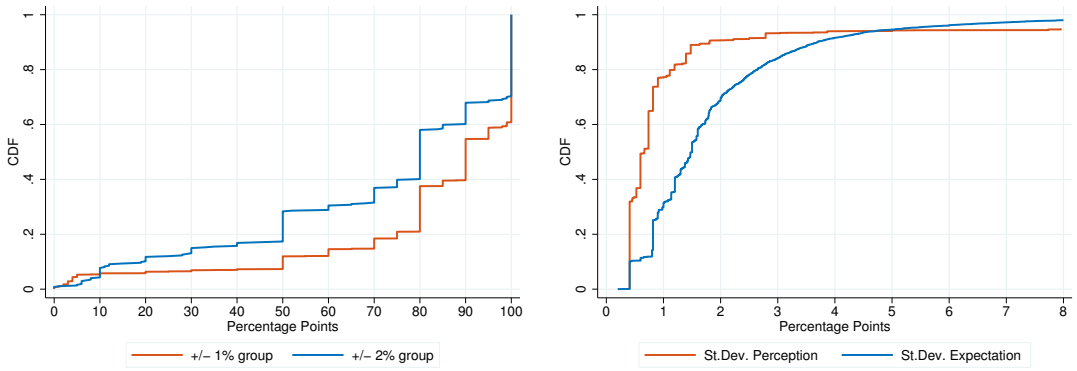
(a) Broader Alternative Measure **(b)** Narrower Alternative Measure

Figure 2.5: IRF of $\hat{c}_t$ under the narrower and broader alternative calculation measures for V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

2.D Additional empirical results

2.D.1 Additional parameter distributions

Figure 2.6a shows the CDF of the raw responses to question 1: respondents' assessment of the probability that current inflation lies within the specified range around their point estimate. We split the data between those with $\tilde{E}_{i,t}\pi_t \in (-5, 5)$, who were shown a $\pm 1\%$ interval, and those with $\tilde{E}_{i,t}\pi_t$ outside of this range, who were shown a $\pm 2\%$ interval. In both distributions, the majority believe there is at least an 80% chance that inflation lies within that range. Note that the $\pm 1\%$ group are more confident, despite seeing a smaller range, consistent with the notion that those who perceive lower rates of inflation or deflation are more certain in their perceptions. Figure 2.6b plots the CDFs of V_i^p and V_i^f. In cases where these are only set-identified, this plots the upper bound from fitting a symmetric triangular distribution. The lower bound in all such cases is 0. On average households are less uncertain about current inflation than about future inflation.



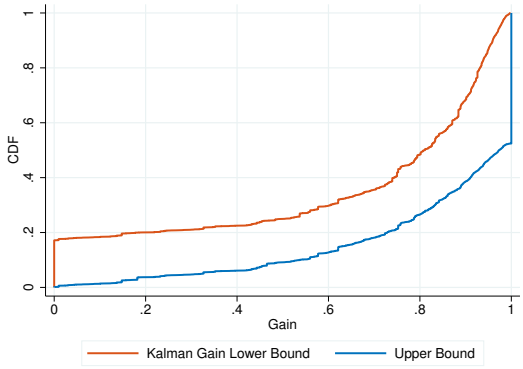
(a) CDFs of responses to question 1.

(b) CDFs of V_i^p and V_i^f.

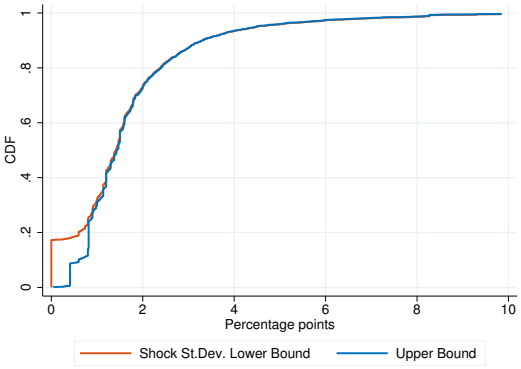
Figure 2.6: CDF of raw responses to question 1, for both the group shown a $\pm 1\%$ range and the group show a $\pm 2\%$ range, and the implied CDFs of V_i^p and V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

For respondents where we can only identify ranges for V_i^p and V_i^f, we can similarly only identify bounds for $\chi_i, \tilde{\sigma}^2_{q,i}, \tilde{\sigma}^2_{\varepsilon,i}$. Figure 2.7 shows the distributions of these parameters if we take the upper or lower bounds of the parameter ranges for those households respectively.

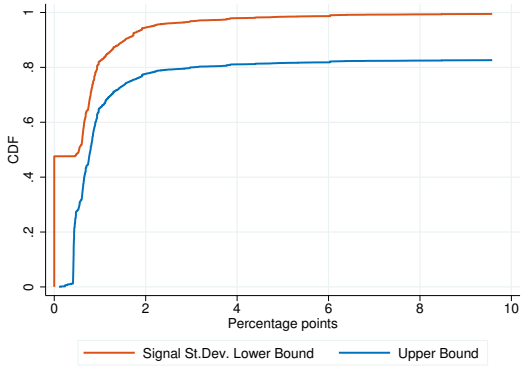
Figure 2.8 plots the distribution of $\tilde{\rho}_i$ when we exclude respondents whose response to the initial inflation expectations question ends in .0 or .5. Even excluding these households with the strongest tendency to round their answers, there is a large mass with $\tilde{\rho}_i = 0$, and substantial heterogeneity.



(a) χ_i



(b) $\tilde{\sigma}_{\varepsilon,i}$



(c) $\tilde{\sigma}_{q,i}$

Figure 2.7: CDFs of upper and lower bounds for inferred parameters. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

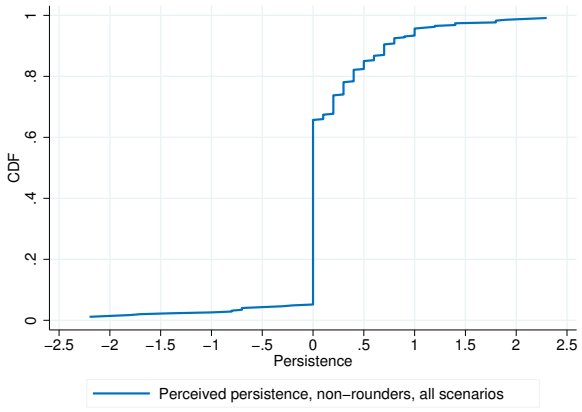


Figure 2.8: CDF of perceived persistence, only respondents whose response to initial inflation expectations question does not end in .0 or .5. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

2.D.2 Relationship between components of expectation formation and point estimates

Table 2.7 shows the means of the elements of the expectation laws of motion, broken down by the respondent's inflation perception ($\tilde{E}_{i,t}\pi_t$). Those with an inflation perception far

away from the actual value (which was approximately 5% at the time of the survey) tend to be the least certain in their perceptions. Those with the highest perceptions are also the least certain in their expectations and perceive the noise in the inflation process to be the highest. Those with perceptions that are either very high or very low also tend to have very low perceived persistence on average, and lower Kalman gains.

Table 2.7: Means of elements of expectation laws of motion, by inflation perception

	$SD_i(\pi_t)$	$SD_i(\pi_{t+1})$	$SD_i(\varepsilon_{t+1})$	χ_i	$\tilde{\rho}_i$
$\tilde{E}_{i,t}\pi_t < 0$	1.07	2.60	2.53	0.67	-0.13
	(0.27)	(0.52)	(0.48)	(0.13)	(0.30)
$\tilde{E}_{i,t}\pi_t \in [0, 2)$	0.77	1.85	1.80	0.70	0.21
	(0.04)	(0.16)	(0.17)	(0.04)	(0.08)
$\tilde{E}_{i,t}\pi_t \in [2, 4)$	0.67	1.84	1.78	0.78	0.24
	(0.01)	(0.03)	(0.03)	(0.01)	(0.02)
$\tilde{E}_{i,t}\pi_t \in [4, 6)$	0.71	2.13	2.07	0.81	0.22
	(0.01)	(0.04)	(0.04)	(0.01)	(0.02)
$\tilde{E}_{i,t}\pi_t \in [6, 8)$	1.44	3.35	3.22	0.74	0.22
	(0.10)	(0.19)	(0.19)	(0.03)	(0.09)
$\tilde{E}_{i,t}\pi_t \in [8, 10)$	1.44	3.95	3.89	0.81	0.17
	(0.17)	(0.44)	(0.41)	(0.07)	(0.17)
$\tilde{E}_{i,t}\pi_t \geq 10$	2.11	4.83	4.73	0.72	0.04
	(0.18)	(0.29)	(0.30)	(0.04)	(0.07)

Note: Bundesbank-Online-panel-Households, November 2021 wave. Standard errors in parentheses. For cases where χ_i is set-identified, respondents are excluded if the parameters are estimated very imprecisely (range$>$ 0.2). For all remaining set-identified parameters, the mid-point of the range is used. Observations of $\tilde{E}_{i,t}\pi_t$, $\tilde{E}_{i,t}\pi_{t+1}$, $SD_i(\pi_{t+1})$, and $SD_i(\pi_t)$ below the 1st or above the 99th percentile of that variable's distribution are also excluded as outliers, as are observations of $\tilde{\rho}_i$ outside $[-5, 5]$ (c.1% of observations).

2.D.3 Relationship between components of expectation formation: further details

Table 2.8 breaks down the different elements of the expectation law of motion according to $\tilde{\rho}_i$, divided into five categories; $\tilde{\rho}_i < 0$ $\tilde{\rho}_i = 0$, $\tilde{\rho}_i \in (0, 1)$, $\tilde{\rho}_i = 1$, and $\tilde{\rho}_i > 1$. This follows the classification for stock return beliefs in Dominitz and Manski (2011). Table 2.8 includes all respondents, with the most extreme 1% of responses for each of the standard deviation variables excluded as outliers.[21]

There is a highly non-linear relationship between the standard deviation of the perception and the persistence type. In particular, those who perceive that inflation is persistent but mean reverting are the most confident in their inflation perceptions. Those who believe that $\tilde{\rho}_i > 1$ have the lowest confidence in their perceptions. This fits with the notion

[21]Note that the first two columns include those for whom $V_i^p > V_i^f$, whose responses are inconsistent with steady-state Kalman filtering.

that those who track inflation most closely are also those who have the best knowledge of its dynamic properties. Those who believe inflation has zero persistence and those who believe it is explosive tend to also believe that the noise in the inflation process is highest.

Consistent with Result 4, those who believe that $\tilde{\rho}_i > 1$ are more likely to be hand-to-mouth than any of the other persistence types, over twice as likely if one only includes those whose responses are consistent with Kalman filtering. They are also much less likely to own securities.

Table 2.8: Means of elements of expectation laws of motion, by persistence type

	$SD_i(\pi_t)$	$SD_i(\pi_{t+1})$	$SD_i(\varepsilon_{t+1})$	χ_i	HTM	Owns Stocks
$\tilde{\rho}_i < 0$	1.58	1.79	1.65	0.85	0.09	0.65
	(0.39)	(0.11)	(0.14)	(0.02)	(0.03)	(0.04)
$\tilde{\rho}_i = 0$	1.80	1.72	1.96	0.79	0.14	0.56
	(0.09)	(0.03)	(0.04)	(0.01)	(0.01)	(0.01)
$\tilde{\rho}_i \in (0,1)$	0.90	1.59	1.68	0.80	0.09	0.67
	(0.15)	(0.07)	(0.08)	(0.02)	(0.02)	(0.03)
$\tilde{\rho}_i = 1$	1.59	1.56	1.51	0.79	0.10	0.59
	(0.20)	(0.05)	(0.06)	(0.01)	(0.01)	(0.02)
$\tilde{\rho}_i > 1$	2.33	2.01	2.21	0.92	0.19	0.47
	(0.35)	(0.09)	(0.20)	(0.01)	(0.02)	(0.03)

Note: Bundesbank-Online-panel-Households, November 2021 wave. Standard errors in parentheses. For cases where χ_i is set-identified, respondents are excluded if the parameters are estimated very imprecisely (range$>$ 0.2). For all remaining set-identified parameters, the mid-point of the range is used. Observations of $\tilde{E}_{i,t}\pi_t, \tilde{E}_{i,t}\pi_{t+1}, SD_i(\pi_{t+1})$, and $SD_i(\pi_t)$ below the 1st or above the 99th percentile of that variable's distribution are also excluded as outliers, as are observations of $\tilde{\rho}_i$ outside $[-5,5]$ (c.1% of observations). Note that the first two columns include those for whom $V_i^p > V_i^f$, whose responses are inconsistent with steady-state Kalman filtering.

Finally, note that those who believe that $\tilde{\rho}_i > 1$ have the highest χ_i on average. However, this is partly mechanical, since if $|\tilde{\rho}_i| > 1$ then that places a lower bound on the values of χ_i that are consistent with steady-state Kalman filtering. Between the groups with $\tilde{\rho}_i \in [0,1]$, the average Kalman filter varies little.

Table 2.9 shows regressions of each component of the expectation laws of motion on $\tilde{\rho}_i$, split in two ways. The first panel splits respondents according to which hypothetical scenario they were shown before Question 2, to explore the role of different shock types. That is, each dependent variable is regressed on $\tilde{\rho}_i$ interacted with a categorical variable reflecting which shock scenario the respondent saw.

The second panel splits households into some of the persistence categories outlined above, specifically those who believe the price level is mean-reverting ($\tilde{\rho}_i < 0$), those who believe inflation is persistent but stationary ($\tilde{\rho}_i \in (0,1)$), and those who believe inflation is non-stationary with positive persistence ($\tilde{\rho}_i \geq 1$). The final panel shows the results of regressing each dependent variable on an indicator equal to 1 if the household does no updating of expectations at all when faced with the hypothetical shock ($\tilde{\rho}_i = 0$).

In the first panel, there are some significant differences between shock types in the relationships of $\tilde{\rho}_i$ with other elements of expectation laws of motion. However, the magnitudes are generally small. For that reason we pool households across shock types for the analysis in Section 2.4.2.

The differences are much larger, however, across persistence types. Panel 2 shows that within households who believe inflation is persistent and stationary, greater perceived persistence is associated with less uncertainty about current and future inflation, less perceived noise in the inflation process, and a greater implied Kalman gain. This is consistent with models of endogenous information acquisition, as with a more persistent inflation process information about the current rate of inflation is more valuable.

Table 2.9: Breakdown of $\tilde{\rho}_i$ relationships with other expectation law of motion components by shock type and persistence category.

	(1) $SD_i(\pi_{t+1})$	(2) $SD_i(\pi_t)$	(3) $SD_i(\varepsilon_{t+1})$	(4) χ_i
Panel A: Shock type				
Shock unspecified $\times$ $\tilde{\rho}_i$	-0.00278	0.0103	-0.162*	-0.00969
	(0.105)	(0.0306)	(0.0913)	(0.0149)
Supply $\times$ $\tilde{\rho}_i$	0.0956	0.00437	-0.118	0.0104
	(0.105)	(0.0246)	(0.102)	(0.00983)
Demand $\times$ $\tilde{\rho}_i$	0.145	-0.0546***	0.0278	0.0418***
	(0.118)	(0.0165)	(0.116)	(0.00872)
Constant	2.077***	0.751***	2.055***	0.789***
	(0.0308)	(0.0115)	(0.0307)	(0.00526)
Panel B: Persistence type				
$\tilde{\rho}_i < 0$ $\times$ $\tilde{\rho}_i$	-0.150	0.0268	0.132	-0.0390***
	(0.117)	(0.0306)	(0.123)	(0.0142)
$\tilde{\rho}_i \in (0,1)$ $\times$ $\tilde{\rho}_i$	-0.459***	-0.321***	-0.532***	0.0497*
	(0.161)	(0.0351)	(0.162)	(0.0301)
$\tilde{\rho}_i \geq 1$ $\times$ $\tilde{\rho}_i$	0.164**	-0.00462	-0.0968	0.0246***
	(0.0795)	(0.0178)	(0.0788)	(0.00727)
Constant	2.082***	0.765***	2.087***	0.784***
	(0.0340)	(0.0130)	(0.0340)	(0.00574)
Panel C: Updating indicator				
$\tilde{\rho}_i \neq 0$	-0.106*	-0.0744***	-0.306***	0.0132
	(0.0594)	(0.0196)	(0.0585)	(0.0102)
Constant	2.129***	0.770***	2.128***	0.789***
	(0.0343)	(0.0136)	(0.0343)	(0.00596)
Observations	2317	2317	2317	2317

Note: Bundesbank-Online-panel-Households, November 2021 wave.
Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Although there are only weak relationships between $\tilde{\rho}_i$ and uncertainty over current and future inflation across the whole sample, the second panel reveals that this is driven by weak relationships among those who believe in inflation processes that are qualitatively different from the data. Among those who believe that inflation is persistent but stationary, the relationships between $\tilde{\rho}_i$ and uncertainty are very strong. Since an AR(1) process estimated on German CPI inflation over the previous 20 years implies a persistence of $\rho = 0.21$, this suggests that the group of households most aware of the time-series properties of inflation behave as predicted by models of rational inattention (e.g. Sims, 2003).

However outside of this group, households behave less in line with those predictions.

2.D.4 Correlations of expectation components with household characteristics: further details

Table 2.10: Probit regressions of $\tilde{\rho}_i \neq 0$ on household characteristics, split by shock type.

	(1) $\tilde{\rho}_i \neq 0$	(2) $\tilde{\rho}_i \neq 0$	(3) $\tilde{\rho}_i \neq 0$	(4) $\tilde{\rho}_i \neq 0$	(5) $\tilde{\rho}_i \neq 0$
Hand-to-mouth	-0.0856	-0.1078	-0.1257	-0.3049**	0.1170
	(0.1079)	(0.0761)	(0.1346)	(0.1353)	(0.1349)
Liquid wealth	0.0001	0.0000	0.0001	0.0001	-0.0000
	(0.0002)	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Illiquid wealth	-0.0000	0.0000	0.0001	0.0000	-0.0000
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Other wealth	0.0007	0.0002	0.0002	-0.0002	0.0010
	(0.0006)	(0.0005)	(0.0008)	(0.0009)	(0.0008)
Debt	0.0004	0.0000	0.0001	-0.0007	0.0006
	(0.0003)	(0.0002)	(0.0004)	(0.0005)	(0.0004)
log(income)	-0.0624	-0.0075	-0.0890	0.0669	-0.0061
	(0.0773)	(0.0574)	(0.1036)	(0.1057)	(0.0992)
HH Controls	Yes	Yes	Yes	Yes	Yes
Shock type	All	All	Unspecified	Supply	Demand
Observations	1899	3194	1053	1051	1068
Pseudo-R^2	0.0285	0.0177	0.0329	0.0405	0.0397

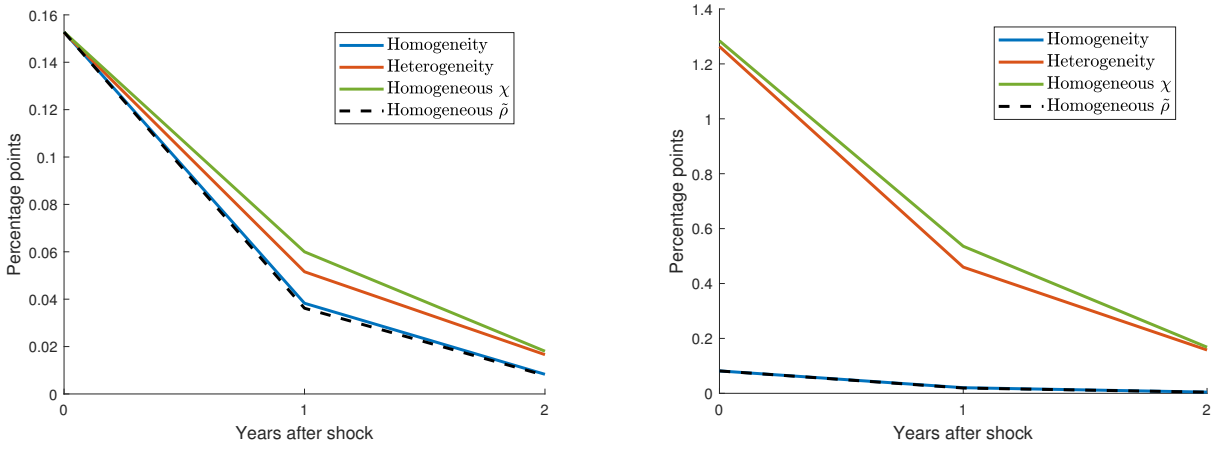
Note: Bundesbank-Online-panel-Households, November 2021 wave. The units of the wealth and debt variables are €1000s. The household controls are age (in years up to a top bin of ≥ 80, coded as 80), age^2, gender, region (north/south/east/west), education, occupation category, and employment status (all categorical, for details see the full questionnaire at https://www.bundesbank.de/en/bundesbank/research/survey-on-consumer-expectations/questionnaires-850746). All controls except age and age^2 are treated as categorical. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.10 column 1 shows the results of a probit regression of of an indicator variable w_i on household characteristics, where $w_i = 1$ if $\tilde{\rho}_i \neq 0$, and $= 0$ otherwise. These results therefore give the estimated relationship between household characteristics and the probability of adjusting expectations in light of hypothetical shocks. None of the characteristics are significantly related to this selection, and the magnitudes are small: the average marginal effect of being hand-to-mouth on the probability of $\tilde{\rho}_i \neq 0$ is less than 3 percentage points.

This result is robust to extending the sample to include those for whom χ_i or $\tilde{\sigma}^2_{\varepsilon i}$ cannot be inferred precisely: as those variables are not used in the calculation of $\tilde{\rho}_i$, the remaining columns of Table 2.10 repeat the exercise for the full sample, and then split by the shock scenario seen by the household. The associations between $\tilde{\rho}_i \neq 0$ and wealth/income is not significantly different from zero throughout, except for the supply shock, for which hand-to-mouth households are somewhat less likely than others to adjust their expectations.

2.E Further impulse response exercises

In Figure 2.9b, we break down the amplification from heterogeneous expectation formation into its components. The impulse response with heterogeneity in $\tilde{\rho}_i$ only, but homogeneous χ_i, is close to that with full heterogeneity. This is therefore the main driver of the amplification we find. Note however that the difference between the IRFs with full heterogeneity and with homogeneous χ_i is small relative to the response with full heterogeneity, it remains large relative to the consumption responses with homogeneous expectation formation, and with homogeneous $\tilde{\rho}_i$. The covariance between $\tilde{\rho}_i$ and χ_i, though small in the data, does still play a non-trivial role in aggregate consumption dynamics.

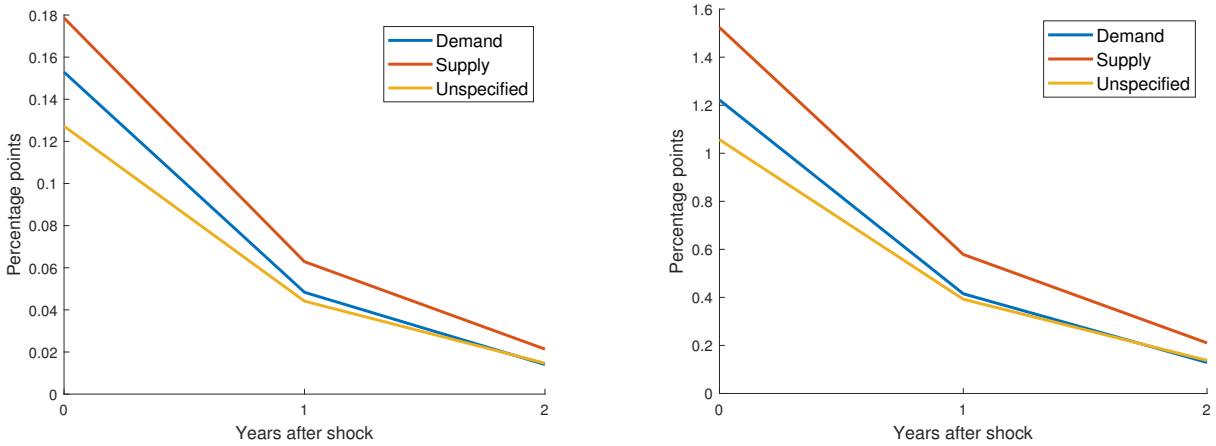


(a) Expectations **(b)** Consumption

Figure 2.9: Implied IRF of aggregate expectations and aggregate consumption. The homogeneity and heterogeneity cases are as described in Section 2.5.2. The remaining cases set χ_i and $\tilde{\rho}_i$ respectively to their average values for all households. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

Using the survey responses to the different hypothetical scenarios in Question 2, we can further compare the effects of heterogeneous expectation laws of motion for different types of shock. We find somewhat greater amplification and persistence in consumption

responses to supply shocks than other types of shock. A comparison of the IRFs between the three cases is shown in Figures 2.10a and 2.10b. This result is consistent with the higher average perceived persistence of supply shocks discussed in Section 2.4.



(a) Expectations **(b)** Consumption

Figure 2.10: Implied IRFs of one-period ahead inflation expectations and consumption by shock. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

As in Section 2.4, we also repeat these exercises with the distributions of subjective models after excluding those whose answers are rounded to a multiple of 0.5. The model with heterogeneity does deliver smaller initial consumption responses in this case, but it is still 4.4× larger than under the homogeneity case calibrated using the average $\tilde{\rho}_i$ and χ_i across all respondents. This rises to 5.5× if one compares to a representative agent model calibrated using the average $\tilde{\rho}_i$ and χ_i for the population of non-rounders (as they have slightly smaller perceived persistence on average). As such, the result that heterogeneity generates very significant amplification of the transmission of inflation shocks to consumption still holds.

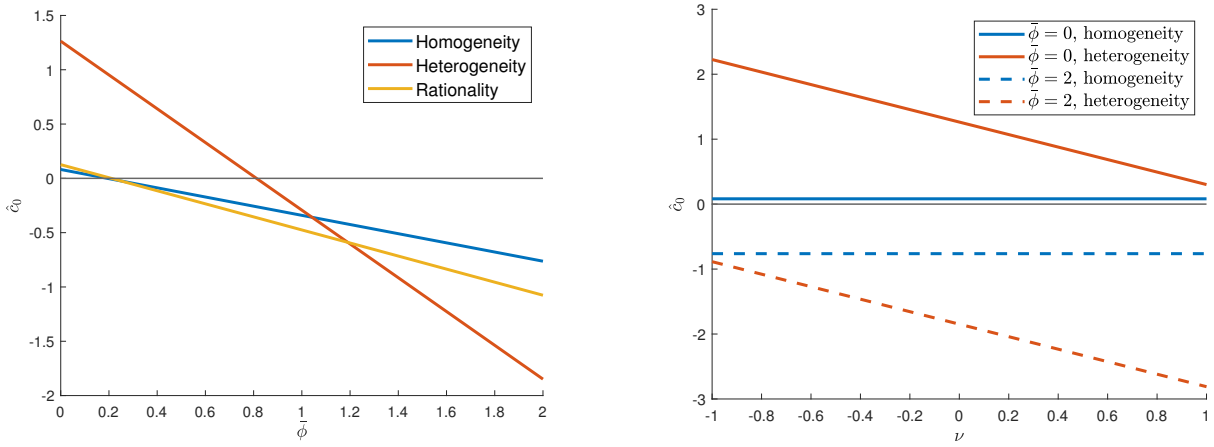
Finally, we allow for households to expect nominal interest rates to respond to inflation. Specifically, we assume that:

$$\tilde{E}_{i,t} i_{t+h} = \tilde{\phi}_i \tilde{E}_{i,t} \pi_{t+h} \tag{2.54}$$

Households are assumed to not observe i_t precisely when they choose $c_{i,t}$, consistent with them not observing π_t. They therefore infer i_t from their perceived current inflation, just as they do for expectations of future periods.

Figure 2.11 shows the initial consumption response $\hat{c}_0$ under homogeneity, heterogeneity, and rationality for a range of assumptions about $\tilde{\phi}_i$. In panel (a), we assume that all households share a common $\tilde{\phi}_i = \bar{\phi}$. When the Taylor principle is expected to be satisfied, higher inflation leads households to expect higher real interest rates, and so to reduce consumption. Heterogeneity amplifies this fall in consumption, just as it

amplifies the consumption increase in the baseline case ($\bar{\phi} = 0$) studied in Section 2.5.2. Indeed, heterogeneity provides substantial amplification of aggregate consumption for all perceived interest rate rules aside from a small region around $\bar{\phi} = 1$, in which inflation shocks are only expected to have very small effects on real interest rates. Heterogeneity also increases the sensitivity of aggregate consumption to the perceived interest rate rule. In panel (b), we show that our results also remain robust even if $\tilde{\phi}_i$ is allowed to covary with $\tilde{\rho}_i$. Such covariances do make some quantitative difference to the results, so future research could consider ways to measure these relationships.



(a) Varying $\bar{\phi}$ **(b)** Varying ν

Figure 2.11: Implied $\hat{c}_0$ under different assumptions on $\tilde{\phi}_i$. In panel (a) all households form interest rate expectations using the same perceived Taylor Rule parameter $\bar{\phi}$. In panel (b) this parameter varies across households according to $\tilde{\phi}_i = \bar{\phi} + \nu(\tilde{\rho}_i - E[\tilde{\rho}_i])$. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

Chapter 3

A Four Equation New Keynesian Model with Many Sectors and Non-Homothetic Preferences

Abstract I study a new Keynesian model with non-homothetic preferences over many final goods. For a broad class of preferences, I show that the textbook IS and Phillips curve hold, but for a differently weighted inflation index: the inflation rate of the marginal consumption basket. If one rewrites the model in terms of consumption-weighted inflation, it features a new variable: the relative price of the marginal consumption basket. I derive a law of motion for this relative price to give an analytically tractable four-equation model. The first-order role of relative prices in this model (i) qualitatively changes optimal policy, (ii) biases the Phillips curve slope obtained when estimating the textbook model, and (iii) in a TANK extension, generates fluctuations in inequality.

3.1 Motivation

In the real world, households allocate consumption across a range of final goods and services. Moreover, the composition of households' consumption spending across these goods and services changes as total spending rises.[1] By way of example, figure 3.1 shows an example of how the budget shares of UK households vary by income decile for selected consumption categories.[2] Food and energy account for a larger share of total spending for lower income households, while higher income households spend proportionally more on transport, recreation, and hospitality. The differences are large in magnitude. Moreover, given that the data shown in figure 3.1 categorizes households according to *equivalized* disposable income, this relationship is evident even after adjusting for household composition.

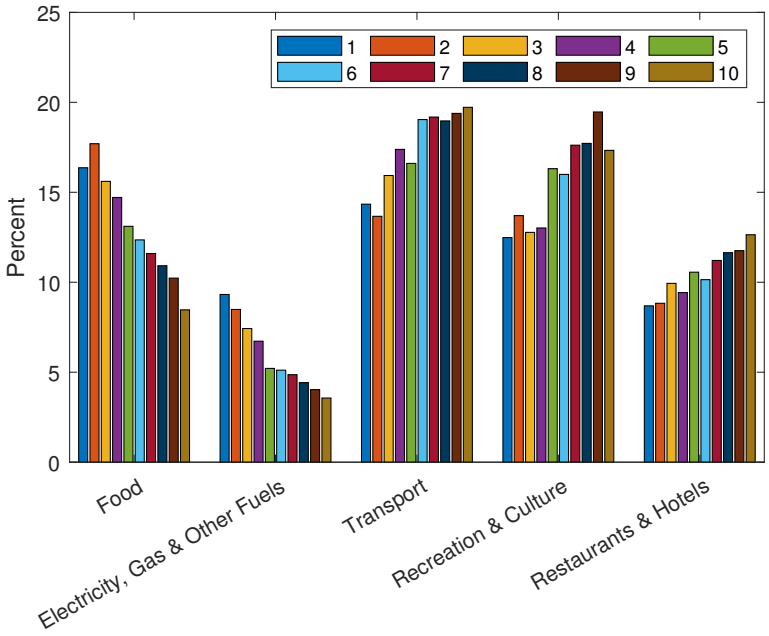


Figure 3.1: Budget Shares for Selected Consumption Categories by Equivalized Disposable Income Decile, Based on 2019 UK Data

The cross-sectional correlation between total spending and the composition of spending shown in figure 3.1 would be consistent with households' preferences being non-homothetic. The existence of non-homothetic preferences is confirmed by a number of empirical studies in the literature on growth and structural transformation. Several studies document that document that the consumption share of services has risen while

[1] This observation dates back as far as Engel (1857).

[2] As the data exclude owner-occupier housing costs such as mortgage interest payments, I have also excluded net rents and repair and maintenance when calculating the budget shares to ensure comparability across income groups. Note that the budget shares do not sum to 100 as only a subset of consumption categories is shown.

the consumption share of services has declined over time in the US as per capita incomes have risen (see for example Kongsamut et al. (2001)). Reallocation across sectors could be driven either by technological preferences causing changes in relative prices and hence substitution effects or by non-homothetic preferences resulting in income effects. However, recent studies indicate that non-homothetic preferences are a quantitatively important driver; Boppart (2014) estimates a model that allows for both income and substitution effects and finds that income effects account for roughly half of structural change in the US, while Comin et al. (2021) show using panel data that income effects account for the majority of sectoral reallocation within countries.[3]

While the implications of non-homothetic preferences for growth and structural transformation have been extensively explored, their effects on business cycle analysis have been less well studied.[4] For example, the textbook new Keynesian model features just one final good, and many multi-sector models feature homothetic preferences.[5] This is all the more surprising given that the evidence indicates that differences in income elasticities across goods do affect consumption patterns at the business cycle frequency; Bils and Klenow (1998) find that spending on luxuries and durables in the US is more cyclical than spending on necessities.

In this paper, I construct a new Keynesian model with many sectors and non-homothetic preferences to study the implications of these features of consumption data for business cycle fluctuations. For a broad class of non-homothetic utility functions, I show that the standard IS and Phillips curve can be obtained, but with a *differently weighted inflation index*. The relevant inflation rate is the inflation rate of the *marginal* consumption basket. The marginal consumption basket describes the additional goods that one would purchase if one were given an additional sum to spend on consumption. Under non-homothetic preferences, the marginal consumption basket differs from the *average* consumption basket. This difference between marginal and average consumption baskets is why the composition of the consumption basket changes with total spending.

I refer to the inflation rate on this marginal consumption basket as π_t^M. The IS and Phillips curves in my model are:

$$\pi_t^M = \beta E_t \pi_{t+1}^M + \kappa y_t + \eta_t^M \tag{3.1}$$

$$y_t = E_t y_{t+1} - \sigma(i_t - E_t \pi_{t+1}^M) \tag{3.2}$$

Here, η_t^M is an AR(1) cost-push shock, with persistence $\rho \in [0, 1)$. This is a weighted

[3]Differences in income elasticities across consumption categories have also been shown to be empirically relevant in the international trade literature. See for example Caron et al. (2014).

[4]Some exceptions are reviewed below.

[5]For example, see Aoki (2001).

average of the cost-push shocks hitting each sector, using the sector weightings in the marginal consumption basket.

Why does this result emerge? The decisions that consumers take in this model are marginal substitution decisions. Consider the intratemporal decision between consumption and labour supply. When the consumer decides whether to marginally increase labour supply in order to increase income and consumption, the relevant marginal rate of transformation is the nominal wage relative to the price of the goods that they will consume at the margin. Consider then the intertemporal decision between consumption today and consumption tomorrow. When the consumer decides whether to marginally reduce consumption today and marginally increase consumption tomorrow, the marginal rate of transformation is governed by the nominal interest rate and expected change in the price of the goods that they consume at the margin.

Although (3.1) and (3.2) feature the inflation rate of the marginal consumption basket, we typically think about central banks as caring about the inflation rate of the average consumption basket. In my simple model, this corresponds to either consumer price inflation or the log change in the GDP deflator.[6] Central bank inflation targets correspond to this quantity. Moreover, concern for the inflation rate of the average consumption basket is justified in my model, given that this turns out to be the welfare-relevant inflation rate. If one rewrites the model in terms of this average inflation rate, one obtains:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa y_t + \lambda \tilde{p}_t^M + \eta_t \tag{3.3}$$

$$y_t = E_t y_{t+1} + \alpha(E_t \tilde{p}_{t+1}^M - \tilde{p}_t^M) - \sigma(i_t - E_t \pi_{t+1}) \tag{3.4}$$

$$\tilde{p}_t^M = \zeta \tilde{p}_{t-1}^M + \frac{\zeta}{1 - \beta\zeta\rho} \tilde{\eta}_t^M \tag{3.5}$$

This looks like the textbook model, but with one additional term, $\tilde{p}_t^M$. This variable is the relative price of the marginal compared to the average consumption basket. The evolution of this relative price is described by (3.5). This law of motion is an AR(1) process, with an AR(1) innovation $\tilde{\eta}_t^M$. This innovation is the difference between η_t^M, which weights the sectoral cost-push shocks according to the weights of the marginal basket, and η_t, which instead uses the weights in the average consumption basket.

This formulation shows that relative prices have a first-order role in inflation and output dynamics. Note the contrast with the textbook model, in which price dispersion only matters to a second-order approximation. The first-order role for relative prices is potentially important, given that we know from the literature that relative price shocks

[6]Note that consumption equals output, so that consumer price inflation and the log change in the GDP deflator are equal.

are large and account for a substantial fraction of the variability of inflation.[7] But how precisely do these relative price shocks affect macroeconomic dynamics? In the baseline representative agent (RANK) model, I study two implications.

First, I solve the model analytically when policy is set using a Taylor rule, and show that a rise in the relative price of the marginal consumption basket leads to a large decline in output, but a much smaller decline in inflation. This disparity arises because $\tilde{p}_t^M$ acts as a negative demand shock, but a positive supply shock, as shown in (3.3) and (3.4). Relative price consequently engender considerable output volatility but far less inflation volatility, which gives the appearance of a flatter Phillips curve. I simulate data from my model, and estimate the textbook model on that data, as would an econometrician who neglects the first-order role of relative price shocks. If relative price shocks are large and persistent, such an econometrician would vastly underestimate the Phillips curve slope.

Second, I solve analytically for optimal policy under discretion and commitment, and show that the nominal rate should optimally respond to relative prices. Moreover, I show that the macroeconomic co-movements obtained under optimal discretion differ qualitatively from those seen in the textbook new Keynesian model. In particular, when $\tilde{p}_t^M$ is high, so that relative prices are far from steady-state, it is optimal for the policymaker to let inflation rise, while lowering the real interest rate. Again, this result obtains because relative prices act as a large negative demand shock, and a smaller positive supply shock. In contrast, in the textbook model with uncorrelated demand and supply shocks, the covariance of the real rate and inflation should always be positive.

I then derive a two agent extension of the model, to study how relative price shifts affect business cycle fluctuations in inequality. Previous studies have analyzed the effects of relative prices on inequality in the case where unconstrained and hand-to-mouth agents have different steady-state consumption baskets. By contrast, I show that relative prices strongly affect inequality even when all agents have the same consumption basket. In particular, I show that hand-to-mouth consumption c_t^H is given by:

$$c_t^H = (1+\phi)y_t + \frac{\phi}{\gamma+\phi}\tilde{p}_t^M \tag{3.6}$$

The first term is standard in two agent new Keynesian (TANK) models: absent redistribution on the part of the fiscal policymaker, hand-to-mouth consumption moves more than one-for-one with output. The second term is novel. When the relative price of the marginal basket is high, it pushes up the nominal wage relative to the average consumption basket, thereby giving hand-to-mouth agents a positive income effect.

[7] See, for example, Reis and Watson (2010) and Ruge-Mucia and Wolman (2022). Also relevant is Borio et al. (2021), who show that the common component of price changes has accounted for a smaller share of total price changes over time, suggesting a more important role for relative price shocks.

I then sketch two main implications of relative prices for inequality in the TANK model. First, I demonstrate that there is a tension between the partial equilibrium and general equilibrium effects of relative prices on inequality. On the one hand, (3.6) shows that for a given level of demand, a rise in $\tilde{p}_t^M$ raises hand-to-mouth consumption. However, once one takes into account the general equilibrium effects of $\tilde{p}_t^M$ on y_t through the demand curve, an increase in the relative price of the marginal consumption basket sharply lowers hand-to-mouth consumption. This finding highlights the importance of taking into account general equilibrium effects on demand when assessing the effects of relative prices on inequality.

Second, welfare in the TANK model is affect by the interaction between output and relative prices; both terms affect hand-to-mouth consumption, and so their covariance matters for fluctuations in inequality. In (3.6), a somewhat negative covariance between y_t and $\tilde{p}_t^M$ can mitigate fluctuations in inequality by attenuating the pro-cyclicality of hand-to-mouth consumption. I show that this effect qualitatively changes how inequality considerations affect optimal policy. Consider a pure cost-push shock, with no relative price changes. In a standard TANK model, one would optimally have a smaller response of output and a larger response of inflation in TANK than in RANK; fluctuations in output are more costly in TANK because they are associated with rising inequality. In response to a relative price shock, however, the optimal output response may be larger in TANK than in RANK, and the inflation response smaller, because the central bank can mitigate fluctuations in inequality by generating a negative covariance between output and the relative price shock.

3.1.1 Related Literature

Other models incorporating non-homothetic preferences into business cycle models include Ravn et al. (2008), Cavallari (2018), Anand et al. (2015), and Blanco and Diz (2021), each of which focus on the special case with Stone-Geary preferences. Anand et al. (2015) considers a two sector model, in which prices are flexible in one sector, which differs from my model in which there are an arbitrary number of sectors and sticky prices. The recent study by Blanco and Diz (2021) is the most closely related paper, also showing that in the special case with Stone-Geary preferences and two sectors a price index weighted by the weights in the marginal consumption basket matter for intratemporal and intertemporal substitution decisions and analysing optimal policy. Compared to this paper, the key result in my paper is to show that this result holds for a surprisingly broad class of nonhomothetic utility functions, and in the case with many sectors. I also study a two-agent model to incorporate the effects of inequality, and derive a more compact representation of the model with just one additional variable, which allows me

to solve analytically for impulse responses and optimal policy prescriptions, leading to novel results.

Other related contributions include Lee (2021), who considers a two agent model in which the two types of agents have different consumption baskets. I also analyze the role of relative prices shocks for business cycle fluctuations in inequality and welfare, but in my model the two agents have the same consumption baskets in steady-state. Relative price shocks affect inequality through a different channel; through changes in the relative price of the marginal consumption basket. Another recent related contribution in an open economy setting is Zhou (2022).

This paper considers the effects of transitory shocks to relative prices (albeit these shocks may be highly persistent). The key new variable in the model will be the difference of the relative price of the marginal basket compared to its steady-state. As such, it can be used to analyze optimal policy when relative prices fluctuate away from their long-run levels. This contrasts with papers such as Wolman (2011) that investigate the optimal rate of inflation in the context of trending relative prices.

My paper is also related to an older literature studying optimal monetary policy in multi-sector new Keynesian models. See for example Aoki (2001) and Mankiw and Reis (2003). The key innovation versus these papers is to show the implications of non-homothetic preferences and their influence on demand. Many papers in this literature emphasize heterogeneity on the production side of the economy, whereas my paper emphasizes cross-sector differences in demand patterns stemming from non-homothetic preferences. That said, the idea that the standard Phillips curve emerges but for a different inflation basket is related to Rubbo (2020), a study which shows that a related result emerges in a multi-sector model with production networks.

The two-agent extension of this model relates to large literature considering the implications of heterogeneity in asset market participation for the transmission of shocks and for optimal policy in new Keynesian models. In particular, I use an analytical two agent framework following Galí et al. (2007) and Bilbiie (2008). This literature is by now very large, and many papers such as Kaplan et al. (2018) consider full heterogeneous agent models. My contribution to this literature is to illustrate a new channel through which relative price shocks can affect business cycle fluctuations in inequality in this class of models; the relative price inflation of the marginal consumption basket.

3.1.2 Outline

Section 3.2 sets out the baseline representative agent model. Section 3.3 shows how to derive the IS and Phillips curve given in (3.1) and (3.2). Section 3.4 shows how to derive the four equation model in terms of consumer price inflation. I then explore implications

of the RANK model in section 3.5. Section 3.6 shows how to extend the model to a two-agent setting, where a fraction of agents have unconstrained access to capital markets and others are hand-to-mouth. Section 3.7 sketches the implications of relative price shocks for inequality in the TANK model. Section 3.8 concludes and suggests a number of extensions that could be considered in future work.

3.2 Baseline Model

This section outlines a representative agent model with non-homothetic preferences and relative price shocks. I start by assuming a specific form of non-homothetic preferences, and then show that the same model emerges under a much broader class of non-homothetic utility functions.

3.2.1 Households

The baseline model assumes that households have an "additive preferences" utility function à la Houthakker (1960). Households buy final goods of a number of types indexed $i \in I$. They also supply labour. Their utility function is:

$$U_t = E_t \sum_{h \geq 0} \beta^h \left(\sum_{i \in I} \frac{C_{i,t}^{1-\gamma_i}}{1-\gamma_i} - \frac{N_t^{1+\phi}}{1+\phi} \right)$$

As such, the exponent on each type of good differs. An important feature of this utility function is that when total expenditure is very low, spending will be concentrated on goods with a high γ_i. As spending increases, spending shifts towards goods with a lower γ_i. Indeed, as Houthakker (1960) shows, the income elasticity of demand for each good is governed by γ_i^{-1}. I will explore this point in more detail below. A related form of utility function is used by Straub (2019). In that model, the consumption choices at different ages have different exponents γ_i.

In the baseline model, all households have unconstrained access to asset markets. Their budget constraint is consequently:

$$\sum_{i \in I} P_{i,t} C_{i,t} + B_t = (1 + i_{t-1}) B_{t-1} + W_t N_t + D_t$$

The first-order conditions are:

$$\begin{aligned}
\forall i,j \in I : C_{i,t}^{-\gamma_i}\frac{1}{P_{i,t}} &= C_{j,t}^{-\gamma_j}\frac{1}{P_{j,t}} \\
\forall i \in I : C_{i,t}^{-\gamma_i}\frac{1}{P_{i,t}} &= N_t^{\phi}\frac{1}{W_t} \\
\forall i \in I : C_{i,t}^{-\gamma_i}\frac{1}{P_{i,t}} &= \beta(1+i_t)E_t C_{i,t+1}^{-\gamma_i}\frac{1}{P_{i,t+1}}
\end{aligned}$$

Define total consumption C_t as:

$$C_t = \sum_{i \in I} C_{i,t}$$

Note that in steady-state:

$$\bar{C} = \sum_{i \in I} \bar{C}_i$$

I will show below that in steady-state all relative prices are equal to unity, and the nominal wage equals the price in each sector. As such, using the intratemporal conditions:

$$\bar{C} = \sum_{i \in I} \bar{C}_j^{\frac{\gamma_j}{\gamma_i}}$$

This equation can be solved numerically to give the weight of good j in the total consumption basket for any steady-state level of consumption. A result that will prove more helpful, though, will be to derive the weight of good j in the *marginal* consumption basket. That is:

$$\psi_j^M = \frac{d\bar{C}_j}{d\bar{C}}$$

Using the implicit function theorem, one has:

$$\begin{aligned}
F &= \sum_{i \in I} \bar{C}_j^{\frac{\gamma_j}{\gamma_i}} - \bar{C} \\
\frac{d\bar{C}_j}{d\bar{C}} &= -\frac{F_{\bar{C}}}{F_{\bar{C}_i}} \\
&= \frac{1}{\sum_{i \in I} \frac{\gamma_j}{\gamma_i} \frac{\bar{C}_j^{\gamma_j} \gamma_i}{\bar{C}_j}} \\
&= \frac{1}{\sum_{i \in I} \frac{\gamma_j}{\gamma_i} \frac{\bar{C}_i}{\bar{C}_j}} \\
&= \frac{\bar{C}_j}{\bar{C}} \frac{1}{\gamma_j} \left(\sum_{i \in I} \frac{\bar{C}_i}{\bar{C}} \frac{1}{\gamma_i} \right)^{-1}
\end{aligned}$$

Define $\bar{\gamma}$ as:

$$\bar{\gamma} = \left(\sum_{i \in I} \frac{\bar{C}_i}{\bar{C}} \frac{1}{\gamma_i} \right)^{-1}$$

Then the weight in the marginal basket can be expressed as:

$$\psi_i^M = \frac{\bar{C}_j}{\bar{C}} \frac{\bar{\gamma}}{\gamma_j}$$

Let us now log-linearize the household's first-order conditions. Starting with the intratemporal conditions:

$$-\gamma_i c_{i,t} = \phi n_t - (w_t - p_{i,t})$$

Now let us sum up over each of the equations $i \in I$ with weights ψ_i^M. One obtains:

$$-\bar{\gamma} \sum_{i \in I} \frac{\bar{C}_i}{\bar{C}} c_{i,t} = \phi n_t - (w_t - p_t^M)$$

Where $p_t^M = \sum_{i \in I} \psi_i^M p_{i,t}$ is the price of the marginal consumption basket. Note that by log-linearizing equation for total consumption, one has:

$$c_t = \sum_{i \in I} \frac{\bar{C}_i}{\bar{C}} c_{i,t}$$

One can then rewrite the intratemporal condition as:

$$\bar{\gamma} c_t + \phi n_t = w_t - p_t^M$$

Denote:

$$\tilde{w}_t = w_t - p_t$$
$$\tilde{p}_t^M = p_t^M - p_t$$

Where p_t is the GDP deflator. That is, it is the cost of the total consumption basket. Then one arrives at the intratemporal condition:

$$\tilde{w}_t = \bar{\gamma} c_t + \phi n_t + \tilde{p}_t^M$$

This resembles the intratemporal condition in a standard new Keynesian model, but with the additional term $\tilde{p}_t^M$. The reason that this arises is because it is the nominal wage relative to the price of the marginal consumption basket that is the relevant marginal rate of transformation when considering additional labour supply.

Now let us focus on the intertemporal conditions. Log-linearization yields:

$$c_{i,t} = E_t c_{i,t+1} - \frac{1}{\gamma_i}(i_t - E_t \pi_{i,t+1})$$

Summing with weights $\bar{C}_i / \bar{C}$ yields:

$$c_t = E_t c_{t+1} - \frac{1}{\bar{\gamma}}(i_t - E_t \pi_{t+1}^M) \tag{3.7}$$

For intertemporal subsitution decisions, it is the price change in the marginal basket that matters.

3.2.2 A More General Result

Suppose that households have a utility function of one of the form:

$$U_t = E_t \sum_{h \geq 0} \beta^h \left(u(C_1, C_2, ..., C_I) - \frac{N_t^{1+\phi}}{1+\phi} \right)$$

Where $u(C_1, C_2, ..., C_I)$ has one of the two following forms:

$$u(C_1, C_2, ..., C_I) = h\left(\sum_{i \in I} f_i(C_i)\right)$$
$$u(C_1, C_2, ..., C_I) = h\left(\prod_{i \in I} f_i(C_i)\right)$$

In 3.A.1, I show that with standard restrictions on f_i and h, one can derive the same intratemporal and intertemporal conditions as above. As such, this result applies to a broad class of non-homothetic preferences, rather than just the additive preferences case.

3.2.3 Final Goods Firms

For each sector i, there is a continuum of competitive final goods firms. They sell the final good $C_{i,t}$. They purchase intermediate goods of type i from a continuum of unit mass of intermediate goods firms of type i. The intermediate goods firms are indexed k. Intermediate goods are transformed into final goods by the Dixit-Stiglitz aggregator. As such, these firms face the standard maximization problem:

$$\max_{C_{i,k,t}} P_{i,t}C_{i,t} - \int_0^1 P_{i,k,t}C_{i,k,t}dk$$

Subject to:

$$C_{i,t} = \left(\int_0^1 C_{i,k,t}^{\frac{\epsilon-1}{\epsilon}}dk\right)^{\frac{\epsilon}{\epsilon-1}}$$

Profit maximization then implies the standard demand curve:

$$C_{i,k,t} = \left(\frac{P_{i,k,t}}{P_{i,t}}\right)^{-\epsilon} C_{i,t}$$

Note that I assume that the elasticity of substitution is the same across sectors.

3.2.4 Intermediate Goods Firms

For each sector i, there is a continuum of unit mass of monopolistically competitive intermediate goods producers, indexed k. These firms have a constant returns to scale production function:

$$Y_{i,k,t} = N_{i,k,t}$$

Firms a subject to a Rotemberg pricing friction, with an adjustment cost φ that is assumed to be homogeneous across sectors.[8] The nominal profits of the firm in period t are given by:

$$D_{i,k,t} = P_{i,k,t}Y_{i,k,t} - (1-\tau)W_t N_{i,k,t} - \frac{\varphi}{2}\left(\frac{P_{i,k,t}}{P_{i,k,t-1}} - 1\right)^2 P_{i,t}Y_{i,t} + T_{i,t}$$

Here, τ represents the standard employment subsidy paid by the fiscal policymaker. It is set to $\tau = 1/\epsilon$ to ensure an efficient steady-state. As usual, the cost of this employment subsidy in sector i is funded by a lump-sum tax $T_{i,t}$, which is given by:

$$T_{i,t} = -\tau W_t N_{i,t}$$
$$N_{i,t} = \int_0^1 N_{i,k,t}dk$$

I assume that firms maximize the expected discounted value of future profits, relative to some price index. The price index chosen matters because it determines which stochastic discount factor is appropriate. For simplicity, I assume that firms maximize the following sum:

$$E_t \sum_{h\geq 0} \beta^h \left(\frac{C_{i,t+h}}{C_{i,t}}\right)^{-\gamma_i} \frac{D_{i,k,t}}{P_{i,t}}$$

Note the choice of stochastic discount factor. If firms are maximizing the discounted value of future profits relative to $P_{i,t}$, then this is the appropriate discount factor. The reason is that a marginal change in $D_{i,t+h}/P_{i,t+h}$ gives the consumer additional income of $P_{i,t+h}$ times that change in time $t+h$. This purchases one unit of $C_{i,t+h}$, and so the appropriate marginal utility to use is $\beta^h C_{i,t+h}^{-\gamma_i}$.

Firms are subject to the demand curve given above. As such, they maximize the standard problem:

$$\max_{\{P_{i,k,t}\}_{t=0}^{\infty}} E_t \sum_{h\geq 0} \beta^h \left(\frac{C_{i,t+h}}{C_{i,t}}\right)^{-\gamma_i} \left(\left(\frac{P_{i,k,t}}{P_{i,t}}\right)^{1-\epsilon} Y_{i,t} - (1-\tau)\frac{W_t}{P_{i,t}}\left(\frac{P_{i,k,t}}{P_{i,t}}\right)^{-\epsilon} Y_{i,t} - \frac{\varphi}{2}\left(\frac{P_{i,k,t}}{P_{i,k,t-1}} - 1\right)^2 Y_{i,t}\right)$$

[8] I use a Rotemberg rather than Calvo pricing friction to simplify the algebra. This does not affect the sectoral Phillips curves, which would have the same form under either friction. That said, the choice would naturally affect the derivation of the welfare function below; under Calvo pricing, price dispersion would affect welfare.

Taking the first-order condition, and imposing a symmetric equilibrium yields:

$$\epsilon - 1 = \epsilon(1-\tau)\frac{W_t}{P_{i,t}} - \varphi\Pi_{i,t}(\Pi_{i,t}-1) + \beta\varphi\left(\frac{C_{i,t+1}}{C_{i,t}}\right)^{-\gamma_i}\frac{Y_{i,t+1}}{Y_{i,t}}\Pi_{i,t+1}(\Pi_{i,t+1}-1)$$

If $\tau = 1/\epsilon$, then the steady-state is:

$$\frac{\bar{W}}{\bar{P}_i} = 1$$

Note that because this relationship holds for all sectors i, it implies that in steady-state relative prices are equal to unity. Log-linearizing around the steady-state, and using the relative price notation above, then yields the Phillips curve for sector i:

$$\pi_{i,t} = \beta\pi_{i,t+1} + \lambda(\tilde{w}_t - \tilde{p}_{i,t})$$

Where $\lambda = (\epsilon-1)/\varphi$. I then add a cost-push shock $\eta_{i,t}$ for each sector, so that the Phillips curve becomes:

$$\begin{aligned}\pi_{i,t} &= \beta\pi_{i,t+1} + \lambda(\tilde{w}_t - \tilde{p}_{i,t}) + \eta_{i,t} \\ \eta_{i,t} &= \rho\eta_{i,t-1} + \varepsilon_{i,t}\end{aligned} \tag{3.8}$$

I assume that cost-push shocks in each sector have persistence ρ. $\varepsilon_{i,t}$ is a mean-zero shock that is independent over time. I impose no restrictions on the covariances between the innovations in each sector. Note also that profits in each sector i are equal to zero in steady-state.

3.2.5 Fiscal Policy

The only role of fiscal policy in this model is to provide employment subsidies to firms in each sector to ensure an efficient steady-state. As noted above, the cost of the subsidy is then rebated to firms as a lump-sum transfer, ensuring a balanced budget.

3.2.6 Monetary Policy

To close the model, I assume that the monetary policymaker sets monetary policy according to a simple Taylor rule. Moreover, I assume that the Taylor rule responds to the GDP deflator.

$$1 + i_t = \beta^{-1}\left(\frac{P_t}{P_{t-1}}\right)^{r_\pi}$$

Log-linearizing then gives:

$$i_t = r_\pi \pi_t$$

In section 3.5, I will study how the economy responds to relative price shocks assuming that the central bank uses this very standard rule, as well as solving for the optimal policy under discretion and commitment.

3.3 An Alternative IS and Phillips Curve

I now show that the standard IS and Phillips curve can be obtained in this model, but in terms of the inflation rate on the marginal consumption basket. I begin by solving for real wages. Recall that the intratemporal condition is:

$$\bar{\gamma} c_t + \phi n_t = \tilde{w}_t - \tilde{p}_t^M$$

The goods market clearing condition and production function imply that:

$$y_t = c_t = n_t$$

As such, real wages are given by:

$$\tilde{w}_t = \tilde{p}_t^M + (\bar{\gamma} + \phi) y_t$$

With an expression for real wages in hand, I can then turn to the Phillips curve. I sum over the sectoral Phillips curves given in (3.8), using the weights in the marginal consumption basket. This gives:

$$\pi_t^M = \beta E_t \pi_{t+1}^M + \lambda(\tilde{w}_t - \tilde{p}_t^M) + \eta_t^M \tag{3.9}$$

Then, using the expression for the real wage, I obtain:

$$\begin{aligned} \pi_t^M &= \beta E_t \pi_{t+1}^M + \kappa y_t + \eta_t^M \\ \kappa &= (\bar{\gamma} + \phi)\lambda \end{aligned} \tag{3.10}$$

Equation (3.10) gives the standard Phillips curve formulation, but in terms of π_t^M, the inflation rate of the marginal consumption basket.

To derive the IS curve, I take the intertemporal condition given in (3.7), and use the

market clearing condition ($c_t = y_t$) to give:

$$y_t = E_t y_{t+1} - \sigma(i_t - E_t \pi^M_{t+1}) \tag{3.11}$$

This is the standard IS curve, but with the real rate here being given by the nominal interest rate less the expected inflation rate of the marginal consumption basket.

3.4 A Four Equation New Keynesian Model

In section 3.3, I showed how one could obtain a standard IS and Phillips curve, but in terms of the rate of inflation on the marginal consumption basket. I now show how to re-write the model in terms of the average consumption basket, that is in terms of consumer price inflation. This shows how relative price shocks affect conventional measures of inflation, which typically correspond to central bank objectives. Moreover, as I will argue below, the rate of inflation on the average consumption basket is the welfare relevant inflation metric in this model.

3.4.1 Demand Side

I take equation (3.7) and re-arrange to give:

$$c_t = E_t c_{t+1} + \frac{1}{\bar{\gamma}}(E_t \tilde{p}^M_{t+1} - \tilde{p}^M_t) - \frac{1}{\bar{\gamma}}(i_t - E_t \pi_{t+1})$$

As such, the Euler equation can be written as in a standard model, with the additional fo the middle term on the right-hand side, which is the expected future change in the relative price of the marginal basket. For any given real interest rate $i_t - E_t \pi_{t+1}$ (where the real interest rate is expressed relative to the GDP deflator), if the marginal basket is expected to get relatively more expensive, then this causes consumers to bring forward purchases and pushes up consumption.

I then combine with the goods market clearing condition:

$$\begin{aligned} y_t &= E_t y_{t+1} + \alpha(E_t \tilde{p}^M_{t+1} - \tilde{p}^M_t) - \sigma(i_t - E_t \pi_{t+1}) \\ \alpha &= \frac{1}{\bar{\gamma}} \\ \sigma &= \frac{1}{\bar{\gamma}} \end{aligned}$$

It is instructive to iterate this equation forward. If $|\rho| < 1$, so that cost-push shocks are

stationary, then relative prices are stationary. As such:

$$\lim_{h\to\infty} \tilde{p}^M_{t+h} = 0$$

One can then iterate the Euler equation forward to obtain:

$$y_t = -\alpha\tilde{p}^M_t - \frac{1}{\bar{\gamma}}\sum_{h\geq 0}(E_t i_{t+h} - E_t\pi_{t+h+1})$$

If relative prices are stationary, and consumers have unconstrained asset market participation, then only current relative prices matter for consumption in partial equilibrium. This stems from absence of discounting in the IS curve. If the relative price of the marginal basket is high, then the undiscounted sum of all future relative marginal inflation rates must be low. Note that in general equilibrium, expected future shocks to relative prices will still affect consumption, because they change expectations of nominal interest rates i_{t+h} and average inflation π_{t+h}.

3.4.2 Supply Side

I begin by aggregating across the sectoral Phillips curves given in equation (3.8), but now using the weights in the consumption basket. This gives:

$$\pi_t = \beta E_t\pi_{t+1} + \lambda\tilde{w}_t + \eta_t$$

The relative price term drops out because it is the price relative to the GDP deflator. The average deviation from a price index, using the weights in that price index, must be zero. One can then substitute in the expression for real wages to obtain:

$$\pi_t = \beta E_t\pi_{t+1} + \kappa y_t + \lambda\tilde{p}^M_t + \eta_t$$

An increase in the relative price of the marginal basket increases the nominal wage relative to the average price index, because of the intratemporal condition given above. This results in higher marginal costs, pushing up inflation. As such, relative price shocks have a first-order effect on inflation.

3.4.3 Relative Price Inflation

I now derive a law of motion for the relative price of the marginal basket. First, define the relative inflation rate of good i as $\tilde{\pi}_{i,t}$. That is:

$$\tilde{\pi}_{i,t} = \pi_{i,t} - \pi_t$$

One can subtract the Phillips curve of good i from the aggregate Phillips curve to give:

$$\tilde{\pi}_{i,t} = \beta E_t \tilde{\pi}_{i,t+1} - \lambda \tilde{p}_{i,t} + \tilde{\eta}_{i,t}$$

The relative price of good i then evolves according to:

$$\tilde{p}_{i,t} = \tilde{p}_{i,t-1} + \tilde{\pi}_{i,t}$$

Note that $\tilde{\eta}_{i,t} = \eta_{i,t} - \eta_t$. Importantly, these expressions imply that relative prices and relative inflation rates do not depend on policy. This stems from the assumption of a constant returns to scale production function, which means that $y_{i,t}$ does not enter the sectoral Phillips curves.[9]

Now let's aggregate over i using the weights in the marginal consumption basket. One obtains:

$$\pi_t^M = \beta E_t \pi_{t+1}^M + \lambda \tilde{w}_t - \lambda \tilde{p}_t^M + \eta_t^M$$

Subtracting the aggregate Phillips curve as before yields:

$$\tilde{\pi}_t^M = \beta E_t \tilde{\pi}_{t+1}^M - \lambda \tilde{p}_t^M + \tilde{\eta}_t^M$$

Note that if the weights in the average consumption basket are ψ_i^A, and the weights in

[9] Of course, constant returns to scale is to some extent restrictive, and may miss some of the channels through which relative prices affect macro dynamics. Nonetheless, it allows me to preserve tractability, and consequently to capture many of the channels through which relative prices are important for inflation and output in a model that can be solved analytically, which in turn yields most of the insights in section 3.5 and 3.7.

the marginal basket are ψ_i^M:

$$\begin{aligned}
\eta_t &= \rho\eta_{t-1} + \varepsilon_t \\
\varepsilon_t &= \sum_{i\in I} \psi_i^A \varepsilon_{i,t} \\
\tilde{\eta}_t^M &= \rho\tilde{\eta}_{t-1}^M + \tilde{\varepsilon}_t^M \\
\tilde{\varepsilon}_t^M &= \sum_{i\in I} (\psi^M - \psi_i^A)\varepsilon_{i,t}
\end{aligned}$$

There can be a positive or negative covariance between ε_t and $\tilde{\varepsilon}_t^M$.

To reduce the number of endogenous variables, it proves useful to re-write the equations above purely in terms of the relative price of the marginal basket $\tilde{p}_t^M$, rather than both the relative price and the relative inflation rate. This helps economize on the number of variables and equations, and is also (in my view) very instructive. To do this, re-write the equation for $\tilde{\pi}_t^M$ using lag notation:

$$\tilde{p}_t^M(1-L) = \beta\tilde{p}_t^M(L^{-1}-1) - \lambda\tilde{p}_t^M + \tilde{\eta}_t^M$$

Using standard techniques, one can then solve for $\tilde{p}_t^M$ as:

$$\begin{aligned}
\tilde{p}_t^M &= \zeta\tilde{p}_{t-1}^M + \frac{\zeta}{1-\beta\zeta\rho}\tilde{\eta}_t^M \\
\zeta &= \frac{1+\beta+\lambda-\sqrt{(1+\beta+\lambda)^2-4\beta}}{2\beta}
\end{aligned}$$

Relative prices are mean-reverting in this model. When the relative price of good i is high, then marginal costs relative to the price of good i will be lower than for other goods, and so by the Phillips curve the inflation rate of good i must be low relative to that of other goods. This situation persists until relative prices are brought back into line. The strength of this self-correcting mechanism is determined by the coefficient ζ. Note that ζ is decreasing in λ. The larger is λ, the stronger is the pass through from marginal costs to inflation, and so the more powerful is the self-correction in relative prices via marginal costs. The persistence of relative price deviations is also decreasing in β. When β is higher, firms are more forward-looking. Future real marginal costs more strongly affect firms decisions today, meaning that the self-correcting mechanism is stronger.

If $\rho > 0$, then $\tilde{\eta}_t^M$ an AR(1). Shocks to relative prices can then lead to hump-shaped responses in $\tilde{p}_t^M$. Specifically, suppose that $\tilde{p}_{t-1}^M = 0$, and suppose that there is a shock

$\tilde{\varepsilon}_0^M$ in time zero. Absent further shocks, the path of $\tilde{p}_t^M$ is given by:

$$\tilde{p}_t^M = \zeta \frac{1}{1-\beta\zeta\rho} \frac{\zeta^{t+1} - \rho^{t+1}}{\zeta - \rho} \tilde{\varepsilon}_0^M$$
$$\Rightarrow \frac{\tilde{p}_1^M}{\tilde{p}_0^M} = \rho + \zeta$$

As such, the model yields hump-shaped impulse responses in $\tilde{p}_t^M$ if $\rho + \zeta > 1$.

3.4.4 The Four Equation Model

Alternatively, if one wants to write the model in terms of the inflation rate of the average consumption basket π_t, one can combine equations

I now combine the equations above to give a four equation new Keynesian model:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa y_t + \lambda \tilde{p}_t^M + \eta_t \tag{3.12}$$
$$y_t = E_t y_{t+1} + \alpha(E_t \tilde{p}_{t+1}^M - \tilde{p}_t^M) - \sigma(i_t - E_t \pi_{t+1}) \tag{3.13}$$
$$\tilde{p}_t^M = \zeta \tilde{p}_{t-1}^M + \zeta \frac{1}{1-\beta\zeta\rho} \tilde{\eta}_t^M \tag{3.14}$$
$$i_t = r_\pi \pi_t \tag{3.15}$$

As noted above, the two shock terms η_t and $\tilde{\eta}_t^M$ are AR(1) processes with persistence ρ, and they may covary positively or negatively. The appealing feature of this model is that even allowing an arbitrary number of sectors with idiosyncratic shocks, and under a range of utility functions, one can summarize the key effects of relative price shocks in a new Keynesian model with just one additional equation.

3.5 Implications of Relative Price Shocks in the RANK Model

In section 3.4, I showed that relative price shocks have first-order effects on macroeconomic dynamics under non-homothetic preferences. I now sketch two implications of this. In section 3.5.1, I solve the model in which policy is set using a simple Taylor rule, and show that the presence of relative price shocks generates considerable output volatility but far less inflation volatility. This implies that an econometrician who estimated the textbook model on data generates by my model would obtain a downwards biased estimate of the Phillips curve slope. In section 3.5.2, I show that the optimal policy response when relative prices are away from steady-state involves allowing inflation to rise above

target while keeping real interest rates low. This contrasts sharply with the response to either a demand or cost-push shock in the textbook model.

The effects described in this section will only be quantitatively important in practice if the relative price of the marginal basket varies meaningfully over time. In appendix 3.B, I assess whether shocks to the relative price of the marginal basket are likely to be large enough to matter for business cycle dynamics in practice by constructing an approximation of the marginal consumption basket and calculating the difference between the inflation rate of the marginal and average consumption basket over time. I argue that the differences between the two inflation rates is likely small most of the time, but can be several percentage points in periods with large shocks to energy or food prices.

3.5.1 An Apparently Flat Phillips Curve

I begin by solving the model assuming that monetary policy is set according to the simple Taylor rule outlined in section 3.2. In section 3.A.3, I show that the model with a Taylor rule can be solved analytically using the method of undetermined coefficients. The solution takes the form:

$$\pi_t = A_1 \tilde{p}_t^M + A_2 \tilde{\eta}_t^M + A_3 \eta_t$$
$$y_t = B_1 \tilde{p}_t^M + B_2 \tilde{\eta}_t^M + B_3 \eta_t$$

I show that the coefficients are given by:

$$A_1 = \frac{(\lambda - \kappa\alpha)(1-\zeta)}{(1-\beta\zeta)(1-\zeta) + \kappa\sigma(r_\pi - \zeta)}$$
$$A_2 = \frac{\zeta\rho}{1-\beta\zeta\rho} \frac{\beta(1-\zeta)(1-\rho) - \kappa\sigma(r_\pi - 1)}{(1-\beta\rho)(1-\rho) + \kappa\sigma(r_\pi - \rho)} \frac{(\lambda - \kappa\alpha)}{(1-\beta\zeta)(1-\zeta) + \kappa\sigma(r_\pi - \zeta)}$$
$$A_3 = \frac{1-\rho}{(1-\rho)(1-\beta\rho) + \kappa\sigma(r_\pi - \rho)}$$
$$B_1 = -\frac{\alpha(1-\zeta)(1-\beta\zeta) + \sigma\lambda(r_\pi - \zeta)}{(1-\zeta)(1-\beta\zeta) + \kappa\sigma(r_\pi - \zeta)}$$
$$B_2 = -\frac{\zeta\rho}{1-\beta\zeta\rho} \frac{\sigma\left((1-\beta\rho)(r_\pi - 1) + \beta(1-\zeta)(r_\pi - \rho)\right)}{(1-\beta\rho)(1-\rho) + \kappa\sigma(r_\pi - \rho)} \frac{(\lambda - \kappa\alpha)}{(1-\beta\zeta)(1-\zeta) + \kappa\sigma(r_\pi - \zeta)}$$
$$B_3 = -\frac{\sigma(r_\pi - \rho)}{(1-\rho)(1-\beta\rho) + \kappa\sigma(r_\pi - \rho)}$$

Note that the sign of the effect of $\tilde{p}_t^M$ on π_t (that is, coefficient A_1) is governed by the sign of $\lambda - \kappa\alpha$. This term represents the net effect of net effect of relative prices on inflation. The first term, λ, is the direct effect through the intratemporal substitution conditions and the increase in real wages. The second term, $\kappa\alpha$, is an indirect effect through output.

For a given path of average inflation, higher $\tilde{p}_t^M$ lower demand as unconstrained households delay purchases. Lower output reduces wages through the intratemporal condition, lowering inflation. One can solve for the net effect as follows:

$$\begin{aligned}\lambda - \kappa\alpha &= \lambda(1 - (\bar{\gamma} + \phi)\alpha) \\ &= \lambda\left(1 - (\bar{\gamma} + \phi)\frac{1}{\bar{\gamma}}\right) \\ &= -\lambda\frac{\phi}{\bar{\gamma}}\end{aligned}$$

As such, the net effect of an increase in the relative price of the marginal consumption basket is to lower inflation. However, inflation only falls slightly, because the total effect is the net of two partially offsetting effects.

By contrast, an increase in the relative price $\tilde{p}_t^M$ leads to a sharp decline in output, as summarized y the coefficient B_1. This coefficient comprises the sum of two effects, both of which push output lower. First, higher relative prices of the marginal basket today implies that the relative price of this basket is expected to fall, because relative prices mean-revert in the long-run. This decreases demand through intertemporal substitution, as summarized by the α coefficient. An additional fall in output results from the real interest rate response to the upwards pressure on inflation resulting from the presence of $\tilde{p}_t^M$ in the Phillips curve. This effect is summarized by the $\sigma\lambda(r_\pi - \zeta)$ coefficient.

The denominators of A_1 and B_1 are standard, and comprise two general equilibrium considerations. First, persistent shocks are amplified by the expectation terms in the Phillips and IS curves; these dynamic general equilibrium effects are summarized by $(1 - \beta\zeta)(1 - \zeta)$. Second, the effects of these shocks on inflation is dampened if real interest rates rise in response. This effect is summarized by $\kappa\sigma(r_\pi - \zeta)$.

If $\rho \neq 0$, then $A_2 \neq 0$ and $B_2 \neq 0$. That is, if cost-push shocks are persistent, then the marginal inflation shock $\tilde{\eta}_t^M$ has an additional effect on output and inflation beyond its effect on $\tilde{p}_t^M$. The reason for this is that the expectation of future relative prices is then a function of both $\tilde{p}_t^M$ and $\tilde{\eta}_t^M$. Formally:

$$\begin{aligned}E_t\tilde{p}_{t+1}^M &= \zeta\tilde{p}_t^M + \frac{\zeta}{1 - \beta\zeta\rho}E_t\tilde{\eta}_{t+1}^M \\ &= \zeta\tilde{p}_t^M + \frac{\zeta\rho}{1 - \beta\zeta\rho}\tilde{\eta}_t^M\end{aligned}$$

These effects may be large if ρ is close to unity.

The responses to the pure cost-push shock, A_3 and B_3, are the same as in a standard

new Keynesian model.[10] It is worth noting, though, that cost-push shocks in this model may be accompanied by a positive or negative marginal inflation shock, depending on which sectors are most affected by the shock. As such, a standard model captures only "pure" cost-push shocks, and does not differentiate between those originating in different sectors. My model highlights that depending on the origin of the shock, the response of macroeconomic variables may be very different.

Figure 3.2 plots the response to a relative price shock in a calibrated example with $\rho = 0$. The baseline calibration sets $\epsilon = 6$ and $\varphi = 100$, so that $\lambda = 0.05$. This in turn implies that the persistence of relative price shocks is $\zeta = 0.80$. I set $\bar{\gamma} = 2$, so that $\alpha = \sigma = 0.5$. I set $\phi = 1$, which implies a Phillips curve slope $\kappa = 0.15$. I calibrate the policy response to inflation as $r_\pi = 1.25$. In appendix 3.C, I show how the IRFs vary with different calibrations.

In this case, the relative price $\tilde{p}_t^M$ follows an AR(1) process with persistence ζ, and output and inflation are simply linear functions of $\tilde{p}_t^M$. As discussed above, a rise in $\tilde{p}_t^M$ reduces inflation, but the decrease is slight because it is the sum of two partially offsetting effects. Output declines by much more. The decline in output and inflation is partially offset by a fall in the real interest rate, which responds to the decline in inflation via the simple Taylor rule.

Figure 3.3 plots the impulse responses in the case with $\rho = 0.6$. Because $\rho + \zeta > 1$, the response of $\tilde{p}_t^M$ to the shock is hump-shaped. The persistence of relative prices also generates hump-shaped responses in the other endogenous variables. Because of the forward-looking nature of the model, the high persistence of the relative price shock leads to a larger effect on output, inflation, and real interest rates. As in the $\rho = 0$ case, though, the response of output is much larger than the response of inflation.

The fact that relative price shocks generate large responses in output but small responses in inflation gives the appearance of a flat Phillips curve. To demonstrate this, I add an AR(1) demand shock to equation (3.13) and an AR(1) policy shock to (3.15).[11] I then vary the standard deviation of the relative price shock innovation $\tilde{\varepsilon}_t^M$, and the

[10] It is also worth noting that in this simple model changes in demand do not affect relative prices, and so the effects of demand shocks are also the same as in the standard new Keynesian model. The reason for this is that for simplicity I have assumed a constant returns to scale production function, so that $Y_{i,k,t} = N_{i,k,t}$. As such, although changes in aggregate demand do have different effects on output across sectors under non-homothetic preferences, they do not generate cross-sector variations in marginal cost. If instead the production function featured diminishing returns to scale, then fluctuations in aggregate demand would affect relative prices. This extension of the model would give additional insights into the transmission of demand shocks, but would come at the expense of the analytical tractability that is an appealing feature of the model considered in this paper. As such, I leave analysis of the effects of demand shocks with diminishing returns to scale for future work.

[11] I set the autocorrelation of demand shocks equal to 0.8, and the autocorrelation of policy shocks at 0.3. I set the standard deviation of the innovation in each shock to 0.5. I set the standard deviation of the supply shock also equal to 0.5, and assume that η_t and $\tilde{\eta}_t^M$ are uncorrelated.

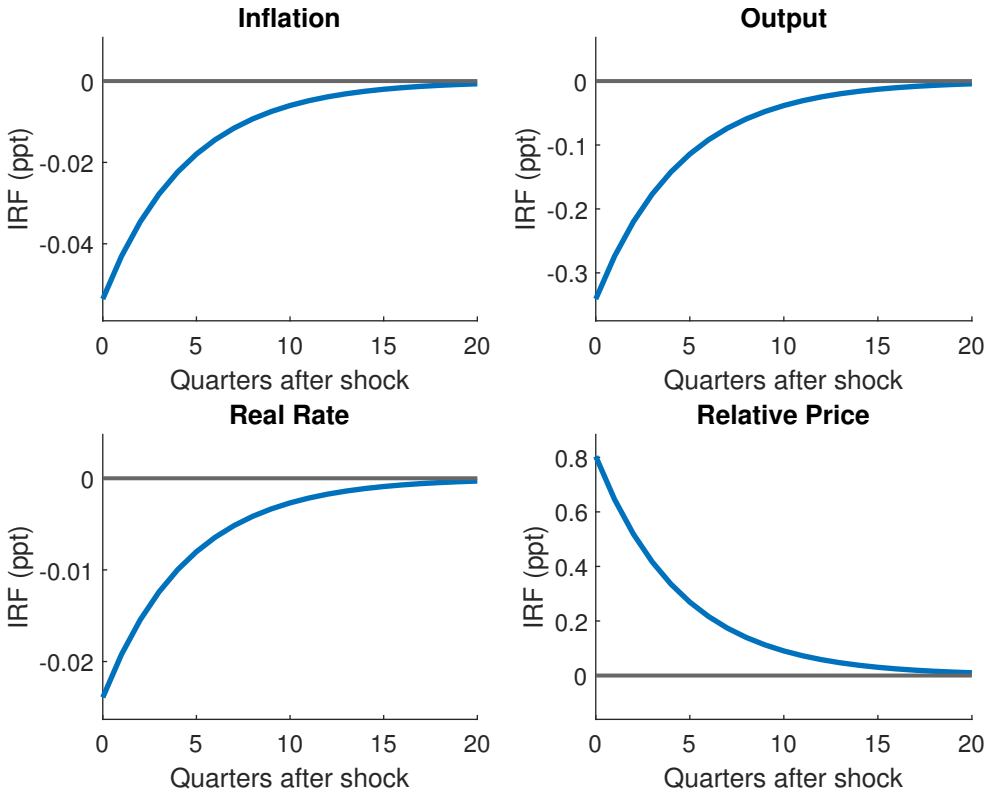


Figure 3.2: Impulse responses to a one percentage point relative price shock ($\tilde{\varepsilon}_t^M$) in the RANK model with $\rho = 0$.

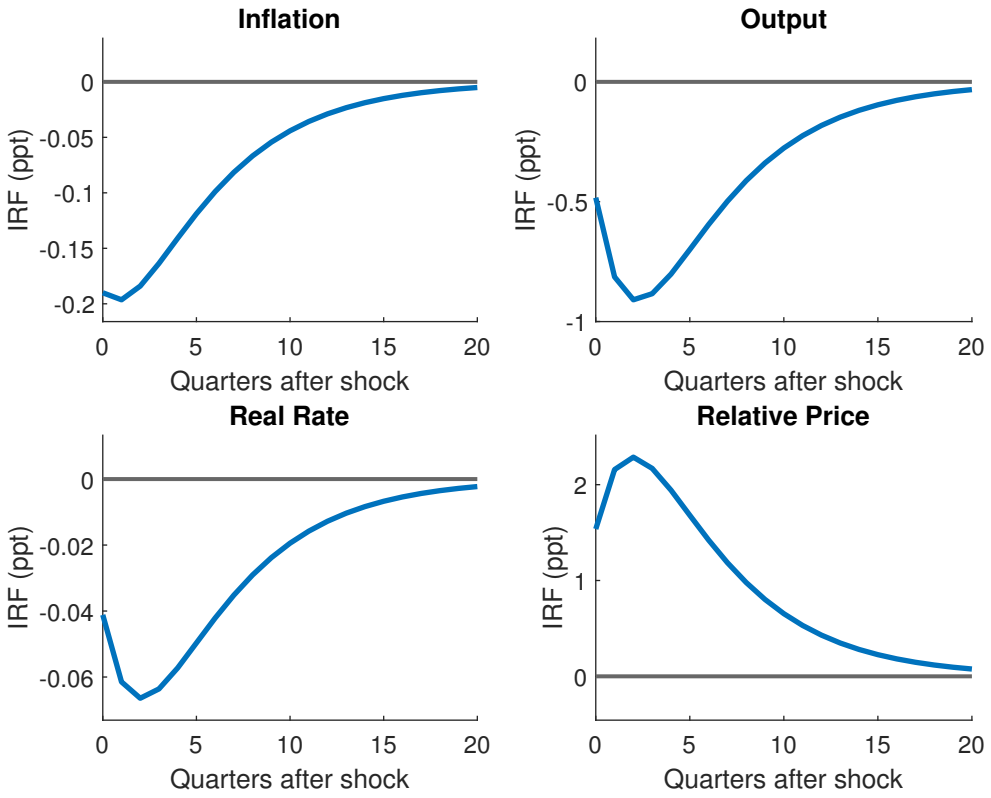


Figure 3.3: Impulse responses to a one percentage point relative price shock ($\tilde{\varepsilon}_t^M$) in the RANK model with $\rho = 0.6$.

persistence of cost-push shocks ρ. When the standard deviation of $\tilde{\varepsilon}_t^M$ is large, and ρ is high, then the standard deviation of $\tilde{p}_t^M$ is high, and this drives a large fraction of variation in demand. For each calibration, I simulate a very long sample of data from the model.[12]

I then estimate the textbook three equation new Keynesian model on the data, using a maximum likelihood approach. In doing so, I use the common assumption that supply, demand, and policy shocks are independent AR(1) processes.[13] Given the true data generating process, this assumption is likely to be misleading; the presence of relative price shocks gives (i) (negatively) correlated shocks in the supply and demand curve, and (ii) a more complex autocovariance structure for the shock processes.

Figure 3.4 shows the estimated Phillips curve slope for a range of calibrations. The dashed line shows the "correct" estimate of 0.15. Notice that where the standard deviation of the relative price shock innovation is large (to the right hand side), or when cost-push shocks are persistent (the purple series), the estimated Phillips curve slope is considerably underestimated. Indeed, in the extreme case with very large and persistent fluctuations in relative prices, the estimated Phillips curve slope falls below 0.02, a vast underestimate. Given the time-variation in the importance of the common versus idiosyncratic component of inflation discussed above, this could lead to a time-varying downwards bias in the Phillips curve slope estimate. This exercise highlights the risk of neglecting the importance of relative price shocks when estimating new Keynesian models. If these shocks do indeed affect both supply and demand, as theory suggests, then that may lead to a considerable bias in the parameter estimates obtained from a DSGE model.

3.5.2 Qualitatively Different Optimal Co-Movements

In 3.5.1, I solved for model dynamics under a simple Taylor rule. I now solve for optimal monetary policy under discretion and commitment. I focus on the discretion case, and show that the optimal policy response when relative prices are away from steady-state is to let inflation rise above target, whilst lowering the real interest rate.[14] This optimal

[12] I simulate 500 years worth of data. As such, the focus here is on the consistency of the estimator of the Phillips curve slope.

[13] Although common, this assumption is not ubiquitous. For example, Lubik and Schorfheide (2004) allow for correlated supply and demand shocks, and Smets and Wouters (2007) allow for two of the shocks in their medium-scale model to be correlated. Cúrdia and Reis (2020) present methods for dealing with correlated shocks with richer autocovariance structures in DSGE models. That paper also contains a fuller review of the related literature relaxing this assumption in estimated DSGE models.

[14] In appendix 3.A.6, I also solve analytically for the optimal policy under commitment. Just as in the discretion case. In appendix 3.C, I plot the optimal policy responses to shocks in the $\rho = 0$ and $\rho = 0.6$ cases for a variety of calibrations.

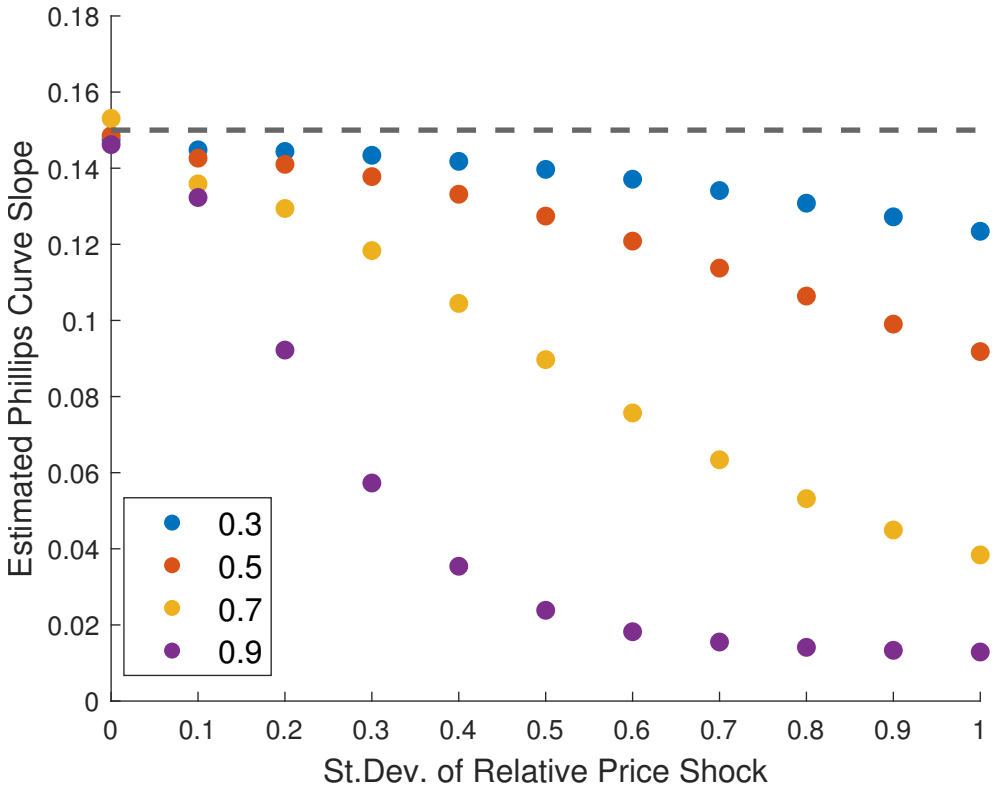


Figure 3.4: Estimated Phillips Curve slope if the econometrician neglects the first-order role of relative price shocks. The x-axis varies the standard deviation of the innovation in the relative price shock process, $\tilde{\varepsilon}_t^M$. The different coloured dots vary the persistence of the relative price shock ρ.

response differs qualitatively from the optimal response to demand or supply shocks in the textbook model.

In appendix 3.A.4, I solve for the welfare loss function, and then for the central bank's objective function in the RANK model. Somewhat surprisingly, the central bank's objective function has a standard form. Losses are given by:

$$L_t = E_t \frac{1}{2} \sum_{h \geq 0} \beta^h \left(\pi_{t+h}^2 + \Omega_y^{RANK} y_{t+h}^2 \right)$$
$$\Omega_y^{RANK} = \frac{1}{\varphi}(\bar{\gamma} + \phi)$$

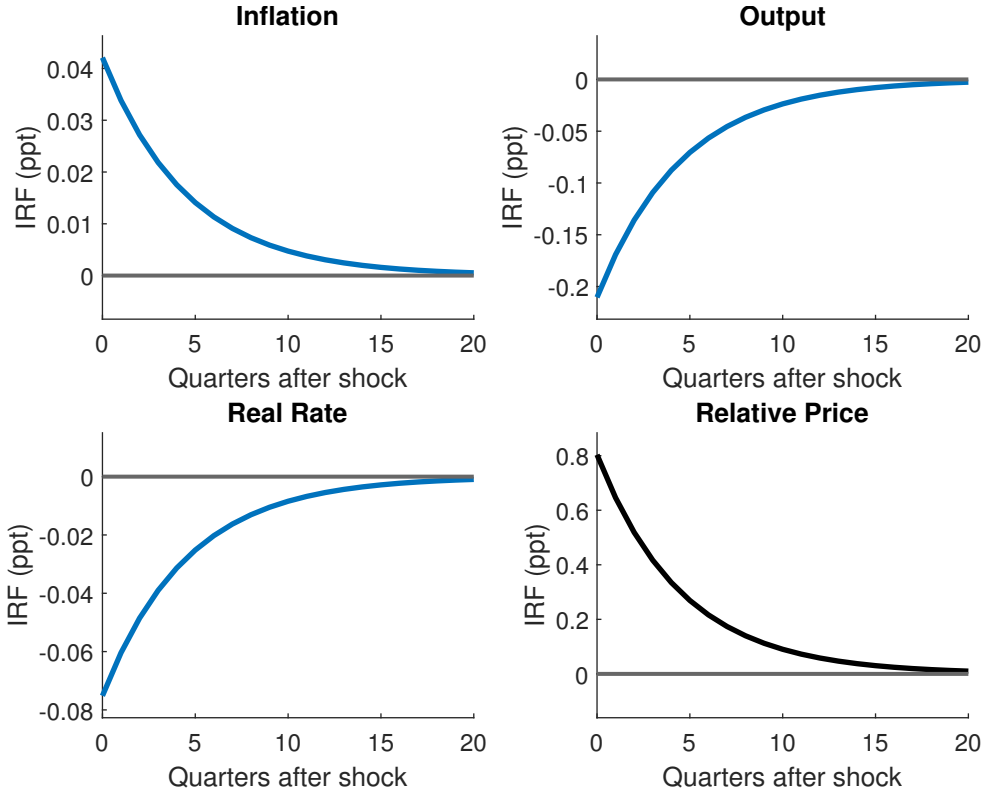
Why is the central bank's objective function standard? There are many relative price and relative inflation terms in the welfare function. However, these become "terms invariant to policy", which the central bank can ignore when optimally choosing policy, because of the exogeneity of relative prices. Again, this results from the simplifying constant returns to scale assumption that allows for tractability.

I then solve analytically for the optimal policy under discretion. I show that the optimal path for inflation takes the following form:

$$\pi_t = \chi_1 \tilde{p}_t^M + \chi_2 \tilde{\eta}_t^M + \chi_3 \eta_t$$
$$\chi_1 = \frac{\lambda \Omega_2 - \kappa \Omega_1}{\kappa^2 + \Omega_2 (1 - \beta \zeta)}$$
$$\chi_2 = \frac{\zeta}{1 - \beta \zeta \rho} \frac{\lambda \Omega_2 - \kappa \Omega_1}{\kappa^2 + \Omega_2 (1 - \beta \zeta)} \frac{\beta \rho \Omega_2}{\kappa^2 + \Omega_2 (1 - \beta \rho)}$$
$$\chi_3 = \frac{\Omega_2}{\kappa^2 + \Omega_2 (1 - \beta \rho)}$$

See appendix 3.A.5 for the derivations, and for the optimal response of output and nominal interest rates. As such, it is optimal for the central bank to respond to fluctuations in relative prices. When $\rho > 0$, then it is also optimal for the central bank to respond to the relative price shock $\tilde{\eta}_t^M$.

Figure shows the optimal policy response to a shock to relative prices in the $\rho = 0$ case. Alternatively, one can think about this as the optimal response to a situation in which there has been a large shock, which has now dissipated, but has left relative prices far away from their steady-state. Remember that a rise in $\tilde{p}_t^M$ in the $\rho = 0$ case acts like an AR(1) shock in the Phillips curve. As such, just as in the case of a standard cost-push shock, it will be optimal for the central bank to let inflation rise, but to partially offset this rise in inflation with a fall in output. Where the optimal policy response differs, however, is the response of the real interest rate. The relative price shock also acts as a large negative shock to output. In order to achieve the optimal path for output, the real



interest rate consequently has to fall.[15]

This set of macroeconomic co-movements differs qualitatively from those that are optimal in response to a pure supply or demand shock in the textbook new Keynesian model. In response to a demand shock, the central bank raises the real interest rate in order to keep the output at its natural level and inflation at target. Demand shocks consequently induce no covariance between inflation and real interest rates. In response to a supply shock, inflation rises, but so do real interest rates in order to lower output and partially offset the increase in inflation. As such, supply shocks lead to a positive covariance between real interest rates and inflation. Using the common assumption that supply and demand shocks are uncorrelated, as discussed in section 3.5.1, the covariance of real interest rates and inflation should be positive in the textbook model. In my model, by contrast, there may be prolonged periods where relative prices are away from steady-state and it is consequently optimal for inflation to rise above target and for real interest rates to be kept low.

In appendix 3.C, I show the optimal policy responses under discretion in the $\rho > 0$ case, and also show the sensitivity of the responses to different calibrations.

[15]Note that this does not depend on the calibration. To see this, note that absent any change in real interest rates output is given by $y_t = -\alpha\tilde{p}_t^M$. Then $\kappa y_t + \lambda\tilde{p}_t^M < 0$, as discussed in section 3.5.1, giving a negative response in inflation. As such, real interest rates must be lowered in order to push output above this level and allow inflation to rise above target.

3.6 A Two Agent Extension

I now extend the model by introducing hand-to-mouth households. Suppose that all households have the utility function described in section 3.2. Now suppose that there are two types of households, indexed $s \in \{U, H\}$. There is a continuum of measure $1 - \xi$ of unconstrained households (type U), who have full access to bond markets, and also own all of the equity in firms. As such, they face the following budget constraint:

$$\sum_{i \in I} P_{i,t} C_{i,t}^U + B_t^U = (1 + i_{t-1}) B_{t-1}^U + W_t N_t^U + \frac{1}{1-\xi} D_t$$

There is a continuum of measure ξ of constrained households (type H), who have no access to bond markets, and own no equity in firms. As such, their budget constraint is:

$$\sum_{i \in I} P_{i,t} C_{i,t}^H = W_t N_t^H$$

Both types have the same steady-state income and labour supply. To see this, note that bonds are in zero net supply, and profits are zero in equilibrium. The steady-state budget constraints are consequently identical. As such, the marginal and average consumption baskets are the same across types.

The first-order conditions for the unconstrained households are as before. That is:

$$\begin{aligned} \phi n_t^U + \bar{\gamma} c_t^U &= \tilde{w}_t - \tilde{p}_t^M \\ c_t^U &= E_t c_{t+1}^U - \frac{1}{\bar{\gamma}} (i_t - E_t \pi_{t+1}^M) \end{aligned}$$

For constrained households, one obtains the following intratemporal conditions:

$$\begin{aligned} \forall i, j \in I : (C_{i,t}^H)^{-\gamma_i} \frac{1}{P_{i,t}} &= (C_{j,t}^H)^{-\gamma_j} \frac{1}{P_{j,t}} \\ \forall i \in I : (C_{i,t}^H)^{-\gamma_i} \frac{1}{P_{i,t}} &= (N_t^H)^{\phi} \frac{1}{W_t} \end{aligned}$$

Log-linearizing these conditions and summing across equations gives the same intratemporal condition as for unconstrained agents:

$$\phi n_t^H + \bar{\gamma} c_t^H = \tilde{w}_t - \tilde{p}_t^M$$

However, consumption is now given by the log-linearized budget constraint:

$$c_t^H = \tilde{w}_t + n_t^H$$

Note that it is the nominal wage relative to price of the average basket (rather than the marginal basket) that enters the budget constraint.

I then aggregate across the two types. Aggregate consumption and labour supply are:

$$C_t = \xi C_t^H + (1-\xi) C_t^U$$
$$N_t = \xi N_t^H + (1-\xi) N_t^U$$

Log-linearizing gives:

$$c_t = \xi c_t^H + (1-\xi) c_t^U$$
$$n_t = \xi n_t^H + (1-\xi) n_t^U$$

One can then aggregate the intratemporal conditions to obtain the same expression for the real wage as in the RANK model:

$$\begin{aligned} \phi n_t + \bar{\gamma} c_t &= \tilde{w}_t - \tilde{p}_t^M \\ \Rightarrow \tilde{w}_t &= \tilde{p}_t^M + (\phi + \bar{\gamma}) y_t \end{aligned}$$

I then solve for the consumption of hand-to-mouth households:

$$c_t^H = (1+\phi) y_t + \frac{\phi}{\bar{\gamma} + \phi} \tilde{p}_t^M \tag{3.16}$$

The first term in (3.16) is standard in TANK models. Absent redistributive fiscal policies, constrained households' consumption rises more than proportionally with output. In general equilibrium, this amplifies the response of output to real interest rate shocks. The second term is novel. A rise in $\tilde{p}_t^M$ increases hand-to-mouth agents' consumption. The reason is that an increase in $\tilde{p}_t^M$ raises the real wage $\tilde{w}_t$ via the intratemporal condition. This effect increases constrained agents' consumption via the budget constraint.[16]

I then use the market clearing condition to express the consumption of unconstrained households in terms of aggregate output relative prices:

$$\begin{aligned} c_t^U &= \frac{1}{1-\xi}(y_t - \xi c_t^H) \\ \Rightarrow c_t^U &= \frac{1-\xi(1+\phi)}{1-\xi} y_t - \frac{\xi}{1-\xi} \frac{\phi}{\bar{\gamma} + \phi} \tilde{p}_t^M \end{aligned}$$

[16] Of course, consumption does not rise one-for-one with $\tilde{p}_t^M$; note the coefficient $\phi/(\bar{\gamma}+\phi)$ on $\tilde{p}_t^M$. Although real wages rise proportionally with the relative price of the marginal basket, this is partially offset by a substitution effect; labour supply falls as consumption increases via the intratemporal condition.

Substituting this condition into the Euler equation for the unconstrained gives:

$$
\begin{aligned}
y_t &= E_t y_{t+1} + \alpha(E_t \tilde{p}^M_{t+1} - \tilde{p}^M_t) - \sigma(i_t - E_t \pi_{t+1}) \\
\alpha &= \frac{1}{\gamma} \frac{1-\xi}{1-\xi(1+\phi)} - \frac{\xi}{1-\xi(1+\phi)} \frac{\phi}{\bar{\gamma}+\phi} \\
\sigma &= \frac{1}{\gamma} \frac{1-\xi}{1-\xi(1+\phi)}
\end{aligned}
$$

The IS curve slope σ is standard. The presence of hand-to-mouth agents whose consumption varies more than one-for-one with output amplifies the effects of real interest rates on output in general equilibrium.

The coefficient α now captures two effects. The first term in α represents an intertemporal substitution effect, as in the representative agent model. When $\tilde{p}^M_t$ rises, unconstrained consumers delay purchases, reducing consumption. The partial equilibrium effect of this is $1/\bar{\gamma}$. However, the presence of hand-to-mouth agents amplifies this effect by a factor of $(1-\xi)/(1-\xi(1+\phi))$. The second term in α is the income effect on hand-to-mouth consumers. When $\tilde{p}^M_t$ rises, hand-to-mouth consumers benefit from a positive income effect, pushing up consumption. As such, the presence of hand-to-mouth consumers has two effects on the transmission of relative price shocks to output. On the one hand, their direct response to an increase in $\tilde{p}^M_t$ is to increase consumption. On the other hand, the presence of these consumers amplifies the decrease in consumption resulting from intertemporal substitution effects on unconstrained households.

Which of the two effects dominates? Combining the two terms yields:

$$\alpha = \frac{(\bar{\gamma}+\phi) - \xi(\bar{\gamma}+\phi+\bar{\gamma}\phi)}{\bar{\gamma}(\bar{\gamma}+\phi)(1-\xi(1+\phi))}$$

The denominator is strictly positive if:

$$\xi < \frac{1}{1+\phi}$$

If this condition does not hold, then $\sigma < 0$, so the IS curve is inverted. This possibility is discussed in Bilbiie (2008). The numerator is strictly positive if:

$$\xi < \frac{\bar{\gamma}+\phi}{\bar{\gamma}+\phi+\bar{\gamma}\phi}$$

If the former condition holds, then so does the latter. As such, α is positive if the IS curve is not inverted. Moreover, if the IS curve is not inverted, then α is increasing in ξ. As such, although the partial equilibrium effect of a rise in $\tilde{p}^M_t$ is to increase hand-to-mouth households' consumption, the presence of more hand-to-mouth households

generates a larger negative effect of $\tilde{p}_t^M$ on aggregate consumption. This is because the partial equilibrium effect is outweighed by the fact that the presence of hand-to-mouth households amplifies the (negative) intertemporal substitution effect of a rise in $\tilde{p}_t^M$ on unconstrained households' consumption.

The supply side of the model and the evolution of relative prices are the same as in the representative agent version of the model. One consequently arrives at the same four equation representation of the model given in section 3.4.

3.6.1 Inequality in the TANK Model

The two agent version of the model allows us to study the effects of relative price fluctuations on consumption inequality. The cross-sectional variance of consumption is:

$$V(c_t^s) = \xi(c_t^H)^2 + (1-\xi)(c_t^U)^2 - y_t^2$$

Substituting in the expressions above and expanding out the resulting terms gives:

$$V(c_t^s) = \frac{\xi}{1-\xi}\left(\phi^2 y_t^2 + \left(\frac{\phi}{\bar{\gamma}+\phi}\right)^2 (\tilde{p}_t^M)^2 + 2\frac{\phi^2}{\bar{\gamma}+\phi} y_t \tilde{p}_t^M\right)$$

As in the textbook TANK model, consumption inequality rises when output deviates from steady-state. In this model, however, consumption inequality also increases when the relative price of the marginal consumption basket increases. Moreover, the covariance of relative prices and output affects inequality. As shown in (3.16), the hand-to-mouth consumption is high relative to average consumption when (i) output is high, and (ii) the relative price of the marginal basket is high. If the two positively covary, that drives larger deviations of hand-to-mouth consumption versus average consumption, and consequently greater business cycle fluctuations in inequality. Note that output in this model is a function of both pure cost-push shocks and relative price shocks. As such, how these shocks covary affects how consumption inequality fluctuates over the business cycle. A cost-push shock has different implications for inequality depending on the associated change in the relative price of the marginal basket. I explore this point in section 3.7.

Relative price shocks matter for inequality in this model despite the fact that unconstrained and hand-to-mouth agents have the same consumption basket.[17] This opens up a new channel for relative prices to cause fluctuations in inequality. As I will show in 3.7,

[17]Of course, one could enrich the model by assuming different steady-state incomes between the two groups. This would be plausible for the poor hand-to-mouth, who likely have lower steady-state incomes than the unconstrained. That said, for the wealthy hand-to-mouth studied by Kaplan et al. (2014), it is not obvious steady-state income would differ from unconstrained households.

this qualitatively changes the way that inequality considerations affect optimal policy responses to shocks.

3.7 Implications of Relative Price Shocks in the TANK Model

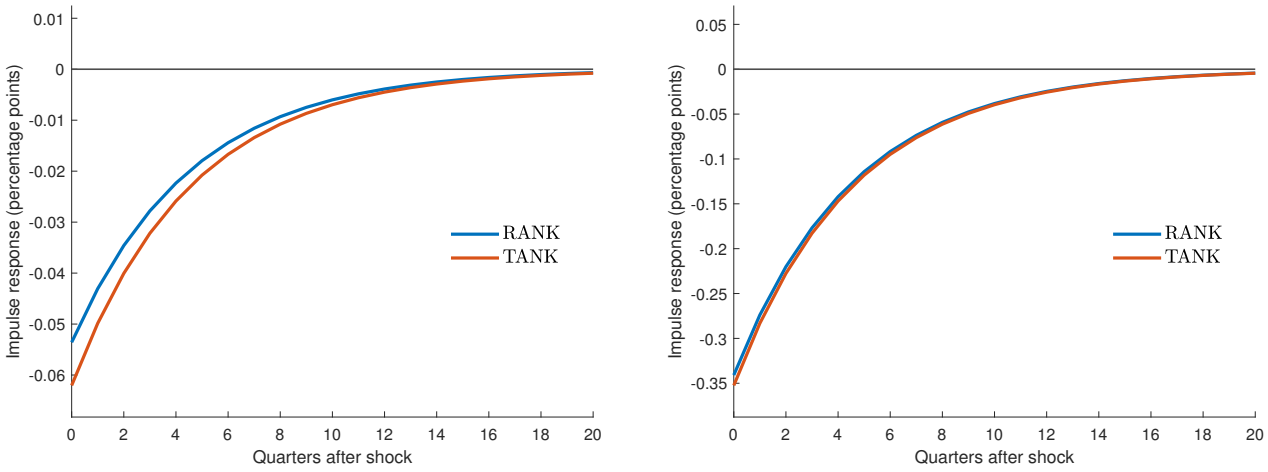
Section 3.6 establishes that relative price shocks generate fluctuations in inequality under non-homothetic preferences, even when unconstrained and hand-to-mouth households have the same steady-state income. I now sketch two main implications of this finding. In section 3.7.1, I explore the model dynamics under a Taylor rule, and demonstrate that there is a partial versus general equilibrium tension in the implications of relative price shocks for consumption inequality. In section 3.7.2, I solve for the optimal policy in the TANK model, and show that inequality considerations have a qualitatively different effect on optimal policy in this model versus the textbook model.

3.7.1 Inequality with a Taylor Rule

I begin by examining the model in which policy is set according to a simple Taylor rule. Because the model has the same four equation formulation as the RANK model, the solution obtained in 3.5 still applies, although the α and σ coefficients are now different. Figure 3.5 shows the impulse responses of the macroeconomic variables to a relative price shock in the TANK model when $\rho = 0$. I also show the responses in the RANK version of the model for comparison.

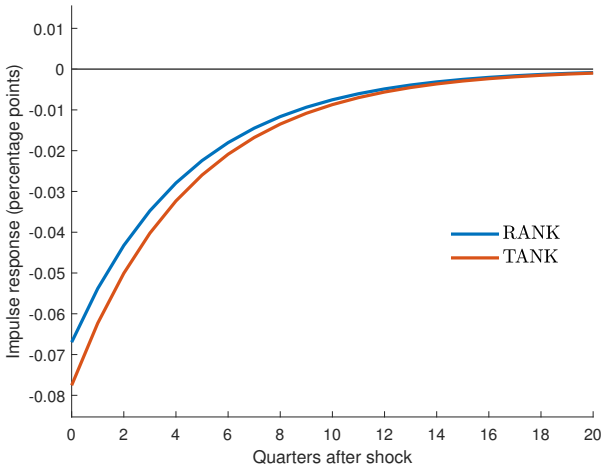
The responses of the macroeconomic variables to the shock are fairly similar. The reason for this is that although the α coefficient is larger than in the RANK model, the difference is slight. The reason for that is that, as discussed in section 3.6, introducing hand-to-mouth agents into the model has two partially-offsetting effects. On the one hand, introducing hand-to-mouth agents amplifies the negative effect of relative prices on demand stemming from the hand-to-mouth agents. On the other hand, for a given level of demand an increase in $\tilde{p}_t^M$ generates a positive income effect for the hand-to-mouth.

The response of inflation to the shock remains negative because in the two-agent

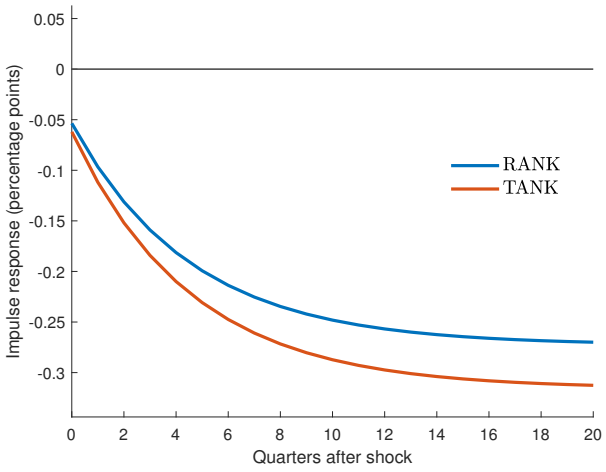


(a) Inflation

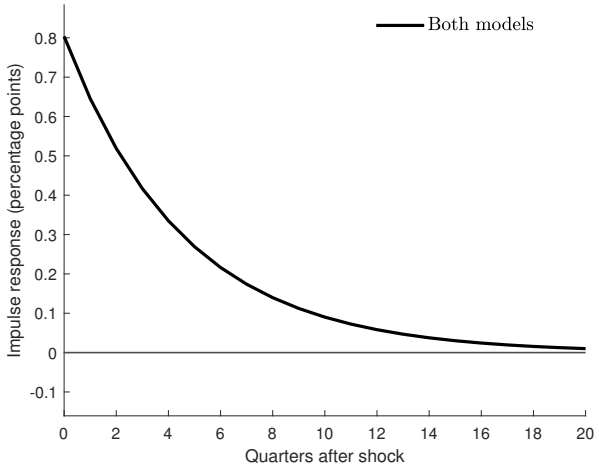
(b) Output



(c) Nominal Rate



(d) Price Level



(e) Relative Price of Marginal Basket

Figure 3.5: Impulse responses to a relative price shock in the RANK and TANK models with $\rho = 0$ under the baseline calibration.

model, one has:

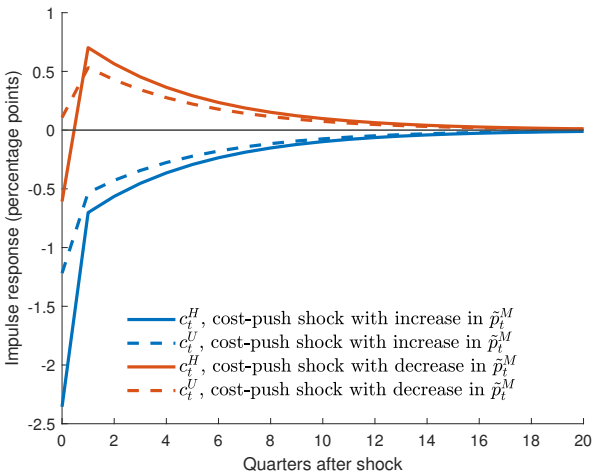
$$\begin{aligned}\lambda - \kappa\alpha &= \lambda(1 - (\bar{\gamma} + \phi)\alpha) \\ &= \lambda\left(1 - \frac{\bar{\gamma} + \phi}{\bar{\gamma}}\left(\frac{1-\xi}{1-\xi(1+\phi)}\right) + \phi\frac{\xi}{1-\xi(1+\phi)}\right) \\ &= -\lambda\frac{\phi}{\bar{\gamma}}\frac{1-\xi}{1-\xi(1+\phi)}\end{aligned}$$

Assuming $\xi(1+\phi) < 1$, so that the IS curve is not inverted, then this is also negative. Moreover, the size of the negative effect of $\tilde{p}_t^M$ on aggregate inflation increases the higher is the fraction of agents that is constrained.

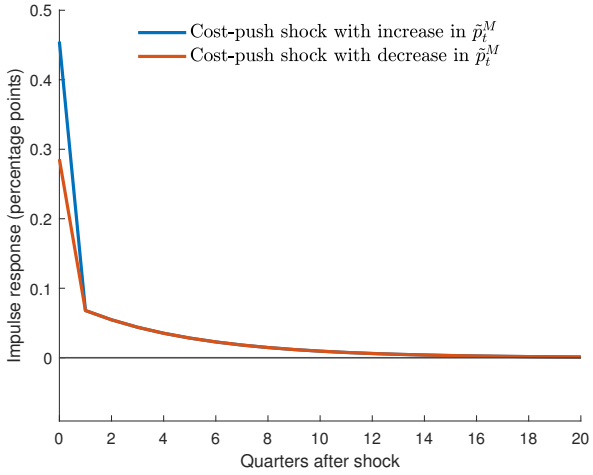
How do the responses of macroeconomic variables translate into responses of consumption inequality? Figure 3.6 explores the effects of two types of shock; the blue lines show a cost-push shock associated with an increase in the relative price of the marginal consumption basket, while the red lines show a cost-push shock associated with a decrease in the relative price of that basket. The left hand side panel shows the effects of both types of shocks on the consumption of the two agents, while the right hand side figure shocks the effects on the cross-sectional standard deviation of consumption. This figure considers the case where $\rho = 0$. The same figure is show in the case with $\rho = 0.6$ in section 3.C.

Let us focus first on the cost-push shock with an increase in the relative price of the marginal basket. As one might anticipate from the IRFs shown above, the overall contraction in output is large in this case; real interest rates rise in response to the pure cost-push shock via the Taylor rule, which depresses output, but there is a further negative effect because the increase in the relative price of the marginal basket further depresses output. Because hand-to-mouth agents' consumption is procyclical, this causes a disproportionate decrease in the consumption of hand-to-mouth agents. This is partially offset by the fact that the increase in the price of the marginal consumption basket generates a positive income effect for hand-to-mouth consumers, but the net effect is still a rather large decrease in hand-to-mouth consumption, and a substantial increase in inequality.

From $t = 1$ onwards, the pure-cost-push shock disappears because $\rho = 0$. One is left with the effect of the relative price of the marginal basket being elevated. Once again, there are two effects to consider. On the one hand, the elevated relative price of the marginal basket depresses output, which disproportionately affects hand-to-mouth output. On the other hand, a higher price of the marginal basket generates a positive income effect for the hand-to-mouth agents. In this calibration, it turns out that the net effect is that hand-to-mouth consumption remains below that of unconstrained agents,



(a) c_t^H and c_t^U



(b) Cross-sectional standard deviation of consumption

Figure 3.6: Impulse responses to a cost-push shock accompanied by positive or negative relative price shock in the TANK model with $\rho = 0$ under the baseline calibration. Formally, I assume $\eta_0 = 1$ in both cases. The blue line shows the responses if $\tilde{\eta}_t^M = 2$, while the red line shows the responses if $\tilde{\eta}_t^M = -2$.

although the cross-sectional standard deviation of consumption is considerably lower than in the first period.

Let us now turn to the case in which the relative price of the marginal consumption basket falls. In this case, the increase in the real interest rate $(i_t - E_t\pi_{t+1})$ in response to the pure cost-push shock is partially offset by the fact that the relative price of the marginal basket has fallen. The latter effect reduces $(i_t - E_t\pi_{t+1}^M)$, which is the relevant variable for unconstrained households' intertemporal substitution condition. Unconstrained households' consumption actually rises slightly. The consumption of hand-to-mouth households falls slightly, because of the negative income effect that arises when the relative price of the marginal basket falls. However, the net of these effects is a much more moderate fall in hand-to-mouth consumption, and a much smaller increase in consumption inequality. From $t = 1$ onwards, one is just left with the effects of the reduced relative price of the marginal consumption basket, and so the effects are just the mirror image of the with the case where $\tilde{p}_t^M$ rises.

These effects are perhaps slightly surprising. One might well expect that a cost-push shock generated by goods that feature heavily in the average but not the marginal consumption basket (such as food and energy) to generate a larger increase in inequality. The fact that the model generates the opposite result shows the value of explicitly incorporating the effects of relative price shifts in a two agent model to guide one's thinking.

3.7.2 Optimal Policy in TANK

I now illustrate how concern for inequality affects the optimal response to relative price shocks. In appendix 3.A.4, I show that the central bank's objective function in the TANK model is given by:

$$L_t = E_t \frac{1}{2} \sum_{h \geq 0} \beta^h \left(\pi_{t+h}^2 + 2\Omega_{yp}^{TANK} \tilde{p}_t^M y_t + \Omega_y^{TANK} y_t^2 \right) \tag{3.17}$$

$$\Omega_{yp}^{TANK} = \frac{1}{\varphi} \bar{\gamma} \phi \frac{\xi}{1-\xi}$$

$$\Omega_y^{TANK} = \frac{1}{\varphi} (\bar{\gamma} + \phi) \left(1 + \bar{\gamma} \phi \frac{\xi}{1-\xi} \right)$$

The Ω_y^{TANK} coefficient is standard. The welfare costs of output fluctuations are higher in the TANK model, because fluctuations in output generate fluctuations in inequality. The reason for this is that hand-to-mouth consumption moves more than one-for-one with output, so when output rises, the consumption of hand-to-mouth agents deviates from the consumption of unconstrained agents. However, there is now a novel term in the welfare function; the interaction between relative prices and output. Why does this term arise? As noted in equation (3.16), both output and relative prices affect hand-to-mouth consumption. The extent to which they covary consequently affects how large are the business cycle fluctuations in inequality, and hence welfare losses. In particular, if the central bank can induce a negative covariance between y_t and $\tilde{p}_t^M$, that can partially mitigate the pro-cyclicality of hand-to-mouth consumption, and consequently attenuate business cycle fluctuations in inequality.

In appendix 3.A.5 and 3.A.6, I show how this new term in the central bank's objective function affects the optimal policy under discretion and commitment. Figure 3.7 illustrates the optimal policy response to a relative price shock under discretion in the RANK and TANK models, in the case where $\rho = 0$. In the TANK model, there is a slightly smaller increase in inflation, and a slightly larger decline in output.

This is perhaps surprising. In response to a standard cost-push shock, the response of output is (perhaps substantially) smaller, and the response of output larger, in the TANK model compared to the RANK model. The reason is that, as noted above, output fluctuations are more costly in TANK than in RANK. Formally, $\Omega_y^{TANK} > \Omega_y^{RANK}$, and so the central bank optimally lets inflation rise more in response to the shock, and induces a smaller reduction in output. In response to the relative price shock, this consideration still applies, but there is now an additional effect working in the opposite direction. The presence of Ω_{yp}^{TANK} in the welfare function causes the central bank to induce a slightly reduction in output than in the RANK model. Doing so induces a larger negative

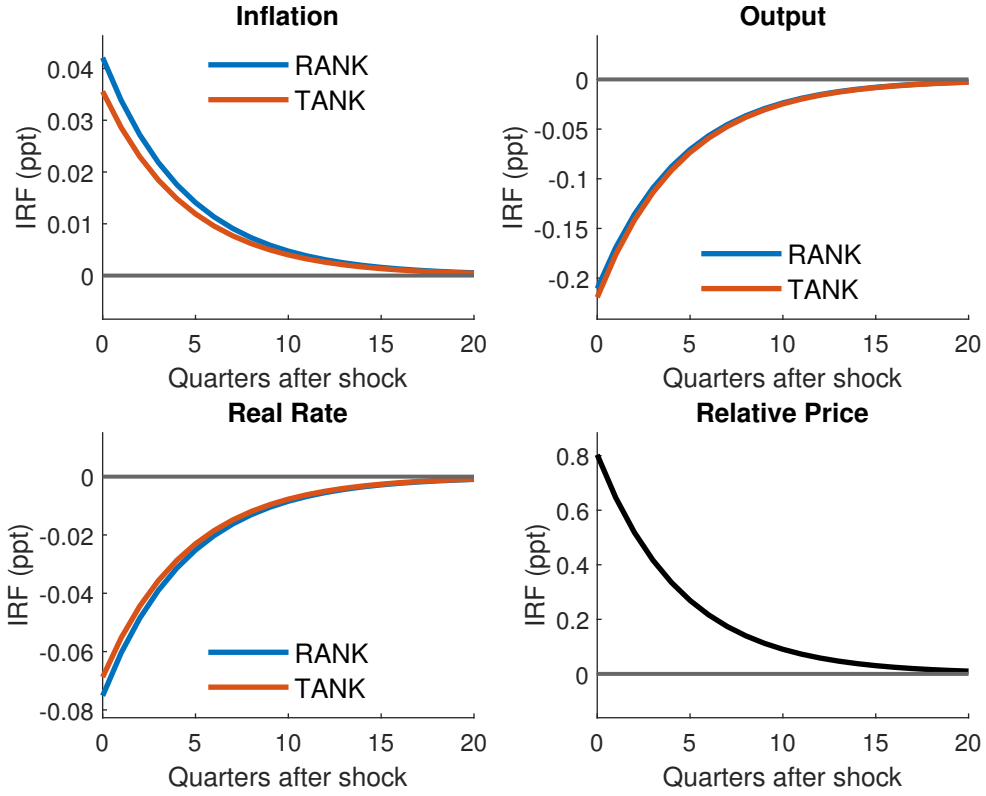


Figure 3.7: Optimal policy responses to a relative price shock under discretion in the $\rho = 0$ case, RANK vs. TANK model.

covariance between y_t and $\tilde{p}_t^M$, which partially mitigates the rise in inequality after the shock.

As such, while inequality considerations in the textbook model induce the central bank to allow smaller fluctuations in output and larger fluctuations in inflation, this is no longer necessarily the case when fluctuations are caused by relative prices fluctuating away from their steady-state values.

3.8 Conclusion

Consumption data suggest that households have non-homothetic preferences over a range of final goods. The importance of this fact in growth models is well known, but the implications for business cycle analysis have been less well studied. In this paper, I constructed a new Keynesian model with many sectors and non-homothetic preferences. I showed that under a broad class of non-homothetic preferences, the textbook new Keynesian IS and Phillips curves hold, but for a differently weighted inflation index: the inflation rate of the marginal consumption basket. I then showed that one can rewrite the model in terms of the (usual) consumption weighted inflation index, but one introduces a new variable: the relative price of the marginal consumption basket. I argued that

shocks to this relative price variable resemble a negative demand shock combined with a positive supply shock. I showed how to solve for the law of motion for this new variable to derive an analytically tractable four equation model.

I sketched two main implications of the baseline representative agent model. First, I showed that the presence of relative price shocks mutes the relationship between inflation and output. An econometrician who neglects the importance of relative price shocks would consequently underestimate the Phillips curve slope. Second, I argued that the optimal policy response to relative price shocks entails a qualitatively different co-movement between inflation and real rates from the response to pure demand or cost-push shocks in the textbook model. I then considered a TANK extension, and showed that relative price shocks strongly affect inequality. Moreover, I showed that there is a tension between the partial equilibrium and general equilibrium effects on inequality. I demonstrated that the output response to a relative price shock under optimal policy is larger in TANK than in RANK. I noted that this differs qualitatively from the case of a pure cost-push shock in the textbook model.

There are many ways in which this theoretical framework could be extended. First, one could incorporate sticky wages, investment, and other medium-scale features, which are presumably important for the propagation of relative price shocks. Second, one could allow for a more sophisticated role for fiscal policy, and assess the welfare effects of different transfer schemes or taxes and subsidies on different goods. Third, one could allow for differences in steady-state incomes between unconstrained and hand-to-mouth households, to study how this changes optimal policy. Finally, one could consider other forms of limited asset market participation, such as the precautionary savings channel considered by Bilbiie (2018), or a quadratic costs of bond market participation as in Cantore and Freund (2021). In addition to potential extensions to the theoretical framework, there is considerable scope for empirical work to test whether households' consumption responses to inflation match the predictions of the model. As such, this model opens up a broad agenda for future research.

Appendices to Chapter 3

3.A Proofs and Derivations

3.A.1 A Broader Class of Non-Homothetic Preferences

I now show that the form of the intratemporal and intertemporal conditions that emerge under additive preferences also arise under a broader class of non-homothetic utility functions. Suppose that households have a utility function of one of the form:

$$U_t = E_t \sum_{h \geq 0} \beta^h \left(u(C_1, C_2, ..., C_I) - V(N_t) \right)$$

Where $u(C_1, C_2, ..., C_I)$ has one of the two following forms:

$$u(C_1, C_2, ..., C_I) = h\left(\sum_{i \in I} f_i(C_i) \right)$$
$$u(C_1, C_2, ..., C_I) = h\left(\prod_{i \in I} f_i(C_i) \right)$$

I impose the following technical restrictions.

- $\forall i \in I$: $f_i(\cdot)$ is twice continuously differentiable, strictly increasing, and strictly concave
- h is twice continuously differentiable, strictly increasing, and weakly concave

Now note that the first-order conditions in the first case are:

$$\forall i, j \in I : f_i'(C_{i,t}) \frac{1}{\tilde{P}_{i,t}} h'\left(\sum_{i \in I} f_i(C_{i,t}) \right) = f_j'(C_{j,t}) \frac{1}{\tilde{P}_{j,t}} h'\left(\sum_{i \in I} f_i(C_{i,t}) \right)$$

In the second case they are given by:

$$\forall i,j \in I : \frac{f_i'(C_{i,t})}{f_i(C_{i,t})} \prod_{i \in I} f_i(C_{i,t}) \frac{1}{\tilde{P}_{i,t}} h' \left(\prod_{i \in I} f_i(C_{i,t}) \right) = \frac{f_j'(C_{j,t})}{f_j(C_{j,t})} \prod_{i \in I} f_i(C_{i,t}) \frac{1}{\tilde{P}_{j,t}} h' \left(\prod_{i \in I} f_i(C_{i,t}) \right)$$

In either case, the first-order condition for goods i and j can be written as:

$$g_i(C_{i,t}) \frac{1}{\tilde{P}_{i,t}} = g_j(C_{j,t}) \frac{1}{\tilde{P}_{j,t}}$$

With $g_i(\cdot)$ given in the first case by:

$$g_i(C_{i,t}) = f_i'(C_{i,t})$$

And in the second case by:

$$g_i(C_{i,t}) = \frac{f_i'(C_{i,t})}{f_i(C_{i,t})}$$

Let's consider the steady-state. As described in the main text, in steady-state all relative prices are equal to unity. As such:

$$\begin{aligned} \forall i,j \in I : g_i(\bar{C}_i) &= g_j(\bar{C}_j) \\ \Rightarrow \bar{C}_i &= g_i^{-1}(g_j(\bar{C}_j)) \end{aligned}$$

Note that the inverse exists in either case, because in either case g_i is strictly decreasing. Then note that:

$$\begin{aligned} \sum_{i \in I} \bar{C}_i - \bar{C} &= 0 \\ \Rightarrow \sum_{i \in I} g_i^{-1}(g_j(\bar{C}_j)) - \bar{C} &= 0 \end{aligned}$$

One can then obtain the derivative of $\bar{C}_j$ with respect to $\bar{C}$ as:

$$\frac{d\bar{C}_j}{d\bar{C}} = \frac{1}{\sum_{i \in I} (g_i^{-1})'(g_j(\bar{C}_j)) g_j'(\bar{C}_j)}$$

By the inverse function theorem, it holds that:

$$(g_i^{-1})'(g_j(\bar{C}_j)) = (g_i'(\bar{C}_i))^{-1}$$

To apply the inverse function theorem, I rely on g_i being continuously differentiable,

which holds provided that f_i is twice continuously differentiable.

As such the weight of good j in the marginal consumption basket is:

$$\frac{d\bar{C}_j}{d\bar{C}} = \frac{(g_j'(\bar{C}_j))^{-1}}{\sum_{i\in I}(g_i'(\bar{C}_i))^{-1}}$$

Denote this marginal weight as ψ_j^M.

Now let us return to the first-order conditions. The intratemporal condition in the additive case is given by:

$$g_i(C_{i,t})h'\left(\sum_{i\in I} f_i(C_{i,t})\right) = V'(N_t)\frac{\tilde{P}_{i,t}}{\tilde{W}_t}$$

Log-linearizing this expression gives:

$$\frac{g_i'(\bar{C}_i)}{g_i(\bar{C}_i)}\bar{C}_i c_{i,t} + \sum_{i\in I}\frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_i(\bar{C}_i)\bar{C}_i c_{i,t} = \frac{V''(\bar{N})\bar{N}}{V'(\bar{N})}n_t + \tilde{p}_{i,t} - \tilde{w}_t$$

Using the fact that $g_i(\bar{C}_i) = g_j(\bar{C}_j)$, one can express this as:

$$\frac{g_i'(\bar{C}_i)}{g_j(\bar{C}_j)}\bar{C}_i c_{i,t} + \frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_j(\bar{C}_j)\sum_{i\in I}\bar{C}_i c_{i,t} = \frac{V''(\bar{N})\bar{N}}{V'(\bar{N})}n_t + \tilde{p}_{i,t} - \tilde{w}_t$$

$$\frac{g_i'(\bar{C}_i)}{g_j(\bar{C}_j)}\bar{C}_i c_{i,t} + \frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_j(\bar{C}_j)c_t = \frac{V''(\bar{N})\bar{N}}{V'(\bar{N})}n_t + \tilde{p}_{i,t} - \tilde{w}_t$$

Then, summing over each equation $i \in I$ with weights ψ_i^M, one obtains:

$$\left(\frac{1}{\sum_{i\in I}(g_i'(\bar{C}_i))^{-1}}\frac{1}{g_j(\bar{C}_j)} + \frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_j(\bar{C}_j)\right)c_t = \frac{V''(\bar{N})\bar{N}}{V'(\bar{N})}n_t + \tilde{p}_t^M - \tilde{w}_t$$

Note that given the assumptions on f and h outlined above, the coefficient on c_t is strictly negative. This gives the intratemporal condition in the same form as with additive preferences.

In the second form of the utility function (the multiplicative case), one has:

$$g_i(C_{i,t})\prod_{i\in I} f_i(C_{i,t})h'\left(\prod_{i\in I} f_i(C_{i,t})\right) = V'(N_t)\frac{\tilde{P}_{i,t}}{\tilde{W}_t}$$

Log-linearizing gives:

$$\frac{g_i'(\bar{C}_i)}{g_i(\bar{C}_i)}\bar{C}_i c_{i,t} + \sum_{i\in I}\frac{f_i'(\bar{C}_i)}{f_i(\bar{C}_i)}\bar{C}_i c_{i,t} + \sum_{i\in I}\frac{f_i'(\bar{C}_i)}{f_i(\bar{C}_i)}\frac{\prod_{i\in I} f_i(\bar{C}_i)h''(\prod_{i\in I} f_i(\bar{C}_i))}{h'(\prod_{i\in I} f_i(\bar{C}_i))}\bar{C}_i c_{i,t} = \frac{V''(\bar{N})\bar{N}}{V'(\bar{N})}n_t + \tilde{p}_{i,t} - \tilde{w}_t$$

Then using the fact that $g_i(\bar{C}_i) = g_j(\bar{C}_j)$:

$$\frac{g_i'(\bar{C}_i)}{g_j(\bar{C}_j)}\bar{C}_i c_{i,t} + g_j(\bar{C}_j)\left(1 + \frac{\prod_{i\in I} f_i(\bar{C}_i)h''(\prod_{i\in I} f_i(\bar{C}_i))}{h'(\prod_{i\in I} f_i(\bar{C}_i))}\right)c_t = \frac{V''(\bar{N})\bar{N}}{V'(\bar{N})}n_t + \tilde{p}_{i,t} - \tilde{w}_t$$

Then, sum across each of the $i \in I$ equations with weights ψ_i^M, to give:

$$\left(\frac{1}{\sum_{i\in I}(g_i'(\bar{C}_i))^{-1}}\frac{1}{g_j(\bar{C}_j)} + g_j(\bar{C}_j)\left(1 + \frac{\prod_{i\in I} f_i(\bar{C}_i)h''(\prod_{i\in I} f_i(\bar{C}_i))}{h'(\prod_{i\in I} f_i(\bar{C}_i))}\right)\right)c_t = \frac{V''(\bar{N})\bar{N}}{V'(\bar{N})}n_t + \tilde{p}_t^M - \tilde{w}_t$$

Now let us turn to the intertemporal conditions. Consider the additive formulation for preferences. The intertemporal condition for each good i is:

$$f_i'(C_{i,t})\frac{1}{P_{i,t}}h'\left(\sum_{i\in I} f_i(C_{i,t})\right) = \beta(1+i_t)f_i'(C_{i,t+1})\frac{1}{P_{i,t+1}}h'\left(\sum_{i\in I} f_i(C_{i,t+1})\right)$$

Log-linearizing yields:

$$\frac{g_i'(\bar{C}_{i,t})}{g_i(\bar{C}_i)}\bar{C}_i c_{i,t} + \sum_{i\in I}\frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_i(\bar{C}_i)\bar{C}_i c_{i,t} = i_t - \pi_{i,t+1} + \frac{g_i'(\bar{C}_{i,t})}{g_i(\bar{C}_i)}\bar{C}_i c_{i,t+1} + \sum_{i\in I}\frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_i(\bar{C}_i)\bar{C}_i c_{i,t+1}$$

Once again noting that $g_i(\bar{C}_i) = g_j(\bar{C}_j)$:

$$\frac{g_i'(\bar{C}_{i,t})}{g_j(\bar{C}_i)}\bar{C}_i c_{i,t} + \frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_i(\bar{C}_i)c_t = i_t - \pi_{i,t+1} + \frac{g_i'(\bar{C}_{i,t})}{g_j(\bar{C}_i)}\bar{C}_i c_{i,t+1} + \frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_i(\bar{C}_i)c_{t+1}$$

Then summing across equations with weights ψ_i^M:

$$-\bar{\gamma}c_t = i_t - \pi_{t+1}^M - \bar{\gamma}c_{t+1}$$

$$\bar{\gamma} = -\left(\frac{1}{\sum_{i\in I}(g_i'(\bar{C}_i))^{-1}}\frac{1}{g_j(\bar{C}_i)} + \frac{h''\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}{h'\left(\sum_{i\in I} f_i(\bar{C}_i)\right)}g_i(\bar{C}_i)\right)$$

This gives the Euler equation as in the additive preferences case.

Now let us turn to the multiplicative case. One has the intertemporal condition:

$$\frac{f_i'(C_{i,t})}{f_i(C_{i,t})}\prod_{i\in I} f_i(C_{i,t})\frac{1}{P_{i,t}}h'\left(\prod_{i\in I} f_i(C_{i,t})\right) = \beta(1+i_t)\frac{f_i'(C_{i,t+1})}{f_i(C_{i,t+1})}\prod_{i\in I} f_i(C_{i,t+1})\frac{1}{P_{i,t+1}}h'\left(\prod_{i\in I} f_i(C_{i,t+1})\right)$$

Following the same log-linearization strategy as above yields:

$$\begin{aligned}-\bar{\gamma}c_t &= i_t - \pi^M_{t+1} - \bar{\gamma}c_{t+1}\\ \bar{\gamma} &= -\left(\frac{1}{\sum_{i\in I}(g_i'(\bar{C}_i))^{-1}}\frac{1}{g_j(\bar{C}_j)} + g_j(\bar{C}_j)\left(1+\frac{\prod_{i\in I} f_i(\bar{C}_i)h''(\prod_{i\in I} f_i(\bar{C}_i))}{h'(\prod_{i\in I} f_i(\bar{C}_i))}\right)\right)\end{aligned}$$

3.A.2 Law of Motion for Relative Prices

I now derive the law of motion for relative prices. I begin with the equation for the relative inflation rate of the marginal basket:

$$\tilde{\pi}^M_t = \beta E_t\tilde{\pi}^M_{t+1} - \lambda\tilde{p}^M_t + \tilde{\eta}^M_t$$

Using lag notation, one can re-write this equation as:

$$\tilde{p}^M_t(1-L) = \beta\tilde{p}^M_t(L^{-1}-1) - \lambda\tilde{p}^M_t + \tilde{\eta}^M_t$$

Re-arranging:

$$\begin{aligned}(1+\beta+\lambda-L-\beta L^{-1})\tilde{p}^M_t &= \tilde{\eta}^M_t\\ \Rightarrow\left(1-\frac{1+\beta+\lambda}{\beta}L+\frac{1}{\beta}L^2\right)\tilde{p}^M_t &= -\frac{1}{\beta}L\tilde{\eta}^M_t\end{aligned}$$

The left-hand side can be factorized to give:

$$(1-\zeta_1 L)(1-\zeta_2 L)\tilde{p}^M_t = -\frac{1}{\beta}L\tilde{\eta}^M_t$$

Here, ζ_1 and ζ_2 must satisfy:

$$\begin{aligned}\zeta_1+\zeta_2 &= \frac{1+\beta+\lambda}{\beta}\\ \zeta_1\zeta_2 &= \frac{1}{\beta}\end{aligned}$$

The second equation implies:

$$\zeta_2 = \frac{1}{\beta\zeta_1}$$

Substituting into the first equation and re-arranging gives:

$$\begin{aligned} 0 &= \beta\zeta_1^2 - (1+\beta+\lambda)\zeta_1 + 1 \\ \Rightarrow \zeta_1 &= \frac{1+\beta+\lambda \pm \sqrt{(1+\beta+\lambda)^2 - 4\beta}}{2\beta} \end{aligned}$$

One can verify that the $+$ root is greater than unity, and the $-$ root is in the interval $(0,1)$, given $\beta \in (0,1)$ and $\lambda > 0$. Denote by ζ the stable root, that is:

$$\zeta = \frac{1+\beta+\lambda - \sqrt{(1+\beta+\lambda)^2 - 4\beta}}{2\beta}$$

The unstable root is then $1/(\beta\zeta)$ given the condition above. Hence the lag polynomial becomes:

$$(1-\zeta L)(1-(\beta\zeta)^{-1}L)\tilde{p}_t^M = -\frac{1}{\beta}L\tilde{\eta}_t^M$$

One can then proceed in the usual fashion:

$$\begin{aligned} (1-\zeta L)\tilde{p}_t^M &= -\frac{1}{\beta}\frac{L}{1-(\beta\zeta)^{-1}L}\tilde{\eta}_t^M \\ (1-\zeta L)\tilde{p}_t^M &= \frac{\zeta}{1-\beta\zeta L^{-1}}\tilde{\eta}_t^M \\ (1-\zeta L)\tilde{p}_t^M &= \zeta\sum_{h\geq 0}(\beta\zeta)^h E_t\tilde{\eta}_{t+h}^M \\ (1-\zeta L)\tilde{p}_t^M &= \zeta\sum_{h\geq 0}(\beta\zeta\rho)^h\tilde{\eta}_t^M \\ \tilde{p}_t^M &= \zeta\tilde{p}_{t-1}^M + \frac{\zeta}{1-\beta\zeta\rho}\tilde{\eta}_t^M \end{aligned}$$

3.A.3 Analytical Solution with a Taylor Rule

I now solve the model with a simple Taylor rule using the method of undetermined coefficients. I guess and verify a solution of the following form:

$$\begin{aligned} \pi_t &= A_1\tilde{p}_t^M + A_2\tilde{\eta}_t^M + A_3\eta_t \\ y_t &= B_1\tilde{p}_t^M + B_2\tilde{\eta}_t^M + B_3\eta_t \end{aligned}$$

A few points are worth mentioning before starting the derivations:

- The solution for A_3 and B_3 will be the same as in a standard new Keynesian model. As such, I will focus on deriving A_1, A_2, B_1, and B_2.
- In the case where $\rho = 0$, so that $\tilde{p}_t^M$ is an AR(1) process, then one can omit $\tilde{\eta}_t^M$. In $\rho \neq 0$, then because $\tilde{p}_t^M$ is not an AR(1) and there are hump-shaped dynamics, the solution will depend on both $\tilde{p}_t^M$ and $\tilde{\eta}_t^M$.
- One could also solve the model with $\tilde{p}_{t-1}^M$ as a state-variable instead of $\tilde{p}_t^M$. That still works, and one can still derive an analytical solution. However, the algebra turns out to be somewhat more cumbersome. As such, I prefer to work with $\tilde{p}_t^M$ as a state variable.
- Using my solution, though, one can find the appropriate coefficients on $\tilde{p}_{t-1}^M$ and $\tilde{\eta}_t^M$ by using the law of motion of $\tilde{p}_t^M$. I show how to do this below.

The model is:

$$\begin{aligned}
\pi_t &= \beta E_t \pi_{t+1} + \kappa y_t + \lambda \tilde{p}_t^M + \eta_t \\
y_t &= E_t y_{t+1} + \alpha(E_t \tilde{p}_{t+1}^M - \tilde{p}_t^M) - \sigma(i_t - E_t \pi_{t+1}) \\
i_t &= r_\pi \pi_t \\
\tilde{p}_t^M &= \zeta \tilde{p}_{t-1}^M + \frac{\zeta}{1 - \beta\zeta\rho} \tilde{\eta}_t^M
\end{aligned}$$

And the shocks η_t and $\tilde{\eta}_t^M$ are, as before, AR(1) processes with persistence ρ. Using the guess above, the Phillips curve becomes:

$$\pi_t = \beta(A_1 E_t \tilde{p}_{t+1}^M + A_2 E_t \tilde{\eta}_{t+1}^M + \beta A_3 E_t \eta_{t+1}) + \kappa(B_1 \tilde{p}_t^M + B_2 \tilde{\eta}_t^M + B_3 \eta_t) + \lambda \tilde{p}_t^M + \eta_t$$

Using the laws of motion for $\tilde{p}_t^M$, $\tilde{\eta}_t^M$, and η_t, this simplifies to:

$$\pi_t = (\beta\zeta A_1 + \kappa B_1 + \lambda)\, \tilde{p}_t^M + \left(\beta \frac{\zeta\rho}{1 - \beta\zeta\rho} A_1 + \beta\rho A_2 + \kappa B_2\right) \tilde{\eta}_t^M + (\beta\rho A_3 + \kappa B_3 + 1)\eta_t$$

Substitute the policy rule into the IS curve, and the substitute in the guessed solution. The IS curve becomes:

$$\begin{aligned}
y_t = (B_1 E_t \tilde{p}_{t+1}^M + B_2 E_t \tilde{\eta}_{t+1}^M + B_3 E_t \eta_{t+1}) + \alpha(E_t \tilde{p}_{t+1}^M - \tilde{p}_t^M) \\
- \sigma r_\pi (A_1 \tilde{p}_t^M + A_2 \tilde{\eta}_t^M + A_3 \eta_t) + \sigma(A_1 E_t \tilde{p}_{t+1}^M + E_t A_2 \tilde{\eta}_{t+1}^M + A_3 E_t \eta_{t+1})
\end{aligned}$$

Using the laws of motion for $\tilde{p}_t^M$, $\tilde{\eta}_t^M$, and η_t and re-arranging yields:

$$\begin{aligned} y_t = &(\zeta B_1 - \alpha(1-\zeta) - \sigma(r_\pi - \zeta)A_1)\tilde{p}_t^M \\ &+ \left(\frac{\zeta\rho}{1-\beta\zeta\rho}B_1 + \rho B_2 + \alpha\frac{\zeta\rho}{1-\beta\zeta\rho} - \sigma(r_\pi - \rho)A_2 + \sigma\frac{\zeta\rho}{1-\beta\zeta\rho}A_1\right)\tilde{\eta}_t^M \\ &+ (\rho B_3 - \sigma(r_\pi - \rho)A_3)\,\eta_t \end{aligned}$$

As such, the guess for the form of the solution is verified, and the coefficients satisfy the following system of equations:

$$\begin{aligned} (1-\beta\zeta)A_1 &= \kappa B_1 + \lambda \\ (1-\beta\rho)A_2 &= \beta\frac{\zeta\rho}{1-\beta\zeta\rho}A_1 + \kappa B_2 \\ (1-\beta\rho)A_3 &= \kappa B_3 + 1 \\ (1-\zeta)B_1 &= -\alpha(1-\zeta) - \sigma(r_\pi - \zeta)A_1 \\ (1-\rho)B_2 &= \frac{\zeta\rho}{1-\beta\zeta\rho}B_1 + \alpha\frac{\zeta\rho}{1-\beta\zeta\rho} - \sigma(r_\pi - \rho)A_2 + \sigma\frac{\zeta\rho}{1-\beta\zeta\rho}A_1 \\ (1-\rho)B_3 &= -\sigma(r_\pi - \rho)A_3 \end{aligned}$$

I begin by solving for A_1 and B_1. Substituting the expression for A_1 into the equation for B_1 yields:

$$\begin{aligned} (1-\zeta)B_1 &= -\alpha(1-\zeta) - \sigma(r_\pi - \zeta)\frac{\kappa B_1 + \lambda}{1-\beta\zeta} \\ \Rightarrow B_1 &= -\frac{\alpha(1-\zeta)(1-\beta\zeta) + \sigma\lambda(r_\pi - \zeta)}{(1-\zeta)(1-\beta\zeta) + \kappa\sigma(r_\pi - \zeta)} \end{aligned}$$

The substituting this expression back into the equation for A_1 gives:

$$\begin{aligned} (1-\beta\zeta)A_1 &= \lambda - \kappa\frac{\alpha(1-\zeta)(1-\beta\zeta) + \sigma\lambda(r_\pi - \zeta)}{(1-\zeta)(1-\beta\zeta) + \kappa\sigma(r_\pi - \zeta)} \\ \Rightarrow A_1 &= \frac{(\lambda - \kappa\alpha)(1-\zeta)}{(1-\beta\zeta)(1-\zeta) + \kappa\sigma(r_\pi - \zeta)} \end{aligned}$$

I now solve for A_2 and B_2. Begin by substituting the equation for A_2 into the equation for B_2.

$$(1-\rho)B_2 = \frac{\zeta\rho}{1-\beta\zeta\rho}B_1 + \alpha\frac{\zeta\rho}{1-\beta\zeta\rho} - \frac{\sigma(r_\pi - \rho)(\beta\frac{\zeta\rho}{1-\beta\zeta\rho}A_1 + \kappa B_2)}{1-\beta\rho} + \sigma\frac{\zeta\rho}{1-\beta\zeta\rho}A_1$$

Re-arranging gives:

$$\frac{1-\beta\zeta\rho}{\zeta\rho}((1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho))B_2 \\ = (1-\beta\rho)B_1+(1-\beta\rho)\alpha-\sigma\left((r_\pi-\rho)\beta-(1-\beta\rho)\right)A_1$$

Then use the expression above for B_1 in terms of A_1, to give:

$$\frac{1-\beta\zeta\rho}{\zeta\rho}((1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho))B_2 \\ = (1-\beta\rho)\frac{-\alpha(1-\zeta)-\sigma(r_\pi-\zeta)A_1}{1-\zeta}+(1-\beta\rho)\alpha-\sigma\left((r_\pi-\rho)\beta-(1-\beta\rho)\right)A_1$$

Notice that the terms in α cancel. Multiplying through by $1-\zeta$ and re-arranging gives:

$$(1-\zeta)\frac{1-\beta\zeta\rho}{\zeta\rho}((1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho))B_2 \\ = \left(-(1-\beta\rho)\sigma(r_\pi-\zeta)-(1-\zeta)\sigma(r_\pi-\rho)\beta+(1-\zeta)\sigma(1-\beta\rho)\right)A_1$$

Simplifying the term bracketed term on the right-hand side yields:

$$(1-\zeta)\frac{1-\beta\zeta\rho}{\zeta\rho}((1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho))B_2 \\ = -\sigma\left((1-\beta\rho)(r_\pi-1)+\beta(1-\zeta)(r_\pi-\rho)\right)A_1$$

Finally, substitute in for A_1 and divide by the coefficient on the left-hand side to give:

$$B_2 = -\frac{\zeta\rho}{1-\beta\zeta\rho}\frac{\sigma\left((1-\beta\rho)(r_\pi-1)+\beta(1-\zeta)(r_\pi-\rho)\right)}{(1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho)}\frac{A_1}{1-\zeta} \\ = -\frac{\zeta\rho}{1-\beta\zeta\rho}\frac{\sigma\left((1-\beta\rho)(r_\pi-1)+\beta(1-\zeta)(r_\pi-\rho)\right)}{(1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho)}\frac{(\lambda-\kappa\alpha)}{(1-\beta\zeta)(1-\zeta)+\kappa\sigma(r_\pi-\zeta)}$$

I now solve for A_2. Recall that:

$$(1-\beta\rho)A_2 = \beta\frac{\zeta\rho}{1-\beta\zeta\rho}A_1+\kappa B_2 \\ = \beta\frac{\zeta\rho}{1-\beta\zeta\rho}A_1-\frac{\zeta\rho}{1-\beta\zeta\rho}\frac{\kappa\sigma\left((1-\beta\rho)(r_\pi-1)+\beta(1-\zeta)(r_\pi-\rho)\right)}{(1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho)}\frac{A_1}{1-\zeta}$$

Simplifying this expression yields:

$$(1-\beta\rho)A_2 = \frac{\zeta\rho}{1-\beta\zeta\rho}\left(\frac{\beta(1-\zeta)(1-\beta\rho)(1-\rho)-\kappa\sigma(1-\beta\rho)(r_\pi-1)}{(1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho)}\right)\frac{A_1}{1-\zeta}$$

Which then gives our final expression for A_2:

$$\begin{aligned} A_2 &= \frac{\zeta\rho}{1-\beta\zeta\rho}\frac{\beta(1-\zeta)(1-\rho)-\kappa\sigma(r_\pi-1)}{(1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho)}\frac{A_1}{1-\zeta} \\ &= \frac{\zeta\rho}{1-\beta\zeta\rho}\frac{\beta(1-\zeta)(1-\rho)-\kappa\sigma(r_\pi-1)}{(1-\beta\rho)(1-\rho)+\kappa\sigma(r_\pi-\rho)}\frac{(\lambda-\kappa\alpha)}{(1-\beta\zeta)(1-\zeta)+\kappa\sigma(r_\pi-\zeta)} \end{aligned}$$

The coefficients A_3 and B_3 are standard. These can be solved simultaneously to give:

$$\begin{aligned} A_3 &= \frac{1-\rho}{(1-\rho)(1-\beta\rho)+\kappa\sigma(r_\pi-\rho)} \\ B_3 &= -\frac{\sigma(r_\pi-\rho)}{(1-\rho)(1-\beta\rho)+\kappa\sigma(r_\pi-\rho)} \end{aligned}$$

If one wants a solution in terms of $\tilde{p}^M_{t-1}$ instead, this is straightforward to obtain in the following way:

$$\begin{aligned} \pi_t &= A_1\tilde{p}^M_t + A_2\tilde{\eta}^M_t + A_3\eta_t \\ &= A_1\left(\zeta\tilde{p}^M_{t-1} + \frac{\zeta}{1-\beta\zeta\rho}\tilde{\eta}^M_t\right) + A_2\tilde{\eta}^M_t + A_3\eta_t \\ &= A_1\zeta\tilde{p}^M_{t-1} + \left(A_1\frac{\zeta}{1-\beta\zeta\rho} + A_2\right)\tilde{\eta}^M_t + A_3\eta_t \end{aligned}$$

The same process can be followed for the solution for y_t.

3.A.4 Welfare Function

I now derive the welfare function. I do this in the two agent version of the model. The representative agent version can then be calculated by setting $\xi = 0$. To a second-order approximation, welfare in time t can be written as:

$$\begin{aligned} U_t - \bar{U} \approx \xi\sum_{i\in I} U_{C_i}\bar{C}^H_i\left(c^H_{i,t} - \frac{\gamma_i-1}{2}(c^H_{i,t})^2\right) + \xi U_N\bar{N}^H\left(n^H_t + \frac{1+\phi}{2}(n^H_t)^2\right) \\ + (1-\xi)\sum_{i\in I} U_{C_i}\bar{C}^U_i\left(c^U_{i,t} - \frac{\gamma_i-1}{2}(c^U_{i,t})^2\right) + (1-\xi)U_N\bar{N}^U\left(n^U_t + \frac{1+\phi}{2}(n^U_t)^2\right) \end{aligned}$$

First, note that steady-state consumption of each good and labour supply is the same across the two types, because they have the same steady-state income. Then note that by the intratemporal condition, and given that $(\bar{W}/\bar{P}_i) = 1$ for all i, one has:

$$\forall i \in I : U_{C_i} = -U_N$$

Moreover, note that: $\bar{C} = \bar{N}$. As such, one has:

$$\frac{U_t - \bar{U}}{U_{C_i}\bar{C}} \approx \xi \sum_{i \in I} \frac{\bar{C}_i}{\bar{C}} \left(c_{i,t}^H - \frac{\gamma_i - 1}{2}(c_{i,t}^H)^2 \right) - \xi \left(n_t^H + \frac{1+\phi}{2}(n_t^H)^2 \right) + (1-\xi) \sum_{i \in I} \frac{\bar{C}_i}{\bar{C}} \left(c_{i,t}^U - \frac{\gamma_i - 1}{2}(c_{i,t}^U)^2 \right) - (1-\xi) \left(n_t^U + \frac{1+\phi}{2}(n_t^U)^2 \right)$$

Next, let us note that for $s \in \{H, U\}$:

$$C_t^s = \sum_{i \in I} C_{i,t}^s$$
$$\Rightarrow \frac{C_t^s - \bar{C}^s}{\bar{C}^s} = \sum_{i \in I} \frac{\bar{C}_i^s}{\bar{C}^s} \frac{C_{i,t}^s - \bar{C}_i^s}{\bar{C}_i^s}$$

Taking a second-order approximation of this expression:

$$c_t^s + \frac{1}{2}(c_t^s)^2 = \sum_{i \in I} \frac{\bar{C}_i}{\bar{C}} \left(c_t^s + \frac{1}{2}(c_t^s)^2 \right)$$

Next, note that by the intratemporal condition, the following holds:

$$c_{i,t}^s = \frac{1}{\gamma_i}(\gamma_j c_{j,t}^s - p_{i,t} + p_{j,t})$$

Summing over equations with weights $\frac{\bar{C}_j}{\bar{C}} \frac{\bar{\gamma}}{\gamma_j}$ yields:

$$\begin{aligned} c_{i,t}^s &= \frac{1}{\gamma_i} \left(\bar{\gamma} \sum_{j \in I} \frac{\bar{C}_j}{\bar{C}} c_{j,t}^s - p_{i,t} + p_t^M \right) \\ &= \frac{1}{\gamma_i} \left(\bar{\gamma} c_t^s - p_{i,t} + p_t^M \right) + \text{second order terms} \end{aligned}$$

As such:

$$\begin{aligned} \gamma_i \left(c_{i,t}^s \right)^2 &= \frac{1}{\gamma_i} \left(\bar{\gamma} c_t^s - (p_{i,t} - p_t^M) \right)^2 + \text{higher order terms} \\ &= \frac{\bar{\gamma}^2}{\gamma_i}(c_t^s)^2 - 2\frac{\bar{\gamma}}{\gamma_i}(p_{i,t} - p_t^M) + \text{t.i.p.} + \text{higher order terms} \end{aligned}$$

Throughout, t.i.p. stands for "terms invariant to policy". Note that relative prices are invariant to policy under the assumption of constant returns to scale, as discussed in the main text, which is why the squared relative price term is invariant to policy. Now sum

across each equation i with consumption weights. Then:

$$\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\gamma_i(c^s_{i,t})^2=\bar{\gamma}^2(c^s_t)^2\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\frac{1}{\gamma_i}-2\bar{\gamma}\left(\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\frac{1}{\gamma_i}p_{i,t}-p^M_t\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\frac{1}{\gamma_i}\right)$$
$$+\text{ t.i.p.}+\text{higher order terms}$$

Note that the bracketed term on the right-hand-side is equal to zero by the definition of p^M_t. As such:

$$\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\gamma_i(c^s_{i,t})^2=\bar{\gamma}(c^s_t)^2+\text{t.i.p.}+\text{higher order terms}$$

Combining with the results above, one has:

$$\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\left(c^s_{i,t}-\frac{\gamma_i-1}{2}(c^s_{i,t})^2\right)=c^s_t-\frac{\bar{\gamma}-1}{2}(c^s_t)^2+\text{higher order terms}+\text{t.i.p.}$$

The intuition here is important. Welfare is lower when dispersion in $c^s_{i,t}$ across i is higher. However, by the intratemporal condition, such dispersion is driven by relative price changes. Since such relative prices changes are invariant to policy in the constant returns to scale model, the dispersion in consumption across goods is also invariant to policy.

Next, note that under Rotemberg pricing, and given constant returns to scale such that $Y_{i,t}=N_{i,t}$, the following relationship must hold in each sector:

$$C_{i,t}=N_{i,t}\left(1-\frac{\varphi}{2}(\Pi_{i,t}-1)^2\right)$$

Summing across sectors:

$$C_t=N_t-\sum_{i\in I}Y_{i,t}\frac{\varphi}{2}(\Pi_{i,t}-1)^2$$
$$\xi C^H_t+(1-\xi)C^U_t=\xi N^H_t+(1-\xi)N^U_t-\frac{\varphi}{2}\sum_{i\in I}Y_{i,t}(\Pi_{i,t}-1)^2$$

Then, acknowledging that $\bar{C}^H=\bar{C}^U=\bar{C}$, $\bar{N}^H=\bar{N}^U=\bar{N}$, and that $\bar{C}=\bar{Y}=\bar{N}$, one has:

$$\xi\frac{C^H_t-\bar{C}^H}{\bar{C}^H}+(1-\xi)\frac{C^U_t-\bar{C}^U}{\bar{C}^U}=\xi\frac{N^H_t-\bar{N}^H}{\bar{N}^H}+(1-\xi)\frac{N^U_t-\bar{N}^U}{\bar{N}^U}-\frac{\varphi}{2}\sum_{i\in I}\frac{Y_{i,t}}{\bar{Y}_i}\frac{\bar{Y}_i}{\bar{Y}}\left(\frac{\Pi_{i,t}-\bar{\Pi}_i}{\bar{\Pi}_i}\right)^2$$

Note that the last term follows from the fact that I consider a zero-inflation steady-state,

and hence $\bar{\Pi}_i = 1$. This can be rearranged to give:

$$\xi \frac{C_t^H - \bar{C}^H}{\bar{C}^H} + (1-\xi)\frac{C_t^U - \bar{C}^U}{\bar{C}^U} = \xi \frac{N_t^H - \bar{N}^H}{\bar{N}^H} + (1-\xi)\frac{N_t^U - \bar{N}^U}{\bar{N}^U} - \frac{\varphi}{2}\sum_{i\in I}\frac{Y_{i,t} - \bar{Y}_i}{\bar{Y}_i}\frac{\bar{Y}_i}{\bar{Y}}\left(\frac{\Pi_{i,t} - \bar{\Pi}_i}{\bar{\Pi}_i}\right)^2 - \frac{\varphi}{2}\sum_{i\in I}\frac{\bar{Y}_i}{\bar{Y}}\left(\frac{\Pi_{i,t} - \bar{\Pi}_i}{\bar{\Pi}_i}\right)^2$$

Then, replacing each fraction with a second-order expansion, one obtains:

$$\xi\left(c_t^H + \frac{1}{2}(c_t^H)^2\right) + (1-\xi)\left(c_t^U + \frac{1}{2}(c_t^U)^2\right) = \xi\left(n_t^H + \frac{1}{2}(n_t^H)^2\right) + (1-\xi)\left(n_t^U + \frac{1}{2}(n_t^U)^2\right) - \frac{\varphi}{2}\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\pi_{i,t}^2$$

Here, I also used the fact that $\bar{C}_i = \bar{Y}_i$ and $\bar{C} = \bar{Y}$. Combining the above results, one obtains:

$$\frac{U_t - \bar{U}}{U_{C_i}\bar{C}} \approx -\xi\frac{\bar{\gamma}}{2}(c_t^H)^2 - (1-\xi)\frac{\bar{\gamma}}{2}(c_t^U)^2 - \xi\frac{\phi}{2}(n_t^H)^2 - (1-\xi)\frac{\phi}{2}(n_t^U)^2 - \frac{\varphi}{2}\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\pi_{i,t}^2$$

I begin by dealing with the last term. I will use the fact that relative prices are exogenous.

$$\begin{aligned}\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\pi_{i,t}^2 &= \sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}(\pi_t + \tilde{\pi}_{i,t})^2 \\ &= \sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}(\pi_t^2 + 2\pi_t\tilde{\pi}_{i,t} + \tilde{\pi}_{i,t}^2) \\ &= \pi_t^2 + 2\pi_t\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\tilde{\pi}_{i,t} + \text{t.i.p.}\end{aligned}$$

Here, I again used the fact that relative inflation rates are exogenous to drop the term in $\tilde{\pi}_{i,t}^2$. Finally, note that by the definition of the price index p_t, one has:

$$\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}} = \pi_t$$

And hence the second term above cancels. As such:

$$\sum_{i\in I}\frac{\bar{C}_i}{\bar{C}}\pi_{i,t}^2 = \pi_t^2 + \text{t.i.p.}$$

I substitute this into the simplified welfare function above, and then divide by $-1/2$.

This yields:

$$-2\frac{U_t-\bar{U}}{U_{C_i}\bar{C}} \approx \xi\bar{\gamma}(c_t^H)^2+(1-\xi)\bar{\gamma}(c_t^U)^2+\xi\phi(n_t^H)^2+(1-\xi)\phi(n_t^U)^2+\varphi\pi_t^2$$

This excludes higher order terms and terms invariant to policy. I then use the following equations from the log-linearized model:

$$\begin{aligned}
c_t^H &= (1+\phi)y_t+\frac{\phi}{\bar{\gamma}+\phi}\tilde{p}_t^M \\
c_t^U &= \frac{1-\xi(1+\phi)}{1-\xi}y_t-\frac{\xi}{1-\xi}\frac{\phi}{\bar{\gamma}+\phi}\tilde{p}_t^M \\
n_t^H &= -(\bar{\gamma}-1)y_t-\frac{\bar{\gamma}}{\bar{\gamma}+\phi}\tilde{p}_t^M \\
n_t^U &= \frac{1+\xi(\bar{\gamma}-1)}{1-\xi}y_t+\frac{\xi}{1-\xi}\frac{\bar{\gamma}}{\bar{\gamma}+\phi}\tilde{p}_t^M
\end{aligned}$$

Substituting these into the welfare function above and simplifying gives the following welfare loss in time t:

$$-2\frac{U_t-\bar{U}}{U_{C_i}\bar{C}} \approx \varphi\pi_t^2+2\bar{\gamma}\phi\frac{\xi}{1-\xi}y_t\tilde{p}_t^M+(\bar{\gamma}+\phi)\left(1+\bar{\gamma}\phi\frac{\xi}{1-\xi}\right)y_t^2+\text{higher order terms}+\text{t.i.p.}$$

Once again, this excludes higher order terms and terms invariant to policy. Note that terms in $(\tilde{p}_t^M)^2$ are terms invariant to policy because of the exogeneity of relative prices. Minimizing the expected present discounted value of welfare losses is then equivalent to minimizing:

$$\begin{aligned}
L_t &= E_t\frac{1}{2}\sum_{h\geq 0}\beta^h\left(\pi_{t+h}^2+2\Omega_1\tilde{p}_t^M y_t+\Omega_2 y_t^2\right) \\
\Omega_1 &= \frac{1}{\varphi}\bar{\gamma}\phi\frac{\xi}{1-\xi} \\
\Omega_2 &= \frac{1}{\varphi}(\bar{\gamma}+\phi)\left(1+\bar{\gamma}\phi\frac{\xi}{1-\xi}\right)
\end{aligned}$$

Notice that in the representative agent case, $\xi=0$, and hence:

$$\begin{aligned}
\Omega_1 &= 0 \\
\Omega_2 &= \frac{1}{\varphi}(\bar{\gamma}+\phi)
\end{aligned}$$

3.A.5 Optimal Policy under Discretion

I now derive the optimal policy under discretion. Once again, I proceed using the two-agent model, noting that one can recover the optimal policy in the representative agent model simply by adjusting Ω_1, Ω_2, and α.

The central bank faces the following problem:

$$\min \frac{1}{2} E_t \sum_{h \geq 0} \beta^h \left(\pi_{t+h}^2 + 2\Omega_1 \tilde{p}_{t+h}^M y_{t+h} + \Omega_2 y_{t+h}^2 \right)$$

Subject to the Phillips curve:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa y_t + \lambda \tilde{p}_t^M + \eta_t$$

As usual, the central bank takes expected inflation as exogenous, as well as acknowledging that relative prices and the cost-push shock are exogenous. The Lagrangian is:

$$\begin{aligned} \mathcal{L} = E_t \sum_{h \geq 0} \beta^h \Bigg(& \frac{1}{2} \left(\pi_{t+h}^2 + 2\Omega_1 \tilde{p}_{t+h}^M y_{t+h} + \Omega_2 y_{t+h}^2 \right) \\ & - \Gamma_{t+h} \left(\pi_{t+h} - \beta \pi_{t+h+1} - \kappa y_{t+h} - \lambda \tilde{p}_{t+h}^M - \eta_{t+h} \right) \Bigg) \end{aligned}$$

Taking first-order conditions in π_t and y_t gives:

$$\begin{aligned} \pi_t - \Gamma_t &= 0 \\ \Omega_1 \tilde{p}_t^M + \Omega_2 y_t + \Gamma_t \kappa &= 0 \end{aligned}$$

These two expressions are familiar from the standard new Keynesian model, but with the additional term in $\tilde{p}_t^M$. This term emerges only in the two agent model when $\Omega_1 \neq 0$. As discussed in the main text, it describes a concern for consumption equality; when y_t and $\tilde{p}_t^M$ covary positively, this increases consumption heterogeneity and so lowers welfare. Re-arranging the above expression gives:

$$y_t = -\frac{\kappa}{\Omega_2} \pi_t - \frac{\Omega_1}{\Omega_2} \tilde{p}_t^M$$

I then substitute this equation into the Phillips curve, which yields:

$$\pi_t = \beta E_t \pi_{t+1} - \frac{\kappa^2}{\Omega_2} \pi_t + \left(\lambda - \kappa \frac{\Omega_1}{\Omega_2} \right) \tilde{p}_t^M + \eta_t$$

To lighten the notation, define:

$$q = \frac{\Omega_2}{\Omega_2 + \kappa^2}$$
$$z = \lambda - \kappa\frac{\Omega_1}{\Omega_2}$$

Note that $q \in (0,1)$. The expression for π_t then becomes:

$$\begin{aligned}
\pi_t &= \beta q E_t \pi_{t+1} + qz\tilde{p}_t^M + q\eta_t \\
\Rightarrow \pi_t &= E_t \sum_{h\geq 0} (\beta q)^h \left(qz\tilde{p}_{t+h}^M + q\eta_{t+h}\right) \\
&= qzE_t \sum_{h\geq 0} (\beta q)^h \tilde{p}_{t+h}^M + \frac{q}{1-\beta q\rho}\eta_t
\end{aligned}$$

Focussing first on the summation term in $\tilde{p}_{t+h}^M$, one has:

$$\begin{aligned}
E_t \sum_{h\geq 0} (\beta q)^h \tilde{p}_{t+h}^M &= E_t \sum_{h\geq 0} (\beta q)^h \left(\zeta^h \tilde{p}_t^M + \frac{\zeta}{1-\beta\zeta\rho}\left(\zeta^{h-1}\tilde{\eta}_{t+1}^M + ... + \tilde{\eta}_{t+h}^M\right)\right) \\
&= \sum_{h\geq 0} (\beta q)^h \left(\zeta^h \tilde{p}_t^M + \frac{\zeta\rho}{1-\beta\zeta\rho}\left(\zeta^{h-1} + ... + \rho^{h-1}\right)\tilde{\eta}_t^M\right) \\
&= \sum_{h\geq 0} (\beta q)^h \left(\zeta^h \tilde{p}_t^M + \frac{\zeta^h\rho}{1-\beta\zeta\rho}\left(1 + ... + \left(\frac{\rho}{\zeta}\right)^{h-1}\right)\tilde{\eta}_t^M\right) \\
&= \sum_{h\geq 0} (\beta q)^h \left(\zeta^h \tilde{p}_t^M + \frac{\zeta^h\rho}{1-\beta\zeta\rho}\frac{1-\left(\frac{\rho}{\zeta}\right)^h}{1-\frac{\rho}{\zeta}}\tilde{\eta}_t^M\right) \\
&= \sum_{h\geq 0} (\beta q)^h \left(\zeta^h \tilde{p}_t^M + \frac{\zeta\rho}{1-\beta\zeta\rho}\frac{\zeta^h-\rho^h}{\zeta-\rho}\tilde{\eta}_t^M\right) \\
&= \frac{1}{1-\beta q\zeta}\tilde{p}_t^M + \frac{\zeta\rho}{1-\beta\zeta\rho}\frac{1}{\zeta-\rho}\left(\frac{1}{1-\beta q\zeta} - \frac{1}{1-\beta q\rho}\right)\tilde{\eta}_t^M \\
&= \frac{1}{1-\beta q\zeta}\tilde{p}_t^M + \frac{\zeta\rho}{1-\beta\zeta\rho}\frac{1}{\zeta-\rho}\frac{\beta q\zeta - \beta q\rho}{(1-\beta q\zeta)(1-\beta q\rho)}\tilde{\eta}_t^M \\
&= \frac{1}{1-\beta q\zeta}\tilde{p}_t^M + \frac{\beta q\zeta\rho}{(1-\beta\zeta\rho)(1-\beta q\zeta)(1-\beta q\rho)}\tilde{\eta}_t^M
\end{aligned}$$

As such, under discretion, the optimal response of inflation can be written as:

$$\pi_t = \frac{qz}{1-\beta q\zeta}\tilde{p}_t^M + \frac{\zeta}{1-\beta\zeta\rho}\frac{qz}{1-\beta q\zeta}\frac{\beta q\rho}{1-\beta q\rho}\tilde{\eta}_t^M + \frac{q}{1-\beta q\rho}\eta_t$$

Denote this expression as:

$$\pi_t = \chi_1 \tilde{p}_t^M + \chi_2 \tilde{\eta}_t^M + \chi_3 \eta_t$$

If desired, using the defintion for q and z, one can rewrite these coefficients as:

$$\begin{aligned}
\chi_1 &= \frac{\lambda\Omega_2 - \kappa\Omega_1}{\kappa^2 + \Omega_2(1-\beta\zeta)} \\
\chi_2 &= \frac{\zeta}{1-\beta\zeta\rho} \frac{\lambda\Omega_2 - \kappa\Omega_1}{\kappa^2 + \Omega_2(1-\beta\zeta)} \frac{\beta\rho\Omega_2}{\kappa^2 + \Omega_2(1-\beta\rho)} \\
\chi_3 &= \frac{\Omega_2}{\kappa^2 + \Omega_2(1-\beta\rho)}
\end{aligned}$$

One can then solve for output in terms of inflation using the optimality condition given above.

$$\begin{aligned}
y_t &= -\frac{\kappa}{\Omega_2}\pi_t - \frac{\Omega_1}{\Omega_2}\tilde{p}_t^M \\
&= -\frac{\kappa}{\Omega_2}(\chi_1 \tilde{p}_{t-1}^M + \chi_2 \tilde{\eta}_t^M + \chi_3 \eta_t) - \frac{\Omega_1}{\Omega_2}\tilde{p}_t^M \\
&= -\left(\frac{\kappa}{\Omega_2}\chi_1 + \frac{\Omega_1}{\Omega_2}\right)\tilde{p}_t^M - \frac{\kappa}{\Omega_2}\chi_2 \tilde{\eta}_t^M - \frac{\kappa}{\Omega_2}\chi_3 \eta_t
\end{aligned}$$

Finally, one can solve for the interest rate using the IS curve:

$$i_t = \sigma^{-1}(E_t y_{t+1} - y_t + \alpha E_t \tilde{p}_{t+1}^M - \alpha \tilde{p}_t^M + \sigma E_t \pi_{t+1})$$

Working through the algebra gives the following response of nominal interest rates:

$$\begin{aligned}
i_t =& \frac{1}{\sigma}\left(\left(\frac{\kappa}{\Omega_2}\chi_1 + \frac{\Omega_1}{\Omega_2} - \alpha\right)(1-\zeta) + \sigma\zeta\chi_1\right)\tilde{p}_t^M \\
&+ \frac{1}{\sigma}\left(\left(\frac{\kappa}{\Omega_2}(1-\rho) + \sigma\rho\right)\chi_2 + \left(\alpha + \sigma\chi_1 - \frac{\kappa}{\Omega_2}\chi_1 - \frac{\Omega_1}{\Omega_2}\right)\frac{\zeta\rho}{1-\beta\zeta\rho}\right)\tilde{\eta}_t^M \\
&+ \frac{1}{\sigma}\left(\frac{\kappa}{\Omega_2}(1-\rho) + \sigma\rho\right)\chi_3 \eta_t
\end{aligned}$$

Optimal Policy under Discretion: RANK Model with $\rho = 0$

The algebraic expressions derived above look complex, but become much simpler in certain special cases. In particular, I consider the representative agent version of the model in which shocks have zero persistence. In the RANK model, $\Omega_1 = 0$. If $\rho = 0$ as well,

then the expression for π_t becomes:

$$\pi_t = \frac{\lambda\Omega_2}{\kappa^2 + \Omega_2(1-\beta\zeta)}\tilde{p}_t^M + \frac{\Omega_2}{\kappa^2+\Omega_2}\eta_t$$

For y_t, one has:

$$y_t = -\frac{\kappa\lambda}{\kappa^2 + \Omega_2(1-\beta\zeta)}\tilde{p}_t^M - \frac{\kappa}{\kappa^2+\Omega_2}\eta_t$$

Finally, for i_t:

$$i_t = \left(\lambda\frac{\kappa\bar{\gamma}(1-\zeta)+\zeta\Omega_2}{\kappa^2+\Omega_2(1-\beta\zeta)} - (1-\zeta)\right)\tilde{p}_t^M + \frac{1}{\sigma}\left(\frac{\kappa}{\Omega_2}(1-\rho)+\sigma\rho\right)\frac{\Omega_2}{\kappa^2+\Omega_2}\eta_t$$

Note that here I used the fact that under RANK $\alpha = \sigma = \bar{\gamma}^{-1}$.

Note that the central bank's optimal response to relative price shocks is determined by two considerations. On the one hand, the optimal choice of inflation-output trade-off requires that the central bank partially offset the inflationary pressure arising from relative prices in the Phillips curve with lower output, and hence a higher nominal rate. This effect is represented by the first-term in brackets, which is strictly positive. On the other hand, higher relative prices directly depress output via the IS curve, which calls for a lower nominal interest rate. This effect is given by the second term in brackets.

3.A.6 Optimal Policy under Commitment

In this case, the central bank's maximization problem is as described above. However, the central bank now does not take $E_t\pi_{t+1}$ as exogenous. As such, the Lagrangian is the same as above, but the first-order conditions become:

$$\pi_t - \Gamma_t + \Gamma_{t-1} = 0$$
$$\Omega_1\tilde{p}_t^M + \Omega_2 y_t + \kappa\Gamma_t = 0$$

As in the textbook model, assume $\Gamma_{-1} = 0$. Then in time $t = 0$, one has:

$$\Omega_1\tilde{p}_0^M + \Omega_2 y_0 + \kappa\pi_0 = 0$$

Define $\hat{p}_t = p_t - p_{-1}$. Using the expression above:

$$\begin{aligned}\hat{p}_t &= \sum_{s=0}^{t} \pi_s \\ &= \Gamma_t - \Gamma_{-1} \\ &= \Gamma_t \\ &= -\frac{\Omega_2}{\kappa} y_t - \frac{\Omega_1}{\kappa} \tilde{p}_t^M\end{aligned}$$

Re-arranging yields the following expression for output:

$$y_t = -\frac{\kappa}{\Omega_2} \hat{p}_t - \frac{\Omega_1}{\Omega_2} \tilde{p}_t^M$$

This is the same as in the textbook model, but with the additional term in $\tilde{p}_t^M$ that appears in the two agent version of the model when $\Omega_1 > 0$. As in the textbook model, one can proceed using the Phillips curve:

$$\begin{aligned}\pi_t &= E_t \beta \pi_{t+1} + \kappa y_t + \lambda \tilde{p}_t^M + \eta_t \\ \hat{p}_t - \hat{p}_{t-1} &= \beta \hat{p}_{t+1} - \beta \hat{p}_t + \kappa y_t + \lambda \tilde{p}_t^M + \eta_t \\ \hat{p}_t - \hat{p}_{t-1} &= \beta \hat{p}_{t+1} - \beta \hat{p}_t + \kappa \left(-\frac{\kappa}{\Omega_2} \hat{p}_t - \frac{\Omega_1}{\Omega_2} \tilde{p}_t^M \right) + \lambda \tilde{p}_t^M + \eta_t \\ \left(1 + \beta + \frac{\kappa^2}{\Omega_2} \right) \hat{p}_t &= \hat{p}_{t-1} + \beta \hat{p}_{t+1} + \left(\lambda - \kappa \frac{\Omega_1}{\Omega_2} \right) \tilde{p}_t^M + \eta_t\end{aligned}$$

To lighten the notation, define:

$$\begin{aligned}x &= \frac{\Omega_2}{(1+\beta)\Omega_2 + \kappa^2} \\ z &= \lambda - \kappa \frac{\Omega_1}{\Omega_2}\end{aligned}$$

Using this notation, the above expression becomes:

$$\hat{p}_t = x \hat{p}_{t-1} + \beta x \hat{p}_{t+1} + x z \tilde{p}_t^M + x \eta_t$$

One can solve this expression using lag notation:

$$\begin{aligned}(1 - xL - \beta x L^{-1}) \hat{p}_t &= x z \tilde{p}_t^M + x \eta_t \\ \left(1 - \frac{1}{\beta x} L + \frac{1}{\beta} L^2 \right) \hat{p}_t &= -\frac{1}{\beta} L \left(z \tilde{p}_t^M + \eta_t \right)\end{aligned}$$

As in the textbook model, the two roots of the lag quadratic equation are:

$$\frac{1 \pm \sqrt{1 - 4\beta x^2}}{2\beta x}$$

The stable root is:

$$\frac{1 - \sqrt{1 - 4\beta x^2}}{2\beta x}$$

I will denote this fraction as δ. The unstable root is then $1/(\beta\delta)$. The equation then becomes:

$$\begin{aligned}
(1 - \delta L)(1 - \beta^{-1}\delta^{-1}L)\hat{p}_t &= -\frac{1}{\beta} L \left(z\tilde{p}_t^M + \eta_t \right) \\
\hat{p}_t &= \delta \hat{p}_{t-1} - \frac{1}{\beta} \frac{L}{1 - \beta^{-1}\delta^{-1}L} \left(z\tilde{p}_t^M + \eta_t \right) \\
\hat{p}_t &= \delta \hat{p}_{t-1} + \frac{\delta}{1 - \beta\delta L^{-1}} \left(z\tilde{p}_t^M + \eta_t \right) \\
\hat{p}_t &= \delta \hat{p}_{t-1} + \delta \sum_{h \geq 0} (\beta\delta)^h \left(z\tilde{p}_t^M + \eta_t \right)
\end{aligned}$$

Using the result derived in the discretion case, we have that:

$$E_t \sum_{h \geq 0} (\beta\delta)^h \tilde{p}_{t+h}^M = \frac{1}{1 - \beta\delta\zeta} \tilde{p}_t^M + \frac{\beta\delta\zeta\rho}{(1 - \beta\zeta\rho)(1 - \beta\delta\zeta)(1 - \beta\delta\rho)} \tilde{\eta}_t^M$$

Substituting this expression into the equation for $\hat{p}_t$ above gives an expression for the evolution of the price level in terms of the exogenous variables:

$$\hat{p}_t = \delta \hat{p}_{t-1} + z\delta \left(\frac{1}{1 - \beta\delta\zeta} \tilde{p}_t^M + \frac{\beta\delta\zeta\rho}{(1 - \beta\zeta\rho)(1 - \beta\delta\zeta)(1 - \beta\delta\rho)} \tilde{\eta}_t^M \right) + \frac{\delta}{1 - \beta\delta\rho} \eta_t$$

Finally, using the expression for $\hat{p}_t$ derived above, one can express output in terms of the exogenous variables:

$$\begin{aligned}
-\frac{\Omega_2}{\kappa} y_t - \frac{\Omega_1}{\kappa} \tilde{p}_t^M = &-\delta \frac{\Omega_2}{\kappa} y_{t-1} - \delta \frac{\Omega_1}{\kappa} \tilde{p}_{t-1}^M \\
&+ z\delta \left(\frac{1}{1 - \beta\delta\zeta} \tilde{p}_t^M + \frac{\beta\delta\zeta\rho}{(1 - \beta\zeta\rho)(1 - \beta\delta\zeta)(1 - \beta\delta\rho)} \tilde{\eta}_t^M \right) + \frac{\delta}{1 - \beta\delta\rho} \eta_t
\end{aligned}$$

Rearranging yields:

$$y_t = \delta y_{t-1} - \frac{\Omega_1}{\Omega_2}(\tilde{p}_t^M - \delta \tilde{p}_{t-1}^M) \\ - \frac{\kappa}{\Omega_2} z\delta \left(\frac{1}{1-\beta\delta\zeta}\tilde{p}_t^M + \frac{\beta\delta\zeta\rho}{(1-\beta\zeta\rho)(1-\beta\delta\zeta)(1-\beta\delta\rho)}\tilde{\eta}_t^M \right) - \frac{\kappa}{\Omega_2}\frac{\delta}{1-\beta\delta\rho}\eta_t$$

Optimal Policy under Commitment: RANK Model with $\rho = 0$

Once again, it is instructive to consider the simpler case of a RANK model with $\rho = 0$. In this case, one can make a few simplifications. First, note that $z = \lambda$. Additionally, the fact that $\rho = 0$ allows one to simplify the terms in $\tilde{p}_{t-1}^M$ and $\tilde{\eta}_t^M$ into a single expression for $\tilde{p}_t^M$, as in the discretion case. The above expressions for the price level and output then become:

$$\hat{p}_t = \delta \hat{p}_{t-1} + \frac{\lambda\delta}{1-\beta\delta\zeta}\tilde{p}_t^M + \delta\eta_t$$
$$y_t = \delta y_{t-1} - \frac{\kappa\lambda\delta}{\Omega_2(1-\beta\delta\zeta)}\tilde{p}_t^M - \frac{\kappa}{\Omega_2}\eta_t$$

3.B How Large are Shocks to the Relative Price of the Marginal Basket?

The effects described in the main paper will only be quantitatively important in practice if the relative price of the marginal basket varies meaningfully over time. In this section, I construct a rough approximation of the marginal consumption basket using UK data, under the assumption that preferences over the main categories of spending (food, energy, recreation, and so on) are common across households. I then argue that relative price of the marginal consumption basket is likely to vary relatively little most of the time, but can vary considerably when there are large shocks to energy or food prices. To illustrate the latter point, I calculate the difference between the inflation rate of the average and marginal baskets over time.

3.B.1 Data and Calculation

I begin by taking data from the UK's Office of National Statistics (ONS) that give the expenditure of UK households by COICOP (Classification of Individual Consumption According to Purpose) classification by equivalised disposable income decile.[18] The data are equivalized using the OECD-modified scale, which takes into account the size and age of the households.[19] I take data for 2019, which I assume to be more representative of the "steady-state" consumption basket than more recent data, given the size of the shocks hitting the economy since early 2020.

I then use data for the average expenditure across household to construct the average consumption basket at the COICOP class level. For reference, I also construct the average consumption basket for households at different income deciles.

To construct an approximation of the marginal consumption basket, I take the difference in spending on each COICOP class between the second and ninth deciles, and divide by the difference in total spending across all classes in those deciles.[20]

[18]The data are given in Table 3.1E of Family Expenditure Workbook 1 (Detailed Expenditure and Trends) from the ONS.

[19]In the non-equivalised data, households in the top disposable income decile tend to be substantially larger than those in the bottom disposable income decile, with 3.5 persons per household in the top decile compared to 1.4 in the bottom decile. The assumption of common preferences would be less accurate with those data, since changes in household age and size would likely lead to differences in consumption patterns for some goods.

[20]Note that when calculating differences in inflation rates between high- and low-income households, the ONS often uses the second and ninth rather than the first and tenth deciles to avoid outliers. Under Stone-Geary preferences, the marginal basket should be the same for households at any income level. This is not the case, however, under other forms of non-homothetic preferences, where the marginal basket differs between low- and high-income households. The simple method that I use gives an approximation of the marginal consumption basket across different income levels. Calculating the marginal consumption basket based on differences between the first and tenth deciles does not meaningfully change the results

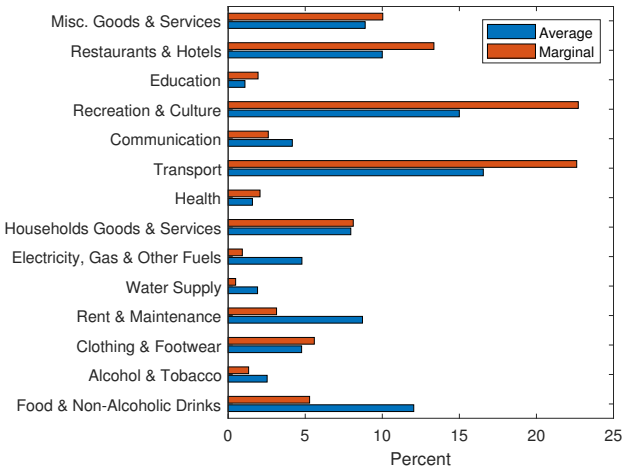


Figure 3.8: Composition of Average and Marginal Baskets, Based on 2019 UK Data

Of course, this method assumes that households of the same age and size have common preferences. Concretely, it assumes that if I took a household currently in the bottom decile for disposable income, and increased their income to the level of the top decile, their new consumption basket would be equal to that of those households (of the same age and size) currently in the top decile. It will give a poor approximation to the marginal basket if preferences over items such as good and energy differ considerably across households of the same age and size. That is, the approximation will be poor if most of the differences in equivalized consumption patterns between poorer and richer households are explained by differences in their preferences rather than differences in incomes. Future work could assess whether this is the case in more detail. Nonetheless, I personally find it far more intuitive that the differences in consumption patterns outlined below stem from certain goods being necessities and others luxuries, rather than households with different income levels having vastly different preferences over basic items such as food and heating.

3.B.2 The Marginal and Average Baskets

Figure 3.8 shows the differences in the weightings of different COICOP classes between the average and marginal basket (in percentage points). A positive number indicates that the good or service in question has a higher weighting in the marginal than the average basket, while a negative number indicates that it has a lower weighting in the marginal than the average basket.

For most goods, the difference in the weighting is quite small. There are, however, several categories which have very different weightings. In particular, food, gas and electricity have substantially lower weights in the marginal relative to the average basket,

outlined below. Nor indeed does construct another a measure of the marginal basket at the middle of the distribution, using differences between the fourth and seventh deciles.

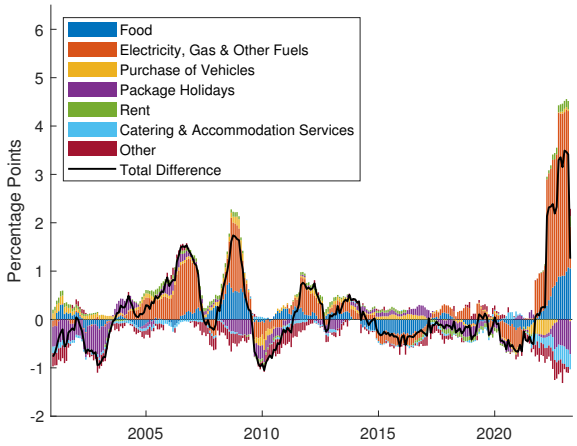


Figure 3.9: Differences Between Inflation On Average and Marginal Basket (Year-on-Year), and Contributions from Different COICOP Categories

whereas packaged holidays have a much larger weight in the marginal than the average basket. This intuitively makes sense; food, gas, and electricity are necessities, and so most households spend similar amounts on those goods irrespective of their income. Packaged holidays are a luxury item, and so spending increases disproportionately with income.

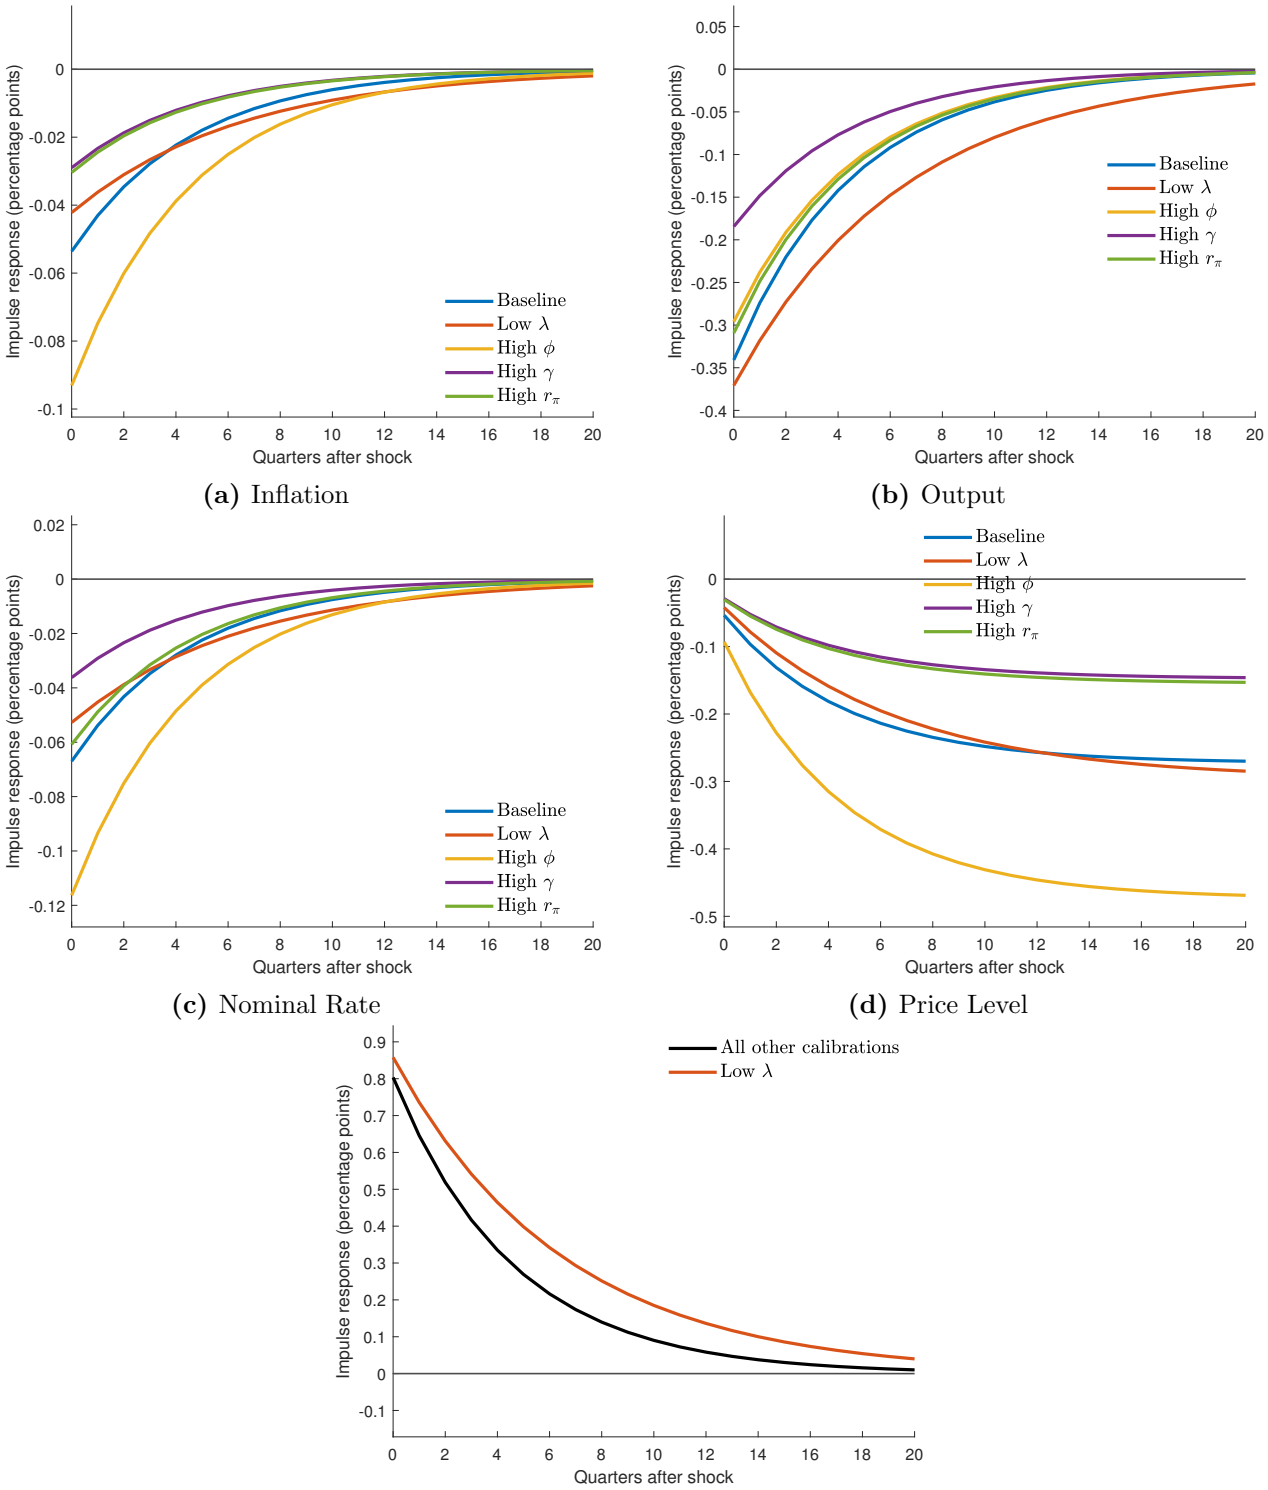
Other items with a lower weight in the marginal basket include water, non-alcoholic drinks, and tobacco. On the other hand, transport services, purchase of vehicles, and catering and accommodation services have a higher weight in the marginal basket.

I then take data on the evolution of prices for each COICOP class to track how the inflation rate of the average and marginal baskets has evolved over time.[21] The differences between the two inflation rates is shown in 3.9. The two inflation rates is relatively similar in most periods, but in periods with very high inflation of energy or food, a large difference can open up between the two. This difference can be as large as several percentage points.

These results suggests that variations in the relative price of the marginal basket are not too large most of the time, but may be considerable in the presence of large food and energy price shocks.

[21] I use data from Table 4 of the Consumer Price Inflation Detailed Reference Tables from the ONS.

3.C Additional Figures



(a) Inflation

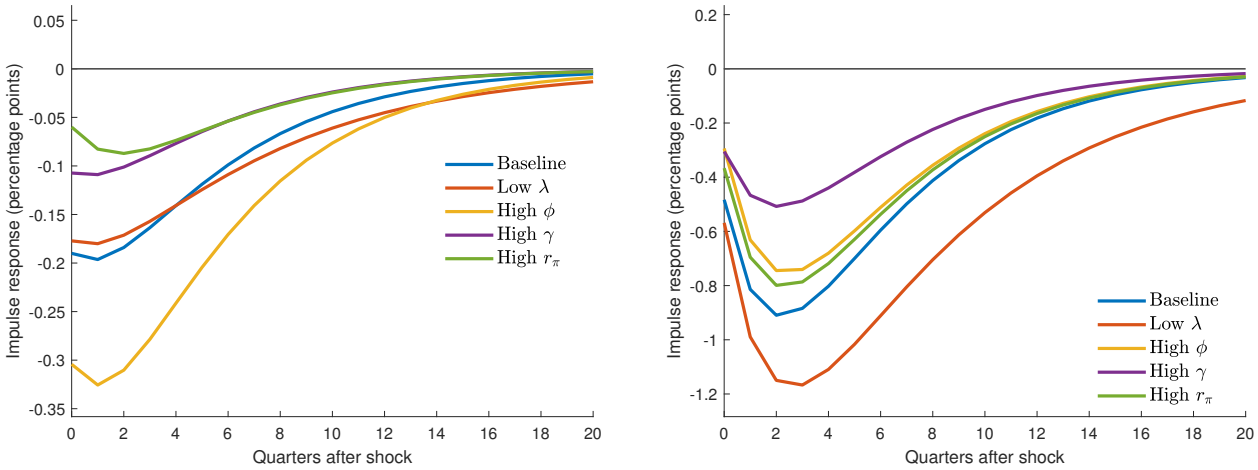
(b) Output

(c) Nominal Rate

(d) Price Level

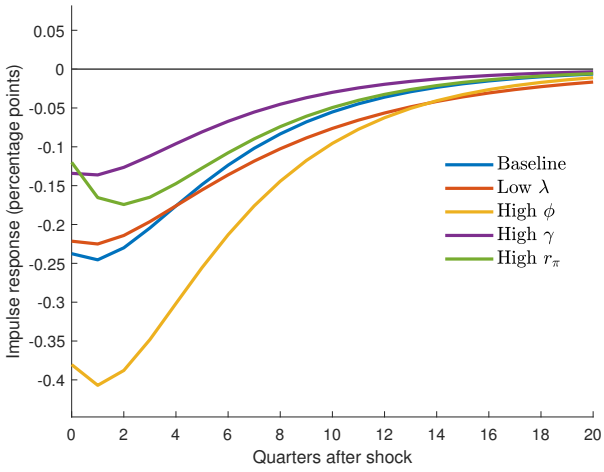
(e) Relative Price of Marginal Basket

Figure 3.10: Impulse responses to a relative price shock in the RANK model with $\rho = 0$ for a variety of calibrations.

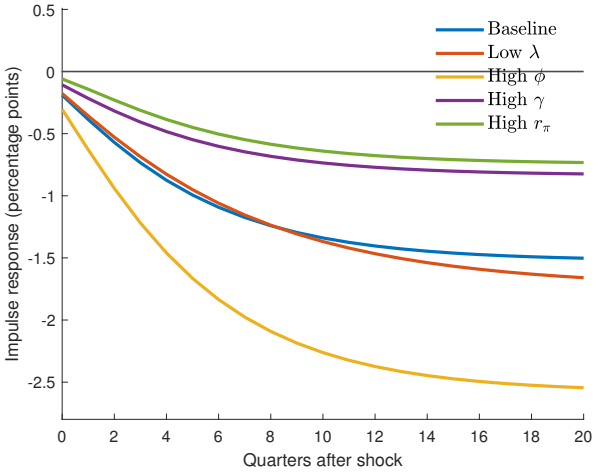


(a) Inflation

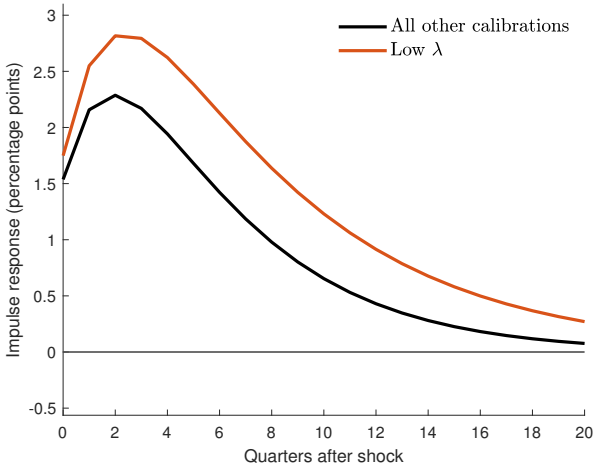
(b) Output



(c) Nominal Rate

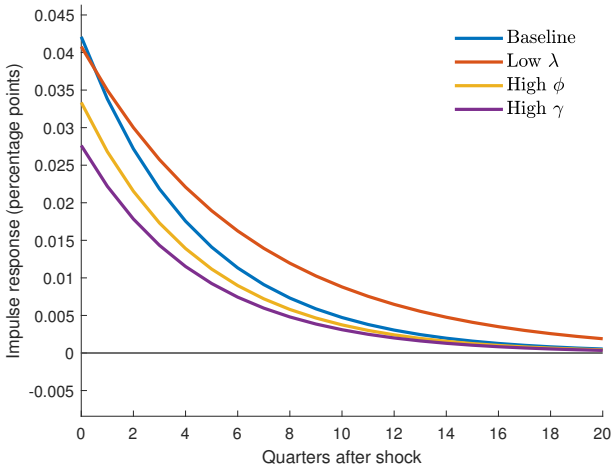


(d) Price Level

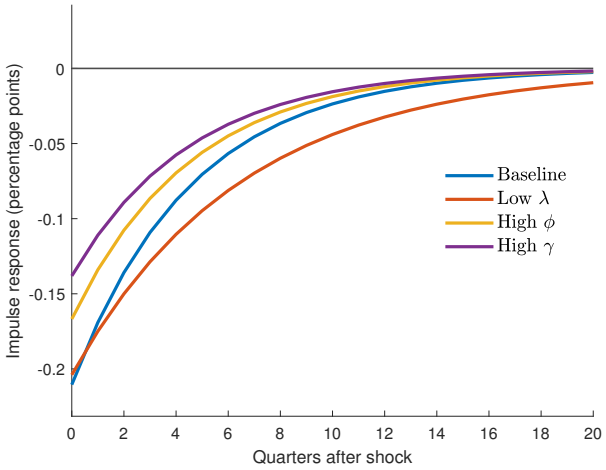


(e) Relative Price of Marginal Basket

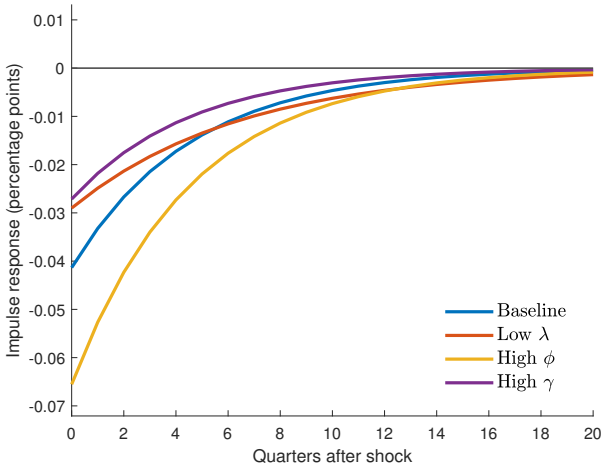
Figure 3.11: Impulse responses to a relative price shock in the RANK model with $\rho = 0.6$ for a variety of calibrations.



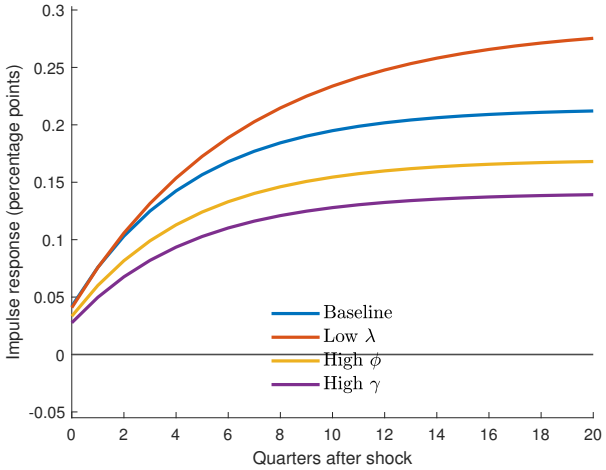
(a) Inflation



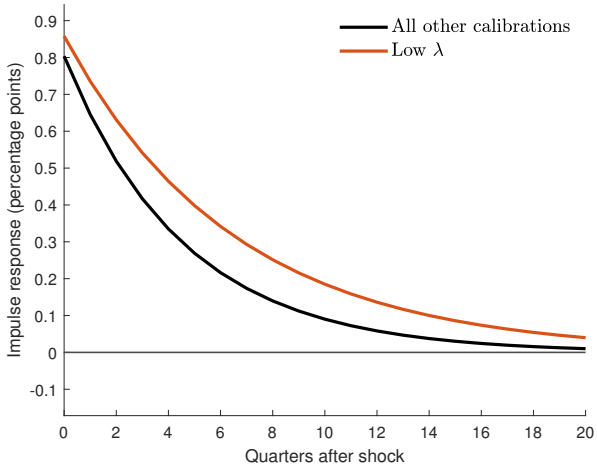
(b) Output



(c) Nominal Rate

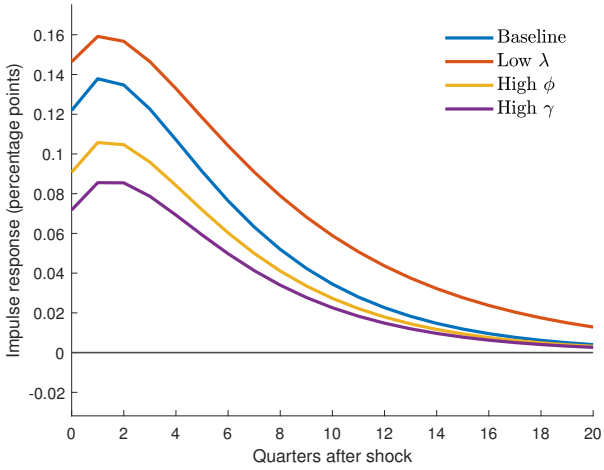


(d) Price Level

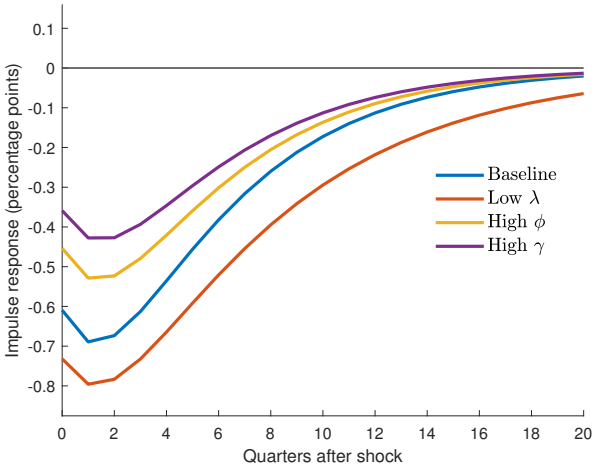


(e) Relative Price of Marginal Basket

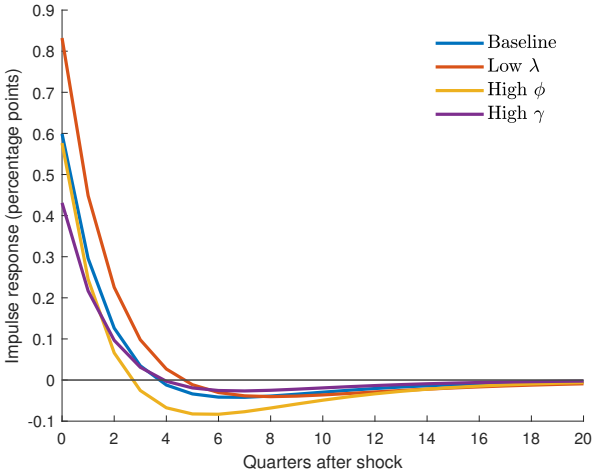
Figure 3.12: Impulse responses to a relative price shock in the RANK model with $\rho = 0$ for a variety of calibrations, with policy set optimally under discretion.



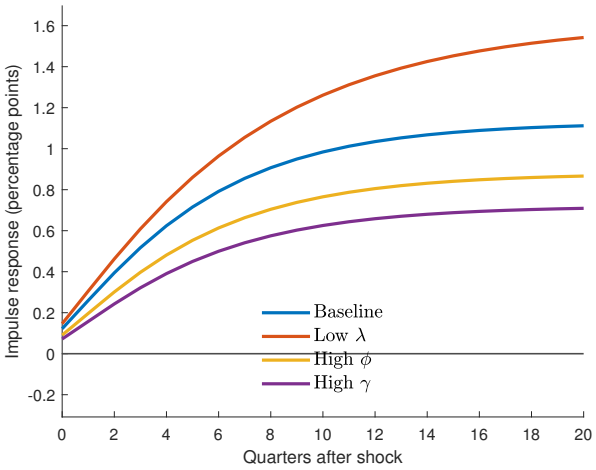
(a) Inflation



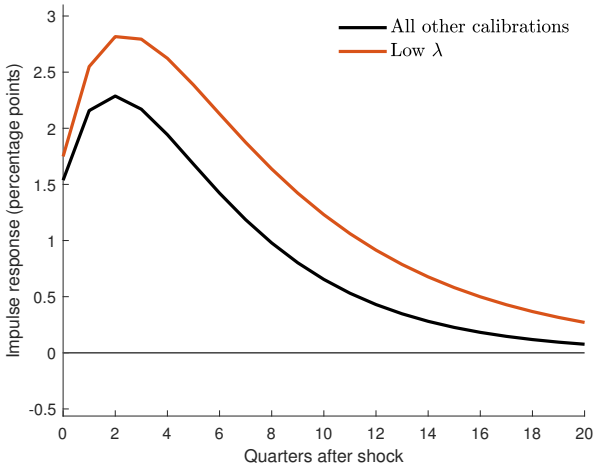
(b) Output



(c) Nominal Rate

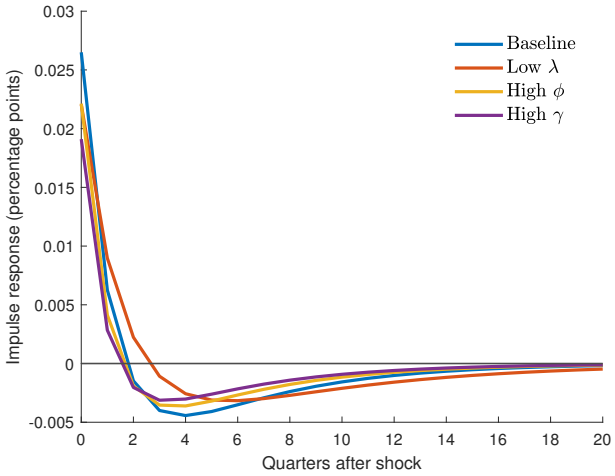


(d) Price Level

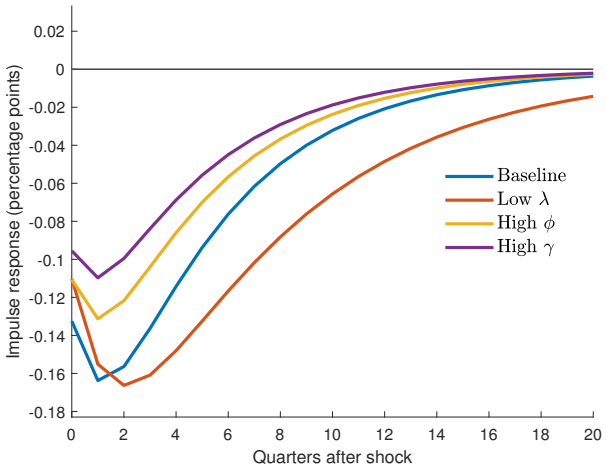


(e) Relative Price of Marginal Basket

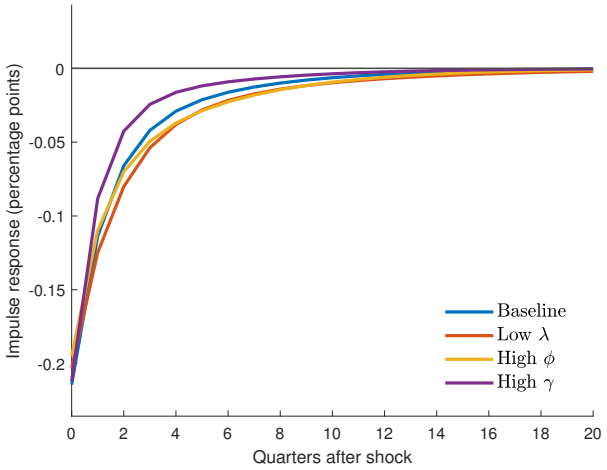
Figure 3.13: Impulse responses to a relative price shock in the RANK model with $\rho = 0.6$ for a variety of calibrations, with policy set optimally under discretion.



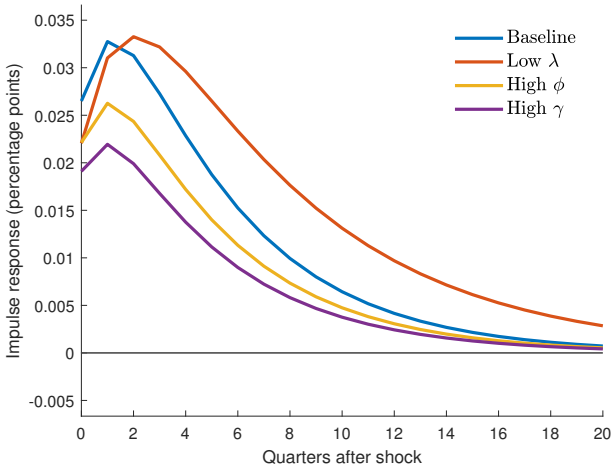
(a) Inflation



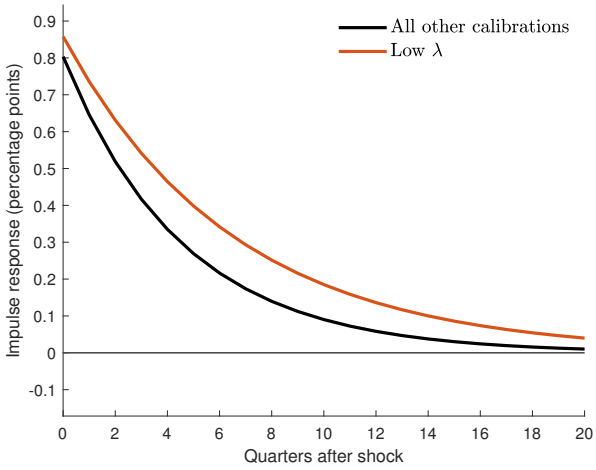
(b) Output



(c) Nominal Rate

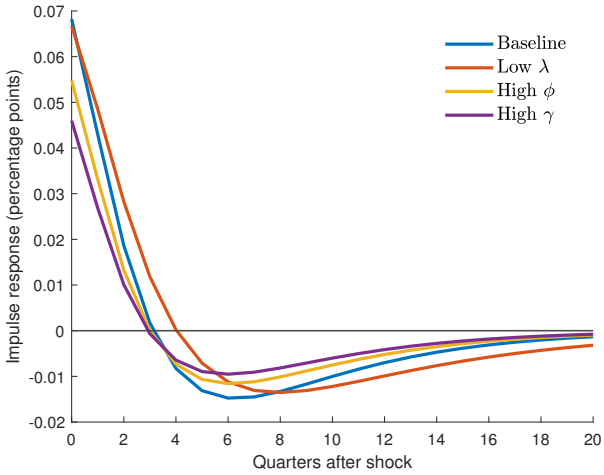


(d) Price Level

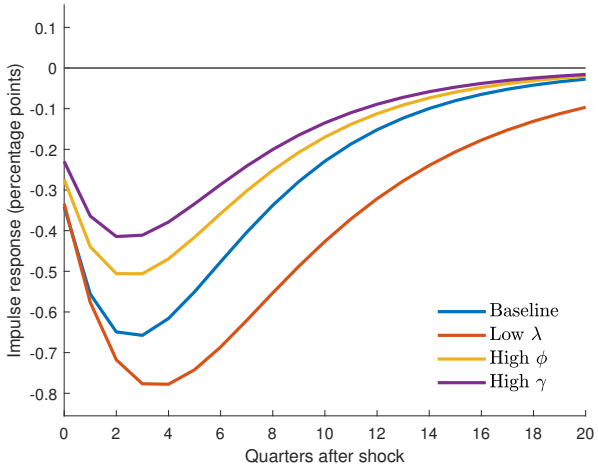


(e) Relative Price of Marginal Basket

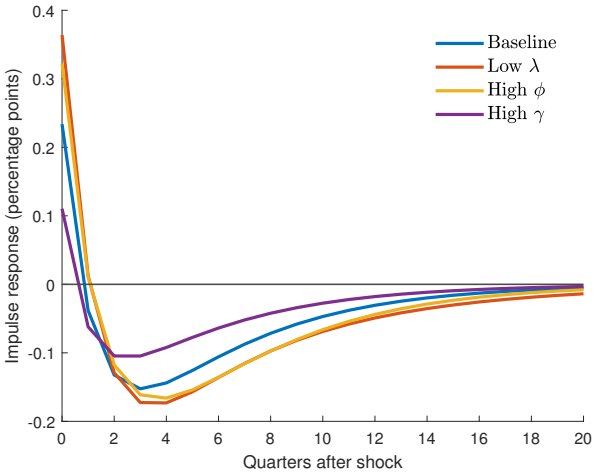
Figure 3.14: Impulse responses to a relative price shock in the RANK model with $\rho = 0$ for a variety of calibrations, with policy set optimally under commitment.



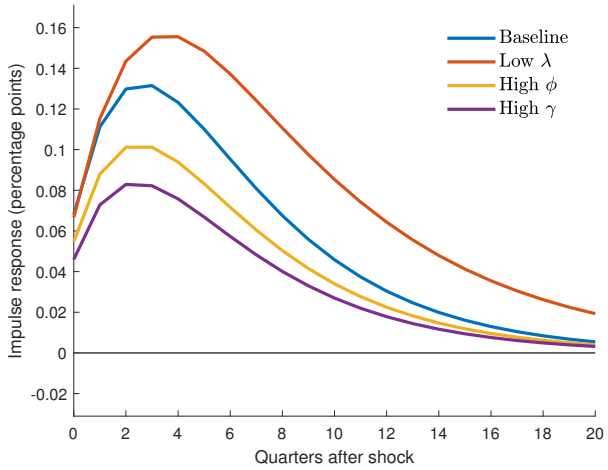
(a) Inflation



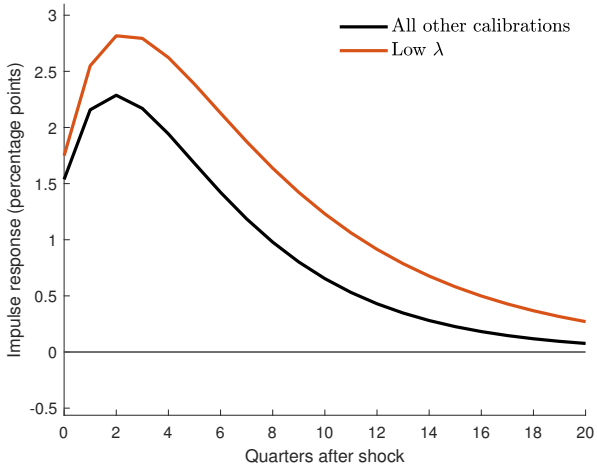
(b) Output



(c) Nominal Rate



(d) Price Level



(e) Relative Price of Marginal Basket

Figure 3.15: Impulse responses to a relative price shock in the RANK model with $\rho = 0.6$ for a variety of calibrations, with policy set optimally under commitment.

Bibliography

Adam, K. (2007). Experimental Evidence on the Persistence of Output and Inflation. *The Economic Journal*, 117(520):603–636.

Afrouzi, H., Kwon, S. Y., Landier, A., Ma, Y., and Thesmar, D. (2023). Overreaction and working memory. *The Quarterly Journal of Economics*, forthcoming.

Anand, R., Prasad, E. S., and Zhang, B. (2015). What measure of inflation should a developing country central bank target? *Journal of Monetary Economics*, 74:102–116.

Andrade, J., Cordeiro, P., and Lambais, G. (2019). Estimating a behavioral new keynesian model. *arXiv preprint arXiv:1912.07601*.

Andrade, P., Crump, R. K., Eusepi, S., and Moench, E. (2016). Fundamental disagreement. *Journal of Monetary Economics*, 83:106–128.

Andre, P., Haaland, I., Roth, C., and Wohlfart, J. (2021). Inflation Narratives. *ECONtribute Discussion Papers*, (127).

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2022). Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample. *Review of Economic Studies*, 89(6).

Andrews, I. (2018). Valid two-step identification-robust confidence sets for gmm. *Review of Economics and Statistics*, 100(2):337–348.

Andrews, I. and Mikusheva, A. (2014). Weak identification in maximum likelihood: A question of information. *American Economic Review*, 104(5):195–99.

Andrews, I. and Mikusheva, A. (2015). Maximum likelihood inference in weakly identified dynamic stochastic general equilibrium models. *Quantitative Economics*, 6(1):123–152.

Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2020). Imperfect Macroeconomic Expectations: Evidence and Theory. *NBER Macroeconomics Annual*, 35(1):2–86.

Aoki, K. (2001). Optimal monetary policy responses to relative-price changes. *Journal of monetary economics*, 48(1):55–80.

Armantier, O., Topa, G., van der Klaauw, W., and Zafar, B. (2017). An overview of

the Survey of Consumer Expectations. *Federal Reserve Bank of New York Economic Policy Review*, 23(2).

Armona, L., Fuster, A., and Zafar, B. (2019). Home price expectations and behaviour: Evidence from a randomized information experiment. *Review of Economic Studies*, 86(4):1371–1410.

Ascari, G., Mavroeidis, S., and McClung, N. (2023). Coherence without rationality at the zero lower bound. *Journal of Economic Theory*, 214:105745.

Bauer, M. D., Lakdawala, A., and Mueller, P. (2022). Market-Based Monetary Policy Uncertainty. *The Economic Journal*, 132(644):1290–1308.

Ben-David, I., Fermand, E., Kuhnen, C. M., and Li, G. (2018). Expectations Uncertainty and Household Economic Behavior. *NBER Working Paper Series*, (25336).

Benati, L. (2008). Investigating inflation persistence across monetary regimes. *The Quarterly Journal of Economics*, 123(3):1005–1060.

Beshears, J., Choi, J. J., Fuster, A., Laibson, D., and Madrian, B. C. (2013). What goes up must come down? experimental evidence on intuitive forecasting. *American Economic Review*, 103(3):570–74.

Bhandari, A., Borovicka, J., and Ho, P. (2019). Survey Data and Subjective Beliefs in Business Cycle Models. *Federal Reserve Bank of Richmond Working Papers*, 19(14):1–60.

Bilbiie, F. O. (2008). Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of economic theory*, 140(1):162–196.

Bilbiie, F. O. (2018). Monetary policy and heterogeneity: An analytical framework. *Available at SSRN 3106805*.

Bilbiie, F. O. (2021). Monetary policy and heterogeneity: An analytical framework.

Bils, M. and Klenow, P. J. (1998). Using consumer theory to test competing business cycles models. *Journal of Political Economy*, 106(2):233–261.

Binder, C. C. (2017). Measuring uncertainty based on rounding: New method and application to inflation expectations. *Journal of Monetary Economics*, 90:1–12.

Blanco, C. and Diz, S. (2021). Optimal monetary policy with non-homothetic preferences.

Boneva, L., Cloyne, J., Weale, M., and Wieladek, T. (2020). Firms' Price, Cost and Activity Expectations: Evidence from Micro Data. *The Economic Journal*, 130(627):555–586.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.

Borio, C., Disyatat, P., Xia, D., and Zakrajšek, E. (2021). Monetary policy, relative prices and inflation control: flexibility born out of success.

Branch, W. A. (2004). The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations. *The Economic Journal*, 114(497):592–621.

Branch, W. A. and Evans, G. W. (2006). Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory*, 127(1):264–295.

Broer, T., Kohlhas, A., Mitman, K., and Schlafmann, K. (2021). Information and Wealth Heterogeneity in the Macroeconomy.

Broer, T. and Kohlhas, A. N. (2019). Forecaster (Mis-)Behavior.

Budianto, F., Nakata, T., and Schmidt, S. (2023). Average inflation targeting and the interest rate lower bound. *European Economic Review*, 152:104384.

Bullard, J. and Suda, J. (2016). The stability of macroeconomic systems with Bayesian learners. *Journal of Economic Dynamics and Control*, 62:1–16.

Canova, F. and Ferroni, F. (2011). Multiple filtering devices for the estimation of cyclical dsge models. *Quantitative Economics*, 2(1):73–98.

Cantore, C. and Freund, L. B. (2021). Workers, capitalists, and the government: fiscal policy and income (re) distribution. *Journal of monetary economics*, 119:58–74.

Caplin, A. and Dean, M. (2015). Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review*, 105(7):2183–2203.

Caron, J., Fally, T., and Markusen, J. R. (2014). International trade puzzles: A solution linking production and preferences. *The Quarterly Journal of Economics*, 129(3):1501–1552.

Cavallari, L. (2018). Monetary policy with non-homothetic preferences.

Cavallo, A., Cruces, G., and Perez-Truglia, R. (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, 9(3):1–35.

Châu, V. T. (2019). Behavioral attention phillips curve.

Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., and van Rooij, M. (2021). Heterogeneous wealth effects. *European Economic Review*, 137:103805.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the

dynamic effects of a shock to monetary policy. *Journal of political Economy*, 113(1):1–45.

Clarida, R., Gali, J., and Gertler, M. (2000). Monetary policy rules and macroeconomic stability: evidence and some theory. *The Quarterly journal of economics*, 115(1):147–180.

Cochrane, J. H. (2016). Comments on 'a behavioral new keynesian model'by xavier gabaix. *mimeo*.

Coibion, O., Georgarakos, D., Gorodnichenko, Y., Kenny, G., and Weber, M. (2021). The Effect of Macroeconomic Uncertainty on Household Spending. *NBER Working Paper Series*, 28625.

Coibion, O. and Gorodnichenko, Y. (2012). What Can Survey Forecasts Tell Us about Information Rigidities? *Journal of Political Economy*, 120(1):116–159.

Coibion, O. and Gorodnichenko, Y. (2015a). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–78.

Coibion, O. and Gorodnichenko, Y. (2015b). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.

Coibion, O., Gorodnichenko, Y., and Ropele, T. (2020). Inflation expectations and firm decisions: New causal evidence. *The Quarterly Journal of Economics*, 135(1):165–219.

Comin, D., Lashkari, D., and Mestieri, M. (2021). Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374.

Cragg, J. G. (1971). Some Statistical Models for Limited Dependent Variables with Application to the Demand for Durable Goods. *Econometrica*, 39(5):829.

Crossley, T., de Bresser, J., Delaney, L., and Winter, J. (2017). Can survey participation alter household saving behaviour? *The Economic Journal*, 127(606):2332–2357.

Cúrdia, V. and Reis, R. (2020). Correlated disturbances and us business cycles. Technical report.

De Grauwe, P. (2012). *Lectures on behavioral macroeconomics*. Princeton University Press.

Debortoli, D., Kim, J., Lindé, J., and Nunes, R. (2019). Designing a simple loss function for central banks: Does a dual mandate make sense? *The Economic Journal*, 129(621):2010–2038.

Del Negro, M., Giannoni, M. P., and Patterson, C. (2012). The forward guidance puzzle. *FRB of New York Staff Report*, (574).

Dominitz, J. and Manski, C. F. (2011). Measuring and Interpreting Expectations of Equity Returns. *Journal of Applied Econometrics*, 370(November 2010):352–370.

Dräger, L., Lamla, M. J., and Pfajfar, D. (2022). How to Limit the Spillover from the 2021 Inflation Surge to Inflation Expectations?

Dräger, L. and Nghiem, G. (2021). Are consumers' spending decisions in line with a euler equation? *Review of Economics and Statistics*, 103(3):580–596.

Engel, E. (1857). Die produktions-und consumtionsverhaltnisse des königreichs sachsen, reprinted with engel (1895), anlage i, 1-54. *TER*, 4(2):216–225.

Evans, G. W. and Honkapohja, S. (2012). Learning and expectations in macroeconomics. In *Learning and Expectations in Macroeconomics*. Princeton University Press.

Farhi, E. and Werning, I. (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, 109(11):3887–3928.

Fernández-Villaverde, J., Guerrón-Quintana, P. A., and Rubio-Ramírez, J. (2010). Reading the recent monetary history of the us, 1959-2007. Technical report, National Bureau of Economic Research.

Fuhrer, J. C. and Rudebusch, G. D. (2004). Estimating the euler equation for output. *Journal of Monetary Economics*, 51(6):1133–1153.

Fuster, A., Kaplan, G., and Zafar, B. (2021). What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans. *The Review of Economic Studies*, 88(4):1760–1795.

Fuster, A., Laibson, D., and Mendel, B. (2010). Natural Expectations and Macroeconomic Fluctuations. *Journal of Economic Perspectives*, 24(4):67–84.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710.

Gabaix, X. (2016). Behavioral macroeconomics via sparse dynamic programming. Technical report, National Bureau of Economic Research.

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review*, 110(8):2271–2327.

Galı, J. and Gertler, M. (1999). Inflation dynamics: A structural econometric analysis. *Journal of monetary Economics*, 44(2):195–222.

Galí, J., López-Salido, J. D., and Vallés, J. (2007). Understanding the effects of gov-

ernment spending on consumption. *Journal of the european economic association*, 5(1):227–270.

Ganong, P. and Noel, P. (2017). Consumer spending during unemployment: Positive and normative implications.

Geweke, J. (1999). Using simulation methods for bayesian econometric models: inference, development, and communication. *Econometric reviews*, 18(1):1–73.

Goldstein, N. and Gorodnichenko, Y. (2022). Expectations Formation and Forward Information. *NBER Working Paper Series*, (29711).

Hajdini, I. (2020). Misspecified Forecasts and Myopia in an Estimated New Keynesian Model.

Hanspal, T., Weber, A., and Wohlfart, J. (2021). Exposure to the covid-19 stock market crash and its effect on household expectations. *Review of Economics and Statistics*, 103(5):994–1010.

Herbst, E. P. and Schorfheide, F. (2016). *Bayesian estimation of DSGE models*. Princeton University Press.

Hommes, C. (2021). Behavioral and experimental macroeconomics and policy analysis: A complex systems approach. *Journal of Economic Literature*, 59(1):149–219.

Hommes, C. and Lustenhouwer, J. (2019). Inflation targeting and liquidity traps under endogenous credibility. *Journal of Monetary Economics*, 107:48–62.

Houthakker, H. S. (1960). Additive preferences. *Econometrica: Journal of the Econometric Society*, pages 244–257.

Ilabaca, F., Meggiorini, G., and Milani, F. (2020). Bounded rationality, monetary policy, and macroeconomic stability. *Economics Letters*, 186:108522.

Iskrev, N. (2010). Local identification in dsge models. *Journal of Monetary Economics*, 57(2):189–202.

Jain, M. (2017). Perceived Inflation Persistence. *Journal of Business and Economic Statistics*, 37(1):110–120.

Jørgensen, P. L. and Lansing, K. J. (2019). Anchored inflation expectations and the slope of the phillips curve. *Federal Reserve Bank of San Francisco Working Paper*, 27.

Justiniano, A. and Primiceri, G. E. (2008). The time-varying volatility of macroeconomic fluctuations. *American Economic Review*, 98(3):604–41.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. Technical report, National Bureau of Economic Research.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the american statistical association*, 90(430):773–795.

Kim, G. and Binder, C. (2022). Learning-through-Survey in Inflation Expectations. *American Economic Journal: Macroeconomics*, forthcoming.

Kongsamut, P., Rebelo, S., and Xie, D. (2001). Beyond balanced growth. *The Review of Economic Studies*, 68(4):869–882.

Koop, G., Pesaran, M. H., and Smith, R. P. (2013). On identification of bayesian dsge models. *Journal of Business & Economic Statistics*, 31(3):300–314.

Kueng, L. (2018). Excess Sensitivity of High-Income Consumers. *The Quarterly Journal of Economics*, 133(4):1693–1751.

Kumar, S., Gorodnichenko, Y., and Coibion, O. (2022). The Effect of Macroeconomic Uncertainty on Firm Decisions. *NBER Working Paper Series*, 30288.

Lagarde, C. (2021). Learning from Lisbon: Recovery and resilience in Europe. *Speech given on the occasion of the 175th anniversary of Banco de Portugal in Lisbon, Lisbon, 03/11/2021*.

Lange, M., Zühlke, D., Holz, O., Villmann, T., and Mittweida, S.-G. (2014). Applications of lp-norms and their smooth approximations for gradient based learning vector quantization. In *ESANN*, pages 271–276.

Laudenbach, C., Weber, A., and Wohlfart, J. (2021). Beliefs About the Stock Market and Investment Choices: Evidence from a Field Experiment. *CEBI Working Paper Series*, (17/21).

Lee, S. (2021). Optimal monetary policy under heterogeneous consumption baskets.

Link, S., Peichl, A., Roth, C., and Wohlfart, J. (2023). Information Frictions Among Firms and Households. *Journal of Monetary Economics*.

Lubik, T. A. and Marzo, M. (2021). Fiscal policy perceptions in a behavioral new keynesian model. *Annals of Economics and Finance*, 22(2):255–287.

Lubik, T. A. and Schorfheide, F. (2004). Testing for indeterminacy: An application to us monetary policy. *American Economic Review*, 94(1):190–217.

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.

Macaulay, A. (2022). Shock Transmission and the Sources of Heterogeneous Expectations. *CESifo Working Paper*, (9733).

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2021). Rational inattention: A review.

Mackowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Maćkowiak, B. and Wiederholt, M. (2015). Business cycle dynamics under rational inattention. *The Review of Economic Studies*, 82(4):1502–1532.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Mankiw, N. G. and Reis, R. (2003). What measure of inflation should a central bank target? *Journal of the European Economic Association*, 1(5):1058–1086.

Mankiw, N. G., Reis, R., and Wolfers, J. (2004). Disagreement about Inflation Expectations. *NBER Macroeconomics Annual*, 18:209–270.

McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, 106(10):3133–58.

McLeay, M. and Tenreyro, S. (2020). Optimal inflation and the identification of the phillips curve. *NBER Macroeconomics Annual*, 34(1):199–255.

Meggiorini, G. (2021). Behavioral new keynesian models: An empirical assessment. *Available at SSRN 3984154*.

Nakamura, E. and Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464.

Pedemonte, M., Toma, H., and Verdugo, E. (2023). Aggregate Implications of Heterogeneous Inflation Expectations: The Role of Individual Experience. *Federal Reserve Bank of Cleveland Working Paper*, (23-04).

Pfäuti, O. and Seyrich, F. (2022). A behavioral heterogeneous agent new keynesian model.

Ravn, M. O., Schmitt-Grohé, S., and Uribe, M. (2008). Macroeconomics of subsistence points. *Macroeconomic Dynamics*, 12(S1):136–147.

Reis, R. (2006a). Inattentive consumers. *Journal of monetary Economics*, 53(8):1761–1800.

Reis, R. (2006b). Inattentive producers. *The Review of Economic Studies*, 73(3):793–821.

Reis, R. (2020). Comment on "Imperfect Expectations: Theory and Evidence". *NBER Macroeconomics Annual*, 35.

Reis, R. and Watson, M. W. (2010). Relative goods' prices, pure inflation, and the phillips correlation. *American Economic Journal: Macroeconomics*, 2(3):128–57.

Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *The quarterly journal of economics*, 100(4):1169–1189.

Rubbo, E. (2020). Networks, phillips curves, and monetary policy. *manuscript, Harvard University*.

Ruge-Mucia, F. J. and Wolman, A. L. (2022). Relative price shocks and inflation.

Ryngaert, J. (2018). What Do (and Don't) Forecasters Know about U.S. Inflation?

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Smets, F. and Wouters, R. (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *American economic review*, 97(3):586–606.

Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *Unpublished manuscript, Harvard University*.

Wolman, A. L. (2011). The optimal rate of inflation with trending relative prices. *Journal of money, credit and banking*, 43(2-3):355–384.

Woodford, M. (2013). Macroeconomic analysis without the rational expectations hypothesis.

Woodford, M. (2019). Monetary policy analysis when planning horizons are finite. *NBER macroeconomics annual*, 33(1):1–50.

Zhou, H. (2022). Open economy, redistribution, and the aggregate impact of external shocks. *Redistribution, and the Aggregate Impact of External Shocks (February 6, 2022)*.