

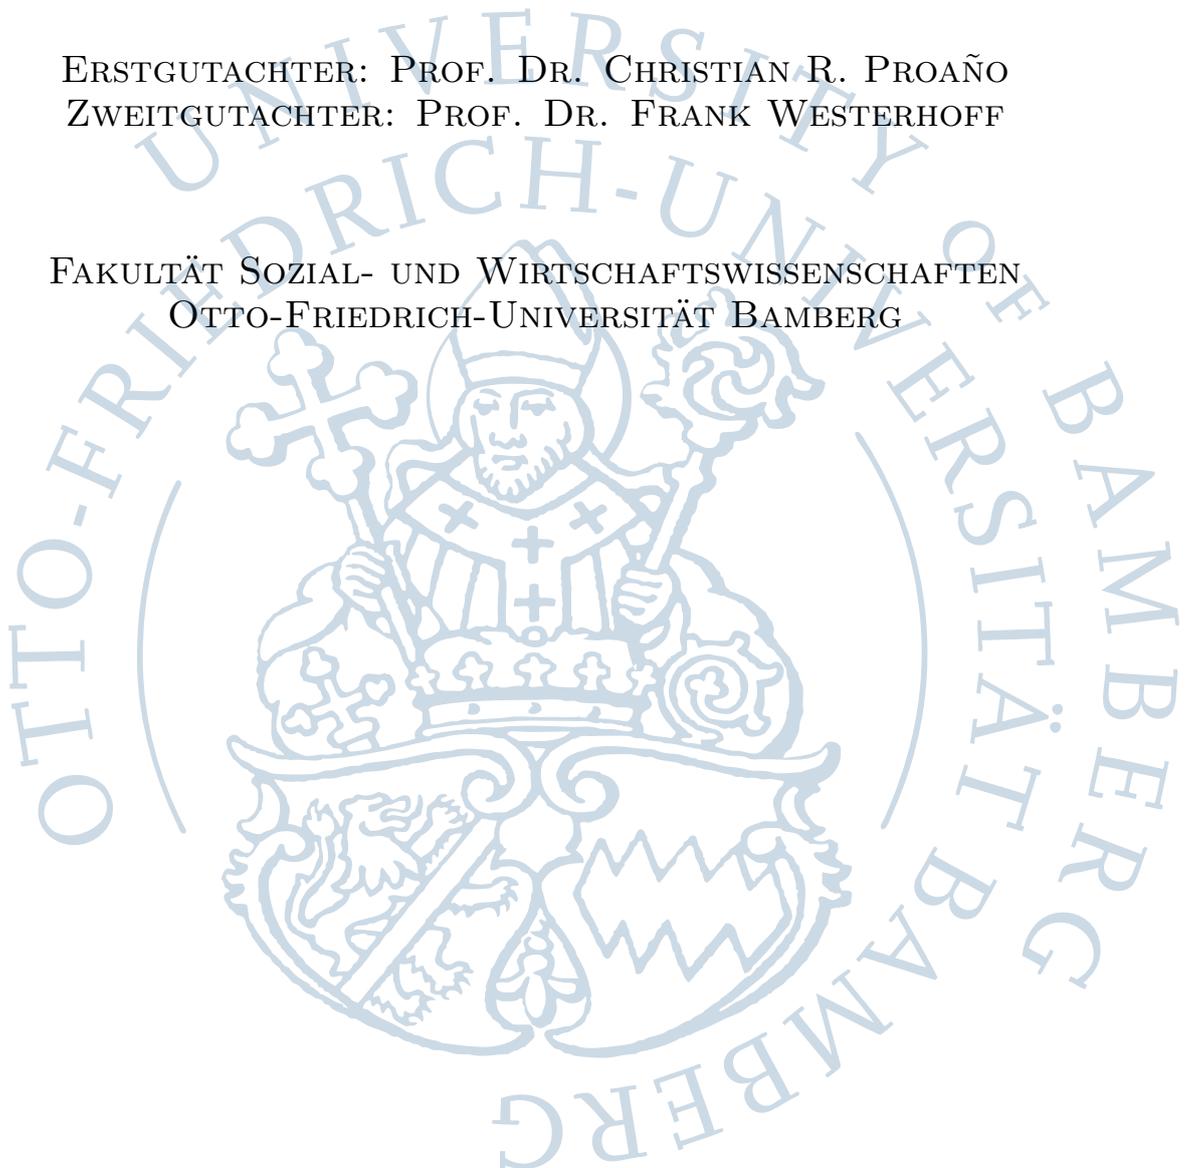
INAUGURALDISSERTATION

ESSAYS ON BOUNDEDLY RATIONAL EXPECTATIONS
IN MACROECONOMICS

AUTOR: TIM A. HAGENHOFF
TAG DER DISPUTATION: 24.03.2021

ERSTGUTACHTER: PROF. DR. CHRISTIAN R. PROAÑO
ZWEITGUTACHTER: PROF. DR. FRANK WESTERHOFF

FAKULTÄT SOZIAL- UND WIRTSCHAFTSWISSENSCHAFTEN
OTTO-FRIEDRICH-UNIVERSITÄT BAMBERG



ERSTGUTACHTER: Prof. Dr. Christian R. Proaño
(*Otto-Friedrich-Universität Bamberg*)
ZWEITGUTACHTER: Prof. Dr. Frank Westerhoff
(*Otto-Friedrich-Universität Bamberg*)
DRITTGUTACHTER: Prof. Mishael Milaković, PhD
(*Otto-Friedrich-Universität Bamberg*)

Dieses Werk ist als freie Onlineversion über das Forschungsinformationssystem (FIS; <https://fis.uni-bamberg.de>) der Universität Bamberg erreichbar. Das Werk steht unter der CC-Lizenz CC-BY.



URN: urn:nbn:de:bvb:473-irb-498030
DOI: <https://doi.org/10.20378/irb-49803>

PK045
PROMOTIONSKOLLEG

Gefördert durch die

Hans Böckler
Stiftung 

Mitbestimmung · Forschung · Stipendien

To Mechthild, Andreas and Jan.

Acknowledgments

First of all, I would like to thank my supervisor Christian Proaño for his continuous support over the last three years which involved writing recommendations and reports for the Hans Böckler Foundation and inviting me to give a talk at the *International Conference on Economics* in Quito. I am, further, grateful for his very valuable comments on each of the chapters of this dissertation.

There is one person that had by far the biggest influence on this dissertation and deserves special recognition: Joep Lustenhouwer - not only because he co-authored two of the four main chapters but also because he can rightfully be called my mentor during the time of my dissertation. Judging by his publication record, it is almost needless to say that Joep is an extraordinarily smart and careful thinker which has certainly benefited the quality of both co-authored chapters. Working with Joep has improved my structural thinking substantially which also helped me to write my single-authored chapters. I also would like to thank Joep for his company while drinking wine and eating fancy meals in Barcelona and having daily Jugós in Ecuador. I am deeply thankful to Joep Lustenhouwer for being my mentor and even more for being such a good friend.

I am also indebted to a few people that gave me valuable advice on individual chapters. Emanuel Gasteiger was the supervisor of my master thesis at the FU Berlin which laid the groundwork for Chapter 2. Emanuel, therefore, has his share in the success of my master thesis and also to some degree in Chapter 2. Further, I did most of the work for Chapter 4 during my three-month internship at the research center of the Deutsche Bundesbank. Jochen Mankert supervised my internship and provided me with excellent help on my project. Also, Tomasz Makarewicz gave me crucial advice on how to deal with some technical difficulties (he was also a very pleasant roommate at conferences in Milan and Ottawa).

Moreover, I am indebted to Frank Westerhoff, who is my second supervisor, for remarks that at least two times led me to consider fundamental changes or improvements of my chapters.

I very much enjoyed the (more or less) daily espresso with my friends and colleagues Michael Neuner and Juan Carlos Peña. I am, further, deeply grateful to Martha Düker, Naira Kotb, Ivonne Schwartz, Carolin Martin, Noemi Schmitt, Benjamin Lojak, Arne Lauber, Jan Schulz, Mishael Milaković, Andrea Gurgone, Alexander Hempfing and all other colleagues that I meet from time to time at the department for the very pleasant and welcoming environment.

Finally, special thanks go to the Hans Böckler Foundation, especially to Jens Becker

and Patrick Tschirner, for providing me financial and non-financial support throughout my dissertation. I am equally grateful to Wiebke Störmann who wrote several reports on my behalf for the foundation.

Co-author

Chapters 2 and 4 are written by myself. Joep Lustenhouwer co-authored Chapters 3 and 5 which is also indicated at the beginning of each of these two chapters.

Joep and I together developed the idea and story of “*The rationality bias*” (Chapter 3) which have both changed several times as we explored the model. The derivation of the model was primarily done by myself while Joep took the equations to the computer and solved them. The first draft was written by me while the revisions were done by both Joep and me.

The initial idea and development of the model in “*The role of stickiness, extrapolation and past consensus forecasts in macroeconomic expectations*” (Chapter 5) primarily came from Joep. For the most part, my role was being the econometrician in this project, i.e. the programming, data manipulation, estimation and robustness analysis using Python were my obligations. The current version of the paper was written both by Joep and me.

Contents

1	Introduction	1
2	Heterogeneous Expectations: Monetary Policy	9
2.1	Introduction	9
2.2	Model	13
2.3	An optimal interest rate rule from a timeless perspective	17
2.3.1	Loss functions	17
2.3.2	An optimal interest rate rule under timeless commitment	19
2.4	Impulse Responses	22
2.5	Welfare evaluation	24
2.5.1	Optimal vs. inflation-targeting rule	24
2.5.2	Inflation and welfare vs. consumption dispersion	28
2.6	Conclusion	30
	Appendix 2.A Implementation under the conventional inflation-targeting objective	31
	Appendix 2.B Optimal monetary policy	31
	2.B.1 Rewriting the model-consistent loss function	31
	2.B.2 Optimal interest rate rule	32
	2.B.3 Taking rational expectations as given	36
	2.B.4 Tables	38
	Appendix 2.C Implementation with model-consistent individual consump- tion expectations	40
3	The Rationality Bias	43
3.1	Introduction	43
3.2	Model	46
3.2.1	The non-linear model	47
3.2.2	Expectations	49

3.2.3	The log-linear economy	49
3.3	The components of the rationality bias	53
3.3.1	Decomposing the rationality bias	54
3.3.2	Individual components	54
3.3.3	Illustrating the components of the Rationality Bias	56
3.4	The magnitude of the rationality bias	63
3.4.1	Eliminating the rationality bias	63
3.4.2	Minimizing the rationality bias and welfare	65
3.5	Conclusion	66
Appendix 3.A	Microfoundations	67
3.A.1	Consumption of rational households	67
3.A.2	Consumption of boundedly rational households	69
3.A.3	Aggregation of individual consumption decisions	71
3.A.4	Phillips-Curve	72
3.A.5	Steady state	73
Appendix 3.B	Impulse responses for different fractions of rational agents . .	74
Appendix 3.C	Derivation of second-order approximated utilitarian welfare loss	77
4	Heterogeneous Expectations: Labor Markets	81
4.1	Introduction	81
4.2	Model	84
4.2.1	Households	84
4.2.2	Production	87
4.2.3	Wage Bargaining	89
4.2.4	Monetary Policy and Resource Constraint	92
4.2.5	The Linearized Model	92
4.3	Results	94
4.3.1	Model Dynamics	95
4.3.2	Comparison with a more “sclerotic” labor market	98
4.4	Conclusion	100
Appendix 4.A	Steady State Values	102
5	Biases in Survey Expectations	103
5.1	Introduction	103
5.2	A behavioral model of expectation formation	108
5.2.1	Stickiness	108

CONTENTS

5.2.2	Forward-looking extrapolation	109
5.2.3	Consensus forecast	110
5.2.4	Combining the three behavioral biases	111
5.3	Empirical model and estimation method	112
5.4	Confronting the model with survey expectations	116
5.4.1	Fitting forecasting revisions	117
5.4.2	Performing a test	118
5.4.3	Forecasting horizon	119
5.5	Extensions	121
5.5.1	Consensus forecast from Livingston survey	122
5.5.2	Forecasts from Livingston survey	124
5.5.3	An alternative deviation from rationality	126
5.6	Discussion: the (ir)rationality of using the lagged consensus forecast . .	128
5.7	Conclusion	130
	Appendix 5.A Variable Construction	132
5.A.1	SPF	132
5.A.2	Livingston	133
	Appendix 5.B News and Consensus Nowcast Regressions	133
	Appendix 5.C Robustness checks	133
5.C.1	Final vintages	133
5.C.2	Time subsamples	135
	Appendix 5.D Additional Tables	136

Chapter 1

Introduction

Expectations matter. How much to consume or save, what price to set, and whether to hire or fire workers are just some of the fundamental decisions underlying macroeconomic dynamics that hinge upon agents' expectations of the future.

– Coibion and Gorodnichenko (2015)

Introduction

The above quote is neither very old nor are its originators famous historical figures, not yet at least. However, this quote very precisely summarizes the reasoning behind the recent explosion of literature in the field of macroeconomic expectations. In essence, expectations are thought to be causal for the dynamic behavior of the economy because individual decisions depend on these expectations. The reason is that economic decisions, in general, have an *intertemporal* dimension where the future - which economic actors will form expectations about - is inherently *uncertain*. Whether to buy a car, build a house, invest in new machines or start an entire new firm, and that perhaps by taking a loan, crucially depends on the personal economic outlook. This personal economic outlook then depends, besides idiosyncratic factors, also on how the economy as a whole is expected to perform. Some questions that one may ask oneself are: *Will I be employed 'tomorrow' (how high is the unemployment rate going to be) and what will my income be? Is the economy going to grow and by how much? How will my profits develop? What will the inflation rate and, therefore, my real interest rate be?* It becomes obvious that economic decisions may drastically vary with respect to the potential answers to such questions. Also, these answers may depend on government policy and the announced trajectory thereof. Once individual actors decide on their

actions, these actions will aggregate into macroeconomic outcomes. Understanding expectations will, therefore, provide a better assessment of how economic actors will change their behavior in response to certain policies or unexpected events. In the end, this will help to explain the dynamics of the economy at the aggregate level.

Already in 1933, Irving Fisher proposed that a sudden change in expectations might have been the starting point of the Great Depression initializing what he calls a debt-deflation (Fisher, 1933). In the same year of Fisher's publication, the unemployment rate in the US rose to a staggering 24.9%. Around 70 years later, over-optimistic price expectations fueled the price boom of real estate and asset-backed securities with questionable underlyings, especially in the US. The financial crisis began when a dramatic *revision of expectations* in the light of new information led to a sharp drop in those asset prices as investors shortened their positions, turning markets from bullish to strongly bearish. This asset revaluation shock propagated through the network of balance sheets causing widespread liquidity and solvency problems which, consequently, brought financial systems across the world to the verge of collapse. Today the COVID-19 pandemic causes unprecedented economic turmoil, for instance, in the US where its GDP dropped by an annualized 31.7% in quarter two of 2020. Additionally, the unemployment rate peaked at 14.7% in April of 2020 in the US. In light of such extreme figures, expectations may very well have a detrimental effect on the economic development as this crisis progresses while economic uncertainty is high. Indeed, it was measured that household expectations react quite strongly in response to the lockdown measures: Coibion et al. (2020) find that households that experienced lockdowns earlier than others also expected significantly higher unemployment rates, lower inflation and reported higher uncertainty. I should emphasize that I do not seek to explain certain historical events or the current crisis with this dissertation. However, such events surely highlight the importance of understanding expectations and their effects on macroeconomic outcomes.

Since the early 1980s, most dynamic macroeconomic models are derived from individual intertemporal decision problems (see for instance Kydland and Prescott, 1982). Hence, expectations have been central to macroeconomic modeling for quite some time now. The predominant assumption on expectations has been (and still seems to be, although this can be debated) the rational expectations hypothesis, first discussed by Muth (1961). This hypothesis assumes that *all* agents populating the economy know the underlying structural equations including the probability distributions of random disturbances. Further, agents endowed with rational expectations can solve the structural equations and use the solution to make unbiased forecasts about all relevant economic variables. The beauty of this assumption lies in its simplicity: it requires

no additional free parameters and the value that an expectation will take is unambiguous. Of course, rational expectations cannot be true literally but the hope was that it is good enough as an approximation - today we know almost certainly that it is not. However, deviating from rational expectations is not an easy task as there are numerous possible ways in which expectations can be non-rational.

As I will discuss evidence against (homogeneous) rational expectations and proposed alternatives several times throughout this dissertation, I will give only a very brief and selective overview here. Already Lovell (1986, p.122) states:

“[...] it seems to me that the weight of empirical evidence is sufficiently strong to compel us to suspend belief in the hypothesis of rational expectations [...]”.

Somewhat later, Ehrbeck and Waldmann (1996) find that deviations from rational expectations among professional forecasters are not merely because of strategic concerns but consistent with certain behavioral explanations. Today, one strand of the literature finds that expectations inefficiently react to new incoming information, both at the aggregate level (Coibion and Gorodnichenko, 2015) and at the individual level (Bordalo et al., 2018a; Fuhrer, 2018). This can be explained by certain biases and/or informational frictions. However, the search for the answer to how exactly expectations react to news is ongoing. One recurrent finding is that survey expectations are inherently persistent, making typical modeling devices such as price indexation and investment adjustment costs to generate observed macroeconomic persistence obsolete (Fuhrer, 2017).

Further, the literature on ‘learning’ proposes that expectations are to some extent backward-looking, i.e. expectations are based on past realizations of economic variables. In the case of homogeneous adaptive learning, agents behave as econometricians that update the parameters of their perceived law of motion when realizations become observable with a lag (Evans and Honkapohja, 2012). Therefore, adaptive learning is entirely backward-looking, introducing persistence in expectations but leaving little room for news to play a role. Another strand of the literature proposes that different economic actors may follow different expectations rules (for an overview, see Branch and McGough, 2018). Mankiw et al. (2003) were one of the first to highlight the disagreement in expectations about future inflation, both among professional forecasters and consumers. One of the many heterogeneous expectations models assumes that there is a certain fraction of the population that has rational expectations while the rest extrapolates the last realization of the variable being forecast into the future (Branch and McGough, 2009). While such a model may be very stylized, it intuitively leads

to aggregate expectations that are persistent and also allows for news to play a role via the rational expectations component. Additionally, this model introduces disagreement of expectations, and thereby allows for the analysis of dispersion effects, and also facilitates analytical tractability due to its simplicity. Some empirical evidence shows that a heterogeneous expectations model, including a certain fraction of relatively sophisticated actors while the rest uses a simple backward-looking heuristic, explains expectations data well (Branch, 2004; Kryvtsov and Petersen, 2019; Cornea-Madeira et al., 2019). For these reasons, this model of heterogeneous expectations will be basis for the theoretical consideration in Chapters 2, 3 and 4.

Monetary policy must take into account and also actively manage private-sector expectations. In an economic downturn, a discretionary one-time drop of the nominal interest rate may not induce much reaction *if* economic actors' expectations do not change. If, however, the interest-rate-setting behavior of the central bank is rule-based and shows some persistence, *forward-looking* actors will expect the interest rate to be lower for some time. Alternatively, the central bank may actively announce to keep interest rates below or equal to a specific value for at least a certain period of time, so-called forward guidance. These policies tend to flatten the yield curve so that actual financing costs are lower and saving less attractive. This may increase production through higher demand for consumption and investment goods, leading to an even more favorable outlook that reinforces the recovery. In the end, this will help the central bank to reach its inflation target.

Having emphasized the importance of expectations in the transmission of monetary policy, it becomes apparent that the design of the policy depends on how expectations are formed. Thus, the first two chapters seek to contribute to the field of monetary policy under heterogeneous, and to some degree, boundedly rational expectations (and boundedly rational decision making in the case of Chapter 3). In particular, Chapter 2 focuses on how to set interest rates to achieve the highest possible welfare given a stylized model of heterogeneous expectations. Chapter 3 further deals with the effects of monetary policy on the boundedly rational biases and the associated dispersion effects under heterogeneity.

Moreover, news articles about the economy give the impression that certain figures are more important to readers than others. When it comes to workers, who are undoubtedly the majority of “agents” in any actual economy, the unemployment rate or the labor market, in general, seems to be of utmost importance - “*Will I be employed ‘tomorrow’ (how high is the unemployment rate going to be) and what will my income be?*”. Several factors determine the unemployment rate. Besides general macroeconomic conditions, wages are one of these factors. Real-world wages are usu-

ally negotiated between workers (or unions) and firms rather than “dictated” by the market, as very simple models of the labor market suggest. Such negotiations involve expectations about the future: what does a worker (or union) believe a firm will gain by employing and keeping this worker for some time, and vice versa? These beliefs depend on future expected productivity, labor market and also macroeconomic conditions. The negotiation outcomes then feed into the unemployment rate and overall macroeconomic performance. Therefore, Chapter 4 introduces the aforementioned model of heterogeneous (household) expectations not only to the consumption decisions of households but also to the wage negotiations on the labor market. This chapter deals with the implications of heterogeneous household expectations for labor market outcomes, dispersion effects and macroeconomic dynamics.

A model of heterogeneous expectations with a certain degree of backward-lookingness undoubtedly improves the fit to the data by introducing persistence. It also gives valuable insights into the interaction between expectations, government policies and the macroeconomy. However, it might be that such a model is too stylized to be interpreted literally and that it rather represents a neat device to explain some important regularities of economic dynamics, such as persistence, in a more plausible and intuitive way. Some deeper questions about expectations are: are expectations extrinsically or intrinsically persistent and how can this persistence be squared with some findings of individual over-reaction, as for instance in Bordalo et al. (2018a)? If any attempt of building an own rational expectation fails, what can be a smart way to deal with it? Therefore, Chapter 5 presents a deeper dive into expectation formation itself and the underlying biases by looking at data of the Survey of Professional Forecasters and the Livingston Survey.

Overview

Having introduced the aim of this dissertation and having outlined the basic questions discussed in each of the four chapters, I now want to provide a brief and slightly more technical overview of the findings presented in these chapters.

In Chapter 2, I discuss the implications of a simple model of heterogeneous expectations for the conduct of monetary policy. In particular, I implement optimal rule-based monetary policy under heterogeneous expectations by deriving an explicit interest rate rule. This implementation requires a more detailed discussion of the consumption decisions of the agents populating the model world. I find that the agents usually defined as “rational” in this specific environment are not sophisticated enough to have model-consistent individual consumption expectations, as assumed in Di Bartolomeo et al.

(2016), even though they forecast aggregate variables correctly on average. This result stems from a particular assumption on higher-order beliefs introduced by Branch and McGough (2009). Further, I show that the optimal interest rate rule yields substantial welfare gains compared to a rule that is derived from a conventional inflation-targeting objective as in Gasteiger (2014). The implementation of the non-optimal inflation-targeting rule already results in welfare losses equivalent to a 14.5 percent loss of steady-state consumption relative to the optimal interest rate rule when only ten percent of the population form (naive) backward-looking expectations. Moreover, the presence of heterogeneous expectations requires the central bank to be extraordinarily hawkish with respect to inflation to achieve optimality. However, consumption dispersion increases with the central bank's aggressiveness towards inflation.

In Chapter 3, Joep Lustenhouwer and I analyze the differences in consumption and wealth in a New Keynesian model with rational households and households that are less sophisticated in more detail. In contrast to Chapter 2, rational households are indeed rational in the conventional sense. Therefore, we can define the rationality bias of boundedly rational households as a deviation from the fully rational benchmark. The goal is to provide a model with a rational benchmark even when other, non-rational households are present. The rationality bias of boundedly rational households systematically explains the emergence of consumption and wealth heterogeneity. This rationality bias can be decomposed into three individual components: the consumption expectation bias, the real interest rate bias and the preference shock expectation bias. For certain specifications of monetary policy, these different components exactly offset each other, so that the rationality bias is eliminated. This might, however, not be desirable from a welfare perspective, since eliminating the rationality bias comes along with high inflation volatility.

In Chapter 4, I build a stylized macroeconomic model with heterogeneous household expectations and labor market frictions. This combination allows diverse expectations to have first-order effects on wages and other labor market variables, in contrast to heterogeneous expectations models that assume frictionless labor markets. Productivity shocks produce persistent hump-shaped responses of relevant labor market variables, pro-cyclical labor market tightness as well as a Beveridge curve. As households of different types adjust their expectations differently in response to shocks, the model predicts distributional consequences through the channel of heterogeneous expectations. The aggregate volatility and cross-sectional dispersion of wages, unemployment and job-finding rates, thereby, depend on the “fluidity” of the labor market.

Finally, in Chapter 5, Joep Lustenhouwer and I utilize data on expectations provided by the Survey of Professional Forecasters and the Livingston survey. We propose

a simple model of expectation formation with three distinct deviations from fully rational expectations. In particular, forecasters' expectations are sticky (intrinsically persistent), extrapolate the most recent news about the current period, and depend on the lagged consensus forecast about the period being forecast. We find that all three biases are present in the data, and that their magnitudes depend on the forecasting horizon. Moreover, in an over-identified econometric specification, we find that the restriction on coefficients implied by our model is always close to being satisfied and in most cases not rejected. We also stress the point that using the past consensus forecast to form expectations is a rather smart thing to do *if* cognitive limitations and biases cause any attempt to build an own fully rational forecast to fail.

Chapter 2

Optimal monetary policy, heterogeneous expectations and consumption dispersion

2.1 Introduction

The New Keynesian model emphasizes the ability of monetary policy to stabilize the macroeconomy by managing agents' expectations. Usually, optimal monetary policy is studied within a framework that assumes all agents to form their expectations rationally (Clarida et al., 1999; Woodford, 1999; McCallum, 1999). However, econometric studies based on inflation expectation surveys show that the data favors heterogeneous expectations with a certain degree of bounded rationality (Branch, 2004, 2007; Pfajfar and Santoro, 2010; Cornea-Madeira et al., 2019). For instance, Branch (2004) and Cornea-Madeira et al. (2019) find evidence for the existence of more sophisticated agents that employ a VAR-heuristic alongside agents that form (naive) backward-looking expectations. Further, Cole and Milani (2019) show that cross-equation restrictions imposed by New Keynesian models are heavily rejected under the assumption of homogeneous rational expectations when trying to match actual survey expectation data. However, these restrictions prove valuable under heterogeneous expectations, indicating that the main source of misspecification in New Keynesian models stems from misspecified (rational) expectations. Additionally, heterogeneity in forecasting schemes is confirmed by examining the expectation formation of actual human subjects in a series of laboratory experiments (Hommes, 2011; Assenza et al., 2019; Pfajfar and Žakelj, 2018). Starting from these observations, several New Keynesian models were designed that include heterogeneous expectations (Branch and McGough, 2009, 2010; De Grauwe,

2011; Massaro, 2013; Hommes and Lustenhouwer, 2019).¹

What are the implications of heterogeneous expectations for monetary policy and, in particular, how should central banks set interest rates optimally given this knowledge? To answer this question I derive an optimal interest rate rule under timeless commitment² based on the Branch and McGough (2009) framework and a model-consistent welfare criterion following Di Bartolomeo et al. (2016). One feature of this interest rate rule is that it is extraordinarily hawkish with respect to inflation. This theoretical finding underpins the experimental finding of Assenza et al. (2019), where a central bank that is quite aggressive with respect to inflation yields more desirable aggregate outcomes in a world with heterogeneous expectations. However, this chapter additionally suggests that there may be undesirable distributional consequences as consumption dispersion increases with the aggressiveness of monetary policy towards inflation.

To describe the micro level explicitly, I derive consumption Euler equations that adequately account for the assumption of Branch and McGough (2009) about higher-order beliefs. More specifically, the underlying model of this chapter incorporates two types of agents. The more sophisticated agents, that I call “rational forecasters”, are able to forecast *aggregate* variables consistent with the model predictions but are not smart enough to understand the micro level fully. In contrast, boundedly rational forecasters use a simple backward-looking heuristic instead for forecasting. Such backward-looking heuristics are broadly consistent with evidence from laboratory experiments (Hommes, 2011; Pfajfar and Žakelj, 2018; Assenza et al., 2019).

While optimal monetary policy under homogeneous rational expectations is well known and extensively studied, the strand of literature dealing with monetary policy under heterogeneous expectations is rather new. Recent advances in the literature are made by Gasteiger (2014, 2018), Di Bartolomeo et al. (2016) and Beqiraj et al. (2019). Beqiraj et al. (2019) investigate optimal discretionary monetary policy under heterogeneous expectations based on the framework developed in Massaro (2013) that includes agents that forecast over all future periods up to infinity. On the other hand, Gasteiger (2014), Gasteiger (2018) and Di Bartolomeo et al. (2016) follow the Euler-equation-learning approach of Branch and McGough (2009). Gasteiger (2014) and Gasteiger (2018) explore interest rate rules derived from an ad-hoc inflation-targeting

¹Another model with heterogeneous expectations can be found in Chapter 3.

²The timeless commitment approach assumes that the optimal commitment policy was implemented in the distant past so as to omit the first period’s optimality condition that is the same as under discretion (see for instance Woodford (2003) and Woodford (2010)). The problem of the latter is that it renders the policy time-inconsistent. Hence, the drop of the initial period’s optimality condition solves this problem.

objective while Di Bartolomeo et al. (2016) provide an extension based on a model-consistent central bank objective.

Although Di Bartolomeo et al. (2016)'s solution algorithm implies that the central bank employs expectations-based interest rate rules, they do not derive them analytically. Thus, the literature has so far not provided an optimal interest rate rule under heterogeneous expectations based on the Branch and McGough (2009) model. (Optimal) Interest rate rules are useful as they allow, in comprehensible way, to identify which variables determine the interest rate and to which degree, especially with varying degrees of heterogeneity in expectations, they should do so. In this chapter, I will focus on the case where the central bank is able to commit to its policy from timeless perspective (Blake et al., 2001; Woodford, 2003) as it elegantly addresses the issue of time-inconsistency under commitment and further allows the central bank to efficiently target private sector (rational) expectations.

Additionally, I explore the role of the higher-order beliefs assumption of Branch and McGough (2009) for agents' individual consumption decisions. The authors impose this assumption as it is *necessary* for aggregation, i.e. to arrive at the same functional forms as in the canonical New Keynesian model but where the rational expectations operator is simply replaced by a weighted sum different expectation operators. Further, consumption decisions have to be made explicit as the central bank's objective function introduced by Di Bartolomeo et al. (2016) depends on consumption dispersion. This approach allows me to clarify the properties of the expectations operator, E_t^R , of "rational" forecasters in Branch and McGough (2009). It is assumed that rational forecasters, by using E_t^R , predict aggregate variables consistent with the model predictions which can, however, not be the case for expectations about the distribution of individual consumption. This is a straightforward consequence of the higher-order beliefs assumption which puts a particular (non-rational) structure on an agent's belief about other agents' individual expectations. In particular, all agents believe that all other agents form the same expectations about their individual consumption as they themselves do. Hence, even "rational forecasters" are not smart enough to sophisticatedly forecast individual consumption. The final consumption equations only depend on aggregate variables and can, therefore, be used to substitute for individual consumption in the optimal interest rate rule.

However, if an Euler equation with model-consistent individual consumption expectations as in Di Bartolomeo et al. (2016) is applied, the first-order conditions of the central bank problem under commitment can only be reduced to a second-order difference equation in one of the Lagrange-multipliers to which the solution is fairly complicated and exponentially depends on time. Consequently, an interpretable in-

terest rate rule under commitment in this case cannot be derived. Further, it would not be possible to substitute for individual consumption so that practical implementation would require consumption of rational and boundedly rational forecasters to be observable. However, applying the consumption equation that appropriately accounts for the higher-order beliefs assumption makes the derivation of a meaningful interest rate rule under commitment possible.

Moreover, I compare the optimal interest rate rule to a micro-founded version of the policy rule in Gasteiger (2014). As already indicated, the author derives an interest rate rule under timeless commitment that recognizes the heterogeneity in expectations in the private sector equations but is optimized under a conventional ad-hoc inflation-targeting objective. The resulting interest rate rule is sub-optimal but also much simpler than the rule derived in this chapter. While it is straight-forward that the optimal rule must yield lower welfare losses than the non-optimal rule derived from the inflation-targeting objective, it is not clear *by how much*.

The welfare analysis shows that the optimal interest rate rule generates substantial welfare gains, given a non-autocorrelated one-standard-deviation cost-push shock. In that case, the implementation of the non-optimal inflation-targeting rule already requires an increase of 14.5 percent of steady-state consumption to compensate for the higher welfare losses relative to the optimal interest rate rule when only ten percent of the population form (naive) backward-looking expectations. The welfare gains of the optimal interest rate rule crucially depend on the relative fraction of agent types. The optimal interest rate rule performs relatively better the higher the fraction of boundedly rational forecasters.

Finally, I find that consumption dispersion is not necessarily lower under the optimal rule compared the non-optimal inflation-targeting rule, even though the former explicitly incorporates consumption heterogeneity as opposed to the latter. This is because consumption dispersion increases with the central bank's aggressiveness towards inflation, as rational and boundedly rational forecasters' consumption decisions become more unequal for stronger reactions of the policy rate. As a consequence, there is also a (local) trade-off between minimizing welfare losses and reducing consumption dispersion.

The remainder of the chapter is organized as follows. The underlying model including the modified consumption rules are presented in Section 2.2. The optimal interest rate rule under heterogeneous expectations is derived in the subsequent section. Section 2.4 shows the impulse responses under optimal monetary policy with an emphasis on the different consumption paths of rational and boundedly rational forecasters, followed by the welfare analysis in Section 2.5. Section 2.6 concludes.

2.2 Model

In this section, I introduce the underlying model and derive individual consumption decisions that incorporate the higher-order beliefs assumption of Branch and McGough (2009). The economy is assumed to be populated by an exogenous fraction α of rational forecasters (R) which have rational (model-consistent) expectations with respect to *aggregate* variables and a fraction $1 - \alpha$ of boundedly rational forecasters (B) that employ a simple backward-looking heuristic. The general forecasting rule of boundedly rational forecasters takes the form of $E_t^B x_{t+1} = \theta^2 x_{t-1}$ for some variable x while rational forecasters simply use the expected value, $E_t^R = E_t$, for forecasting aggregate variables. Backward-looking expectations for $\theta < 1$ are called steady-state-reverting, for $\theta = 1$ naive and for $\theta > 1$ trend-setting. Steady-state-reverting expectations constitute a stabilizing force while trend-setting expectations imply a further amplification of macroeconomic variables.

Assuming perfect consumption insurance *within* each of the two agent groups the model can be expressed in terms of two representative agents (RA). This approach, therefore, abstracts from further individual characteristics in order to isolate the effects of heterogeneous expectations on individual consumption (dispersion) and aggregate variables, as well as their implications for monetary policy. Both RA's aim at maximizing their individual expected discounted lifetime utility $E_0^\tau \sum_{t=0}^{\infty} \beta^t U_t$ given their subjective expectations E_0^τ with $\tau \in \{R, B\}$. However, as in Branch and McGough (2009) agents follow Euler-equation learning, i.e. they disregard their intertemporal budget constraint as an optimality condition and solely base their consumption decision on the variational intuition of the consumption Euler equation.

The period utility function is of CES-form and is given by

$$U_t^\tau = \frac{(C_t^\tau)^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \frac{(Y_t^\tau)^{1+\eta}}{1+\eta} \quad (2.1)$$

with C_t^τ being consumption of the RA of type τ , Y_t^τ the output that each type- τ RA produces, $\frac{1}{\sigma}$ the coefficient of relative risk aversion and η the elasticity of marginal dis-utility of producing output. Agents must satisfy their real budget constraint

$$C_t^\tau + B_t^\tau = \frac{1+i_{t-1}}{\Pi_t} B_{t-1}^\tau + \Psi_t^\tau \quad (2.2)$$

with B_t^τ being real bond holdings, i_{t-1} the nominal interest rate in $t-1$, Π_t gross inflation and Ψ_t^τ real income of agent τ .

All agents in this economy are assumed to believe that all other agents will form the

same expectations as they themselves do. Branch and McGough (2009) explicitly emphasize that this assumption is *necessary* for aggregation. I explicitly show below that this assumption implies that “rational” forecasters are not fully rational in the conventional sense and, therefore, cannot have model-consistent individual consumption expectations. Assuming rational forecasters to possess rational individual consumption expectation, as in Di Bartolomeo et al. (2016), *implies too much rationality* to be consistent with the *necessary* higher-order beliefs assumption. Incorporating the higher-order beliefs assumption into the consumption decisions of agents is, further, crucial as an interpretable interest rate rule under timeless commitment cannot be derived otherwise (see Appendix 2.C).

I will show later on that the central bank’s welfare criterion can be re-written using market clearing in a way that it only depends on the consumption decision of rational forecasters. Hence, for now I focus on the consumption Euler equation for $\tau = R$ which is given by

$$(C_t^R)^{-\frac{1}{\sigma}} = \beta E_t^R \left[(C_{t+1}^R)^{-\frac{1}{\sigma}} \frac{1 + i_t}{\Pi_{t+1}} \right]. \quad (2.3)$$

Log-linearizing (2.3) gives

$$c_t^R = E_t^R c_{t+1}^R - \sigma(i_t - E_t^R \pi_{t+1}) \quad (2.4)$$

where lower case letters indicate log-deviations from the zero-inflation steady state. Forward iteration yields

$$c_t^R = E_t^R c_\infty^R - \sigma \sum_{k=0}^{\infty} (i_{t+k} - E_t^R \pi_{t+k+1}). \quad (2.5)$$

Rational forecasters assume that the Euler equation of boundedly rational forecasters will also be satisfied and that market clearing $y_t = \alpha c_t^R + (1 - \alpha)c_t^B$ holds. Writing market clearing one period forward and inserting equation (2.5), and equivalently the forward-iterated consumption Euler equation for boundedly rational forecasters, yields

$$\begin{aligned} E_t^R y_{t+1} = E_t^R & \left[\alpha (E_t^R c_\infty^R - \sigma E_{t+1}^R \sum_{k=1}^{\infty} (i_{t+k} - \pi_{t+k+1})) \right. \\ & \left. + (1 - \alpha) (E_t^B c_\infty^B - \sigma E_{t+1}^B \sum_{k=1}^{\infty} (i_{t+k} - \pi_{t+k+1})) \right]. \end{aligned} \quad (2.6)$$

Note that (2.6) contains higher-order beliefs, i.e. beliefs of rational forecasters E_t^R about the beliefs of boundedly rational forecasters E_t^B and E_{t+1}^B . In order to arrive at the IS

curve that has the same functional form as in the model under homogeneous rational expectations, Branch and McGough (2009) impose a specific (non-rational) structure on higher-order beliefs. In the context of this chapter, the assumption states that agents' believe that all other agents will forecast their individual consumption in the same way they themselves do. Mathematically and in the case of rational forecasters: $E_t^R E_{t+k}^B c_{t+l} = E_t^R c_{t+l}$ with $l > k$. Hence, boundedly rational expectations just drop out under this assumption. Further, note that making an alternative assumption, e.g. allowing rational forecasters to be fully rational, would result in a different system of aggregate equations (see Chapter 3).

Using the higher-order beliefs assumption and the law of iterated expectations at the individual level yields

$$E_t^R y_{t+1} = E_t^R y_\infty - \sigma \sum_{k=1}^{\infty} (i_{t+k} - E_t^R \pi_{t+k+1}). \quad (2.7)$$

It becomes obvious that (2.7) cannot hold under conventional rational expectations, i.e. when $E_t^R = E_t$ would hold, as boundedly rational expectations, E_t^B and E_{t+1}^B , would *not drop out* and thus show up in (2.7). Mathematically written: $E_t E_t^B = E_t^B$ and $E_t E_{t+1}^B = E_{t+1}^B$, which would contradict the higher-order beliefs assumption of Branch and McGough (2009). Equation (2.7) would only hold under conventional rational expectations when boundedly rational forecasters were absent, i.e. under *homogeneous rational expectations*. In this case (2.6) would collapse to (2.7) without any further assumption.

The main crucial point that follows from this consideration is that rational forecasters in this model are not smart enough to have model-consistent expectations with respect to the distribution of individual consumption *when there is heterogeneity*. Thus, assuming model-consistent individual consumption expectations in the Euler equation of rational forecasters, as in Di Bartolomeo et al. (2016), is inconsistent with the underlying framework.

Using (2.7) to replace the infinite sum in (2.5), one obtains

$$c_t^R = E_t^R y_{t+1} + E_t^R (c_\infty^R - y_\infty) - \sigma (i_t - E_t^R \pi_{t+1}) \quad (2.8)$$

which is the true consumption decision of rational forecasters satisfying the higher-order beliefs assumption from above.

Equation (2.8) could have been derived for the general case of agent τ as the assumption on higher-order beliefs holds for both agent types. In the general case

(2.8) reads

$$c_t^\tau = E_t^\tau y_{t+1} + E_t^\tau (c_\infty^\tau - y_\infty) - \sigma(i_t - E_t^\tau \pi_{t+1}). \quad (2.9)$$

To be consistent with Branch and McGough (2009), I assume that agents believe to be back in steady state in the long-run.³ In this case $E_t^\tau (c_\infty^\tau - y_\infty) = 0$ holds and, thus, (2.9) becomes

$$c_t^\tau = E_t^\tau y_{t+1} - \sigma(i_t - E_t^\tau \pi_{t+1}). \quad (2.10)$$

From (2.10) one can infer that agents only forecast aggregate variables when making consumption decisions. Note that, as rational forecasters have rational expectations with respect to aggregate variables, E_t^R can be replaced by E_t in the consumption decision of rational forecasters.

Using goods market clearing and (2.10) the IS curve is given by

$$y_t = \alpha E_t y_{t+1} + (1 - \alpha)\theta^2 y_{t-1} - \sigma(i_t - \alpha E_t \pi_{t+1} - (1 - \alpha)\theta^2 \pi_{t-1}). \quad (2.11)$$

Further, output is produced under monopolistic competition. Calvo pricing is assumed where a fixed fraction ξ_p of yeoman farmers cannot reset their prices in a given period (Calvo, 1983). Price dispersion arises because, first, optimal prices are different *between* expectation types since they depend on *expected* future marginal costs and, second, they differ *within* each type due to the fact that only a fraction of firms can reset prices. Thus, the Phillips-curve can be derived as

$$\pi_t = \alpha\beta E_t \pi_{t+1} + (1 - \alpha)\beta\theta^2 \pi_{t-1} + \kappa y_t + e_t \quad (2.12)$$

with $\kappa = \frac{(1-\xi_p)(1-\beta\xi_p)(\eta+\sigma^{-1})}{\xi_p(1+\epsilon\eta)}$ where ϵ is the price elasticity of demand for a differentiated good.⁴ As in Di Bartolomeo et al. (2016) the Phillips curve is augmented with a random cost-push shock e_t .⁵

Note that inflation and output exhibit some degree of *persistence* due to the presence of backward-looking expectations. The degree of persistence depends on the fraction of boundedly rational forecasters $1 - \alpha$ and their forecasting coefficient θ . The higher the

³Branch and McGough (2009) assume that agents agree on expected differences in expected limiting consumption, so that long-run expectations drop out when aggregating the individual consumption decisions. As the authors are interested only in aggregate dynamics, there is no need to explicitly specify what these expectations actually are. However, this model differs because optimal monetary policy depends on consumption dispersion. Branch and McGough (2009, p.1041) mention that one assumption, consistent with their own, would be to assume that agents believe to be back in the steady state in the long-run which I employ here.

⁴For the derivation of the Phillips curve (2.12) please refer to Branch and McGough (2009).

⁵This shock can, for instance, be micro-founded by assuming an exogenous time-varying wage mark-up as in Galí (2015).

two parameters the higher the degree of persistence. This persistence is also recognized by rational forecasters which feeds back into the aggregate equations (2.11) and (2.12) via rational expectations. Hence, the interaction between rational and boundedly rational expectations leads, *ceteris paribus*, to an *amplification* of shocks as identified by Gasteiger (2018). However, the central bank's ability to manipulate (rational) expectations by setting interest rates accordingly is a powerful tool to counteract this amplification effect.

The model is calibrated as in Di Bartolomeo et al. (2016) for the US economy following Rotemberg and Woodford (1997) with the time unit being one quarter.

$$\alpha = 0.7 \quad \theta = 1 \quad \beta = 0.99 \quad \sigma = 6.25 \quad \epsilon = 7.84 \quad \eta = 0.47 \quad \xi_p = 0.66$$

Table 2.1: Baseline calibration

2.3 An optimal interest rate rule from a timeless perspective

In this section I derive, first, an optimal timeless interest rate rule from a model-consistent welfare criterion and, second, a non-optimal rule under a conventional inflation-targeting objective as in Gasteiger (2014).

2.3.1 Loss functions

As in Gasteiger (2014, 2018) and Di Bartolomeo et al. (2016) the central bank aims at maximizing social welfare. I follow the approach of Di Bartolomeo et al. (2016) where the central bank exploits its detailed knowledge about the heterogeneity in expectations and minimizes a social welfare loss that is a second-order approximation of household utility (2.1). The intertemporal second-order approximated aggregate welfare loss can be derived as

$$\mathcal{W} = -\frac{\bar{C}U_C}{2} \sum_{t=0}^{\infty} \beta^t L_t + t.i.p. \quad (2.13)$$

with

$$L_t = \left(\eta + \frac{1}{\sigma} \right) y_t^2 + (\epsilon^2 \eta) \text{var}_i(p_t(i)) + \frac{1}{\sigma} \text{var}_i(c_t(i)). \quad (2.14)$$

and

$$var_i(p_t(i)) = \delta\pi_t^2 + \frac{\delta\xi_p(1-\alpha)}{\alpha} \left[\pi_t - \beta\theta^2\pi_{t-1} - \kappa \frac{c_t^B + \eta\sigma y_t}{1 + \eta\sigma} \right]^2 \quad (2.15)$$

$$var_i(c_t(i)) = \alpha(1-\alpha)(c_t^R - c_t^B)^2. \quad (2.16)$$

where $\delta = \frac{\xi_p}{(1-\beta\xi_p)(1-\xi_p)}$ is a measure of price stickiness and *t.i.p.* are the terms independent of policy.

Using market clearing to eliminate c_t^B as well as (2.15) and (2.16), loss function (2.14) can be rewritten as

$$\begin{aligned} L_t = & \Gamma_1 y_t^2 + \Gamma_2 \pi_t^2 + \Gamma_3 \pi_{t-1}^2 + \Gamma_4 (c_t^R)^2 \\ & + \Gamma_5 y_t c_t^R + \Gamma_6 \pi_t c_t^R + \Gamma_7 \pi_{t-1} c_t^R + \Gamma_8 \pi_t \pi_{t-1} + \Gamma_9 \pi_t y_t + \Gamma_{10} \pi_{t-1} y_t \end{aligned} \quad (2.17)$$

where the Γ_x -coefficients are given in the Appendix 2.B.1.

Under homogeneous rational expectations, i.e for $\alpha = 1$, price dispersion reduces to $var_i(p_t(i)) = \delta\pi_t^2$ and $var_i(c_t(i))$ to zero. Hence, in this case (3.34) reduces to

$$L_t^{\alpha=1} = \left(\eta + \frac{1}{\sigma} \right) y_t^2 + \epsilon^2 \eta \delta \pi_t^2. \quad (2.18)$$

Equation (3.34) can be called the model-consistent loss function, i.e. it is consistently microfounded under heterogeneous expectations as assumed in this chapter, and (2.18) the conventional⁶ inflation-targeting loss function.

In general, agents dislike volatility of consumption and, therefore, output due to the concave nature of their utility function. However, the weight that is placed on inflation in second-order approximated utility functions in the canonical New Keynesian model such as (2.18) is usually considerably higher compared to the weight on output (see Woodford, 2003 or Galí, 2015). This reflects that price dispersion, due to inflation and sticky prices, is *the* source of inefficiency in the baseline model, quickly resulting in relatively large welfare losses. In case of the conventional inflation-targeting loss (2.18), the weight on inflation is roughly 160 times the weight on output under baseline calibration.⁷

However, when producers have heterogeneous expectations, price dispersion arises not only due to sticky prices but also because they have different expectations regarding

⁶It is conventional in the sense that it can be derived from the conventional three-equation NK model under homogeneous rational expectations.

⁷The fact that the weight on inflation is high relative to the weight on output is general and robust with respect to the calibration.

future inflation and marginal costs, as reflected by (2.15). The weight on contemporaneous inflation in the model-consistent loss function (2.17) is roughly 230 times of the weight on contemporaneous output under baseline calibration with a 70 percent of rational forecasters, and increases to 270 times of the weight on contemporaneous output for 50 percent of rational forecasters. Additionally, the weights on lagged inflation and on the interaction between contemporaneous and lagged inflation are non-negligible. Hence, inflation results in even higher welfare losses through the price dispersion channel under heterogeneous expectations. However, even though inflation explains most of the results in the the welfare analysis in Section 2.5, there is still a trade-off between inflation and output (and consumption dispersion) under a cost-push shock that can be important in some cases.

Further, an interesting novelty of (2.14) is the appearance of consumption dispersion, i.e. the cross-sectional variance in consumption $var_i(c_t(i))$. This is due to the fact that the central bank recognizes that each individual agent dislikes volatility in her own individual consumption. Under homogeneity volatility in aggregate consumption would be equally distributed across agents and it would, therefore, be sufficient for the central bank to stabilize aggregate consumption (or output in this context). Thus, if consumption decisions are heterogeneous, the central bank additionally needs to minimize consumption dispersion. It should be noted, however, that the weight on consumption dispersion in (2.14) is even smaller compared to the weight on output. This indicates that agents might accept a certain degree of heterogeneity when the economy is relatively stable instead (see also Chapter 3). This finding is also in line with Debortoli and Galí (2017) who derive a model-consistent loss function in a Two Agent New Keynesian (TANK) model and show that it depends on a measure of heterogeneity but where the corresponding weight relative to inflation and output is also very low.

2.3.2 An optimal interest rate rule under timeless commitment

The central bank is assumed to set its interest rate so as to minimize (2.13) with L_t given by the model-consistent loss function (2.17) or, respectively, the conventional inflation-targeting objective (2.18) subject to the private sector equations

$$y_t = \alpha E_t y_{t+1} + (1 - \alpha)\theta^2 y_{t-1} - \sigma[i_t - \alpha E_t \pi_{t+1} - (1 - \alpha)\theta^2 \pi_{t-1}] \quad (2.19)$$

$$\pi_t = \alpha\beta E_t \pi_{t+1} + (1 - \alpha)\beta\theta^2 \pi_{t-1} + \kappa y_t + e_t \quad (2.20)$$

$$c_t^R = E_t y_{t+1} - \sigma(i_t - E_t \pi_{t+1}). \quad (2.21)$$

Minimizing (2.13), with the period loss function given by the conventional inflation-targeting objective (2.18), subject to the Phillips curve (2.20) gives

$$i_t = \gamma_1 y_{t-1} + \gamma_2 E_t y_{t+1} + \gamma_3 \pi_{t-1} + \gamma_4 E_t \pi_{t+1} + \gamma_5 e_t \quad (2.22)$$

with

$$\gamma_1 = \frac{(1 - \alpha)\theta - \alpha}{\sigma} + \alpha \frac{\delta\epsilon^2\eta\kappa^2}{1 + \eta\sigma + \delta\epsilon^2\eta\kappa^2\sigma} \quad (2.23)$$

$$\gamma_2 = \alpha \frac{1}{\sigma} - \frac{(1 - \alpha)}{\sigma} \left[\frac{\beta^2\theta^2(1 + \eta\sigma)}{1 + \eta\sigma + \delta\epsilon^2\eta\kappa^2\sigma} \right] \quad (2.24)$$

$$\gamma_3 = (1 - \alpha) \left[\frac{\theta^2(1 + \eta(\sigma + \delta\epsilon^2\kappa(\beta + \kappa\sigma)))}{1 + \eta\sigma + \delta\epsilon^2\eta\kappa^2\sigma} \right] \quad (2.25)$$

$$\gamma_4 = \alpha \left[1 + \frac{\beta\delta\epsilon^2\eta\kappa}{1 + \eta\sigma + \delta\epsilon^2\eta\kappa^2\sigma} \right] \quad (2.26)$$

$$\gamma_5 = \frac{\delta\epsilon^2\eta\kappa}{1 + \eta\sigma + \delta\epsilon^2\eta\kappa^2\sigma}. \quad (2.27)$$

Equation (2.22) is similar to the rule derived by Gasteiger (2018) and Gasteiger (2014).⁸ Note that timeless commitment introduces persistence even in the absence of boundedly rational forecasters, i.e. γ_1 reduces to $\frac{\delta\epsilon^2\eta\kappa^2}{1 + \eta\sigma + \delta\epsilon^2\eta\kappa^2\sigma} - \frac{1}{\sigma}$ and does not vanish for $\alpha = 1$.

The *optimal* interest rate rule under timeless commitment can be obtained by minimizing (2.13), with the period loss function given by (2.17), subject to the private sector equation (2.19), (2.20) and (2.21) as

$$i_t = \Omega_1 E_t \pi_{t+1} + \Omega_2 E_t \pi_{t+2} + \Omega_3 \pi_{t-3} + \Omega_4 \pi_{t-2} + \Omega_5 \pi_{t-1} + \Omega_6 E_t y_{t+1} + \Omega_7 E_t y_{t+2} \quad (2.28)$$

$$+ \Omega_8 y_{t-2} + \Omega_9 y_{t-1} + \Omega_{10} E_t c_{t+1}^R + \Omega_{11} E_t c_{t+2}^R + \Omega_{12} c_{t-2}^R + \Omega_{13} c_{t-1}^R + \Omega_{14} e_t$$

where the reaction coefficients Ω_x and derivations are given in the Appendix 2.B.2. A first inspection of (2.28) shows that the central bank reacts to output and inflation as usual but also to individual consumption of rational forecasters due to the consumption

⁸The author uses a non-micro-founded version of (2.18), i.e. $L_t = \frac{1}{2}(\pi_t^2 + \omega y_t^2)$. Setting $\omega = \frac{\sigma\eta+1}{\epsilon^2\eta\delta\sigma}$ and calculating through the optimization problem yields the interest rate rule (2.22).

dispersion dimension. However, a more striking observation is that the central bank finds it optimal to react to lags and leads of all variables ranging from $t - 2$ to $t + 2$ (and additionally $t - 3$ for inflation). This will be clarified below. Note that under homogeneous rational expectations, $\alpha = 1$, all coefficients associated with heterogeneous expectations vanish as well as the additional coefficients due to commitment, except for y_{t-1} which can be seen in table (2.5) in the Appendix 2.B.4.

Further, efficiently controlling private sector rational expectations requires the central bank to induce substantial persistence by including terms with timing $t - 2$ and $t - 3$ in the optimal rule (2.28). This becomes clear when one would construct a hypothetical case in which the central bank observes and reacts to expectations but neglects the endogenous nature of rational expectations with respect to its policy.⁹ In such a hypothetical case, the interest rate rule would be given by¹⁰

$$i_t = \Omega_1^* y_{t-1} + \Omega_2^* E_t y_{t+1} + \Omega_3^* E_t y_{t+2} + \Omega_4^* \pi_{t-1} + \Omega_5^* E_t \pi_{t+1} + \Omega_6^* E_t \pi_{t+2} + \Omega_7^* E_t c_{t+1}^R + \Omega_8^* E_t c_{t+2}^R + \Omega_9^* e_t \quad (2.29)$$

where all terms with timing $t - 2$ and $t - 3$ drop out.

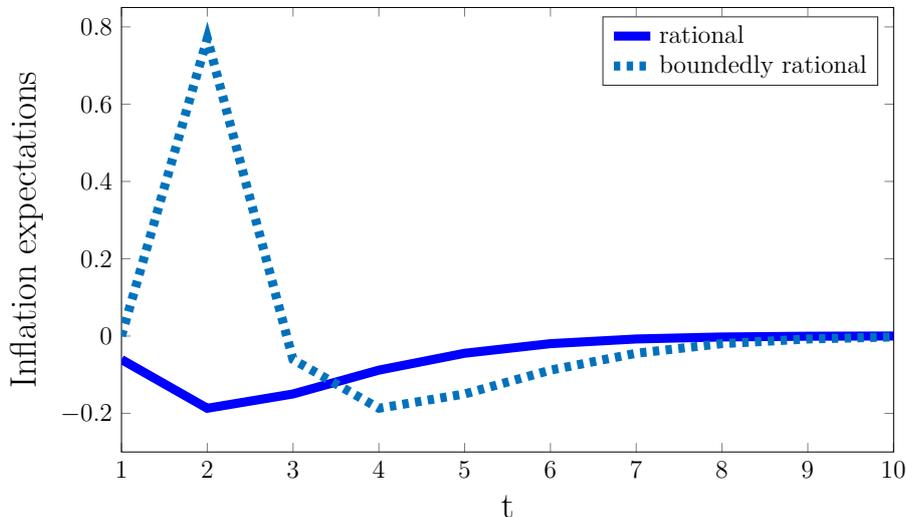


Figure 2.1: Individual inflation expectations in percentage deviation from steady state following a single, non-autocorrelated one-percent cost-push shock with monetary policy given by interest rate rule (2.28).

Moreover, to understand the appearance of the $t + 2$ terms consider Figure (2.1) which displays the reaction of the inflation expectations of both agent types following

⁹Such a behavior would, for instance, be consistent with a boundedly rational central bank that operates under some sort of “limited” commitment. However, I use this only for the sake of exposition.

¹⁰The reaction coefficients are given in the Appendix 2.B.2

a one standard deviation i.i.d. cost-push shock under the policy rule (2.28). Since all subsequent shock realizations are zero and rational forecasters know the true aggregate equations, they have de facto perfect foresight. Thus, rational forecasters' expectations in $t = 1$ about inflation in $t + 1$ will be correct, i.e. $E_t \pi_{t+1} = \pi_{t+1}$. However, the backward-looking expectations of boundedly rational forecasters in t about inflation in $t + 1$ will be zero, i.e. $E_t^B \pi_{t+1} = \pi_{t-1} = 0$ (where $\theta = 1$ for simplicity). In period $t + 1$ boundedly rational forecasters will expect inflation to *increase* drastically in $t + 2$ as their expectations are based on the period where the shock hits, i.e. $E_{t+1}^B \pi_{t+2} = \pi_t$. On the contrary, rational forecasters correctly expect inflation to *decrease* further as the central bank increases the nominal interest rate a second time (see impulse responses in Figure 2.2 in Section 2.4). Thus, the central bank induces the different expectations in $t + 1$ about $t + 2$ to diverge transitorily. Consequently, rational expectations counteract boundedly rational ones which stabilizes inflation. However, as price dispersion is additionally caused by *differences* in expectations about inflation and marginal costs, the central bank includes terms with $t + 2$ timing into its interest rate rule to prevent sub-optimally high dispersion in expectations.

Moreover, it seems, at first glance, that for practical implementation the optimal interest rate rule (2.28) requires to observe individual consumption of rational forecasters which is, however, not observable in reality. As already indicated, an advantage of the consumption decision (2.10) is that it only depends on aggregate variables as a result of the explicit incorporation of the higher-order beliefs assumption. Therefore, it is possible to substitute for individual consumption so that the optimal interest rate rule is merely a function of several leads and lags of aggregate variables.

I do not discuss determinacy issues as the model is determinate for all considered parameter constellations. There are two reasons for this. First, an expectations-based interest rate rule is derived, i.e. it properly accounts for private sector expectations which are known to perform exceptionally well as opposed to fundamentals-based reaction functions (Evans and Honkapohja, 2006). Second, the interest rate rule is derived from the model-consistent loss function and is, therefore, optimal.¹¹

2.4 Impulse Responses

This section briefly describes the simulation outcomes under baseline calibration given in Table 1. The impulse responses of a one percent i.i.d cost-push shock with monetary policy given by (2.28) are depicted in Figure 2.2. The aggregate behavior of the model

¹¹Given the restrictions imposed by the timeless commitment assumption.

is straightforward: taking into account subjective expectations, the real interest rates of both agent types, $r_t^r = i_t - E_t^r \pi_{t+1}$, increase due to an increase of the nominal rate by the central bank. Hence, both agent types cut their individual consumption which leads to a quite severe recession which counteracts the cost-push shock to some extent. Consequently, inflation increases by less than one percent. Thus, the central bank finds it optimal to be extraordinarily hawkish with respect to inflation which comes at the cost of a significant recession.

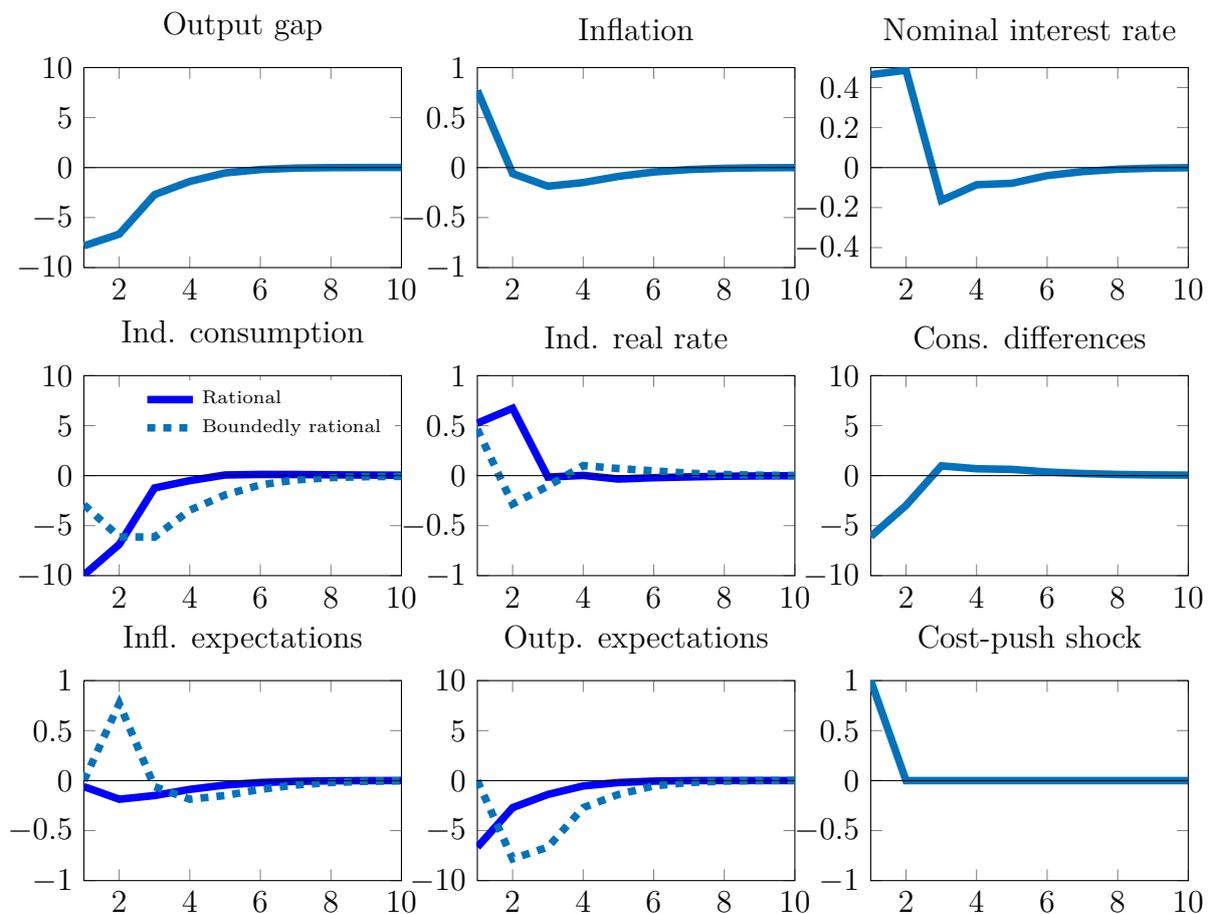


Figure 2.2: Impulse responses in %-deviations from steady state following a single, non-autocorrelated cost-push shock.

On the individual level, the disparity between the consumption adjustment paths of both agent types becomes obvious. While boundedly rational forecasters cut their consumption by approximately three percent on impact, rational forecasters decrease consumption by almost ten percent. This is, first, because of substantially negative rational output gap expectations and, second, due to a slightly higher subjective real interest rate. On the other hand, as boundedly rational forecasters are backward-looking, their output gap expectations are zero on impact and, therefore, cut their

consumption because of the increase in the subjective real interest rate only. This results in a consumption cut that is far smaller compared to rational forecasters and thereby in significant differences in individual consumption on impact.¹²

Thus, boundedly rational forecasters seem to be better off than rational ones at first. However, boundedly rational forecasters make less smart decisions than rational forecasters *by definition*. This becomes clear when looking at the following periods where boundedly rational forecasters have to pay for their initially higher consumption by giving up a lot of future consumption. Specifically, one can observe that boundedly rational output gap expectations in the second period ($t + 1$) drop drastically to $E_{t+1}^B y_{t+2} = y_t$ which is approximately minus 8 percent, resulting in a further cut of consumption. This is the case even though the subjective real interest rate of boundedly rational forecasters becomes negative which is due to the jump of their inflation expectations to $E_{t+1}^B \pi_{t+2} = \pi_t$. At the same time, output gap expectations of rational forecasters *increase* as the output gap recovers. From this period onwards rational forecasters are able to consume more than boundedly rational ones for a prolonged time.

2.5 Welfare evaluation

This section provides a comparison between the optimal interest rate rule (2.28) and the non-optimal inflation-targeting rule (2.22) in terms of welfare and a brief discussion on consumption dispersion. In particular, I analyze the welfare consequences of these rules following a one-percent i.i.d. cost-push shock as before. It should be noted that shocks to inflation directly (and, hence, to price dispersion) induce high welfare losses. The reason is, first, that price dispersion leads to dispersion in imperfectly substitutable individual production and, therefore, to losses in the final consumption bundle and, second, that output itself needs to be contracted in order to bring down inflation. Further, I will restrict the analysis in this Section to the case of naive expectations, i.e. $\theta = 1$, of boundedly rational forecasters for simplicity.

2.5.1 Optimal vs. inflation-targeting rule

In the following, I show *to what extent* the optimal interest rate rule (2.28) yields lower welfare losses compared to rule (2.22) depending on the fraction of rational forecasters. To that end, I compute consumption equivalent welfare losses following Ravenna and

¹²In this context, I simply define a weighted difference between consumption of rational and boundedly rational agents as $\alpha c_t^R - (1 - \alpha)c_t^B$.

Walsh (2011). Let

$$\mathcal{W}^O = -\frac{\bar{C}U_c}{2}E \sum_{t=0}^{\infty} \beta^t L_t^O + t.i.p. = -\frac{\bar{C}U_c}{2(1-\beta)}L^O + t.i.p. \quad (2.30)$$

be the welfare loss under the optimal commitment policy (2.28), and

$$\mathcal{W}^{IT} = -\frac{\bar{C}U_c}{2}E \sum_{t=0}^{\infty} \beta^t L_t^{IT} + t.i.p. = -\frac{\bar{C}U_c}{2(1-\beta)}L^{IT} + t.i.p. \quad (2.31)$$

be the welfare loss under the non-optimal inflation-targeting objective (2.22), where instantaneous losses are measured by (2.17) in both cases. The welfare loss of implementing policy (2.22) instead of (2.28) can be measured as the percentage increase of steady state consumption, CE , satisfying

$$\frac{U((1+CE) \times \bar{C})}{1-\beta} + \mathcal{W}^{IT} = \frac{U(\bar{C})}{1-\beta} + \mathcal{W}^O. \quad (2.32)$$

Inserting consumption utility $U(\bar{C}) = \frac{\bar{C}^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}}$, (2.30), (2.31) and $U_c = \bar{C}^{-\frac{1}{\sigma}}$, and solving for CE gives¹³

$$CE = \left(1 - \frac{\sigma-1}{2\sigma}(L^O - L^{IT})\right)^{\frac{\sigma}{\sigma-1}} - 1. \quad (2.33)$$

Table 2.2 shows absolute losses L^O and L^{IT} (second and third column) as well as the consumption equivalent welfare costs, CE , (fifth column) for different fractions of rational agents, α . The other columns of Table 2.2 are discussed further below.

A first, more trivial, observation is that absolute welfare losses increase with the fraction of boundedly rational forecasters for both interest rate rules. The higher the fraction of naive forecasters, the more persistent the deviations of variables from steady state and, therefore, the higher the long-run variances. Further, welfare losses are naturally lowest under the optimal interest rate rule (2.28). Deriving the interest rate rule from the conventional inflation-targeting (IT) objective (2.18) is costly in terms of consumption equivalents, CE . The implementation of the non-optimal inflation-targeting rule (2.22) already requires an increase of 14.5 percent of steady-state consumption to compensate for the higher welfare losses relative to the optimal interest rate rule (2.28) when only ten percent of the population form (naive) backward-looking expectations. Welfare costs become substantially higher for higher fractions of boundedly rational forecasters (lower α). Hence, the optimal interest rate rule (2.28) yields considerable

¹³Note that the terms independent of policy (*t.i.p.*) are the same for both \mathcal{W}^{IT} and \mathcal{W}^O and, therefore, drop out.

Table 2.2: Absolute welfare losses and consumption equivalent welfare costs

α	L^{T^*}	L^O	L^{IT}	L^{IT^*}	CE	CE^*
0.95	161.65	127.077	127.138	127.08	0.030	0.006
0.9	183.82	144.988	145.274	144.99	0.145	0.034
0.7	277.56	243.455	246.632	244.02	1.744	0.286
0.5	400.36	394.547	402.428	395.45	4.693	0.47

Column 2-4 show absolute welfare losses for *numerically optimized* simple Taylor rule (T^*), the optimal (O) interest rate rule (2.28), the inflation-targeting (IT) rule (2.22) and the inflation-targeting rule where the all reaction coefficients are *numerically optimized* (IT^*). Column 5-6 depict consumption equivalent welfare costs of implementing the non-optimal inflation-targeting (IT) interest rate rule (2.22) relative to the optimal (O) commitment policy (2.28) in percent, CE , and also for the *numerically optimized* inflation-targeting (IT^*) rule relative to the optimal rule, CE^* . All shown values are calculated for different fractions of rational agents, α .

welfare gains when the underlying economy features high degree of bounded rationality.

The difference between the two rules can be explained by their relative ability to stabilize inflation. The first two rows in Table 2.3 depict the variances of inflation and output for different values of rational forecasters, α , for both interest rate rules. In general, the optimal interest rate rule (first row) yields lower inflation but higher output volatility across all fractions of rational forecasters, α . This is a straightforward implication of the relatively higher weights on inflation in the model-consistent loss function (2.17) compared to the conventional inflation-targeting objective (2.18), as discussed in Section 2.3.1. While the differences in output and inflation volatility are relatively low for $\alpha = 0.95$, they become substantial for higher fractions of boundedly rational forecasters, resulting in quite extreme consumption equivalents.

Further, as discussed in Section 2.3.2, the optimal interest rate rule (2.28) depends on much more leads and lags of all variables compared to the non-optimal inflation-targeting rule (2.22). This raises the question whether it is the absence of these leads and lags that makes (2.22) sub-optimal or if it is rather an issue of weighting the different variables in the interest rate rule, or both. To answer this question, I *numerically optimize* the reaction coefficients in (2.22) with respect to welfare (2.17). The corresponding absolute welfare losses, L^{IT^*} , and consumption equivalent welfare costs relative to the optimal commitment policy, CE^* , can be found in columns four and six in Table 2.2, respectively. Consumption-equivalent welfare costs, CE^* , *substantially* decrease relative to the consumption-equivalent welfare costs, CE . This indicates that the inflation-targeting interest rate rule (2.22) is sub-optimal especially because of the

sub-optimal weighting of its different terms.¹⁴ However, for $\alpha = 0.7$ and $\alpha = 0.5$ the welfare gains of the optimal interest rate rule (2.28) are still non-negligible, suggesting that its additional leads and lags are, at least to some extent, important when heterogeneity increases.

rule	$Var(y)$				$Var(\pi)$			
	$\alpha = 0.95$	$\alpha = 0.9$	$\alpha = 0.7$	$\alpha = 0.5$	$\alpha = 0.95$	$\alpha = 0.9$	$\alpha = 0.7$	$\alpha = 0.5$
Optimal	41.23	50.66	115.17	231.69	0.58	0.60	0.66	0.67
IT	38.62	44.80	88.98	179.89	0.59	0.62	0.75	0.83
IT^*	41.37	50.93	116.82	232.55	0.58	0.60	0.66	0.66
T^*	42.00	59.98	149.45	240.37	0.78	0.78	0.70	0.65

Table 2.3: Theoretical variances of inflation and output under the optimal (O) rule (2.28), the inflation-targeting (IT) objective (2.22), the *numerically optimized* inflation-targeting (IT^*) rule and the *numerically optimized* simple Taylor rule (T^*) for different values of rational forecasters α .

Finally, for the sake of comparison, I add a numerically optimized simple Taylor rule with contemporaneous inflation and output to the analysis. Absolute welfare losses, L^{T^*} , are shown in the first column of Table 2.2 and the corresponding variances of inflation and output can be found in the last row of Table 2.3. For a fraction of 70 percent of rational forecasters, or higher, the numerically optimized simple Taylor rule is the worst performing among all alternatives. However, interestingly, for 50 percent of rational forecasters it trumps the analytically derived inflation-targeting rule (2.22). Note that the Taylor rule is numerically optimized under the model-consistent loss function (2.17) which, in particular, recognizes the complex nature of price dispersion which is highest around $\alpha = 0.5$. On the other hand, the inflation-targeting rule is derived from the conventional inflation-targeting objective (2.18) which abstracts from the complex nature of price dispersion.¹⁵ Further, in the case of $\alpha = 0.5$, the numerically optimized Taylor rule generates lower inflation volatility than the optimal

¹⁴I started by optimizing the coefficients on inflation (expectations) and the cost-push shock while “fixing” the coefficients on output to the analytically derived ones, (2.23) and (2.24). The resulting absolute welfare losses and consumption equivalents are nearly the same as in Table 2.2, indicating that the sub-optimal weighting in the analytically derived inflation-targeting rule (2.22) is mainly because of *sub-optimal weighting of inflation expectations and the shock*.

¹⁵For $\alpha \neq 0.5$ (where the case of $\alpha < 0.5$ is not explicitly shown here) the numerically optimized Taylor rule performs relatively worse as it does not react to rational and boundedly rational expectations (and the shock) while the inflation-targeting rule (2.22) does. Hence, it seems that reacting to the contemporaneous values of aggregate variables works relatively well only when *fractions are equal*.

commitment policy. This is, however, not welfare maximizing as the Taylor rule implies the highest output volatility among all interest rate rules and for all α . This indicates that, although inflation is the most important driver of welfare, there is still a welfare-relevant trade-off between inflation and output.

2.5.2 Inflation and welfare vs. consumption dispersion

In this section, I briefly discuss the issue of distribution, measured by the cross-sectional variance in consumption (2.16), and monetary policy. It should be noted, however, that *welfare losses* due to consumption dispersion are in general very low under the model-consistent loss function (2.17), where individual utilities are weighted equally and simply summed up. Of course, different conclusions may be reached when a social planner would attach higher weights to consumption dispersion.

Table 2.4 shows consumption dispersion (CD) and the variance of inflation for both the optimal interest rate rule (O) and inflation-targeting rule (IT) for different fractions of rational forecasters, α . Interestingly, consumption dispersion is *higher* under the optimal interest rate rule (2.28) compared to the (analytically derived) inflation-targeting rule (2.22) for all α . This is the case even though the former explicitly incorporates consumption heterogeneity as opposed to the latter. At the same time, the optimal interest rate rule yields lower inflation volatility, as discussed in the previous section. This indicates that the model additionally implies a trade-off between stabilizing inflation and consumption heterogeneity.

α	CD^O	CD^{IT}	$var(\pi)^O$	$var(\pi)^{IT}$
0.95	1.904	1.104	0.580	0.590
0.9	4.071	2.567	0.597	0.618
0.7	17.805	12.700	0.664	0.750
0.5	39.071	25.812	0.670	0.833

Table 2.4: Consumption dispersion (CD) $var_i(c_t(i))$ under the optimal (O) interest rate rule (2.28) and the non-optimal inflation-targeting (IT) rule (2.22) for different fractions of rational forecasters α .

The trade-off between stabilizing inflation and consumption dispersion in this model becomes more evident when considering Figure 2.3. Figure 2.3 depicts consumption dispersion, the inflation variance and absolute welfare losses in case of a simple Taylor

rule against different values of the coefficient on inflation. The higher the coefficient on inflation, the lower inflation volatility and the higher consumption dispersion. Therefore, it is not surprising that this also implies a local trade-off between minimizing welfare losses and reducing consumption dispersion.¹⁶ Minimizing welfare losses requires the central bank to get a tight grip on inflation causing a substantial drop of individual consumption over time. As both agents react quite differently to the increase in the policy rate, as discussed in Section 2.4, consumption dispersion increases with the central bank's aggressiveness towards inflation.

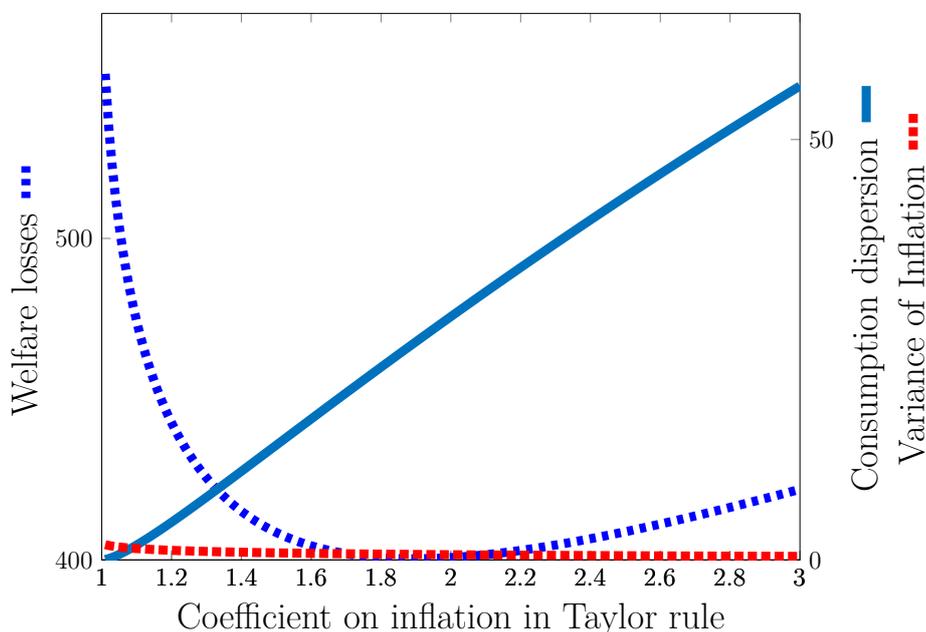


Figure 2.3: Trade-off between minimizing welfare, inflation and consumption dispersion under a simple Taylor rule and 70% of rational agents.

Another observation worth mentioning is that consumption dispersion is minimized when the coefficient on inflation in the Taylor rule is one, i.e. when $i_t = \pi_t$ holds. This will analytically be shown in Chapter 3 using a model with fully rational agents and boundedly rational agents similar to this chapter. Thus, the appearance of the same finding in this model serves as a robustness check for the findings in Chapter 3.

¹⁶On the left side of the minimum, welfare losses decrease because of decreasing inflation. At the same time, consumption dispersion increases. At some point, however, welfare losses increase again as output volatility (not shown in Figure 2.3) becomes substantially higher as the central bank needs to contract output further to achieve further reductions in inflation. Therefore, the trade-off between welfare and consumption dispersion arises only locally, i.e. on the left side of the minimum where welfare losses and inflation decrease simultaneously.

2.6 Conclusion

In this chapter, I propose an optimal interest rate rule under heterogeneous expectations where the central bank commits to its policy from a timeless perspective. This rule incorporates the more complex nature of price dispersion and consumption dispersion under heterogeneous expectations as identified by Di Bartolomeo et al. (2016). Further, this rule performs considerably better than a microfounded version of the interest rate rule as in Gasteiger (2014). The implementation of the non-optimal inflation-targeting rule already requires an increase of 14.5 percent of steady-state consumption to compensate for the higher welfare losses relative to the optimal interest rate rule when only ten percent of the population form (naive) backward-looking expectations.

I additionally explore the properties of the expectations operator of “rational” agents in the Branch and McGough (2009) framework and find that the consumption Euler equation that includes model-consistent individual consumption expectations as in Di Bartolomeo et al. (2016) is inconsistent with the higher-order beliefs assumption of Branch and McGough (2009). This assumption puts a specific (non-rational) structure on higher-order beliefs which implies that not even “rational forecasters” understand the micro level fully. Therefore, I derive consumption decisions that account for this particular assumption which makes it the implementation of the optimal commitment policy by an interest rate rule possible in the first place.

Finally, I illustrate that the model implies a local trade-off between maximizing welfare and reducing consumption dispersion. The reason is that consumption dispersion increases with the central bank’s aggressiveness towards inflation, as rational and boundedly rational forecasters’ consumption decisions become more unequal with more aggressive inflation-targeting. Because inflation is the most important determinant of welfare, the central bank has to allow for a certain heterogeneity in consumption to maximize welfare.

Appendix 2.A Implementation under the conventional inflation-targeting objective

The policy problem under commitment and the conventional inflation-targeting objective is given by

$$\begin{aligned} \mathcal{L} = E_t \sum_{s=0}^{\infty} \beta^s \frac{1}{2} \left[\left(\eta + \frac{1}{\sigma} \right) y_{t+s}^2 + \epsilon^2 \eta \delta \pi_{t+s}^2 \right] \\ + \lambda_{t+s} [\pi_{t+s} - \alpha \beta E_t \pi_{t+s+1} - (1 - \alpha) \beta \theta^2 \pi_{t+s-1} - \kappa y_{t+s} - e_{t+s}] \end{aligned} \quad (2.34)$$

The first order conditions are

$$\frac{\partial \mathcal{L}}{\partial \pi_{t+s}} : E_t \left\{ \epsilon^2 \eta \delta \pi_{t+s} + \frac{\lambda_{t+s}}{2} - (1 - \alpha) \beta^2 \theta^2 \frac{\lambda_{t+s+1}}{2} - \alpha \frac{\lambda_{t+s-1}}{2} \right\} \stackrel{!}{=} 0 \quad (2.35)$$

$$\frac{\partial \mathcal{L}}{\partial y_{t+s}} : E_t \left\{ \beta^s \left[\left(\eta + \frac{1}{\sigma} \right) y_{t+s} - \frac{\kappa}{2} \lambda_{t+s} \right] \right\} \stackrel{!}{=} 0. \quad (2.36)$$

Combining and solving for inflation gives

$$\pi_t = -\frac{1 + \eta \sigma}{\sigma \epsilon^2 \eta \delta \kappa} [y_t - \alpha y_{t-1} - (1 - \alpha) \beta^2 \theta^2 E_t y_{t+1}] \quad (2.37)$$

where the index s was dropped as the central bank commits from timeless perspective. Combining with the Phillips and IS curve yields (2.22).

Appendix 2.B Optimal monetary policy

2.B.1 Rewriting the model-consistent loss function

The period loss function L_t is given by

$$\begin{aligned} L_t = \frac{\sigma \eta + 1}{\sigma} y_t^2 + \frac{\alpha (y_t - c_t^R)^2}{(1 - \alpha) \sigma} \\ + \epsilon^2 \eta \delta \left\{ \pi_t^2 + \frac{\xi_p (1 - \alpha)}{\alpha} \left[\pi_t - \beta \theta^2 \pi_{t-1} - \kappa y_t - \frac{\alpha \kappa (y_t - c_t^R)}{(1 + \eta \sigma) (1 - \alpha)} \right]^2 \right\}. \end{aligned} \quad (2.38)$$

where market clearing has been used to eliminate c_t^B .

By multiplying out, we get

$$L_t = \Gamma_1 y_t^2 + \Gamma_2 \pi_t^2 + \Gamma_3 \pi_{t-1}^2 + \Gamma_4 (c_t^R)^2 \quad (2.39)$$

$$+ \Gamma_5 y_t c_t^R + \Gamma_6 \pi_t c_t^R + \Gamma_7 \pi_{t-1} c_t^R + \Gamma_8 \pi_t \pi_{t-1} + \Gamma_9 \pi_t y_t + \Gamma_{10} \pi_{t-1} y_t$$

with

$$\Gamma_1 = \frac{((\alpha - 1)\eta\sigma - 1)(\alpha(\eta^2\sigma^2(\delta\epsilon^2\kappa^2\xi_p - 1) - 1 - 2\eta\sigma) - \delta\epsilon^2\eta\kappa^2\xi_p\sigma(1 + \eta\sigma))}{(1 - \alpha)\alpha\sigma(1 + \eta\sigma)^2} \quad (2.40)$$

$$\Gamma_2 = \frac{\delta\epsilon^2\eta(\alpha + \xi_p - \alpha\xi_p)}{\alpha} \quad (2.41)$$

$$\Gamma_3 = \frac{(1 - \alpha)\beta^2\delta\epsilon^2\eta\theta^4\xi_p}{\alpha} \quad (2.42)$$

$$\Gamma_4 = \frac{\alpha(1 + \eta\sigma(2 + \delta\epsilon^2\kappa^2\xi_p) + \eta^2\sigma^2)}{(1 - \alpha)\sigma(1 + \eta\sigma)^2} \quad (2.43)$$

$$\Gamma_5 = \frac{2(\alpha + 2\alpha\eta\sigma + \alpha\eta^2\sigma^2(1 - \delta\epsilon^2\kappa^2\xi_p) + \delta\epsilon^2\eta\kappa^2\xi_p\sigma(1 + \eta\sigma))}{(\alpha - 1)\sigma(1 + \eta\sigma)^2} \quad (2.44)$$

$$\Gamma_6 = \frac{2\delta\epsilon^2\eta\kappa\xi_p}{1 + \eta\sigma} \quad (2.45)$$

$$\Gamma_7 = -\frac{2\beta\delta\epsilon^2\eta\theta^2\kappa\xi_p}{1 + \eta\sigma} \quad (2.46)$$

$$\Gamma_8 = \frac{2(\alpha - 1)\beta\delta\epsilon^2\eta\theta^2\xi_p}{\alpha} \quad (2.47)$$

$$\Gamma_9 = \frac{2\delta\epsilon^2\eta\kappa\xi_p((\alpha - 1)\eta\sigma - 1)}{\alpha + \alpha\eta\sigma} \quad (2.48)$$

$$\Gamma_{10} = \frac{2\beta\delta\epsilon^2\eta\theta^2\kappa\xi_p(1 + \eta\sigma(1 - \alpha))}{\alpha + \alpha\eta\sigma}. \quad (2.49)$$

2.B.2 Optimal interest rate rule

The policy problem under commitment takes the following form:

$$\mathcal{L} = E_t \sum_{s=0}^{\infty} \beta^s \left[\Gamma_1 y_{t+s}^2 + \Gamma_2 \pi_{t+s}^2 + \Gamma_3 \pi_{t+s-1}^2 + \Gamma_4 (c_{t+s}^R)^2 \quad (2.50)$$

$$+ \Gamma_5 y_{t+s} c_{t+s}^R + \Gamma_6 \pi_{t+s} c_{t+s}^R + \Gamma_7 \pi_{t+s-1} c_{t+s}^R + \Gamma_8 \pi_{t+s} \pi_{t+s-1} + \Gamma_9 \pi_{t+s} y_{t+s} + \Gamma_{10} \pi_{t+s-1} y_{t+s}$$

$$+ \lambda_{1,t+s} [y_{t+s} - \alpha E_t y_{t+s+1} - (1 - \alpha)\theta^2 y_{t+s-1} + \sigma [i_{t+s} - \alpha E_t \pi_{t+s+1} - (1 - \alpha)\theta^2 \pi_{t+s-1}]]$$

$$+ \lambda_{2,t+s} [\pi_{t+s} - \alpha \beta E_t \pi_{t+s+1} - (1 - \alpha)\beta\theta^2 \pi_{t+s-1} - \kappa y_{t+s} - e_{t+s}]$$

$$+ \lambda_{3,t+s} [c_{t+s}^R - E_t y_{t+s+1} + \sigma (i_{t+s} - E_t \pi_{t+s+1})] \right].$$

The first order conditions are

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial y_{t+s}} : E_t \left\{ \beta^s [2\Gamma_1 y_{t+s} + \Gamma_5 c_{t+s}^R + \Gamma_9 \pi_{t+s} + \Gamma_{10} \pi_{t+s-1} + \lambda_{1,t+s} - \kappa \lambda_{2,t+s}] \right. \\ \left. - \beta^{s+1} (1 - \alpha) \theta^2 \lambda_{1,t+s+1} - \beta^{s-1} [\alpha \lambda_{1,t+s-1} + \lambda_{3,t+s-1}] \right\} \stackrel{!}{=} 0 \end{aligned} \quad (2.51)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \pi_{t+s}} : E_t \left\{ \beta^s [2\Gamma_2 \pi_{t+s} + \Gamma_6 c_{t+s}^R + \Gamma_8 \pi_{t+s-1} + \Gamma_9 y_{t+s} + \lambda_{2,t+s}] \right. \\ \left. + \beta^{s+1} [2\Gamma_3 \pi_{t+s} + \Gamma_7 c_{t+s+1} + \Gamma_8 \pi_{t+s+1} + \Gamma_{10} y_{t+s+1} - (1 - \alpha) \theta^2 \sigma \lambda_{1,t+s+1}] \right. \\ \left. - (1 - \alpha) \beta \theta^2 \lambda_{2,t+s+1} - \beta^{s-1} [\alpha \sigma \lambda_{1,t+s-1} + \alpha \beta \lambda_{2,t+s-1} + \sigma \lambda_{3,t+s-1}] \right\} \stackrel{!}{=} 0 \end{aligned} \quad (2.52)$$

$$\frac{\partial \mathcal{L}}{\partial c_{t+s}^R} : E_t \left\{ \beta^s [2\Gamma_4 c_{t+s}^R + \Gamma_5 y_{t+s} + \Gamma_6 \pi_{t+s} + \Gamma_7 \pi_{t+s-1} + \lambda_{3,t+s}] \right\} \stackrel{!}{=} 0 \quad (2.53)$$

$$\frac{\partial \mathcal{L}}{\partial i_{t+s}} : E_t \left\{ \beta^s \sigma \lambda_{1,t+s} + \beta^s \sigma \lambda_{3,t+s} \right\} \stackrel{!}{=} 0. \quad (2.54)$$

Again, the index s can be dropped assuming commitment from a timeless perspective. Using $\lambda_{3,t} = -\lambda_{1,t}$ the FOCs can equivalently be written as

$$\begin{aligned} 2\Gamma_1 y_t + \Gamma_5 c_t^R + \Gamma_9 \pi_t + \Gamma_{10} \pi_{t-1} + \lambda_{1,t} - \kappa \lambda_{2,t} - (1 - \alpha) \beta \theta^2 \lambda_{1,t+1} - \\ \beta^{-1} (\alpha - 1) \lambda_{1,t-1} \stackrel{!}{=} 0 \end{aligned} \quad (2.55)$$

$$\begin{aligned} 2\Gamma_2 \pi_t + \Gamma_6 c_t^R + \Gamma_8 \pi_{t-1} + \Gamma_9 y_t + \lambda_{2,t} \\ + 2\Gamma_3 \beta \pi_t + \Gamma_7 \beta c_{t+1} + \Gamma_8 \beta \pi_{t+1} + \Gamma_{10} \beta y_{t+1} - (1 - \alpha) \beta \theta^2 \sigma \lambda_{1,t+1} \\ - (1 - \alpha) \beta^2 \theta^2 \lambda_{2,t+1} + \beta^{-1} (1 - \alpha) \sigma \lambda_{1,t-1} - \alpha \lambda_{2,t-1} \stackrel{!}{=} 0 \end{aligned} \quad (2.56)$$

$$2\Gamma_4 c_t^R + \Gamma_5 y_t + \Gamma_6 \pi_t + \Gamma_7 \pi_{t-1} - \lambda_{1,t} \stackrel{!}{=} 0. \quad (2.57)$$

Eliminating the Lagrange multipliers yields the reduced-form FOC

$$\begin{aligned} \Delta_1^c \pi_t + \Delta_2^c \pi_{t+1} + \Delta_3^c \pi_{t+2} + \Delta_4^c \pi_{t-3} + \Delta_5^c \pi_{t-2} + \Delta_6^c \pi_{t-1} + \Delta_7 y_t + \Delta_8^c y_{t+1} \\ + \Delta_9^c y_{t+2} + \Delta_{10}^c y_{t-2} + \Delta_{11}^c y_{t-1} + \Delta_{12}^c c_t^R + \Delta_{13}^c c_{t+1}^R + \Delta_{14}^c c_{t+2}^R + \Delta_{15}^c c_{t-2}^R + \Delta_{16}^c c_{t-1}^R \stackrel{!}{=} 0 \end{aligned} \quad (2.58)$$

with

$$\Delta_1 = \Gamma_6 + \Gamma_9 + (1 - \alpha)(-1 + 2\alpha)\beta\Gamma_6\theta^2 + 2\Gamma_2\kappa + 2\beta\Gamma_3\kappa \quad (2.59)$$

$$+ (\alpha - 1)\beta\theta^2(\Gamma_7 + \beta(\Gamma_{10} + \Gamma_7) + \Gamma_7\kappa\sigma) \quad (2.60)$$

$$\Delta_2 = \beta(\Gamma_8\kappa + (\alpha - 1)\theta^2(\beta(\Gamma_9 + (\alpha - 1)\beta\Gamma_7\theta^2) + \Gamma_6(1 + \beta + \kappa\sigma))) \quad (2.61)$$

$$\Delta_3 = (\alpha - 1)^2\beta^3\Gamma_6\theta^4 \quad (2.62)$$

$$\Delta_4 = \frac{(\alpha - 1)\alpha\Gamma_7}{\beta} \quad (2.63)$$

$$\Delta_5 = \frac{\alpha^2\Gamma_6 + \Gamma_7 + \Gamma_7\kappa\sigma - \alpha(\Gamma_6 + \Gamma_7 + \beta(\Gamma_{10} + \Gamma_7) + \Gamma_7\kappa\sigma)}{\beta} \quad (2.64)$$

$$\Delta_6 = \Gamma_{10} + \Gamma_7 - \alpha(\Gamma_6 + \Gamma_9) + (-1 + (3 - 2\alpha)\alpha)\beta\Gamma_7\theta^2 + \Gamma_8\kappa \quad (2.65)$$

$$- \frac{(\alpha - 1)\Gamma_6(1 + \kappa\sigma)}{\beta} \quad (2.66)$$

$$\Delta_7 = 2\Gamma_1 + \Gamma_5 + (1 - \alpha)(2\alpha - 1)\beta\Gamma_5\theta^2 + \Gamma_9\kappa \quad (2.67)$$

$$\Delta_8 = \beta(\Gamma_{10}\kappa + (\alpha - 1)\theta^2(\Gamma_5 + \beta(2\Gamma_1 + \Gamma_5) + \Gamma_5\kappa\sigma)) \quad (2.68)$$

$$\Delta_9 = (\alpha - 1)^2\beta^3\Gamma_5\theta^4 \quad (2.69)$$

$$\Delta_{10} = \frac{(\alpha - 1)\alpha\Gamma_5}{\beta} \quad (2.70)$$

$$\Delta_{11} = \frac{\Gamma_5 + \Gamma_5\kappa\sigma - \alpha(\Gamma_5 + \beta(2\Gamma_1 + \Gamma_5) + \Gamma_5\kappa\sigma)}{\beta} \quad (2.71)$$

$$\Delta_{12} = 2\Gamma_4 + \Gamma_5 - 2(\alpha - 1)(2\alpha - 1)\beta\Gamma_4\theta^2 + \Gamma_6\kappa \quad (2.72)$$

$$\Delta_{13} = \beta(\Gamma_7\kappa + (\alpha - 1)\theta^2(\beta\Gamma_5 + 2\Gamma_4(1 + \beta + \kappa\sigma))) \quad (2.73)$$

$$\Delta_{14} = 2(\alpha - 1)^2\beta^3\Gamma_4\theta^4 \quad (2.74)$$

$$\Delta_{15} = \frac{2(\alpha - 1)\alpha\Gamma_4}{\beta} \quad (2.75)$$

$$\Delta_{16} = -\alpha\Gamma_5 - \frac{2\Gamma_4(-1 + \alpha + \alpha\beta + (\alpha - 1)\kappa\sigma)}{\beta}. \quad (2.76)$$

Solving (2.58) for π_t and setting it equal to the NK Phillips curve yields

$$\begin{aligned} y_t = & -\frac{1}{\Delta_7 + \Delta_1\kappa}((\alpha\beta\Delta_1 + \Delta_2)\pi_{t+1} + \Delta_3\pi_{t+2} + \Delta_4\pi_{t-3} + \Delta_5\pi_{t-2} \\ & + (\Delta_6 + (1 - \alpha)\beta\theta^2\Delta_1)\pi_{t-1} + \Delta_8y_{t+1} + \Delta_9y_{t+2} + \Delta_{10}y_{t-2} + \Delta_{11}y_{t-1} \\ & + \Delta_{13}c_{t+1}^R + \Delta_{14}c_{t+2}^R + \Delta_{15}c_{t-2}^R + \Delta_{16}c_{t-1}^R) \end{aligned} \quad (2.77)$$

Setting (2.77) equal to the New IS curve, substituting c_t^R for consumption demand and solving for i_t gives the central bank's reaction function under timeless commitment

(2.28):

$$i_t = \Omega_1 E_t \pi_{t+1} + \Omega_2 E_t \pi_{t+2} + \Omega_3 \pi_{t-3} + \Omega_4 \pi_{t-2} + \Omega_5 \pi_{t-1} + \Omega_6 E_t y_{t+1} + \Omega_7 E_t y_{t+2} \quad (2.78)$$

$$+ \Omega_8 y_{t-2} + \Omega_9 y_{t-1} + \Omega_{10} E_t c_{t+1}^R + \Omega_{11} E_t c_{t+2}^R + \Omega_{12} c_{t-2}^R + \Omega_{13} c_{t-1}^R + \Omega_{14} e_t$$

with

$$\Omega_1 = \frac{\alpha\beta\Delta_1 + \Delta_2 + \sigma\Delta_{12} + \alpha\sigma(\Delta_7 + \Delta_1\kappa)}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.79)$$

$$\Omega_2 = \frac{\Delta_3}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.80)$$

$$\Omega_3 = \frac{\Delta_4}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.81)$$

$$\Omega_4 = \frac{\Delta_5}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.82)$$

$$\Omega_5 = \frac{\Delta_6 + (1 - \alpha)\theta^2(\beta\Delta_1 + \sigma(\Delta_7 + \Delta_1\kappa))}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.83)$$

$$\Omega_6 = \frac{\Delta_{12} + \alpha\Delta_7 + \Delta_8 + \alpha\kappa\Delta_1}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.84)$$

$$\Omega_7 = \frac{\Delta_9}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.85)$$

$$\Omega_8 = \frac{\Delta_{10}}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.86)$$

$$\Omega_9 = \frac{\Delta_{11} + (1 - \alpha)\theta^2(\Delta_7 + \Delta_1\kappa)}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.87)$$

$$\Omega_{10} = \frac{\Delta_{13}}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.88)$$

$$\Omega_{11} = \frac{\Delta_{14}}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.89)$$

$$\Omega_{12} = \frac{\Delta_{15}}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.90)$$

$$\Omega_{13} = \frac{\Delta_{16}}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)} \quad (2.91)$$

$$\Omega_{14} = \frac{\Delta_1}{\sigma(\Delta_7 + \Delta_{12} + \Delta_1\kappa)}. \quad (2.92)$$

2.B.3 Taking rational expectations as given

Defining $f_{t+s} = E_t y_{t+s+1}$, $g_{t+s} = E_t \pi_{t+s+1}$ and $h_{t+s} = E_t c_{t+s+1}^R$, the policy problem is

$$\begin{aligned} \mathcal{L} = E_t \sum_{s=0}^{\infty} \beta^s & \left[\Gamma_1 y_{t+s}^2 + \Gamma_2 \pi_{t+s}^2 + \Gamma_3 \pi_{t+s-1}^2 + \Gamma_4 (c_{t+s}^R)^2 \right. \\ & + \Gamma_5 y_{t+s} c_{t+s}^R + \Gamma_6 \pi_{t+s} c_{t+s}^R + \Gamma_7 \pi_{t+s-1} c_{t+s}^R + \Gamma_8 \pi_{t+s} \pi_{t+s-1} + \Gamma_9 \pi_{t+s} y_{t+s} + \Gamma_{10} \pi_{t+s-1} y_{t+s} \\ & + \lambda_{1,t+s} [y_{t+s} - \alpha f_{t+s} - (1-\alpha)\theta^2 y_{t+s-1} + \sigma [i_{t+s} - \alpha g_{t+s} - (1-\alpha)\theta^2 \pi_{t+s-1}]] \\ & + \lambda_{2,t+s} [\pi_{t+s} - \alpha \beta g_{t+s} - (1-\alpha)\beta \theta^2 \pi_{t+s-1} - \kappa y_{t+s} - e_{t+s}] \\ & \left. + \lambda_{3,t+s} [c_{t+s}^R - h_{t+s} + \sigma (i_{t+s} - g_{t+s})] \right]. \end{aligned} \quad (2.93)$$

The first order conditions are

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial y_{t+s}} : E_t \left\{ \beta^s [2\Gamma_1 y_{t+s} + \Gamma_5 c_{t+s}^R + \Gamma_9 \pi_{t+s} + \Gamma_{10} \pi_{t+s-1} + \lambda_{1,t+s} - \kappa \lambda_{2,t+s}] \right. \\ \left. - \beta^{s+1} (1-\alpha)\theta^2 \lambda_{1,t+s+1} \right\} \stackrel{!}{=} 0 \end{aligned} \quad (2.94)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \pi_{t+s}} : E_t \left\{ \beta^s [2\Gamma_2 \pi_{t+s} + \Gamma_6 c_{t+s}^R + \Gamma_8 \pi_{t+s-1} + \Gamma_9 y_{t+s} + \lambda_{2,t+s}] \right. \\ \left. + \beta^{s+1} [2\Gamma_3 \pi_{t+s} + \Gamma_7 c_{t+s+1}^R + \Gamma_8 \pi_{t+s+1} + \Gamma_{10} y_{t+s+1} - (1-\alpha)\theta^2 \sigma \lambda_{1,t+s+1} \right. \\ \left. - (1-\alpha)\beta \theta^2 \lambda_{2,t+s+1}] \right\} \stackrel{!}{=} 0 \end{aligned} \quad (2.95)$$

$$\frac{\partial \mathcal{L}}{\partial c_{t+s}^R} : E_t \left\{ \beta^s [2\Gamma_4 c_{t+s}^R + \Gamma_5 y_{t+s} + \Gamma_6 \pi_{t+s} + \Gamma_7 \pi_{t+s-1} + \lambda_{3,t+s}] \right\} \stackrel{!}{=} 0 \quad (2.96)$$

$$\frac{\partial \mathcal{L}}{\partial i_{t+s}} : E_t \left\{ \beta^s \sigma \lambda_{1,t+s} + \beta^s \sigma \lambda_{3,t+s} \right\} \stackrel{!}{=} 0. \quad (2.97)$$

Since the central bank acts under timeless commitment, the index s can be dropped. Using $\lambda_{3,t} = -\lambda_{1,t}$ the FOCs can equivalently be written as

$$2\Gamma_1 y_t + \Gamma_5 c_t^R + \Gamma_9 \pi_t + \Gamma_{10} \pi_{t-1} + \lambda_{1,t} - \kappa \lambda_{2,t} - (1 - \alpha) \beta \theta^2 E_t \lambda_{1,t+1} \stackrel{!}{=} 0 \quad (2.98)$$

$$2\Gamma_2 \pi_t + \Gamma_6 c_t^R + \Gamma_8 \pi_{t-1} + \Gamma_9 y_t + \lambda_{2,t} + 2\Gamma_3 \beta \pi_t + \Gamma_7 \beta E_t c_{t+1}^R + \Gamma_8 \beta E_t \pi_{t+1} + \Gamma_{10} \beta E_t y_{t+1} - (1 - \alpha) \beta \theta^2 \sigma E_t \lambda_{1,t+1} - (1 - \alpha) \beta^2 \theta^2 E_t \lambda_{2,t+1} \stackrel{!}{=} 0 \quad (2.99)$$

$$2\Gamma_4 c_t^R + \Gamma_5 y_t + \Gamma_6 \pi_t + \Gamma_7 \pi_{t-1} - \lambda_{1,t} \stackrel{!}{=} 0. \quad (2.100)$$

Eliminating the Lagrange multipliers yields the reduced-form FOC

$$\begin{aligned} &\Delta_1 \pi_t + \Delta_2 E_t \pi_{t+1} + \Delta_3 E_t \pi_{t+2} + \Delta_4 \pi_{t-1} + \Delta_5 y_t + \Delta_6 E_t y_{t+1} + \Delta_7 E_t y_{t+2} \\ &+ \Delta_8 c_t^R + \Delta_9 E_t c_{t+1}^R + \Delta_{10} E_t c_{t+2}^R \stackrel{!}{=} 0 \end{aligned} \quad (2.101)$$

with

$$\Delta_1 = -\frac{\Gamma_6 + \Gamma_9 + 2\Gamma_2 \kappa + 2\beta \Gamma_3 \kappa - (1 - \alpha) \beta \theta^2 (\Gamma_7 + \beta(\Gamma_{10} + \Gamma_7) + \Gamma_7 \kappa \sigma)}{\kappa} \quad (2.102)$$

$$\Delta_2 = \frac{(1 - \alpha) \beta \theta^2 (\beta(\Gamma_9 - (1 - \alpha) \beta \Gamma_7 \theta^2) + \Gamma_6 (1 + \beta + \kappa \sigma))}{\kappa} - \beta \Gamma_8 \quad (2.103)$$

$$\Delta_3 = -\frac{(\alpha - 1)^2 \beta^3 \theta^4 \Gamma_6}{\kappa} \quad (2.104)$$

$$\Delta_4 = -\frac{\Gamma_{10} + \Gamma_7 + \Gamma_8 \kappa}{\kappa} \quad (2.105)$$

$$\Delta_5 = -\frac{2\Gamma_1 + \Gamma_5 + \Gamma_9 \kappa}{\kappa} \quad (2.106)$$

$$\Delta_6 = -\frac{\beta((\alpha - 1) \beta (2\Gamma_1 + \Gamma_5) \theta^2 + \Gamma_{10} \kappa + (\alpha - 1) \Gamma_5 \theta^2 (1 + \kappa \sigma))}{\kappa} \quad (2.107)$$

$$\Delta_7 = -\frac{(\alpha - 1)^2 \beta^3 \theta^4 \Gamma_5}{\kappa} \quad (2.108)$$

$$\Delta_8 = -\frac{2\Gamma_4 + \Gamma_5 + \Gamma_6 \kappa}{\kappa} \quad (2.109)$$

$$\Delta_9 = -\frac{\beta((\alpha - 1) \beta \Gamma_5 \theta^2 + \Gamma_7 \kappa + 2(\alpha - 1) \Gamma_4 \theta^2 (1 + \beta + \kappa \sigma))}{\kappa} \quad (2.110)$$

$$\Delta_{10} = -\frac{2(\alpha - 1)^2 \beta^3 \theta^4 \Gamma_4}{\kappa}. \quad (2.111)$$

Solving (2.101) for π_t and setting it equal to the NK Phillips curve yields

$$y_t = -\frac{1}{\Delta_5 + \Delta_1\kappa}(\Delta_6 y_{t+1} + \Delta_7 y_{t+2} + (\Delta_2 + \alpha\beta\Delta_1)\pi_{t+1} + \Delta_3\pi_{t+2} + (\Delta_4 + (1-\alpha)\beta\theta^2\Delta_1)\pi_{t-1} + \Delta_8 c_t^R + \Delta_9 c_{t+1}^R + \Delta_{10} c_{t+2}^R + \Delta_1 e_t). \quad (2.112)$$

Setting (2.112) equal to the New IS curve, substituting c_t^R for consumption demand and solving for i_t gives the reaction function under the assumption that the central bank takes rational expectations as given

$$i_t = \Omega_1^* y_{t-1} + \Omega_2^* E_t y_{t+1} + \Omega_3^* E_t y_{t+2} + \Omega_4^* \pi_{t-1} + \Omega_5^* E_t \pi_{t+1} + \Omega_6^* E_t \pi_{t+2} + \Omega_7^* E_t c_{t+1}^R + \Omega_8^* E_t c_{t+2}^R + \Omega_9^* e_t \quad (2.113)$$

with

$$\Omega_1^* = \frac{(1-\alpha)\theta^2(\Delta_5 + \Delta_1\kappa)}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)} \quad (2.114)$$

$$\Omega_2^* = \frac{\Delta_6 + \Delta_8 + \alpha(\Delta_5 + \Delta_1\kappa)}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)} \quad (2.115)$$

$$\Omega_3^* = \frac{\Delta_7}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)} \quad (2.116)$$

$$\Omega_4^* = \frac{\Delta_4 + (1-\alpha)\theta^2(\beta\Delta_1 + \sigma(\Delta_5 + \Delta_1\kappa))}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)} \quad (2.117)$$

$$\Omega_5^* = \frac{\alpha\beta\Delta_1 + \Delta_2 + \sigma\Delta_8 + \alpha\sigma(\Delta_5 + \Delta_1\kappa)}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)} \quad (2.118)$$

$$\Omega_6^* = \frac{\Delta_3}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)} \quad (2.119)$$

$$\Omega_7^* = \frac{\Delta_9}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)} \quad (2.120)$$

$$\Omega_8^* = \frac{\Delta_{10}}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)} \quad (2.121)$$

$$\Omega_9^* = \frac{\Delta_1}{\sigma(\Delta_5 + \Delta_8 + \Delta_1\kappa)}. \quad (2.122)$$

2.B.4 Tables

Table 2.5 shows the reaction coefficients in interest rate rule (2.28) computed under baseline calibration but with varying α . It can be seen that the central bank puts relatively high weights on inflation (expectations) and the shock. In general, the weights that are placed on lagged inflation increase in absolute value when the

fraction of rational forecasters, α , decreases as the central bank puts a higher weight on backward-looking expectations. The opposite holds true for the one-period ahead rational expectations. For the homogeneous rational expectations case, $\alpha = 1$, all reaction coefficients associated with heterogeneous expectations vanish.

Ω_x	$\alpha = 0.1$	$\alpha = 0.3$	$\alpha = 0.5$	$\alpha = 0.7$	$\alpha = 0.9$	$\alpha = 1$ (RE)
y_{t-2}	0.001	0.004	0.01	0.018	0.029	0
y_{t-1}	0.115	0.052	-0.016	-0.085	-0.153	-0.139
$E_t y_{t+1}$	-0.08	-0.02	0.045	0.111	0.175	0.160
$E_t y_{t+2}$	-0.005	-0.008	-0.009	-0.007	-0.003	0
π_{t-3}	0.012	0.028	0.034	0.029	0.013	0
π_{t-2}	-0.609	-0.502	-0.38	-0.241	-0.085	0
π_{t-1}	2.747	2.082	1.420	0.792	0.238	0
$E_t \pi_{t+1}$	-0.482	0.115	0.697	1.229	1.639	1.851
$E_t \pi_{t+2}$	0.103	0.064	0.033	0.012	0.001	0
$E_t c_{t-2}^R$	0	-0.003	-0.009	-0.017	-0.029	0
$E_t c_{t-1}^R$	0.004	0.013	0.022	0.031	0.041	0
$E_t c_{t+1}^R$	-0.004	-0.012	-0.021	-0.031	-0.040	0
$E_t c_{t+2}^R$	0.003	0.007	0.008	0.007	0.003	0
e_t	1.275	1.227	1.154	1.056	0.932	0.859

Table 2.5: Values of reaction coefficients Ω_x in the interest rate rule (2.28) for different values of the share of rational forecasters α .

Appendix 2.C Implementation with model-consistent individual consumption expectations

The policy problem under commitment and the conventional Euler equation with model-consistent individual consumption expectations takes the following form:

$$\begin{aligned} \mathcal{L} = E_t \sum_{s=0}^{\infty} \beta^s & \left[\Gamma_1 y_{t+s}^2 + \Gamma_2 \pi_{t+s}^2 + \Gamma_3 \pi_{t+s-1}^2 + \Gamma_4 (c_{t+s}^R)^2 \right. \\ & + \Gamma_5 y_{t+s} c_{t+s}^R + \Gamma_6 \pi_{t+s} c_{t+s}^R + \Gamma_7 \pi_{t+s-1} c_{t+s}^R + \Gamma_8 \pi_{t+s} \pi_{t+s-1} + \Gamma_9 \pi_{t+s} y_{t+s} + \Gamma_{10} \pi_{t+s-1} y_{t+s} \\ & + \lambda_{1,t+s} [y_{t+s} - \alpha E_t y_{t+s+1} - (1-\alpha)\theta^2 y_{t+s-1} + \sigma [i_{t+s} - \alpha E_t \pi_{t+s+1} - (1-\alpha)\theta^2 \pi_{t+s-1}]] \\ & + \lambda_{2,t+s} [\pi_{t+s} - \alpha \beta E_t \pi_{t+s+1} - (1-\alpha)\beta\theta^2 \pi_{t+s-1} - \kappa y_{t+s} - e_{t+s}] \\ & \left. + \lambda_{3,t+s} [c_{t+s}^R - E_t c_{t+s+1}^R + \sigma (i_{t+s} - E_t \pi_{t+s+1})] \right]. \end{aligned} \quad (2.123)$$

The first order conditions are

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial y_{t+s}} : E_t \left\{ \beta^s [2\Gamma_1 y_{t+s} + \Gamma_5 c_{t+s}^R + \Gamma_9 \pi_{t+s} + \Gamma_{10} \pi_{t+s-1} + \lambda_{1,t+s} - \kappa \lambda_{2,t+s}] \right. \\ \left. - \beta^{s+1} (1-\alpha)\theta^2 \lambda_{1,t+s+1} - \beta^{s-1} \alpha \lambda_{1,t+s-1} \right\} \stackrel{!}{=} 0 \end{aligned} \quad (2.124)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \pi_{t+s}} : E_t \left\{ \beta^s [2\Gamma_2 \pi_{t+s} + \Gamma_6 c_{t+s}^R + \Gamma_8 \pi_{t+s-1} + \Gamma_9 y_{t+s} + \lambda_{2,t+s}] \right. \\ \left. + \beta^{s+1} [2\Gamma_3 \pi_{t+s} + \Gamma_7 c_{t+s+1} + \Gamma_8 \pi_{t+s+1} + \Gamma_{10} y_{t+s+1} - (1-\alpha)\theta^2 \sigma \lambda_{1,t+s+1} \right. \\ \left. - (1-\alpha)\beta\theta^2 \lambda_{2,t+s+1}] - \beta^{s-1} [\alpha \sigma \lambda_{1,t+s-1} + \alpha \beta \lambda_{2,t+s-1} + \sigma \lambda_{3,t+s-1}] \right\} \stackrel{!}{=} 0 \end{aligned} \quad (2.125)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial c_{t+s}^R} : E_t \left\{ \beta^s [2\Gamma_4 c_{t+s}^R + \Gamma_5 y_{t+s} + \Gamma_6 \pi_{t+s} + \Gamma_7 \pi_{t+s-1} \right. \\ \left. + \lambda_{3,t+s}] - \beta^{s-1} \lambda_{3,t+s-1} \right\} \stackrel{!}{=} 0 \end{aligned} \quad (2.126)$$

$$\frac{\partial \mathcal{L}}{\partial i_{t+s}} : E_t \left\{ \beta^s \sigma \lambda_{1,t+s} + \beta^s \sigma \lambda_{3,t+s} \right\} \stackrel{!}{=} 0. \quad (2.127)$$

Again, the index s can be dropped assuming commitment from a timeless perspective. Using $\lambda_{3,t} = -\lambda_{1,t}$ the FOCs can equivalently be written as

$$2\Gamma_1 y_t + \Gamma_5 c_t^R + \Gamma_9 \pi_t + \Gamma_{10} \pi_{t-1} + \lambda_{1,t} - \kappa \lambda_{2,t} - (1 - \alpha) \beta \theta^2 \lambda_{1,t+1} - \beta^{-1} \alpha \lambda_{1,t-1} \stackrel{!}{=} 0 \quad (2.128)$$

$$\begin{aligned} & 2\Gamma_2 \pi_t + \Gamma_6 c_t^R + \Gamma_8 \pi_{t-1} + \Gamma_9 y_t + \lambda_{2,t} \\ & + 2\Gamma_3 \beta \pi_t + \Gamma_7 \beta c_{t+1} + \Gamma_8 \beta \pi_{t+1} + \Gamma_{10} \beta y_{t+1} - (1 - \alpha) \beta \theta^2 \sigma \lambda_{1,t+1} \\ & - (1 - \alpha) \beta^2 \theta^2 \lambda_{2,t+1} + \beta^{-1} (1 - \alpha) \sigma \lambda_{1,t-1} - \alpha \lambda_{2,t-1} \stackrel{!}{=} 0 \end{aligned} \quad (2.129)$$

$$2\Gamma_4 c_t^R + \Gamma_5 y_t + \Gamma_6 \pi_t + \Gamma_7 \pi_{t-1} - \lambda_{1,t} + \beta^{-1} \lambda_{1,t-1} \stackrel{!}{=} 0. \quad (2.130)$$

(2.130) can be used to replace $\lambda_{1,t-1}$ and $\lambda_{1,t+1}$ with $\lambda_{1,t}$ in (2.55). Then, solving (2.55) for $\lambda_{1,t}$ and inserting in (2.129) yields a second-order difference equation in $\lambda_{2,t}$. A solution to this equation can in principle be substituted back into the difference equation, which would give a targeting rule. However, this solution is fairly complicated in which some parameter terms exponentially depend on time. The solution is available upon request. The resulting targeting rule and, hence, a reaction function would also be of such a complicated form where parameters exponentially depend on time. Consequently, no interpretable interest rate rule under commitment can be derived in this case.

Chapter 3

The Rationality Bias

Co-authored by Joep Lustenhouwer

3.1 Introduction

There is a growing consensus that models with homogeneous rational expectations cannot adequately approximate actual human behavior at the microeconomic or macroeconomic level. Even though there may be some highly sophisticated and well informed economic actors, a non-negligible fraction of the population might not be nearly as rational as assumed in theoretical models. The presence of agents with different levels of cognitive ability may lead to non-trivial interactions. An unforeseen economic shock will, for example, affect the expectations and, therefore, the economic decisions of different actors differently. Well-informed and sophisticated actors revise their expectations instantly and act according to these in an appropriate manner. However, less sophisticated actors may only observe the effects of the shock with a delay and make mistakes accordingly. Finally, the sophisticated actors may internalize the behavior of less sophisticated actors and change their own behavior accordingly. These differences and interactions between agents will have distributional as well as macroeconomic implications.

In response to a growing need for models that take account of cognitive imperfections, different strands of literature that incorporate bounded rationality have gained popularity (see, e.g. Evans and Honkapohja, 2012 and Gabaix, 2020). One of these strands of literature incorporates heterogeneous, possibly non-rational expectations in macro models (see, e.g. Branch and McGough, 2018). However, papers in this tradition have focused on aggregate dynamics, largely ignoring dynamics at the individual level. We want to fill this gap in the literature and explicitly consider how differences in

cognitive ability lead to differences in consumption and wealth, and how the magnitude of these differences is determined.

For this purpose, we set up a micro-founded model where a given fraction of agents is *fully rational* in the conventional sense (Muth, 1961), while the other fraction is *boundedly rational* as we will explain in detail below. Further, we keep track of the individual bond holdings of both groups. We identify the “rationality bias” of boundedly rational agents, defined as the deviation from the fully rational benchmark, as the driver of consumption and wealth heterogeneity. Moreover, we find a strong interaction between the size of the rationality bias and monetary policy.

Empirical evidence based on US inflation survey data suggests that, in the real world, expectations are indeed heterogeneous. Moreover, a sizable fraction of the population seems to follow simple backward-looking heuristics. For instance, Branch (2004) finds evidence of the presence of different expectation types in the Michigan Survey of Consumer Attitudes and Behavior. In particular, both naive expectations, where the last observed value of a variable is used as a best guess for the future, and a more sophisticated VAR heuristic are present in the data. Cornea-Madeira et al. (2019), further, find that a model with naive agents and agents that use a VAR approach to predict inflation based on its forward-looking relation with marginal costs fits actual inflation data well. Moreover, simple backward-looking heuristics are consistent with evidence from laboratory experiments (Assenza et al., 2019; Pfajfar and Žakelj, 2018), where a large degree of heterogeneity is found. Also, Fuhrer (2017) identifies slow-moving expectations as a source of macroeconomic persistence. He finds that ad-hoc assumptions like indexation and habit-formation to generate observed persistence become obsolete when survey expectations are incorporated into an otherwise standard DSGE model.

In our model, boundedly rational households form their expectations in a naive manner, consistent with Branch (2004) and Cornea-Madeira et al. (2019), and following earlier literature on heterogeneous expectations (Branch and McGough, 2009, 2010; De Grauwe, 2011; Gasteiger, 2014; Hommes and Lustenhouwer, 2019). The great advantage of the naive rule compared to more sophisticated backward-looking heuristics is that it allows for analytical tractability. Moreover, naive expectations perform well when the variable being forecast is highly persistent and are optimal when the variable follows a random walk.

Our boundedly rational households, further, base their consumption decision only on the variational intuition of the consumption Euler equation (Euler-equation learning), following Branch and McGough (2009). Boundedly rational households, hence, do not make decisions according to the infinite horizon learning approach of Preston

(2005) and Massaro (2013). Infinite horizon learning would require boundedly rational households to form expectations about all variables over an infinite horizon, and make optimal decisions based on these expectations. We believe that this would require too much cognitive ability of boundedly rational households.

The other fraction of households, on the other hand, is fully rational in the conventional sense. These households are fully aware of the presence of boundedly rational households. Further, rational households choose the optimal intertemporal allocation of consumption, labor supply and bonds based on all their optimality conditions, including their intertemporal budget constraint. Thus, we combine the Euler-equation learning approach of Branch and McGough (2009) and Honkapohja et al. (2011) for boundedly rational households with full rationality, in the conventional sense, for rational households. To the best of our knowledge, this is the first research project to do so. With this approach, we can contrast boundedly rational households that have considerable limitations in their cognitive ability with fully rational households. The presence of rational households, further, allows us to construct an intuitive measure of the “rationality bias” of boundedly rational households, as the deviation from the fully rational benchmark.

We find that considerable consumption and wealth heterogeneity arises when the economy is hit by shocks. The reason is that boundedly rational households respond differently with their consumption decision to shocks than rational households. In particular, the rationality bias of boundedly rational agents leads them to make non-optimal choices.

Moreover, it turns out that the rationality bias can be decomposed into three different components. One component, the consumption expectation bias, reflects a bias in expectations about an household’s own future consumption. A second component, the real interest rate bias, captures the differences in the subjective real interest rates. Finally, the third component, the preference shock expectation bias, arises whenever boundedly rational households have biased expectations about future realizations of the preference shock. We, further, find that the consumption expectation bias is mainly driven by boundedly rational household’s expectations about future aggregate output. The real interest rate bias, in turn, appears solely due to biased expectations about future inflation.

Each of the three biases leads boundedly rational households to deviate from the rational consumption decision. However, depending on the shocks that have hit the economy, the different biases may also work in opposite directions and counteract each other.

Next, we are interested in the magnitude of the rationality bias and consider how

the mistakes of boundedly rational households can be reduced. We show that the rationality bias completely disappears when the percentage deviation from steady state of the nominal interest rate and the inflation rate are equal in every period. This requires that the economy was initialized without consumption and wealth heterogeneity (as in steady state). In this case, the three components of the rationality biases exactly offset each other no matter what the realizations of shocks are.

However, eliminating the rationality bias comes along with high inflation volatility. From a welfare perspective, it might, therefore, not be desirable to completely eliminate the rationality bias. Even though consumption heterogeneity enters the utilitarian welfare loss of this model, the weight on inflation suggests that the households populating this economy prefer stable prices over homogeneous consumption (and output stability) by far.

As previously indicated, the literature on heterogeneous expectations (see for instance Branch and McGough (2018)) mainly focuses on aggregate dynamics and does not explicitly consider differences in consumption and wealth of the different agent types. Recent exceptions are Beqiraj et al. (2019) and Annicchiarico et al. (2019). The former calculate a measure of consumption inequality in their economy. The latter present differences in bond holdings that arise between some of their agents. The main interest of these studies, however, remains aggregate dynamics. To the best of our knowledge, there is no research project in this tradition that *focuses* on individual agent's decisions in a macro-setting.

The remainder of this chapter is organized as follows. The model is introduced in Section 3.2 where we put emphasis on the assumptions on bounded rationality. In Section 3.3, we study in detail the different components that comprise the rationality bias. Finally, in Section 3.4 we show that the rationality bias can be eliminated with certain monetary policy and briefly consider the welfare implications of such a policy. Section 3.5 concludes.

3.2 Model

In this section, we introduce heterogeneous rationality into a standard New Keynesian model. In particular, we assume a unit-mass continuum of households $i \in [0, 1]$ and firms $j \in [0, 1]$ as well as a monetary and fiscal authority. The population of households splits into two groups with a fixed size: rational households, who make up a fraction α of the population, and boundedly rational households, making up the remaining $1 - \alpha$ of the population.

There is an ongoing discussion in the literature (see e.g. Honkapohja et al., 2011) whether boundedly rational agents should be assumed to follow Euler-equation learning or rather infinite horizon learning. In the former case boundedly rational agents base their consumption decision only on the variational intuition of the consumption Euler equation. The latter case assumes boundedly rational agents to make optimal decisions based on their subjective forecasts of the entire future paths of all variables. Since we study differences in consumption and wealth that arise due to heterogeneity in cognitive ability, we let our two types of agents differ considerably in their degrees of rationality. For this reason, we stick with the less cognitive demanding Euler-equation learning approach for our boundedly rational agents. Rational agents are modelled as being fully rational instead. Thus, we do not follow Branch and McGough (2009) and the related literature on heterogeneous expectations by letting agents with rational expectations also use Euler-equation learning. Instead, we combine the Euler-equation learning of boundedly rational agents with fully rational optimization in the conventional sense of rational agents. Detailed descriptions of both household types, including their individual expectation formation schemes, are included in Sections 3.2.2 and 3.2.3.

Further, we assume perfect consumption insurance within the groups, so that the two types of households can be interpreted as two different representative agents. In order to focus on heterogeneity in households only, we assume that all firms are run by rational managers.

3.2.1 The non-linear model

First, we introduce the basic model elements in non-linear form. The log-linearized version, including a specification of each household type and policy rules, is given in a subsequent subsection.

Households

Households of type i with $i \in \{R, B\}$ optimize their expected lifetime utility $E_t^i \sum_{t=0}^{\infty} \beta^t \Upsilon_t U_t^i$ subject to their real flow budget constraint

$$C_t^i + \bar{Y} b_t^i \leq W_t H_t^i + \bar{Y} \frac{I_{t-1}}{\Pi_t} b_{t-1}^i + D_t - T_t, \quad (3.1)$$

where E_t^i is the type-specific expectations operator, $\beta < 1$ the subjective discount factor, U_t^i period-utility of type i and Υ_t is a preference shock. C_t^i denotes individual consumption, T_t is a lump-sum tax and W_t the real wage rate, which are equal across groups.

$b_t^i = B_t^i / (P_t \bar{Y})$ are the real bond holdings of household i as a fraction of steady state output, where \bar{Y} is the steady state of output, Y_t . B_t^i refer to nominal government bonds, which are the only asset in the economy. Both rational and boundedly rational households can hold positive and negative positions of these bonds and can trade them with each other as well as directly with the government. Note that the price of the nominal government bonds is given by the inverse of the gross nominal interest rate, I_t , which is controlled by the central bank.

Further, H_t^i denotes individual hours, $\Pi_t = P_t / P_{t-1}$ the inflation rate where P_t is the aggregate price level in period t and D_t dividends that households obtain as shareholders of the firms.

Period utility is of CES-form and given by

$$U_t^i = \frac{(C_t^i)^{1-\sigma}}{1-\sigma} - \frac{(H_t^i)^{1+\gamma}}{1+\gamma}, \quad (3.2)$$

where σ is the coefficient of relative risk aversion and γ the labor supply elasticity.

Firms

We assume that firms are run by rational managers that are appointed by households (who are the shareholders of firms). In the end, this assumption allows us to isolate the effect of *heterogeneous rationality of households* on differences in consumption and wealth. The assumption that all firm managers are rational can be justified by the fact that rational managers will always perform better than boundedly rational ones. Therefore, if given the chance, all firm owners would prefer to appoint rational managers.

Our further assumptions about the firm side of this model are entirely standard. We assume that production is linear in labor and that total factor productivity is constant over time and equal to one. We, further, introduce Calvo pricing where a firm j is able to re-set its price only with a given probability of $1 - \omega$ in each period. Thus, each firm j maximizes its current and future discounted flow of profits by setting its price $P_t(j)^*$ subject to the demand of differentiated goods. We present the derivation of the Phillips curve in the Appendix 3.A.4 and give the log-linear version in the next section.

Government and market clearing

The government is assumed to finance its spending via government debt and collecting lump-sum taxes. The real government budget constraint reads

$$b_t = g_t - \frac{T_t}{\bar{Y}} + s \frac{H_t}{\bar{Y}} W_t + \frac{I_{t-1}}{\Pi_t} b_{t-1}, \quad (3.3)$$

where $g_t = G_t/\bar{Y}$ is government spending relative to steady state output. Further, $s = 1/\eta$ is a subsidy provided to firms (as in Galí, 2015) that makes the steady state of our model efficient. We provide the description of the evolution of government spending, taxes and the nominal interest rate in the log-linear version of the economy.

Goods and bond markets clear according to

$$Y_t = (\alpha C_t^R + (1 - \alpha) C_t^B) + \bar{Y} g_t, \quad (3.4)$$

$$b_t = \alpha b_t^R + (1 - \alpha) b_t^B. \quad (3.5)$$

3.2.2 Expectations

As indicated above, our two types of households employ different forecasting rules. Rational households are sophisticated enough to compute an optimal linear forecast, i.e. they use the expected value $E_t^R x_{t+1} = E_t x_{t+1}$. Boundedly rational households, on the other hand, use the last observed value (x_{t-1}) of a variable as their best guess for the future with the following naive heuristic: $E_t^B x_{t+1} = x_{t-1}$.

3.2.3 The log-linear economy

For our policy analysis, we use a log-linear version of the model that is obtained by approximating around a non-stochastic zero-inflation steady state.

Individual consumption, labor and bonds

Households first-order conditions, implied by the optimization problem outlined in Section 3.2.1, are given in log-linearized terms by

$$\hat{c}_t^i = E_t^i \hat{c}_{t+1}^i - \frac{1}{\sigma} (i_t - E_t^i \pi_{t+1} - v_t + E_t^i v_{t+1}), \quad (3.6)$$

$$w_t = \gamma h_t^i + \sigma \hat{c}_t^i, \quad (3.7)$$

$$\hat{b}_t^i = h_t^i + w_t + \beta^{-1} \hat{b}_{t-1}^i + \bar{b} \beta^{-1} (i_{t-1} - \pi_t) + \eta^{-1} d_t - (1 - \bar{g}) \hat{c}_t^i - \frac{\bar{T}}{\bar{Y}} \tau_t, \quad (3.8)$$

where lower-case letter indicate a log-deviation from steady state (and τ_t indicates log-deviation of taxes from steady state). Further, we denote $\hat{b}_t^i = b_t^i - \bar{b}$ and use $\bar{D} = \bar{Y} - (1 - s)\bar{W}\bar{H} = \eta^{-1}\bar{Y}$ where $\bar{W} = 1$. All steady state values are derived in Appendix 3.A.5.

Equation (3.6) determines the households' optimal consumption path depending on its subjective real interest rate and preference shock (expectations). The optimal labor supply is pinned down by the optimality condition (3.7). Equation (3.8) shows the household's log-linear budget constraint.

For rational households the budget constraint (3.8) is given by

$$\hat{b}_t^R = h_t^R + w_t + \beta^{-1}\hat{b}_{t-1}^R + \bar{b}\beta^{-1}(i_{t-1} - \pi_t) + \eta^{-1}d_t - (1 - \bar{g})c_t^R - \frac{\bar{T}}{\bar{Y}}\tau_t. \quad (3.9)$$

Iterating (3.9) until infinity, substituting for all choice variables and plugging in expectations, $E_t^R = E_t$, gives the consumption decision of rational agents:

$$\begin{aligned} c_t^R = & \zeta\hat{b}_{t-1}^R + \zeta\bar{b}(i_{t-1} - \pi_t) + \zeta\beta E_t \sum_{s=t}^{\infty} \beta^{s-t} [\Gamma_y y_s - \Gamma_g \hat{g}_s - \Gamma_\tau \tau_s] \\ & - \frac{(1 - \zeta\bar{b}\sigma)\beta}{\sigma} E_t \sum_{s=t}^{\infty} \beta^{s-t} (i_s - \pi_{s+1}) + \frac{\beta}{\sigma} E_t \sum_{s=t}^{\infty} \beta^{s-t} (v_s - v_{s+1}) \end{aligned} \quad (3.10)$$

where composite parameters ζ , Γ_y , Γ_g , Γ_τ and the derivation of (3.10) are given in Appendix 3.A.1. Thus, fully rational households make predictions over the *entire paths* of output, government spending, taxes, real interest rates and preference shocks *until infinity*. Further, rational households make optimal consumption decisions based on these predictions.

For boundedly rational households, we consider such an approach to require too much cognitive load. Therefore, our boundedly rational agents follow Euler-equation learning and believe that all other agents will form the same beliefs as they do (higher-order beliefs assumption) as in Branch and McGough (2009). This implies that boundedly rational agents will neglect their intertemporal budget constraint as an *optimality condition*. Boundedly rational agents are, further, assumed to know that market clearing holds. Also, we assume that boundedly rational agents believe that all other agents will satisfy the consumption Euler equation in the same way they themselves do. Finally, we require boundedly rational agents to be able to iterate their consumption Euler equation until an arbitrary period N which they consider to be the "long-run". The horizon N could in principle be infinite but it can also be finite. This will not change the result.

In Appendix 3.A.2 it is shown that under the above assumptions the consumption decision of boundedly rational agents is given by

$$c_t^B = \frac{1}{1-\bar{g}} E_t^B y_{t+1} + E_t^B (c_N^B - c_N) - \frac{1}{\sigma} [i_t - E_t^B \pi_{t+1} - v_t + E_t^B v_{t+1}] - \frac{1}{1-\bar{g}} E_t^B \hat{g}_{t+1}. \quad (3.11)$$

Further, we assume that when boundedly rational agents have more current (beginning-of-period) wealth than the average, they will expect to be able to consume more than the average in the “long-run” N . In mathematical terms: $E_t^B (c_N^B - c_N) = \psi(\hat{b}_{t-1}^B - \hat{b}_{t-1})$. This assumption ensures that we can rule out equilibria in which the bond holdings of individual households are explosive. Put differently, households are not able to roll over their debt until infinity. This approach is similar to including bonds in the utility function as in Kurz et al. (2013) and to quadratic adjustment costs of bonds in the budget constraint as proposed by Schmitt-Grohé and Uribe (2003). Further, this assumption re-introduces concerns for the intertemporal life-time budget constraint and a transversality condition to the decision making procedure of boundedly rational agents but in a rather behavioral manner. We choose the parameter ψ to be high enough to rule out explosiveness but also low enough to have a small impact on short-run dynamics.

Plugging in expectations, $E_t^B x_{t+1} = x_{t-1}$, and the assumption about “long-run” beliefs gives

$$c_t^B = \frac{1}{1-\bar{g}} y_{t-1} + \psi(\hat{b}_{t-1}^B - \hat{b}_{t-1}) - \frac{1}{\sigma} [i_t - \pi_{t-1} - v_t + v_{t-1}] - \frac{1}{1-\bar{g}} \hat{g}_{t-1}. \quad (3.12)$$

Comparing (3.12) with (3.10) shows that we impose far less cognitive load on boundedly rational households than on rational households as discussed above. While boundedly rational households are forward looking in the sense that they forecast one-period-ahead (and the “long-run”), they become ultimately backward-looking due to their behavioral heuristics.

Bond holdings of rational and boundedly rational agents are given by (3.8), where we substitute for hours and wages. This gives

$$\hat{b}_t^i = \beta^{-1} \hat{b}_{t-1}^i - \Gamma_c c_t^i + \Gamma_y y_t + \bar{b} \beta^{-1} (i_{t-1} - \pi_t) - \Gamma_g \hat{g}_t - \Gamma_\tau \tau_t \quad (3.13)$$

where the composite parameter Γ_c is given in the Appendix 3.A.1.

Aggregate equations

Inserting (3.10) and (3.12) into the log-linearized goods market clearing equation, $y_t = (1 - \bar{g})(\alpha c_t^R + (1 - \alpha)c_t^B) + \hat{g}_t$, and writing this equation recursively yields aggregate output:

$$\begin{aligned} y_t = & \Phi_1 E_t y_{t+1} + \Phi_2 y_{t-1} + \Phi_3 E_t \pi_{t+1} - \Phi_4 \pi_t + \Phi_5 \pi_{t-1} + \Phi_6 i_{t-1} - \Phi_7 i_t \\ & + \Phi_8 E_t i_{t+1} - \Phi_9 \hat{b}_{t-1}^R + \Phi_{10} b_t^R - \Phi_{11} \hat{g}_{t-1} + \Phi_{12} \hat{g}_t - \Phi_{13} E_t \hat{g}_{t+1} \\ & - \Phi_{14} \tau_t + \Phi_{15} \hat{b}_{t-1} - \Phi_{16} \hat{b}_t + \Phi_{17} v_t - \Phi_{18} E_t v_{t+1} - \Phi_{19} v_{t-1} \end{aligned} \quad (3.14)$$

where coefficients and the derivation of (3.14) are given in the Appendix 3.A.3.

Note that (3.14) depends on individual bonds of rational households and on aggregate bonds. This stems from the fact that both the consumption decision of rational households (3.10) and the consumption decision of boundedly rational households (3.12) depend on individual bonds. Bond holdings of boundedly rational households are eliminated in (3.14) using bond market clearing (3.5). All other lagged terms in Equation (3.14) are due to the backward-looking expectations of boundedly rational households. Forward-looking terms are in turn rational expectations of rational households.

Since all firm managers are rational, inflation follows a standard forward-looking Phillips-curve. As shown in Appendix 3.A.4, this implies

$$\pi_t = \delta \left(\gamma + \frac{\sigma}{1 - \bar{g}} \right) y_t - \delta \sigma (1 - \bar{g})^{-1} \hat{g}_t + \beta E_t [\pi_{t+1}] + \mu_t, \quad (3.15)$$

where $\delta = ((1 - \omega\beta)(1 - \omega))/\omega$ and μ_t being a cost-push shock.

The log-linearized government budget constraint is given by

$$\hat{b}_t = \hat{g}_t - \Gamma_\tau \tau_t + s(w_t + h_t) + \bar{b}\beta^{-1}(i_{t-1} - \pi_t) + \beta^{-1}\hat{b}_{t-1}. \quad (3.16)$$

For simplicity, we assume that government spending remains at its steady state level ($\hat{g}_t = 0$). Further, the costs of the subsidy, $s(w_t + h_t)$, are directly financed by lump-sum taxes. Additionally, taxes are assumed to respond to beginning-of-period debt. Lump-sum taxes, therefore, evolve as

$$\tau_t = \phi_{b,\tau} \hat{b}_{t-1} + \frac{s}{\Gamma_\tau} \left(1 + \gamma + \frac{\sigma}{1 - \bar{g}} \right) y_t. \quad (3.17)$$

where we used $h_t = y_t$ and $w_t = (\gamma + \sigma/(1 - \bar{g})) y_t$ (since $\hat{g}_t = 0$).

The evolution of aggregate debt can, thus, be written as

$$\hat{b}_t = (\beta^{-1} - \Gamma_\tau \phi_{b,\tau}) \hat{b}_{t-1} + \bar{b} \beta^{-1} (i_{t-1} - \pi_t). \quad (3.18)$$

Moreover, log-linearizing the bond market clearing equation gives

$$\hat{b}_t = \alpha \hat{b}_t^R + (1 - \alpha) \hat{b}_t^B. \quad (3.19)$$

For central bank policy we assume a standard Taylor rule based on contemporaneous output and inflation, i.e.

$$i_t = \phi_\pi \pi_t + \phi_y y_t. \quad (3.20)$$

The preference shock and cost push-shock in the linearized model are assumed to follow AR(1) processes:

$$v_t = \rho_v v_{t-1} + \epsilon_{v,t} \quad (3.21)$$

$$\mu_t = \rho_\mu \mu_{t-1} + \epsilon_{\mu,t} \quad (3.22)$$

where $\epsilon_{\mu,t}$ and $\epsilon_{v,t}$ are uncorrelated IID shocks.

3.3 The components of the rationality bias

In this section, we analyze differences in consumption and wealth that arise due to different degrees of rationality. Differences in wealth between boundedly rational and rational households arise as a consequences of differences in their consumption. We, hence, focus on the latter.

The difference in individual consumption arises because boundedly rational households follow Euler-equation learning, believe that all other households form the same expectations as they themselves do, and use a simple naive forecast heuristic to form these expectations. If a boundedly rational household would not have these limitations, it would act as a rational household. Hence, the difference $\Delta_i c_t^i = c_t^B - c_t^R$ can be interpreted as the *bias of a boundedly rational household*. We, therefore, label this difference the *Rationality bias*.

We first decompose the rational bias into three different components, in Section 3.3.1. Next, in Section 3.3.2, we study the different components in some detail. In Section 3.3.3, we use impulse response functions to provide further intuition into the rationality bias and its components.

3.3.1 Decomposing the rationality bias

Using the individual consumption Euler equation, (3.6), of both household types we can write

$$\Delta_i c_t^i = c_t^B - c_t^R = \underbrace{(E_t^B c_{t+1}^B - E_t c_{t+1}^R)}_{\text{consumption exp. bias}} - \underbrace{\frac{1}{\sigma}(rr_t^B - rr_t)}_{\text{real int. bias}} - \underbrace{\frac{1}{\sigma}(E_t^B v_{t+1} - E_t v_{t+1})}_{\text{preference shock exp. bias}}. \quad (3.23)$$

Equation (3.23) shows that we are able to decompose this bias into three individual sources: the consumption expectation bias, the real interest rate bias and the preference shock expectation bias.

The consumption expectation bias arises whenever the expectations of boundedly rational households about their own future consumption deviate from the rational expectation about future rational consumption. A positive consumption expectation bias leads, in isolation, to higher consumption of boundedly rational households relative to the consumption of rational households.

The real interest rate bias emerges whenever the subjective (ex-ante) real interest rate of boundedly rational households, rr_t^B , is different from the objective (ex-ante) real interest rate. The latter coincides with the subjective real interest rate of rational households. When the real interest rate bias is positive, the consumption of boundedly rational households, *ceteris paribus*, will be lower compared to the consumption of rational households.

Finally, the preference shock expectation bias arises when boundedly rational households expect the next periods' preference shock to be different from the rational expectations about this shock. A positive preference shock expectation bias induces boundedly rational households, *ceteris paribus*, to lower their own consumption relative to the consumption of rational households.

To summarize: the rationality bias $\Delta_i c_t^i$ depends *positively* on the consumption expectation bias, *negatively* on the real interest rate bias and *negatively* on the preference shock expectation bias.

3.3.2 Individual components

Now that we know how the rationality bias depends on its individual components, we can further study the individual components themselves. In order to do so, we need to clarify what, for instance, boundedly rational consumption expectations actually are.

In Section 3.2.3 we derived the consumption decision of boundedly rational agents

(3.12) that can be rearranged as

$$c_t^B = \underbrace{\frac{1}{1-\bar{g}}(y_{t-1} - \hat{g}_{t-1}) + \psi(\hat{b}_{t-1}^B - \hat{b}_{t-1})}_{=E_t^B c_{t+1}^B} - \frac{1}{\sigma} \underbrace{[i_t - \pi_{t-1}]}_{=rr_t^B} + \frac{1}{\sigma} [v_t - \underbrace{v_{t-1}}_{=E_t^B v_{t+1}}]. \quad (3.24)$$

Comparing Equation (3.24) with the Euler equation of boundedly rational agents (3.6) shows explicitly that their consumption expectations, their subjective real interest rate and their preference shock expectations are given by

$$E_t^B c_{t+1}^B = \frac{1}{1-\bar{g}} \underbrace{(y_{t-1} - \hat{g}_{t-1}) + \psi(\hat{b}_{t-1}^B - \hat{b}_{t-1})}_{=c_{t-1}} \quad (3.25)$$

$$rr_t^B = i_t - \pi_{t-1} \quad (3.26)$$

$$E_t^B v_{t+1} = v_{t-1}. \quad (3.27)$$

These equations generally reflect that expectations of boundedly rational households are equal to the last observed value of the variable being forecast. Moreover, equation (3.25) shows that when boundedly rational households expect aggregate consumption next period to be high, they also expect their own individual consumption to be high in the next period. Additionally, when boundedly rational households own more wealth than the average, they expect to be able to consume more in the future, as discussed in Section 3.2.3. Note, however, that the main component of boundedly rational consumption expectations (3.25) is boundedly rational expectations about output, y_{t-1} . The reason is that as ψ will assumed to be small and that $\hat{g}_t = 0$ holds for all periods.

Inserting (3.25)-(3.27) into (3.23) yields

$$\Delta_i c_t^i = c_t^B - c_t^R = \underbrace{\left(\frac{1}{1-\bar{g}} (y_{t-1} - \hat{g}_{t-1}) + \psi(\hat{b}_{t-1}^B - \hat{b}_{t-1}) - E_t c_{t+1}^R \right)}_{\text{consumption exp. bias}} \quad (3.28)$$

$$+ \underbrace{\frac{1}{\sigma} (\pi_{t-1} - E_t \pi_{t+1})}_{\text{neg. real int. bias}} - \underbrace{\frac{1}{\sigma} (v_{t-1} - E_t v_{t+1})}_{\text{preference shock exp. bias}}.$$

The second term in (3.28) reveals that a negative real interest rate bias merely is a positive inflation expectation bias. Whenever the inflation expectations of boundedly rational households are above the corresponding rational expectations, the subjective real interest rate of boundedly rational households will be lower than the real interest rate of rational households. As a consequence, consumption of boundedly rational households will increase compared to consumption of rational households, ceteris

paribus.

In principle, the individual biases can amplify or counteract each other, which we describe in detail in Section 3.3.3. In Section 3.4, we will show that these biases exactly off-set each other when the percentage deviations from steady state of the inflation rate and the nominal interest rate are equal and when the economy was initialized without heterogeneity (as in steady state).

3.3.3 Illustrating the components of the Rationality Bias

In order to get further insight into the rationality bias, we numerically study impulse responses of our model. In particular, we study in detail the different components of the rationality bias of boundedly rational agents in a dynamic setting. We, further, study the consequences of the rationality bias for consumption and wealth heterogeneity and aggregate dynamics.

Parameterization

For this exercise, we parameterize the model as shown in Table 3.1. We emphasize, however, that this parameterization only serves the purpose of illustration, and that our main results are derived analytically without parameterizing the model. We choose most of the parameter values for our calibration to be in line with standard literature.

Micro parameter	$\beta = 0.99$	$\sigma = 2$	$\eta = 7.84$	$\gamma = 2$	$\omega = 0.75$
Expectations	$\alpha = 0.5$	$\psi = 0.05$			
Policy parameter	$\phi_y = 0.2$	$\phi_\pi = 1.5$	$\phi_{b,\tau} = 1$	$\bar{g} = 0.21$	$\bar{\tau} = 0.375$
Shock parameters	$\rho_\mu = 0.85$	$\rho_v = 0.85$	$sd(\epsilon_{\mu,t}) = 0.025$	$sd(\epsilon_{v,t}) = 0.025$	

Table 3.1: Baseline parameterization

As indicated in Section 3.2.3, we pick a small number for the responsiveness of boundedly rational consumption with respect to the difference of individual and aggregate wealth ψ , i.e., $\psi = 0.05$. This value of ψ is high enough to rule out explosive equilibria but also low enough to have a small impact on the short-run dynamics.

In the baseline calibration, we assume a fraction of 0.5 of rational and boundedly rational agents. In Appendix 3.B, we show that the results described in this section are robust to the calibration of α (the fraction of rational agents in the economy). In particular, we find that, although the size of the rationality bias and differences in bond holdings depend on α , the general patterns described below always emerge.

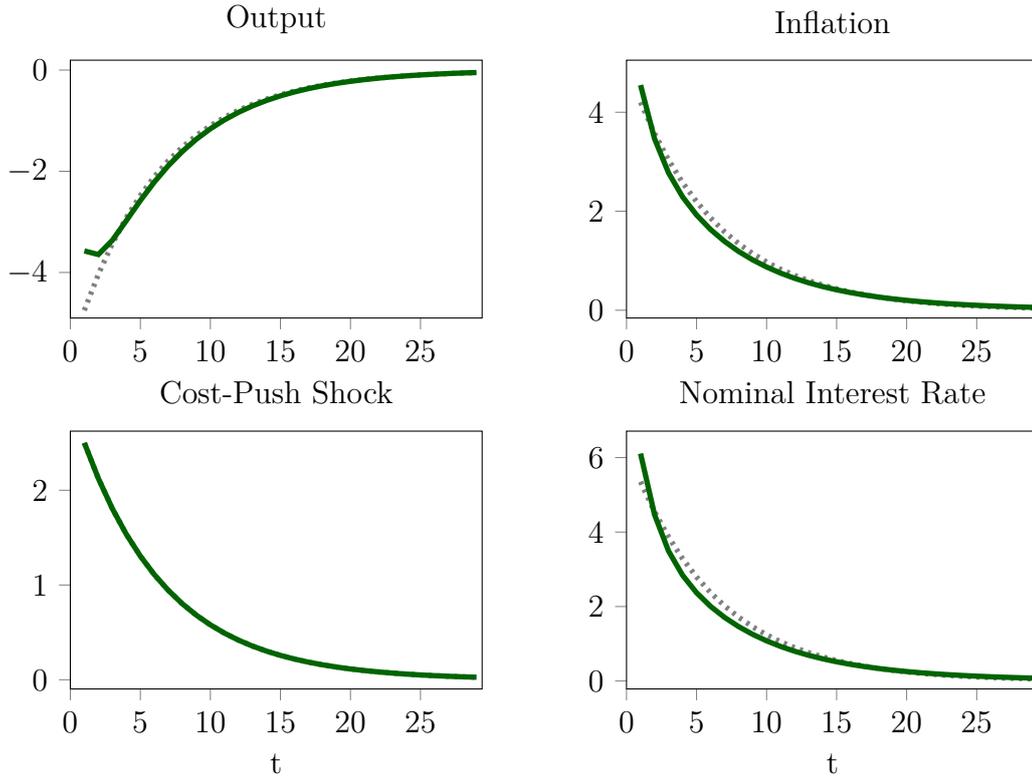


Figure 3.1: Impulse responses of aggregate variables to a persistent 2.5% cost-push-shock under the baseline calibration. The gray-dotted curve in each panel shows the homogeneous rational benchmark, i.e., $\alpha = 1$. The green curves show the aggregate dynamics under heterogeneity, for $\alpha = 0.5$.

Impulse responses to a cost-push shock

Figure 3.1 depicts the impulse responses of aggregate variables following a persistent 2.5 percent (one standard deviation) cost-push under this parameterization. The impulse responses show that the cost-push shock raises inflation substantially more than one-for-one under both homogeneous rational expectations (dotted gray) and heterogeneous rationality (solid green). This is due to the persistence of the shock.

The central bank reacts to the shock by increasing the nominal interest rate. This reduces output, but limits the increase in inflation. Note that, in the initial period, output is not as low in the model with heterogeneous rationality as under homogeneous rational expectations. This is because boundedly rational households consume considerably more than rational households in that period, as will be discussed below.

Now, let's turn in detail to the different components of the rationality bias. In the impulse responses to a cost push shock, there is no preference shock, and, hence, the preference shock expectation bias is zero. Thus, we need to consider only two different components of the rationality bias: first, the real interest rate bias and, second, the

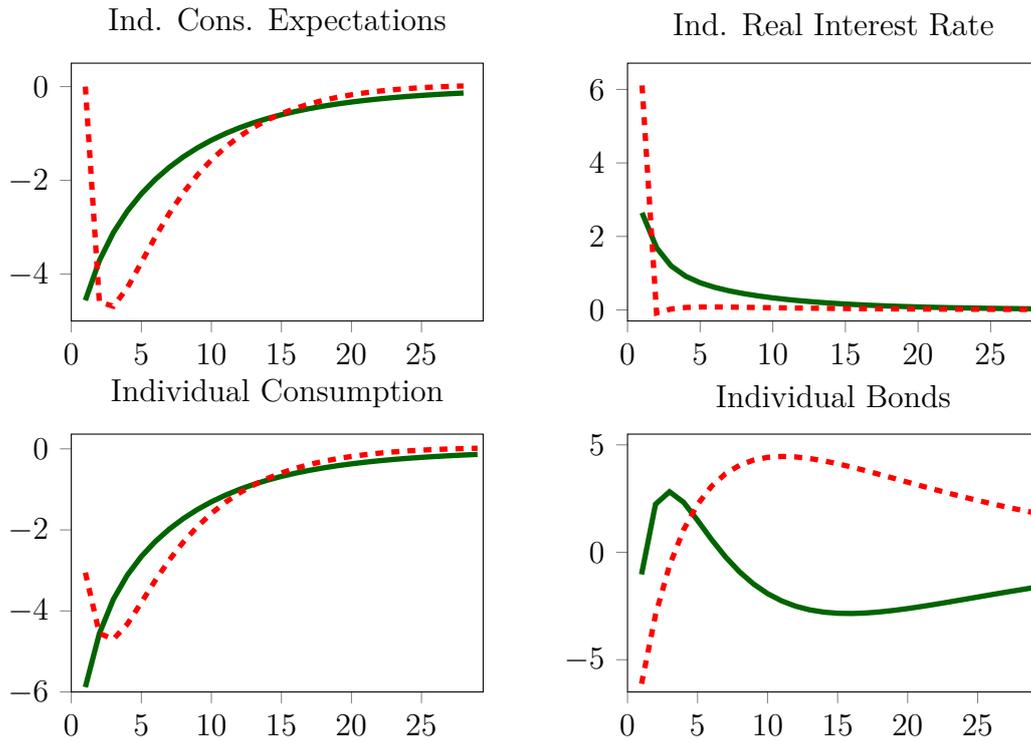


Figure 3.2: Impulse responses of individual variables to a persistent 2.5% cost-push-shock under the baseline calibration with $\alpha = 0.5$. The green curves depict the individual variables of rational agents and the red-dashed curve the individual variables of boundedly rational agents.

consumption expectation bias. In order to see these biases at work, we need to consider the individual consumption expectations and subjective real interest rates of boundedly rational and rational households. Figure 3.2 plots these variables, as well as individual consumption and bond holdings.

Starting with the top right panel, it can be seen that the subjective real interest rate of boundedly rational households differs considerably from that of rational households. While the rational households' real interest rate (solid green) increases by less than three percent on impact, the subjective real interest rate of boundedly rational households (dashed red) initially increases by *six percent*. This is because boundedly rational households base their inflation expectations on the previous period and have not observed the effects of the shock yet when they form expectations in the beginning of the shock period. Thus, they do not anticipate an increase in future inflation in that period. The subjective real interest rate of boundedly rational households, hence, increases as much as the nominal interest rate.

Rational households, on the other hand, already anticipate future inflation in the shock period, which reduces their real interest rate. We can conclude that the real interest rate bias is positive in the first period.

However, from period two onward, inflation is falling. As a consequence, boundedly rational households consistently overestimate future inflation. This results in a subjective real interest rate that is persistently below the real rate of rational households, and, hence, a negative real interest rate bias in these periods. Eventually, inflation expectations align when model variables approach the steady state, which results also in an alignment of the subjective real interest rates.

These difference in subjective real interest rates partly explain the differences in consumption and bond holdings observed in the bottom two panels of Figure 3.2. But, to fully explain the consumption and bond dynamics, we need to consider the consumption expectation bias of boundedly rational households as well.

The top left panel of Figure 3.2 plots the expectations of rational and boundedly rational households about their own future consumption. In the shock period, rational households already expect lower future own consumption. This is because they correctly anticipate positive future real interest rates due to the shock persistence and the response of the central bank. Boundedly rational households, on the other hand, do not expect any deviations from steady state yet and, hence, expect zero future consumption. Boundedly rational households, hence, have a positive consumption expectation bias in the initial period.

It turns out that this positive consumption expectation bias outweighs the positive real interest rate bias of boundedly rational households in the first period. As a consequence, boundedly rational households consume more than rational agents in that period (see bottom-left panel). For this reason, boundedly rational households lose more bonds than rational households (see bottom-right panel).¹⁷

In the next period, boundedly rational households observe the recession caused by the shock and expect low future output and consumption. In fact, boundedly rational households now have a negative consumption expectation bias. However, in the period after the shock, the negative real interest rate bias is larger than the negative consumption expectation bias. Therefore, boundedly rational households still consume more than rational households.

This is different in later periods. From period three onward, the negative consumption expectation bias is larger than the negative real interest rate bias, and boundedly rational households consume less than rational households. In the medium-run, boundedly rational households are, therefore, accumulating more bonds than rational households.

It follows that wealth differences in the short-run arise because rational households

¹⁷Even though households cut consumption, initially their real bond holdings go down due to high inflation.

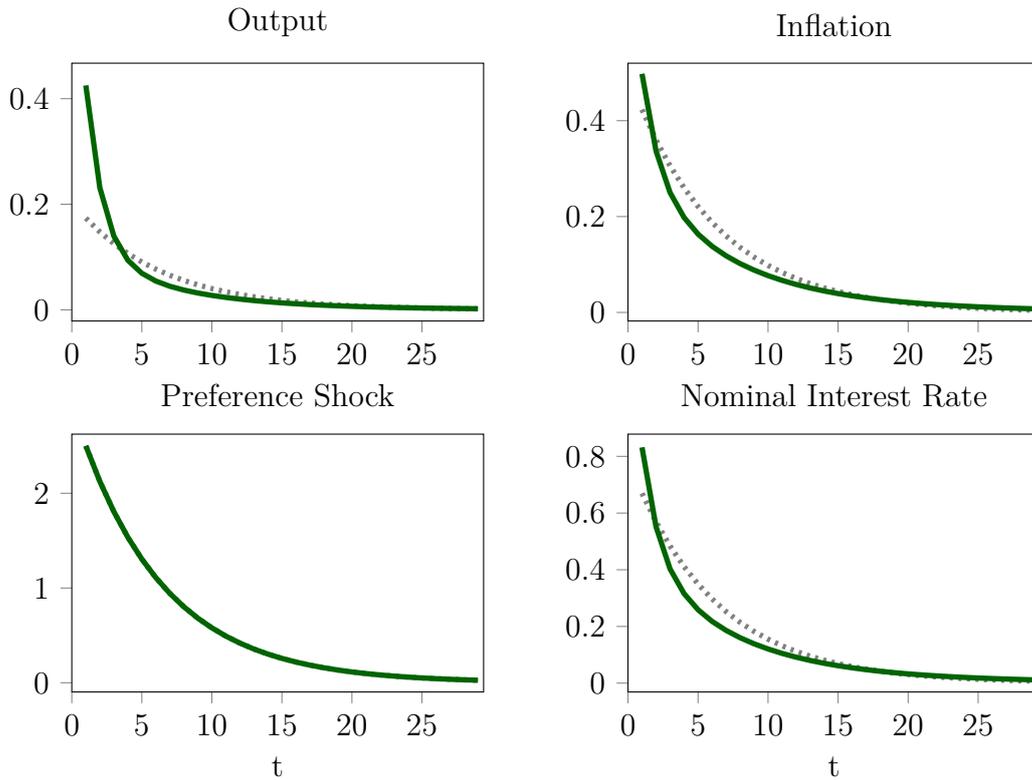


Figure 3.3: Impulse responses for a persistent 2.5% preference shock under baseline calibration. The gray-dotted curve in each panel shows the homogeneous rational benchmark, i.e. $\alpha = 1$. The green curves in the upper and lower row show the aggregate dynamics under heterogeneity, for $\alpha = 0.5$. The green curves in the middle row depict the individual variables of rational households and the red-dashed curve the individual variables of boundedly rational households for $\alpha = 0.5$.

own more bonds, while in the medium-run boundedly rational households are more wealthy. The latter allows boundedly rational households to consume out of wealth in the medium-run so that their consumption can pick up again. Note, though, that since rational households are fully optimizing using optimal linear forecasts, the consumption utility losses of boundedly rational households over the life-cycle are higher than that of rational agents by definition.¹⁸

¹⁸As one can see, individual related variables behave quite differently between household types to the cost-push shock. Yet, differences between the homogeneous rational benchmark and the heterogeneous rationality model at the aggregate level seem to be somewhat less pronounced. The reason is that the shock hits firms directly which have homogeneous rational expectations while household are affected more indirectly. The preference shock further below which hit households directly and, therefore, aggregate demand, one can observe a more distinct initial reaction of output, for instance, between the homogeneous rational benchmark and the heterogeneous rationality model.

Impulse responses to a preference shock

Next, we consider the impulse responses to a persistent preference shock. Figure 3.3 plots the responses of aggregate variables. The shock shifts the preferences of both household types such that they want to consume more and save less in the current period. This increases aggregate output and puts upward pressure on inflation and the nominal interest rate.

Note that in the initial period, output and inflation increase considerably more under heterogeneous rationality (solid green) than under homogeneous rational expectations (dotted gray). This is because of the rationality bias of boundedly rational households, which is discussed in detail below.

As can be seen in Equation (3.6), households' consumption does not only depend on today's shock realization but also on their expectations about future preference shock realizations. Rational households expect a lower but still substantial shock realization tomorrow due to the shock persistence. On the other hand, the expectations of boundedly rational households are not yet affected by the shock in the first period. Thus, boundedly rational households expect a zero future shock realization and have a considerable negative preference shock expectation bias in the shock period. This puts upward pressure on their consumption.

From period two onward, boundedly rational households consistently overestimate the future shock realization. Boundedly rational households then have a negative preference shock expectations bias, which puts downward pressure on their consumption in these periods.

However, the actual difference in consumption between rational and boundedly rational households also depends on the real interest rate bias and the consumption expectation bias. Figure 3.4 plots the individual real interest rate and consumption expectations of both household types, as well as their individual consumption and bond holdings.

As can be seen in the top right panel of Figure 3.4, the subjective real interest rate of boundedly rational households increases substantially more than the objective (rational) one in the first period. This is so because boundedly rational households expect next period's inflation to be zero, just as in the previous section.

Similarly, output expectations of boundedly rational households are zero as well in the initial period and thus below rational expectations (top left panel). Hence, both the real interest rate bias and consumption expectations bias counteract the preference shock expectation bias. The net effect of all three biases on consumption of boundedly rational households is, however, still substantially positive, as can be seen in the bottom

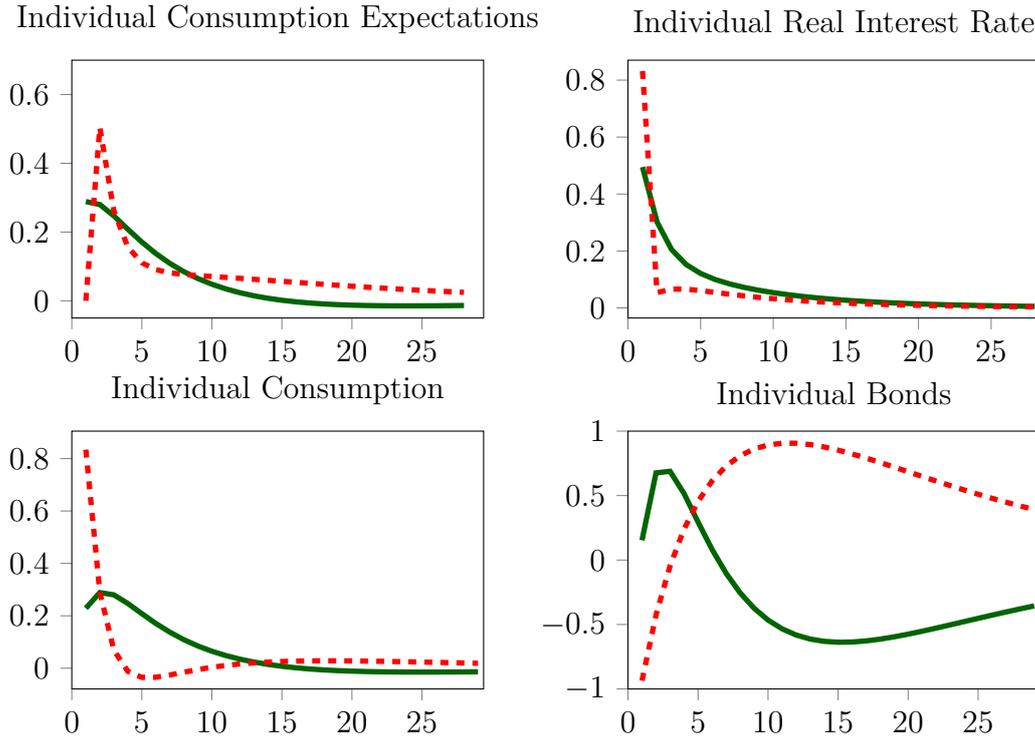


Figure 3.4: Impulse responses for a persistent 2.5% preference shock under baseline calibration. The gray-dotted curve in each panel shows the homogeneous rational benchmark, i.e. $\alpha = 1$. The green curves in the upper and lower row show the aggregate dynamics under heterogeneity, for $\alpha = 0.5$. The green curves in the middle row depict the individual variables of rational households and the red-dashed curve the individual variables of boundedly rational households for $\alpha = 0.5$.

left panel of Figure 3.4.

Further, from period two onward, the subjective real interest rate of boundedly rational households is lower than that of rational households. The reason is that boundedly rational agents overestimate inflation consistently as it falls. Additionally, output expectations of boundedly rational households, and, therefore, their own consumption expectations, are biased upwards as output falls. Both the real interest rate bias and the consumption expectations bias, thus, put upward pressure on consumption in these periods, again counteracting the preference shock expectation bias. However, also in these periods, the preference shock expectation bias dominates so that consumption of boundedly rational households is lower than that of rational households.

This consumption pattern has similar implications for individual bonds as under the cost-push shock, as can be seen in the bottom right panel of Figure 3.4. Also after a preference shock, boundedly rational households have less wealth than rational households in the short run, but more bond holdings in the medium-run. Boundedly rational

households are then able to consume their savings later on so that their consumption eventually slightly overshoots the consumption of rational households. However, note again that boundedly rational households must lose relatively more utility compared to rational households over the life-cycle.

3.4 The magnitude of the rationality bias

So far, we have decomposed the rationality bias in three individual components, and have analyzed the effects of these biases on the micro and macro dynamics. In this section, we study how the magnitude of the rationality bias interacts with monetary policy.

3.4.1 Eliminating the rationality bias

It turns out that the magnitude of the rationality bias directly depends on monetary policy. In particular, depending on monetary policy, the rationality bias can completely be eliminated. This is stated in Proposition 3.1.

Proposition 3.1. *If the log-deviations from steady state of the nominal interest rate and the inflation rate are equal in every period, then $c_{t-1}^R = c_{t-1}^B$ and $b_{t-1}^B = b_{t-1}$ imply that $c_{t+s}^B = c_{t+s}^R$ and $b_{t+s-1}^B = b_{t+s-1}$, $s > 0$. That is, shocks to the economy do not lead to a rationality bias for any parameterization of the model.*

Proof. When the nominal interest rate equals inflation in every period, we can write the Euler equation, (3.6), of rational agents as

$$c_t^R = E_t c_{t+1}^R - \frac{1}{\sigma} (\pi_t - E_t \pi_{t+1} - v_t + E_t v_{t+1}) = E_t c_{t+s}^R - \frac{1}{\sigma} (\pi_t - E_t \pi_{t+s} - v_t + E_t v_{t+s}), s > 0, \quad (3.29)$$

where the second equality follows from iterating this equation forward until period s .

Taking the limit $s \rightarrow \infty$, and using that $\lim_{s \rightarrow \infty} E_t c_{t+s}^R = \lim_{s \rightarrow \infty} E_t \pi_{t+s} = \lim_{s \rightarrow \infty} E_t v_{t+s} = 0$, it follows that

$$c_t^R = -\frac{1}{\sigma} (\pi_t - v_t). \quad (3.30)$$

For boundedly rational consumption, we use (3.12). When $\pi_t = i_t$, and assuming $c_{t-1}^R = c_{t-1}^B$ and $b_{t-1}^B = b_{t-1}$, we can write this as

$$c_t^B = \frac{1}{1 - \bar{g}} (y_{t-1} - \hat{g}_{t-1}) - \frac{1}{\sigma} [\pi_t - \pi_{t-1} - v_t + v_{t-1}]. \quad (3.31)$$

Since it is assumed that $c_{t-1}^B = c_{t-1}^R = -\sigma^{-1}(\pi_{t-1} - v_{t-1})$, it follows from market clearing that $y_{t-1} = -\sigma^{-1}(1 - \bar{g})(\pi_{t-1} - v_{t-1}) + \hat{g}_{t-1}$. Therefore, (3.31) reduces to

$$c_t^B = -\frac{1}{\sigma}(\pi_t - v_t). \quad (3.32)$$

This implies that when $c_{t-1}^R = c_{t-1}^B$ and $b_{t-1}^B = b_{t-1}$, it holds that $c_t^R = c_t^B$, so that boundedly rational and rational agents make identical decisions and also $b_t^B = b_t$. That is, no differences in consumption and wealth in period $t - 1$ imply no consumption and wealth differences in period t . This then holds for all periods $s \geq t$, which proves the proposition. \square

To understand the intuition of Proposition 3.1, recall the different components of the rationality bias depicted in Equation (3.28). As derived in the proof of the proposition, the rational consumption expectation is $E_t c_{t+1}^R = -\sigma^{-1}E_t(\pi_{t+1} - v_{t+1})$. Also, using that initially there were no consumption and wealth heterogeneity, market clearing and Equation (3.30) imply $y_{t-1} = -\sigma^{-1}(1 - \bar{g})(\pi_{t-1} - v_{t-1}) + \hat{g}_{t-1}$. Equation (3.28) then reduces to

$$\begin{aligned} \Delta_i c_t^i = c_t^B - c_t^R = & \underbrace{-\frac{1}{\sigma}(\pi_{t-1} - v_{t-1} - E_t(\pi_{t+1} - v_{t+1}))}_{\text{consumption exp. bias}} \\ & + \underbrace{\frac{1}{\sigma}(\pi_{t-1} - E_t \pi_{t+1})}_{\text{neg. real int. bias}} - \underbrace{\frac{1}{\sigma}(v_{t-1} - E_t v_{t+1})}_{\text{preference shock exp. bias}} = 0. \end{aligned} \quad (3.33)$$

In Equation (3.33) it can be seen that the consumption expectation bias, the real interest rate bias and the preference shock expectation bias *exactly off-set* each other, even though they are different from zero *individually*. Hence, the above proposition implies that boundedly rational agents act *as if* they were rational, so that the difference $\Delta_i c_t^i$ is zero, but that they do so for different reasons than rational agents. Moreover, this holds for any fraction of rational agents, α .

Note, that the assumption that the model is initialized with no consumption and wealth heterogeneity is not as restrictive as it may seem. In particular, when the well-known Taylor principle is satisfied, our model is stable and determinate. Heterogeneity would, in the absence of shocks, then always disappear over time as the model converges to steady state. Moreover, Proposition 3.1 states that, when its conditions are satisfied, shocks do not lead to a rationality bias. It is, therefore, intuitive that, also in the presence of shocks, consumption and wealth heterogeneity would disappear over time when $i_t = \pi_t$ in every period.

We find that this is indeed the case. When we initialize simulations of the model with wealth and consumption heterogeneity and shock the model in every period, differences in consumption and bond holdings between the two household types quickly go to zero.

3.4.2 Minimizing the rationality bias and welfare

A possible implementation of Proposition 3.1 is setting $\phi_\pi = 1$ and $\phi_y = 0$ in the Taylor rule.¹⁹ This would, however, imply that the model is infinitely close to being indeterminate. Indeterminacy issues can be solved by setting ϕ_π slightly bigger than 1 or ϕ_y slightly bigger than 0, which would still approximately implement Proposition 3.1.

Note, though, that such a weak response of monetary policy implies increased aggregate volatility in the economy. It may, therefore, not be desirable for a central bank to implement a policy where the bias is completely eliminated.

In Appendix 3.C we show that, under the assumption that individual preferences can be aggregated by simply summing up their utilities, the second-order approximated aggregate utility loss is given by

$$L_t \simeq a_1 \text{var}(y_t) + a_2 \text{var}(\pi_t) + a_3 \text{var}(c_t^B - c_t^R) \quad (3.34)$$

with

$$a_1 = \left[\gamma + \frac{\sigma}{(1 - \bar{g})} \right] \quad (3.35)$$

$$a_2 = \frac{\eta}{\delta} \quad (3.36)$$

$$a_3 = \alpha(1 - \alpha) \left[(1 - \bar{g})\sigma + \frac{\sigma^2}{\gamma} \right]. \quad (3.37)$$

Under the parameterization of Section 3.3.3, the weight a_3 on the variance of the rationality bias is approximately three percent of the weight on inflation volatility.²⁰ This reflects a very strong distaste of agents in this economy for inflation volatility and a rather minor distaste for consumption heterogeneity (and output volatility). Thus, even though the third term in Equation (3.34) disappears when Proposition

¹⁹It can be shown that when the linear model in this chapter is hit by only one type of shock, it is also possible to implement a Taylor rule that results in an interest rate equal to the inflation rate with different combinations of ϕ_π and ϕ_y . Results are available on request. This, however, no longer works in medium scale DSGE models and in the presence of multiple shocks.

²⁰the weight on output volatility, a_1 , is approximately five percent of that on inflation volatility under the benchmark calibration.

3.1 is implemented, the corresponding increase in inflation volatility will lead to a considerably higher welfare loss. Hence, based on (3.34), agents in this economy will accept a certain magnitude of the rationality bias and thus a certain heterogeneity in consumption (and wealth) if this comes with a policy that achieves price stability. Note that the strong dislike for inflation volatility is model-inherent and robust with respect to the parameterization.²¹ Our findings regarding welfare are very much in line with the results in Debortoli and Galí (2017), who consider a two-agent New Keynesian (TANK) model where the weight in the central bank's micro-founded loss function on a certain heterogeneity index is equivalently low.

Of course, if a social planner believes that reducing the rationality bias and, therefore, consumption and wealth heterogeneity is more important than is reflected in the utility functions of the households populating the economy, a different conclusion may be reached. While we consider the analysis of welfare under alternative loss functions to be interesting, it is beyond the scope of this chapter.

3.5 Conclusion

We build a macroeconomic model with different degrees of rationality. While rational households are indeed fully rational, boundedly rational households are assumed to be considerably less sophisticated. Boundedly rational households follow the less cognitive demanding Euler-equation learning and use a simple naive forecast heuristic to form expectations.

Because both household types make different decisions, substantial consumption and wealth heterogeneity arises when the economy is hit by shocks. We show that the rationality bias of boundedly rational households is the driver of consumption and wealth heterogeneity and can be decomposed into three different components: the consumption expectation bias, the real interest rate bias and the preference shock expectation bias. Further, we show that certain monetary policy can eliminate the rationality bias independent of the shocks that hit the economy. In this case the components of the rationality bias exactly offset each other while they are non-zero individually.

However, it might not be desirable from a welfare perspective to completely eliminate the rationality bias and, therefore, consumption and wealth heterogeneity, as this comes along with high inflation volatility.

²¹See e.g. Galí (2015), Woodford (2003), Di Bartolomeo et al. (2016) and Debortoli and Galí (2017) for loss functions in similar models.

Appendix 3.A Microfoundations

3.A.1 Consumption of rational households

Recall the log-linearized budget constraint of rational households

$$\hat{b}_t^R = h_t^R + w_t + \beta^{-1}\hat{b}_{t-1}^R + \bar{b}\beta^{-1}(i_{t-1} - \pi_t) + \eta^{-1}d_t - (1 - \bar{g})c_t^R - \frac{\bar{T}}{\bar{Y}}\tau_t. \quad (3.38)$$

Iterating until infinity and rearranging gives

$$\begin{aligned} E_t^R \sum_{s=t}^{\infty} \beta^{s-t} c_s^R &= (\beta(1 - \bar{g}))^{-1} \hat{b}_{t-1}^R \\ &+ E_t^R \sum_{s=t}^{\infty} \beta^{s-t} (1 - \bar{g})^{-1} \left[h_s^R + w_s + \bar{b}\beta^{-1}(i_{s-1} - \pi_s) + \eta^{-1}d_s - \frac{\bar{T}}{\bar{Y}}\tau_s \right]. \end{aligned} \quad (3.39)$$

Using $h_t^R = \gamma^{-1}(w_t - \sigma c_t^R)$ and solving for c_s^R yields

$$\begin{aligned} \left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{\gamma} \right] E_t^R \sum_{s=t}^{\infty} \beta^{s-t} c_s^R &= (\beta(1 - \bar{g}))^{-1} \hat{b}_{t-1}^R \\ &+ \sum_{s=t}^{\infty} \beta^{s-t} (1 - \bar{g})^{-1} \left[\frac{\gamma + 1}{\gamma} w_s + \bar{b}\beta^{-1}(i_{s-1} - \pi_s) + \eta^{-1}d_s - \frac{\bar{T}}{\bar{Y}}\tau_s \right]. \end{aligned} \quad (3.40)$$

For now we are focusing on the left-hand side of (3.40) which can be written as

$$\left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{\gamma} \right] c_t^R + \left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{\gamma} \right] E_t^R \sum_{s=t+1}^{\infty} \beta^{s-t} c_s^R. \quad (3.41)$$

Then we can iterate the Euler equation until period s

$$E_t^R c_s^R = c_t^R + \frac{1}{\sigma} E_t^R \sum_{k=t}^{s-1} (i_k - \pi_{k+1} - v_k + v_{k+1}). \quad (3.42)$$

Using this, we can write

$$\begin{aligned} \left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{\gamma} \right] E_t^R \sum_{s=t}^{\infty} \beta^{s-t} c_s^R &= \left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{(1 - \beta)\gamma} \right] c_t^R \\ &+ \left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{\gamma\sigma} \right] E_t^R \sum_{s=t+1}^{\infty} \sum_{k=t}^{s-1} \beta^{s-t} (i_k - \pi_{k+1} - v_k + v_{k+1}). \end{aligned} \quad (3.43)$$

Simplifying the double sum yields

$$\begin{aligned} & \left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{\gamma} \right] E_t^R \sum_{s=t}^{\infty} \beta^{s-t} c_s^R = \left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{(1 - \beta)\gamma} \right] c_t^R \\ & + \left[\frac{\gamma + \sigma(1 - \bar{g})^{-1}}{\gamma\sigma} \right] \frac{\beta}{1 - \beta} E_t^R \sum_{s=t}^{\infty} \beta^{s-t} (i_s - \pi_{s+1} - v_s + v_{s+1}). \end{aligned} \quad (3.44)$$

Inserting into (3.40) and solving for c_t^R gives

$$\begin{aligned} c_t^R &= \zeta \hat{b}_{t-1}^R + \zeta \bar{b}(i_{t-1} - \pi_t) \\ &+ \zeta \beta E_t^R \sum_{s=t}^{\infty} \beta^{s-t} \left[\frac{\gamma + 1}{\gamma} w_s + \eta^{-1} d_s - \frac{\bar{T}}{\bar{Y}} \tau_s \right] \\ &- \frac{(1 - \zeta \bar{b} \sigma) \beta}{\sigma} E_t^R \sum_{s=t}^{\infty} \beta^{s-t} (i_s - \pi_{s+1}) + \frac{\beta}{\sigma} E_t^R \sum_{s=t}^{\infty} \beta^{s-t} (v_s - v_{s+1}) \end{aligned} \quad (3.45)$$

with $\zeta = (1 - \beta)\gamma / ((\gamma(1 - \bar{g}) + \sigma)\beta)$. Taking $w_t = \gamma h_t^i + \sigma c_t^i$, integrating over agent types and using $y_t = h_t$ and $c_t = y_t / (1 - \bar{g}) - \hat{g}_t / (1 - \bar{g})$ gives $w_t = (\gamma + \sigma / (1 - \bar{g})) y_t - \sigma \hat{g}_t / (1 - \bar{g})$. Further, total dividends are defined as $D_t = Y_t - (1 - s)W_t H_t$. Log-Linearizing and using $y_t = h_t$ gives

$$d_t = \frac{\bar{Y}}{\bar{D}} [(1 - (1 - s)\bar{W})y_t - (1 - s)\bar{W}w_t]. \quad (3.46)$$

Using $w_t = (\gamma + \sigma / (1 - \bar{g})) y_t - \sigma \hat{g}_t / (1 - \bar{g})$, $(1 - s)\bar{W} = 1 - \eta^{-1}$ and $\bar{D} = \bar{Y} \eta^{-1}$ yields

$$d_t = [1 - (\eta - 1)(\gamma + \frac{\sigma}{1 - \bar{g}})] y_t + (\eta - 1) \sigma \frac{\hat{g}_t}{1 - \bar{g}}. \quad (3.47)$$

Using the expression for d_t and w_t and inserting into (3.45) gives

$$\begin{aligned} c_t^R &= \zeta \hat{b}_{t-1}^R + \zeta \bar{b}(i_{t-1} - \pi_t) + \zeta \beta E_t^R \sum_{s=t}^{\infty} \beta^{s-t} [\Gamma_y y_s - \Gamma_g \hat{g}_s - \Gamma_\tau \tau_s] \\ &- \frac{(1 - \zeta \bar{b} \sigma) \beta}{\sigma} E_t^R \sum_{s=t}^{\infty} \beta^{s-t} (i_s - \pi_{s+1}) + \frac{\beta}{\sigma} E_t^R \sum_{s=t}^{\infty} \beta^{s-t} (v_s - v_{s+1}) \end{aligned} \quad (3.48)$$

with

$$\Gamma_y = \frac{1 + \gamma}{\eta} + 1 + \frac{(\gamma + \eta)\sigma}{\gamma\eta(1 - \bar{g})} \quad (3.49)$$

$$\Gamma_g = \frac{(\gamma + \eta)\sigma}{(1 - \bar{g})\gamma\eta} \quad (3.50)$$

$$\Gamma_\tau = \frac{\bar{T}}{\bar{Y}}. \quad (3.51)$$

Further, Equation (3.38) can be rewritten as

$$\hat{b}_t^R = \beta^{-1}\hat{b}_{t-1}^R - \Gamma_c c_t^R + \Gamma_y y_t + \bar{b}\beta^{-1}(i_{t-1} - \pi_t) - \Gamma_g \hat{g}_t - \Gamma_\tau \hat{\tau}_t \quad (3.52)$$

with $\Gamma_c = (1 - \bar{g}) + (\sigma/\gamma)$. Equation (3.52) can equivalently be obtained for boundedly rational households.

3.A.2 Consumption of boundedly rational households

The consumption Euler equation for boundedly rational households reads

$$c_t^B = E_t^B c_{t+1}^B - \frac{1}{\sigma}(i_t - E_t^B \pi_{t+1} - v_t + v_{t+1}). \quad (3.53)$$

Iterated forward until period N which is far enough in the future for the boundedly rational households to consider it as “long-run” (which can, in principle, also be infinity) and using the law of iterated expectations (LIE) gives

$$c_t^B = E_t^B c_N^B - \frac{1}{\sigma} \sum_{k=0}^N E_t^B [i_{t+k} - \pi_{t+k+1} - v_{t+k} + v_{t+k+1}]. \quad (3.54)$$

Also, agents are assumed to know about market clearing $y_t = (1 - \bar{g})c_t + \hat{g}_t$ which can be written one period ahead from the point of view of boundedly rational households as

$$E_t^B y_{t+1} = (1 - \bar{g})E_t^B (\alpha c_{t+1}^R + (1 - \alpha)c_{t+1}^B) + E_t^B \hat{g}_{t+1}. \quad (3.55)$$

Further, we assume that boundedly rational households believe that all other households will satisfy the consumption Euler equation in the same way they themselves do, including the iteration of the consumption Euler equation until period N . Inserting

both forward-iterated individual consumption Euler equations yields

$$E_t^B y_{t+1} = (1 - \bar{g})E_t^B \left(\alpha(E_{t+1}^R c_N^R - \frac{1}{\sigma} \sum_{k=1}^N E_{t+1}^R [i_{t+k} - \pi_{t+k+1} - v_{t+k} + v_{t+k+1}]) + (1 - \alpha)(E_{t+1}^B c_N^B - \frac{1}{\sigma} \sum_{k=1}^N E_{t+1}^B [i_{t+k} - \pi_{t+k+1} - v_{t+k} + v_{t+k+1}]) \right) + E_t^B \hat{g}_{t+1}. \quad (3.56)$$

This can be written using Branch and McGough (2009)'s assumptions on higher-order beliefs as

$$E_t^B \frac{y_{t+1}}{1 - \bar{g}} - E_t^B c_N - E_t^B \frac{\hat{g}_{t+1}}{1 - \bar{g}} = -\frac{1}{\sigma} E_t^B \sum_{k=1}^N [i_{t+k} - \pi_{t+k+1} - v_{t+k} + v_{t+k+1}]. \quad (3.57)$$

Further, writing (3.54) as

$$c_t^B - E_t^B c_N^B + \frac{1}{\sigma} [i_t - E_t^B \pi_{t+1} - v_t + v_{t+1}] = -\frac{1}{\sigma} E_t^B \sum_{k=1}^N [i_{t+k} - \pi_{t+k+1} - v_{t+k} + v_{t+k+1}] \quad (3.58)$$

and equating (3.57) and (3.58) gives

$$c_t^B = \frac{1}{1 - \bar{g}} E_t^B y_{t+1} + E_t^B (c_N^B - c_N) - \frac{1}{\sigma} [i_t - E_t^B \pi_{t+1} - v_t + E_t^B v_{t+1}] - \frac{1}{1 - \bar{g}} E_t^B \hat{g}_{t+1}. \quad (3.59)$$

Assuming that, when boundedly rational households have more wealth than the average, they will be able to consume more than the average in the long-run, i.e. $E_t^B (c_N^B - c_N) = \psi(\hat{b}_{t-1}^B - \hat{b}_{t-1})$, yields

$$c_t^B = \frac{1}{1 - \bar{g}} E_t^B y_{t+1} + \psi(\hat{b}_{t-1}^B - \hat{b}_{t-1}) - \frac{1}{\sigma} [i_t - E_t^B \pi_{t+1} - v_t + v_{t+1}] - \frac{1}{1 - \bar{g}} E_t^B \hat{g}_{t+1}. \quad (3.60)$$

3.A.3 Aggregation of individual consumption decisions

Plugging individual consumption decisions (3.11) and (3.48) into market clearing $y_t = (1 - \bar{g})(\alpha c_t^R + (1 - \alpha)c_t^B) + \hat{g}_t$ and using $E_t^R = E_t$ and $E_t^B x_{t+1} = x_{t-1}$ gives

$$\begin{aligned}
 y_t &= (1 - \bar{g}) \left[\alpha \zeta \hat{b}_{t-1}^R + \alpha \zeta \bar{b}(i_{t-1} - \pi_t) + \zeta \alpha \beta E_t \sum_{s=t}^{\infty} \beta^{s-t} [\Gamma_y y_s - \Gamma_g \hat{g}_s - \Gamma_\tau \tau_s] \right. \\
 &\quad - \frac{(1 - \zeta \bar{b} \sigma) \alpha \beta}{\sigma} E_t \sum_{s=t}^{\infty} \beta^{s-t} (i_s - \pi_{s+1}) + \frac{\alpha \beta}{\sigma} E_t \sum_{s=t}^{\infty} \beta^{s-t} (v_s - v_{s+1}) \\
 &\quad + (1 - \alpha)(1 - \bar{g})^{-1} y_{t-1} + (1 - \alpha) \psi (\hat{b}_{t-1}^B - \hat{b}_{t-1}) \\
 &\quad \left. - \frac{(1 - \alpha)}{\sigma} [i_t - \pi_{t-1} - v_t + v_{t-1}] - (1 - \alpha)(1 - \bar{g})^{-1} \hat{g}_{t-1} \right] + \hat{g}_t. \tag{3.61}
 \end{aligned}$$

Writing (3.61) one period ahead (from the point of view of rational households) and multiplying by β and subtracting the resulting equation from (3.61) gives after collecting terms and solving for y_t

$$\begin{aligned}
 y_t &= \Phi_1 E_t y_{t+1} + \Phi_2 y_{t-1} + \Phi_3 E_t \pi_{t+1} - \Phi_4 \pi_t + \Phi_5 \pi_{t-1} + \Phi_6 i_{t-1} - \Phi_7 i_t \\
 &\quad + \Phi_8 E_t i_{t+1} - \Phi_9 \hat{b}_{t-1}^R + \Phi_{10} b_t^R - \Phi_{11} \hat{g}_{t-1} + \Phi_{12} \hat{g}_t - \Phi_{13} E_t \hat{g}_{t+1} \\
 &\quad - \Phi_{14} \tau_t + \Phi_{15} \hat{b}_{t-1} - \Phi_{16} \hat{b}_t + \Phi_{17} v_t - \Phi_{18} E_t v_{t+1} - \Phi_5 v_{t-1} \tag{3.62}
 \end{aligned}$$

with

$$\Phi_1 = \frac{\beta}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \tag{3.63}$$

$$\Phi_2 = \frac{(1 - \alpha)}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \tag{3.64}$$

$$\Phi_3 = \frac{(1 - \bar{g})\alpha\beta}{(1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta)\sigma} \tag{3.65}$$

$$\Phi_4 = \frac{(1 - \bar{g})((1 - \alpha)\beta + \bar{b}\alpha\zeta\sigma)}{(1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta)\sigma} \tag{3.66}$$

$$\Phi_5 = \frac{(1 - \bar{g})(1 - \alpha)}{(1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta)\sigma} \tag{3.67}$$

$$\Phi_6 = \frac{(1 - \bar{g})\alpha\zeta\bar{b}}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \tag{3.68}$$

$$\Phi_7 = \frac{(1 - \bar{g})(1 - \alpha(1 - \beta))}{(1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta)\sigma} \tag{3.69}$$

$$\Phi_8 = \frac{(1 - \bar{g})(1 - \alpha)\beta}{(1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta)\sigma} \tag{3.70}$$

$$\Phi_9 = \frac{(1 - \bar{g})\alpha(\psi - \zeta)}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \quad (3.71)$$

$$\Phi_{10} = \frac{(1 - \bar{g})\alpha\beta(\psi - \zeta)}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \quad (3.72)$$

$$\Phi_{11} = \frac{(1 - \alpha)}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \quad (3.73)$$

$$\Phi_{12} = \frac{1 - (1 - \bar{g})\alpha\beta\Gamma_g\zeta + (1 - \alpha)\beta}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \quad (3.74)$$

$$\Phi_{13} = \frac{\beta}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \quad (3.75)$$

$$\Phi_{14} = \frac{(1 - \bar{g})\alpha\beta\zeta\Gamma_\tau}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \quad (3.76)$$

$$\Phi_{15} = \frac{(1 - \bar{g})\alpha\psi}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \quad (3.77)$$

$$\Phi_{16} = \frac{(1 - \bar{g})\alpha\psi\beta}{1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta} \quad (3.78)$$

$$\Phi_{17} = \frac{(1 - \bar{g})((1 - \alpha)(1 + \beta) + \alpha\beta)}{(1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta)\sigma} \quad (3.79)$$

$$\Phi_{18} = \frac{(1 - \bar{g})\beta}{(1 - (1 - \bar{g})\alpha\beta\Gamma_y\zeta + (1 - \alpha)\beta)\sigma}. \quad (3.80)$$

3.A.4 Phillips-Curve

The production function is linear and given by

$$Y_t(j) = H_t(j) \quad (3.81)$$

where $H_t(j)$ is labor supply for good j .

Firms maximize their profits by choosing their individual price $P_t(j)^*$ according to

$$\max_{P_t(j)^*} E_t \sum_{s=0}^{\infty} \omega^s Q_{t|t+s} [P_t(j)^* Y_{t+s}(j) - Y_{t+s}(j) P_{t+s} MC_{t+s}] \quad (3.82)$$

s.t.

$$Y_{t+s}(j) = \left(\frac{P_t(j)^*}{P_{t+s}} \right)^{-\eta} Y_{t+s} \quad (3.83)$$

where

$$Q_{t|t+s}^s = \beta^s \frac{\Upsilon_{t+s}}{\Upsilon_t} \left(\frac{C_{t+s}}{C_t} \right)^{-\sigma} \frac{P_t}{P_{t+s}} \quad (3.84)$$

is the stochastic discount factor with C_t being aggregate consumption.

The optimization procedure gives firm j 's first-order condition

$$\tilde{P}_t(j)E_t \sum_{s=0}^{\infty} \omega^s \beta^s \Upsilon_{t+s} C_{t+s}^{-\sigma} \left(\frac{P_{t+s}}{P_t}\right)^{\eta-1} Y_{t+s} = \frac{\eta}{\eta-1} E_t \sum_{s=0}^{\infty} \omega^s \beta^s \Upsilon_{t+s} C_{t+s}^{-\sigma} \left(\frac{P_{t+s}}{P_t}\right)^{\eta} Y_{t+s} MC_{t+s}. \quad (3.85)$$

where $\tilde{P}_t(j) = P_t(j)^*/P_t$.

Log-linearizing firms optimality condition (3.85) around the zero-inflation steady state gives

$$\tilde{p}_t = (1 - \omega\beta)mc_t + \omega\beta E_t[\tilde{p}_{t+1} + \pi_{t+1}] \quad (3.86)$$

where marginal costs are given by

$$mc_t = w_t = \left(\gamma + \frac{\sigma}{1-\bar{g}}\right)y_t - \sigma \frac{\hat{g}_t}{1-\bar{g}} + \mu_t \quad (3.87)$$

and where μ_t is a time varying exogenous wage markup (as in Galí, 2015), which we refer to as a cost-push shock. Plugging in marginal costs in (3.86) and combining with the log-linearized evolution of aggregate prices $\pi_t = ((1-\omega)/\omega)\tilde{p}_t$ yields the Phillips-Curve given by (3.15).

3.A.5 Steady state

In this section, we derive the steady state around which the model is log-linearized, where gross inflation equals 1.

Evaluating (3.85) at the zero inflation steady state gives

$$\bar{MC} = \frac{\eta-1}{\eta}. \quad (3.88)$$

From the first order conditions of the households it follows that in this steady state we must have

$$1 + \bar{i} = \frac{1}{\beta}. \quad (3.89)$$

Furthermore, it follows from Equation 3.81 that

$$\bar{H} = \bar{Y}. \quad (3.90)$$

Next, we solve the steady state aggregate resource constraint for consumption, and write

$$\bar{C} = \bar{Y}(1 - \bar{g}). \quad (3.91)$$

Plugging in these steady state labor and consumption levels in the non-linear optimality

condition for labor supply, $W_t = (H_t^i)^\gamma (C_t^i)^\sigma$, and using $s = 1/\eta$ gives

$$\bar{W} = \bar{Y}^\gamma (\bar{Y}(1 - \bar{g}))^\sigma = \bar{Y}^{\gamma+\sigma} (1 - \bar{g})^\sigma = \frac{1}{(1 - s)} \bar{M}C = 1 \quad (3.92)$$

where the last equality follows from $\bar{M}C = (1 - s)\bar{W}$ and (3.88). We can thus write

$$\bar{Y} = \left(\frac{1}{1 - \bar{g}} \right)^{\frac{\sigma}{\gamma+\sigma}}. \quad (3.93)$$

Then we turn to the government budget constraint. In steady state, (3.3) reduces to

$$\bar{b} = \frac{\bar{g} - \frac{\bar{T}}{\bar{Y}} + \eta^{-1}}{1 - \beta^{-1}} \quad (3.94)$$

where we use $s = 1/\eta$ and substitute for the interest rate using (3.89).

Appendix 3.B Impulse responses for different fractions of rational agents

In this appendix, we study the robustness of the impulse responses and the effects of the rationality bias to the calibration of the fraction of rational households (α). All other parameters are parameterized as in Section 3.3.3. The first two columns in Figure 3.5 show the impulse responses of individual consumption and bond holdings to a cost-push shock, as well as the difference between individual consumption (the rationality bias) and bond holdings of boundedly rational households and rational households. The right column of Figure 3.5 shows the impulse responses of aggregate inflation, output and the nominal interest rate. The black-dotted curves correspond to the benchmark case of $\alpha = 0.5$. The yellow-solid curves depict the case of more boundedly rational households ($\alpha = 0.25$), while the blue-dashed curves plot the case of more rational households ($\alpha = 0.75$).

In the figure, it can be seen that for lower α (yellow-solid) inflation initially responds stronger to the cost-push shock, and then falls more steeply during periods 2-5. This causes boundedly rational households to overestimate future inflation more during these periods, resulting in a more negative real interest rate bias. Moreover, aggregate output is flatter than for higher values of α so that the consumption expectation bias becomes less negative. Both these effects lead to less negative consumption of boundedly rational households in periods 2-5.

However, the general findings of Section 3.3.3 remain valid for different values of

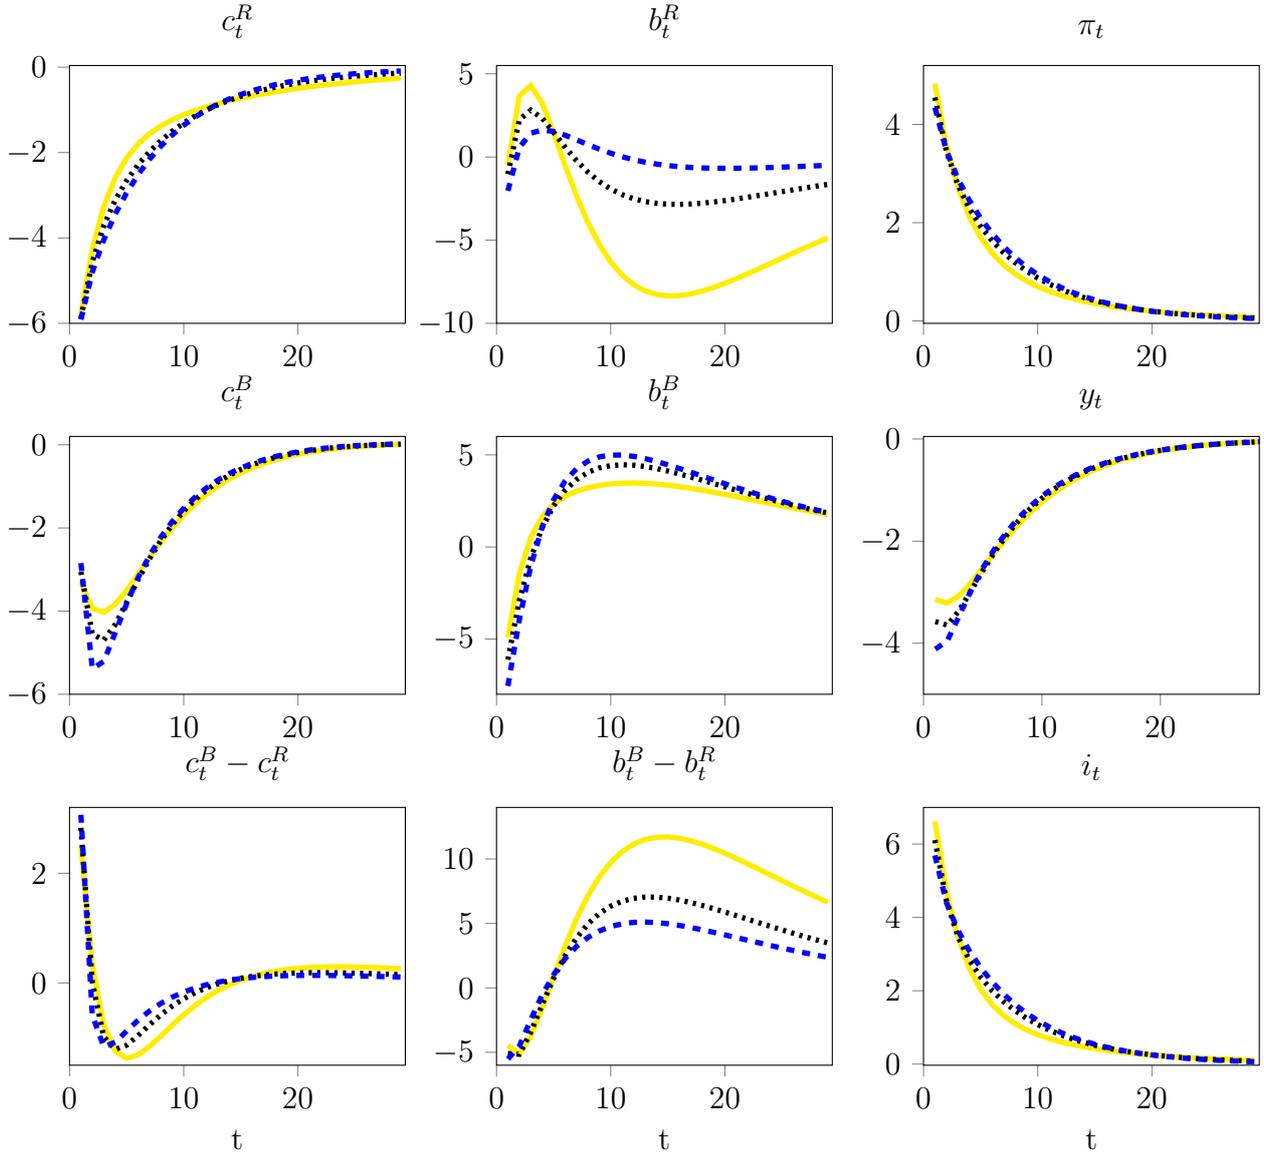


Figure 3.5: Impulse responses in case of a 2.5% autocorrelated cost-push shock for $\alpha = 0.25$ (yellow-solid), $\alpha = 0.5$ (black-dotted) and $\alpha = 0.75$ (blue-dashed).

α . In particular, after a cost-push shock, the real interest rate bias and consumption expectations bias lead boundedly rational households to first consume more than rational households. After one or two periods this reverses and the rationality bias becomes positive. Boundedly rational households, therefore, first have less wealth than rational households, but accumulate more bond holdings in the medium run. Note, however, that the level of wealth that both households accumulate depends on the fractions of agents in the economy (i.e. on α), through bonds market clearing.

In a similar manner, Figure 3.6 depicts impulse responses to a preference shock for three different values of α . Here it can be seen that the general results of Section

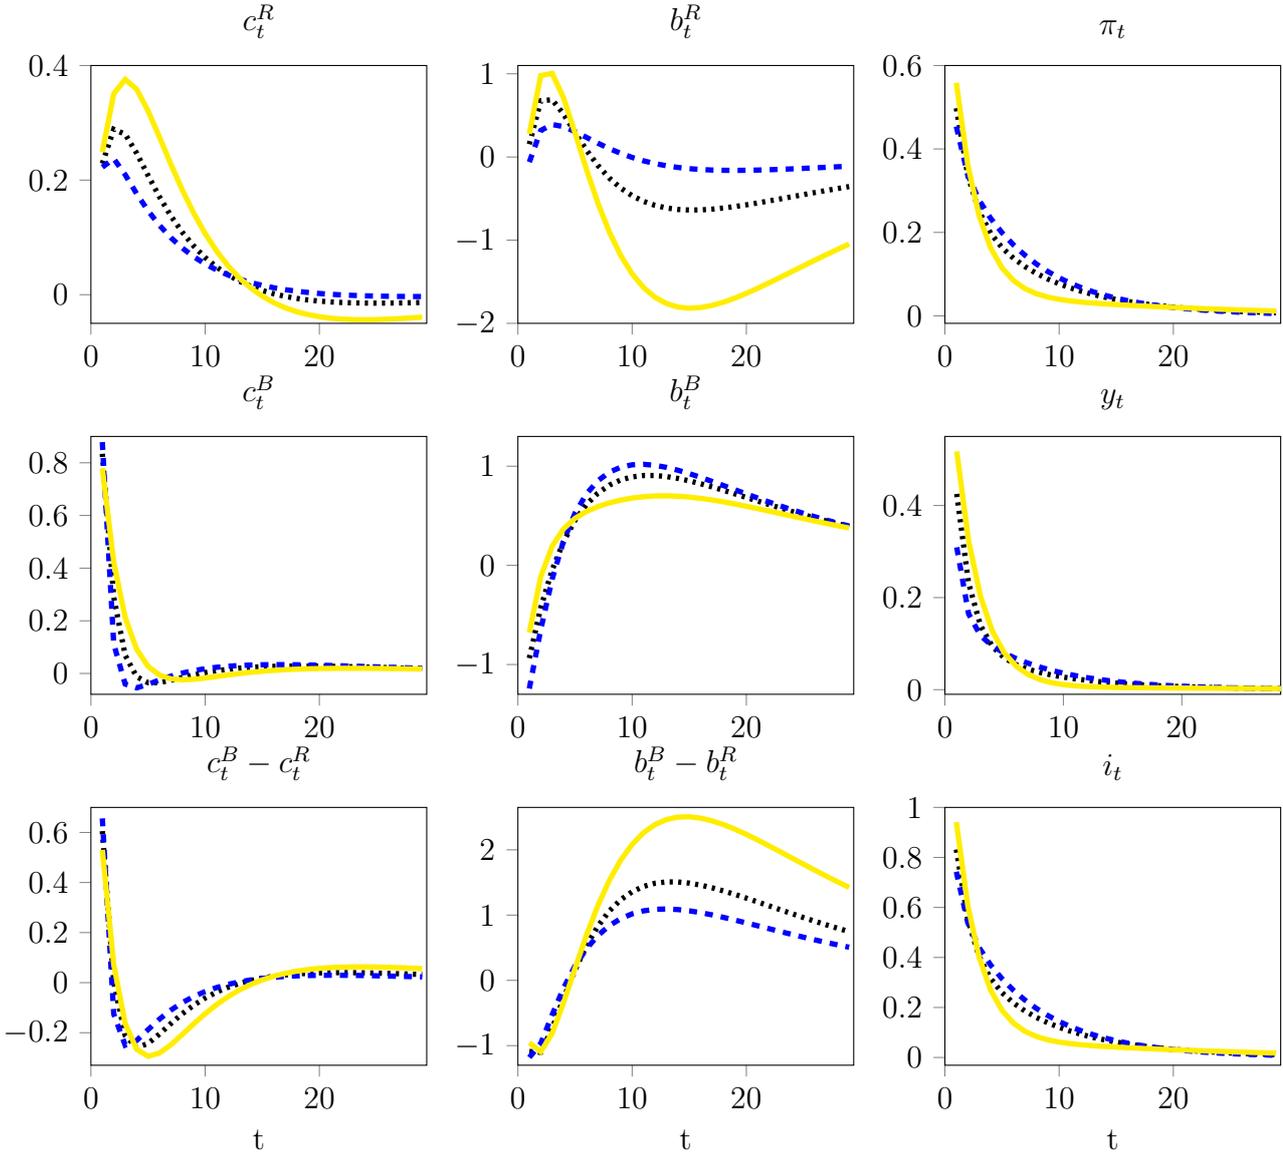


Figure 3.6: Impulse responses in case of a 2.5% autocorrelated preference shock for $\alpha = 0.25$ (yellow-solid), $\alpha = 0.5$ (black-dotted) and $\alpha = 0.75$ (blue-dashed).

3.3.3 are robust to the calibration of α , also under the preference shock. Moreover, the rationality bias and differences in bond holdings of boundedly rational and rational households are affected by the value of α in the same way as in Figure 3.5.

Appendix 3.C Derivation of second-order approximated utilitarian welfare loss

Household utility is given by

$$\Upsilon_t U_t^i = U(C_t^i, \Upsilon_t) - V(H_t^i, \Upsilon_t) = \frac{\Upsilon_t (C_t^i)^{1-\sigma}}{1-\sigma} - \frac{\Upsilon_t (H_t^i)^{1+\gamma}}{1+\gamma} \quad (3.95)$$

Therefore, $U_C = (\bar{C})^{-\sigma}$, $\bar{C}U_{CC} = -\sigma U_C$, $U_{C\Upsilon} = U_C$ and $V_H = (\bar{Y})^\eta$, $\bar{Y}V_{HH} = \eta V_H$, $V_{H\Upsilon} = V_H$. Further, we will use $U_C = \bar{W}V_H$ (labor-leisure condition) where $\bar{W} = 1$ due to a subsidy to marginal costs to firms so that the steady state is efficient. Thus, $U_C = V_H$. As before, lower case letter indicated a log-deviation from steady state. We will frequently use the fact that

$$\frac{X_t - \bar{X}}{\bar{X}} = x_t + \frac{1}{2}x_t^2 \quad (3.96)$$

holds up to second-order for any variable X . Hence, the second-order approximated goods market clearing condition in terms of log-deviations from steady state becomes $y_t + (1/2)y_t^2 = (1-\bar{g})(c_t + (1/2)c_t^2)$ where we already assume that government spending will not deviate from its steady state value. Taking squares, the latter can be written as $y_t^2 = (1-\bar{g})^2 c_t^2 + \mathcal{O}(3)$. Plugging back in yields $y_t = (1-\bar{g})c_t + (1/2)\bar{g}(1-\bar{g})c_t^2 + \mathcal{O}(3)$ which can be rewritten as

$$(1-\bar{g})c_t = y_t - \frac{1}{2}\frac{\bar{g}}{(1-\bar{g})}y_t^2 + \mathcal{O}(3). \quad (3.97)$$

Now, we turn to approximating the utility function (3.95) using the results above. The consumption utility part in (3.95) can be approximated as (see Woodford, 2003 or Di Bartolomeo et al., 2016)

$$U(\cdot) = \bar{C}U_C \left(c_t - \frac{\sigma}{2}\text{var}_i(c_t^i) + \frac{1-\sigma}{2}(c_t)^2 + v_t c_t \right) + t.i.p. + \mathcal{O}(3) \quad (3.98)$$

where we used that $c_t = E_i c_t^i + (1/2)\text{var}_i(c_t^i)$ holds up to a second order and where $t.i.p.$ contains steady state values and shock terms that are not interacting with policy dependent variables.

Using market clearing, Equation (3.98) can be written as

$$U(\cdot) = \bar{Y}U_C \left(y_t - \frac{1}{2}\frac{\bar{g}}{(1-\bar{g})}y_t^2 - (1-\bar{g})\frac{\sigma}{2}\text{var}_i(c_t^i) + \frac{1-\sigma}{2(1-\bar{g})}(y_t)^2 + v_t y_t \right) + t.i.p. + \mathcal{O}(3). \quad (3.99)$$

which can be rewritten as

$$U(\cdot) = \bar{Y}U_C\left(y_t - (1 - \bar{g})\frac{\sigma}{2}\text{var}_i(c_t^i) + \frac{1 - \frac{\sigma}{(1-\bar{g})}}{2}y_t^2 + v_t y_t\right) + t.i.p. + \mathcal{O}(3). \quad (3.100)$$

Now, we turn to the approximation of the labor-utility term which is analogously given by

$$V(\cdot) = \bar{Y}V_H\left(h_t + \frac{\gamma}{2}\text{var}_i(h_t^i) + \frac{1 + \gamma}{2}h_t^2 + v_t h_t\right) + t.i.p. + \mathcal{O}(3). \quad (3.101)$$

Now, our goal is to write aggregate labor in terms of aggregate output and price dispersion. To that end, we follow Galí (2015) where we use $H_t = \int_0^1 H_t(j) dj = \int_0^1 Y_t(j) dj$ which is equal to

$$H_t = Y_t \int_0^1 \left(\frac{P_t(j)}{P_t}\right)^{-\eta} dj. \quad (3.102)$$

Making a second-order approximation and using that $(\eta/2)\text{var}_i(p_t(i)) = \log\left(\int_0^1 (P_t(j)/P_t)^{-\eta} dj\right)$ (Lemma 1 in Appendix A in Chapter 4 of Galí, 2015) yields

$$h_t = y_t + \frac{\eta}{2}\text{var}_j(p_t(j)). \quad (3.103)$$

Inserting gives

$$v(\cdot) = \bar{Y}V_H\left(y_t + \frac{\eta}{2}\text{var}_j(p_t(j)) + \frac{\gamma}{2}\text{var}_i(h_t^i) + \frac{1 + \gamma}{2}(y_t)^2 + v_t y_t\right) + t.i.p. + \mathcal{O}(3). \quad (3.104)$$

Now, using the linear labor-leisure condition (which we can do as we insert it into a variance which is already second-order) we can write $(\gamma/2)\text{var}_i(h_t^i) = (\sigma^2/(2\gamma))\text{var}_i(c_t^i)$.²² Hence, (3.104) becomes (also using $V_h = U_c$ from above)

$$V(\cdot) = \bar{Y}U_H\left(y_t + \frac{\eta}{2}\text{var}_j(p_t(j)) + \frac{\sigma^2}{2\gamma}\text{var}_i(c_t^i) + \frac{1 + \gamma}{2}(y_t)^2 + v_t y_t\right) + t.i.p. + \mathcal{O}(3). \quad (3.105)$$

Inserting (3.100) and (3.105) into (3.95) gives

$$\begin{aligned} & \bar{Y}U_C\left(y_t - (1 - \bar{g})\frac{\sigma}{2}\text{var}_i(c_t^i) + \frac{1 - \frac{\sigma}{(1-\bar{g})}}{2}(y_t)^2 + v_t y_t\right) \\ & - \left[\bar{Y}U_C\left(y_t + \frac{\eta}{2}\text{var}_j(p_t(j)) + \frac{\sigma^2}{2\gamma}\text{var}_i(c_t^i) + \frac{1 + \gamma}{2}(y_t)^2 + v_t y_t\right)\right] + t.i.p. + \mathcal{O}(3) \end{aligned} \quad (3.106)$$

²² $\text{var}_i(\frac{1}{\gamma}w_t - \frac{\sigma}{\gamma}c_t^i) = \text{var}_i(\frac{1}{\gamma}w_t) + \text{var}_i(\frac{\sigma}{\gamma}c_t^i) - 2\text{Cov}_i(\frac{1}{\gamma}w_t, \frac{\sigma}{\gamma}c_t^i) = \frac{\sigma^2}{\gamma^2}\text{var}_i(c_t^i)$.

and combining

$$- \frac{\bar{Y}U_C}{2} \left[a_1(y_t)^2 + \tilde{a}_2 \text{var}_j(p_t(j)) + \tilde{a}_3 \text{var}_i(c_t^i) \right] + t.i.p. + \mathcal{O}(3) \quad (3.107)$$

where $a_1 = [\gamma + (\sigma/(1-\bar{g}))]$, $\tilde{a}_2 = \eta$ and $\tilde{a}_3 = [(1-\bar{g})\sigma + (\sigma^2/\gamma)]$. Further, $\text{var}_j(p_t(j)) = \delta^{-1}\pi_t^2$ with $\delta = ((1-\omega)(1-\omega\beta))/\omega$ (see Woodford, 2003, Chapter 6) and $\text{var}_i(c_t^i) = \alpha(1-\alpha)(c_t^B - c_t^R)^2$. Thus, the aggregated intertemporal utility can be written as

$$\mathcal{W} = -\frac{\bar{Y}U_C}{2} \sum_{t=0}^{\infty} \beta^t [a_1(y_t)^2 + a_2\pi_t^2 + a_3(c_t^B - c_t^R)^2 + t.i.p. + \mathcal{O}(3)]. \quad (3.108)$$

where $a_2 = \frac{\eta}{\delta}$ and $a_3 = \alpha(1-\alpha)[(1-\bar{g})\sigma + (\sigma^2/\gamma)]$. Taking unconditional expectations yields

$$\mathcal{W} \equiv -\frac{\bar{Y}U_C}{2(1-\beta)} [a_1 \text{var}(y_t) + a_2 \text{var}(\pi_t) + a_3 \text{var}(c_t^B - c_t^R) + t.i.p. + \mathcal{O}(3)]. \quad (3.109)$$

Hence, the approximated utility loss is given by

$$L \simeq a_1 \text{var}(y_t) + a_2 \text{var}(\pi_t) + a_3 \text{var}(c_t^B - c_t^R). \quad (3.110)$$

Chapter 4

A New Keynesian model with heterogeneous household expectations and labor market frictions

4.1 Introduction

The standard New Keynesian model abstracts from realistic features along several dimensions. Movements in employment, for instance, are the result of voluntary adjustments in hours worked. Therefore, the basic model is not useful to analyze labor market variables albeit their obvious importance for the business cycle. By now, there has emerged a vast literature on incorporating labor market frictions into the New Keynesian model, greatly enhancing the models' empirical performance.²³

Another crucial abstraction is the assumption of homogeneous rational expectations. This assumption is at odds with empirical literature on (heterogeneous) expectations and, therefore, denies to importance of possibly non-rational expectations for individual decisions and macroeconomic outcomes. Evidence for heterogeneous expectations based on survey expectations can be found in Branch (2004) and Cornea-Madeira et al. (2019). Also, laboratory experiments show that expectation formation is heterogeneous (Kryvtsov and Petersen, 2019; Pfajfar and Žakelj, 2018; Assenza et al., 2019). One general finding of this literature is that a certain fraction of forecasters is quite sophisticated while the rest follows simple backward-looking heuristics. Several

²³Some notable examples are Gertler et al. (2008), Gertler and Trigari (2009), Trigari (2009) and Christiano et al. (2016).

strands of the literature have developed models to accommodate this shortcoming as well. Some models incorporating heterogeneous expectations can be found in Branch and McGough (2009), De Grauwe (2011), Kurz et al. (2013), Gasteiger (2014), Hommes and Lustenhouwer (2019) and in Chapter 3.

However, these models assume frictionless labor markets where wages are not explicitly negotiated but are dictated by the market, and where involuntary unemployment is not possible. Therefore, in these models heterogeneous expectations only have *indirect* equilibrium effects on wages and employment. Allowing for labor market frictions results in first-order effects of expectations on labor market variables.

The danger of increasing a models' realism is that it may become far more complex, potentially reducing the benefit of the additional explanatory power. Therefore, I build a macroeconomic model with labor market frictions that is highly stylized but, as soon as augmented with an equally stylized model of heterogeneous household expectations, generates qualitatively realistic responses of relevant labor market variables. It, further, predicts distributional consequences through the channel of heterogeneous expectations along the business cycle. The degree of cross-sectional dispersion, thereby, depends on the “fluidity” of the labor market. To the best of my knowledge, this research project is the first to combine heterogeneous expectations and labor market frictions in a macro-setting.

The central idea for modeling the labor market is that hiring of workers is costly, as in Blanchard and Galí (2010). Therefore, as in the real world, firms have to spend resources to find workers that produce their goods. The tighter the labor market, i.e. the larger the hiring activity relative to the size of the unemployment pool, the larger the cost. Thus, this approach shares the same basic feature with usual search-and-matching models á la Mortensen and Pissarides (1994): the cost of hiring is increasing in labor-market-tightness. However, this model abstracts from an explicit matching technology on the grounds of simplicity.²⁴

For every job that continues in the future, firms save hiring costs. The fact that hiring is costly will result in unemployment of workers that are willing to work. As a result, existing jobs inhere surpluses for both workers and employers. The whole “pie” of surpluses is then split between workers and firms through wage negotiations. Therefore, the negotiated wages do not merely equal the marginal revenue of workers as in the frictionless case but more realistically also depend on expected future surpluses.

Firms may hire workers that differ by expectation type. As expectations enter the wage negotiations, wages itself will be a function of type-specific expectations. The

²⁴For a more detailed discussion, please refer to Blanchard and Galí (2010).

reason is that the expected surpluses of a job continuation differs between expectation-types which is naturally absent under homogeneous rational expectations. Consequently, unemployment and job-finding rates will differ between workers of different expectation types, too. This is the case even if workers are equally productive and, therefore, have the same marginal revenue, an assumption that I shall maintain throughout this chapter.

I derive the model under the possibility that there may be an arbitrary amount of different expectation types. I will, however, illustrate the model using two types of expectations only, i.e. rational and naive expectations. The latter model of expectations assumes that when boundedly rational households (or its individual workers) form expectations at the beginning of the period, they rely on the last observed value to forecast future values. There is already some evidence that expectations may well be described by such a model (Branch, 2004; Cornea-Madeira et al., 2019; Kryvtsov and Petersen, 2019). In particular, Kryvtsov and Petersen (2019) find that expectations of participants in a learning-to-forecast experiment can well be described by a fixed fraction of conventional rational and naive agents, both making up half of the population. Thus, aggregate expectations exhibit a certain degree of persistence affecting the model's dynamics.

At the aggregate level, the persistence in expectations leads to qualitatively realistic dynamics of relevant labor market variables, especially hump-shaped impulse responses of hiring, unemployment rates and labor market tightness. This feature is not generated by the homogeneous rational expectations version of this model. Admittedly, there are more complex rational expectations models that generate observed empirical patterns through different mechanisms (see for instance Christiano et al. 2016). Therefore, heterogeneous expectations with a certain degree of persistence should be seen as an alternative, or additional, explanation for some labor market stylized facts.

Moreover, a model of heterogeneous expectations can give insights in the dispersion effects of shocks on type-specific wages, unemployment and job-finding rates along the business cycle. In a “fluid” labor market with a low level of unemployment and high job-finding rates, hiring is relatively more costly. The reason is that the unemployment pool from which firms can hire is small. Therefore, firms have a higher value from continuing the employment relationship as this saves future hiring costs, *ceteris paribus*. The latter information is exploited by the workers in the wage negotiations. Thus, expectations about future hiring costs are more important the higher the level of hiring costs are. This then leads to more wage dispersion between different expectation types. In contrast, a more “sclerotic” labor market with a higher level of unemployment implies lower hiring costs for firms. Hence, expectations about future hiring costs matter less,

leading to less wage dispersion.

As all workers are remunerated equally for the same marginal revenue, wage dispersion in this framework is relatively low, and approaches zero the more “sclerotic” the labor market is. However, as the model of heterogeneous expectations implies persistent dynamics, already small wage differences may lead to significant and persistent differences in type-specific unemployment and job-finding rates. Further, a more “sclerotic” labor market facilitates larger volatility and persistence of aggregate activity.

The chapter is organized as follows. Section 4.2 derives the model, putting an emphasis on the labor market frictions, the assumptions on heterogeneous expectations and the wage bargaining. Subsequently, Section 4.3 illustrates the model using a simple model of heterogeneous expectations for “fluid” and “sclerotic” calibrations of the labor market. Section 4.4 concludes.

4.2 Model

In this section, I introduce a New Keynesian model where households may differ in the way they form expectations comparable to Kurz et al. (2013) and Branch and McGough (2009). However, I simplify the analysis by assuming that firms will be rational throughout. Moreover, the model will be subject to labor market frictions by assuming that hiring is costly, as in Blanchard and Galí (2010).

4.2.1 Households

Households of different type τ have different expectation operators.²⁵ The fraction of a specific type- τ household is a_τ where $\sum_\tau a_\tau = 1$. Each household consumes a fraction of the aggregate consumption basket, i.e. $c_t = \sum_\tau a_\tau c_t^\tau = \int_0^1 [(c_{i,t})^{\frac{\epsilon-1}{\epsilon}} di]^{\frac{\epsilon}{\epsilon-1}}$, where c_t^τ is type-specific consumption, $c_{i,t}$ is the differentiated good that firm i produces and ϵ is the elasticity of substitution. The corresponding optimal (aggregate) expenditure on the differentiated good $c_{i,t}$ is $c_{i,t} = \left(\frac{p_{i,t}}{p_t}\right)^{-\epsilon} c_t$ where $p_{i,t}$ is the retail price of $c_{i,t}$ and c_t is aggregate demand. Further, each representative household of type τ inelastically

²⁵Following Merz (1995), the demand side of this economy is composed of consumers (workers) that are uniformly distributed over $[0,1]$. Consumers (workers) of type τ are grouped in large family households. Every member of a specific household engages in perfect consumption insurance within the type- τ family household. This assumption allows a group of type τ consumers (workers) to be modeled as a representative agent of type τ .

supplies a fraction a_τ of the total unit measure of labor to labor market.²⁶

The real budget constraint of a type- τ household reads

$$c_t^\tau + b_t^\tau \leq n_t^\tau w_t^\tau + \frac{i_{t-1}}{\pi_t} b_{t-1}^\tau + d_t. \quad (4.1)$$

b_t^τ is type-specific real bond holdings, n_t^τ is the fraction of the household that is employed²⁷ and w_t^τ is the wage that employed workers of the household receive. Further, i_t is the nominal interest rate set by the central bank, $\pi_t = \frac{p_t}{p_{t-1}}$ is inflation and d_t are real dividends that households obtain as shareholders of the firms. Further, bond market clearing $0 = \sum_\tau a_\tau b_t^\tau$ has to hold, implying that aggregate bonds are in zero net supply.

Households are assumed to maximize

$$E_0^\tau \sum_{t=0}^{\infty} \beta^t \left[\frac{(c_t^\tau)^{1-\sigma}}{1-\sigma} - \frac{\tilde{\psi}}{2} (b_t^\tau)^2 \right], \quad (4.2)$$

where E_0^τ is the type-specific expectations operator. Moreover, β is the subjective discount factor and σ the coefficient of relative risk aversion that are both equal across types. Households also derive quadratic dis-utility from holding bonds which is scaled by the parameter $\tilde{\psi} \geq 0$.

I borrow the latter assumption from Kurz et al. (2013) as it elegantly enforces a transversality condition on *individual* borrowing. It is convenient to do so as I will apply the Euler approach, discussed in detail in Honkapohja et al. (2011), where agents use Euler equations as usual but with subjective expectations to make decisions. I will chose the value of the parameter $\tilde{\psi}$ to be small, but sufficiently high to rule out unstable borrowing and equilibria with unrealistic high volatility.

All expectations operators satisfy assumptions A1-A5 of Branch and McGough (2009). That is, the expectations operators fix observables while the expectation in steady state of some variable is the steady state value of that variable itself. The latter assumption implies that households of different expectation types are observationally equivalent in steady state. Further, usual linearity assumptions apply and the law of

²⁶This assumption is a prominent one in the literature on labor market frictions, see for instance Christiano et al. (2016), Krause and Lubik (2007), Faia (2008) and Gertler and Trigari (2009). Yet, the literature does not seem to have settled on whether to allow for dis-utility out of labor or not. For instance, in Blanchard and Galí (2010) and Thomas (2008) the wage does also depends on the marginal rate of substitution (labor-leisure condition) due to dis-utility out of working. However, as discussed in Blanchard and Galí (2010), this assumption aggravates the Shimer Puzzle (Shimer, 2005) by causing wages to be even more pro-cyclical. Also, assuming dis-utility out of labor is not a necessary condition for the existence of the household's value function for labor (4.8).

²⁷Thus, the fraction of type- τ workers of the total unit measure of labor that is employed is $a_\tau n_t^\tau$.

iterated expectations needs to hold at the individual level. If $E_t^\tau \neq E_t$, i.e. E_t^τ is boundedly rational, I assume that these type of agents believe that all other agents form the same beliefs as they themselves do.

Denoting the Lagrange-multiplier on (4.1) as λ_t^τ and maximizing (4.2) subject to (4.1) gives the consumption Euler equation

$$(c_t^\tau)^{-\sigma} + \tilde{\psi} b_t^\tau = \beta E_t^\tau \left[(c_{t+1}^\tau)^{-\sigma} \frac{i_t}{\pi_{t+1}} \right] \quad (4.3)$$

of type- τ household and $\lambda_t^\tau = (c_t^\tau)^{-\sigma}$.

Employment of a type- τ household evolves according to

$$n_t^\tau = (1 - \delta)n_{t-1}^\tau + h_t^\tau \quad (4.4)$$

where δ denotes the exogenous separation rate, and where h_t^τ is the amount of hiring of type- τ worker in that period.²⁸ Further, the beginning-of-period unemployment pool of type- τ worker available for hiring is given by

$$u_t^\tau = 1 - (1 - \delta)n_{t-1}^\tau. \quad (4.5)$$

Moreover, the probability of an individual type- τ worker of finding a job in period t is defined as the ratio between hiring and the corresponding unemployment pool,

$$x_t^\tau = \frac{h_t^\tau}{u_t^\tau}. \quad (4.6)$$

Definition (4.6) can, from the point of view of a firm, be interpreted as the type-specific labor-market-tightness index. Using (4.6) and (4.5), equation (4.4) can be rewritten as

$$n_t^\tau = (1 - \delta)n_{t-1}^\tau + x_t^\tau(1 - (1 - \delta)n_{t-1}^\tau). \quad (4.7)$$

Equation (4.7) states that the fraction of a type- τ household that is currently employed is equal to the non-separated fraction of the previous period, plus the fraction of the specific unemployment pool that finds a job in the current period.

Let γ_t^τ be the Lagrange-Multiplier on (4.7). The value of a marginal job to the type- τ household is then given by

$$\gamma_t^\tau = \lambda_t^\tau w_t^\tau + \beta(1 - \delta)E_t^\tau(1 - x_{t+1}^\tau)\gamma_{t+1}^\tau. \quad (4.8)$$

²⁸Equation (4.4) highlights an important property of the steady state equilibrium: in steady state, hiring is equal to job destruction, i.e. $h = \delta n$.

Thus, the value of a marginal job to a specific household is equal to the wage, in terms of marginal consumption utility, and the subjectively expected continuation value. Equation (4.8) will play an important role in the wage negotiations later on.

4.2.2 Production

The production side of the economy consists of a continuum of individual firms i that are uniformly distributed over $[0, 1]$. All households are shareholders of each individual firm.

The production function of each firm is given by $f_{i,t} = z_t n_{i,t}$ where z_t is productivity, which is a mean-one AR(1)-process with persistence ρ_z , and $n_{i,t}$ is labor input of firm i . Each firm hires workers of different expectation-types, i.e. $n_{i,t} = \sum_{\tau} a_{\tau} n_{i,t}^{\tau}$. Aggregate employment is, therefore, given by $n_t = \sum_{\tau} a_{\tau} n_t^{\tau} = \int_0^1 n_{i,t} di$.

As indicated earlier, hiring is assumed to be costly as in Blanchard and Galí (2010). The cost per individual hire is given by

$$g_t^{\tau} = z_t B (x_t^{\tau})^{\varphi} \quad (4.9)$$

where $\varphi \geq 0$ and where B is a positive constant.²⁹ This approach shares the same basic feature with usual search-and-matching models á la Mortensen and Pissarides (1994): the cost of hiring is increasing in labor-market-tightness. Further, hiring costs are taken as given by each individual firm. One implication of this approach is that firms fill newly created vacancies immediately, i.e. in that same quarter, by paying the hiring-costs. Therefore, the model does not rely on the typical assumption that a formed match becomes productive only in the subsequent quarter to generate observed persistence. Christiano et al. (2016) already pointed out that such a long delay between a match and start of employment is rather unrealistic.

For the purpose of exposition, I will make a distinction between the wholesale and retail branch of each firm i . The wholesale branch (WB) produces the differentiated good, hires workers and makes zero profits (perfect competition). The retail branch (RB), in turn, prices the differentiated good under monopolistic competition and quadratic price adjustment costs. The real price $\tilde{p}_{i,t}^{WB} = \frac{p_{i,t}^{WB}}{p_t}$ of the wholesale branch is thereby the marginal cost of the retail branch.

²⁹I follow Blanchard and Galí (2010) in assuming the dependence of (4.9) on productivity to avoid effects of productivity shocks on the cost of production versus the cost of hiring.

Wholesale Branch

The wholesale branch of firm i chooses $\{n_{i,t}^\tau, h_{i,t}^\tau\} \forall \tau$ taking as given $\{w_{i,t}^\tau, \tilde{p}_{i,t}^{WB}\} \forall \tau$ to maximize the current and future discounted flow of real profits³⁰

$$\Pi_{i,t}^{WB} = E_0 \sum_{t=0}^{\infty} \beta^t \frac{\lambda_t}{\lambda_0} \left[\tilde{p}_{i,t}^{WB} z_t \sum_{\tau} a_{\tau} n_{i,t}^{\tau} - \sum_{\tau} a_{\tau} w_{i,t}^{\tau} n_{i,t}^{\tau} - \sum_{\tau} a_{\tau} g_t^{\tau} h_{i,t}^{\tau} \right] \quad (4.10)$$

subject to

$$n_{i,t}^{\tau} = (1 - \delta)n_{i,t-1}^{\tau} + h_{i,t}^{\tau}. \quad (4.11)$$

Let μ_t^τ be the Lagrange-multiplier on (4.11). All firms will make identical choices. Hence, the first-order conditions, for all τ , are given by

$$n_t^\tau : \mu_t^\tau = z_t \tilde{p}_t^{WB} - w_t^\tau + (1 - \delta)\beta E_t \frac{\lambda_{t+1}}{\lambda_t} \mu_{t+1}^\tau \quad (4.12)$$

$$h_t^\tau : \mu_t^\tau = g_t^\tau \quad (4.13)$$

for all τ . Equation (4.12) shows that the value of a marginal job to a firm is the marginal revenue less the wage plus the expected continuation value. Moreover, Equation (4.13) is an arbitrage condition on hiring: firms will hire type τ workers until the marginal value of a job is equal to the hiring costs.

Combining (4.12) and (4.13), and recognizing that \tilde{p}_t^{WB} is equal to the marginal costs mc_t of the retail branch, yields

$$mc_t = \tilde{g}_t^\tau - (1 - \delta)\beta E_t \frac{\lambda_{t+1}}{\lambda_t} \frac{z_{t+1}}{z_t} \tilde{g}_{t+1}^\tau + \frac{w_t^\tau}{z_t} \quad \forall \tau \quad (4.14)$$

where $\tilde{g}_t^\tau = B(x_t^\tau)^\varphi$. Marginal costs are, thus, hiring costs of type τ net of savings in hiring costs in the future plus the wage (scaled by productivity). Note that, in equilibrium, firms will at the margin be indifferent between hiring workers of different type as marginal costs are equal across types.

Retail Branch

The retail branch of firm i prices the differentiated good $y_{i,t}$ subject to the demand curve of households. The retail branch, further, pays quadratic price adjustment costs which depend on aggregate activity (Rotemberg, 1982). Hence, taking marginal costs

³⁰Where $\lambda_t = \sum_{\tau} a_{\tau} \lambda_t^\tau$.

as given, retailer i chooses $p_{i,t}$ to maximize real profits

$$\Pi_{i,t}^{RB} = E_0 \sum_{t=0}^{\infty} \beta^t \frac{\lambda_t}{\lambda_0} \left[\frac{p_{i,t}}{p_t} y_{i,t} - mc_t y_{i,t} - \frac{\chi}{2} \left(\frac{p_{i,t}}{p_{i,t-1}} - 1 \right)^2 y_t \right] \quad (4.15)$$

subject to

$$y_{i,t} = \left(\frac{p_{i,t}}{p_t} \right)^{-\epsilon} y_t. \quad (4.16)$$

Equation (4.16) assumes that market clearing, $c_t = y_t$, holds. Recognizing that all firms make identical choices, the optimality condition for pricing is given by

$$\begin{aligned} 1 - \chi(\pi_t - 1)\pi_t + \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \left[\chi(\pi_{t+1} - 1)\pi_{t+1} \frac{y_{t+1}}{y_t} \right] \\ = (1 - mc_t)\epsilon. \end{aligned} \quad (4.17)$$

In linearized form, (4.17) gives the standard New Keynesian Phillips curve which coincides with the Phillips curve under Calvo pricing for particular values of the scaling parameter χ , see Section 4.2.5.

4.2.3 Wage Bargaining

The wage that workers receive is an outcome of bargaining process between workers and firms that splits the joint surplus of an existing employment relationship.

Let $S_t^{F,\tau} = \mu_t$ denote the firms' surplus from an existing employment relationship with a type- τ worker. Hence,

$$S_t^{F,\tau} = z_t \tilde{p}_t^{WB} - w_t^\tau + (1 - \delta)\beta E_t \frac{\lambda_{t+1}}{\lambda_t} S_{t+1}^{F,\tau}. \quad (4.18)$$

Further, let $S_t^{H,\tau} = \frac{\gamma_t^\tau}{\lambda_t^\tau}$ be the surplus, in terms of the consumption basket, for household of type τ , so

$$S_t^{H,\tau} = w_t^\tau + \beta(1 - \delta)E_t^\tau (1 - x_{t+1}^\tau) \frac{\lambda_{t+1}^\tau}{\lambda_t^\tau} S_{t+1}^{H,\tau}. \quad (4.19)$$

I assume that there exists a union for each type τ of workers that negotiates the type-specific wage with the firms on behalf of their members. However, in the special case where one type has rational expectations, the strategic interaction with the firms can equivalently be characterized by *bilateral* (generalized) Nash bargaining.

From now on, I distinct between rational (R) and boundedly rational (B) types of workers, i.e. $\tau \in \{R, B\}$. In the case of rational expectations, the wage will be given

by the maximizing argument of weighted product of the respective surpluses

$$w_t^R = \operatorname{argmax} \left(S_t^{H,R} \right)^\zeta \left(S_t^{F,R} \right)^{1-\zeta} \quad (4.20)$$

where ζ measures the bargaining power of workers. Equation (4.20) implies the Nash sharing rule $\zeta(S_t^{H,R} + S_t^{F,R}) = S_t^{H,R}$, i.e. the rational household gets a share ζ of the whole “pie” ($S_t^{H,R} + S_t^{F,R}$). The implied wage, in turn, is given by

$$w_t^R = \zeta \left[z_t \tilde{p}_t^{WB} + (1 - \delta) \beta E_t g_{t+1}^R \left(\frac{\lambda_{t+1}}{\lambda_t} - (1 - x_{t+1}^R) \frac{\lambda_{t+1}^R}{\lambda_t^R} \right) \right]. \quad (4.21)$$

Thus, the wage that a rational worker receives is a share of real revenues, $z_t \tilde{p}_t^{WB}$, and a share of expected savings in tomorrow's hiring costs, weighted by the difference in the expected marginal rates of substitution between the shareholders of the firm and the rational workers. Recognizing that the real price of the wholesale branch equals the marginal costs of retailer, $\tilde{p}_t^{WB} = mc_t$, the wage becomes

$$w_t^R = \zeta \left[z_t mc_t + (1 - \delta) \beta E_t g_{t+1}^R \left(\frac{\lambda_{t+1}}{\lambda_t} - (1 - x_{t+1}^R) \frac{\lambda_{t+1}^R}{\lambda_t^R} \right) \right]. \quad (4.22)$$

In the case of boundedly rational workers, I assume that the union of boundedly rational workers knows the value function (4.18) of firms. Further, the union also understands that a firm's surplus from an existing employment relationship next period is equal to the hiring costs that it saves. In other words, the union knows that the firm's surplus from hiring a boundedly rational worker is equal to the marginal revenue and expected next period's savings in hiring costs less the wage. Therefore, the union evaluates the following equation:

$$S_t^{F,B} = z_t \tilde{p}_t^{WB} - w_t^B + (1 - \delta) \beta E_t \frac{\lambda_{t+1}}{\lambda_t} g_{t+1}^B. \quad (4.23)$$

In doing so, the union needs to forecast the third term in equation (4.23), i.e.³¹

$$(1 - \delta) \beta E_t^B \left[E_t \frac{\lambda_{t+1}}{\lambda_t} g_{t+1}^B \right]. \quad (4.24)$$

In this case, I follow Branch and McGough (2009) and the associated literature in as-

³¹Here it is implicit that the expectation formation of the union is equal to the expectations of its members.

suming that boundedly rational agents believe that all other agents will form the same expectations they themselves do. Therefore, the union's boundedly rational forecast of the expected surplus of firms next period is $(1 - \delta)\beta E_t^B \frac{\lambda_{t+1}}{\lambda_t} g_{t+1}^B$. Further, the union knows how productive its members currently are and what the associated marginal revenue is. This assumption ensures that differences in wages between types are ultimately due to expectations about the future only. Hence, the union's *boundedly rational perception* of the firms' current surplus is

$$E_t^B S_t^{F,B} = z_t \tilde{p}_t^{WB} - w_t^B + (1 - \delta)\beta E_t^B \frac{\lambda_{t+1}}{\lambda_t} g_{t+1}^B. \quad (4.25)$$

Moreover, the union believes to engage in a Nash bargain with the firms. As a result, the *ex-ante perceived* Nash sharing rule of the union is

$$\zeta E_t^B (S_t^{H,B} + S_t^{F,B}) = S_t^{H,B}. \quad (4.26)$$

Thus, the union demands a share ζ of its own *boundedly rational perception* of the whole "pie" of surpluses. The wage on which the union insists is, therefore,

$$w_t^B = \zeta \left[z_t \tilde{p}_t^{WB} + (1 - \delta)\beta E_t^B g_{t+1}^B \left(\frac{\lambda_{t+1}}{\lambda_t} - (1 - x_{t+1}^B) \frac{\lambda_{t+1}^B}{\lambda_t^B} \right) \right] \quad (4.27)$$

where $z_t \tilde{p}_t^{WB}$ is equal to $z_t m c_t$.³² As the union will quit negotiations otherwise ("take it or leave it"), the firms will accept as long as their share of the whole "pie" of surpluses is positive

$$E_t (S_t^{H,B} + S_t^{F,B}) - \zeta E_t^B (S_t^{H,B} + S_t^{F,B}) = S_t^{F,B} > 0. \quad (4.28)$$

As the arbitrage condition (4.13) states, the firms' surplus must equal the hiring cost in equilibrium. Hence, for condition (4.28) to be satisfied, hiring (and productivity) must be positive. Given that hiring is positive in steady state, condition (4.28) will be satisfied for small enough productivity shocks. I shall assume the latter to be true throughout the analysis. Under this assumption, the ex-ante perceived sharing rule (4.26) is equal to the ex-post realized one. Also, requiring shocks to be small rules out scenarios where wage demands of the union of boundedly rational households are unrealistically extreme relative to the prevailing economic conditions.

Moreover, Shimer (2005) famously pointed out that real wages are too pro-cyclical

³²Equation (4.27) can alternatively be seen as a wage contract, that is written after the union forms its expectations in the beginning of a period, and where that contract is contingent on the marginal revenue that will realize later that period. In other words, the wage implied by (4.27) conditions on the "job performance" in any given period.

to be consistent with business cycle data. Therefore, I allow for real wage rigidity by assuming a weighted sum of the Nash-bargained wage and a wage norm (Hall, 2005). As in Krause and Lubik (2007) and Faia (2008), I assume the wage norm to be steady state of the Nash-bargained wage w . Hence, the individual wage becomes

$$w_t^\tau = \gamma \zeta \left[z_t m c_t + (1 - \delta) \beta E_t^\tau g_{t+1}^\tau \left(\frac{\lambda_{t+1}}{\lambda_t} - (1 - x_{t+1}^\tau) \frac{\lambda_{t+1}^\tau}{\lambda_t^\tau} \right) \right] + (1 - \gamma) w \quad \forall \tau \quad (4.29)$$

where $\gamma \in [0, 1]$ is the weight attached to the Nash wage.

4.2.4 Monetary Policy and Resource Constraint

I assume monetary policy to aim at stabilizing inflation according to

$$\ln \frac{\dot{i}_t}{i} = \phi_\pi \ln \frac{\pi_t}{\pi}. \quad (4.30)$$

where $\phi_\pi > 1$. The aggregate resource constraint is given by

$$\sum_\tau a_\tau c_t^\tau = z_t \sum_\tau a_\tau n_t^\tau - \sum_\tau a_\tau g_t^\tau h_t^\tau - \frac{\chi}{2} (\pi_t - 1)^2 y_t. \quad (4.31)$$

4.2.5 The Linearized Model

To proceed with an analysis of the dynamic behavior of the model, I linearize it around the zero-inflation steady state. Variables with a hat present log-deviations from steady state, except for the unemployment rate which is presented as an absolute deviation from steady state.

Linearizing equation (4.17) gives the standard Phillips-Curve

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \frac{\epsilon - 1}{\chi} \hat{m} c_t. \quad (4.32)$$

depending on the price-adjustment-cost parameter χ . Using $\hat{g}_t^\tau = \varphi \hat{x}_t^\tau$ and $\hat{\lambda}_t = -\sigma \hat{c}_t$, marginal costs can be obtained from equation (4.14),

$$\begin{aligned} \hat{m} c_t &= \frac{mw}{z} (\hat{w}_t^\tau - \hat{z}_t) + m\tilde{g}\varphi (\hat{x}_t^\tau - \beta(1 - \delta) E_t \hat{x}_{t+1}^\tau) \\ &\quad - m\tilde{g}(1 - \delta)\beta [(\sigma \hat{c}_t - \hat{z}_t) - E_t(\sigma \hat{c}_{t+1} - \hat{z}_{t+1})]. \end{aligned} \quad (4.33)$$

Further, the individual labor market tightness/job-finding rate follows from linearizing

(4.6),

$$\delta \hat{x}_t^\tau = \hat{n}_t^\tau - (1 - \delta)(1 - x)\hat{n}_{t-1}^\tau. \quad (4.34)$$

Using $\hat{g}_t^\tau = \hat{z}_t + \varphi \hat{x}_t^\tau$, $\hat{\lambda}_t = -\sigma \hat{c}_t$ and $\hat{\lambda}_t^\tau = -\sigma \hat{c}_t^\tau$, the individual wage (4.29) in log-linear form is given by

$$\begin{aligned} \hat{w}_t^\tau = \frac{\gamma \zeta}{w} \left[\frac{z}{m} (\hat{z}_t + \hat{m} c_t) + \Gamma E_t^\tau [\hat{z}_{t+1} + (1 + \varphi) \hat{x}_{t+1}^\tau] \right. \\ \left. + (1 - \delta) \beta g \sigma [(\hat{c}_t - E_t^\tau \hat{c}_{t+1}) - (1 - x)(\hat{c}_t^\tau - E_t^\tau \hat{c}_{t+1}^\tau)] \right] \end{aligned} \quad (4.35)$$

where $\Gamma = \frac{w}{\zeta} - \frac{z}{m} > 0$. From equation (4.4) I obtain an expression for individual hiring

$$\hat{h}_t^\tau = \delta^{-1} \hat{n}_t^\tau - \frac{1 - \delta}{\delta} \hat{n}_{t-1}^\tau \quad (4.36)$$

while individual end-of-period unemployment rate, $u_t^\tau = 1 - n_t^\tau$, in terms of deviations from steady state is given by

$$\hat{u}_t^\tau = -n \hat{n}_t^\tau. \quad (4.37)$$

Moreover, the linearized individual consumption Euler equation reads

$$\hat{c}_t^\tau = E_t^\tau \hat{c}_{t+1}^\tau - \frac{1}{\sigma} [\hat{i}_t - E_t^\tau \hat{i}_{t+1}] + \psi b_t^\tau \quad (4.38)$$

where $\psi = c\tilde{\psi}$. From the individual budget constraint (4.1), I get an expression for the evolution of individual bonds

$$b_t^\tau = wn(\hat{w}_t^\tau + \hat{n}_t^\tau) + d\hat{d}_t + \beta^{-1} b_{t-1}^\tau - c\hat{c}_t^\tau \quad (4.39)$$

where dividends are given by

$$\begin{aligned} d\hat{d}_t = zn(1 - \delta\tilde{g})\hat{z}_t - wn \sum_{\tau} a_{\tau} \hat{w}_t^\tau - \varphi gn\delta \sum_{\tau} a_{\tau} \hat{x}_t^\tau \\ + (z - w - g)n \sum_{\tau} a_{\tau} \hat{n}_t^\tau + gn(1 - \delta) \sum_{\tau} a_{\tau} \hat{n}_{t-1}^\tau. \end{aligned} \quad (4.40)$$

Finally, linearizing the aggregate resource constraint (4.31) yields

$$\begin{aligned} \sum_{\tau} a_{\tau} \hat{c}_t^\tau = \hat{z}_t + \frac{1 - \tilde{g}}{1 - \delta\tilde{g}} \sum_{\tau} a_{\tau} \hat{n}_t^\tau \\ + \frac{(1 - \delta)\tilde{g}}{1 - \delta\tilde{g}} \sum_{\tau} a_{\tau} \hat{n}_{t-1}^\tau - \frac{\varphi\tilde{g}\delta}{1 - \delta\tilde{g}} \sum_{\tau} a_{\tau} \hat{x}_t^\tau. \end{aligned} \quad (4.41)$$

Monetary policy is given by (4.30) and productivity follows $\hat{z}_t = \rho_z \hat{z}_{t-1} + \epsilon_{z,t}$.

4.3 Results

I illustrate the dynamics of the model by assuming two types of expectations: conventional rational expectations and naive expectations. In the latter case, boundedly rational agents use the last observed value of any variable v_{t+1} being forecast, i.e. $E_t^B v_{t+1} = v_{t-1}$. Hence, when boundedly rational agents form expectations in the beginning of period t , contemporaneous variables are generally not included in their information set. This assumption is typical in the adaptive learning literature, and it is convenient as it avoids the simultaneous determination of boundedly rational forecasts and endogenous variables.

Moreover, empirical evidence from several studies support the fragmentation of the population into more sophisticated expectations, where I view rational expectations as an appropriate approximation, and naive backward-looking expectations (Branch, 2004; Cornea-Madeira et al., 2019; Kryvtsov and Petersen, 2019). In particular, Kryvtsov and Petersen (2019) find that expectations of participants in a learning-to-forecast experiment can well be described by a fixed fraction of conventional rational and naive agents, both making up half of the population. I employ the latter specification to simulate the model at hand.

The baseline calibration underlying the simulation exercises in the next sections is shown in Table 4.1. I borrow parameter values related to the labor market from Blanchard and Galí (2010). These describe a relatively “fluid” labor market which is characterized by a low unemployment rate, relatively high separations and, implied by that, a high job-finding rate.³³ Later on I will contrast this by a calibration that tries to capture a more “sclerotic” labor market with higher steady-state unemployment and lower job-finding rates. Further, the fraction of rational households, α , is chosen to be one half, in accordance to Kryvtsov and Petersen (2019). As there is no hard evidence on the transversality parameter, $\tilde{\psi}$, I pick a number that rules out unstable borrowing and equilibria with unrealistic high volatility. At the same time, $\tilde{\psi}$ should be low enough to have a small impact on short run dynamics. Finally, all remaining parameter values are in line with standard literature.

³³The *quarterly* job-finding-rate under baseline calibration is 70%, consistent with US Data, which follows from the steady state relation (4.43) in Appendix 4.A.

Table 4.1: Baseline calibration

Parameter	Description	Value	Reference
β	subjective discount factor	0.99	standard
σ	coef. of rel. risk aversion	1	standard
ϵ	price elasticity (diff. goods)	10	standard
χ	price adjustment costs	40	standard
δ	separation rate	0.12	Blanchard and Galí (2010)
n	steady state employment	0.95	Blanchard and Galí (2010)
B	hiring cost parameter	0.12	Blanchard and Galí (2010)
ζ	workers bargaining power	0.5	Blanchard and Galí (2010)
φ	elasticity of hiring cost	1	Blanchard and Galí (2010)
γ	real wage rigidity	0.5	Blanchard and Galí (2010)
α	fraction of rational households	0.5	Kryvtsov and Petersen (2019)
$\tilde{\psi}$	transversality parameter	0.05	no hard evidence
ϕ_π	inflation coef. in Taylor	1.5	standard
ρ_z	persistence of tech. shock	0.9	standard
ρ_m	persistence of mon. shock	0.5	standard

4.3.1 Model Dynamics

Figures 4.1 and 4.3 display impulse responses of aggregate and individual variables in the case of heterogeneous expectations to an adverse, one-percent productivity shock, respectively.

Marginal costs increase as labor is less productive. Therefore, inflation increases while the increase is limited by contractionary monetary policy. The drop in consumption indicates a decrease in aggregate activity. As a consequence, firms hire less workers and unemployment increases.

In the context of this model, vacancy creation by firms is equal to hiring since firms immediately hire workers by paying the hiring cost. Hence, the lower-left panel of Figure 4.1 shows the model's ability to generate a Beveridge curve. Also, the model generates persistent hump-shaped responses of labor market variables as well as procyclical labor market tightness, consistent with empirical evidence. Especially, the former stylized fact seems to be less well replicated by a homogeneous rational expectations version of the model, as can be seen in Figure 4.2. In that case, labor market variables seem to be less inherently persistent but rather display the persistence of the AR(1)-shock process.

The model also implies *differences* in the responses of type-related variables, as can be seen in Figure 4.3. These differences stem from differences in household expectations entirely. As unemployment increases, individual job-finding rates drop. The job finding rate of boundedly rational workers will keep decreasing for ten quarters. As long as a variable being forecast decreases, boundedly rational workers will *systematically* expect

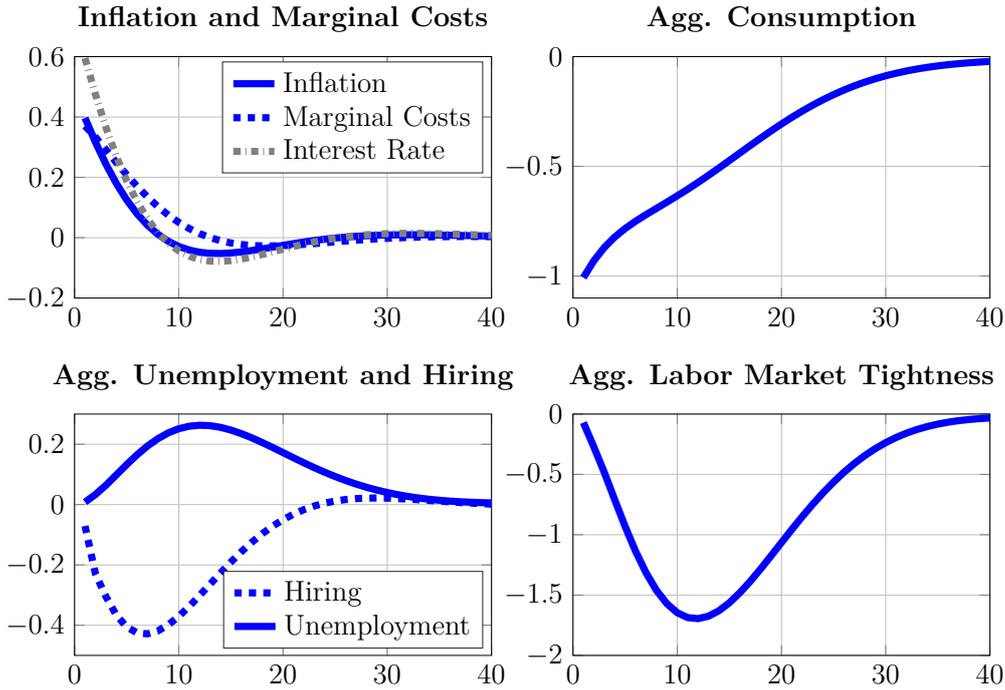


Figure 4.1: Aggregate Dynamics (Baseline) following a one percent aggregate productivity shock.

the next period's realization to be higher than it will actually turn out to be. The reason is that boundedly rational workers will always project the last observed value into the future. This leads them to be *less pessimistic* relative to rational workers with respect to their own job finding rate, and thus also to the firms' hiring costs. As the union of boundedly rational workers overestimates the firms' saving in future hiring costs, it will demand a higher wage *relative* to rational workers (although in *absolute terms* wage demands decrease for both expectation-types in response to the shock). As the firms will accept this wage by assumption, boundedly rational workers receive higher wages *relative* to rational workers for that time span. Even though firms accept to pay the relatively higher wage, they will, at the same time, hire less boundedly rational workers relative to rational workers. This results in a higher unemployment rate and lower job-finding rate relative to rational workers for some time. These effects then reverse after approximately ten periods.

A similar pattern can be observed for individual consumption. Boundedly rational households will initially be overly optimistic as they do not anticipate the effect of the shock instantaneously but only recognize that the central bank increases their policy rate. This, however, changes when their own individual consumption expectations become pessimistic while the nominal interest rate is still elevated. Then, for some time, boundedly rational households consume considerably less than rational house-

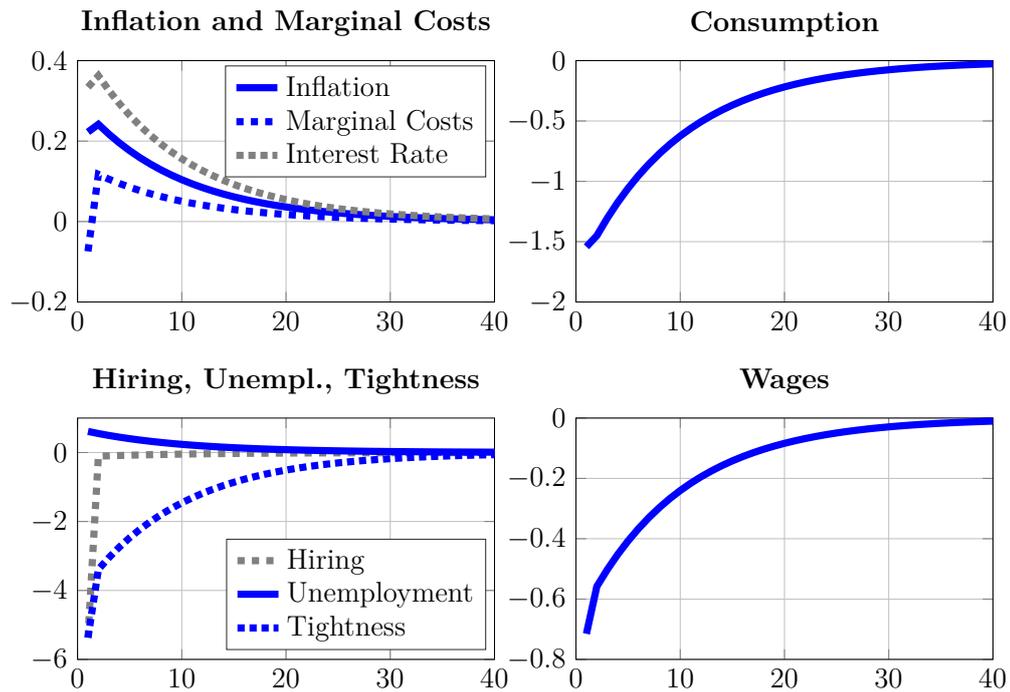


Figure 4.2: Rational Expectations Dynamics (Baseline) following a one percent aggregate productivity shock.

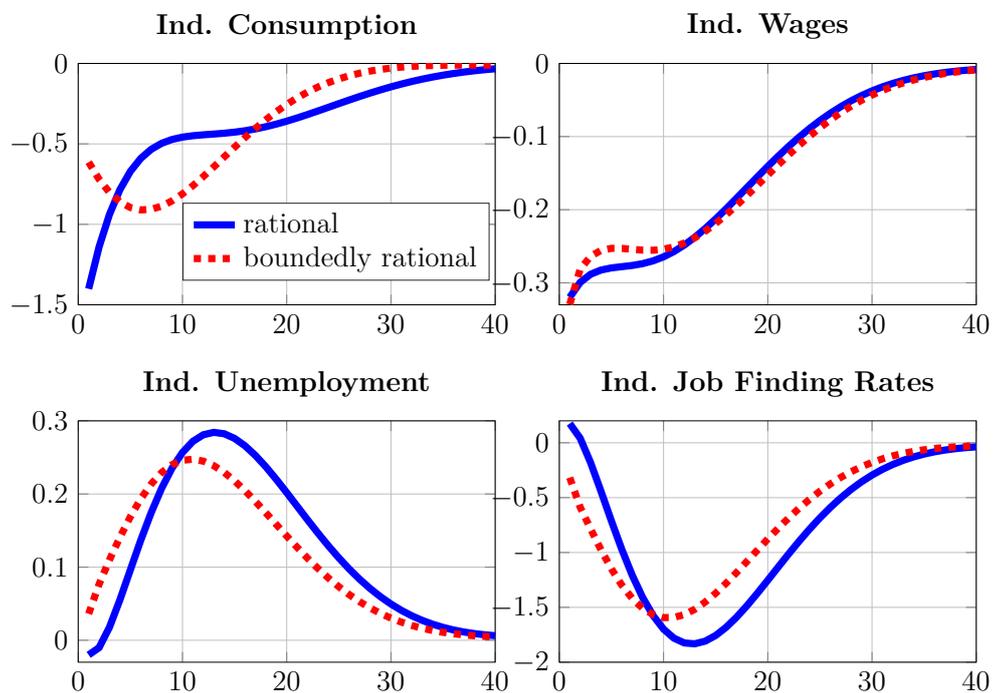


Figure 4.3: Micro-Dynamics (Baseline) following a one percent aggregate productivity shock.

holds leading them to accumulate savings.

The model with heterogeneous expectations generates empirically plausible move-

ments in aggregate variables in response to a productivity shock. Further, it predicts distributional consequences of shocks through the channel of heterogeneous expectations, even when workers of different expectation types are equally productive.

4.3.2 Comparison with a more “sclerotic” labor market

In this section, I compare the dynamic properties of the model under baseline calibration with a calibration that captures a more “sclerotic” labor market, i.e. a labor market with a higher steady-state unemployment rate and a lower job-finding rate.

As in Blanchard and Galí (2010), I assume the steady state end-of-period unemployment rate to be ten percent, i.e. $n = 0.9$, the separation rate to be four percent and, implied by that, a relatively lower quarterly job-finding rate of approximately 26%. The remaining parameters are unchanged.

A lower quarterly separation rate implies a higher expected lifetime of a job. Therefore, given some shock, reactions in the hiring decision of firms and, consequently, movements in the unemployment rate will be stronger and more persistent. This can be seen in the lower-left panel of Figure 4.4. Moreover, a relatively stronger and persistent response of hiring leads to to an approximately four times more volatile aggregate labor market tightness, see table 4.2.

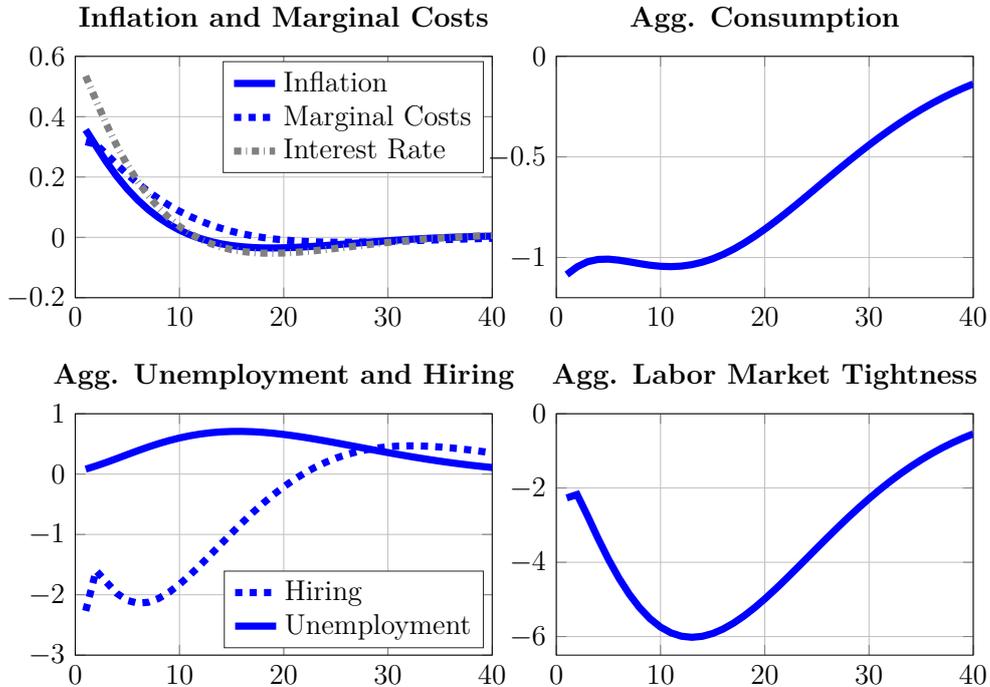


Figure 4.4: Aggregate Dynamics following a one percent aggregate productivity shock with “sclerotic” calibration.

An interesting and at first sight counter-intuitive picture emerges when looking at individual wages in the upper-right panel of Figure 4.5. First, wages do not seem to react much stronger relative to the fluid labor market. Second, the difference in wages between the two types seems to be very small. The first three rows of Table 4.2 show the theoretical cross-sectional standard deviations, $sd_{\tau}(\cdot)$, of wages, unemployment and job-finding rates. The cross-sectional standard deviation of wages indicates that wage dispersion has indeed dropped substantially relative to the fluid labor market. One might assume that, given a higher lifetime of a job, expectations of future hiring costs might matter more and, thereby, lead to more and not less wage dispersion. However, by assumption, the *level* of labor market tightness decreases substantially and, therefore, the hiring costs of the firms decrease accordingly. Albeit higher volatility, this effect counteracts the higher expected lifetime of a job, leading expectations to be less important for wage determination. Hence, individual wages are almost fully determined by the marginal revenue of workers which is the same across expectation types. Thus, as the behavior of marginal costs is not significantly different than under the more fluid labor market and the process of productivity and the shock is the same, the initial reaction in wages is not that different as well. Still, the small remaining wage dispersion leads to persistent and, in terms of the cross-sectional standard deviation, also relatively higher absolute dispersion in unemployment and job-finding rates. Relative to aggregate volatility, however, dispersion generally decreases.

While the initial response of wages is not that different from before and the sharp drop in hiring takes some time to fully transmit into unemployment, the initial drop of consumption is also similar to the fluid version of the labor market. However, consumption is depressed more persistently as unemployment is persistently elevated.

This exercise suggests that if the institutional design affecting the labor market is such that it is more “sclerotic”, an adverse productivity shock might cause less wage dispersion, a less persistent and a similar initial drop in wages relative to a more “fluid” labor market. This property of the model is largely robust with respect to varying degrees of labor market “fluidity” and also to alternative calibrations of the remaining parameters. While this effect in isolation may be desirable, it might also facilitate larger volatility and persistence in other labor market measures as unemployment and job-finding rates, translating into a more persistent decline in aggregate activity.

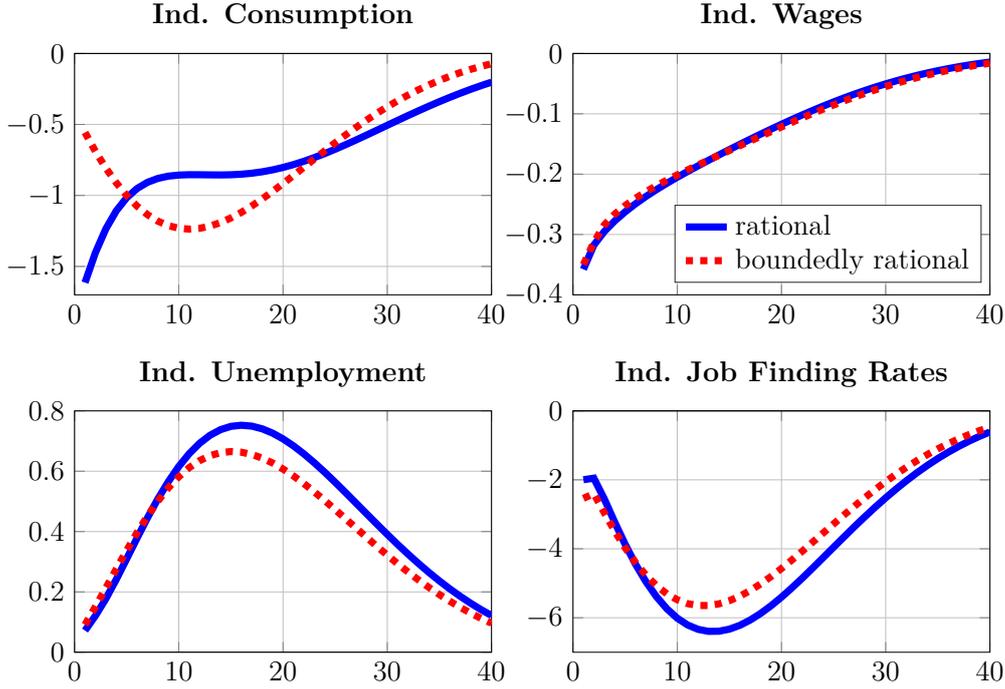


Figure 4.5: Micro-Dynamics following a one percent aggregate productivity shock with “sclerotic” calibration.

	Fluid	Sclerotic
$sd_{\tau}(w_t^{\tau})$	0.1	0.03
$sd_{\tau}(u_t^{\tau})$	0.33	0.42
$sd_{\tau}(x_t^{\tau})$	2.16	3.59
$sd(w_t)$	1.05	1.02
$sd(u_t)$	18.3	30.43
$sd(x_t)$	5.9	25.47
$corr(w_t, w_{t-5})$	0.84	0.75
$corr(u_t, u_{t-5})$	0.83	0.93
$corr(x_t, x_{t-5})$	0.81	0.91

Table 4.2: Labor Mark: Fluid vs. Sclerotic (Some theoretical Moments)

4.4 Conclusion

I develop a macroeconomic model with labor market frictions and heterogeneous expectations. The central idea of modeling the labor market is that hiring is costly which gives rise to involuntary unemployment. I, further, employ a simple model of heterogeneous expectations where a fixed fraction of households has rational expectations while the rest follows a naive heuristic. This implies persistence of aggregate expectations

which generate hump-shaped responses of relevant labor market variables. In contrast, a rational expectations version of the model does not deliver the latter stylized fact but rather follows the persistence of the AR(1) shock process.

Moreover, as expectations enter the wage negotiations, the model implies dispersion of wages, unemployment and job-finding rates following some shock. Given the model's persistent dynamics, already small differences in wages lead to a pro-longed dispersion in type-specific unemployment and job-finding rates. A more "fluid" labor market, thereby, implies that differences in expectations lead to more wage dispersion. The reason is that a generally low level of unemployment require higher future hiring costs which workers form expectations about. Thus, while a more "sclerotic" labor market entails less wage dispersion, it facilitates larger volatility and persistence in other labor market measures as unemployment and job-finding rates, translating into a more persistent response of aggregate activity.

In contrast to earlier models with heterogeneous expectations, assuming labor market frictions allows for first-order effects of heterogeneity in expectations on wages and other labor market variables, not just through equilibrium. This was achieved by an axiomatic approach of modeling the wage bargain between the union of boundedly rational workers and rational firms. This approach requires firms to accept the boundedly rational wage demand as long as the firms' share of the surplus is non-zero. This might not be too unrealistic in the case of small shocks, as assumed in this chapter, but is not easily extended too cases with larger disturbances. Future research could, therefore, develop a model of bargaining between rational players and players with boundedly expectations that needs less *ad hoc* assumptions and at the same time is simple enough to fit in a macro model.

Appendix 4.A Steady State Values

This appendix presents relevant steady state values:

$$z = 1 \tag{4.42}$$

$$x = \frac{\delta n}{1 - (1 - \delta)n} \tag{4.43}$$

$$\tilde{g} = g = zBx^\varphi \tag{4.44}$$

$$m = \frac{\epsilon}{\epsilon - 1} \tag{4.45}$$

$$c = (1 - \tilde{g}\delta)zn \tag{4.46}$$

$$w = \zeta \left[\frac{z}{m} + (1 - \delta)\beta gx \right]. \tag{4.47}$$

Chapter 5

The role of stickiness, extrapolation and past consensus forecasts in macroeconomic expectations

Co-authored by Joep Lustenhouwer

5.1 Introduction

How are macroeconomic expectations formed? There is ample evidence that forecasters do not make fully efficient forecasts and have systematic biases.³⁴ That is, real-world expectations deviate (at least to some extent) from the fully rational, model consistent benchmark. The question of how exactly they deviate from full rationality is less easily answered. Answering this question is, however, crucial for judging the validity of macroeconomic models and their policy implications. First of all, because most macroeconomic models rely on the assumption that all agents form fully rational expectations, and secondly because empirical evidence for models that deviate from full-information rational expectations are often not fully convincing, especially when one considers expectations data of individual forecasters.

In this paper, we propose a simple model of expectation formation that allows for three distinct deviations from full rationality. We then confront this model with survey expectations to test whether our model is in line with the data, and to investigate to what extent the three proposed biases are present in survey data. We mainly focus on data from the US Survey of Professional Forecasters (SPF) which has become the

³⁴See e.g. Lovell (1986); Ehrbeck and Waldmann (1996); Loungani (2001); Fildes and Stekler (2002); Coibion and Gorodnichenko (2015).

benchmark data source to test models against expectations data.³⁵ However, we also consider extensions with the Livingston survey (Carlson, 1977).

The first two biases that we allow for are stickiness in expectations and the extrapolation of recent news about the current period. A combination of stickiness (under-reaction) and extrapolation (over-reaction) can explain why different studies have found conflicting evidence as to whether forecasters under-react or over-react to news. In particular, Coibion and Gorodnichenko (2015) find clear under-reaction at the aggregate level by showing that aggregate forecast errors are positively correlated with aggregate forecast revisions. Bordalo et al. (2018a), with a similar regression at the individual level, find negative relation between forecast errors and revisions, pointing toward over-reaction. Fuhrer (2018), on the other hand, finds evidence of expectations smoothing and stickiness also at the individual level. If individual forecasters have sticky expectations but also extrapolate recent news about the current period, then one can indeed find either under-reaction or over-reaction, depending on the methodology and the data used. In fact, whether forecasts over-react or under-react to a particular piece of news will then depend on which quarter is affected the most by the news. Due to extrapolation, forecasts about future periods will overreact to news that has a large impact on the current period. On the other hand, when the news has little effect on the current period and mainly concerns future periods, extrapolation will be weak. In this case, under-reaction to news due to stickiness in expectations will arise.

Furthermore, in a laboratory experiment where subjects form expectations about a simple AR(1) process, Landier et al. (2019) find that an expectations rule that combines stickiness and extrapolation fits well to their experimental data and outperforms other expectations rules. The estimated coefficients in this expectation rule are robust to the autocorrelation coefficient in the AR(1) process that is being forecast and to other parameters of the model, as well as to the framing of the experiment (e.g. as forecasts about a macroeconomic process).

An important difference between the experimental setup of Landier et al. (2019) and the real world is that, in the experiment, subjects have no information at all about what others in the economy have been forecasting. In the real world, there are many sources of information on what experts have recently been forecasting about future periods. For example, the consensus forecast of the survey of professional forecasters is published on the website of the Federal Reserve Bank of Philadelphia each quarter. Also, the

³⁵The advantages of using the SPF include the long time span of data availability at a quarterly frequency and the fact that one can observe forecasts made at different time points for different forecasting horizons for the same individual forecaster. The latter is necessary to calculate revisions and to investigate stickiness in expectations.

SPF has a large audience: the Philadelphia’s FED external webpages counted more than 45,000 unique hits because of the SPF in 2018 (Croushore et al., 2019). Further, numerous (economic) institutions frequently publish their forecasts which also find their way into newspapers and, therefore, to the large public. A lagged consensus forecast may, thus, be seen as a public signal that forecasters observe and could base their current expectations on. Fuhrer (2018) finds that there, indeed, is a strong dependence of current individual forecasts on the lagged consensus forecast, in the SPF and the Livingston survey as well as in the Michigan survey of consumer expectations.

For this reason, we allow for a further deviation of rational expectations. In addition to stickiness to past individual expectations and extrapolation of current news, expectations may further deviate from the fully rational benchmark by partly being based on the most recently observed consensus forecast about the period that is being forecast. Such a bias can be interpreted as meaning that individual forecasters do not have the skills or the time to efficiently use all available pieces of information to come up with a fully rational forecast all by themselves. Therefore, they instead only partly base their forecast on new information and form only part of the rational expectations solution. For the other part, they make use of the “wisdom of the crowd” and turn to an easy to use piece of information: the most recently observed consensus forecast.

When we fit our model, we find evidence that all three deviations from rationality are present in the SPF data. Moreover, larger forecasting horizons turn out to imply more sticky expectations. All in all, we find (across different macroeconomic variables) estimates of stickiness in forecasts that lie between 0.41 and 0.76. Furthermore, estimates of the extrapolation parameter of current news lie between 0.18 and 0.91, with larger forecasting horizons implying smaller estimates for most variables. Finally, expectations are for a significant part formed based on the most recent past consensus forecast about the period being forecast. In particular, forecasters are found to put weights between 0.23 and 0.83 on the most recently observed consensus forecast rather than on a rational expectations forecast that they would have to build by themselves. This indicates that individual forecasters have a hard time processing all available (new) information and building an own rational expectations forecast.

Interestingly, we further show that 20-50% of the forecasters in the survey would have been better off in terms of forecasting performance if they would have abandoned all attempts to build their own forecast and instead blindly submitted the most recently observed consensus forecast in every period. Basing a forecast on the lagged consensus can, hence, be seen as a reasonable thing to do. We further document that individual forecasts of the SPF can equivalently be modeled as depending on the lagged consensus forecast of the Livingston survey rather than the lagged consensus of their own survey.

This supports the hypothesis that forecasters base their forecasts on a more general consensus forecast that can be obtained from different sources and that can reach a large audience.

Moreover, our results are not a specific artifact of the forecasting behavior of the SPF. We find qualitatively similar estimates when we perform our regression for the individual forecasts of the Livingston survey. The stickiness estimates are, however, lower than for the SPF. This can be explained by the semi-annual structure of the Livingston survey compared to the quarterly structure of the SPF. For instance, a forecaster may regard her last individual forecast, which is half a year old, as less informative compared to an individual forecast that would have been only three months old. This then results in lower estimates of the stickiness bias.

In addition to fitting our model of expectation formation to survey data, we also test a restriction on the coefficients of an over-identified econometric specification of our model. In particular, the sum of three of the coefficients in that regression should be equal to 1 according to the theoretical formulation of the model. We can, hence, estimate this over-identified specification and see whether the restriction holds. We find for all our selected variables and forecasting horizons that the sum of the coefficients is very close to what the restriction implies. Statistically, the restriction is not rejected for most of the cases. We, therefore, conclude that our proposed model of expectation formation is not rejected by the data.

Many alternative models of expectations that deviate from full rationality have been proposed. A considerable body of literature assumes that expectations are formed in a purely backward-looking manner, either by all or by a fraction of agents in the economy. That is, (some) households and firms consider past realizations of macroeconomic variables and base their expectations about future variables solely on this. The most well-known example of such expectations is found in the *adaptive learning* literature (Evans and Honkapohja, 2012), where agents use statistical learning based on past realizations to learn the relations between economic variables. Further, in the *heterogeneous expectations* literature, backward-looking expectations take the form of a simple adaptive or extrapolative heuristic, where past realizations of a particular variable are used to predict future realizations of that variable (see e.g. Branch and McGough, 2010; De Grauwe, 2012; Hommes and Lustenhouwer, 2019).

A disadvantage of modeling expectations as purely backward-looking is that there can be little to no role for news and for announcements about the future. Backward-looking expectations can, hence, not meaningfully contribute to, e.g., the question of whether forecasts over-react or under-react to news. Moreover, the literature on empirical validation of backward-looking expectation formation is not large and not

fully convincing. Backward-looking expectations generate persistence in expectations and in the macroeconomic model. This type of expectation is, hence, a potential explanation for the larger persistence that is found in survey data compared to rational expectations. This explains why expectation models that include backward-looking components improve the fit with the data (Branch, 2004; Milani, 2007; Slobodyan and Wouters, 2012; Cornea-Madeira et al., 2019). However, there is no convincing evidence that backward-lookingness is the correct explanation for higher persistence in survey expectations and that expectations are actually formed by considering past observations.

In fact, when we add to our model a potential fourth bias, where forecasters partly use the lagged realization of the variable being forecast to form expectations, the estimated magnitude of this backward-looking bias is small and almost always statistically insignificant. This is a clear indication that the persistence found in survey data is not caused by agents basing their expectations on past observations, but instead by other biases such as stickiness. Similar conclusions are drawn by Fuhrer (2018) from survey data and by Landier et al. (2019) in a laboratory experiment.

Other deviations from full-information rational expectations that have been proposed include models with sticky information or noisy information and rational inattention (Mankiw and Reis, 2002; Woodford, 2001; Sims, 2003; Mackowiak and Wiederholt, 2009). Moreover, Gabaix (2020) proposes a framework of sparsity to model deviations from fully rational expectations. These theoretical frameworks may offer partial explanations for e.g. why expectations are sticky. These explanations are, however, not always in line with survey data of *individual* expectations (Bordalo et al., 2018a; Fuhrer, 2018).

In this paper, we do not focus on *why* forecasters have certain deviations from full rationality. Instead, we seek to converge on a model of expectation formation that combines different higher-level individual biases and that can explain actual expectations from survey data. Such a model of expectation formation can, in future research, be used as a building block to improve the realism of the assumptions made on expectations in macroeconomic modeling. In particular, our findings stress the importance of modeling the interaction between expectations of individual agents in the economy and of explicitly modeling (over-)reaction to news. We believe that deriving policy implications from models where expectation formation is more in line with individual expectation data can greatly help the field forward.

The rest of the paper is organized as follows. In Section 5.2, we outline our proposed model of expectations formation. In Section 5.3, we derive two empirical formulations of the model, and our main estimation results are presented in Section 5.4. Next, we

discuss some extensions that contribute to the interpretation and robustness of results in Section 5.5. Section 5.6 presents a discussion on the role of the lagged consensus forecast and Section 5.7 concludes.

5.2 A behavioral model of expectation formation

We first discuss three behavioral deviations from fully rational expectations that have already found some empirical support. We then combine these three biases into a behavioral model of expectation formation. In later sections, we will confront this model with data from survey expectations to see whether the proposed biases are present in the data.

5.2.1 Stickiness

Fuhrer (2017) replaces, in a macroeconomic model, rational expectations with actual expectations from survey data. He shows that this makes usual modeling elements to generate macroeconomic persistence such as habit formation and price indexation obsolete. Moreover, Coibion et al. (2012) and Angeletos et al. (2020) show that expectations initially under-react to economic shocks. Also, Coibion and Gorodnichenko (2015) document that forecast errors of aggregate expectations in survey data are forecastable by aggregate revisions. This implies under-reaction to new information and considerable persistence in expectations.

There is also empirical evidence of sticky expectations at the *individual level* in the case of financial analysts (Bouchaud et al., 2019), firm managers (Ma et al., 2018) and participants in laboratory experiments (Landier et al., 2019). Furthermore, Fuhrer (2018) finds that individual forecasts are *intrinsically persistent*, supporting his early findings on the macroeconomic level.

Stickiness in individual expectations implies that expectations of individuals depend on their most recent own past forecast about the same period. From a rational perspective, there is no reason why this should be the case.³⁶ When an individual i forms expectations with stickiness as the only deviation from rationality, her expectations h

³⁶A psychological explanation for why forecasts may be formed in such a way could e.g. be a confirmation bias (Nickerson, 1998). See also the discussion at the end of Section 5.5.2. As mentioned in the introduction, our focus does, however, not lie on explaining *why* certain biases arise. Instead, we aim to come up with a higher-level model of expectation formation that is in line with individual survey data.

periods ahead about some variable x can be written as

$$F_{i,t}x_{t+h} = \lambda F_{i,t-1}x_{t+h} + (1 - \lambda)E_t x_{t+h}. \quad (5.1)$$

where $F_{i,t}$ denotes individual, possibly non-rational, expectations, and where E_t is the rational expectations operator. This is the first behavioral bias that we allow for and will let the data speak to.

5.2.2 Forward-looking extrapolation

The second bias we consider is forward-looking extrapolation. This is a variation of diagnostic expectations (Bordalo et al., 2018b). In both cases, individual expectations contain a “kernel of truth” as they depend one-for-one on the associated rational expectations forecast. However, additionally, agents overreact to recent news. In the case of diagnostic expectations, this news is taken to be the most recent rational expectations revision about the period being forecast. Landier et al. (2019) find, however, that extrapolation of news about the *current* period better and more robustly fits their experimental data.

The motivation for such a behavioral bias is that people may be particularly affected by surprises and news that concern the *current* state of affairs. This explanation fits well with the notion of “representativeness” of Kahneman and Tversky (1973) where some new information that is regarded as representative is extrapolated when making a prediction. That is, news about today may be viewed as more “representative” about both current and future economic conditions than news that only concerns the future. Thus, if forecasters are surprised by a higher realization of a variable today than they were expecting before, they extrapolate this surprise also into the future. As a consequence, forecasters then adjust their expectations about the future upward, more than would be rational. Further evidence for this can be found in the randomized control survey experiment of Coibion et al. (2019). In this study, when households are presented with news about *current* interest rates, their expectations about *future* interest rates are considerably more affected than when they are presented with news about *future* interest rates.

When extrapolation is the only behavioral bias, expectations are given by

$$F_{i,t}x_{t+h} = E_t x_{t+h} + \gamma News_{t,t}, \quad (5.2)$$

where $News_{t,t}$ represents news arriving in period t that concerns realizations in period t .

If the current period would be fully observed, as it is in the experiment of Landier et al. (2019), this news term would take the form

$$News_{t,t} = x_t - E_{t-1}x_t. \quad (5.3)$$

In practice, however, current values are not fully observed at the time that forecasts are made. To be consistent with the empirical data, we, therefore, do not assume actual realizations in period t to be observable in period t in our model of expectations. Consequently, current news about the current period is equal to today's *rational nowcast* minus the previous rational forecast about today:

$$News_{t,t} = E_t x_t - E_{t-1} x_t. \quad (5.4)$$

Plugging this in into (5.2) gives

$$F_{i,t}x_{t+h} = E_t x_{t+h} + \gamma(E_t x_t - E_{t-1} x_t). \quad (5.5)$$

5.2.3 Consensus forecast

The above two biases assume that forecasters, in principle, are able to build something like fully rational expectations, but that they systematically deviate from the rational forecast because of a behavioral bias.

Forming rational expectations, in a world where the true model and its parameterization are unknown and may vary over time is, however, not an easy task for an individual forecaster. In particular, it may be infeasible for an individual forecaster to try to combine all information on all relevant factors and variables in the economy in a rational way. Instead of trying to construct a fully rational expectations forecast all by herself, an individual forecaster may, therefore, also consider information on expectations of other forecasters when forming her own forecasts.

A consensus forecast of a large group of (professional) forecasters can be seen as a very valuable source of information for an individual forecaster. In particular, this is an example of the well known “wisdom of the crowd” (Surowiecki, 2004; Davis-Stober et al., 2014). As a consequence, we find that mean squared forecast errors of consensus forecasts are considerably lower than those of individual forecasters.

In real-time, a typical forecaster has only very little information available about what other forecasters are expecting and a general consensus forecast cannot be derived from this. However, with a time lag, there are plenty of direct and indirect sources of information that reflect a consensus forecast. For example, the consensus forecast of

the survey of professional forecasters is published on the website of the Federal Reserve Bank of Philadelphia each quarter. Further, numerous institutions frequently publish their forecasts which also find their way into newspapers and blogs and, thus, to the larger public.

When forecasters partly base their forecasts on a lagged consensus forecast, they do not fully take into account all information that is available now but was not yet available in the previous period. We find, however, that even a *lagged* consensus forecast can still be seen as a useful piece of information for many forecasters. We will discuss this in detail in Section 5.6. Moreover, Fuhrer (2018) finds evidence that forecasters indeed base their forecasts significantly on a lagged consensus forecast. This holds for the Survey of Professional Forecasters and the Livingston survey, as well as for the Michigan survey of consumer expectations.

When forecasters only partially are able to form rational expectations themselves, and for the other part base their expectations on a lagged consensus forecast, their expectations become

$$F_{i,t}x_{t+h} = (1 - \delta)E_t x_{t+h} + \delta C_{t-1}x_{t+h}, \quad (5.6)$$

where C_{t-1} is a lagged consensus forecast (e.g. the median of the expectations of all expert forecasters).

5.2.4 Combining the three behavioral biases

Next, we will combine the three biases introduced above into a behavioral model of expectation formation. Suppose that when agents form expectations, they are only partially able to do this rationally by themselves. For the other part, they chose to base their expectations on a lagged consensus forecast, as in (5.6). Suppose further that they also have the stickiness bias of Section 5.2.1, so that they anchor their newly formed expectations to their previously formed own expectations. Expectations would then become

$$F_{i,t}x_{t+h} = \lambda F_{i,t-1}x_{t+h} + (1 - \lambda) [(1 - \delta)E_t x_{t+h} + \delta C_{t-1}x_{t+h}]. \quad (5.7)$$

Next, assume that forecasters also are subject to the forward-looking extrapolation bias of Section 5.2.2. That is, agents base their expectations on (5.7) and then deviate from this by extrapolating the most recent news. Replacing $E_t x_{t+h}$ in (5.5) with the

RHS of (5.7) gives our final model of behavioral expectations,

$$F_{i,t}x_{t+h} = \lambda F_{i,t-1}x_{t+h} + (1 - \lambda) [(1 - \delta)E_t x_{t+h} + \delta C_{t-1}x_{t+h}] + \gamma(E_t x_t - E_{t-1}x_t) + \mu_{i,t}, \quad (5.8)$$

where we have added an idiosyncratic white noise term, $\mu_{i,t}$, to allow for additional heterogeneity among forecasters.

Under this model of expectation formation, there can be both over-reaction and under-reaction to news. The deciding factor, here, is whether the piece of news only concerns the period that is being forecast ($t + h$) or that the news also considerably impacts realizations in the current period (t). If the news is only relevant for period $t + h$, then forecasters will under-react to it. This is because the news then only shows up in the term $E_t x_{t+h}$, to which forecasters respond in a dampened manner with coefficient $(1 - \lambda)(1 - \delta)$.

In general, however, one expects there to be a correlation between current news about $t+h$ ($News_{t,t+h}$) and current news about t ($News_{t,t}$). In that case, there can also be over-reaction to news. This is because the news will now also affect $E_t x_t - E_{t-1}x_t (= News_{t,t})$. Forecasters will then, on the one hand, respond to the news in a dampened manner, but at the same time over-react to the news with coefficient γ . This can result in a net over-reaction to the news if $News_{t,t}$ is enough affected relative to $News_{t,t+h}$.

Note, further, that even though expectations are no longer one-for-one based on rational expectations, rational expectations still are a component of “base” expectations (5.7). Therefore, (although a smaller one) the behavioral expectations in (5.8) still contain a “kernel of truth”. Moreover, since we will estimate the parameters of our model using survey data, we are not assuming that any of the above biases *must* be present. Instead, we will let the data decide which bias is present and with what magnitude. If it turns out that $\lambda = \gamma = \delta = 0$, then (in the absence of idiosyncratic noise) fully rational expectations are obtained in (5.8).

5.3 Empirical model and estimation method

Before we can confront the model of expectation formation derived in the previous section with survey data, we need to rewrite it somewhat. This is because rational expectations and the news term in Equation (5.8) are not directly observable.

As mentioned in Section 5.2.3, the consensus forecasts of the SPF are considerably better than the individual forecasts. This indicated that, for our purposes, it might be possible to approximate the news variable, $E_t x_t - E_{t-1}x_t$, with the revision in the

consensus forecast about the current period, $C_t x_t - C_{t-1} x_t$. Here, $C_t x_t$ is the consensus nowcast and $C_{t-1} x_t$ the lagged consensus forecast about today.

Coibion and Gorodnichenko (2015) find, however, that for a longer forecasting horizon, there is a systematic bias in the consensus forecast that is reflected in a dependence of forecast errors on forecast revisions. We should, therefore, check whether this is also the case for consensus nowcast revisions and possibly control for this when approximating news.

In particular, modifying the steps of Coibion and Gorodnichenko (2015) to the case of nowcasts, we have the following: whereas nowcast errors under rational expectations are given by

$$x_t - E_t x_t = \varepsilon_{t,t}, \quad (5.9)$$

nowcast errors of the consensus might be given by

$$x_t - C_t x_t = \psi(C_t x_t - C_{t-1} x_t) + \varepsilon_{t,t}, \quad (5.10)$$

for some value of ψ . Here, $\varepsilon_{t,t}$ is the rational expectations nowcast error that reflects information about period t that is not yet available when forecasts in period t are made.

Combining (5.9) and (5.10), we can write

$$E_t x_t = C_t x_t + \psi(C_t x_t - C_{t-1} x_t). \quad (5.11)$$

Taking the same equation in terms of expectations formed one period earlier and subtracting it from (5.11) gives, after rearranging,

$$News_{t,t} = E_t x_t - E_{t-1} x_t = (1 + \psi)(C_t x_t - C_{t-1} x_t) - \psi(C_{t-1} x_t - C_{t-2} x_t). \quad (5.12)$$

That is, we obtain an expression for news in terms of the observable consensus forecasts that is valid even when consensus revisions are inefficient.³⁷

The estimation results of (5.10) are presented in Appendix 5.B. It turns out that for inflation and nominal GDP the resulting estimated value of ψ is small and not statistically different from zero. For these variables, we therefore set $\psi = 0$ in (5.12) which

³⁷Note that we have implicitly assumed that the value of ψ obtained from the nowcast regression, (5.10), is also valid for one-period-ahead forecasts. For inflation and nominal GDP, the estimates of ψ in both the nowcast regressions and an analogue one-period-ahead forecast regression are not statistically significantly different from zero. Here, it is, therefore, sensible to stick with $\psi = 0$. For real GDP and unemployment, it cannot be rejected that the estimates in the one-period-ahead forecast regressions are equal to values of ψ that are obtained from the respective nowcast regressions. Therefore, we choose to stick with the estimates from the nowcast regressions for these two variables.

then reduces to $News_{t,t} = C_t x_t - C_{t-1} x_t$. However, for real GDP and unemployment we find statistically significant values of ψ so that $News_{t,t}$ depends on the current as well as on the lagged nowcast revision for these two variables.

Next, we need to replace $E_t x_{t+h}$ in Equation (5.8) with an observable. Taking the consensus forecast about period $t+h$ would not be a good enough approximation here. However, similar to Clarida et al. (1998) and Galí and Gertler (1999) and others, we can use the analog of (5.9) for h -period ahead rational expectations,

$$x_{t+h} = E_t x_{t+h} + \varepsilon_{t,t+h}, \quad (5.13)$$

where $\varepsilon_{t,t+h}$ is the current rational expectations error regarding period $t+h$. Again, this term reflects information about period $t+h$ that is not yet available when forecasts are made in period t . Moreover, to be as close as possible to the information set of forecasters, we use the first vintage realizations of the actuals in $t+h$. When considering actual realizations in this literature, using first vintages has become a standard procedure (see Bordalo et al. (2018a) and Coibion and Gorodnichenko (2015)). However, we also provide robustness in Appendix 5.C.1 where we use final vintages. Further, a detailed description of our variables is given in Appendix 5.A.

Using (5.13) and (5.12), we can write (5.8) as

$$\begin{aligned} F_{i,t} x_{t+h} = & \lambda F_{i,t-1} x_{t+h} + (1 - \lambda) [(1 - \delta)x_{t+h} + \delta C_{t-1} x_{t+h}] + \gamma News_{t,t} \\ & - (1 - \lambda)(1 - \delta)\varepsilon_{t,t+h} + \mu_{i,t}, \end{aligned} \quad (5.14)$$

which can be estimated with empirical data. Note that the rational expectations error term $\varepsilon_{t,t+h}$ is orthogonal to all information dated t and earlier by definition. Therefore, this term does not correlate with the variables on the RHS of (5.14), except for the term x_{t+h} .

Hence, in the empirical formulation of (5.14),

$$F_{i,t} x_{t+h} = \beta_1 F_{i,t-1} x_{t+h} + \beta_2 x_{t+h} + \beta_3 C_{t-1} x_{t+h} + \beta_4 News_{t,t} + e_{i,t}, \quad (5.15)$$

we follow the approach of Clarida et al. (1998) and Galí and Gertler (1999) and instrument for x_{t+h} using a GMM estimation. As instruments we take four lags of actual inflation, unemployment, the T-Bill rate and a measure of GDP.³⁸ As these instruments are in forecasters' time- t information set, they are, by definition, not correlated with

³⁸In the estimations where x is inflation, unemployment or real GDP, we use real GDP as the relevant GDP measure, whereas in the estimations where x is nominal GDP we use nominal GDP instead.

the current rational expectations error. Moreover, as all our actual variables are highly persistent, these lagged terms explain x_{t+h} well.³⁹ Further, we allow for clustering across individuals and time periods (two-way clustering) when calculating standard errors and when calculating the GMM weighting matrix.⁴⁰

Another specification of the model can be obtained by subtracting $F_{i,t-1}x_{t+h}$ from both sides of (5.14). We can then write

$$F_{i,t}x_{t+h} - F_{i,t-1}x_{t+h} = (1 - \lambda) [(1 - \delta)(x_{t+h} - F_{i,t-1}x_{t+h}) + \delta(C_{t-1}x_{t+h} - F_{i,t-1}x_{t+h})] + \gamma News_{t,t} - (1 - \lambda)(1 - \delta)\varepsilon_{t,t+h} + \mu_{i,t}. \quad (5.16)$$

Here, the revision in forecaster's i forecast about period $t + h$ depends on how the actual realization deviates from the individuals' lagged forecast. This term indicates how the forecast is revised (rationally) because of news about $t + h$ in period t , which was not yet available in period $t - 1$ (when the previous forecast was made). However, this rational updating to news is only partial, and, for another part, the revision depends on how the most recently observed consensus forecast deviates from the forecaster's own previous forecast. Forecasters, hence, partly update their forecasts towards the consensus. Stickiness causes adjustments to be even more partial (i.e. smaller revisions), as can be seen from the scaling with $1 - \lambda$. Finally, there still is the extrapolation term, just as in (5.14), that leads to larger forecast revisions when there is positive news about the current period.

The empirical version of (5.16) is

$$F_{i,t}x_{t+h} - F_{i,t-1}x_{t+h} = a_1(x_{t+h} - F_{i,t-1}x_{t+h}) + a_2(C_{t-1}x_{t+h} - F_{i,t-1}x_{t+h}) + a_3 News_{t,t} + e_{i,t}. \quad (5.17)$$

Analogue to (5.15), we estimate this equation using IV-GMM to instrument for $(x_{t+h} - F_{i,t-1}x_{t+h})$. We adjust the above mentioned instruments accordingly by subtracting $F_{i,t-1}x_{t+h}$ from each of them.

³⁹The R^2 in regressions of x_{t+h} on the instruments always lie above 0.9.

⁴⁰Note that our empirical formulation does not contain an individual constant. The reason is that there is, to the best of our knowledge, no evidence in the literature suggesting an individual *time-invariant* bias. Therefore, our estimations do not involve first difference (FD) or fixed-effects (FE) transformations of the data. Also, employing FD or FE would lead to serious practical problems here. In particular, FE creates a dependence of the error term with the past, so that no valid instruments would be available. Further, FD would introduce the last periods' rational expectations error, which is not orthogonal to period t information, and, thus, correlates with the news term. Hence, FD would create an additional endogeneity problem.

5.4 Confronting the model with survey expectations

For our main results, we utilize expectations provided by the Survey of Professional Forecasters (SPF). The SPF is a quarterly survey published first in the fourth quarter of 1968 by the National Bureau of Economic Research and since 1990 published by the Federal Reserve Bank of Philadelphia. Around 40 participants provide forecasts about a variety of variables in the current and next four quarters.⁴¹ The survey is conducted at the end of the second month in each given quarter. Hence, forecasters know the realizations of quarterly data up to quarter $t - 1$ and for unemployment up to the previous month. In our sample, forecasters stay in the survey for about 41 quarters on average.

As can be seen in Table 5.1, we calculate expectations about inflation, growth rates for nominal and real GDP, and the unemployment rate.⁴² As the consensus forecast, we take the median forecast of all forecasters in the panel. Further, for actual realizations of variables, we use the Real-Time Data Set for Macroeconomists of the Philadelphia FED. A detailed description of our variable construction is given in Appendix 5.A. Finally, we follow Angeletos et al. (2020) by trimming outliers where forecast revisions lie more than 4 times the inter-quartile range from the median revision. Here, the inter-quartile range and median of revisions are calculated over the entire sample underlying the respective regression.

Table 5.1: Expectation data based on SPF

Variable	Period
Inflation	1970:Q1 - 2019:Q3
Nominal GDP Growth Rate	1970:Q1 - 2019:Q3
Real GDP Growth Rate	1970:Q1 - 2019:Q3
Unemployment Rate	1970:Q1 - 2019:Q4

Note: The period covered for each variable may slightly vary with the forecasting horizon.

Using this data, we fit both forecast revisions based on Equation (5.17) (in Section 5.4.1) and individual forecasts as in Equation (5.15) (in Section 5.4.2). Note that the

⁴¹For some years, around the time where the Philadelphia FED took over, the SPF had exceptionally few respondents. However, this only amounts to roughly 170 observations of more than 5000. Excluding these observations does not significantly affect our results.

⁴²We chose to calculate inflation expectations from expectations about the GDP price index as it covers the same time period as the rest of our variables. Expectations about the CPI inflation rate are available from 1981:Q3 only which we, therefore, do not consider.

coefficients in equation (5.17) identify the underlying theoretical model parameters exactly. This is equivalent to imposing a restriction on the estimation of Equation (5.15), which we subsequently test.

We start with presenting the results for one-period-ahead forecasts ($h = 1$) and show result for $h = 2$ and $h = 3$ in Section 5.4.3. For a forecasting horizon of $h = 4$, no lagged individual forecasts about the same period are available, so we cannot perform our estimations for that case.

5.4.1 Fitting forecasting revisions

In Table 5.2, we presents the results of the revisions regression, (5.17), for the SPF for the four different macroeconomic variables introduced above for the case of $h = 1$.

Table 5.2: Regression of Equation (5.17) with $h = 1$ based on SPF
Dependent Var.: $F_{i,t}x_{t+1} - F_{i,t-1}x_{t+1}$

	Inflation	Nominal GDP	Real GDP	Unemp. Rate
$x_{t+1} - F_{i,t-1}x_{t+1}$	0.095*** (0.024)	0.155*** (0.054)	0.235*** (0.055)	0.184*** (0.060)
$C_{t-1}x_{t+1} - F_{i,t-1}x_{t+1}$	0.453*** (0.035)	0.417*** (0.056)	0.311*** (0.060)	0.407*** (0.059)
$News_{t,t}$	0.908*** (0.045)	0.840*** (0.073)	0.630*** (0.076)	0.634*** (0.086)
R^2	0.66	0.70	0.66	0.70
n	5581	5613	5682	5850

Note: *** p<0.01, ** p<0.05, * p<0.1. Estimates are rounded to the third digit after the comma.

In the table, we see that all included regressors significantly influence the individual revisions. Not surprisingly, and in line with earlier literature, this shows that the rational expectations hypothesis, i.e. in our case $\lambda = \gamma = \delta = 0$, is clearly rejected. Further, these results substantiate our proposed alternative: SPF forecasters indeed anchor their forecasts to their previous own forecast, extrapolate news about the current period into the future and make use of the lagged consensus forecast when forming their own expectations.

Next, we want to obtain estimates of λ , δ and γ . By comparing Equation (5.17) with (5.14) it follows that $\lambda = 1 - a_1 - a_2$, $\delta = \frac{a_2}{a_1 + a_2}$ and $\gamma = a_3$. For each variable, an estimate of λ , δ and γ can so be obtained. We find estimates of stickiness (λ) lying between 0.41 and 0.45. The estimates of the extent to which individuals are not able to form their own rational expectations but instead use the most recent consensus forecast (δ) vary from 0.57 to 0.83. Finally, the estimated extrapolation parameter (γ) lies between 0.63 and 0.91. More detailed results are presented in Table 5.4 which will be discussed in Section 5.4.3.

5.4.2 Performing a test

Next, we turn to the regression based on Equation (5.15). This specification is over-identified, as, according to our theoretical model, $\beta_1 + \beta_2 + \beta_3 = \lambda + (1 - \lambda)\delta + (1 - \lambda)(1 - \delta) = 1$. Hence, this specification provides a good opportunity to test the validity of the restriction that our model of expectation formation implies. When the sum of these three coefficients in this regression is close to 1, this indicates that expectation formation in the data is, at least to some extent, in line with our model.

The estimation results are presented in Table 5.3. Also, the sums of these three coefficients, $\beta_1 + \beta_2 + \beta_3$, are shown at the bottom of the table. Interestingly, these sums are surprisingly close to 1 across all variables and range from 0.997 to 1.007.

Further, to conduct a formal test, we use a Wald test to test whether $\beta_1 + \beta_2 + \beta_3 = 1$ can be rejected. As can be seen there, this restriction cannot be rejected by the data at the 5% significance level for inflation, Nominal GDP and real GDP. Given the large amount of observations (over 5500 for these three variables), this is a strong indication that the restriction seems to have some validity, and that our model captures some important features of the data.⁴³

In the case of unemployment, the Wald test does reject the restriction that $\beta_1 + \beta_2 + \beta_3 = 1$. Still, the sum of the three coefficients is quite close to 1 here (1.007). Therefore, the rejection reflects even smaller standard errors rather than a clear economically significant misspecification of the model.

It can further be seen in Table 5.3 that the estimated values of β_2 , β_3 and β_4 are close to the estimates of α_1 , α_2 and α_3 in Table 5.2. This is further indication that the theoretical model is not rejected by the data, as these coefficients are equal to respectively $(1 - \lambda)(1 - \delta)$, $(1 - \lambda)\delta$ and γ in both Equation (5.14) and Equation (5.16).

⁴³Using a (robust) F-Test for the validity of the restriction gives identical results, which is not surprising given a large number of observations and the asymptotic equivalence with the (robust) Wald-test in this case.

Table 5.3: Regression of Equation (5.17) with $h = 1$ based on SPF
Dependent Var.: $F_{i,t}x_{t+1}$

	Inflation	Nominal GDP	Real GDP	Unemp. Rate
$\beta_1 (F_{i,t-1}x_{t+1})$	0.444*** (0.027)	0.427*** (0.019)	0.458*** (0.019)	0.414*** (0.012)
$\beta_2 (x_{t+1})$	0.111*** (0.022)	0.162*** (0.042)	0.226*** (0.051)	0.270*** (0.057)
$\beta_3 (C_{t-1}x_{t+1})$	0.447*** (0.036)	0.408*** (0.045)	0.323*** (0.057)	0.323*** (0.056)
$\beta_4 (News_{t,t})$	0.878*** (0.040)	0.831*** (0.062)	0.637*** (0.074)	0.535*** (0.083)
Wald-Test on restriction $\beta_1 + \beta_2 + \beta_3 = 1$				
$\beta_1 + \beta_2 + \beta_3 =$	1.002	0.997	1.007	1.007
p-value	0.505	0.227	0.142	0.000
R^2	0.991	0.992	0.975	0.999
n	5581	5613	5682	5850

Note: *** p<0.01, ** p<0.05, * p<0.1. Estimates are rounded to the third digit after the comma. The R^2 are extremely high because $F_{i,t-1}x_{t+1}$ explains a very large portion of $F_{i,t}x_{t+1}$. This is because period-to-period changes in expectations about a particular period are much smaller than period-to-period changes in the variable being forecast. Hence, the high R^2 are here not an indication of spurious correlations. This is confirmed by the fact that the estimates of β_1 , β_2 and β_3 are very close to the estimates of the revision regression in Table 5.2. Moreover, the results even become equivalent when we estimate an 'over-identified' version of the revision regression by adding the lagged individual forecast as an additional regressor.

5.4.3 Forecasting horizon

Now, we consider results for different forecasting horizons. That is, we first estimate Equation (5.17) for $h = 2$ and for $h = 3$. The estimation results can be found in Tables 5.15 and 5.16, respectively, in Appendix 5.D.

From the estimated coefficients, which are all statistically significant at the 1%-level, we again calculate estimates of the parameters λ , δ and γ , as described in Section 5.4.1. These are presented in Table 5.4 for the three different forecasting horizons.

In Table 5.4, it can be seen that the results for $h = 2$ and $h = 3$ are qualitatively in line with those discussed in Section 5.4.1. However, there are quantitative differences

Table 5.4: Estimates of behavioral parameters for different forecasting horizons.

SPF	λ			δ			γ		
	h=1	h=2	h=3	h=1	h=2	h=3	h=1	h=2	h=3
Inflation	0.45	0.52	0.76	0.83	0.71	0.29	0.91	0.76	0.22
NGDP	0.43	0.54	0.72	0.73	0.40	0.23	0.84	0.73	0.31
RGDP	0.45	0.58	0.74	0.57	0.36	0.49	0.63	0.54	0.18
Unemployment	0.41	0.54	0.60	0.69	0.56	0.69	0.63	0.56	0.64
Average	0.44	0.55	0.71	0.71	0.51	0.43	0.75	0.65	0.34

in the estimates of λ , δ and γ . In fact, there seems to be a clear pattern, where the estimates of λ become higher the larger the forecasting horizon, and the estimates of γ become lower the larger the forecasting horizon. The latter, however, with the exception of unemployment. Although, individual estimates of δ fluctuate considerably, there does not seem to be a clear pattern that holds for all variables there.

It is intuitive that boundedly rational extrapolation of current news into future periods becomes weaker the further these periods lie in the future. That is, a forecaster may be considerably biased in her one-quarter-ahead forecast when she is surprised by a shock or news in the current period, but her three-quarter-ahead forecast may be considerably less distorted by what is currently happening in the economy. A lower value of the extrapolation parameter γ for larger forecasting horizons is, therefore, understandable.

The finding that expectations become more sticky (λ higher) as the forecasting horizons increases might be explained by a corresponding increase in forecasting uncertainty. In particular, there may be less (accurate) new information available about periods that lie further in the future than about periods that lie less far in the future. Forecasters have, therefore, less reason to revise their three-period-ahead forecast than they have to revise their forecast about the next quarter. As a consequence, they may respond to this by updating even less than would be rational and have a higher stickiness parameter for larger forecasting horizons. We obtain estimates of λ that range from 0.41 and 0.76 across variables and horizons. The literature on stickiness and expectation smoothing finds values between 0.16 and 0.65 across different measures, forecasting horizons and data sources (Fuhrer, 2018; Ma et al., 2018; Bouchaud et al., 2019; Landier et al., 2019).

These results are qualitatively robust to using final vintages instead of first vintages for actual realizations in Equation (5.17), as well as to limiting the sample to post-1984 (i.e. where the Great Moderation started) or to the Great Moderation period (1984-

2006). This is shown in, respectively, Appendix 5.C.1 and Appendix 5.C.2. However, individual estimates do quantitatively differ. Generally, the estimates of λ tend to be a little bit lower in the more stable time-subsamples, whereas the estimates of γ tend to be somewhat larger when final vintages are used.

Additionally, we run the unrestricted regression in Equation (5.15) for $h = 2$ and $h = 3$. We present the sum of the estimated coefficients β_1 , β_2 and β_3 as well as the p-values of the Wald test on the restriction $\beta_1 + \beta_2 + \beta_3 = 1$ in Table 5.5. For comparison, we also summarize the results for $h = 1$ here again.

Table 5.5: Wald-Test on restriction $\beta_1 + \beta_2 + \beta_3 = 1$

	Inflation	Nominal GDP	Real GDP	Unemp. Rate
h=1 (SPF)				
$\beta_1 + \beta_2 + \beta_3 =$	1.002	0.997	1.007	1.007
p-value	0.505	0.227	0.142	0.000
h=2 (SPF)				
$\beta_1 + \beta_2 + \beta_3 =$	0.992	0.990	0.993	1.008
p-value	0.028	0.025	0.329	0.000
h=3 (SPF)				
$\beta_1 + \beta_2 + \beta_3 =$	1.001	0.993	0.989	1.008
p-value	0.851	0.134	0.133	0.000

Note: Robust Wald-Test on the restriction $\beta_1 + \beta_2 + \beta_3 = 1$ in equation (5.15) for $h = 1, 2, 3$ (SPF). Test statistics are based on two-way clustered standard errors. Values are rounded to the third digit after the comma.

The results for $h = 2$ and $h = 3$ are generally in line with those of $h = 1$. The sums of coefficients are again close to 1, also for unemployment where the Wald test rejects an exact equality to 1. It can be seen in Table 5.5, though, that, for $h = 2$, the restriction $\beta_1 + \beta_2 + \beta_3 = 1$ is rejected at the 5%-level for inflation and nominal GDP (but not at the 1%-level). However, also here, the sums are quite close to 1.

5.5 Extensions

Below, we consider several extensions of the main estimations of the previous section. These help to provide more intuition and robustness. In Section 5.5.1, we study how

we can interpret forecasters using the past consensus. In particular, we investigate whether this finding is limited to forecasters considering the consensus of the survey that they are participating in, or that, instead, our results indicate that forecasters may use a more general consensus in the economy. That is, we consider whether the consensus that forecasters base their expectations on might also have been obtained from other sources such as a different survey. Subsequently, we study the robustness of our results by using data of individual forecasts from the Livingston survey in Section 5.5.2. Finally, in Section 5.5.3, we consider an extension of the model of expectation formation where forecasts are also partly based on lagged observations of the variables being forecast (learning).

5.5.1 Consensus forecast from Livingston survey

In Section 5.4, it was assumed that the lagged consensus forecasts that the SPF forecasters base their expectations on is the median of the lagged individual forecasts of the SPF. It is, therefore, not clear whether the result that forecasters use this lagged median forecast should be interpreted as participants of the survey checking the latest publications of the survey that they are participating in, or that it can be interpreted more broadly. In this section, we investigate whether similar results would be obtained if the lagged consensus forecast is *not* obtained from the same survey that forecasters are participating in.

In particular, we perform similar estimations as in Section 5.4, but we let the lagged consensus forecast about the period being forecast, $C_{t-1}x_{t+h}$, no longer be equal to the lagged median forecast of the SPF. Instead, we set it equal to the median forecast of a different survey: the Livingston survey.

The Livingston survey was launched in June 1946 and is today published by the Federal Reserve Bank of Philadelphia twice a year. Forecasts relevant to this paper are made in June and December. Thus, only forecasts made in Q2 and in Q4 are available. This means we can only use the Livingston consensus as a lagged consensus forecast in the SPF estimations if we restrict ourselves to SPF individual forecasts made in Q1 and Q3. Moreover, forecasters in the Livingston survey are only asked to make predictions about two quarters and four quarters ahead. We can, hence, use the two-quarter-ahead consensus estimate in the SPF regression of $h = 1$ and the four-quarter-ahead consensus forecast for the SPF regression of $h = 3$ (since a *lagged* forecast in the regression must always have a horizon that is one period longer than the *current* individual forecast). For the $h = 2$ regressions, on the other hand, no suitable lagged consensus forecast is available from the Livingston survey.

A further restriction of data availability is that nowcasts are only reported from 1992 onward in the Livingston survey. For the two quarter ahead forecasts used in the $h = 1$ regression this is of no concern. However, for the $h = 3$ regressions we need to limit the sample to start in 1992, since no meaningful yearly growth rates can be calculated from a four-quarter-ahead index forecast without having a nowcast (see Appendix 5.A for further details on how we calculate growth rates). Finally, we focus on nominal GDP, real GDP and unemployment only, as forecasts about the GDP price index are not available.

To see directly the effect of replacing the lagged consensus forecast of the SPF with the lagged consensus forecast of the Livingston survey, we first re-estimate the SPF model under the same restrictions on the sample. That is, we throw out all observations with forecasts made in Q2 and Q4 and additionally limit the sample to post-1992 for $h = 3$. Next, we replace the consensus forecast with that of the Livingston survey and estimate the model with exactly the same samples.

We then calculate the values of λ , δ and γ that are implied by these estimations. In Tables 5.6 and 5.7 we present the results, comparing the estimates of the SPF consensus specification with the estimates of the Livingston consensus specification. Table 5.6 corresponds to the case of $h = 1$, whereas Table 5.7 corresponds to the case of $h = 3$.

In the tables, it can immediately be seen that the estimates of δ for the Livingston consensus specification are still quite large, ranging from 0.41 to 0.96. If anything, these estimates are larger than in the corresponding SPF specification, and, in most cases, both estimates lie quite close to each other. This implies that replacing the SPF lagged consensus forecast with the lagged consensus forecast from the Livingston survey does not overturn the result that individual SPF forecasters, for a considerable part, use a lagged consensus rather than their own rational expectations.

Moreover, when making pairwise comparisons of the λ and γ estimates in case of SPF consensus and Livingston consensus for $h = 1$, it can be seen in Table 5.6 that the estimates of λ are very close to each other, but that the estimates of γ are somewhat more affected. However, qualitatively this does not change much. Finally, Table 5.7 shows that, for $h = 3$, the estimates of both λ and δ are robust to the choice of consensus forecast.

All in all, it can, hence, be concluded that our estimation is quite robust to replacing the SPF consensus with the Livingston consensus. This is a clear indication that forecasters in the SPF are not responding to the lagged consensus from their own survey per se, but rather to a more general consensus among economic experts across the whole economy. Individual forecasters can obtain this general consensus from the lagged consensus of the SPF, but they might just as well obtain it from other sources.

Table 5.6: Estimates of behavioral parameters for different lagged consensus forecasts for h=1.

h=1 (Q1 & Q3)	λ		δ		γ	
	SPF cons.	Liv. cons.	SPF cons.	Liv. cons.	SPF cons.	Liv. cons.
NGDP	0.46	0.47	0.70	0.74	0.89	0.83
RGDP	0.46	0.46	0.58	0.54	0.75	0.70
Unemployment	0.41	0.44	0.61	0.96	0.53	0.70
Average	0.44	0.46	0.63	0.75	0.72	0.74

Table 5.7: Estimates of behavioral parameters for different lagged consensus forecasts for h=3.

h=3 (Q1 & Q3)	λ		δ		γ	
	SPF cons.	Liv. cons.	SPF cons.	Liv. cons.	SPF cons.	Liv. cons.
NGDP	0.69	0.68	0.80	0.81	0.61	0.57
RGDP	0.75	0.74	0.40	0.41	0.32	0.32
Unemployment	0.59	0.61	0.68	0.76	0.59	0.63
Average	0.68	0.68	0.63	0.66	0.51	0.51

5.5.2 Forecasts from Livingston survey

To further study the robustness of the results of Section 5.4, we next perform our estimations for the *individual forecasts* of the Livingston survey. As mentioned in the previous section, the survey is semi-annual, so that we only have observations of forecasts made in the second and fourth quarter. This also implies that the most recent lagged individual forecast was made two quarters ago rather than in the previous quarter. Equation (5.17), hence, becomes

$$F_{i,t}x_{t+h} - F_{i,t-2}x_{t+h} = a_1(x_{t+h} - F_{i,t-2}x_{t+h}) + a_2(C_{t-1}x_{t+h} - F_{i,t-2}x_{t+h}) + a_3News_{t,t} + e_{i,t}. \quad (5.18)$$

As we have shown in the previous section, the estimations are quite robust to whether we use the lagged consensus forecast of the SPF or of the Livingston survey. We, therefore, use the SPF lagged consensus in the above estimation equation, rather than the 2-quarter lagged consensus of the Livingston survey.⁴⁴ We also still proxy news with the SPF consensus revision rather than a Livingston consensus revisions, since

⁴⁴The later would bias the estimation because the consensus forecast in the estimation would then not be the most recent consensus forecast available to individual forecasters.

the latter would be a revision over two quarters and would, hence, include old news in addition to current news.

Data for (5.18) are in the Livingston survey only available for $h=2$ and only from 1992 onward.⁴⁵ In order to be able to make a fair comparison between estimations based on individual forecasts of the Livingston survey with those based on individual forecasts of the SPF, we, therefore, first estimate (5.17) with SPF individual data for a similar sample. That is, we keep only observations from the second and fourth quarter and limit the sample to post-1992. Again, we focus on nominal GDP, real GDP and the unemployment rate, as forecasts about the GDP price index are not available.

Table 5.8: Estimates of behavioral parameters for individual expectations in SPF and Livingston survey.

h=2 (Q2 & Q4)	λ		δ		γ	
	SPF	Livingston	SPF	Livingston	SPF	Livingston
NGDP	0.50	0.32	0.47	0.50	0.70	0.69
RGDP	0.53	0.31	0.54	0.75	0.58	0.74
Unemployment	0.52	0.31	0.67	0.83	0.69	0.79
Average	0.52	0.31	0.56	0.69	0.66	0.74

In Table 5.8, we compare the resulting estimates of λ , δ , and γ . Here, it can be seen that the estimates of δ and γ , if anything, are even somewhat larger in the Livingston survey than in the SPF. That is, also in the Livingston survey, forecasters considerably use the lagged consensus forecast and extrapolate current news.

If we turn to the estimates of λ , a different picture arises. For all three variables, the extent to which forecasters in the Livingston survey anchor to their own lagged forecast seems to be around 40% smaller than for the SPF forecasters. A reasonable explanation for this could lie in the fact that participants in the Livingston survey made their previous forecast two quarters ago rather than one quarter ago. This, first of all, makes their previous forecasts more outdated and less relevant, and forecasters may realize this. Secondly, if stickiness in expectations is partly caused by a confirmation bias (Nickerson, 1998), then this bias is likely to be weaker when the previous forecast was formed a longer time ago. This is because a forecaster may be more inclined to defend, and stick with, a previously formed opinion or forecast (i.e., may have a stronger confirmation bias) if she formed this opinion/forecast in the more recent past.

⁴⁵To be able to compute the lagged expected yearly growth rate for $h = 2$ implied by the level forecast of GDP, we need individual nowcasts. See also Appendix 5.A. These nowcasts are available from 1992 only.

Along these lines, Zhu et al. (2012) find, for example, that subjects in their laboratory experiment were more likely to revise their opinion in response to new information when more time had passed between the moment of the possible revision and the moment that the original opinion was formed.

5.5.3 An alternative deviation from rationality

As a final extension, we consider adding an additional bias to our model of expectation formation. In addition to stickiness, extrapolation and the dependence of expectations on lagged consensus forecasts, one could imagine that forecasters also partly base their expectations on the most recent observation of the variable being forecast.

Such behavior in expectation formation is found in laboratory experiments in macroeconomic settings by, e.g., Pfajfar and Žakelj (2018) and Assenza et al. (2019). Moreover, when forecasters were using adaptive learning as in Evans and Honkapohja (2012), one would also expect to find that forecasts depend on, a.o., the lag of the variable being forecast.⁴⁶

Furthermore, these different kinds of ‘backward-looking’ expectation formation processes can make expectations more persistent. Therefore, macroeconomic models that include some form of backward-looking expectation formation can improve the fit to the data compared to a specification with fully rational expectation (Branch, 2004; Milani, 2007; Slobodyan and Wouters, 2012; Cornea-Madeira et al., 2019). Based on these findings, it might be the case that basing forecasts on lagged observations is also an important deviation from full rationality that can explain forecasting behavior in survey data.

We, therefore, could consider a possible extension of the model, where forecasters may also partly base their forecast on the most recent observation of the variable that is being forecast. Modeling this in a similar way as the use of the lagged consensus forecast, (5.16) would become

$$\begin{aligned}
 F_{i,t}x_{t+h} - F_{i,t-1}x_{t+h} = & (1 - \lambda)[(1 - \delta_1 - \delta_2)(x_{t+h} - F_{i,t-1}x_{t+h}) + \delta_1(C_{t-1}x_{t+h} - F_{i,t-1}x_{t+h}) \\
 & + \delta_2(x_{t-1} - F_{i,t-1}x_{t+h})] + \gamma News_{t,t} \\
 & - (1 - \lambda)(1 - \delta)\varepsilon_{t,t+h} + \mu_{i,t},
 \end{aligned} \tag{5.19}$$

where we now denote the original δ as δ_1 , and where δ_2 is the extent to which fore-

⁴⁶Additionally, forecasts might then also depend on observations further in the past than the most recent lag and on lags of other variables. We have also checked for this. When adding other lags to our regressions, estimated coefficients on these lags are always very small and in most cases statistically insignificant.

casters form their forecast based on the most recent observation of the variable that is being forecast. If forecasters base their decisions on this variable in addition to the components of our proposed model of expectation formation, this should be reflected in the coefficient on the additional regressor, a_4 , in the following extension of (5.17):

$$F_{i,t}x_{t+h} - F_{i,t-1}x_{t+h} = a_1(x_{t+h} - F_{i,t-1}x_{t+h}) + a_2(C_{t-1}x_{t+h} - F_{i,t-1}x_{t+h}) + a_3News_{t,t} + a_4(x_{t-1} - F_{i,t-1}x_{t+h}) + e_{i,t}. \quad (5.20)$$

Table 5.9: Estimates of a_4 (the coefficient on $x_{t-1} - F_{i,t-1}x_{t+h}$) in (5.20) for different forecasting horizons.

	$h = 1$	$h = 2$	$h = 3$
Inflation	0.011 (0.508)	-0.016 (0.364)	0.011 (0.617)
NGDP	0.007 (0.501)	0.002 (0.905)	0.016 (0.251)
RGDP	0.020 (0.174)	0.028 (0.112)	0.014 (0.312)
Unemployment	-0.019 (0.597)	0.050 (0.229)	0.066 (0.032)

Note: P-values are in brackets.

In Table 5.9 we present the estimates of a_4 , together with the corresponding P-value in brackets. From the latter, it can be seen that for 11 of the 12 cases, the estimated coefficient is not statistically significant at the 5%-level (or even at the 10%-level). Moreover, the point estimates of a_4 are always small in size and even take on negative values in two cases. This confirms the findings of Fuhrer (2018) that stickiness in expectations (or expectations smoothing) can better explain the persistence found in survey expectations than backward-looking adaptive expectations or adaptive learning. Landier et al. (2019) draw similar conclusions from the expectations data of their laboratory experiment. Based on the above exercise, we see no convincing reason to extend our baseline model of expectation formation with an additional (backward-looking) bias.

5.6 Discussion: the (ir)rationality of using the lagged consensus forecast

In the previous sections, we found that individual forecasters, to a large extent, are not able to build their own rational expectations and, instead, partly base their forecasts on the lagged consensus forecast. This raises the question of how useful the lagged consensus forecast is as a source of information for individual forecasters.

Of course, the lagged consensus forecast might contain some important new information that agents should rationally take into account. Given that not all agents in the economy have rational expectations, expectations of other economic participants are an important determinant of what will happen to economic variables.

However, this does not explain the positive (and relatively large) values of δ that we have found in our estimations. If an individual forecaster was building rational expectations based on all information, including the lagged consensus, then her value of δ would be zero. After all, the rational expectation forecast efficiently incorporates all information, including the lagged consensus forecast and the lagged consensus would not show up as an additional determinant of the individual forecast. A positive value of δ , hence, means that *instead* of forming rational expectations, a forecaster is partly basing her forecast on what the most recent consensus forecast is.

To get some insight into the extent to which this is a reasonable thing to do, we compare mean squared forecast error (MSFE) of the individual forecasters in the surveys with the corresponding MSFE of the *lagged* consensus forecast. Thus, for each individual i we take the MSFE, $\text{Mean}(F_{i,t}x_{t+h} - x_{t+h})^2$, over all periods (t) for which she was in the panel. We then compare this with the MSFE of the *lagged* consensus forecast, $\text{Mean}(C_{t-1}x_{t+h} - x_{t+h})^2$, about *the same* periods. This gives an indication of how well the individual performed relative to the lagged consensus forecast.

In particular, when the MSFE of the individual is larger than that of the lagged consensus forecast, then the individual was outperformed by the lagged consensus. In the top rows of the different panels of Table 5.10 we summarize, for different forecasting horizons, for how many forecasters in the sample this was the case.

By looking at the table as a whole, it immediately stands out that the percentage of individuals that was outperformed by the lagged consensus forecast lies between 18% and 50%. This is strikingly high, as the individuals can observe this lagged consensus forecast when making their forecast. Hence, they obviously have the opportunity to use this forecast when making their own forecast. Moreover, in the previous sections, we found evidence that they indeed, for a considerable part, do base their forecasts on the

Table 5.10: Percentage of forecasters outperformed by lagged consensus forecast

	Infl.	Nom. GDP	Real GDP	Unemp.
h=1 (SPF)				
Outperformed by $C_{t-1}x_{t+h}$	30%	29%	29%	21%
Mean $(F_{i,t}x_{t+h} - x_{t+h})^2$	0.6	1.4	1.29	0.18
h=2 (SPF)				
Outperformed by $C_{t-1}x_{t+h}$	34%	37%	30%	25%
Mean $(F_{i,t}x_{t+h} - x_{t+h})^2$	1.23	2.70	2.58	0.39
h=3 (SPF)				
Outperformed by $C_{t-1}x_{t+h}$	44%	50%'	43%'	33%
Mean $(F_{i,t}x_{t+h} - x_{t+h})^2$	2.11	4.4	4.2	0.68
h=2 (LIV)				
Outperformed by $C_{t-1}x_{t+h}$		50%	35%	18%
Mean $(F_{i,t}x_{t+h} - x_{t+h})^2$		1.76	1.18	0.38

Note: A ' indicates that individual forecasts are on average (across individuals and time) outperformed by the lagged consensus forecast.

past consensus forecast. However, many forecasters are apparently still outperformed by this forecast.

The strong performance of the lagged consensus forecast relative to individual forecasts first of all confirms that using the consensus forecast as an alternative to rational expectations is a good choice for an individual that is not able to fully build a rational expectations forecast all by herself. The cognitive bias of Section 5.2.3 should, therefore, really be interpreted in this way, and not as irrational overweighting of the past consensus forecast.

This becomes even more clear when we consider the following. For many individuals, the lagged consensus forecast has a better forecasting performance than the individual forecasts that are made one period later. If such an “outperformed” forecaster would have decided to abandon all attempts to build her own forecast and instead always blindly submitted the observed past consensus forecast (for example from the SPF website), then her forecasting performance would on average have been *better!*

Considering the top three panels of Table 5.10 in some more detail, it can be seen that the larger the forecasting horizon, the more individuals of the SPF would have been better off by always blindly copying the lagged consensus forecast. This indicates that the more difficult the forecasting task, the larger individual errors might become

and the worse the performance of individuals relative to the lagged consensus forecast. The bottom rows of the first three panels confirm that the MSFE across individuals *and* time periods indeed increases considerably as the forecasting horizon increases. This holds for all variables. While the MSFE of the consensus forecast also increases as the horizons increases, it does so to a lesser extent, so that more individuals are outperformed by the lagged consensus forecast for larger forecasting horizons.

Moreover, for nominal GDP and real GDP, for a forecasting horizon of three periods (marked with a '), we even find that, when we compare the MSFE of *all* individuals and *all* time periods with the MSFE of the corresponding lagged consensus forecasts, the individual forecasts are, on average, outperformed by the lagged consensus forecasts.

Finally, turn to the bottom panel of Table 5.10 that displays the results for the Livingston survey for a forecasting horizon of $h = 2$. The MSFE in the Livingston survey (bottom row) are smaller than those of the SPF (bottom row of the second panel). This is because the sample starts in 1992 where all variables were generally more stable and better forecastable. However, even in these more stable times, even more forecasters of the Livingston survey would have been better off by always copying the lagged (SPF) consensus nominal and real GDP forecasts than the forecasters in the SPF for the same forecasting horizon.

5.7 Conclusion

We have constructed a theoretical model of expectation formation and confronted it with survey expectations. We consider forecasting data of four important macroeconomics variables: inflation, nominal GDP, real GDP and the unemployment rate. We confirm that forecasters in the Survey of Professional Forecasters and the Livingston survey have stickiness in their expectations in the sense that their forecasts are biased toward their own past forecast. Moreover, we find that forecasters overly extrapolate news (surprises) about the current period into the future, which further biases their forecasts.

Also, we find that forecasters are not able to come up with their own rational forecasts but, instead, partly rely on the most recent consensus forecast about the period that they are forecasting. They can observe this consensus of economic experts from the website of their survey, but could also get it from other sources. In particular, we have shown that they are not basing their forecast on the survey that they are participating in per se, but that the past consensus forecast rather is a more general public signal that can also be obtained from, e.g., a different survey.

According to our proposed model of expectation formation, the sum of the coefficients in front of an individual's lagged own forecast, the lagged consensus forecast and the lagged realization of the variable being forecast should be equal to 1. When we estimate an over-identified specification of our model, the sum of these coefficients always lies very close to 1. Moreover, the null hypothesis that the sum is equal to 1 can, in most cases, not be rejected. This gives support to the validity of the proposed model.

We also obtained estimates of the behavioral parameters in our model with a revision regression that is not over-identified. For the SPF, we find stickiness estimates between 0.41 and 0.76, depending on the forecasting horizon. Larger forecasting horizons seem to imply more stickiness in expectations. Meanwhile, estimates of the extrapolation parameter of current news lie between 0.18 and 0.91, with larger forecasting horizons implying smaller estimates for most variables.

Regarding the lagged consensus forecast, we find that forecasters put weights between 0.23 and 0.83 on the most recently observed consensus forecast rather than on a rational expectations forecast that they would have to build themselves. It, therefore, seems that individual forecasters have a hard time processing all available information and news and building a rational expectations forecast. We find, however, that using the lagged consensus forecast to form expectations then is a reasonable thing to do. In fact, depending on the forecasting horizon and the variable being forecast, up to 50% of the forecasters in the survey would have been better off in terms of forecasting performance if they would have abandoned all attempts to build their own forecast and instead blindly submitted the most recently observed consensus forecast in every period.

In the individual forecasts of the Livingston survey, all three biases are also present. However, the estimates of the stickiness parameter lie here between 0.30 and 0.32 which is considerably smaller than the stickiness found in the SPF for the same forecasting horizon. This can be explained by the fact that the survey is semi-annual so that the previous forecast that a forecaster made was half a year ago rather than one quarter ago. This could make this previous forecast more outdated and could reduce the magnitude of a confirmation bias.

Finally, we find that lagged realizations of variables being forecast have little to no explanatory power and are often statistically insignificant when added to our model. This indicates that the biases in our proposed model can better explain individual survey expectations than alternative, backward-looking biases.

Appendix 5.A Variable Construction

For the price deflator, nominal GDP and real GDP, the data is about levels (indexes), and we turn these into growth rates. To construct yearly growth rates of actual realizations, we take the first vintages of the level and divide it by the corresponding lag (four quarters earlier) of the *same vintage*. For individual forecasts, we divide the forecast by the actual realization four quarters before the period that is being forecast. Here, we use the most up-to-date vintage at the time the forecast was made, to be as close as possible to the information sets of forecasters. For four-quarter-ahead forecasts, the actual realization of the corresponding lag is not available yet at the time the forecast was made. For that case, we divide the individual's forecast by her nowcast. Below, we give detailed formulas for all cases.

5.A.1 SPF

- Constructing growth rates x for $y \in \{\text{Price Deflator, Nominal GDP, Real GDP}\}$, where $y_{a|b}$ denotes vintage b data about the period a realization of y .

- actual realizations in $t + h$ for $h = -1, 1, 2, 3$: $x_{t+h} = \frac{y_{t+h|t+h+1}}{y_{t+h-4|t+h+1}} - 1$
- individual forecast in t for $h = 1, 2, 3$: $F_{i,t}x_{t+h} = \frac{F_{i,t}y_{t+h}}{y_{t+h-4|t}} - 1$
- individual forecast in $t - 1$ for $h = 1, 2$: $F_{i,t-1}x_{t+h} = \frac{F_{i,t-1}y_{t+h}}{y_{t+h-4|t-1}} - 1$
- individual forecast in $t - 1$ for $h = 3$: $F_{i,t-1}x_{t+3} = \frac{F_{i,t-1}y_{t+3}}{F_{i,t-1}y_{t-1}} - 1$
- consensus forecast in $t - 1$ for $h = 1, 2$: $C_{t-1}x_{t+h} = \frac{C_{t-1}y_{t+h}}{y_{t+h-4|t-1}} - 1$
- consensus forecast in $t - 1$ for $h = 3$: $C_{t-1}x_{t+3} = \frac{C_{t-1}y_{t+3}}{C_{t-1}y_{t-1}} - 1$
- consensus nowcast: $C_t x_t = \frac{C_t y_t}{y_{t-4|t}} - 1$
- consensus forecast in $t - 1$ about t : $C_{t-1} x_t = \frac{C_{t-1} y_t}{y_{t-4|t-1}} - 1$
- consensus forecast in $t - 2$ about t (used for Real GDP only): $C_{t-2} x_t = \frac{C_{t-2} y_t}{y_{t-4|t-2}} - 1$

- Unemployment

- actuals in $t + h$ for $h = -1, 1, 2, 3$: $x_{t+h|t+h+1}$
- for $h = 1, 2, 3$: $F_{i,t}x_{t+h}, F_{i,t-1}x_{t+h}, C_{t-1}x_{t+h}$
- consensus nowcast and forecasts about t : $C_t x_t, C_{t-1} x_t, C_{t-2} x_t$

5.A.2 Livingston

- Constructing growth rates x for $y \in \{\text{Nominal GDP, Real GDP}\}$
 - actuals: $x_{t+2} = \frac{y_{t+2|t+3}}{y_{t-2|t+3}} - 1$
 - individual forecast in t : $F_{i,t}x_{t+2} = \frac{F_{i,t}y_{t+2}}{y_{t-2|t}} - 1$
 - individual forecast in $t - 2$: $F_{i,t-2}x_{t+2} = \frac{F_{i,t-2}y_{t+2}}{F_{i,t-2}y_{t-2}} - 1$
 - consensus forecasts in $t - 1$ used in regressions with individual SPF forecasts:
 $C_{t-1}^{LIV} x_{t+1} = \frac{C_{t-1}^{LIV} y_{t+1}}{y_{t-3|t-1}} - 1, C_{t-1}^{LIV} x_{t+3} = \frac{C_{t-1}^{LIV} y_{t+3}}{C_{t-1}^{LIV} y_{t-1}} - 1$
- Unemployment
 - actuals: $x_{t+2|t+3}$
 - individual forecasts: $F_{i,t}x_{t+2}, F_{i,t-1}x_{t+2}$
 - consensus forecasts used in regressions with individual SPF forecasts: $C_{t-1}^{LIV} x_{t+1}, C_{t-1}^{LIV} x_{t+3}$

Appendix 5.B News and Consensus Nowcast Regressions

In order to approximate $News_{t,t}$, we estimate (5.10), the results of which are given in Table 5.11. For inflation and nominal GDP, the estimation shows that the consensus nowcast is unbiased. For real GDP and unemployment there is a bias. The magnitude of this bias is, however, relatively small compared to the estimates of Coibion and Gorodnichenko (2015) for the three quarter ahead consensus forecast which range from 1.02 to 1.23 for inflation. Nonetheless, we correct for the bias in the consensus nowcast revisions of real GDP and unemployment by using equation (5.12), with ψ taken from Table 5.11.

Appendix 5.C Robustness checks

5.C.1 Final vintages

In this appendix, we consider the robustness of our results if the actual realization of the respective variable in (5.17) is based on final vintages instead of first vintages. Correspondingly, we also take final vintages of the lagged actual realizations that are used as instruments in the IV-GMM estimations. Even though most of the literature

Table 5.11: SPF Consensus Nowcast Regressions
 Dependent Var.: $x_t - C_t x_t$

	Inflation	Nominal GDP	Real GDP	Unemp. Rate
constant	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	-0.034*** (0.009)
$C_t x_t - C_{t-1} x_t$	0.056 (0.098)	0.030 (0.083)	0.165** (0.074)	0.23*** (0.031)
R^2	0.006	0.002	0.046	0.17
n	204	204	204	205

Note: OLS estimation of equation (5.10) including a constant. Newey-West standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Estimates are rounded to the third decimal.

so far uses first vintages, it can be argued that the econometrician should use *final* vintages of the *actual* data that agents are trying to forecast (Angeletos et al., 2020).

Table 5.12: Estimates of behavioral parameters for different forecasting horizons (final vintage actual).

SPF	λ			δ			γ		
	h=1	h=2	h=3	h=1	h=2	h=3	h=1	h=2	h=3
Inflation	0.46	0.52	0.73	0.89	0.79	0.36	0.98	0.85	0.25
NGDP	0.44	0.55	0.73	0.85	0.65	0.21	0.93	0.86	0.36
RGDP	0.47	0.58	0.75	0.96	0.54	0.48	0.85	0.68	0.20
Unemployment	0.41	0.54	0.61	0.67	0.63	0.63	0.62	0.62	0.62
Average	0.45	0.55	0.71	0.84	0.65	0.42	0.85	0.75	0.36

Table 5.12 provides the resulting estimates of λ , δ and γ . Comparing the results in Table 5.12 with the results in Table 5.4, it can be seen that the estimates of λ are very much the same. The estimates for δ and γ are quantitatively somewhat different, but no general pattern arises that holds for all variables and forecasting horizons. If anything, estimates of γ are somewhat larger when final vintages are used. In any case, the qualitative differences between forecasting horizons of Table 5.4 remain when we use final vintages.

5.C.2 Time subsamples

In this appendix, we consider robustness with respect to different time subsamples. In particular, we are interested in answering the question of whether behavioral biases in expectation formation are different in a more stable economic environment. Therefore, we investigate whether the estimates of our model parameters are different if we limit the sample to the post-1984 (exclusion of pre - Great Moderation period) and, secondly, to 1984-2006 (Great Moderation).

Tables 5.13 and 5.14 present the estimates of λ , δ and γ when we limit the sample to start at 1984Q1 and to the Great Moderation, respectively.

Table 5.13: Estimates of behavioral parameters for different forecasting horizons (post-1984).

SPF	λ			δ			γ		
	h=1	h=2	h=3	h=1	h=2	h=3	h=1	h=2	h=3
Inflation	0.40	0.50	0.71	0.86	0.67	0.49	0.92	0.84	0.28
NGDP	0.40	0.48	0.68	0.77	0.37	0.29	0.95	0.69	0.32
RGDP	0.43	0.53	0.70	0.64	0.32	0.31	0.72	0.53	0.18
Unemployment	0.40	0.51	0.60	0.62	0.65	0.66	0.59	0.65	0.65
Average	0.41	0.51	0.67	0.72	0.50	0.44	0.80	0.68	0.36

Table 5.14: Estimates of behavioral parameters for different forecasting horizons (Great Moderation, 1984-2006).

SPF	λ			δ			γ		
	h=1	h=2	h=3	h=1	h=2	h=3	h=1	h=2	h=3
Inflation	0.36	0.49	0.67	0.85	0.72	0.57	0.93	0.87	0.31
NGDP	0.42	0.48	0.64	0.84	0.69	0.36	0.98	0.77	0.19
RGDP	0.41	0.54	0.69	0.81	0.63	0.60	0.78	0.59	0.17
Unemployment	0.42	0.52	0.59	0.70	0.71	0.74	0.62	0.65	0.64
Average	0.40	0.51	0.65	0.80	0.69	0.57	0.83	0.72	0.33

In the tables, it can be seen that the individual estimates quantitatively differ from those of the full sample in Table 5.4. Generally, the estimates of λ tend to be a little bit lower in the more stable time-subsamples. However the estimates of the three model parameters remain of a similar order of magnitude, and the dependence of λ and γ on the forecasting horizon that was discussed in Section 5.4.3 remains.

Appendix 5.D Additional Tables

Table 5.15: Regression of Equation (5.17) with $h = 2$ based on SPF
 Dependent Var.: $F_{i,t}x_{t+2} - F_{i,t-1}x_{t+2}$

	Inflation	Nominal GDP	Real GDP	Unemp. Rate
$x_{t+2} - F_{i,t-1}x_{t+2}$	0.137*** (0.021)	0.276*** (0.065)	0.270*** (0.062)	0.204*** (0.048)
$C_{t-1}x_{t+2} - F_{i,t-1}x_{t+2}$	0.342*** (0.035)	0.181*** (0.069)	0.151*** (0.064)	0.258*** (0.049)
$News_{t,t}$	0.755*** (0.051)	0.726*** (0.083)	0.543*** (0.077)	0.561*** (0.082)
R^2	0.57	0.56	0.49	0.59
n	5545	5560	5566	5869

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Estimates are rounded to the third digit after the comma.

Table 5.16: Regression of Equation (5.17) with $h = 3$ based on SPF
 Dependent Var.: $F_{i,t}x_{t+3} - F_{i,t-1}x_{t+3}$

	Inflation	Nominal GDP	Real GDP	Unemp. Rate
$x_{t+3} - F_{i,t-1}x_{t+3}$	0.175*** (0.019)	0.211*** (0.043)	0.133*** (0.033)	0.126*** (0.032)
$C_{t-1}x_{t+3} - F_{i,t-1}x_{t+3}$	0.070*** (0.031)	0.064*** (0.047)	0.126*** (0.041)	0.274*** (0.034)
$News_{t,t}$	0.223*** (0.045)	0.307*** (0.073)	0.176*** (0.046)	0.643*** (0.062)
R^2	0.28	0.24	0.23	0.52
n	5301	5327	5268	5558

Note: *** p<0.01, ** p<0.05, * p<0.1. Estimates are rounded to the third digit after the comma.

Bibliography

- Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2020). Imperfect macroeconomic expectations: Evidence and theory. NBER Working Paper, No. 27308.
- Annicchiarico, B., Surrinchio, S., and Waldmann, R. J. (2019). A behavioral model of the credit cycle. Journal of Economic Behavior & Organization, 166:53–83.
- Assenza, T., Heemeijer, P., Hommes, C., and Massaro, D. (2019). Managing self-organization of expectations through monetary policy: a macro experiment. Journal of Monetary Economics, forthcoming.
- Beqiraj, E., Di Bartolomeo, G., and Serpieri, C. (2019). Rational vs. long-run forecasters: Optimal monetary policy and the role of inequality. Macroeconomic Dynamics, 23(1):9–24.
- Blake, A. P., Beltramo, J., and Paul, J. (2001). A timeless perspective on optimality in forward-looking rational expectations models. NIESR Discussion Papers, 188.
- Blanchard, O. and Galí, J. (2010). Labor markets and monetary policy: A new keynesian model with unemployment. American economic journal: macroeconomics, 2(2):1–30.
- Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2018a). Over-reaction in macroeconomic expectations. NBER Working Paper, No. 24932.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2018b). Diagnostic expectations and credit cycles. The Journal of Finance, 73(1):199–227.
- Bouchaud, J.-P., Krueger, P., Landier, A., and Thesmar, D. (2019). Sticky expectations and the profitability anomaly. The Journal of Finance, 74(2):639–674.
- 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. (2007). Sticky information and model uncertainty in survey data on inflation expectations. Journal of Economic Dynamics and Control, 31(1):245–276.
- Branch, W. A. and McGough, B. (2009). A new keynesian model with heterogeneous expectations. Journal of Economic Dynamics and Control, 33(5):1036 – 1051.
- Branch, W. A. and McGough, B. (2010). Dynamic predictor selection in a new keynesian model with heterogeneous expectations. Journal of Economic Dynamics and Control, 34(8):1492 – 1508.
- Branch, W. A. and McGough, B. (2018). Heterogeneous expectations and micro-foundations in macroeconomics. Handbook of computational economics, 4:3–62.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. Journal of monetary Economics, 12(3):383–398.
- Carlson, J. A. (1977). A study of price forecasts. Annals of Economic and Social Measurement, 6:27–56.
- Christiano, L. J., Eichenbaum, M. S., and Trabandt, M. (2016). Unemployment and business cycles. Econometrica, 84(4):1523–1569.
- Clarida, R., Gali, J., and Gertler, M. (1998). Monetary policy rules in practice: Some international evidence. European Economic Review, 42(6):1033–1067.
- Clarida, R., Gali, J., and Gertler, M. (1999). The Science of Monetary Policy: A New Keynesian Perspective. Journal of Economic Literature, 37(4):1661–1707.
- Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review, 105(8):2644–78.
- Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2012). Innocent bystanders? monetary policy and inequality in the us. National Bureau of Economic Research.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2019). Monetary policy communications and their effects on household inflation expectations. NBER Working Paper, No. 25482.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). The cost of the covid-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. NBER Working Paper, No. 27141.

- Cole, S. J. and Milani, F. (2019). The misspecification of expectations in new keynesian models: a dsge-var approach. Macroeconomic Dynamics, 23(3):974–1007.
- Cornea-Madeira, A., Hommes, C., and Massaro, D. (2019). Behavioral heterogeneity in us inflation dynamics. Journal of Business & Economic Statistics, 37(2):288–300.
- Croushore, D., Stark, T., et al. (2019). Fifty years of the survey of professional forecasters. Economic Insights, 4(4):1–11.
- Davis-Stober, C. P., Budescu, D. V., Dana, J., and Broomell, S. B. (2014). When is a crowd wise? Decision, 1(2):79.
- De Grauwe, P. (2011). Animal spirits and monetary policy. Economic theory, 47(2-3):423–457.
- De Grauwe, P. (2012). Lectures on behavioral macroeconomics. Princeton University Press.
- Debortoli, D. and Galí, J. (2017). Monetary policy with heterogeneous agents: Insights from tank models. mimeo.
- Di Bartolomeo, G., Di Pietro, M., and Giannini, B. (2016). Optimal monetary policy in a new keynesian model with heterogeneous expectations. Journal of Economic Dynamics and Control, 73:373 – 387.
- Ehrbeck, T. and Waldmann, R. (1996). Why are professional forecasters biased? agency versus behavioral explanations. The Quarterly Journal of Economics, 111(1):21–40.
- Evans, G. W. and Honkapohja, S. (2006). Monetary policy, expectations and commitment. The Scandinavian Journal of Economics, 108(1):15–38.
- Evans, G. W. and Honkapohja, S. (2012). Learning and expectations in macroeconomics. Princeton University Press.
- Faia, E. (2008). Optimal monetary policy rules with labor market frictions. Journal of Economic dynamics and control, 32(5):1600–1621.
- Fildes, R. and Stekler, H. (2002). The state of macroeconomic forecasting. Journal of macroeconomics, 24(4):435–468.
- Fisher, I. (1933). The debt-deflation theory of great depressions. Econometrica, 1(4):337–357.

- Fuhrer, J. (2017). Expectations as a source of macroeconomic persistence: Evidence from survey expectations in a dynamic macro model. Journal of Monetary Economics, 86:22–35.
- Fuhrer, J. C. (2018). Intrinsic expectations persistence: evidence from professional and household survey expectations. FRB of Boston Working Paper, No. 18-9.
- Gabaix, X. (2020). A behavioral new keynesian model. American Economic Review, 110(8):2271–2327.
- Galí, J. (2015). Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications. Princeton University Press.
- Galí, J. and Gertler, M. (1999). Inflation dynamics: A structural econometric analysis. Journal of monetary Economics, 44(2):195–222.
- Gasteiger, E. (2014). Heterogeneous expectations, optimal monetary policy, and the merit of policy inertia. Journal of Money, Credit and Banking, 46(7):1535–1554.
- Gasteiger, E. (2018). Optimal constrained interest-rate rules under heterogeneous expectations. mimeo.
- Gertler, M., Sala, L., and Trigari, A. (2008). An estimated monetary dsge model with unemployment and staggered nominal wage bargaining. Journal of Money, Credit and Banking, 40(8):1713–1764.
- Gertler, M. and Trigari, A. (2009). Unemployment fluctuations with staggered nash wage bargaining. Journal of political Economy, 117(1):38–86.
- Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. American economic review, 95(1):50–65.
- Hommes, C. (2011). The heterogeneous expectations hypothesis: Some evidence from the lab. Journal of Economic dynamics and control, 35(1):1–24.
- Hommes, C. and Lustenhouwer, J. (2019). Inflation targeting and liquidity traps under endogenous credibility. Journal of Monetary Economics, 107:48–62.
- Honkapohja, S., Mitra, K., Evans, G. W., et al. (2011). Notes on agents’ behavioral rules under adaptive learning and studies of monetary policy. Macroeconomics at the Service of Public Policy, pages 63–79.

- Kahneman, D. and Tversky, A. (1973). On the psychology of prediction. Psychological Review, 80(4):237.
- Krause, M. U. and Lubik, T. A. (2007). The (ir)relevance of real wage rigidity in the new keynesian model with search frictions. Journal of Monetary Economics, 54(3):706–727.
- Kryvtsov, O. and Petersen, L. (2019). Expectations and monetary policy: experimental evidence. mimeo.
- Kurz, M., Piccillo, G., and Wu, H. (2013). Modeling diverse expectations in an aggregated new keynesian model. Journal of Economic Dynamics and Control, 37(8):1403–1433.
- Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. Econometrica, 50(6):1345–1370.
- Landier, A., Ma, Y., and Thesmar, D. (2019). Biases in expectations: Experimental evidence. mimeo.
- Loungani, P. (2001). How accurate are private sector forecasts? cross-country evidence from consensus forecasts of output growth. International journal of forecasting, 17(3):419–432.
- Lovell, M. C. (1986). Tests of the rational expectations hypothesis. The American Economic Review, 76(1):110–124.
- Ma, Y., Ropele, T., Sraer, D., and Thesmar, D. (2018). Do managerial forecasting biases matter. mimeo.
- Mackowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. American Economic Review, 99(3):769–803.
- 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., Reis, R., and Wolfers, J. (2003). Disagreement about inflation expectations. NBER macroeconomics annual, 18:209–248.
- Massaro, D. (2013). Heterogeneous expectations in monetary DSGE models. Journal of Economic Dynamics and Control, 37(3):680–692.

- McCallum, B. T. (1999). Issues in the design of monetary policy rules. Handbook of macroeconomics, 1:1483–1530.
- Merz, M. (1995). Search in the labor market and the real business cycle. Journal of monetary Economics, 36(2):269–300.
- Milani, F. (2007). Expectations, learning and macroeconomic persistence. Journal of monetary Economics, 54(7):2065–2082.
- Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. The review of economic studies, 61(3):397–415.
- Muth, J. F. (1961). Rational expectations and the theory of price movements. Econometrica, 29(3):315–335.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. Review of general psychology, 2(2):175–220.
- Pfajfar, D. and Santoro, E. (2010). Heterogeneity, learning and information stickiness in inflation expectations. Journal of Economic Behavior & Organization, 75(3):426–444.
- Pfajfar, D. and Žakelj, B. (2018). Inflation expectations and monetary policy design: Evidence from the laboratory. Macroeconomic Dynamics, 22(4):1035–1075.
- Preston, B. (2005). Learning about monetary policy rules when long-horizon expectations matter. International Journal of Central Banking, 1(2):81–126.
- Ravenna, F. and Walsh, C. E. (2011). Welfare-based optimal monetary policy with unemployment and sticky prices: a linear-quadratic framework. American Economic Journal: Macroeconomics, 3(2):130–162.
- Rotemberg, J. J. (1982). Monopolistic price adjustment and aggregate output. The Review of Economic Studies, 49(4):517–531.
- Rotemberg, J. J. and Woodford, M. (1997). An optimization-based econometric framework for the evaluation of monetary policy. NBER macroeconomics annual, 12:297–346.
- Schmitt-Grohé, S. and Uribe, M. (2003). Closing small open economy models. Journal of international Economics, 61(1):163–185.
- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. American economic review, 95(1):25–49.

- Sims, C. A. (2003). Implications of rational inattention. Journal of Monetary Economics, 50(3):665–690.
- Slobodyan, S. and Wouters, R. (2012). Learning in a medium-scale dsge model with expectations based on small forecasting models. American Economic Journal: Macroeconomics, 4(2):65–101.
- Surowiecki, J. (2004). The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations. Little, Brown, London.
- Thomas, C. (2008). Search and matching frictions and optimal monetary policy. Journal of Monetary Economics, 55(5):936–956.
- Trigari, A. (2009). Equilibrium unemployment, job flows, and inflation dynamics. Journal of money, credit and banking, 41(1):1–33.
- Woodford, M. (1999). Optimal monetary policy inertia. The Manchester School, 67:1–35.
- Woodford, M. (2001). Imperfect common knowledge and the effects of monetary policy. NBER Working Paper, No. 8673.
- Woodford, M. (2003). Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press.
- Woodford, M. (2010). Optimal monetary stabilization policy. Handbook of Monetary Economics, 3:723–828.
- Zhu, H., Huberman, B., and Luon, Y. (2012). To switch or not to switch: understanding social influence in online choices. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pages 2257–2266.