Utrecht University

School of Economics

The roles of uncertainty and beliefs in the economy

Thomas Gomez

U.S.E. 066

The roles of uncertainty and beliefs in the economy

Thomas Gomez

Manuscript committee:	Prof. dr. C.H. Hommes
	Prof. dr. W. Kets
	Prof. dr. M. Motolese
	Prof. dr. A.M. Salomons
	Dr. S. Dirksen

This work was supported by the Complex Systems Fund, with special thanks to Peter Koeze.

ISBN: 978-94-91870-52-1
U.S.E. Dissertation Series
USE 066
Printed by Ridderprint

The roles of uncertainty and beliefs in the economy

De rollen van onzekerheid en verwachtingen in de economie
(met een samenvatting in het Nederlands)

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit Utrecht
op gezag van de rector magnificus,
prof. dr. H.R.B.M. Kummeling,
ingevolge het besluit van het college voor promoties
in het openbaar te verdedigen op
vrijdag 2 december 2022 des middags te 2.15 uur

door

Thomas Gomez

geboren op 12 november 1991
te Alkmaar

Promotoren:	Prof. dr. G.U. Weitzel Prof. dr. J.E. Frank
Copromotoren:	Dr. G. Piccillo Dr. S. Lugo

To my mom and my wife

Acknowledgments

First and foremost, I would like to thank my supervisors for their support and guidance. Thank you, Giulia, for being my mentor. You helped shape my master thesis, PhD proposal, and dissertation. Utz, thank you for pointing me in the right direction when I was looking for the topic for my master thesis and PhD proposal. I also appreciate your continued support and advice. Jason, thank you for your guidance and for checking in. The meetings we had were very valuable to me. Stefano, your comments and insights have been very helpful. Thank you for reinforcing the team.

My thanks go out to my office mates turned friends of room 2.11 who were always there for advice or a nice chat: Bora, Fujin, Lucía, Timo, and Vincent. I have fond memories of the parties and dinners we had. The online calls we had during lockdown were welcome distractions when working from home. I also enjoyed my time with the other PhDs at USE during lunches in the canteen and dinners.

I am grateful to Peter Koeze for his generosity which helped fund my research. His genuine interest in the project was stimulating. I admire your continued efforts to facilitate academic research.

I thank the Centre for Complex Systems Studies for providing a workspace and inspiring surroundings and events. Thanks to all its junior members for the interesting conversations and helpful suggestions. Special thanks to Adriana and Joris. I am happy to be part of the first-generation PhDs of the Complex Systems Fund.

I would also like to thank all colleagues at the Utrecht School of Economics and Mathematical Institute, including the support staff. Those I would discuss my research with or helped me with my teaching.

Many thanks to my mom, Judith, for all our discussions that sparked my scientific interests. And as you said yourself, my brain could not have done it without all your healthy, homecooked meals during my childhood. *Gracias por todo.*

Lastly, I would like to thank my wife, Saskia, for her unwavering love and support during the entire PhD. You have played a big part in finishing this dissertation, and I am very grateful for that.

Thomas Gomez

Contents

List of Tables

List of Figures

Chapter 1

Introduction

In this dissertation, I investigate the roles of uncertainty and beliefs in the economy.

There are various sources of uncertainty. At a fundamental level, we do not know the current state of the economy with certainty. Many macroeconomic indicators, like economic growth and the inflation rate, are revised after their first release. These revisions can be substantial, especially in turbulent times like the COVID-19 pandemic or the 2008 financial crisis. Even if we had perfect knowledge about the current state of the economy, we would not know with certainty what the economy would look like in the future because we do not know its laws of motion.

Uncertainty forces us to form beliefs about the current state and future path of the economy. Consider a scenario where someone has to decide whether to buy a house now or wait. Essential factors in her decision are the current state of the housing market and her future income. As both these factors involve uncertainty, she has to form beliefs about them to make a decision. For example, if she believes that an economic crisis is on the horizon, increasing the probability of losing her job, she might decide to postpone buying a house.

While uncertainty leads to belief formation, beliefs themselves also involve uncertainty. First, good forecasts consist of a range of possible outcomes and a probability distribution over those outcomes.[1] Second, we do not know the beliefs of all other economic agents, adding to the uncertainty about the economic system itself.

1.1 Economic impact

Given that uncertainty is ubiquitous in the economy, leading to belief formation, the question remains whether uncertainty and beliefs are important determinants of economic outcomes.

1.1.1 Uncertainty

The theoretical literature suggests that uncertainty leads to lower economic growth. One strand of the literature (e.g., Bernanke, 1983) focuses on a wait-and-see effect of

[1] Economic forecasts, especially as reported in the media, are often point-forecasts. One might predict year-on-year economic growth of 2%, for example, without specifying how confident one is about the forecast.

uncertainty. In uncertain times, firms prefer to postpone investments in capital and employees, and consumers tend to postpone purchases of consumer durables like cars. They prefer to wait until the uncertainty is resolved. Agents might also want to build a savings buffer to withstand any economic difficulties (Leland, 1968). Lower investment and consumption dampen economic growth.

It is challenging to find empirical evidence for these theoretical mechanisms because it is inherently difficult to measure uncertainty. Ideally, one would like to know each agent's subjective probability distribution over all possible economic outcomes (Born et al., 2018). Even if these distributions are known, it is not straightforward to aggregate them into a single measure of uncertainty. To make the problem even harder, one can distinguish between risk on one side and ambiguity (also called deep or Knightian uncertainty) on the other. While risk refers to situations with known outcomes and a known probability distribution over those outcomes, ambiguity refers to situations where the distribution or even the set of possible outcomes is unknown. Both types of uncertainty are important in theory, but especially the deep kind is tricky to operationalize.[2]

A conclusion that is reasonably robust to different operationalizations is that uncertainty rises during recessions (e.g., Bloom et al., 2018; Jurado et al., 2015). However, this conclusion does not necessarily support the wait-and-see channel because causation can run both ways. Indeed, another stream of the theoretical literature shows that lower growth may cause uncertainty through the implementation of new policies and risky behavior of firms and investors. Uncertainty and lower growth might also feed off each other in a vicious cycle. Different approaches to disentangling the two causal directions exist. Bloom et al. (2018) use a theoretical modeling approach to conclude that uncertainty shocks lead to drops in economic output of around 2.5%. Ludvigson et al. (2015) distinguish between macroeconomic and financial uncertainty, concluding that macro-uncertainty responds to output fluctuations while financial uncertainty drives them.

1.1.2 Beliefs

As illustrated by the above scenario of deciding whether to buy a house, it is intuitively clear that beliefs influence behavior. Experiments support this intuition, showing that expectations have a distinct impact on the behavior of consumers (Roth and Wohlfart, 2018) and firms (Enders et al., 2019). Because this behavior lies at the heart of the economy, one readily concludes that beliefs affect economic outcomes. As with uncertainty, it is not easy to operationalize beliefs. Studies of the impact of beliefs use different measures and draw different conclusions. Some find that beliefs play a prominent role in economic fluctuations (e.g., Chahrour and Jurado, 2018), while Fève and Guay (2019), for example, find that beliefs have little explanatory power.

[2] The literature specifies many other types of uncertainty, like model uncertainty, innovation uncertainty, and policy uncertainty. Especially policy uncertainty plays a prominent role in the uncertainty literature with the seminal paper by Baker et al. (2016) showing that it is associated with reduced investment and employment.

1.2 Modeling beliefs

Suppose we accept that beliefs are important determinants of economic outcomes. In that case, one might ask a further question: Do we have to consider them explicitly when testing economic policies or making economic predictions? Econometric modelers in the tradition of Tinbergen (1937, 1939) built macroeconomic models that only implicitly involved beliefs. They based their models on historical relationships between economic aggregates without an explicit role for beliefs. However, stagflation in the 1970s showed that historical relationships could break down: inflation skyrocketed in times of high unemployment, contradicting the historical inverse relationship between the two (the Phillips curve). Policy played an essential role in the explanations of this breakdown. This episode sparked criticism of the econometric approach to macroeconomics. A famous formulation is now known as the Lucas critique (Lucas, 1976): Lucas argued that when policy changes, economic agents form beliefs about the impact of these changes and alter their behavior accordingly. These behavioral shifts can break the historical relationships that the macroeconometric models relied on at the time. He concluded that these models are unfit for forecasting or policy analysis.[3]

The Lucas critique ushered in an immense shift in macroeconomic methodology. To respond to the critique, economists have introduced models founded on the behavior of the agents that make up the economy (microfoundations). Because these agents use beliefs to guide their behavior, beliefs play a central role in these models. They mostly rely on Rational Expectations, as introduced by Muth (1961), to model beliefs. The idea is that agents' predictions are the same as those implied by the model they inhabit. Rational Expectations are therefore also known as model-consistent expectations.

The literature has proposed theoretical and empirical arguments to move away from Rational Expectations toward a Bounded Rationality approach (e.g., Conlisk, 1996; Simon, 1957). Especially since the 2008 financial crisis and the ensuing recession, the idea has gained traction that self-fulfilling beliefs (also referred to as 'animal spirits' or sentiment) play a critical role in financial and economic fluctuations (e.g., Akerlof and Shiller, 2009; Kocherlakota, 2010). When you look for alternatives, you enter the so-called 'Bounded Rationality wilderness.' Contrary to model-consistent expectations, which have a unique specification for each model, there are infinitely many ways of specifying boundedly rational beliefs. The challenge thus is to discipline the beliefs in some way. Various answers have been formulated to this problem, offering some well-defined paths through the wilderness. In this dissertation, I explore three of these paths: Bayesian learning, heuristic switching, and Rational Beliefs.

Bayesianism embraces the idea that all statistical statements (e.g., relating to probabilities of outcomes, tests of significance, or distributions of model parameters) are to some extent subjective.[4] The idea is that one forms subjective beliefs based on pre-

[3]Although this critique is powerful in theory, the question of its relevance in practice remains. For example, Leeper and Zha (2003) find that what they call 'modest' policy interventions do not change beliefs significantly. Moreover, they conclude that one can reliably forecast their impact with precisely the type of models that Lucas criticizes.

[4]The main opposing interpretation of probability is frequentism, which asserts that the probability that an experiment gives a particular outcome is the proportion of outcomes in the limit that the experiment

vious experience and knowledge, called priors. When confronted with data, one updates these priors according to Bayes' law to derive posterior beliefs.[5] Under Bayesian learning, agents assume that the economy follows a specific model but do not know its parameters. They form prior beliefs about the parameters and continually update them as new economic data is released. They make predictions by iterating the model forward. Bayesian learning is part of the broader learning literature (e.g., Evans and Honkapohja, 2001). One of the main conclusions of this literature is that learning can converge to Rational Expectations under certain assumptions, thereby providing a foundation for Rational Expectations.[6]

Milani (2011) investigates the extent to which learning can describe real expectations and how significant any deviations are in explaining economic fluctuations. To measure real expectations, he uses data from surveys that ask professional forecasters to predict various economic quantities (Croushore, 1993). Milani incorporates these survey forecasts in a standard macroeconomic model that allows for learning by its agents.[7] He observes substantial departures of the survey forecasts from the learning model, which he attributes to waves of optimism and pessimism. He finds that these animal spirits explain roughly half of economic fluctuations.

In Milani's study, as in most learning and Rational Expectations models, one agent embodies the whole economy (the representative agent). The representative agent has received criticism about its failure to represent the heterogeneous preferences of its underlying agents (e.g., Kirman, 1992; Chiappori et al., 2009). A critical shortcoming of the representative agent in the context of this dissertation is that it does not allow for heterogeneous beliefs. Empirical evidence indicates that beliefs are indeed heterogeneous. Surveys show that professional forecasters' predictions diverge (Croushore, 1993), and Mankiw et al. (2003) show that consumers also have heterogeneous forecasts. Muth (1961) already noted that belief heterogeneity only matters when deviations from rationality are correlated. Otherwise, these deviations would disappear in the aggregate, and Rational Expectations would be appropriate. Studies show that correlated deviations from rationality indeed play an essential role in experimental asset markets (e.g., Smith et al., 1988) and laboratory economies (e.g., Hommes, 2021).

Convinced by the significance of heterogeneity and animal spirits in the financial and real economy, the heuristic switching literature, pioneered by Brock and Hommes (1997), attempts to incorporate them in a model of expectation formation. It adopts the idea that agents use simple forecasting rules (heuristics) to make predictions. They choose between these rules based on their past performance. As the performance of each rule changes over time, agents may switch to a different rule. They

is repeated infinitely many times.

[5] Say that based on our experience, we believe that the a priori probability of event A is $P(A)$. Now new data comes in: event B occurred. As Bayesians, we would use Bayes' law to derive that the probability $P(A|B)$ of event A given event B is given by

$$P(A|B) = \frac{P(B|A)P(A)}{P(B|A)P(A) + P(B|\neg A)P(\neg A)},$$

where $\neg A$ denotes the event that A does not occur.

[6] For convergence to Rational Expectations, agents need to have prior knowledge of some aspects of the truth (e.g., Evans and Honkapohja, 2001; Kalai and Lehrer, 1993).

[7] Milani (2011) uses constant-gain learning instead of Bayesian learning, but one can expect similar results.

do not necessarily choose the same rules, giving rise to time-varying heterogeneity. The heuristic switching literature could easily get lost in the wilderness, as agents could use many different rules. However, it has used laboratory experiments to isolate four heuristics that best describe subjects' forecasts. These heuristics vary from a strong-trend-following rule that can inflate asset bubbles to more moderate adaptive expectations. With these four forecasting rules, the heuristic switching model can explain correlated, non-rational expectations in different laboratory markets (e.g., Hommes, 2021). The heuristic switching literature provides further empirical support for the model, both for the underlying switching mechanism based on past performance (Branch, 2004) and for its ability to replicate patterns of business cycles (De Grauwe, 2012) and stock returns (Hommes, 2001).

Like the heuristic switching literature, Rational Beliefs theory, established by Kurz (1994), recognizes that individuals in a complex economy do not know its exact structure. However, instead of resorting to simple rules to deal with this uncertainty, the theory posits that agents form beliefs consistent with the data. Their beliefs can temporarily deviate from the distribution implied by the data (the empirical distribution), allowing for the waves of optimism and pessimism that Milani (2011) identifies. However, beliefs should align with the data in the long run, thereby taming the Bounded Rationality wilderness. Because different belief models are consistent with the data, agents deviate from the empirical distribution differently.[8] The theory assumes that these deviations are correlated across agents, meaning that they do not disappear in the aggregate. The average deviation constitutes a new macro-level variable about which agents have to form beliefs, which plays an essential role in the theory.

1.3 Overview of dissertation

1.3.1 Chapter 2

There is no consensus about the direct economic impact of uncertainty, but agents' beliefs about this impact can also influence the economy. Especially when uncertainty affects the decisions of policymakers, the impact could be substantial. In the second chapter of this dissertation, based on joint work with Giulia Piccillo, I investigate whether macroeconomic uncertainty affects monetary-policy decisions in the US. Eight times per year, the Federal Open Market Committee (FOMC) meets to review monetary policy in light of current economic conditions. I assume that the FOMC members use a standard macroeconomic model to make sense of the economic conditions. It captures the relationships between economic growth, inflation, and the interest rate. I further assume that the policymakers are Bayesian learners: as new data comes in, they update their beliefs about the model's parameters. These beliefs are represented by a probability distribution. I derive a measure of macroeconomic uncertainty from its dispersion. In constructing this uncertainty measure, I use macroeconomic data as it was available at each FOMC meeting.

I estimate the impact of this real-time, Bayesian measure of macroeconomic uncertainty on the FOMC's interest rate decisions. Here, I control for the impact of economic forecasts prepared by the Federal Reserve Board for each meeting (the

[8] Motolese and Nielsen (2007) explain this point in detail.

so-called Greenbook forecasts). I also include a measure of financial uncertainty to disentangle the roles played by financial and macroeconomic uncertainty in monetary policy. I find that policymakers set a significantly lower interest rate in times of higher macroeconomic uncertainty.

1.3.2 Chapter 3

Individuals have different attitudes to uncertainty. Some are risk-loving, while others shy away from risky situations (e.g., Choi et al., 2007; Cohen and Einav, 2007; Falk et al., 2018). In the third chapter of this dissertation, based on joint work with Giulia Piccillo, I investigate whether risk attitudes influence beliefs. I adopt the heuristic switching model, where I introduce a role for risk aversion in agents' choice between forecasting rules: they choose a rule based on its performance and the variability of that performance. As agents have different risk preferences, they choose different rules, leading to heterogeneous expectations. To empirically validate the model, I draw the agents' risk aversions from a distribution based on survey data.

I incorporate this belief-formation model in a stylized financial market, which consists of one risky and one riskless asset. Agents form beliefs about the risky asset's price to determine how to divide their wealth between both assets. I prove that a representative agent cannot capture this model. Simulations show that the resulting belief dynamics can drive unpredictable booms and busts in the asset price. Introducing small stochastic price shocks leads to larger asset price bubbles and can destabilize markets.

1.3.3 Chapter 4

In chapter 4, I propose an explanation for the mixed results from studies about the role of animal spirits in economics: these studies measure different dimensions of sentiment that have distinct macroeconomic impacts. To test this hypothesis, I rely on Rational Beliefs theory. It offers a clear definition of animal spirits as temporary deviations from the empirical distribution. This definition also suggests measuring them as the difference between observed forecasts and forecasts implied by the empirical distribution. Survey data on the projections of professional forecasters covers the first half of the equation. I approximate the empirical forecasts by collecting a large panel of real-time data covering all relevant aspects of the economy and using a statistical model to produce predictions.

I use the full extent of the 50 years of survey data, covering approximately 40 forecasters per survey, various economic variables (e.g., output, prices, interest rates, housing), and multiple forecasting horizons. Subtracting the corresponding empirical forecasts gives a panel of sentiment for all surveys and forecasters across all combinations of variables and horizons. I use a statistical procedure to extract three dimensions that together capture about 50% of forecasters' animal spirits. I find that sentiment is indeed multidimensional, with the first dimension explaining only about a fifth of forecasters' animal spirits. I furthermore find that each dimension has a distinct macroeconomic impact, supporting my hypothesis.

To guide the empirical analysis, I develop a theoretical framework based on Rational Beliefs theory. This framework allows me to derive testable implications of the theory. This chapter thus also serves as a test of its empirical relevance. My results support some of the theory's implications while undermining others.

Chapter 2

Does US monetary policy respond to macroeconomic uncertainty?[1]

Abstract

We find that macroeconomic uncertainty plays a significant role in U.S. monetary policy. First, we construct a measure of uncertainty as felt by policymakers at the time of making their rate-setting decisions. This measure is derived from a real-time, Bayesian estimation of a small monetary VAR with time-varying parameters. We use it to calculate the probability of being in a high-uncertainty regime. Second, we estimate a monetary policy reaction function that, apart from macroeconomic uncertainty, includes Greenbook forecasts, revisions of those forecasts, and a measure of stock market volatility. Using data for the period 1969 – 2008, we find that policymakers set an interest rate that is significantly lower in a high-uncertainty regime, compared to a low-uncertainty regime.

2.1 Introduction

Monetary policymakers often emphasize the importance of uncertainty. Greenspan (2004), for example, notes that "uncertainty is not just a pervasive feature of the monetary policy landscape; it is the defining characteristic of that landscape" (p. 36). Uncertainty has also been proposed as an important driver of business cycles (e.g., Bloom et al., 2018; Caldara et al., 2016). The idea is that uncertainty causes cautious behaviour by consumers and firms, leading to decreasing output and increasing unemployment.[2] To counteract these effects, monetary policy should be looser in

[1] This chapter is based on joint work with Giulia Piccillo. Thomas Gomez is the lead author, and has been the main contributor in most phases of this work, including data management, data analysis, results interpretation and writing the chapter. Giulia Piccillo played a major role in the idea generation phase of the study. This chapter was supported by the Complex Systems Fund, with special thanks to Peter Koeze.

[2] Theoretically, the link between uncertainty and the behaviour of economic actors has been established for years. Leland (1968) shows how income uncertainty can lower consumption through precautionary savings. Batra and Ullah (1974) shows that firms decrease their output in reaction to increases in price uncertainty. The real options literature describes how (partly) irreversible consumption and investment can be postponed in uncertain times: It can be preferable to wait for more information than to make a costly mistake (e.g., Bernanke, 1983).

high-uncertainty regimes.[3] Uncertainty has resurfaced in the monetary policy debate in the context of the COVID-19 pandemic (e.g., Panetta, 2020).

We set out to empirically determine the response of monetary policy to macroeconomic uncertainty. Our analysis consists of two parts: (1) constructing an appropriate measure of macroeconomic uncertainty, and (2) estimating a monetary policy reaction function that includes this uncertainty measure.

An appropriate macroeconomic uncertainty measure should capture uncertainty as felt by policymakers at the time they make rate-setting decisions. In the case of the United States, this means uncertainty felt by the members of the Federal Open Market Committee (FOMC) during their meetings. We pursue this goal by incorporating three key elements.

First, we treat uncertainty as inherently subjective, and therefore take a Bayesian approach. This is in line with the description by Greenspan (2004) of the risk-management approach to monetary policy as being an application of Bayesian decision-making. Additionally, a Bayesian approach to uncertainty quantification is standard in the statistics literature (e.g., Berger and Smith, 2019), and supported by evidence on human decision-making and learning (Kording, 2014; Viscusi, 1985; Yu, 2007).

Second, we base our uncertainty measure on the data that is available at the time of each FOMC meeting. This data differs from the currently available data because of data revisions, which can be substantial, especially in times of uncertainty. By using real-time data, we can measure uncertainty as perceived at the time that policymakers make rate-setting decisions. It has been known for some time that it is important to use real-time data in the analysis of monetary policy (e.g., Orphanides, 2001).

Third, we measure the uncertainty surrounding the relationships among a small set of key macroeconomic variables, namely, real output growth, inflation, and the effective federal funds rate. Our aim is to incorporate the idea formulated by Greenspan (2004, p. 37) as follows: "A critical result has been the identification of a relatively small set of key relationships that, taken together, provide a useful approximation of our economy's dynamics."

We capture these relationships in a Bayesian time-varying parameter VAR (TVP-VAR) model as in Koop and Korobilis (2013).[4] For each FOMC meeting, we estimate a Bayesian TVP-VAR on the associated real-time data. We derive the posterior density of one-quarter-ahead forecasts and compute uncertainty as its differential entropy. Entropy has previously been used to measure uncertainty (e.g., Rich and Tracy, 2010). Additionally, we check whether the uncertainty felt by policymakers, and its influence on the interest rate, might be better captured in terms of low- and high-uncertainty regimes than as a linear response. A small change in the uncertainty measure might lead to a shift in regimes and have a large interest rate effect. To take this into account, we also compute the probability of being in a high-uncertainty regime for every meeting.

[3] Other reasons for monetary policy to respond to uncertainty have been proposed in the literature. Evans et al. (2016) provide an overview.

[4] Their approach uses a forgetting factor, which allows for the model's coefficients to change over time. Furthermore, it is computationally efficient, which is useful because we have to recalculate the uncertainty measure (i.e., repeat the TVP-VAR estimation) for each FOMC meeting. Bayesian VARs are standard forecasting tools with a long history in macroeconometrics (e.g., Litterman, 1986). Forgetting factor approaches go as far back as the 1960s (Koop and Korobilis, 2013).

In the second part of this chapter, we investigate how macroeconomic uncertainty affects monetary policy decisions by the FED. We begin by estimating an extended version of the monetary policy reaction function used by Romer and Romer (2004). They combine quantitative and narrative sources to identify the intended change to the federal funds rate surrounding each FOMC meeting, and use this as dependent variable in their reaction function. As explanatory variables, they include Greenbook forecasts for output growth, inflation, and the unemployment rate at various horizons. They also include revisions of those forecasts compared to the previous meeting. We augment this reaction function with our macroeconomic uncertainty measure, and we also add the VXO measure of stock market volatility as a proxy for financial uncertainty.

Following Romer and Romer (2004), we estimate the reaction function meeting by meeting. Compared to a monthly or quarterly specification, this has the advantage that it prevents endogeneity issues related to the impact of monetary policy on uncertainty that is evidenced by Mumtaz and Theodoridis (2019): at the meeting frequency there is simply no time for the policy decisions to affect uncertainty. Our baseline sample covers FOMC meetings in the period 1969 – 2008. We also estimate policy reaction functions separately on the periods 1969 – 1979 and 1987 – 2008.

We have three main results. First, U.S. monetary policy is significantly affected by macroeconomic uncertainty. On the full sample, the linear response to a one-standard-deviation increase in macroeconomic uncertainty is a decrease in the intended funds rate of 7 basis points.

Second, we find that instead of a linear response, the role of macroeconomic uncertainty is best captured in terms of low- and high-uncertainty regimes. The Fed sets a funds rate that is 14 basis points lower in a high-uncertainty regime than in a low-uncertainty regime. When calculating the probability of being in a high-uncertainty regime with output or inflation uncertainty, the differences between the two regimes are 17 and 16 basis points, respectively.

Third, financial uncertainty, as proxied by the VXO index, plays a significant role in monetary policy that is separate from the one played by macroeconomic uncertainty. On the full sample, a one-standard-deviation increase in the VXO index is associated with a decrease in the intended funds rate of 4.3 basis points.

We contribute to two strands of literature. The first is on risk management in monetary policy. We discuss the two contributions that are closest to ours. First, Evans et al. (2016) also investigate the role of uncertainty in monetary policy by estimating a battery of monetary policy reaction functions. They first estimate reaction functions at the meeting frequency, using uncertainty indicators that are based on FOMC meeting notes.[5] They find that the funds rate is about 8 basis points higher under uncertainty. This stands in stark contrast with our result of a lower funds rate under uncertainty over a similar sample (their sample is comparable to our later subsample, starting with the onset of Greenspan's tenure as chair). Evans et al. also estimate quarterly reaction functions with various uncertainty proxies.[6] Note that such a specification

[5]Evans et al. (2016) also use forecast revisions as an uncertainty proxy for meeting-by-meeting estimates. They argue that revisions are often caused by unusual events that are difficult to interpret, thereby raising uncertainty. We would argue however, that revisions reflect newly available information, which could either increase or decrease uncertainty. We include revisions in the reaction function together with the uncertainty measures.

[6]These include financial uncertainty proxies, like the VXO index, an uncertainty measure by Jurado

might suffer from endogeneity issues due to the simultaneous impact of monetary policy on uncertainty (as we discussed earlier). They find mixed results, with some uncertainty measures associated with sizeable decreases and other proxies with significant increases in the funds rate.

We highlight three important differences in the identification procedure that may explain the contrast between our results and those of Evans et al. (2016). First of all, we include a specific measure of macroeconomic uncertainty, while their FOMC-based indicator (as well as some of their uncertainty proxies) potentially mixes financial and macroeconomic uncertainty. Second, we include more explanatory variables in the reaction function, and financial and macroeconomic uncertainty measures at the same time, which allows us to disentangle their effects. Evans et al. on the other hand, enter one uncertainty measure at a time. Third, we use the intended funds rate instead of the realized, effective rate. This accounts for the fact that the Fed does not have full control over the interest rate.

Second, Gnabo and Moccero (2015) estimate non-linear monetary policy reaction functions that allow for different responses in regimes of high and low uncertainty. They measure uncertainty as dispersion in inflation forecasts from the Survey of Professional Forecasters (SPF) or as the level of the VXO index. Their sample spans the period from the start of Greenspan's tenure to the end of 2005. They find that monetary policy reacts more aggressively to the output gap in high-uncertainty regimes (for both uncertainty measures), but they find no difference for inflation forecasts. We do not identify the effect of uncertainty on the response to specific forecasts, but we complement their study by disentangling the direct response to different types of uncertainty. Together, these studies support the broader notion that uncertainty affects monetary policy.

Methodologically, we also contribute to the literature that studies macroeconomic uncertainty. In this literature, our work comes closest to Orlik and Veldkamp (2014), who also use a Bayesian approach and real-time data to study uncertainty. Their goal is not to introduce the most appropriate measure, but to explain why macroeconomic uncertainty fluctuates. They use non-normal priors for a simple model of GDP growth, and show that changing estimates of disaster risk lead to large and countercyclical uncertainty fluctuations. Our model includes inflation and the interest rate in addition to output growth, and we find substantial countercyclical fluctuations in our macroeconomic uncertainty measure even though we use normal priors.

2.2 Method

2.2.1 Data

For the monetary policy reaction function, we use data on changes to the intended federal funds rate around FOMC meetings. Before 1997, this data is provided by Romer and Romer (2004). The period from 1997 onwards is covered by data published by St. Louis Federal Reserve Bank Economic Data (FRED), which is based on

et al. (2015) that combines macroeconomic and financial uncertainty, and a number of measures based on the Survey of Professional Forecasters (SPF). The uncertainty proxy used by Evans et al. that comes closest to our measure, is the one introduced by Jurado et al. (2015). Contrary to our measure, it uses revised data, takes a frequentist approach and combines financial and macroeconomic uncertainty.

FOMC meeting transcripts and statements. We also collect data on economic projections produced by the staff at the Board of Governors of the Federal Reserve System (Greenbook forecasts). These are the forecasts available to the FOMC meeting members. Specifically, we include forecasts for the quarterly average of the unemployment rate and for annualized quarter-on-quarter real output growth and inflation, as published by the Federal Reserve Bank of Philadelphia.[7]

Our baseline sample covers the FOMC meetings in the period January 1969 – October 2008. The January 1969 meeting is the first meeting for which intended funds rate data is provided by Romer and Romer. The October 2008 meeting is the last meeting before the ZLB was hit. We exclude the ZLB period, because it cannot be modelled by a reaction function intended to describe conventional monetary policy. The period after the funds rate moved away from the zero lower bound is excluded from the analysis, because it is not covered by our Greenbook data.

For the calculation of our macroeconomic uncertainty measure, we construct a real-time dataset. This dataset consists of vintages of the available data at the time of each FOMC meeting. Each vintage can be different because new data is released, or old data is revised. It contains quarterly data for the effective federal funds rate and for quarter-on-quarter annualized real output growth and inflation.

The Philadelphia Fed provides real-time output and price index data. This data consists of monthly vintages that reflect the data available in the middle of the associated month. Vintages for the first month of quarter t have observations for 1947Q1 up to and including quarter $t-2$. Vintages associated with the second and third month of the quarter span the period from 1947Q1 up to and including quarter $t-1$.[8]

Apart from forecasts, the Greenbook data contains historical values for up to four quarters before each Greenbook is released. We carefully match vintages to meetings by comparing these historical values to the corresponding observations in the real-time dataset. However, because most vintage dates do not match exactly with the Greenbook release dates, some differences remain, especially in the most recent observation at the time (the quarter before the Greenbook is released). To ensure that our data accurately reflects the information available to the FOMC meeting participants, we replace observations in the vintage with historical Greenbook values as far as they are available.

More recent information about the state of the economy may be available at the time of each meeting than is captured by its associated vintage. This is due to higher-frequency data like the unemployment rate and industrial production. This additional information may affect the uncertainty felt by the meeting members. To take this into account, while keeping the model simple, we include the information implicitly. We do this by adding to each vintage the Greenbook projections for all quarters up to and including the quarter in which the meeting takes place. As a robustness test, we repeat our analysis without adding the Greenbook projection for the current quarter.

[7] Wieland and Yang (2020) report some errors in this data as published by the Philadelphia Fed. We adopt their corrections after double checking, and include some of our own corrections based on Greenbook and supplement documents.

[8] Some vintages have a later starting date. Specifically, the December 1991 – December 1992 and November 1999 – March 2000 vintages start in 1959Q1, while the January 1996 – April 1997 vintages start in 1959Q3.

Funds rate data is provided by FRED. Funds rate projections are not included in the Greenbook dataset, but the Philadelphia Fed does publish a separate dataset with the funds rate assumptions underlying the Greenbook forecasts. This data is available for meetings in the period 1981Q1 – 2008Q3. For the remainder of the sample, we estimate the current funds rate by using the available monthly funds rate data from before the meeting, and assume that it stays at the target rate level immediately before the meeting for the remainder of the quarter.[9]

The inflation and federal funds rate series are differenced to make them approximately stationary. All data is standardized using the mean and standard deviation computed over the vintage after the historical values have been replaced. Their calculation excludes the two most recent observations however, so they do not use projections for the current and previous quarters. Funds rate data is available only from 1954Q3 onwards, which means that after differencing and including Greenbook projections, each vintage covers data from 1954Q4 up to and including the quarter the meeting takes place. Henceforth, when we refer to the vintage associated with a meeting, we refer to these adjusted vintages, and not to the original vintages provided by the Philadelphia Fed.

We proxy financial uncertainty with the VXO index on the Greenbook release day. VXO data is only available from 1986 onwards. Following standard practice (e.g., Bloom, 2009), we approximate it by the 30-day standard deviation of S&P500 daily returns before 1986. The realized volatility series is standardized to have the same mean and variance as the VXO index over the period where they overlap.

2.2.2 Construction of macroeconomic uncertainty measure

For each FOMC meeting, we estimate a simple monetary VAR with time-varying parameters (TVP-VAR) on the associated real-time data vintage. Using a Bayesian approach, we derive the posterior density of one-quarter-ahead forecasts. We compute uncertainty as the differential entropy of this density. The priors for the VAR coefficients and the degree of their time-variability are determined by three hyperparameters. We find reasonable values for these hyperparameters by fitting the TVP-VAR forecasts for output growth and inflation at every FOMC meeting to the corresponding Greenbook forecasts.

Bayesian time-varying parameter VAR

Consider one of the real-time data vintages (corresponding to one of the FOMC meetings), with T quarterly observations indexed by $t = 1, 2, \ldots, T$. The TVP-VAR consists of annualized quarter-on-quarter real output growth (y_t), the first difference of annualized quarter-on-quarter inflation ($\Delta\pi_t$), and the first-differenced effective federal funds rate (Δi_t). As is standard in quarterly VARs, we include four lags.[10]

[9] For the period where the Greenbook assumptions are available, our estimate for the current quarter funds rate is very similar, with a correlation coefficient between the two of 0.9994.

[10] An alternative to using a fixed number of lags would be to use an information criterion like BIC to select a lag length. However, Stock and Watson (2002) show that a similar VAR performs better with a fixed lag length of four, than with a BIC-selected lag length.

We define $x_t = \left(y_t, \Delta\pi_t, \Delta i_t\right)'$ and write the TVP-VAR as

$$x_t = Z_t \beta_t + \epsilon_t, \tag{2.1a}$$
$$\beta_t = \beta_{t-1} + q_t, \tag{2.1b}$$

with $\epsilon_t \sim \mathcal{N}(0, \Sigma_t)$, $q_t \sim \mathcal{N}(0, Q_t)$, β_t the time-varying coefficient vector, and

$$Z_t = I_3 \otimes X_{t-1}, \quad X_{t-1} = (1, x'_{t-1}, \ldots, x'_{t-4}).$$

We adopt the Bayesian estimation approach of Koop and Korobilis (2013). It uses the Kalman filter, which is a recursive estimator. Consider $\tau \in \{1, \ldots, T\}$. Given a normal prior for the coefficient vector β_0, and Σ_t and Q_t for $t = 1, \ldots, \tau$, the Kalman filter gives expressions for the posterior mean $\beta_{\tau|\tau}$ and covariance matrix $V_{\tau|\tau}$ of the coefficient vector, conditional on the observations through time τ:[11]

$$\beta_\tau | x_1, x_2, \ldots, x_\tau \sim \mathcal{N}\left(\beta_{\tau|\tau}, V_{\tau|\tau}\right). \tag{2.2}$$

For $\tau = T$, this corresponds to the posterior conditional on all observations in the vintage. This is the posterior we use in the uncertainty calculation later on.

We denote the prior mean and covariance matrix by $\beta_{0|0}$ and $V_{0|0}$, respectively:

$$\beta_0 \sim \mathcal{N}\left(\beta_{0|0}, V_{0|0}\right).$$

We use the same prior as Koop and Korobilis, which is a variant of the classical Minnesota prior (Doan et al., 1984; Litterman, 1986). Because we transform the variables in our VAR to approximate stationarity in our data setup, we set $\beta_{0|0} = 0$. We define $V_{0|0}$ to be diagonal, with diagonal elements $V^{ii}_{0|0}$:

$$V^{ii}_{0|0} = \begin{cases} \dfrac{\gamma}{l^2} & \text{for coefficients on lag } l = 1, \ldots, 4; \\ 100 & \text{for the intercepts.} \end{cases} \tag{2.3}$$

The overall tightness of the prior is determined by the parameter γ. The prior on coefficients of older lags is tighter than that for more recent lag coefficients: The assumption is that older lags have a smaller impact on the current value, compared to more recent ones.

We follow Koop and Korobilis in replacing Σ_t and Q_t by estimates, denoted by $\widehat{\Sigma}_t$ and $\widehat{Q}_t$. This has the advantage that no priors have to be defined for these covariance matrices, and that no computationally expensive MCMC methods are required. We denote the resulting estimates for the posterior mean and covariance matrix by $\widehat{\beta}_{t|t}$ and $\widehat{V}_{t|t}$.

First, we estimate Σ_t with an exponentially weighted moving average with decay factor $0 \leq \kappa \leq 1$:

$$\widehat{\Sigma}_t = \kappa\, \widehat{\Sigma}_{t-1} + (1-\kappa)\, \widehat{\epsilon}_t \widehat{\epsilon}'_t. \tag{2.4}$$

Here, the Kalman filter gives the estimated residual $\widehat{\epsilon}_t = x_t - Z_t \beta_{t|t}$. The decay factor determines the degree of time-variability of Σ_t: The larger κ, the slower the dynamics.

[11] We use the standard Kalman filtering formulae (see, e.g., Durbin and Koopman, 2012).

In the extreme case where $\kappa = 1$, the covariance matrix is constant. Following Koop and Korobilis, we use the sample covariance matrix of the whole vintage as the initial value $\widehat{\Sigma}_0$.

Second, we posit that $\widehat{Q}_t = \left(\lambda^{-1} - 1\right) \widehat{V}_{t-1|t-1}$, where $0 < \lambda \leq 1$ is called the forgetting factor. The forgetting factor determines the amount of weight put on older observations compared to the current one. A larger λ implies that the coefficient vector β_t changes more slowly.

Calculating uncertainty

The Kalman filter provides us with estimates for the posterior mean $\widehat{\beta}_{T|T}$ and covariance matrix $\widehat{V}_{T|T}$ conditional on all the data contained in the vintage. The estimated density for the coefficient vector in the next quarter (β_{T+1}) is given by

$$\widehat{\beta}_{T+1}|x_1, \ldots, x_T \sim \mathcal{N}\left(\widehat{\beta}_{T+1|T}, \widehat{V}_{T+1|T}\right),$$

with mean and covariance matrix given by (2.1b):

$$\widehat{\beta}_{T+1|T} = \widehat{\beta}_{T|T}, \quad \widehat{V}_{T+1|T} = \widehat{V}_{T|T} + \widehat{Q}_{T+1} = \lambda^{-1}\widehat{V}_{T|T}.$$

Since Z_{T+1} is known, it follows that the posterior forecast for the state of the economy one quarter ahead (x_{T+1}) is normal, with mean x^F_{T+1} and covariance matrix $\widehat{R}_{T+1}$ given by

$$x^F_{T+1} = Z_{T+1}\,\widehat{\beta}_{T+1|T}, \quad \widehat{R}_{T+1} = Z_{T+1}\,\widehat{V}_{T+1|T}\,Z'_{T+1}.$$

Recall however, that the inflation and interest rate series are differenced, and that all variables in the VAR are standardized. We transform back to non-standardized forecasts of the series that are included in the Greenbook, because these are the forecasts the policymakers are concerned with. We denote the non-standardized, non-differenced version of x_t by $\widetilde{x}_t$. We define $\widehat{\mu}_x$ to be the mean used to standardize x, and $\widehat{\Sigma}_x$ to be the matrix with the standard deviations used to standardize x on its diagonal. Furthermore defining $\chi_t = (0, \pi_t, i_t)'$, it follows that the estimated posterior density of forecasts for $\widetilde{x}_{T+1}$ is normal, with mean $\widetilde{x}^F_{T+1}$ and covariance matrix $\widetilde{R}_{T+1}$ given by

$$\widetilde{x}^F_{T+1} = \chi_T + \widehat{\Sigma}_x\,x^F_{T+1} + \widehat{\mu}_x, \quad \widetilde{R}_{T+1} = \widehat{\Sigma}_x\widehat{R}_{T+1}\widehat{\Sigma}'_x.$$

We now denote the length of the vintage corresponding to FOMC meeting m by T_m, and its associated posterior mean and covariance matrix by $\widetilde{x}^F_{m,T_m+1}$ and $\widetilde{R}_{m,T_m+1}$. We define the macroeconomic uncertainty measure, denoted by U_{Mm}, as the differential entropy[12] of the posterior density of one-quarter-ahead forecasts:

$$U_{Mm} = \frac{3}{2} + \frac{3}{2}\log 2\pi + \frac{1}{2}\log\det\widetilde{R}_{m,T_m+1}. \tag{2.5}$$

[12]Differential entropy, a concept from information theory, is defined as

$$-\int f(x)\log f(x)\,dx,$$

where $f(x)$ is the probability density function. An expression in the case of a multivariate normal distribution is given by Ahmed and Gokhale (1989).

This entropy takes into account the covariance between the forecasts of different variables. The last term in equation (2.5) can be rewritten as

$$\frac{1}{2}\log\det\widetilde{R}_{m,T_m+1} = \sum_{i=1}^{3}\log\left(\sqrt{eig_{im}}\right), \tag{2.6}$$

where eig_{im} are the eigenvalues of the covariance matrix $\widetilde{R}_{m,T_m+1}$. The square roots of these eigenvalues give the standard deviations along the orthogonal components of maximum variation (the principal components, given by the corresponding eigenvectors). Larger covariances lead to a smaller differential entropy.[13]

In addition to this aggregate measure of macroeconomic uncertainty, we assess the policy reaction to the uncertainty surrounding the forecasts for individual series. In particular, we consider the posterior standard deviation of the one-quarter-ahead forecasts of output growth (U_{ym}) and inflation ($U_{\pi m}$). These are computed as the square root of the corresponding diagonal element in the covariance matrix of the posterior density of forecasts ($\widetilde{R}_{m,T_m+1}$).

Choosing hyperparameters

Before we move to the monetary policy reaction function, let us note that the TVP-VAR contains three hyperparameters: the forgetting factor λ, the decay factor κ, and the tightness parameter γ. We fit these hyperparameters to the Greenbook data by constrained maximum likelihood estimation of the following model:

$$\begin{pmatrix} y^{GB}_{m,T_m+1} \\ \pi^{GB}_{m,T_m+1} \end{pmatrix} = \begin{pmatrix} \widetilde{y}^{F}_{m,T_m+1}(\lambda,\kappa,\gamma) \\ \widetilde{\pi}^{F}_{m,T_m+1}(\lambda,\kappa,\gamma) \end{pmatrix} + e_m, \tag{2.7a}$$

$$e_m = \overline{e} + P\, e_{m-1} + \eta_m, \tag{2.7b}$$

where $\eta_m \sim \mathcal{N}(0,S)$. Here, we emphasize the dependence of the TVP-VAR forecasts on the hyperparameters, and indicate the one-quarter-ahead forecasts for output and inflation included in the Greenbook for meeting m by y^{GB}_{m,T_m+1} and π^{GB}_{m,T_m+1}, respectively. This model assumes that differences between the Greenbook and TVP-VAR forecasts follow a VAR(1) process. This means that it can take into account any first-order serial and cross-correlations in these differences. This approach is similar to how Milani (2011) fits a comparable learning model to forecasts from the SPF.

We let $0.94 \le \kappa \le 0.98$, following Koop and Korobilis (2013), while λ and γ are constrained to the unit interval. We only include meetings from the period January 1969 – October 2008 in the estimation (the same period we use for estimation of the monetary policy reaction function). The meetings before the starting date are excluded because we do not have intended funds rate data for them, and hence have less reliable estimates for the current federal funds rate. We also exclude the meetings after October 2008 for which the ZLB is binding.

[13] This entropy measure is similar to a measure of forecasting performance that is sometimes used, namely, the log determinant of the covariance matrix of forecast errors (e.g., Smets and Wouters, 2007).

As a robustness exercise, we repeat our analysis for deviations from the maximum likelihood hyperparameters. We consider the range of values that Koop and Korobilis (2013) find to be optimal hyperparameters for a small monetary TVP-VAR.[14]

2.2.3 Monetary policy reaction function

To identify the impact of macroeconomic uncertainty on monetary policy, we estimate the monetary policy reaction function used by Romer and Romer (2004), augmented with uncertainty measures:

$$\begin{aligned} \Delta ff_m = {} & \gamma_0 + \gamma_b ffb_m + \sum_{i=-1}^{2} \gamma_{yi}\, y_{mi}^{GB} + \sum_{i=-1}^{2} \gamma_{\pi i}\, \pi_{mi}^{GB} + \gamma_{u0}\, u_{m0}^{GB} \\ & + \sum_{i=-1}^{2} \delta_{yi}\, \Delta y_{mi}^{GB} + \sum_{i=-1}^{2} \delta_{\pi i}\, \Delta\pi_{mi}^{GB} + \phi_{VXO}\, VXO_m + \phi_U\, U_m + v_m. \end{aligned} \tag{2.8}$$

Here, Δff_m indicates the change in the intended federal funds rate around FOMC meeting m, and ffb_m refers to the intended funds rate before any changes related to that meeting. This lagged interest rate can capture any mean-reverting behaviour of the intended funds rate.

We indicate Greenbook forecasts for meeting m by y_{mi}^{GB} (output growth), π_{mi}^{GB} (inflation) and u_{mi}^{GB} (unemployment), where i indicates the horizon relative to meeting m: $i = -1$ corresponds to the previous quarter, $i = 0$ to the current quarter, and $i = 1$ and $i = 2$ to the the one-quarter- and two-quarters-ahead forecasts, respectively.[15] Projections for the previous quarter account for lagged economic conditions, while unemployment forecasts are included because maximum sustainable employment is one of the explicit goals of the Fed.

The reaction function also includes forecast revisions $\Delta y_{mi}^{GB} = y_{mi}^{GB} - y_{m-1,i}^{GB}$, with a similar definition for $\Delta\pi_{mi}^{GB}$. Here, the forecasts refer to the same quarter: For example, if meeting m takes place in quarter t, while meeting $m-1$ took place in the quarter before $(t-1)$, the forecast revision Δy_{m0}^{GB} is calculated using the meeting m forecast for the current quarter and the one-quarter-ahead forecast from meeting $m-1$. Both are forecasts for quarter t. These forecast revisions are likely to impact the intended interest rate. Suppose that the rate setting decision at the previous meeting was based on forecasts that now have been extensively revised. This might lead to a change in the interest rate that is larger than just explained by the level of the new forecasts and the interest rate lag.

Most importantly for our analysis, we include financial and macroeconomic uncertainty in the monetary policy reaction function. Including both allows us to disentangle the response to financial uncertainty, which has already been established in

[14] Koop and Korobilis (2013) estimate (among other models) a similar small TVP-VAR, but with inflation measured by the consumer price index, and allowing the hyperparameters to change over time. They find that $\lambda = 1$ is optimal most of the time, with some periods of $\lambda = 0.99$ before 1985, and a brief period of $\lambda = 0.98$ in the early 1980s. They find an optimal $\gamma = 0.05$, apart from two periods in the mid-1980s, where $\gamma = 0.1$ is optimal. We explore the robustness of our results to changing the hyperparameters within these bounds. Koop and Korobilis do not report how the optimal decay factor evolves over time, but they allow for $\kappa \in \{0.94, 0.96, 0.98\}$.

[15] In many cases, a first release for previous quarter data is already available, and hence is not a forecast.

the literature, from that to macroeconomic uncertainty. The VXO at the time of meeting m is denoted by VXO_m. We consider several variants of the reaction function that differ in terms of the macroeconomic uncertainty measure (U_m) that is included. Our main specification uses the entropy measure of aggregate uncertainty ($U_m = U_{Mm}$), but we also estimate the reaction to uncertainty related to output growth ($U_m = U_{ym}$) and inflation ($U_m = U_{\pi m}$). We do not consider interest rate uncertainty, because funds rate forecasts are not as relevant for monetary policymakers (since they are the ones who determine the future rate).

In addition to estimating a linear relationship between the uncertainty measures and the interest rate, we consider a non-linear transformation that aims to measure the probability of being in a high-uncertainty regime. We adopt the logistic function used by Falck et al. (2019) in the context of regimes of disagreement among forecasters. It gives the probability of being in a high-uncertainty regime as a function of the uncertainty measure U_m:

$$F(U_m) = \frac{\exp\left(\theta \frac{U_m - c}{\sigma_U}\right)}{1 + \exp\left(\theta \frac{U_m - c}{\sigma_U}\right)}, \tag{2.9}$$

where c is the median and σ_U the standard deviation of U. The parameter θ controls how strongly the probability responds to changes in uncertainty. Such a specification is also considered by Gnabo and Moccero (2015) to identify uncertainty regimes. However, instead of estimating the parameters, we follow Falck et al. (2019) in choosing $\theta = 5$ and assessing our results' robustness to changing its value.[16] The function values lie between zero and one, and the corresponding coefficient in the reaction function gives the effect of being in a high-uncertainty regime.

Lastly, v_m is an error term.

We estimate the reaction function (2.8) by least squares on the full sample, as well as on two subsamples. The first subsample is characterized by the tenure of Martins, Burns, and Miller as chairmen of the Fed, and covers meetings in the period 14 January 1969 – 11 July 1979. The second subsample covers Greenspan's tenure, as well as part of Bernanke's. It spans meetings between 18 August 1987 and 29 October 2008.[17] For these subsamples, we separately re-standardize the uncertainty measures and re-calculate the probabilities of being in a high-uncertainty regime.

These two subsamples exclude the period October 1979 – October 1982 in which the Fed stopped targeting the funds rate, and targeted non-borrowed reserves instead. While Romer and Romer (2004) include this period in their estimation, others have argued that it can lead to biased results (Coibion and Gorodnichenko, 2011; Coibion, 2012).[18] By both including and excluding this period, we can investigate to what extent the results are driven by this period.

[16] We also consider the special case where $F(U_m)$ is a dummy variable that equals one when the uncertainty measure is above its median and zero otherwise.

[17] These two samples are also used by Caggiano et al. (2018), and like our second subsample, the sample of Evans et al. (2016) starts with the beginning of Greenspan's tenure as chairman.

[18] Romer and Romer (2004) note that even in this period, the Fed was concerned about the federal funds rate, and discussed its behaviour. They argue that the change in the intended funds rate therefore is the easiest indicator of monetary policy over a long timespan where monetary policy has changed. However, Coibion (2012) suggests that this is the period where the identification is most likely to be misspecified. He shows that the results of Romer and Romer (2004) are highly sensitive to excluding the period of

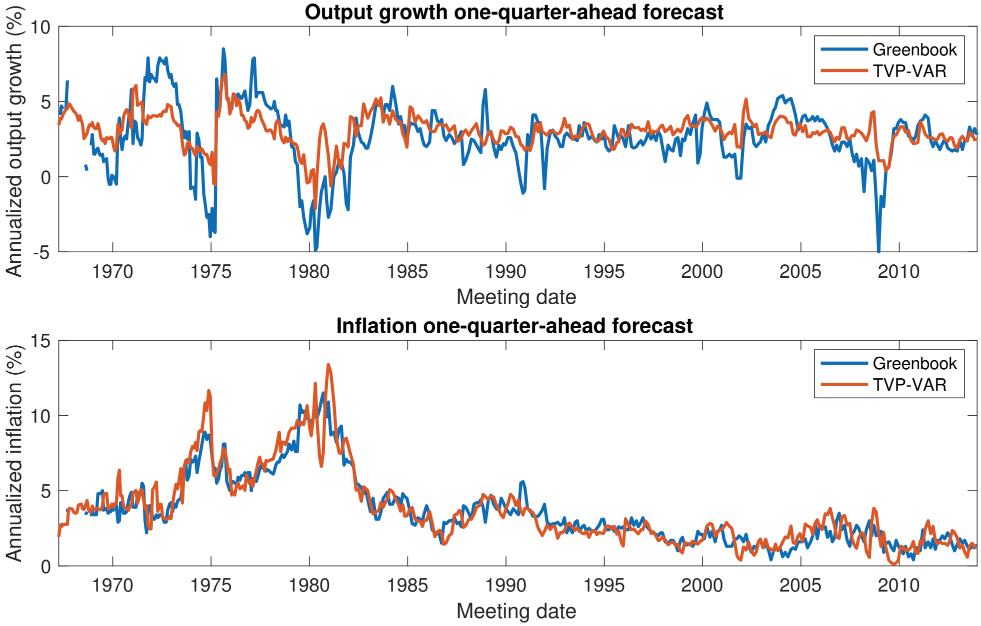

Figure 2.1: Comparison between one-quarter-ahead forecasts provided by the TVP-VAR and by the Greenbook for output growth (*Top*) and inflation (*Bottom*).

2.3 Results

2.3.1 Real-time macroeconomic uncertainty

We find the following maximum likelihood hyperparameters:

$$\lambda^* = 1.000, \ 95\% \text{ CI } [0.993, 1.000];$$
$$\kappa^* = 0.980, \ 95\% \text{ CI } [0.949, 0.980];$$
$$\gamma^* = 0.031, \ 95\% \text{ CI } [0.016, 0.038].$$

The confidence intervals are based on a bootstrap procedure with 500 replications.[19] This value for λ corresponds to the constant coefficient case: Older observations receive the same weight as current observations. The value for κ means that the covariance matrix of the VAR error term changes relatively slowly. These hyperparameters are in line with the optimal parameters that Koop and Korobilis (2013) find for a small monetary TVP-VAR (see footnote 14).

We provide a comparison of the TVP-VAR forecasts conditional on these parameters with those provided by the Greenbook in Figure 2.1. The model forecasts seem to be a decent approximation of the Greenbook forecasts. The correlation between the two forecast series is 0.72 ($p = 0$) for output growth and 0.92 ($p = 0$) for inflation.

non-borrowed reserves targeting. This is partly due to the fact that this period contains the largest funds rate changes in the sample.

[19]The estimated values for λ and κ lie on the boundary of the interval to which they are constrained, and therefore on the boundary of their confidence intervals.

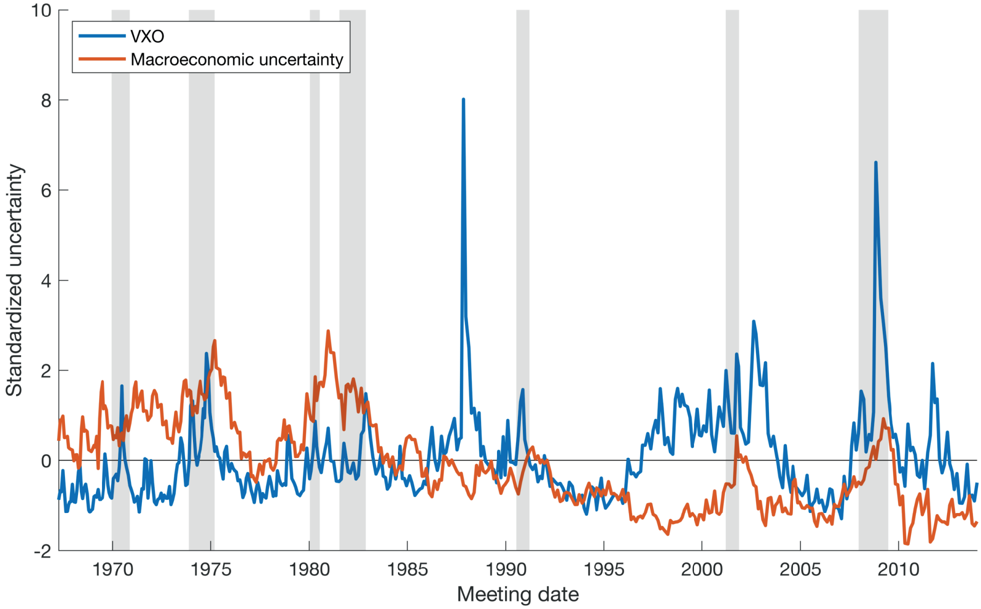

Figure 2.2: Financial uncertainty proxy (VXO index) and entropy measure of macroeconomic uncertainty. Both series are standardized to have zero mean and unit standard deviation. The grey bars indicate NBER recessions.

The largest deviations between the model and Greenbook forecasts occur in periods of rapid, large change in the economic outlook, especially for the output growth forecasts. This is understandable given the fact that the VAR models the economy as a mean-reverting process, which is not a good model for those periods.

In Figure 2.2, we plot the resulting entropy-based macroeconomic uncertainty measure, together with the VXO index, which we use as financial uncertainty proxy. Figure A.1 in the appendix plots the uncertainties for the individual series included in the TVP-VAR. Note that although our baseline sample for the policy reaction estimation runs from January 1969 to October 2008, we calculate macroeconomic uncertainty on the whole Greenbook sample, spanning the period March 1967 – December 2013. Several characteristics stand out.

First, and most importantly for our analysis, there seems to be no clear relationship between the two types of uncertainty. The correlation coefficient between financial and macroeconomic uncertainty is small (-0.06) and insignificant ($p = 0.18$).[20] This helps us disentangle the role that both types of uncertainty play in monetary policy.

Second, macro-uncertainty is persistent: A first-order autoregression yields an AR(1) coefficient of 0.96 ($p = 0$). The AR(1) coefficient for financial uncertainty is 0.66 ($p = 0$). Third, there are few times where macro-uncertainty peaks. Most prominent are the recession of 1974–75, the early 1980s recession, the 2001 recession, and the Great Recession of 2008–09. This leads to a fourth observation: Uncertainty tends to

[20]The correlation of the VXO with output growth uncertainty is -0.07 ($p = 0.15$), and with inflation uncertainty -0.12 ($p = 0.01$).

Table 2.1: Reaction function estimates on full sample using entropy uncertainty

	Linear			Regime		
	Coeff	SE	p	Coeff	SE	p
Intercept	−0.063	0.103	0.543	0.037	0.093	0.691
γ_b	−0.012	0.010	0.216	−0.013	0.010	0.218
$\sum_{i=-1}^{2} \gamma_{yi}$	0.053	0.020	0.007	0.055	0.020	0.006
$\sum_{i=-1}^{2} \gamma_{\pi i}$	0.055	0.016	0.001	0.051	0.017	0.003
γ_{u0}	−0.038	0.014	0.007	−0.042	0.015	0.005
$\sum_{i=-1}^{2} \delta_{yi}$	0.173	0.047	0.000	0.173	0.048	0.000
$\sum_{i=-1}^{2} \delta_{\pi i}$	0.018	0.076	0.808	0.024	0.075	0.751
ϕ_{VXO}	−0.043	0.015	0.005	−0.044	0.015	0.004
ϕ_U	−0.074	0.034	0.032	−0.137	0.059	0.021
N		360			360	
RMSE		0.344			0.345	
$\overline{R}^2$		0.261			0.258	

Note. The full sample spans FOMC meetings in the period 14 January 1969 – 29 October 2008. The left panel presents estimates for our reaction function (2.8) with the entropy uncertainty measure: $U_m = U_{Mm}$. The right panel presents estimates with the probability of being in a high-uncertainty regime given by (2.9): $U_m = F(U_{Mm})$. The VXO and entropy uncertainty series are standardized prior to estimation. Reported standard errors and p-values are robust to heteroscedasticity and autocorrelation, estimated with Bartlett kernel and data-driven bandwidth estimated with AR(1) model by maximum likelihood (Andrews, 1991).

rise before recessions, and peak during or shortly after them.

Lastly, the figure shows a downward trend in macroeconomic uncertainty, with most observations above the mean in the first half of the sample. This is partly explained by the decreasing volatility of the economic time series, and partly by the constant-coefficient nature of the underlying model ($\lambda = 1$): Every additional observation gives new information about the fixed coefficients, thereby lowering the uncertainty surrounding the coefficient estimates. As the real-time sample grows over time, uncertainty decreases.

2.3.2 Monetary policy response to uncertainty

In Table 2.1, we present our full sample estimates for the monetary policy reaction function with the entropy measure of macroeconomic uncertainty (U_{Mm}). Both the VXO index and the macro-uncertainty measure are standardized to have zero mean and unit standard deviation.[21] This means that their respective coefficients measure the policy response in terms of standard deviations.

For clarity, we only report the net total effect for some groups of variables.[22] For example, we report for the level of the output growth forecasts that the sum of their

[21] We standardize over the whole Greenbook sample, so that the series used in the estimation are the same as those plotted in Figure 2.2.

[22] Full results are available upon request.

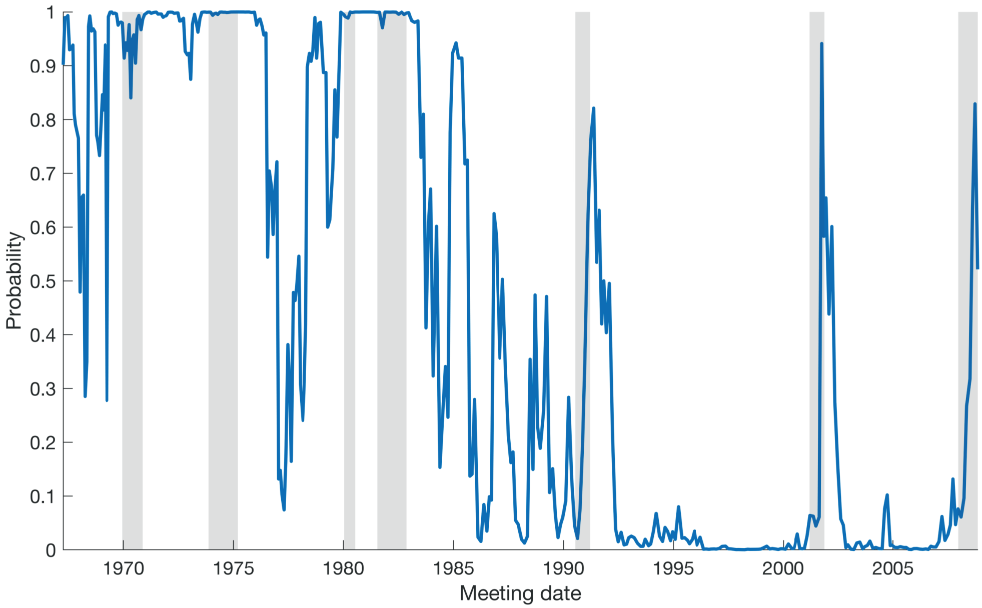

Figure 2.3: Probability of being in a high-uncertainty regime, as defined in terms of entropy uncertainty. The grey bars indicate NBER recessions.

coefficients (γ_{yi}) over the different horizons (previous quarter, current quarter, one-quarter-ahead, two-quarters-ahead) is 0.053. The sum of the coefficients for revisions to the output growth forecasts is 0.173. This means that the net effect of a one percentage point increase compared to the previous meeting in the output growth forecast for each horizon, is an increase in the intended funds rate of about 23 basis points.

Starting with the left panel of Table 2.1, we see that the policy reaction to financial uncertainty, as measured by the VXO index, is negative and significant at the 1% level. The effect size is about 4 basis points per standard deviation. The reaction to macroeconomic uncertainty is also significantly negative (at the 5% level), with a one-standard-deviation increase leading to an intended fall in the funds rate of 7.4 basis points. This is comparable to the response to an increase in the unemployment rate of two percentage points.

The right panel of Table 2.1 shows our estimates when we replace the macro-uncertainty measure in the reaction function by the probability of being in a high-uncertainty regime ($U_m = F(U_{Mm})$). In Figure 2.3, we plot the evolution of that probability over the sample. Most high-uncertainty episodes lie in the first half of the sample. The reaction to the VXO index is virtually the same in the specification with regime probabilities. Furthermore, the estimates indicate that the intended interest rate lies almost 14 basis points lower in the high-uncertainty regime compared to the low-uncertainty regime. This is relatively large compared to the most common policy move of 25 basis points.

Because our entropy measure captures multiple components of macroeconomic

Table 2.2: Reaction function estimates on full sample using output growth uncertainty

	Linear			Regime		
	Coeff	SE	p	Coeff	SE	p
Intercept	−0.001	0.089	0.991	0.060	0.090	0.505
γ_b	−0.014	0.010	0.153	−0.013	0.010	0.191
$\sum_{i=-1}^{2} \gamma_{yi}$	0.053	0.020	0.009	0.056	0.020	0.005
$\sum_{i=-1}^{2} \gamma_{\pi i}$	0.053	0.016	0.001	0.055	0.017	0.002
γ_{u0}	−0.046	0.015	0.002	−0.045	0.015	0.003
$\sum_{i=-1}^{2} \delta_{yi}$	0.176	0.047	0.000	0.173	0.048	0.000
$\sum_{i=-1}^{2} \delta_{\pi i}$	0.018	0.075	0.814	0.021	0.075	0.777
ϕ_{VXO}	−0.044	0.015	0.004	−0.043	0.015	0.005
ϕ_U	−0.067	0.038	0.080	−0.169	0.066	0.011
N		360			360	
RMSE		0.344			0.344	
$\overline{R}^2$		0.261			0.261	

Note. The full sample spans FOMC meetings in the period 14 January 1969 – 29 October 2008. The left panel presents estimates for our reaction function (2.8) with the output growth uncertainty measure: $U_m = U_{ym}$. The right panel presents estimates with the probability of being in a high-uncertainty regime given by equation (2.9): $U_m = F(U_{ym})$. The VXO and output growth uncertainty series are standardized prior to estimation. Reported standard errors and p-values are robust to heteroscedasticity and autocorrelation, estimated with Bartlett kernel and data-driven bandwidth estimated with AR(1) model by maximum likelihood (Andrews, 1991).

uncertainty, we study whether the Fed reacts differently to each of those components. To investigate this possibility, we look at the response to output growth and inflation uncertainty separately.

We start with the reaction function that includes the output growth uncertainty measure (U_{ym}). The full sample estimates (Table 2.2) show that the linear relationship between macro-uncertainty and the intended funds rate is significant at the 10% level. This may partly be explained by the logarithmic transformation in the calculation of the entropy uncertainty measure. This transformation dampens the high uncertainty peaks in the 1970s and 1980s compared to those later in the sample. Something similar happens with the calculation of the regime probabilities, which might explain why the regime specification still leads to a significant response. In fact, the effect is even larger than for the entropy-based probabilities: being in a high-uncertainty regime is associated with a decrease in the intended funds rate of almost 17 basis points.

We present the subsample estimates in Table 2.3. In the Greenspan-Bernanke sample, a high-uncertainty regime is associated with a highly significant decrease in the intended funds rate of 9 basis points. The response is the same as that to a one-percentage-point increase in the unemployment rate forecast for the current quarter. The response to uncertainty is insignificant on the earlier (Martins-Burns-Miller) sample.

Lastly, we discuss the estimates that use the inflation uncertainty measure ($U_{\pi m}$) in the reaction function. On the full and Martins-Burns-Miller sample, the responses

Table 2.3: Reaction function estimates on subsamples using output growth uncertainty

	Linear			Regime		
	Coeff	SE	p	Coeff	SE	p
Martins-Burns-Miller sample: 14 January 1969 – 11 July 1979						
Intercept	−0.158	0.366	0.668	−0.136	0.367	0.712
γ_b	0.002	0.033	0.960	0.003	0.033	0.936
$\sum_{i=-1}^{2} \gamma_{yi}$	0.068	0.024	0.005	0.067	0.023	0.004
$\sum_{i=-1}^{2} \gamma_{\pi i}$	0.048	0.031	0.125	0.043	0.033	0.198
γ_{u0}	−0.060	0.047	0.206	−0.058	0.047	0.225
$\sum_{i=-1}^{2} \delta_{yi}$	0.089	0.067	0.187	0.090	0.065	0.167
$\sum_{i=-1}^{2} \delta_{\pi i}$	0.218	0.095	0.024	0.227	0.094	0.017
ϕ_{VXO}	−0.070	0.040	0.081	−0.068	0.041	0.098
ϕ_U	−0.000	0.025	0.999	−0.032	0.065	0.621
N		123			123	
RMSE		0.241			0.240	
$\overline{R}^2$		0.321			0.323	
Greenspan-Bernanke sample: 18 August 1987 – 29 October 2008						
Intercept	0.084	0.174	0.630	0.106	0.160	0.508
γ_b	−0.070	0.020	0.001	−0.071	0.018	0.000
$\sum_{i=-1}^{2} \gamma_{yi}$	0.106	0.019	0.000	0.107	0.018	0.000
$\sum_{i=-1}^{2} \gamma_{\pi i}$	0.200	0.044	0.000	0.213	0.041	0.000
γ_{u0}	−0.090	0.027	0.001	−0.091	0.026	0.001
$\sum_{i=-1}^{2} \delta_{yi}$	0.030	0.030	0.313	0.029	0.028	0.293
$\sum_{i=-1}^{2} \delta_{\pi i}$	−0.050	0.066	0.445	−0.058	0.065	0.373
ϕ_{VXO}	−0.028	0.013	0.029	−0.026	0.012	0.038
ϕ_U	−0.019	0.014	0.187	−0.091	0.033	0.006
N		171			171	
RMSE		0.153			0.151	
$\overline{R}^2$		0.526			0.541	

Note. The left panel presents estimates for our reaction function (2.8) with the output growth uncertainty measure: $U_m = U_{ym}$. The right panel presents estimates with the probability of being in a high-uncertainty regime given by equation (2.9): $U_m = F(U_{ym})$. The VXO and output growth uncertainty series are standardized on each subsample prior to estimation. Reported standard errors and p-values are robust to heteroscedasticity and autocorrelation, estimated with Bartlett kernel and data-driven bandwidth estimated with AR(1) model by maximum likelihood (Andrews, 1991). The sample dates indicate the first and last FOMC meeting in the sample.

Table 2.4: Reaction function estimates on full sample using inflation uncertainty

	Linear			Regime		
	Coeff	SE	p	Coeff	SE	p
Intercept	−0.020	0.096	0.837	0.035	0.094	0.711
γ_b	−0.013	0.010	0.182	−0.012	0.010	0.240
$\sum_{i=-1}^{2} \gamma_{yi}$	0.052	0.020	0.010	0.057	0.020	0.005
$\sum_{i=-1}^{2} \gamma_{\pi i}$	0.053	0.016	0.001	0.055	0.018	0.003
γ_{u0}	−0.043	0.014	0.003	−0.044	0.015	0.004
$\sum_{i=-1}^{2} \delta_{yi}$	0.175	0.047	0.000	0.172	0.048	0.000
$\sum_{i=-1}^{2} \delta_{\pi i}$	0.014	0.076	0.857	0.019	0.075	0.797
ϕ_{VXO}	−0.046	0.016	0.004	−0.046	0.016	0.004
ϕ_U	−0.062	0.043	0.143	−0.159	0.072	0.028
N		360			360	
RMSE		0.345			0.345	
$\overline{R}^2$		0.258			0.258	

Note. The full sample spans FOMC meetings in the period 14 January 1969 – 29 October 2008. The left panel presents estimates for our reaction function (2.8) with the inflation uncertainty measure: $U_m = U_{\pi m}$. The right panel presents estimates with the probability of being in a high-uncertainty regime given by equation (2.9): $U_m = F(U_{\pi m})$. The VXO and inflation uncertainty series are standardized prior to estimation. Reported standard errors and p-values are robust to heteroscedasticity and autocorrelation, estimated with Bartlett kernel and data-driven bandwidth estimated with AR(1) model by maximum likelihood (Andrews, 1991).

are qualitatively similar as those using output growth uncertainty (see Table 2.4 and Table 2.5). On the Greenspan-Bernanke sample however, the effect of being in a high-uncertainty regime, as measured by inflation uncertainty, is statistically indistinguishable from zero (see Table 2.5).[23] The significant response to output growth uncertainty may be muted by the inflation component of the entropy measure.[24]

2.3.3 Robustness

Most of our results are robust to deviations from our benchmark specification along various dimensions. The most sensitive result is the linear response to entropy uncertainty on the full sample.

First, we investigate the effect of changing the TVP-VAR hyperparameters. Specifically, we recalculate macroeconomic uncertainty for values of the hyperparameters that lie on a 3×3×3-grid:

$$\lambda \in \{0.98, 0.99, 1\}, \quad \kappa \in \{0.94, 0.96, 0.98\}, \quad \gamma \in \{0.01, 0.05, 0.1\}.$$

[23] We compare the high-uncertainty probabilities based on output growth uncertainty with those based on inflation uncertainty for the subsamples in Figure A.3 in the appendix.

[24] Using interest rate uncertainty (the third component of the entropy measure) to calculate the probability of being in a high-uncertainty regime, estimates indicate that being in that regime leads to an intended decrease in the interest rate of almost 10 basis points ($p = 0.005$).

Table 2.5: Reaction function estimates on subsamples using inflation uncertainty

	Linear			Regime		
	Coeff	SE	p	Coeff	SE	p
Martins-Burns-Miller sample: 14 January 1969 – 11 July 1979						
Intercept	−0.158	0.366	0.667	−0.139	0.366	0.704
γ_b	0.002	0.033	0.964	0.002	0.033	0.943
$\sum_{i=-1}^{2} \gamma_{yi}$	0.068	0.024	0.005	0.067	0.023	0.004
$\sum_{i=-1}^{2} \gamma_{\pi i}$	0.048	0.031	0.121	0.044	0.033	0.186
γ_{u0}	−0.060	0.047	0.204	−0.058	0.048	0.226
$\sum_{i=-1}^{2} \delta_{yi}$	0.088	0.067	0.191	0.090	0.065	0.167
$\sum_{i=-1}^{2} \delta_{\pi i}$	0.217	0.095	0.024	0.225	0.094	0.019
ϕ_{VXO}	−0.071	0.040	0.078	−0.068	0.040	0.095
ϕ_U	0.002	0.025	0.926	−0.030	0.064	0.635
N		123			123	
RMSE		0.241			0.240	
$\overline{R}^2$		0.321			0.322	
Greenspan-Bernanke sample: 18 August 1987 – 29 October 2008						
Intercept	0.050	0.160	0.755	0.064	0.167	0.701
γ_b	−0.065	0.018	0.000	−0.065	0.018	0.000
$\sum_{i=-1}^{2} \gamma_{yi}$	0.110	0.018	0.000	0.111	0.018	0.000
$\sum_{i=-1}^{2} \gamma_{\pi i}$	0.188	0.043	0.000	0.185	0.041	0.000
γ_{u0}	−0.085	0.025	0.001	−0.087	0.026	0.001
$\sum_{i=-1}^{2} \delta_{yi}$	0.026	0.030	0.400	0.025	0.030	0.416
$\sum_{i=-1}^{2} \delta_{\pi i}$	−0.050	0.066	0.448	−0.050	0.066	0.450
ϕ_{VXO}	−0.032	0.013	0.013	−0.032	0.012	0.010
ϕ_U	−0.002	0.017	0.904	0.009	0.039	0.818
N		171			171	
RMSE		0.154			0.154	
$\overline{R}^2$		0.521			0.522	

Note. The left panel presents estimates for our reaction function (2.8) with the output growth uncertainty measure: $U_m = U_{\pi m}$. The right panel presents estimates with the probability of being in a high-uncertainty regime given by equation (2.9): $U_m = F(U_{\pi m})$. The VXO and inflation uncertainty series are standardized on each subsample prior to estimation. Reported standard errors and p-values are robust to heteroscedasticity and autocorrelation, estimated with Bartlett kernel and data-driven bandwidth estimated with AR(1) model by maximum likelihood (Andrews, 1991). The sample dates indicate the first and last FOMC meeting in the sample.

These values are reasonable in terms of what Koop and Korobilis (2013) find to be optimal hyperparameters for a small monetary VAR (see footnote 14). The only result that is relatively sensitive to changes in the hyperparameters is the linear response to entropy uncertainty on the full sample. It is significant at the 5% level only for $\lambda = 1$ (the constant-parameter version of the TVP-VAR). However, the full-sample response to being in a high-uncertainty regime is significant at the 5% level across all hyperparameter values that we consider. The effect sizes vary between 9 and 14 basis points, where the effect gets larger as λ approaches 1. The estimates for the Martins-Burns-Miller sample are consistent across hyperparameters, as are those for the Greenspan-Bernanke sample. The size of the response to output growth uncertainty on the latter sample varies between 7 and 10 basis points, where the largest effects correspond to $\lambda = 1$.

Second, we vary the parameter θ, which determines the shape of the probability function in equation (2.9). We focus on the regime estimates that use entropy uncertainty on the full sample and that use output growth uncertainty on the Greenspan-Bernanke sample. For values $\theta \in \{1, 5, 10, 25\}$, we find that the results are largely consistent. The estimated full-sample coefficients for the entropy-based probability variable lie between -0.33 and -0.099, with p-values between 0.020 and 0.033. The coefficients for the Greenspan-Bernanke sample (based on output growth uncertainty) lie between -0.15 and -0.069, with p-values between 0.006 and 0.033. The effect sizes get smaller as θ gets larger. We also consider the special case of a dummy that equals one when the uncertainty is above its median (as computed on the respective sample), and zero otherwise.[25] This case gives coefficients of -0.083 ($p = 0.043$) on the full sample, and -0.066 ($p = 0.012$) on the Greenspan-Bernanke sample.

Third, we exclude the Greenbook nowcast from the real-time vintages, meaning that we only use data up to and including the quarter before the meeting takes place. We calculate macroeconomic uncertainty using the posterior density of nowcasts for the quarter the meeting takes place, instead of using the density of forecasts for next quarter. The most notable difference with the baseline results is that the full-sample estimated responses (both linear and regime) to entropy uncertainty are only significant at the 10% level. However, the response to being in a high-uncertainty regime, as calculated using output growth uncertainty, is still significant at the 5% level on the full sample. The other results are in line with our baseline estimates.

2.4 Conclusion

Using a new, Bayesian, real-time measure of macroeconomic uncertainty, we find for the period 1969 – 2008 that monetary policy responds to macroeconomic uncertainty with a significant decrease in the intended federal funds rate. Specifically, being in a high-uncertainty regime leads to a decrease in the intended funds rate of about 14 basis points.

We split the sample into subsamples before and after the period of non-borrowed reserves targeting, and zoom in on different components of macroeconomic uncertainty. These estimates also indicate that a specification with low- versus high-uncertainty regimes best captures the relationship between uncertainty and mone-

[25]This corresponds to the case $\theta \to \infty$, apart from the times when the uncertainty measure exactly equals the median, in which case the limit equals 1/2.

tary policy. On the full sample, the effect of being in a high uncertainty regime also holds when measured by output growth uncertainty or inflation uncertainty. The results furthermore imply that the reaction function of the Fed changes over time. In particular, macroeconomic uncertainty only plays a significant role in the later subsample, covering Greenspan's and part of Bernanke's tenure as chairman. Being in a high-output-growth-uncertainty regime leads to a decrease in the intended funds rate of 9.1 basis points. We find no significant funds rate response to uncertainty in the period 1969 – 1979.

These responses to macroeconomic uncertainty are orthogonal to the effect of financial uncertainty, proxied by the VXO index. On the full 1969 – 2008 period, as well as on the later subsample, an increase in the VXO index also leads to a significant decrease in the intended funds rate.

Our measure could further be used to investigate whether it is optimal for monetary policy to respond to uncertainty. One approach would be to revisit Romer and Romer's (2004) analysis of monetary policy effectiveness, distinguishing between macroeconomic uncertainty regimes and taking into account the policy response to uncertainty. Such a study would be similar to the one by Falck et al. (2019) in the context of disagreement about inflation expectations. One could also use our new uncertainty measure to contribute to the stream of literature that looks at uncertainty and its impact at the macro level, which has taken off after the seminal study of Bloom (2009). If macroeconomic uncertainty plays a role in recessions, as some studies in this stream have suggested, it would make sense for monetary policy to respond accordingly.

Chapter 3

Diverse risk preferences and heterogeneous expectations in an asset pricing model[1]

Abstract

We propose a heuristic switching model of an asset market where the agents' choice of heuristic is consistent with their individual risk aversion. They choose between a fundamentalist and a trend-following rule to form expectations about the price of a risky asset. Given their risk aversion, agents make a deterministic trade-off between mean and variance both in choosing a forecasting heuristic and determining the number of risky assets to buy. Heterogeneous risk preferences can lead to diverse choices of heuristic. Using empirical estimates for the distribution of risk aversion, simulations show that the resulting time-varying heterogeneity of expectations can give rise to chaotic dynamics: irregular booms and busts in the asset price without exogenous shocks. Small, stochastic price shocks lead to larger asset price bubbles, and can make stable solutions explosive. We prove that a representative agent cannot capture our model.

3.1 Introduction

From historical examples like the Tulip Mania to more recent ones like the US housing bubble, it seems that bubbles are a recurring characteristic of asset markets. While it is hard to identify bubbles in real markets because fundamental values cannot be observed, experimental asset markets, with known fundamentals, can also exhibit price bubbles and crashes (e.g., Smith et al., 1988).

Asset price bubbles cannot always be explained by rational, speculative motives. Lei et al. (2001) show that bubbles can emerge even when speculation is not possible, suggesting that the behaviour can be inherently irrational. Hommes et al. (2005, 2008) specifically look at irrationality in expectations. In their learning-to-forecast

[1]This chapter is based on joint work with Giulia Piccillo. Thomas Gomez is the lead author, and has been the main contributor in all phases of this work, including idea generation, theory development, running simulations, results interpretation and writing the chapter. This chapter was supported by the Complex Systems Fund, with special thanks to Peter Koeze.

experiments, subjects are rewarded for accurately forecasting future price realizations of a risky asset. At the individual level, their results indicate that the subjects use different, simple rules to predict future asset prices. At the aggregate level, they find that the interaction between the forecasting rules used by the subjects can lead to different asset price dynamics in the same experimental setting. These dynamics include slow convergence to, and significant deviations from the fundamental.

Anufriev and Hommes (2012) show that these different price dynamics can be explained by modelling expectations with the heuristic switching framework introduced by Brock and Hommes (1997). In line with the learning-to-forecast results, this framework assumes that individuals choose between a number of simple forecasting rules (heuristics) to form their expectations. They base their choice on the past performance of those rules. As the performance changes over time, the individuals update their choice, switching between heuristics. The heuristic switching framework is in line with a large literature on boundedly rational, heterogeneous expectations, where heterogeneity changes over time. Mankiw et al. (2003), for example, show that consumers as well as professional forecasters have diverse inflation expectations, and Wieland and Wolters (2011) find time-varying heterogeneity in model forecasts. Heuristic switching models are able to explain certain stylized facts of financial markets (e.g., Anufriev and Panchenko, 2009; Gaunersdorfer et al., 2008; Schmitt, 2020). These models have also been applied in macroeconomics, investigating the effects of monetary policy (e.g., Hommes et al., 2019), fiscal policy (e.g., De Grauwe et al., 2019), and the financial sector (De Grauwe and Macchiarelli, 2015). Branch (2004) estimates a heuristic switching model using inflation forecasting data. His results indicate that consumers do indeed switch between forecasting strategies based on past performance.

In the context of this heuristic switching literature, our main contribution is to study the natural consequence of the standard asset pricing assumption that agents are mean-variance optimizers.[2] Specifically, they make a trade-off between forecasting performance and variability of that performance when choosing a heuristic. This means that risk aversion plays a key role in expectation formation: Agents with different risk preferences might choose a different heuristic, because their risk aversion determines the importance of performance variability for their choice. As the heuristics' performance changes over time, agents reconsider their choice. This can lead them to switch between forecasting rules. In this way, diverse risk preferences can lead to time-varying heterogeneity of expectations. We model this diversity by using risk preference estimates for the general population provided by Kimball et al. (2008) and Aarbu and Schroyen (2009).[3]

We argue that diverse risk preferences are a natural source of heterogeneous expectations for two reasons.[4] First, risk preferences play a central role in economic theory and have been linked empirically to a variety of outcomes related to, for ex-

[2]This assumption is central in modern portfolio theory, which goes back to Markowitz (1952). The benchmark for asset pricing models with heuristic switching (Brock and Hommes, 1998) also assumes that agents are mean-variance optimizers.

[3]The estimate by Kimball et al. (2008) is also used by Xiouros and Zapatero (2010) to empirically validate an asset pricing model, but assuming rational expectations.

[4]In a previous contribution to the heterogeneous switching literature, Pfajfar (2013) also pinpoints potential sources of expectations heterogeneity. He links heterogeneity to the computing capabilities and information sets of agents.

ample, career choice, financial decision making, and migration (Dohmen et al., 2012; Falk et al., 2018; Jaeger et al., 2010). Second, risk preferences are indeed heterogeneous, as shown by a large body of empirical research, consisting of surveys (e.g., Falk et al., 2018), experiments (e.g., Choi et al., 2007; von Gaudecker et al., 2011), and decision-making by actual market participants (Cohen and Einav, 2007; Paravisini et al., 2016). Diverse risk preferences have previously been used to explain differences in asset allocation across households (Kimball et al., 2008), and to generate a number of empirical regularities of stock market returns in a general equilibrium model (Chan and Kogan, 2002).

We have three main results. First, we analytically characterize the model's dynamics around its steady state, which corresponds to the fundamental price. We do this for the case that expectations heterogeneity is constant. This heterogeneity is measured by the fractions of aggregate risk tolerance that are represented by fundamentalists and trend-followers. These risk-tolerance fractions measure the importance of the fundamentalist and trend-following rules in the market, and play a central role in the asset price dynamics. We show that a threshold risk-tolerance fraction of fundamentalists exists that determines the stability of the steady state. This threshold depends on the strength of the trend-following rule relative to the risk-free rate. Depending on the combination of parameter values, the constant-fractions model can exhibit oscillatory or exponential asset price dynamics, where the price converges towards the fundamental if the steady state is stable, and otherwise diverges away from it.

Second, we show numerically that time-varying risk-tolerance fractions give rise to rich dynamics. For some parameter values, our model exhibits deterministic chaos. Chaos, which can only arise in non-linear dynamical systems, is characterized by sensitive dependence on initial conditions: Small perturbations in the initial state of the system blow up exponentially. This makes long-term forecasting difficult even when the laws governing the system are known. Another characteristic of deterministic chaos is endogenous variability: Even without stochastic shocks, irregular fluctuations can occur. In our model, chaos manifests itself in the form of irregular, unpredictable booms and busts in the deviations from the fundamental asset price (excessive volatility). To investigate the resilience of the system, we introduce noise traders, which leads to stochastic price shocks. We find that small price shocks can be amplified by the chaotic dynamics, turning deterministically stable solutions into explosive ones.

Third, we show that a representative agent aggregation does not exist for our model with heterogeneous agents. However, our mean-variance modelling of the choice of forecasting heuristic does allow for such an aggregation in the case of the benchmark asset pricing model with heuristic switching (Brock and Hommes, 1998), which does not include heterogeneous risk preferences. The representative agent uses a weighted average of two heuristics to form expectations, where the weights are determined by mean-variance maximization. We derive an analytical expression for the representative agents' time-varying risk aversion.

We highlight two previous contributions to the heuristic switching literature that have incorporated heterogeneous risk preferences. First, Park (2014) links the rate of risk aversion to the forecasting heuristic that is chosen. He assumes that fundamentalists have a constant rate of risk aversion, while chartists have a time-varying risk aversion that is inspired by prospect theory's reflection effect: Chartist's risk aversion increases (decreases) if the risky asset generates positive (negative) returns. Note that this approach implies that the heuristic that agents use to form expectations determines whether they are affected by the reflection effect, and that they change their risk preferences when they switch heuristics. Moreover, it does not take into account the effect of risk aversion on the choice of forecasting rule. We choose to model risk preferences as constant, and explicitly incorporate these in the heuristic choice process. This means that the focus is on differences in risk preferences and their implications for expectations, instead of time-varying risk preferences.

Second, Chiarella and He (2002, 2003) study a heuristic switching model where risk aversion is again linked to the chosen heuristic, but the risk aversion parameter for a given heuristic is constant. They study this model, which also incorporates learning, in two institutional settings: Chiarella and He (2002) use a Walrasian scenario, while Chiarella and He (2003) use a market-maker. They find that the diversity of risk preferences matters for the asset pricing dynamics. Like Park (2014), they do not take into account the implications of diverse risk preferences for expectation formation.

3.2 The model

3.2.1 An asset market with heuristic switching mean-variance optimizers

Following Brock and Hommes (1998) and Park (2014), we consider a discrete-time asset market that consists of a risk-free and a risky asset. The risk-free asset has perfectly elastic supply and pays a gross return $R^f > 1$. The risky asset has price p_t and pays a dividend d_t in period $t \in \mathbb{N}_0$. We denote by R^e_{t+1} the excess return per risky asset in period $t+1$: $R^e_{t+1} = p_{t+1} + d_{t+1} - R^f\, p_t$.

Agents differ only in their constant relative risk aversion (CRRA) parameter γ, drawn from a distribution defined on $(0, \infty)$. We assume a population size normalized to 1 and denote the density function by g, with cumulative distribution function G. Focusing on relative risk aversion as it appears in the CRRA utility function allows us to use the empirical estimates of its distribution that can be found in the literature (see subsection 3.2.3).

We give all variables that differ across agents a subscript γ to indicate their dependence on risk aversion. We write $z_{\gamma,t}$ for agents' real demand for the risky asset in period t. We let $\mathbb{I}_t = \{p_t, p_{t-1}, \ldots, d_t, d_{t-1}, \ldots\}$ be the information set in period t and denote agents' beliefs about expectation and variance, conditional on $\mathbb{I}_t$, by $\widetilde{E}_{\gamma,t}$ and $\widetilde{V}_{\gamma,t}$.

The agents are myopic mean-variance maximizers of their expected wealth:

$$U_{\gamma,t} = \widetilde{E}_{\gamma,t}[W_{\gamma,t+1}] - \frac{\gamma}{2}\,\widetilde{V}_{\gamma,t}[W_{\gamma,t+1}], \tag{3.1}$$

with period $t+1$ wealth $W_{\gamma,t+1} = z_{\gamma,t}\, R^e_{t+1} + R^f\, W_{\gamma,t}$. Note that we make the simplifying assumption that the value of risk aversion that appears in the mean-variance utility is the CRRA parameter γ. This assumption is supported by Ang (2014), who argues that mean-variance and CRRA utility are closely related, and even converge under certain conditions.[5]

[5]We could also explicitly consider CRRA utility maximization, but mean-variance optimization is more tractable, and is more in line with the idea of boundedly rational agents using simple forecasting strategies.

We have that[6]

$$\widetilde{E}_{\gamma,t}[W_{\gamma,t+1}] = z_{\gamma,t}\, \widetilde{E}_{\gamma,t}[R^e_{t+1}] + R^f\, W_{\gamma,t} \tag{3.2a}$$

$$\widetilde{V}_{\gamma,t}[W_{\gamma,t+1}] = z^2_{\gamma,t}\, \widetilde{V}_{\gamma,t}[R^e_{t+1}]. \tag{3.2b}$$

Since $U_{\gamma,t}$ is concave in $z_{\gamma,t}$, the first-order condition gives the risky asset demand that maximizes utility:

$$z_{\gamma,t} = \frac{\widetilde{E}_{\gamma,t}\left[R^e_{t+1}\right]}{\gamma\, \widetilde{V}_{\gamma,t}\left[R^e_{t+1}\right]}. \tag{3.3}$$

Aggregate real demand for the risky asset in period t, denoted by Z_t, is then obtained by integrating over the population with density g:

$$Z_t = \int_0^\infty \frac{\widetilde{E}_{\gamma,t}\left[R^e_{t+1}\right]}{\gamma\, \widetilde{V}_{\gamma,t}\left[R^e_{t+1}\right]}\, g(\gamma)\, d\gamma. \tag{3.4}$$

In line with previous literature (e.g., Brock and Hommes, 1998; Park, 2014), we assume that agents have homogeneous beliefs on variance ($\widetilde{V}_{\gamma,t}\left[R^e_{t+1}\right] = \widetilde{V}_t\left[R^e_{t+1}\right]$) and that the net supply of the risky asset is zero. The former assumption can be justified by previous studies that indicate that conditional variances are easier to estimate than conditional means (e.g., Nelson, 1992). The latter assumption means that no shares in the risky asset are issued or withdrawn: A fixed number of shares is traded in the market. These assumptions increase the tractability of the model.

Equating aggregate demand and supply, $Z_t = 0$, gives an expression for the risky asset price in period t:

$$p_t = \frac{1}{R^f\, \Theta} \int_0^\infty \frac{g(\gamma)}{\gamma}\, \widetilde{E}_{\gamma,t}\left[p_{t+1} + d_{t+1}\right]\, d\gamma, \tag{3.5}$$

where we have defined the aggregate risk tolerance

$$\Theta = \int_0^\infty \frac{g(\gamma)}{\gamma}\, d\gamma. \tag{3.6}$$

Risk tolerance θ is the reciprocal of risk aversion: $\theta = \gamma^{-1}$.

The fundamental price for the risky asset, denoted by p^*_t, is the price that would arise in a homogeneous, rational market. We denote expectation and variance, conditional on information set $\mathbb{I}_t$, by E_t and V_t. Equation (3.5) implies that

$$p^*_t = \frac{E_t\left[p^*_{t+1} + d_{t+1}\right]}{R^f}. \tag{3.7}$$

[6]In line with previous literature (e.g., Brock and Hommes, 1998; Park, 2014), we assume that agents' beliefs about conditional expectation and variance have certain properties in common with the standard expectation and variance operators:

$$\widetilde{E}_{\gamma,t}\left[a\, X_{t+p+1} + b\, Y_{t+q+1} + c\, Z_{t-r} + d\right] = a\, \widetilde{E}_{\gamma,t}\left[X_{t+p+1}\right] + b\, \widetilde{E}_{\gamma,t}\left[Y_{t+q+1}\right] + c\, Z_{t-r} + d;$$

$$\widetilde{V}_{\gamma,t}\left[a\, X_{t+p+1} + b\, Y_{t-q} + c\right] = a^2\, \widetilde{V}_{\gamma,t}\left[X_{t+p+1}\right],$$

with X, Y, Z stochastic processes, $a, b, c, d \in \mathbb{R}$, and $p, q, r \in \mathbb{N}_0$.

In the case of an i.i.d. dividend process with mean μ_d, for example, we have a single non-explosive solution given by $p^* = \mu_d/(R^f - 1)$, a constant. This fundamental price corresponds to the present value of the expected future dividend stream, in accordance with the dividend discount model. In the remainder of this section, we describe the the system in terms of price deviations from the fundamental $x_t = p_t - p_t^*$.

Following Brock and Hommes (1998) and Park (2014), we assume that agents form beliefs about deviations from the fundamental. First, all agents have rational beliefs about the fundamental price and dividend process:

$$\widetilde{E}_{\gamma,t}\left[p_{t+1}^* + d_{t+1}\right] = E_t\left[p_{t+1}^* + d_{t+1}\right]. \tag{3.8}$$

This is reasonable in the case of a simple dividend process, like the i.i.d. example that we discussed earlier. All possible fundamental price and dividend processes can be captured by the model however, since we do not have to specify these processes to model the deviations from the fundamental.

Second, the agents form beliefs about future price deviations from the fundamental by using a forecasting heuristic. These heuristics are simple rules that predict price deviations based on a specific number of observed previous deviations. Formally, the heuristics are indexed by the finite set $\mathcal{I}$, and for $s \in \mathcal{I}$, we define a heuristic as a function $h_s : \mathbb{R}^L \to \mathbb{R}$, $(x_{t-1}, \ldots, x_{t-L}) \mapsto h_s(x_{t-1}, \ldots, x_{t-L})$, where $L \in \mathbb{N}$ indicates the number of lags of price deviations that are taken into account.[7] In period t, agents of type γ use heuristic $s_{\gamma,t} \in \mathcal{I}$, and their beliefs about price deviations from the fundamental are given by

$$\widetilde{E}_{\gamma,t}\left[x_{t+1}\right] = h_{s_{\gamma,t}}\left(x_{t-1}, \ldots, x_{t-L}\right). \tag{3.9}$$

An example is the naive heuristic, which uses only one lag, and predicts that next period's price deviation from the steady state will be the same as last period's deviation: $h(x_{t-1}) = x_{t-1}$. The agents decide which heuristic to use based on their past performance. This will be discussed in more detail below.

Together, equations (3.8) and (3.9) pin down the agents' beliefs about the gross return per risky asset: $\widetilde{E}_{\gamma,t}\left[p_{t+1} + d_{t+1}\right] = \widetilde{E}_{\gamma,t}\left[p_{t+1}^* + d_{t+1}\right] + \widetilde{E}_{\gamma,t}\left[x_{t+1}\right]$. This means that we can solve for the price deviations from the fundamental using equations (3.5) and (3.7):

$$x_t = \frac{1}{R^f\,\Theta}\int_0^\infty \frac{g(\gamma)}{\gamma}\, h_{s_{\gamma,t}}\left(x_{t-1}, \ldots, x_{t-L}\right)\, d\gamma. \tag{3.10}$$

Usually in the heuristic switching literature, a central role is played by the fractions of agents that use the different heuristics. In our set-up with diverse risk preferences however, this role is played by the fraction of aggregate risk tolerance that is represented by the agents that use the heuristics. This is explained by the fact that agents with a larger risk tolerance buy and sell more of the risky asset, because of the inverse relationship between risky asset demand and risk aversion (see equation 3.3).

We define $\Gamma_{s,t}$ as the set of risk aversion rates for which agents use heuristic s in period t, and $\Theta_{s,t}$ as the fraction of aggregate risk tolerance (defined in equation 3.6)

[7] Note that these heuristics are functions of past asset price deviations only. Omitting period t deviations in forming beliefs about period $t+1$ is in accordance with previous studies in the heuristic switching literature (e.g., Brock and Hommes, 1998; Park, 2014). It is possible to include x_t in the specification of h_s, but this leads to simultaneity issues where x_t is undefined in some cases.

represented by agents that use heuristic s in period t:

$$\Theta_{s,t} = \Theta^{-1} \int_{\Gamma_{s,t}} \frac{g(\gamma)}{\gamma}\, d\gamma. \tag{3.11}$$

Note that these fractions add up to one: $\sum_{s\in\mathcal{I}} \Theta_{s,t} = 1$.

Now we can rewrite equation (3.10) as

$$x_t = \frac{1}{R^f} \sum_{s\in\mathcal{I}} \Theta_{s,t}\, h_s\left(x_{t-1}, \ldots, x_{t-L}\right). \tag{3.12}$$

We can interpret $\Theta_{s,t}$ as a measure of the importance of heuristic s in the market in period t. In the extreme case that $\Theta_{s,t} = 0$, no one is using it, and it has no impact on the price. At the other extreme, when $\Theta_{s,t} = 1$ everyone is using the heuristic, and it fully determines the price. If $\Theta_{s,t} = \frac{1}{2}$ for example, its impact on the asset price is the same as that of all other heuristics combined.

In the following, we explain how the agents choose a forecasting heuristic in each period. The agents are mean-variance maximizers, who therefore base their choice on the past performance of the heuristics, and on the variability of that performance. Because their rate of risk aversion determines the weight that the variability carries in their decision, agents with different risk preferences may use different heuristics in a given period.

The agents measure the performance of heuristic s in period t, denoted by $u_{s,t}$, in terms of their squared forecasting errors:

$$u_{s,t} = -\left[x_t - h_s(x_{t-2}, \ldots, x_{t-L-1})\right]^2. \tag{3.13}$$

They calculate a weighted average $\langle u_s\rangle_t$, a weighted squared average $\langle u_s^2\rangle_t$, and a weighted variance $\widetilde{\sigma}^2_{s,t}$ of the observed performance of heuristic s, where past performance is weighted by a memory parameter η that satisfies $0 < \eta < 1$:[8]

$$\langle u_s\rangle_t = \eta\, \langle u_s\rangle_{t-1} + (1-\eta)\, u_{s,t} \tag{3.14a}$$

$$\left\langle u_s^2\right\rangle_t = \eta\, \left\langle u_s^2\right\rangle_{t-1} + (1-\eta)\, u_{s,t}^2 \tag{3.14b}$$

$$\widetilde{\sigma}^2_{s,t} = \left\langle u_s^2\right\rangle_t - \langle u_s\rangle_t^2. \tag{3.14c}$$

The weighted variance satisfies $\widetilde{\sigma}^2_{s,t} \geq 0$ if $\widetilde{\sigma}^2_{s,0} \geq 0$.[9] The average and variance are combined in a mean-variance performance measure, defined as

$$\Psi_{\gamma,s,t} = \langle u_s\rangle_t - \frac{\gamma}{2}\, \widetilde{\sigma}^2_{s,t}. \tag{3.15}$$

For simplicity, we assume that we can use the same risk aversion coefficient γ as in the mean-variance utility in terms of their wealth: The same agents have the same risk aversion in these two contexts.

[8] Many heuristic switching studies use such a weighted average (e.g., Brock and Hommes, 1998). The weighted variance is a natural extension.

[9] A proof is available upon request. One can also show that this is a sensible definition of variance, in the sense that it is consistent with the definition of weighted average.

At the end of each period, the agents choose the heuristic with the largest mean-variance performance measure to use in the next period. This means that the risk-tolerance fraction represented by heuristic s in period t, $\Theta_{s,t}$, will be determined by the values of the mean-variance performance measure in period $t-1$: An agent with risk aversion γ uses the heuristic in period t if $\Psi_{\gamma,s,t-1} > \Psi_{\gamma,r,t-1}$ for all $r \in \mathcal{I} \setminus \{s\}$. We assume that when two or more rules have the same mean-variance utility, agents are equally likely to choose any one of those rules.

3.2.2 Fundamentalists and momentum traders

Following previous literature (Anufriev and Panchenko, 2009; Gaunersdorfer et al., 2008), and in line with empirical evidence (e.g., Chiarella et al., 2014), we consider the case in which agents choose between a fundamentalist and a trend-following (chartist) rule. We focus on just two rules to keep the model simple, but one could extend it to include more rules.

The fundamentalist heuristic predicts that the price of the risky asset converges to the fundamental, where the speed of convergence is measured by the fundamentalist parameter f, which satisfies $0 \leq f < 1$:

$$h_F(x_{t-1}) = f\, x_{t-1}. \tag{3.16}$$

The smaller f, the faster the asset price will converge. In the special case that $f = 0$ (the usual choice in the heuristic switching literature), the rule predicts that any deviations from the fundamental will disappear in the next period.

The trend-following heuristic, also referred to as momentum rule, assumes that price deviations from the fundamental follow a trend. It predicts that the deviation in the next period equals the one observed in the previous period, corrected for the last observed change in price deviations. A momentum parameter m, which satisfies $m > 0$, determines the strength of the correction:

$$h_M(x_{t-1}, x_{t-2}) = x_{t-1} + m\,(x_{t-1} - x_{t-2})\,. \tag{3.17}$$

The larger m, the stronger the past momentum will be extrapolated into future deviations from the fundamental.[10]

To analyse the dynamics of the price of the risky asset, we have to derive the risk-tolerance fractions represented by the two heuristics. Because the fractions add up to 1, the fundamentalist fraction $\Theta_{F,t}$ will also give us the fraction of momentum traders: $\Theta_{M,t} = 1 - \Theta_{F,t}$.

Looking at the definition of the mean-variance performance measure for choosing between the heuristics (3.15), we see that the choice in period t is independent of risk aversion if the variances of the rules' performance are equal ($\widetilde{\sigma}_{F,t} = \widetilde{\sigma}_{M,t}$). This means that all agents use the same rule in period $t+1$: If the fundamentalist rule has performed better on average ($\langle u_F \rangle_t > \langle u_M \rangle_t$), everyone uses the fundamentalist rule in period $t+1$ and $\Theta_{F,t+1} = 1$. On the other hand, if the momentum rule has performed better on average, we have that $\Theta_{F,t+1} = 0$. In the case that the heuristics

[10]The trend-following rule is often defined without reference to the fundamental price. We choose to define it in terms of price deviations from the fundamental however, to avoid the need to explicitly specify a fundamental price process. Note that the two definitions are equivalent when the fundamental price is constant.

have performed equally well on average, the agents are evenly divided between the two heuristics in period $t+1$, and $\Theta_{F,t+1} = \frac{1}{2}$.

If the weighted variances of the heuristics' performance are not equal ($\widetilde{\sigma}_{F,t} \neq \widetilde{\sigma}_{M,t}$), risk aversion does play a role in the choice between forecasting rules. The risk aversion coefficient for which the mean-variance performance of the forecasting heuristics are equal in period t, denoted by $\overline{\gamma}_t$, is given by

$$\overline{\gamma}_t = 2\,\frac{\langle u_F\rangle_t - \langle u_M\rangle_t}{\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t}}. \tag{3.18}$$

If $\widetilde{\sigma}_{F,t} > \widetilde{\sigma}_{M,t}$, agents of type $\gamma < \overline{\gamma}_t$ choose the fundamentalist heuristic for period $t+1$. Only agents with a low enough risk aversion are willing to accept the larger variance in order to profit from the better average performance. When $\langle u_F\rangle_t < \langle u_M\rangle_t$ in this case, $\overline{\gamma}_t < 0$, and all agents use the momentum rule in period $t+1$: When the risk is bigger and average performance worse, none of the agents use the fundamentalist heuristic. If $\widetilde{\sigma}_{F,t-1} > \widetilde{\sigma}_{M,t-1}$ and $\langle u_F\rangle_{t-1} \geq \langle u_M\rangle_{t-1}$, $\Theta_{F,t}$ is given by equation (3.11), with $s = F$ and $\Gamma_{F,t} = (0, \overline{\gamma}_{t-1})$.

The opposite holds when $\widetilde{\sigma}_{F,t-1} < \widetilde{\sigma}_{M,t-1}$: Agents of type $\gamma > \overline{\gamma}_{t-1}$ choose the fundamentalist heuristic for period t, so that $\Theta_{F,t} = 1$ when $\langle u_F\rangle_{t-1} > \langle u_M\rangle_{t-1}$, and $\Theta_{F,t}$ is given by equation (3.11), with $s = F$ and $\Gamma_{F,t} = (\overline{\gamma}_{t-1}, 0)$ if $\langle u_F\rangle_{t-1} \leq \langle u_M\rangle_{t-1}$.

We now have that

$$\Theta_{F,t+1} = \begin{cases} \Theta^{-1}\int_0^{\overline{\gamma}_t} \frac{g(\gamma)}{\gamma}\,d\gamma & \text{if } \widetilde{\sigma}_{F,t} > \widetilde{\sigma}_{M,t} \text{ and } \langle u_F\rangle_t \geq \langle u_M\rangle_t; \\ \Theta^{-1}\int_{\overline{\gamma}_t}^{\infty} \frac{g(\gamma)}{\gamma}\,d\gamma & \text{if } \widetilde{\sigma}_{F,t} < \widetilde{\sigma}_{M,t} \text{ and } \langle u_F\rangle_t \leq \langle u_M\rangle_t; \\ 1 & \text{if } \widetilde{\sigma}_{F,t} \leq \widetilde{\sigma}_{M,t} \text{ and } \langle u_F\rangle_t > \langle u_M\rangle_t; \\ 0 & \text{if } \widetilde{\sigma}_{F,t} \geq \widetilde{\sigma}_{M,t} \text{ and } \langle u_F\rangle_t < \langle u_M\rangle_t; \\ \frac{1}{2} & \text{if } \widetilde{\sigma}_{F,t} = \widetilde{\sigma}_{M,t} \text{ and } \langle u_F\rangle_t = \langle u_M\rangle_t. \end{cases} \tag{3.19}$$

The function $\Theta_{F,t+1}$ is discontinuous at points where $\widetilde{\sigma}_{F,t} = \widetilde{\sigma}_{M,t}$ and $\langle u_F\rangle_t = \langle u_M\rangle_t$: If we start in the $\langle u_F\rangle_t > \langle u_M\rangle_t$ area for example, and move towards the $\langle u_F\rangle_t < \langle u_M\rangle_t$ area while keeping $\widetilde{\sigma}_{F,t} = \widetilde{\sigma}_{M,t}$, $\Theta_{F,t+1}$ jumps from 1 to 1/2 to 0 as we pass the $\langle u_F\rangle_t = \langle u_M\rangle_t$ point. This discontinuity has consequences for the analysis of the dynamics that we discuss in subsection 3.3.2.

Combining equations (3.12) and (3.14), the dynamical system with fundamentalists and momentum traders is given by

$$\begin{cases} x_t = \dfrac{1}{R^f}\left\{\Theta_{F,t}\, f\, x_{t-1} + (1-\Theta_{F,t})\left[x_{t-1} + m\,(x_{t-1} - x_{t-2})\right]\right\} \\ \langle u_F\rangle_t = \eta\,\langle u_F\rangle_{t-1} - (1-\eta)\left(x_t - f\,x_{t-2}\right)^2 \\ \langle u_M\rangle_t = \eta\,\langle u_M\rangle_{t-1} - (1-\eta)\left[x_t - x_{t-2} - m\,(x_{t-2} - x_{t-3})\right]^2 \\ \left\langle u_F^2\right\rangle_t = \eta\left\langle u_F^2\right\rangle_{t-1} + (1-\eta)\left(x_t - f\,x_{t-2}\right)^4 \\ \left\langle u_M^2\right\rangle_t = \eta\left\langle u_M^2\right\rangle_{t-1} + (1-\eta)\left[x_t - x_{t-2} - m\,(x_{t-2} - x_{t-3})\right]^4, \end{cases} \tag{3.20}$$

where the risk-tolerance fractions $\Theta_{F,t}$ are given by equation (3.19).

3.2.3 Empirical grounding

To empirically ground our model, we draw from the rich literature on risk aversion, which includes studies on experiments, surveys, and decision-making in the field. Kimball et al. (2008) estimate the CRRA parameter for risk tolerance θ. They determine the distribution of risk tolerance by using hypothetical income gambles by respondents to the US Health and Retirement Study.[11] Kimball et al. find that a log-normal distribution fits the individual-level risk tolerance data well, because it has a fat right tail and imposes non-negative risk aversion. An additional advantage is that this distribution is computationally efficient. They estimate that $\log \theta \sim \mathcal{N}(-1.84, 0.73)$. This implies that $\log \gamma \sim \mathcal{N}(1.84, 0.73)$ and that the relative rate of risk aversion has mean 8.22 and standard deviation 6.90.

A limitation of Kimball et al.'s estimated distribution is that their respondents are all between 51 and 61 years of age. Aarbu and Schroyen (2009) take a similar approach, but use Norwegian survey data that covers ages 18 – 74. They find a lower average rate of risk aversion of 3.92, which is in line with earlier studies, with standard deviation 2.95. Figure 3.1 shows the density functions for both estimates. We find similar results with both distributions, but we present the results using Aarbu and Schroyen's estimate, since they cover a broader age group.

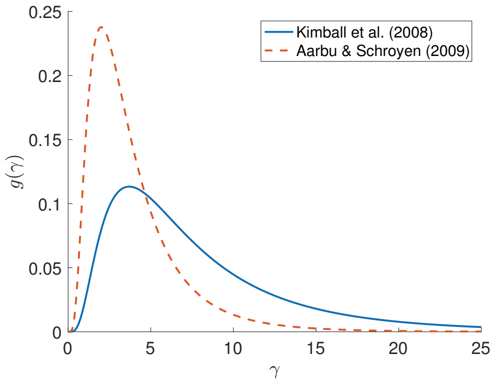

Figure 3.1: Distribution of relative rate of risk aversion, as estimated by Kimball et al. (2008): $\log \gamma \sim \mathcal{N}(1.84, 0.73)$, and Aarbu and Schroyen (2009): $\log \gamma \sim \mathcal{N}(1.14, 0.67)$.

A disadvantage of (most) surveys and experiments is that incentives are hypothetical or small. However, we have not been able to find studies that explicitly estimate the distribution of CRRA parameters using actual decision-making in the field. Cohen and Einav (2007) and Paravisini et al. (2016) do elicit relative risk aversion in this way, but these studies use income-based relative risk aversion: Absolute risk aversion is multiplied by a measure of income to obtain relative risk aversion. Reassuringly, both studies find the same distributional characteristics as laboratory experiments and survey studies.

[11] Surveys often have the drawback that they produce ordinal instead of cardinal data on risk aversion. The German Socio-Economic Panel Study (SOEP), which has been used in studies of risk aversion and its implications (e.g., Jaeger et al., 2010; Dohmen et al., 2012), falls into this category and is therefore not considered here.

3.3 Analytical results

3.3.1 Does a representative agent exist?

In finance, the behaviour of many agents is often modelled to be resulting from a single, representative agent. As a consequence, the heterogeneity in the market is not 'fundamental', and there is no need to explicitly model that heterogeneity. A natural question would be to ask whether the behaviour of our heterogeneous, switching agents can be captured by a representative agent.

We consider a representative agent with time-varying risk aversion who diversifies across heuristics. She observes the performance and variance of the two heuristics (and possibly their correlation). Instead of choosing only one, she can combine the two heuristics in a weighted average to achieve a better balance between expected performance and risk. The agent's forecast will be somewhere between the forecasts of both heuristics. She assigns weights to the heuristics by mean-variance optimization, where we allow the representative agent's risk aversion to change over time. Correlations are incorporated in the same way as in standard mean-variance analysis.

In Appendix B.1, we show that our market with heterogeneous agents cannot be captured by a representative agent with time-varying risk aversion. We furthermore prove that a representative agent aggregation does exist for the asset pricing model by Brock and Hommes (1998). Specifically, we show that their heterogeneous market with heuristic switching agents is equivalent to one with a diversifying representative agent whose risk aversion is given by

$$\gamma_{R,t+1} = \begin{cases} \dfrac{6\beta^2(U_{F,t} - U_{M,t})}{\pi^2} \dfrac{\exp(\beta\, U_{F,t}) + \exp(\beta\, U_{M,t})}{\exp(\beta\, U_{F,t}) - \exp(\beta\, U_{M,t})} & \text{if } U_{F,t} \neq U_{M,t}; \\ \dfrac{12\beta}{\pi^2} & \text{if } U_{F,t} = U_{M,t}. \end{cases} \tag{3.21}$$

Here, $\beta > 0$ is the intensity of choice parameter. It measures how quickly agents switch between heuristics, and plays a key role in previous heuristic switching models. The heuristic performance measures $U_{F,t}$ and $U_{M,t}$ for period t can be equal to the weighted average of squared forecasting errors that we use in our model, or any other measure of performance.[12]

3.3.2 Dynamics

In this section, we present analytical results on the dynamics of the asset pricing model as defined by the dynamical system (3.20). These results give an idea of the type of behaviour that the model can exhibit, and help us understand the simulations that we present in section 3.4. First, there is only one steady state, with the risky asset price at the fundamental and the weighted averages and variances at zero:

Proposition 1 (Existence of steady state). *The system* (3.20) *has precisely one steady state, namely* $(0, 0, 0, 0, 0)$.

[12] Although we refer to the fundamentalist and momentum rules here, we note that our conclusions regarding representative agent aggregations hold for any two heuristics, including those explored by Brock and Hommes (1998).

Proof. See Appendix B.3. □

In steady state, both fundamentalists and momentum traders believe that the risky asset price stays at the fundamental. Because we assume that all traders have rational beliefs about the fundamental price and dividend process, it follows that for all agents:

$$\widetilde{E}_{\gamma,t}\left[R^e_{t+1}\right] = E_t\left[p^*_{t+1} + d_{t+1}\right] - R^f p^*_t. \tag{3.22}$$

Because we furthermore assume that agents have homogeneous beliefs on variance, we get that individual, real demand for the risky asset is given by

$$z_{\gamma,t} = \frac{E_t\left[p^*_{t+1} + d_{t+1}\right] - R^f p^*_t}{\gamma\, \widetilde{V}_t\left[R^e_{t+1}\right]}, \tag{3.23}$$

and aggregate, real demand for the risky asset in the steady state is given by

$$\frac{E_t\left[p^*_{t+1} + d_{t+1}\right] - R^f p^*_t}{\gamma\, \widetilde{V}_t\left[R^e_{t+1}\right]}\,\Theta. \tag{3.24}$$

We cannot use standard linearization techniques to study the stability of the steady state, because the risk-tolerance fraction $\Theta_{F,t}$ has a discontinuity at this point. However, we can investigate its stability under the assumption that agents stick to one of the two forecasting rules, so that the risk-tolerance fractions are constant and equal to an exogenously given value Θ_F. This assumption simplifies the analysis even further, because we can ignore the part of the system that governs the switching, namely the weighted averages and variances of the heuristics' performance. In what follows, we derive three propositions that characterize the behaviour under the constant, exogenous fraction assumption. After this analysis, we relax the assumption to return to the original model, and use these propositions to understand its behaviour.

Exogenous, constant fractions

Under the assumption of exogenous, constant risk-tolerance fractions Θ_F, the relevant dynamics is described by a two-dimensional, linear system. If we write $v_t = x_{t-1}$, it is given by

$$\begin{pmatrix} v_t \\ x_t \end{pmatrix} = A \begin{pmatrix} v_{t-1} \\ x_{t-1} \end{pmatrix} \tag{3.25a}$$

$$A = \begin{pmatrix} 0 & 1 \\ -\frac{1}{R^f}\,(1-\Theta_F)\,m & \frac{1}{R^f}\left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right] \end{pmatrix}. \tag{3.25b}$$

Proposition 2 (Existence of steady state with constant fractions)**.** *The system* (3.25) *has* $(0,0)^T$ *as its only steady state.*

Proof. See Appendix B.3. □

This steady state corresponds to the risky asset price being at the fundamental, like that of the unrestricted system (3.20).

Turning to the dynamics around the steady state, we start with the special case of a homogeneous market:

Proposition 3 (Dynamics in homogeneous market). *If* $\Theta_F = 1$, *the steady state of* (3.25) *is stable, and the eigenvalues of A are real. Let*

$$m_{\pm} = 2\left(R^f \pm \sqrt{R^f \ \left(R^f - 1\right)}\right) - 1, \tag{3.26}$$

which satisfy $0 < m_- < 1$ *and* $m_+ > R^f$. *For* $\Theta_F = 0$, *we have:*

1. *If* $0 < m < R^f$, *the steady state is stable.*
2. *If* $m > R^f$, *the steady state is unstable.*
3. *If* $m \leq m_-$ *or* $m \geq m_+$, *the eigenvalues of A are real.*
4. *If* $m_- < m < m_+$, *the eigenvalues of A are complex.*[13]

Proof. See Appendix B.3. □

This proposition tells us that the steady state is stable when all agents use the fundamentalist rule, while the real eigenvalues in this case imply exponential price dynamics: The risky asset price exponentially converges to the fundamental. The intuition is that the fundamentalist rule is stabilizing, because it implies a move of asset price deviations towards zero, the steady state.

If all agents are momentum traders, stability depends on the momentum parameter: The larger m, the stronger past momentum is extrapolated into future asset price deviations. The stability threshold lies at the risk-free rate R^f: If the momentum parameter is smaller than the risk-free rate, the steady state is stable, while it is unstable if the parameter is larger than the risk-free rate.

The type of dynamics also depends on the momentum parameter. In the extreme cases of $m \leq m_-$ and $m \geq m_+$, we have exponential convergence towards, and divergence from the fundamental, respectively. In between these extremes ($m_- < m < m_+$), the eigenvalues are complex, indicating a rotation in the $v_t - x_t$ plane, and hence an oscillation of the asset price around the fundamental.

If the market consists of a constant mix of fundamentalists and momentum traders, the dynamics is described by the following proposition, which is represented graphically in Figure 3.2:

Proposition 4 (Dynamics with constant fractions). *Define* $m_{\pm}$ *as in Proposition 3. Furthermore define*

$$\overline{\Theta}_F = 1 - \frac{R^f}{m},$$

$$f_{\pm} = \frac{1 + m \pm \sqrt{(1+m)^2 - 4mR^f}}{2},$$

and

$$\Theta_{F,\pm} = \frac{(1+m-f)(1+m) - 2m\,R^f \pm 2\sqrt{m^2\ \left(R^f\right)^2 - f\,m\,R^f(1+m-f)}}{(1+m-f)^2}.$$

[13]We refer to an eigenvalue λ as complex if it has non-zero imaginary part: We can write $\lambda = a + b\,i$, with $a, b \in \mathbb{R}$, $b \neq 0$, and $i^2 = -1$.

1. *The steady state of* (3.25) *is unstable if* $\Theta_F < \overline{\Theta}_F$.

2. *The steady state of* (3.25) *is stable if* $\Theta_F > \overline{\Theta}_F$.

3. *If* $m \le m_-$ *and* $f_- \le f \le f_+$*, the eigenvalues of* A *are real. Otherwise, we have:*

   a) *The eigenvalues are real if* $\Theta_F \le \Theta_{F,-}$ *or* $\Theta_F \ge \Theta_{F,+}$.

   b) *The eigenvalues are complex if* $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$.

4. *If* $m > 1$*, it follows that* $\Theta_{F,-} < \overline{\Theta}_F < \Theta_{F,+}$.

5. *If* $m > m_-$*, it follows that* $0 < \Theta_{F,+} \le 1$*, where* $\Theta_{F,+} = 1$ *if and only if* $f = 0$.

6. *If* $m_- < m < m_+$*, it follows that* $\Theta_{F,-} < 0$*, and if* $m \ge m_+$*, we have that* $0 \le \Theta_{F,-} < \Theta_{F,+}$*, where* $\Theta_{F,-} = 0$ *if and only if* $m = m_+$.

Proof. See Appendix B.3. □

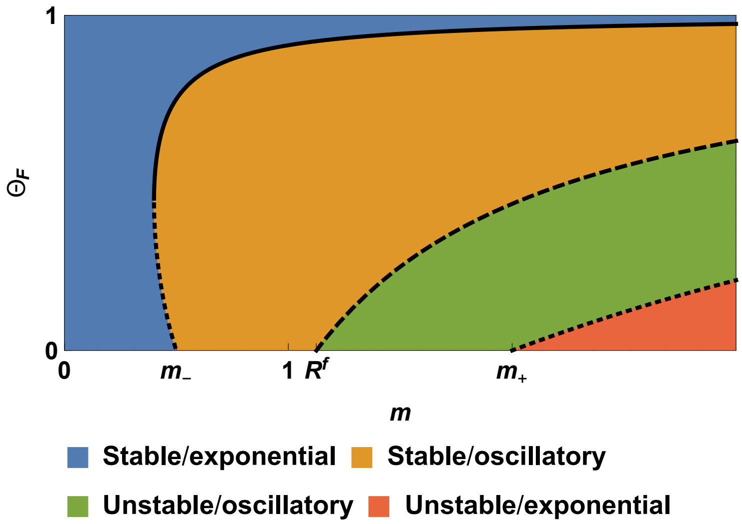

Figure 3.2: Graphical representation of Proposition 4, showing risky asset price dynamics with fixed fractions of fundamentalists and momentum traders, for different values of the momentum parameter m and the risk-tolerance fraction of fundamentalists Θ_F. Each of the four regions is characterized by two properties: the stability of the steady state and the type of dynamics (exponential or oscillatory). The solid and dotted lines represent $\Theta_{F,+}$ and $\Theta_{F,-}$, respectively, while the dashed line represents $\overline{\Theta}_F$. The shapes of these lines depend on the values of R^f and f, but the four regions remain. As R^f gets larger (smaller), the blue and red areas become smaller (larger). Changes in f have the opposite effect. The boundaries between the different regions cross the m-axis at m_-, R^f, and m_+, regardless of the parameter values.

This proposition tells us that in a heterogeneous market with constant fractions, the stability of the steady state depends on the risk-tolerance fraction of fundamentalists, relative to a threshold $\overline{\Theta}_F$. If the fundamentalist fraction is relatively low ($\Theta_F < \overline{\Theta}_F$), the steady state is unstable, while a relatively high fraction ($\Theta_F > \overline{\Theta}_F$) implies a stable steady state. The size of the momentum parameter m also plays an important role. First, note that if the momentum parameter is small enough, namely $m < R^f$, the steady state is stable, independent of the fraction of fundamentalists: In this case we have that $\overline{\Theta}_F < 0$, which means that $\Theta_F > \overline{\Theta}_F$, since Θ_F is positive. Second, a larger m implies a larger $\overline{\Theta}_F$, meaning that a larger risk-tolerance fraction of fundamentalists is required for the steady state to be stable. These results are consistent with the intuition behind Proposition 3 for the homogeneous case: The fundamentalist traders bring stability in the market, while the momentum rule is destabilizing, where a larger momentum parameter implies a stronger destabilizing effect.

The type of dynamics depends on the complexity of the eigenvalues of A, which, according to Proposition 4, is determined by combination of the risk-tolerance fraction of fundamentalists and the parameters m, f, and R^f: Real eigenvalues imply exponential behaviour, and complex eigenvalues imply oscillatory behaviour.

We can use above results to describe the behaviour of the risky asset price for specific ranges of the momentum parameter and the risk-tolerance fraction of fundamentalists. As shown in Figure 3.2, we can distinguish between four regions. First, if Θ_F is large or m small, the asset price exponentially converges towards the fundamental (the blue region). Second, the system exhibits converging oscillations when m is relatively small or Θ_F relatively large (the orange region). Third, if $m > R^f$, the asset price can exhibit oscillations diverging from the fundamental, depending on the value of Θ_F (the green region). Fourth, if $m \geq m_+$, the price exponentially diverges from the fundamental when Θ_F is small enough (the red region).

Endogenous, time-varying fractions

We return to the unrestricted system (3.20), in which agents are allowed to switch between forecasting strategies. Proposition 3 and Proposition 4 help us characterize its behaviour, because we can see the constant-fractions system (3.25) as part of the switching system. In each period t, $\Theta_{F,t}$ and the lagged asset price deviations x_{t-1} and x_{t-2} determine the current deviation x_t through (3.25). Given x_t, the weighted average and variances are pinned down by (3.20), which in turn determine $\Theta_{F,t+1}$ through (3.19). This new fraction then serves as the input for the constant-fractions system in period $t + 1$, and the cycle repeats.

Based on our analytical results for the system with constant fractions, we expect the risky asset price to exhibit a combination of oscillatory behaviour around, exponential convergence towards, and exponential divergence away from the fundamental. Proposition 4 tells us that the asset price will converge to the fundamental as long as the risk-tolerance fraction of fundamentalists $\Theta_{F,t}$ lies above the threshold $\overline{\Theta}_F$. As soon as $\Theta_{F,t}$ falls below the threshold, the asset price starts diverging away from the fundamental. The proposition also explains how depending on the fundamentalist fraction and the parameters of the system, the dynamics is exponential or oscillatory.

We cannot describe the precise dynamics of the switching system, as the fractions of fundamentalist and momentum traders are endogenously determined in a non-linear way. The evolution of the asset price deviations from the steady state depend on the non-linear interaction between the stabilizing and destabilizing forces of the fundamentalist and momentum investors.

In one particular case however, we can to a certain extent describe the behaviour of the switching system: As discussed earlier, the steady state of the constant-fractions system is stable regardless of the fraction of fundamentalists if $m < R^f$. This means that solutions of the switching system always converge to the steady state in this case. Whether the dynamics is exponential or oscillatory depends on the specific combination of parameters, but we know that the asset price eventually converges to the fundamental.

3.4 Numerical results

We now turn to numerical methods to study the dynamics of the asset pricing model with switching. The interesting dynamics occurs when the fundamentalist heuristic is attractive even when all agents use the momentum rule. This is the case when $m_- < m < m_+$, so that the asset price oscillates around the fundamental when $\Theta_F < \Theta_{F,+}$: During the oscillations, the fundamentalist rule becomes more attractive as the asset price moves towards the fundamental. For certain parameter combinations, the fundamentalist rule becomes attractive enough to make agents switch from the momentum rule. We use $R^f = 1.01$, $\eta = 0.2$, $f = 0.6$, and $m = 1.1$ as the baseline parameter values in our simulations. The distribution of relative risk aversion follows the empirically validated density $g(\gamma)$ that was introduced in subsection 3.2.3.

These parameter values correspond to $m_- \approx 0.82$, $m_+ \approx 1.22$, $\Theta_{F,-} \approx -0.02$, $\Theta_{F,+} \approx 0.84$, and $\overline{\Theta}_F \approx 0.08$. It follows that indeed $m_- < m < m_+$, so that $\Theta_{F,-} < 0$. Using the insights from subsection 3.3.2, we know that the asset price oscillates away from the fundamental if the risk-tolerance fraction of fundamentalists is smaller than 8%, oscillates towards it if the fundamentalist fraction lies between 8% and 84%, and exponentially converges towards the fundamental if the fraction is larger than 84%.

3.4.1 Chaotic dynamics

In the two left panels of Figure 3.3, we show the evolution of the price deviations from the fundamental, and of the risk-tolerance fraction of fundamentalists for the baseline parameters. The stabilizing and destabilizing forces of the two heuristics keep each other in check, leading to oscillations in the asset price that neither converge to, nor diverge from the fundamental in the long run. In period 1,003, one of the first periods shown, all traders follow the momentum rule, and the asset price follows an explosive, oscillatory path. This path first leads the asset price away from the fundamental, but it slows down and eventually moves back towards the fundamental starting in period 1,025. The return towards the fundamental makes the fundamentalist rule more attractive, and some investors start abandoning the momentum rule in period 1,035. As more investors follow, the system first moves from the unstable and oscillatory region into the stable and oscillatory region (crossing the dashed line), and then into the stable and exponential region (crossing the dotted line), meaning that the asset price starts on an exponential convergence towards the fundamental. The momentum rule picks up on this convergence, and becomes attractive again. Almost half of the investors have switched back to the momentum rule just before the asset price reaches the fundamental, so that it overshoots. It starts moving away from the fundamental again on a diverging oscillatory path.

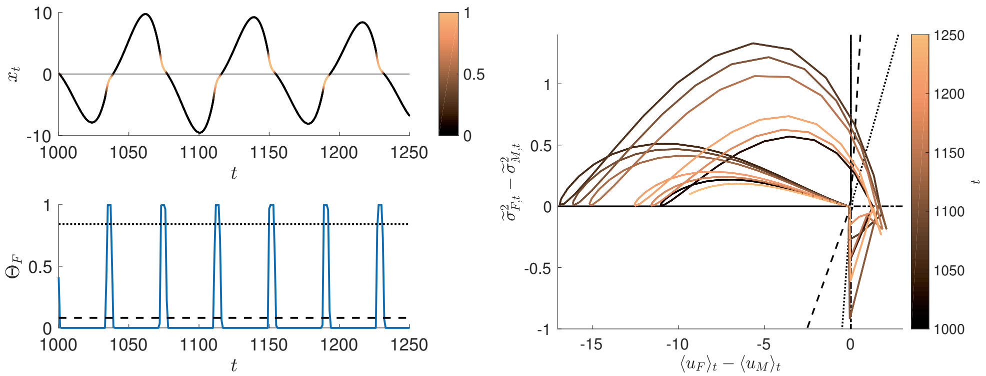

Figure 3.3: *Top left:* Price deviations from the fundamental for periods 1000 – 1250, where the colour gradient represents the value of Θ_F. *Bottom left:* Risk-tolerance fraction choosing the fundamentalist strategy for periods 1000 – 1250. The dashed and dotted lines indicate $\overline{\Theta}_F$ and $\Theta_{F,+}$, respectively. *Right:* Evolution of the dynamical system in the $(\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ plane for periods 1000 – 1250. The solid, dashed, dotted, and dot-dashed black lines indicate the 0, $\overline{\Theta}_F$, $\Theta_{F,+}$, and 1 contour lines of $\Theta_{F,t+1}$, respectively. *Parameter values:* $R^f = 1.01$, $\eta = 0.2$, $f = 0.6$, $m = 1.1$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $x_{-2} = x_{-1} = x_0 = 0.1$, $\langle u_F \rangle_0 = -1$, $\langle u_M \rangle_0 = -0.90$, $\langle u_F^2 \rangle_0 = 2$, $\langle u_M^2 \rangle_0 = 1.91$ (implying $\Theta_{F,1} = 0.5$).

The right panel of Figure 3.3 shows the behaviour of the system over the same period in the $(\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ plane. Its position in this plane determines the risk-tolerance fraction $\Theta_{F,t+1}$ in the next period, through equation (3.19). We have added several contour lines for $\Theta_{F,t+1}$ to the figure, in particular those corresponding to $\Theta_{F,t+1} = 0$ (solid) and $\Theta_{F,t+1} = 1$ (dot-dashed): In the second quadrant, corresponding to $\langle u_F \rangle_t - \langle u_M \rangle_t < 0$ and $\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t} > 0$, we have that $\Theta_{F,t+1} = 0$, and in the fourth quadrant, corresponding to $\langle u_F \rangle_t - \langle u_M \rangle_t > 0$ and $\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t} < 0$, we have that $\Theta_{F,t+1} = 1$. Note that all contour lines come together in the discontinuity at the origin. Each cycle, the system stays in the second quadrant relatively long, which means that all agents follow the momentum rule. As the fundamentalist rule becomes more attractive, it moves towards the first quadrant, where the variance of the fundamentalist rule is still larger than that of the momentum rule. This means that the most risk-tolerant investors are the first to switch to the fundamentalist rule. When the variance of the fundamentalist rule drops below that of the momentum rule, the system moves towards the fourth quadrant, where all agents are fundamentalists. As the momentum rule starts picking up on the asset price movement towards the fundamental and becomes more attractive, the system briefly visits the third quadrant. It moves back to the second quadrant, with only momentum traders in the market, when the variance of the momentum rule becomes smaller than that of the fundamentalist rule. This cycle repeats itself indefinitely, each time in a slightly different way.

Although the price deviations look regular, it is difficult to make long-term predictions about its state. The amplitude of the oscillations, as well as their length, changes with each cycle. This apparently random behaviour is called chaotic. Various definitions of chaos exist, but we follow the intuitive definition by Gros (2015, p. 66): "A deterministic dynamical system that shows exponential sensibility of the time development on the initial conditions is called chaotic." More rigorous definitions exist, but we are mostly interested in sensitive dependence on initial conditions and the apparent randomness that can arise in deterministic, chaotic systems, because these have implications for forecasting and policy-making.

The dynamics of a chaotic system can be characterized by a spectrum of Lyapunov exponents that indicate how nearby solutions converge or diverge over time. A negative exponent indicates convergence, while a positive exponent indicates divergence. The largest, or maximal, Lyapunov exponent, denoted by λ_1, determines whether the dynamics is chaotic (see, e.g., Gros, 2015). The average factor by which the distance between two neighbouring points increases or decreases in each time-step is given by e^{λ_1} (Boccara, 2010, ch. 5). Using the algorithm proposed by Benettin et al. (1980), we find for the asset pricing model that $\lambda_1 = 0.0035$.[14] This positive exponent implies exponential divergence: Solutions that start close to each other drift apart over time. This sensitive dependence on initial conditions is the defining characteristic of a chaotic system. Given the value $\lambda_1 = 0.0035$, the distance between neighbouring points doubles after approximately 200 time-steps.

Figure 3.4 shows the long-term behaviour of the asset pricing system. It plots the points of a solution consisting of one million time-steps (after an initialization period of 1000 steps) in the $(\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ plane (left panel) and in $x_t - (\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ space (right panel). In the left panel, we also show a contour plot for the fraction of fundamentalists Θ_F. It shows that the behaviour in Figure 3.3 is representative of the long-term dynamics. The system often comes close to the discontinuity at the origin, and it spends most time in the second quadrant, where all traders use the momentum rule. The right panel shows that these points correspond to large asset price deviations, while those points where x_t is small correspond to larger fractions of fundamentalists.

Both panels show that the solution is confined to a specific set of points: the strange attractor. Strange attractors arise in chaotic dynamical systems when trajectories are attracted to a subset of their phase space, but have chaotic dynamics within this subset: Nearby solutions that enter the attractor drift apart over time, while staying within the confines of the attractor. The dimension of a strange attractor is a useful tool in studying chaotic systems. A strange attractor's dimension quantifies the complexity of the system by representing its effective number of degrees of freedom. It can also be compared with other theoretical studies and empirically tested. The strange attractor in our model has a box-counting (fractal) dimension (Falconer,

[14]The other Lyapunov exponents are $\lambda_2 = 0.0010$, $\lambda_3 = -1.4917$, $\lambda_4 = -1.6059$, $\lambda_5 = -1.6094$, $\lambda_6 = -2.4687$, $\lambda_7 = -23.8493$. To calculate the exponents, we have used parameter values equal to those used in Figure 3.3, and an initialization period of 2000 time-steps.

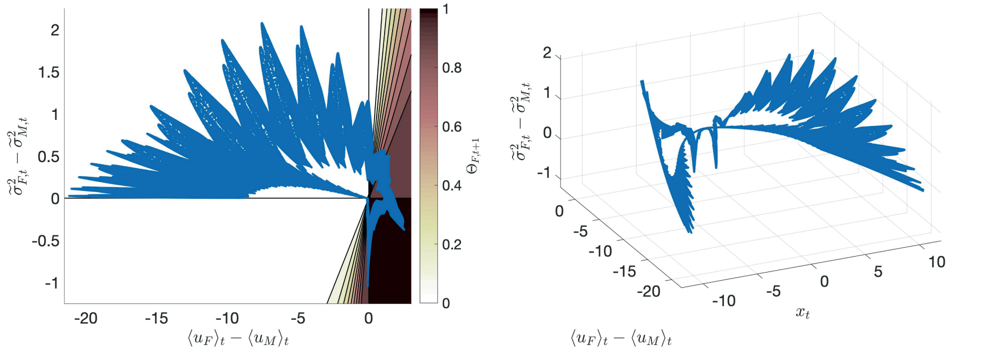

Figure 3.4: *Left:* Attractor of the asset pricing system (blue) and contour plot of $\Theta_{F,t+1}$ (colour gradient) in the $(\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ plane. *Right:* Attractor of the asset pricing system in $x_t - (\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ space. *Parameter values:* $R^f = 1.01$, $\eta = 0.2$, $f = 0.6$, $m = 1.1$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $x_{-2} = x_{-1} = x_0 = 0.1$, $\langle u_F \rangle_0 = -1$, $\langle u_M \rangle_0 = -0.90$, $\left\langle u_F^2 \right\rangle_0 = 2$, $\left\langle u_M^2 \right\rangle_0 = 1.91$ (implying $\Theta_{F,1} = 0.5$).

2004, ch. 3) of about 1.8. This value is close to the dimensions reported by Brock and Hommes (1998).

The dynamics of the asset pricing model crucially depends on the parameter values. One example is the threshold $\overline{\Theta}_F$, which is determined by the ratio between R^f and m. Here we focus on the memory parameter η. Figure 3.5 shows long-term asset price deviations in the deterministic system for different values of this parameter: For each value of η, this diagram plots 1000 asset price deviations after an initialization period of 1000 time-steps. Solutions with deviations that are larger than 1000 in absolute value are omitted.

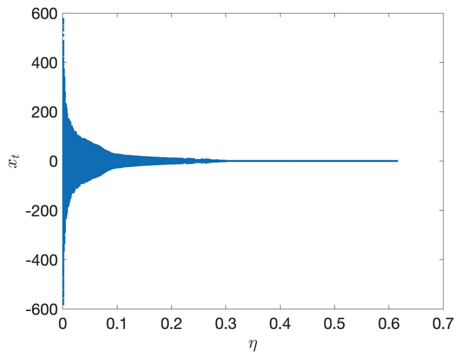

Figure 3.5: Long-term behaviour of the asset price deviations from the fundamental for different values of the memory parameter η. Deviations that are larger than 1000 in absolute value are omitted. *Parameter values:* $R^f = 1.01$, $f = 0.6$, $m = 1.1$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $x_{-2} = x_{-1} = x_0 = 0.1$, $\langle u_F \rangle_0 = -1$, $\langle u_M \rangle_0 = -0.90$, $\left\langle u_F^2 \right\rangle_0 = 2$, $\left\langle u_M^2 \right\rangle_0 = 1.91$ (implying $\Theta_{F,1} = 0.5$).

For η larger than 0.63, the system is unstable, and the amplitude of the asset price oscillations increases indefinitely. In this case, the momentum traders take over the market, and $\Theta_{F,t}$ stays constant at 0. For some values between 0.40 and 0.63, the solution converges to the steady state, with the asset price at the fundamental. For other values between 0.40 and 0.63, the dynamics is chaotic, with small asset price deviations from the fundamental. The dynamics is also chaotic if η is smaller than 0.40, with larger deviations for smaller η.

Note on the wealth distribution While the risk aversion distribution does not change in the asset pricing model, the wealth distribution can shift over time. More risk-tolerant traders invest a larger share of their wealth in the risky asset, which means that they can earn bigger returns. They also make bigger losses, but without a bankruptcy mechanism in place, they stay in the market. Shifts of wealth between traders does not affect the asset price however, because their demand for the risky asset is independent of their wealth (see equation 3.3). We have chosen not to include a bankruptcy mechanism to keep the asset pricing model as simple as possible, and focus on the dynamics resulting from the heterogeneous expectations.

Cobweb application To investigate to what extent the richness of the dynamics resulting from our expectation-formation framework depends on the choice of model, we also study its dynamics in the cobweb model (see Appendix B.4). The cobweb model is used by Muth (1961) to introduce the rational expectations hypothesis, and by Brock and Hommes (1997) to introduce their heuristic switching model. Contrary to the asset pricing model, this market exhibits negative expectations feedback: A higher expected price will lead to a lower realized price. Like in Brock and Hommes (1997), agents can choose between a stable, rational strategy, and an unstable, naive strategy, where the former is costlier. In this different context, our mean-variance switching model gives rise to a dynamics that is just as rich as that observed in the asset pricing model, including deterministic chaos.

3.4.2 Resilience to price shocks

To assess the resilience of our asset market, we also consider a version of the model that adds stochasticity in the form of noise traders. As we show in Appendix B.2, the pricing equation in this case includes a price shock, denoted by ϵ_t:

$$x_{t+1} = \frac{1}{R^f} \left\{ \Theta_{F,t}\, f\, x_t + (1 - \Theta_{F,t}) \left[x_t + m\,(x_t - x_{t-1}) \right] \right\} + \epsilon_t. \tag{3.27}$$

We simulate this stochastic version of the model with an i.i.d. price shock $\epsilon_t \sim \mathcal{N}(0, 0.05)$, which is small compared to the range of price deviations that occur in the deterministic system. The results are presented in the top panels of Figure 3.6.

We do not just observe noise around the deterministic solution of Figure 3.3, but the dynamics is different. The heterogeneous beliefs, especially the trend-following heuristic, amplify the small noise, leading to larger and less regular asset price deviations. Periods where the fundamentalists take over the market also last longer. The changed dynamics is reflected in the long-term behaviour, which is presented in the top-right panel of Figure 3.6: The shape of the attractor is preserved, but it is bigger and less well-defined.

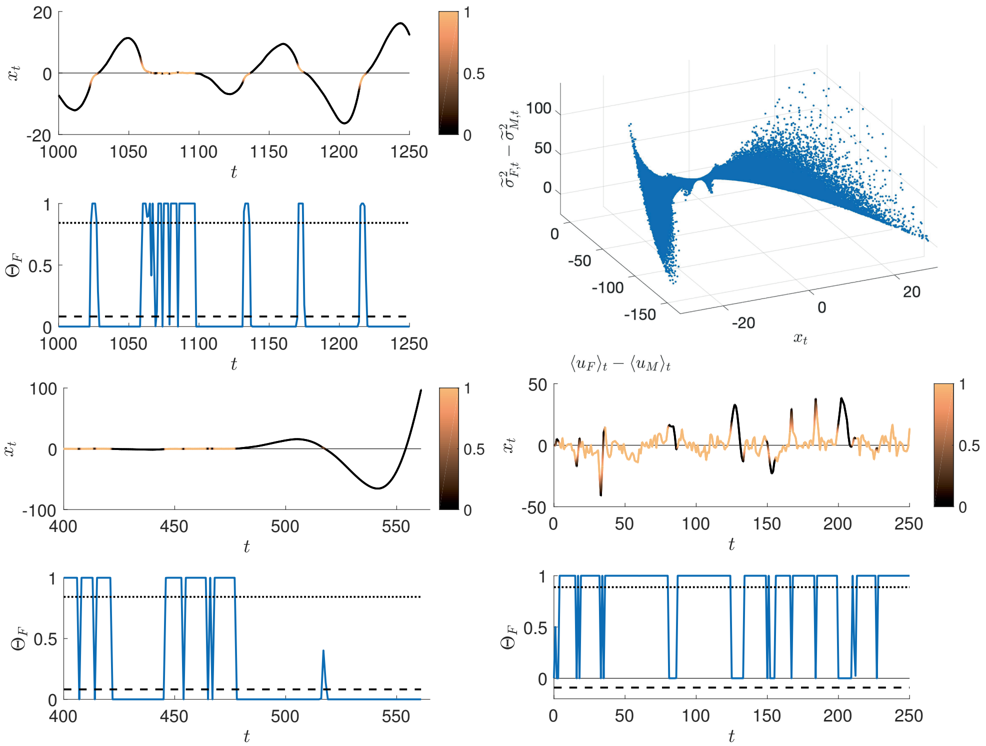

Figure 3.6: *Two top-left panels:* Price deviations from the fundamental (top) and risk-tolerance fraction choosing the fundamentalist strategy (bottom) for periods 1000 – 1250, with $R^f = 1.01$, $\eta = 0.2$, and x_{t+1} given by (3.27) with an i.i.d. price shock $\epsilon_t \sim \mathcal{N}(0, 0.05)$. *Top-right panel:* Attractor in $x_t - (\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ space of the noisy asset pricing system depicted in the top-left panels. *Two bottom-left panels:* Price deviations from the fundamental (top) and risk-tolerance fraction choosing the fundamentalist strategy (bottom) for periods 400 – 561, with $R^f = 1.01$, $\eta = 0.4$, and x_{t+1} given by (3.27) with an i.i.d. price shock $\epsilon_t \sim \mathcal{N}(0, 0.05)$. *Two bottom-right panels:* Price deviations from the fundamental (top) and risk-tolerance fraction choosing the fundamentalist strategy (bottom) for periods 0 – 250, with $R^f = 1.2$, $\eta = 0.4$, and x_{t+1} given by (3.27) with an i.i.d. price shock $\epsilon_t \sim \mathcal{N}(0, 5)$. *Note:* A colour gradient represents the value of Θ_F, and dashed and dotted lines indicate $\overline{\Theta}_F$ and $\Theta_{F,+}$, respectively. *Parameter values:* $f = 0.6$, $m = 1.1$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $x_{-2} = x_{-1} = x_0 = 0.1$, $\langle u_F \rangle_0 = -1$, $\langle u_M \rangle_0 = -0.90$, $\langle u_F^2 \rangle_0 = 2$, $\langle u_M^2 \rangle_0 = 1.91$ (implying $\Theta_{F,1} = 0.5$).

In the following, we investigate the stochastic system in different parameter ranges. When $\eta = 0.4$, the deterministic system converges to the steady state.[15] However, when we introduce a small normally distributed price shock with a standard deviation of 0.01, the solution is pushed out of equilibrium, and exhibits behaviour similar to that shown in the top panels of Figure 3.6.

Increasing the standard deviation of the shock to 0.05, the solution becomes explosive. In the bottom-left panels of Figure 3.6, we show its behaviour for periods 400-561. It starts close to the fundamental, with long periods of only fundamentalists in the market. After about 80 time-steps, the momentum traders take over, and the asset price starts oscillating around the fundamental with ever-increasing amplitude. Although the deterministic system is stable, it is less resilient than the chaotic $\eta = 0.2$ system, because it is less able to absorb small price shocks.

Following Proposition 4, we showed that the deterministic asset pricing model always converges to the steady state when $m < R^f$, because in this case $\Theta_{F,t} \geq 0 > \overline{\Theta}_F$. If we adjust the $\eta = 0.4$ system by setting $R^f = 1.2 > m = 1.1$, it is much more resilient. The bottom-right panels of Figure 3.6 show the solution with large price shocks $\epsilon_t \sim \mathcal{N}(0, 5)$. Note that indeed, the dashed line, corresponding to $\overline{\Theta}_F$, lies below zero, and hence $\Theta_{F,t}$ never comes below it. Even though the price shocks are bigger, the solution stays closer to the fundamental than in the case where $R^f < m$. The market is dominated by fundamentalists, which stabilize it. Momentum traders only take over from time to time to give rise to short-lived booms and busts.

3.5 Conclusion

We have incorporated diverse risk preferences into an asset pricing model with fundamentalists and momentum traders. Agents switch between the two forecasting rules based on a trade-off between forecasting performance and variability, taking into account their risk preferences. Heterogeneous expectations result from heterogeneous risk aversion. By using estimates for the risk aversion distribution, we have empirically grounded our switching mechanism.

We have proven that this model cannot be captured by a representative agent. Furthermore, we have shown that a steady state exists, and have characterized the asset price dynamics under the assumption of constant fractions: Stability of the steady state depends on the risk-tolerance fraction of fundamentalists relative to a threshold. This threshold depends on the strength of the momentum rule relative to the risk-free rate. If the risk-free rate is larger than the strength of the momentum rule ($R^f > m$), the steady state is always stable. A lower interest rate implies a larger threshold, which means that a larger fundamentalist fraction is needed for the steady state to be stable. This result is particularly interesting given the current low-interest-rate climate, and suggests that monetary stimulus can lead to less stable financial markets.

Numerically, we have shown that allowing for time-varying fractions, our asset pricing model can exhibit chaotic dynamics with a strange attractor. Bubbles

[15] This is true for a wide range of initial values, as long as the initial asset price deviations do not become too large. The convergence is not exponential and relatively slow: Starting from the initial conditions used in Figure 3.3, the solution oscillates towards the steady state in approximately 2500 time-steps.

and crashes emerge endogenously and unpredictably. By including stochastic asset price shocks, we have shown that heterogeneous beliefs and the resulting non-linearities have implications for resilience. Even when the endogenous dynamics (without stochastic shocks) is minimal or non-existent, small price shocks can cause large asset price bubbles.

We highlight three promising avenues for further research. First, agents could be modelled to form expectations by combining heuristics in a weighted average, in the same way that the representative agent does in subsection 3.3.1. While exploring the implications of this assumption, we have found that a representative agent does exist when the heterogeneous agents diversify across heuristics in this way, given that she has a time-varying risk aversion.[16] This could be linked to the time-varying risk aversion that is used in the more traditional finance literature that uses representative agents (e.g., Campbell and Cochrane, 1999).

Second, one may include a bankruptcy mechanism in our model, which we have not included for the sake of simplicity. Such a model allows for the study of wealth implications of heterogeneous beliefs. Combined with a demand function that depends on wealth, it can also be linked to aforementioned time-varying risk aversion in the representative agent framework, which can be explained by underlying wealth-shifts between agents with different risk preferences (Chan and Kogan, 2002; Xiouros and Zapatero, 2010).

Third, further empirical validation and calibration of our model could be achieved by comparing the numerical results, like the fractal dimension of the attractor, with data. It would be interesting to include different specifications of the forecasting strategies in this analysis.

[16] A proof is available upon request.

Chapter 4

The multidimensionality of sentiment[1]

Abstract

In line with Rational Beliefs theory, I identify sentiment as the difference between survey forecasts and non-judgmental forecasts that I derive from a large, real-time dataset. I use over half a million individual forecast differences covering various aspects of the economy in 1968Q4 - 2019Q4 to extract three underlying belief factors: (1) a long-term supply-side factor, (2) a short-term supply-side factor, and (3) a demand-side factor. Together, these three dimensions of sentiment account for roughly half of the variation in forecast differences. The belief factors' cross-sectional averages have distinct macroeconomic impacts, playing significant roles in interest rate and output dynamics. The factors are persistent, partly determined by learning, and correlated across forecasters. These results support core assumptions in Rational Beliefs theory. In some cases, however, beliefs violate its rationality assumption.

4.1 Introduction

The idea that optimism and pessimism play a significant role in macroeconomic fluctuations goes back to at least Pigou (1927). More recently, especially since the 2007-2008 financial crisis and ensuing recession, this idea has gained traction. However, linking sentiment (also referred to as animal spirits) to business cycles leads to mixed results. Some studies find that it plays an important role (e.g., Chahrour and Jurado, 2018; Milani, 2011; Winkler, 2019), while others find that it has little explanatory power (e.g., Fève and Guay, 2019). Beaudry and Willems (2018) show that over-optimism about growth prospects can be harmful in the long run.

These studies use various measures of sentiment. For example, they consider sentiment about different economic variables or over different horizons. Some studies use an ex-ante approach, comparing agents' expectations with benchmark forecasts to identify sentiment, while others use an ex-post approach, comparing them with realized values. This observation hints at an explanation for the inconsistent results: they measure different aspects of sentiment that all have a distinct macroeconomic impact. For example, fear of high inflation leaves room for being either optimistic

[1] This chapter was supported by the Complex Systems Fund, with special thanks to Peter Koeze.

about real activity (in a demand-pull scenario) or pessimistic (in a cost-push scenario). Animal spirits regarding these two scenarios might have different macroeconomic implications.

On the other hand, previous research indicates that different sentiment measures are related. The expectation-formation theory put forward by Winkler (2019) predicts that an agent's forecast bias is correlated across variables. Kurz and Motolese (2011) find that differences between survey projections and non-judgmental forecasts are correlated across variables and horizons. Intuitively, it makes sense that optimism about one aspect of the economy implies optimism (or pessimism) about other aspects. For example, optimism about technological progress might influence forecasts about economic growth, employment, or prices.

This chapter aims to combine the distinctness of some and the relatedness of other aspects of animal spirits by identifying a small number of orthogonal factors that together summarize an agent's sentiment. It furthermore intends to assess the macroeconomic impact of each of these factors (or dimensions) of sentiment. The analysis consists of three steps: (1) finding a clear, operationalizable definition of sentiment, (2) extracting the factors from appropriate data, and (3) assessing the macroeconomic impact of each factor.

For the first step, I rely on Rational Beliefs theory (Kurz, 1994). In this theory, agents do not know the exact structure of a complex, non-stationary economy. Instead, they form beliefs about its structure that are disciplined by rationality principles. The most important principle is that beliefs have to converge to the empirical distribution in the long run. The empirical distribution describes the non-stationary economic data over long periods and is known to all agents. Crucial for matching expectations data, this principle allows for heterogeneity: because many stochastic processes are consistent with the empirical one, agents can have diverse beliefs.[2] Furthermore, it implies fluctuations: after a period of optimism, agents have to become pessimistic for their beliefs to equal the empirical distribution on average.[3] Rational Beliefs theory thus offers a clear definition of sentiment: temporary deviations from the empirical distribution. In the remainder of this chapter, I use the terms beliefs and sentiment interchangeably.

For the second step of the analysis, Rational Beliefs theory provides explicit instructions for measuring beliefs about any given variable: computing the difference between observed forecasts for this variable and a non-judgmental forecast.[4] This forecast difference isolates beliefs. To ensure that my empirical exercise is adequately grounded in theory and pinpoint which beliefs I estimate, I develop a theoretical framework based on Rational Beliefs theory. This framework also allows me to derive testable implications from the assumptions underlying the theory. However, this chapter does not aim to develop a microfounded economic model that links beliefs to economic observables. Instead, I assume that beliefs play a role in the economy that is in line with earlier contributions to the Rational Beliefs literature that do use rigorous microfoundations. Various such contributions exist, both in macroeconomics

[2] See, for example, Motolese and Nielsen (2007) for a thorough discussion of this point.

[3] In Rational Beliefs theory, optimism and pessimism are defined relative to the empirical distribution. For example, agents are optimistic about economic growth when their forecast lies above the empirical forecast.

[4] This non-judgmental forecast is supposed to approximate the expected value under the empirical distribution.

and finance (e.g., Kurz and Motolese, 2011; Kurz et al., 2013). In this framework, a small number of factors drive all economic observables. These factors are subject to structural breaks about which agents form beliefs. The intertwined dynamics of the factors, observables, and beliefs are based on the log-linearized economy of Kurz et al. (2013).

I obtain data about observed forecasts from the Survey of Professional Forecasters (SPF). The SPF has been conducted quarterly since 1968Q4 and contains forecasts provided by about 40 professional forecasters per survey. It includes forecasts for various economic variables, including output, prices, housing, and interest rates. For each variable, projections are provided over multiple horizons, from nowcasts to 1-year ahead forecasts. To avoid biases caused by the massive volatility following the COVID-19 pandemic, I exclude the surveys after 2019Q4. The resulting dataset contains more than half a million individual forecasts. In line with the theory, most SPF panelists indicate that they combine a mathematical model with subjective beliefs to formulate their forecast (Stark, 2013).

It is crucial to use real-time data when comparing historical forecasts (e.g., Croushore and Stark, 2001; Orphanides, 2001). For each survey, I collect a large panel of real-time data covering all relevant aspects of the economy. It contains the data that was available to the forecasters at the time of the survey deadline. I employ the diffusion index approach of Stock and Watson (2002) to forecast the SPF variables. This approach is especially suitable for forecasting macroeconomic time series using many predictors. It extracts a small number of diffusion indexes (factors) from the real-time dataset to produce forecasts. These indexes correspond to the macroeconomic factors in the theoretical framework. To best approximate the forecast under the empirical distribution, I do not try to estimate temporary economic phenomena and use all available observations without judgment.

Subtracting from the SPF projections the corresponding non-judgmental forecasts gives a panel of beliefs for all variable-horizon combinations for all agents included in each survey. I use principal component analysis to extract the underlying belief factors from this panel. These factors capture the dimensions of maximum correlation between beliefs. Agents' beliefs for a particular variable and horizon are given by a linear combination of the factors up to an error term.[5] In line with earlier applications of Rational Beliefs theory (e.g., Kurz and Motolese, 2011; Kurz et al., 2013), I assume that the coefficients associated with each factor in this linear combination (the factor loadings) are constant across agents and time. The values of the factors (the factor scores) do vary across agents and time. The Rational Beliefs literature refers to the cross-sectional distribution of beliefs as market belief. Its mean is one of the determinants of aggregate fluctuations. In this chapter, market belief is multidimensional, with dimension equal to the number of belief factors.

For the final step of the analysis, I use two vector autoregressive (VAR) approaches. The first captures the dynamics of mean market belief's various dimensions and the endogenous variables of the classical New-Keynesian model: the output gap, inflation, and the interest rate. Kurz et al. (2013) study these same variables with their microfounded Rational Beliefs model. I compute impulse response functions and forecast error variance decompositions to assess how market belief contributes

[5] I refer to the idiosyncratic errors as variable-specific beliefs. The properties and implications of these beliefs are beyond the scope of this chapter.

to the fluctuations of the three macroeconomic variables. The second approach is closer to the theoretical framework and adds macroeconomic factors to the VAR model. Such a factor augmented VAR (FAVAR) has been applied in monetary policy research (e.g., Bernanke et al., 2005).[6]

The analysis leads to three main results. First, I find that a single factor only explains about a fifth of the variation in beliefs, showing that sentiment is indeed multidimensional. I identify three factors that together explain approximately half of the variation in the forecasters' beliefs. The first factor is positively related to inflation and interest rate beliefs and negatively to beliefs about housing, corporate profits, and output. Because it moves inflation and output in opposite directions and mainly affects long-term beliefs, I refer to this factor as a long-term supply-side inflation factor.[7] Analogously, I characterize the second factor as a short-term supply-side factor. A key difference with the first factor, apart from affecting short-term beliefs more, is that it moves output and interest rates in the same direction. It increases beliefs related to real activity, employment, profits, and interest rates while decreasing inflation beliefs. The third factor is positively related to inflation beliefs and beliefs regarding economic activity. Hence, I describe it as a demand-side factor. It decreases interest rate beliefs, which suggests that it concerns demand-side sentiment related to monetary policy. These three factors represent the dimensions of sentiment. The coordinates in this three-dimensional space together determine an agent's sentiment. As her sentiment changes, the coordinates also change.

The second important finding is that each dimension of market belief has a distinct macroeconomic impact. The impulse responses indicate that a one-standard-deviation positive shock to the first factor's cross-sectional average increases the interest rate by one to two standard policy moves. The second factor negatively impacts the output gap and the interest rate over more than a year, where the effect sizes are 0.3 and 0.5 percentage points, respectively. Shocks to the third dimension of mean market belief result in transient increases in the output gap (15 basis points) and the interest rate (25 basis points). Together, the three dimensions explain 40% of interest rate fluctuations, 15% of output gap fluctuations, and 6% of inflation fluctuations. Mean market belief also explains a substantial part of the macroeconomic factors included in the FAVAR, most notably of a real activity (15%) and interest rate factor (20%).

Third, I find that the sentiment factors have similar characteristics but also differ in meaningful ways. They exhibit the same degree of persistence, and they are correlated across agents. Learning contributes significantly to the factors' dynamics, with forecasters learning from real-time innovations to the macroeconomic factors. They are interrelated over short periods: lags of one factor partly determine other factors' realizations. The role of learning and the influence of other factors is different for each dimension. These characteristics are in line with Rational Beliefs theory. However, the beliefs about some variables and horizons violate the implications of its main rationality principle. Specifically, this principle implies that beliefs have zero long-term means and that the empirical variance bounds their variance. Average beliefs about many variables and horizons significantly deviate from zero, and while the majority

[6] I do not directly estimate the theoretical framework to stay close to the data.

[7] Some of the previous literature refers to horizons of about ten years as long-term. However, I use the term for a one-year horizon because it is the longest in my dataset.

of forecast differences satisfy the variance bound, short-term belief variances are up to five times as large.[8]

This chapter contributes to three streams of literature. The first investigates the role of animal spirits in macroeconomic fluctuations. To my knowledge, this is the first study that fully appreciates the multidimensionality of sentiment. The belief factors that I estimate have distinct, in some cases opposite, impacts. Studies that implicitly put different weights on each factor can reach contradictory conclusions about sentiment's role in business cycles. I discuss the two contributions to this literature that are most related to mine.

First, Milani (2011) identifies sentiment (which he refers to as expectation shocks) as the difference between forecasts from the SPF and forecasts resulting from a learning model. He estimates expectation shocks regarding the output gap, inflation, and the interest rate in the context of a DSGE model and finds that they explain approximately half of business cycle fluctuations. I use the same three macroeconomic series as Milani to assess the impact of beliefs, but I use a non-structural VAR approach to stay close to the data. I also use a more comprehensive range of variables across more horizons to identify sentiment. My results are qualitatively similar in that I find a significant impact of beliefs on macroeconomic fluctuations. However, I estimate the contribution of sentiment to be smaller. The different reference forecasts that I use to extract beliefs from the SPF forecasts might explain this discrepancy. Perhaps Milani's learning model misses some information that is included in my empirical forecasts and affects both the SPF forecasts and macroeconomic outcomes. The learning model is quarterly and uses only three macroeconomic time series, whereas I base the empirical forecast on a large, monthly dataset. Milani shows that macroeconomic factors do not explain the estimated expectation shocks, but he uses factors extracted from a revised panel instead of real-time data.

Second, Bhandari et al. (2019) use survey data on consumer expectations to calibrate a business cycle model with subjective beliefs. They identify these subjective beliefs as the one factor explaining most variation in what they call belief wedges: differences between benchmark forecasts (from the SPF or a VAR model) and survey forecasts for unemployment and inflation. They find that subjective beliefs play an important role in macroeconomic fluctuations. I extend their empirical contribution by identifying the factors underlying the belief wedges for more variables and across a range of horizons.

This study also contributes to the literature strand that strives to model expectation-formation. It provides characteristics regarding persistence, learning, and cross-sectional correlation that can serve as a benchmark for expectation theories. As I explicitly incorporate Rational Beliefs in the theoretical framework, the empirical results are particularly significant for the Rational Beliefs literature. Earlier studies provide empirical support for some assumptions using survey forecast data (e.g., Kurz and Motolese, 2011), but they look at beliefs for variables and horizons separately. Theoretical contributions hint at multidimensional beliefs, but this is the first study that investigates them. It sheds light on the number and nature of factors needed to capture the various aspects of agents' beliefs. Furthermore, this is

[8] These are violations of Rational Beliefs only within the context of my model (including, e.g., linearized transition functions for the observables) and to the extent that the data covers a sufficiently long period of time (the rationality principle is only guaranteed to hold in the limit $t \to \infty$). For a further discussion of this point, see this chapter's conclusion (section 4.4).

the first contribution that empirically supports learning and cross-sectional belief correlations in the Rational Beliefs setting. Lastly, this chapter provides an empirical underpinning for the impact of beliefs in macroeconomics, where Rational Beliefs theory is well-developed (Kurz et al., 2005, 2013, 2018), but its empirical relevance remains mostly unexplored.

Finally, this chapter contributes to the literature that studies forecasts by consumers, professionals, and government institutions. Earlier contributions look at biases in these forecasts (e.g., Elliott and Timmermann, 2008). For the SPF, Orlik and Veldkamp (2014) report that forecasters systematically underestimate GDP growth by 0.4 percentage points. This finding is in line with the negative average GDP beliefs that I observe. Instead of comparing survey forecasts to realizations, however, I compare them to non-judgmental forecasts. Similarly, Liebermann (2014) employs factors extracted from a comprehensive real-time panel to forecast GDP growth and concludes that these forecasts compare well to SPF predictions.

Herbst and Winkler (2020) also leverage the correlation of survey forecasts. They analyze the multidimensionality of disagreement by estimating a factor model on survey forecast data. Instead of using the factors to explain differences with a non-judgmental forecast, they use differences with the consensus forecast. They identify two factors, which they refer to as supply-side and demand-side. These factors are similar to my long-term supply-side and demand-side factors. A related paper by Dovern (2015) notes that the correlation between forecasts for different variables is low. One could argue that this finding is in line with the belief factors explaining about 50% of forecast differences. Like Herbst and Winkler (2020), I consider the glass half-full rather than half-empty: just three underlying factors explain over half of the variation in all 105 variable-horizon beliefs that I investigate.

4.2 Method

4.2.1 Theoretical framework

Factors

I consider an economy that is driven by several underlying factors that are subject to structural breaks. Agents in the economy form beliefs about the economy's structure based on quantitative and qualitative information. The cross-sectional distribution of these beliefs, which I refer to as market belief, plays a crucial role in the model: I assume that market belief partly explains the factor dynamics that drive observable economic variables. I sometimes refer to these factors as macroeconomic factors to distinguish between macroeconomic and belief factors.

I assume that there are r observable[9] factors that drive the economy and combine their month t values in an $r \times 1$ vector F_t. It evolves according to the following process:

$$F_{t+1} = A_F F_t + A_F^Z Z_{t+1} + A_F^s s_t + \widetilde{\epsilon}_{t+1}^F, \quad \widetilde{\epsilon}_{t+1}^F \sim \mathcal{N}(0, \widetilde{\Sigma}_F). \tag{4.1}$$

Here, Z_{t+1} denotes mean market belief in month $t+1$, a $d \times 1$ vector capturing average beliefs about the economy's structure across agents.[10] I will specify its dynamics and

[9] The factors are not observable in practice, but they can be estimated based on a large panel of data.

[10] As I show below, agents also form beliefs about mean market belief.

the underlying individual beliefs below. The $w \times 1$ vector s_t represents the economic structure at time t and is unobservable. The distribution of s_t is unknown and may vary over time. What is known, however, is that s_t averages to zero in the long run. As explained in the Rational Beliefs literature, s_t captures shifts in the economy due to, for example, the industrial or IT revolution. The coefficient matrices are A_F $(r \times r)$, A_F^Z $(r \times d)$, and A_F^s $(r \times w)$.

In Rational Beliefs theory, the empirical distribution plays an important role. Kurz et al. (2013, p 1410) define it as "the distribution one computes from a long series of observations by computing relative frequencies or moments of all past data together and where such computations are made without judgement or attempts to estimate the effect of any transitory short term events." It is denoted by the letter m and is known to all agents. I also refer to empirical transition functions and empirical forecasts, both based on the empirical distribution. The empirical transition function of F_t is given by

$$F_{t+1}^m = A_F F_t + A_F^Z Z_{t+1} + \epsilon_{t+1}^F, \quad \epsilon_{t+1}^F \sim \mathcal{N}(0, \Sigma_F). \tag{4.2}$$

Note that the transitory effect of the vector of structure parameters s_t disappears in the empirical distribution.

Observables

Let $\widetilde{x}_t$ be some economic observable, transformed such that its difference $\widehat{x}_t = \widetilde{x}_t - \widetilde{x}_{t-1}$ is stationary when considered over long periods. I assume that, in reality, observable variables are subject to structural shifts. One reason is that the r factors above, which undergo structural changes due to s_t, drive their dynamics. Moreover, I assume that observables are subject to variable-specific structural shifts, parametrized by the scalars s_t^x, which average to zero over time. I introduce these variable-specific shifts to be able to reconcile the theory with the data. In practice, the forecasters' belief factors do not entirely pin down their expectations. Variable-specific structure parameters, about which the agents form separate beliefs, allow for idiosyncratic deviations from the expectations implied by the belief factors. However, these parameters are not necessary for developing the theory, and previous literature does not use them.

The transition function of economic observables is given by

$$\widehat{x}_{t+1} = c_x + \lambda_x \widehat{x}_t + A_x^F F_{t+1} + \lambda_x^s s_t^x + \widetilde{\epsilon}_{t+1}^x, \quad \widetilde{\epsilon}_{t+1}^x \sim \mathcal{N}(0, \widetilde{\sigma}_x^2). \tag{4.3}$$

Here, c_x, λ_x, and λ_x^s are scalars, and A_x^F is a $1 \times r$ vector. This transition function is based on the log-linearized New-Keynesian economy with Rational Beliefs (Kurz et al., 2013), where time t aggregates are linear functions of mean market belief and exogenous shocks at time t. I add a lag of the observable to match the data better.

I assume that the empirical transition function is given by

$$\widehat{x}_{t+1}^m = c_x + \lambda_x \widehat{x}_t + A_x^F F_{t+1} + \epsilon_{t+1}^x, \quad \epsilon_{t+1}^x \sim \mathcal{N}(0, \sigma_x^2). \tag{4.4}$$

Beliefs

The state of agent i's beliefs is represented by the d-dimensional vector g_t^i. In the empirical application, it collects the agent's factor scores for d belief factors. It has the following transition function:

$$g_{t+1}^{i} = A_Z g_t^i + A_Z^F \left(F_{t+1} - A_F F_t - A_F^Z A_Z Z_t\right) + \epsilon_{t+1}^{ig}, \quad \epsilon_{t+1}^{ig} \sim \mathcal{N}(0, \Sigma_g). \tag{4.5}$$

Here, A_Z and A_Z^F are $d \times d$ and $d \times r$ matrices, respectively. I do not formally derive this transition function but assume that it is a multidimensional equivalent of the one-dimensional function used by Kurz et al. (2013). In my empirical analysis, I show that it fits the data well. To derive their transition function, Kurz et al. assume that agents form beliefs about the economy's structure by using a combination of Bayesian updating based on quantitative information and judgment of qualitative information. The specific functional form of the dynamics does not matter for the empirical identification of the belief factors, but introducing some structure helps with the interpretation.[11] For example, it enables me to investigate the role of learning in the formation of beliefs.

Learning is represented by the second right-hand-side term in equation (4.5). As I show later, the expression in parentheses captures the difference between the realization of the macroeconomic factors and their empirical forecast based on data for the previous month. This difference is determined by stochastic shocks and the structure parameter s_t, and therefore provides a noisy signal about any structural breaks.

The correlation of beliefs across agents plays an essential role in Rational Beliefs theory. It is represented by the correlation of the stochastic terms ϵ_{t+1}^{ig} across i, which I test on the data. Considered across agents and time, they follow a multivariate normal distribution with covariance matrix Σ_g. Rational Beliefs theory requires that the g_t^i have a long-run mean of zero.

I refer to the cross-sectional distribution of g_t^i as market belief and denote its mean by Z_t. It is observable, and its transition function follows from (4.5):

$$Z_{t+1} = A_Z Z_t + A_Z^F \left(F_{t+1} - A_F F_t - A_F^Z A_Z Z_t\right) + \widetilde{\epsilon}_{t+1}^{Z}. \tag{4.6}$$

Note that the agent-level terms ϵ_{t+1}^{ig} do not average to zero because they are correlated across agents. Instead, they average to the stochastic term $\widetilde{\epsilon}_{t+1}^{Z}$, whose distribution is unknown and possibly time-varying.

For simplicity, I assume that the empirical transition function of Z_{t+1} is given by

$$Z_{t+1}^{m} = A_Z Z_t + A_Z^F \left(F_{t+1} - A_F F_t - A_F^Z A_Z Z_t\right) + \epsilon_{t+1}^{Z}, \quad \epsilon_{t+1}^{Z} \sim \mathcal{N}(0, \Sigma_Z). \tag{4.7}$$

In general, agents do not believe that the empirical distribution is the truth. Instead, agent i's perceived transition function of Z_{t+1} is given by

$$Z_{t+1}^{i} = A_Z Z_t + A_Z^F \left(F_{t+1} - A_F F_t - A_F^Z A_Z Z_t\right) + A_Z^g g_t^i + \epsilon_{t+1}^{iZ}, \quad \epsilon_{t+1}^{iZ} \sim \mathcal{N}(0, \widehat{\Sigma}_Z). \tag{4.8}$$

The term $A_Z^g g_t^i$ captures the agent's belief about the aggregate stochastic term $\widetilde{\epsilon}_{t+1}^{Z}$.[12] She believes that F_{t+1} follows the following process:

$$F_{t+1}^{i} = A_F F_t + A_F^Z Z_{t+1} + A_F^g g_t^i + \epsilon_{t+1}^{iF}, \quad \epsilon_{t+1}^{iF} \sim \mathcal{N}(0, \widehat{\Sigma}_F). \tag{4.9}$$

[11] Some contributions to the Rational Beliefs literature use a different transition function. Kurz and Motolese (2011), for example, do not (explicitly) incorporate learning in the belief dynamics.

[12] Unfortunately, I am not able to identify A_Z^g, because the SPF does not contain forecasts of mean market belief.

Agents also form beliefs about the variable-specific structure parameters s_t^x. These beliefs are not persistent and are based solely on qualitative information.[13] They are represented by the random variable ρ_t^{ix} that satisfies

$$\rho_t^{ix} \sim \mathcal{N}(0, \sigma_{x\rho}^2). \tag{4.10}$$

Like the innovations to the belief factors (ϵ_t^{ig}), the variable-specific beliefs are correlated across agents. This correlation implies that the mean of the cross-sectional distribution, which I denote by P_t^x, is non-zero. For simplicity, I assume that it is given by

$$P_t^x \sim \mathcal{N}(0, \sigma_{xP}^2). \tag{4.11}$$

I assume that P_t^x does not play a role in the dynamics of economic observables.[14]

Agent i's perceived transition function of the observable variable $\widehat{x}_{t+1}$ is given by

$$\widehat{x}_{t+1}^i = c_x + \lambda_x \widehat{x}_t + A_x^F F_{t+1} + \rho_t^{ix} + \epsilon_{t+1}^{ix}, \quad \epsilon_{t+1}^{ix} \sim \mathcal{N}(0, \widehat{\sigma}_x^2). \tag{4.12}$$

Forecasts

Combining the transition function of the macroeconomic factors (4.1) with that of mean market belief (4.6) gives

$$F_{t+1} = A_F F_t + A_F^Z A_Z Z_t + \left(I_r - A_F^Z A_Z^F\right)^{-1} \left(A_F^s s_t + A_F^Z \widetilde{\epsilon}_{t+1}^Z + \widetilde{\epsilon}_{t+1}^F\right), \tag{4.13}$$

where I_r denotes the $r \times r$ identity matrix. This derivation is valid as long as $\left(I_r - A_F^Z A_Z^F\right)$ is invertible, which the data indicates it is.[15] A direct implication is that

$$Z_{t+1} = A_Z Z_t + A_Z^F \left(I_r - A_F^Z A_Z^F\right)^{-1} \left(A_F^s s_t + A_F^Z \widetilde{\epsilon}_{t+1}^Z + \widetilde{\epsilon}_{t+1}^F\right) + \widetilde{\epsilon}_{t+1}^Z. \tag{4.14}$$

It also has implications for the transition function of observable variables, which can now be written as

$$\begin{aligned} \widehat{x}_{t+1} = {} & c_x + \lambda_x \widehat{x}_t + \lambda_x^s s_t^x + \widetilde{\epsilon}_{t+1}^x + \\ & A_x^F \left[A_F F_t + A_F^Z A_Z Z_t + \left(I_r - A_F^Z A_Z^F\right)^{-1} \left(A_F^s s_t + A_F^Z \widetilde{\epsilon}_{t+1}^Z + \widetilde{\epsilon}_{t+1}^F\right)\right]. \end{aligned} \tag{4.15}$$

[13] In principle, I could introduce a dynamics similar to the one for the belief factors (4.5). Because I focus on the dynamics and role of the belief factors, I assume that the variable-specific beliefs are drawn from a normal distribution. Moreover, this assumption fits the procedure that I use to extract the belief factors well.

[14] I could add an extra term that involves P_t^x to the empirical dynamics of economic observables (4.4). I leave it out for two reasons. First, I only include the variable-specific structure parameter and beliefs to reconcile the theory with the empirical application. I do this in a way that distracts as little as possible from the main focus of the chapter: the dynamics of and the role played by the belief factors g_t^i and their cross-sectional mean Z_t. Second, unlike the belief factors, variable-specific beliefs have not been linked to economic fluctuations in a microfounded theoretical model.

[15] Using the final estimates for the real-time macroeconomic factors and mean market belief, I estimate the following equation:

$$F_{t_s} = A_Z F_{t_s - 1} + A_F^Z Z_{t_s} + \epsilon_{t_s}^F.$$

Here t_s indicates the month associated with survey s (see subsection 4.2.3). I combine the resulting estimate for A_F^Z with that for $\bar{A}_Z^F$ from (4.43) to numerically approximate the reciprocal condition number of $\left(I_r - A_F^Z A_Z^F\right)$. The resulting reciprocal condition number of about 0.59 suggests that it is indeed invertible.

Assuming that the empirical innovations $\epsilon^{\bullet}_{t+1}$ are independent, I can write the empirical transition functions as follows:

$$
\begin{aligned}
F^m_{t+1} &= A_F F_t + A^Z_F A_Z Z_t + \bar{\epsilon}^F_{t+1}, & \bar{\epsilon}^F_{t+1} &\sim \mathcal{N}(0, \bar{\Sigma}_F), && (4.16)\\
Z^m_{t+1} &= A_Z Z_t + \bar{\epsilon}^Z_{t+1}, & \bar{\epsilon}^Z_{t+1} &\sim \mathcal{N}(0, \bar{\Sigma}_Z), && (4.17)\\
\widehat{x}^m_{t+1} &= c_x + \lambda_x \widehat{x}_t + A^F_x A_F F_t + A^F_x A^Z_F A_Z Z_t + \bar{\epsilon}^x_{t+1}, & \bar{\epsilon}^x_{t+1} &\sim \mathcal{N}(0, \bar{\sigma}^2_x). && (4.18)
\end{aligned}
$$

I write the transition functions, as perceived by agent i, as

$$
\begin{aligned}
F^i_{t+1} &= A_F F_t + A^Z_F A_Z Z_t + \bar{A}^g_F g^i_t + \bar{\epsilon}^{iF}_{t+1}, & \bar{\epsilon}^{iF}_{t+1} &\sim \mathcal{N}(0, \check{\Sigma}_F), && (4.19)\\
Z^i_{t+1} &= A_Z Z_t + \bar{A}^g_Z g^i_t + \bar{\epsilon}^{iZ}_{t+1}, & \bar{\epsilon}^{iZ}_{t+1} &\sim \mathcal{N}(0, \check{\Sigma}_Z), && (4.20)\\
\widehat{x}^i_{t+1} &= c_x + \lambda_x \widehat{x}_t + A^F_x A_F F_t + A^F_x A^Z_F A_Z Z_t + \\
&\qquad A^g_x g^i_t + \rho^{ix}_t + \bar{\epsilon}^{ix}_{t+1}, & \bar{\epsilon}^{iZ}_{t+1} &\sim \mathcal{N}(0, \check{\sigma}^2_x). && (4.21)
\end{aligned}
$$

Here, I assume that the individual innovations $\epsilon^{i\bullet}_{t+1}$ are independent.

The perceived distribution (4.21) implies the following for agent i's expectation of $\widehat{x}_{t+1}$ at time t:

$$E^i_t[\widehat{x}_{t+1}] = c_x + \lambda_x \widehat{x}_t + A^F_x A_F F_t + A^F_x A^Z_F A_Z Z_t + A^g_x g^i_t + \rho^{ix}_t. \tag{4.22}$$

The implied empirical forecast is given by

$$E^m_t[\widehat{x}_{t+1}] = c_x + \lambda_x \widehat{x}_t + A^F_x A_F F_t + A^F_x A^Z_F A_Z Z_t. \tag{4.23}$$

The difference between the individual and empirical forecasts isolates beliefs:

$$E^i_t[\widehat{x}_{t+1}] - E^m_t[\widehat{x}_{t+1}] = A^g_x g^i_t + \rho^{ix}_t. \tag{4.24}$$

This expression can be interpreted as a factor model, where g^i_t is the factor score of agent i at time t, A^g_x captures the factor loadings of variable x on the d belief factors, and the variable-specific beliefs ρ^{ix}_t designate the idiosyncratic error terms. The g^i_t thus represent coordinates in the d-dimensional space spanned by the belief factors. They summarize agent i's overall sentiment at time t. However, these coordinates do not fully determine sentiment because the variable-horizon-specific beliefs also account for part of it.[16]

In the this chapter's empirical part, I use SPF data on individual forecasts and construct empirical forecasts to isolate beliefs as in equation (4.24). To include as much data as possible in the PCA, I exploit the SPF forecasts over all available quarterly horizons, from nowcasts to one-year ahead forecasts. To facilitate this approach, I derive an expression for forecasts over multiple horizons.

Proposition 5 (Distributions of h-month change to observable)**.** *Under the empirical and perceived distributions, the h-month change to observable $\widetilde{x}_t$ ($h \geq 1$) is given by*

$$
\begin{aligned}
[\widetilde{x}_{t+h} - \widetilde{x}_t]^m &= c_{hx} + \lambda_{hx} \widehat{x}_t + A^F_{hx} F_t + A^Z_{hx} Z_t + u^{hx}_{t+h}, & u^{hx}_{t+h} &\sim \mathcal{N}(0, \sigma^2_{hx}), && (4.25)\\
[\widetilde{x}_{t+h} - \widetilde{x}_t]^i &= c_{hx} + \lambda_{hx} \widehat{x}_t + A^F_{hx} F_t + A^Z_{hx} Z_t + \\
&\qquad A^g_{hx} g^i_t + \rho^{ihx}_t + u^{ihx}_{t+h}, & u^{ihx}_{t+h} &\sim \mathcal{N}(0, \widehat{\sigma}^2_{hx}). && (4.26)
\end{aligned}
$$

[16] As I show in the empirical part of this chapter, the belief factors and idiosyncratic components each explain approximately half of the variance of beliefs.

Here, c_{hx} and λ_x are scalars, both linear functions of c_x and λ_x. The $1 \times r$ vector A^F_{hx}, and $1 \times d$ vectors A^Z_{hx} and A^g_{hx} are linear combinations of the $A^{\bullet}_{\bullet}$ operators introduced earlier. Agent i's variable-horizon-specific belief, denoted by ρ^{ihx}_t, is a scalar multiple of ρ^{ix}_t. Its cross-sectional mean is P^{hx}_t, which is a scalar multiple of P^x_t. The error variances are linear combinations of the (co)variances $\bar{\Sigma}_{\bullet}$, $\bar{\sigma}^2_x$, $\check{\Sigma}_{\bullet}$, $\check{\sigma}^2_x$, Σ_g, and $\sigma^2_{x\rho}$ introduced earlier.

Proof. See Appendix C.2. □

The h-period ahead forecasts of $\widetilde{x}_{t+h}$ directly follow from Proposition 5:

$$E^m_t[\widetilde{x}_{t+h}] = \widetilde{x}_t + c_{hx} + \lambda_{hx}\widehat{x}_t + A^F_{hx}F_t + A^Z_{hx}Z_t, \tag{4.27}$$

$$E^i_t[\widetilde{x}_{t+h}] = \widetilde{x}_t + c_{hx} + \lambda_{hx}\widehat{x}_t + A^F_{hx}F_t + A^Z_{hx}Z_t + A^g_{hx}g^i_t + \rho^{ihx}_t. \tag{4.28}$$

The difference between the two forecasts isolates beliefs about a given economic observable x and horizon h:

$$E^i_t[\widetilde{x}_{t+h}] - E^m_t[\widetilde{x}_{t+h}] = A^g_{hx}g^i_t + \rho^{ihx}_t. \tag{4.29}$$

The factor model interpretation carries over from equation (4.24), where the factor loadings and error terms are variable-horizon-specific.

Implications of rationality

The central axiom of Rational Beliefs theory posits that beliefs should be compatible with the available data (Kurz, 1994). It implies that the perceived distributions for the factors, market belief, and observables should converge to the empirical distributions in the long run. I derive two testable implications.

First, the axiom implies that over long periods, beliefs should have a vanishing mean. By comparing the perceived distribution for observables over horizon h (4.26) with the empirical one (4.25), for example, it follows that

$$E\left[A^g_{hx}g^i_t + \rho^{ihx}_t\right] = 0. \tag{4.30}$$

This equation should hold for every variable x and horizon h (and every agent i). Because ρ^{ihx}_t has zero mean, this implies that g^i_t should have zero mean. Because this has to be valid for all agents, I also have that Z_t has a vanishing long-run mean.

The rationality axiom also implies that the variance of the empirical distribution bounds the variance of beliefs. Specifically, it follows from equations (4.25) and (4.26) that

$$\begin{aligned}\operatorname{Var}\left[A^g_{hx}g^i_t + \rho^{ihx}_t + u^{ihx}_{t+h}\right] &= \operatorname{Var}\left[u^{hx}_{t+h}\right]\\ \operatorname{Var}\left[A^g_{hx}g^i_t + \rho^{ihx}_t\right] + \widehat{\sigma}^2_{hx} &= \sigma^2_{hx},\end{aligned}$$

which in turn implies that for any x, h, and i:

$$\operatorname{Var}\left[A^g_{hx}g^i_t + \rho^{ihx}_t\right] \le \sigma^2_{hx}. \tag{4.31}$$

4.2.2 Data

I use data from the Survey of Professional Forecasters, covering each quarter in 1968Q4 – 2019Q4. The surveys have between 9 and 87 respondents (not counting respondents with only missing forecasts). They have 40 respondents on average, with 90% of the surveys having between 19 and 62 respondents.

I include data for 21 variables that are part of the SPF.[17] Some variables were added several years after the start of the SPF. In Table C.1 in the appendix, I list a description of the variables included at each point in time. For each variable, survey respondents provide forecasts over six horizons: a backcast for the quarter before the survey, a nowcast for the current quarter, and quarterly forecasts for up to one year ahead. I exclude the backcast from the analysis because usually there is data available for this horizon, which means that most forecasters simply report the value given by the latest data. Several forecasts are missing, with some forecasters providing forecasts for only a subset of the variables or horizons. Not counting missing observations, the SPF data covers about 526,000 individual forecasts.

To construct empirical forecasts, I use snapshots (also called vintages) of the information available at the time of each survey deadline. Specifically, I use a large panel of monthly real-time data with observations starting in 1959:01. It is based on the FRED-MD monthly database (McCracken and Ng, 2016) provided by the St. Louis Fed. I include all variables that are not revised and those for which real-time data is available. Variables for which this is not the case, I try to find real-time replacements. The panel also includes all variables that are forecast in the SPF, some of which are quarterly. It covers all relevant aspects of the economy: (a) output and income, (b) the labor market, (c) housing, (d) consumption, investment, orders, and inventories, (e) money and credit, (f) interest and exchange rates, (g) prices, (h) the stock market. It contains hard as well as soft (survey) data. For a given month t, the corresponding real-time panel contains data available at the time, up to and including the previous month ($t-1$).

Most data comes from the ALFRED database provided by the St. Louis Fed. For some variables, real-time data is only available from a specific date after the start of the SPF onwards, which means that the panel becomes wider over time. It varies between 66 for the first survey and 104 for the last survey. Moreover, the panel is unbalanced in general. These properties reflect the forecasters' changing information set but are also due to limitations of the real-time data. I describe the panel in detail in Appendix C.1. Among other things, it provides information about the first available observation for each variable and about a seasonal adjustment that I use for some variables.

I carefully match the real-time data with the information set of the SPF participants. For the surveys since 1990Q3, survey deadlines are known.[18] For those surveys, I start with the data available on the day of the deadline, when many participants return their surveys (Stark, 2010). Before 1990Q3, I start with the data available on the 12th of the second month of the quarter.[19] For each variable and each survey, I

[17] I exclude long-term forecasts, probability forecasts, and forecasts for real net exports and the change in real private inventories.

[18] As a robustness test, I carry out the analysis using only surveys with a known deadline.

[19] This choice is also guided by Stark (2010), who states that the industrial production data for the first month of the survey quarter is generally not available to the forecasters. Since in my sample the earliest

check the differences between the SPF backcasts and the corresponding entry in the real-time dataset. When for a given variable, all forecasters report a backcast that differs from the data, I change the vintage date to ensure that the backcasts match the real-time data. In some cases, this means setting an earlier date. In others, this means setting a later date that reflects more recent releases. I then recompile the data for that quarter and repeat the exercise until backcast and data match for at least one forecaster for all variables. There are cases where some forecasters report the latest data, but others report older data in their backcasts. In these cases, I use the most recent data to compute the empirical forecast since this best reflects the available information at the time.

For assessing the macroeconomic impact of beliefs, I use quarterly data on the output gap (using the U.S. Congressional Budget Office estimate of real potential output), the level of the chain-weighted GDP price index, and the federal funds rate. The St. Louis Fed provides this data through its FRED database. I use revised data to get the most accurate estimate currently possible of the economy's reaction to belief shocks. I also use factors extracted from the FRED-MD monthly database. It consists of 128 variables covering the same aspects of the economy as the real-time panel I use to construct the empirical forecast, also in the period 1959:01 – 2019:12. It covers a broader set of variables because it only contains revised data.

4.2.3 Translation of theory to empirical setting

The forecasts included in the SPF data concern quarterly variables and quarterly averages of monthly variables. The theoretical framework is defined in terms of monthly dynamics, however. To harmonize the two, I transform all data into monthly time series, where the SPF forecasts refer to the values in the third month. In the case of quarterly variables, the values in the third month of a given quarter correspond to quarterly observations. In the case of monthly variables, those values correspond to quarterly averages. I explain the two cases in more detail below.

Let x_t indicate a monthly variable (e.g., industrial production). I transform it to an average over a 3-month window, denoted by $\bar{x}_t$: $\bar{x}_t = \frac{1}{3}(1 + L + L^2)x_t$, where L indicates the lag operator. The observation in the third month of the quarter then corresponds to the quarterly average of the monthly series.

Now let $\overset{*}{x}_q$ be a time series for some quarterly variable, where $q = 1, 2, \ldots$ indicates the quarter. I assume that there is an underlying monthly series $\check{x}_t$, and a transformation $\bar{x}_t = f(L)\check{x}_t$ such that the value in the third month of each quarter is equal to the corresponding quarterly observation: $\bar{x}_{\tau_q} = \overset{*}{x}_q$, where τ_q denotes the third month of quarter q. In the case of real GDP, for example, $\check{x}_t$ would be 'monthly GDP,' and the transformation would be given by $f(L) = (1 + L + L^2)$. Note that observations for the first two months of a quarter are missing since only data for $\overset{*}{x}_q$ is available.

I transform each variable such that its 3-month difference is stationary. The transformed variable corresponds to $\widetilde{x}_t$ in the theory. For example, if x_t is an interest rate, I define $\widetilde{x}_t = \bar{x}_t$. In the case of industrial production, I use $\widetilde{x}_t = 100 \log \overline{x}_t$. This transformation implies that the 3-month difference corresponds to a quarterly log-growth

release date for this data is the 13th of the second month, I use the 12th as the starting point.

rate in the third month of a quarter. In Appendix C.1, I specify the transformation I use for each variable included in the analysis.

In the empirical setting, I denote the forecast based on data for the period up to and including month t by $E_t[\bullet]$. Note that in my real-time setting, this data is only available in month $t+1$. I index the surveys by $s = 1, 2, \ldots$, and denote the first month of the quarter in which survey s is conducted by t_s (i.e., $t_s = t_1 + 3s - 3$, where t_1 corresponds to October 1968). The SPF deadline lies in the second month of each quarter, which means that forecasts for survey s are based on data up to and including month t_s. As values in the third month of the quarter correspond to quarterly observations, the SPF provides observations of the individual forecasts $E^i_{t_s}[\widetilde{x}_{t_s+h}]$ for horizons $h = 2, 5, \ldots, 14$. For $h = 2$, this corresponds to a nowcast of the quarter in which the survey is conducted. For $h = 14$, this corresponds to a 1-year ahead forecast.

4.2.4 Computing empirical forecasts and beliefs

I can obtain estimates for the coefficients in the empirical forecasting equation (4.27) by estimating the empirical transition function (4.25). However, the macroeconomic factors and mean market belief are not observable in practice. To best reproduce the information set available to forecasters at the time, I estimate both in real-time.

To estimate the factors, I use the real-time panel described in subsection 4.2.2. First, I transform all variables to approximate stationarity, using 3-month differences or growth rates. (The transformations are described in Appendix C.1.) I then remove outliers, identified as having a distance to the median that is more than ten times the interquartile range.[20] Next, I standardize each variable to have zero mean and unit standard deviation. Because of the unbalancedness of the panel, I cannot directly use PCA. Instead, I use a variant of the expectation-maximization (EM) algorithm (Stock and Watson, 2002), which starts from a balanced panel with chosen values in place of the missing observations. After identifying principal components using ordinary PCA, these components are used to predict the missing values by least squares. A new panel with the predicted values in place of the chosen starting values for the missing observations is then used for another round of PCA (after restandardization). This procedure is repeated until convergence. Instead of running the algorithm until the predicted missing values converge, I stop when the principal components reach their maximum explanatory power of the non-missing values.[21] To find appropriate starting values, I use short EM runs with random initialization (Biernacki et al., 2003; Schumacher and Breitung, 2008). I indicate the estimate of month t factors based on data up to and including month τ by $E^m_\tau[F_t]$, and assume it is known to all agents.

At the start of the SPF, no data on market belief was available to the forecasters. One could argue that forecasters were not yet aware of (the impact of) the beliefs of their colleagues at the time. I assume that forecasters only started taking into account market belief when a reasonable amount of forecasting data was available from the SPF. For the first ten years of the survey (until 1978Q3), I calculate the empirical forecast without considering market beliefs.

[20] I use this outlier definition throughout the chapter.

[21] Convergence takes a long time, while a regression of the non-missing values on the principal components reaches the maximum average R^2 much quicker. This significantly increases the computational efficiency of the analysis, since I have to run the algorithm at least once for each survey.

Consider a survey s in those first ten years ($1 \leq s \leq 40$). I first calculate the factor estimates $E^m_{t_s}[F_t]$ for $t \leq t_s$. I then have to make a distinction between monthly and quarterly variables. For monthly variables, I use data up to and including month t_s to estimate by OLS[22] the empirical transition function (4.25) for horizons $h = 2, 5, \ldots, 14$, where I replace F_t by its estimate:[23]

$$\widetilde{x}_{t+h} - \widetilde{x}_t = c_{hx} + \lambda_{hx}\widehat{x}_t + A^F_{hx}E^m_{t_s}[F_t] + u^x_{t+h}. \tag{4.32}$$

To better approximate the empirical distribution, I remove outliers before estimation. I do this for all empirical transition function estimations. I then compute the empirical forecasts

$$E^m_{t_s}[\widetilde{x}_{t_s+h}] = \widetilde{x}_{t_s} + \widehat{c}_{hx} + \widehat{\lambda}_{hx}\widehat{x}_{t_s} + \widehat{A}^F_{hx}E^m_{t_s}[F_{t_s}], \tag{4.33}$$

where I denote the estimated parameters by $\widehat{c}_{hx}$, $\widehat{\lambda}_{hx}$, and $\widehat{A}^F_{hx}$.[24]

For quarterly variables, $\widetilde{x}_t$ is only observable in the third month of each quarter. Slightly abusing notation, I use $\widehat{x}_{\tau_q} = \widetilde{x}_{\tau_q} - \widetilde{x}_{\tau_{q-1}}$ for quarterly variables, where τ_q is the third month of quarter q. I estimate

$$\widetilde{x}_{\tau_q+h+1} - \widetilde{x}_{\tau_q} = c_{hx} + \lambda_{hx}\widehat{x}_{\tau_q} + A^F_{hx}E^m_{t_s}[F_{\tau_q+1}] + u^x_{\tau_q+h+1} \tag{4.34}$$

for horizons $h = 2, 5, \ldots, 14$. Note that I incorporate the factor estimate for the first month of the survey quarter, which is available at the time. I compute the empirical forecast for these quarterly variables as

$$E^m_{t_s}[\widetilde{x}_{t_s+h}] = \widetilde{x}_{t_s-1} + \widehat{c}_{hx} + \widehat{\lambda}_{hx}\widehat{x}_{t_s-1} + \widehat{A}^F_{hx}E^m_{t_s}[F_{t_s}], \tag{4.35}$$

Now consider a survey conducted after the initial period of ten years ($s \geq 41$). I use the empirical forecasts and individual forecasts for all past surveys to compute

$$E^i_{t_k}[\widetilde{x}_{t_k+h}] - E^m_{t_k}[\widetilde{x}_{t_k+h}] = A^g_{hx}g^i_{t_k} + \rho^{ihx}_{t_k}, \quad k = 1, \ldots, s-1. \tag{4.36}$$

This computation results in a panel of forecast differences with columns corresponding to variable-horizon combinations (105 in total) and rows corresponding to forecaster-survey combinations (about 40 forecasts per survey, so $40(s-1)$ rows in total). Because some forecasts are missing, I cannot use PCA directly. Instead, I use the EM algorithm described above to estimate the individual belief factors $g^i_{t_k}$.[25] Before running the algorithm, I remove outliers and standardize all forecast differences to have zero mean and unit standard deviation. I can use the whole panel instead of running the algorithm for each survey or agent separately, since the coefficients A^g_{hx} are constant across agents and time.

[22] All estimations in this chapter are by OLS.

[23] I estimate the transition function for each horizon separately instead of iterating forward the one-step-ahead forecast. I motivate this choice at the end of this subsection.

[24] For some monthly variables, $\widetilde{x}_{t_s}$ is not observable at the time of the survey. For those variables, I replace $\widetilde{x}_t$ and $\widehat{x}_t$ in (4.32) by $\widetilde{x}_{t-1}$ and $\widehat{x}_{t-1}$, and calculate the empirical forecast with $\widetilde{x}_{t_s-1}$, $\widehat{x}_{t_s-1}$, and $E^m_{t_s}[F_{t_s}]$ as in (4.35).

[25] I do not use the factor loadings A^g_{hx} and variable-horizon-specific beliefs ρ^{ihx}_t identified by the EM algorithm for further analysis, because it fills in missing values with predictions from the factor model. Instead, I obtain them by regressing the standardized forecast differences on the estimated belief factors without the missing values. Because I do not need these parameters for constructing the empirical forecast, I only do this for the final estimate (see subsection 4.3.1).

Averaging the belief factors across agents for each survey results in a time series for mean market belief:

$$Z_{t_k} = \frac{1}{N_k}\sum_i g^i_{t_k}, \quad k = 1, \ldots, s-1, \tag{4.37}$$

where N_k is the number of participants for survey k. I denote the estimate of month t mean market belief, based on data up to and including month τ, by $E^m_\tau[Z_t]$. It is known to all agents. I then estimate the following transition functions for monthly (4.38) and quarterly (4.39) variables:

$$\widetilde{x}_{\tau_q+h+1} - \widetilde{x}_{\tau_q+1} = c_{hx} + \lambda_{hx}\widehat{x}_{\tau_q+1} + A^F_{hx}E^m_{t_s}[F_{\tau_q+1}] + A^Z_{hx}E^m_{t_s}[Z_{\tau_q-2}] + u^x_{\tau_q+h+1}, \tag{4.38}$$

$$\widetilde{x}_{\tau_q+h+1} - \widetilde{x}_{\tau_q} = c_{hx} + \lambda_{hx}\widehat{x}_{\tau_q} + A^F_{hx}E^m_{t_s}[F_{\tau_q+1}] + A^Z_{hx}E^m_{t_s}[Z_{\tau_q-2}] + u^x_{\tau_q+h+1}. \tag{4.39}$$

I use the coefficient estimates, indicated with hat superscripts, to compute the monthly (4.40) and quarterly (4.41) empirical forecasts:

$$E^m_{t_s}[\widetilde{x}_{t_s+h}] = \widetilde{x}_{t_s} + \widehat{c}_{hx} + \widehat{\lambda}_{hx}\widehat{x}_{t_s} + \widehat{A}^F_{hx}E^m_{t_s}[F_{t_s}] + \widehat{A}^Z_{hx}E^m_{t_s}[Z_{t_s-3}], \tag{4.40}$$

$$E^m_{t_s}[\widetilde{x}_{t_s+h}] = \widetilde{x}_{t_s-1} + \widehat{c}_{hx} + \widehat{\lambda}_{hx}\widehat{x}_{t_s-1} + \widehat{A}^F_{hx}E^m_{t_s}[F_{t_s}] + \widehat{A}^Z_{hx}E^m_{t_s}[Z_{t_s-3}]. \tag{4.41}$$

Notes on the empirical distribution In constructing the empirical forecast, I mostly adopt the approach of Kurz and Motolese (2011). First, they note that in theory, the empirical distribution should be time-invariant, and could be estimated over any long period of time. Because real datasets are relatively short, they re-estimate the coefficients of the forecasting equation in real-time instead of estimating them once over the full sample. Given the changing nature of my real-time dataset, I adopt the same approach of computing empirical forecasts in real-time.

Second, Kurz and Motolese highlight the importance of using only a small number of predictors. One way I minimize the amount of variables is by using real-time macroeconomic factors to summarize the information set available at the time. Recognizing the persistence of most economic time series, I include a lag of the dependent variable, but only one. This does mean that the model is more likely to be misspecified, meaning that a direct forecasting model might perform better than an iterated multi-step forecast (Marcellino et al., 2006). I therefore estimate the transition function for each horizon separately, even though the coefficients in equation (4.25) are all defined in terms of the coefficients in the one-month ahead model (see Appendix C.2). Kurz and Motolese use the same approach.

Similarly, I do not restrict A^g_{hx} and ρ^{ihx}_t in terms of their $h = 1$ values, but use the data for all horizons to get a more accurate estimate of the belief factors. I also estimate A^F_{hx} and A^Z_{hx} as separate coefficients, even though strictly speaking, A^F_x determines the impact of both the factors and mean market belief on x over all horizons (given A_F, A_Z, and A^Z_F).

The inclusion of real-time estimates of mean market belief in the empirical transition function is in line with the theory, but increases the number of predictors, decreases the sample size, and differs from the approach by Kurz and Motolese. As a robustness check, I follow their approach, and derive beliefs only at the end of the sample. This means that mean market belief is not explicitly included in the construction of the empirical forecast: I use equation (4.35) to compute empirical forecasts for all surveys and then extract beliefs using equation (4.36) once. The assumption is that it is implicitly included in the factors extracted from the real-time macroeconomic panel, for example through the consumer sentiment index.

The linearity of the empirical transition function might not be well-suited for the period in which the zero lower bound (ZLB) on interest rates is binding. As a further robustness exercise, I exclude surveys conducted after 2008Q3 to avoid this period.

4.2.5 Analysis of beliefs

Dynamics

After constructing the empirical forecast for all surveys, I compute a final estimate of individual (variable-horizon-specific) beliefs. I denote them by $\widehat{g}^i_{t_s}$ and $\widehat{\rho}^{ihx}_{t_s}$. Mean market belief $\widehat{Z}_{t_s}$ directly follows.

For each survey s, I use the real-time estimates of the macroeconomic factors and mean market belief to estimate the following version of the empirical transition function of the factors (4.16):

$$E^m_{t_s}[F_{t_{k+1}}] = \bar{A}_F E^m_{t_s}[F_{t_k+2}] + \bar{A}^Z_F E^m_{t_s}[Z_{t_k}] + u^F_{t_{k+1}}. \tag{4.42}$$

Again, I only include mean market belief after the SPF has run for ten years.[26] I then estimate a quarterly version of the transition function of individual beliefs (4.5):

$$\widehat{g}^i_{t_s} = \bar{A}_Z \widehat{g}^i_{t_{s-1}} + \bar{A}^F_Z \widehat{u}^F_{t_s} + u^{ig}_{t_s}, \tag{4.43}$$

where $\widehat{u}^F_{t_s}$ denotes the time t_s residual resulting from the estimation of equation (4.42) for survey s. The estimated coefficients provide insight into the dynamics of beliefs. In particular, the size and significance of the entries of $\bar{A}^F_Z$ provide insight in the importance of learning.

Correlations

A crucial role in Rational Beliefs theory is played by the correlation of beliefs across agents, or more precisely, the correlations between the innovations ϵ^{ig}_{t+1} in the transition function of beliefs (4.5). To investigate these correlations, I collect the residuals $\widehat{u}^{ig}_{t_s}$ resulting from the estimation of the empirical transition function (4.43) in a panel. Because the forecasters participate in only part of the surveys included in the sample, the panel contains many missing values. I remove forecasters that participated in less than ten surveys, which leaves 176 forecasters with three residuals each (one for each dimension). This results in a panel of 528 forecaster-residual variables with observations in the period 1969Q1 – 2019Q4. The number of forecaster-residual observations for each survey varies between 21 and 141, with a mean of 84.

Instead of separately computing correlations for each dimension and for each pair of forecasters, I use the EM algorithm to extract three principal components. The

[26] When mean market belief is not included, I run a monthly estimation of

$$E^m_{t_s}[F_{t+1}] = \bar{A}_F E^m_{t_s}[F_t] + u^F_{t+1}.$$

.

variance explained by each of these components summarizes the correlation along each of the three belief dimensions across all agents. The more variance is explained by the components, the more correlated the residuals are.

Testing implications of rationality

I investigate the extent to which the survey participants satisfy the main axiom of Rational Beliefs theory by checking the validity of the implications that I introduce in section 4.2.1. One could test the rationality of individual forecasters. However, because forecasters participate in only a limited number of surveys, this would mean that distributional properties have to be computed over a limited timespan. Since the axiom refers to long-run properties, I choose to look at rationality of the market as a whole in order to use the full 1968Q4 – 2019Q4 range of surveys.

First, I test whether beliefs have zero long-run averages. Because this should hold for each survey participant, it should also hold for the cross-sectional mean. Denoting by N_s the number of participants of survey s that forecast x over horizon h, and by S the total number of surveys that include variable-horizon combination xh, this implies that

$$\frac{1}{S}\sum_{s=1}^{S}\frac{1}{N_s}\sum_{i=1}^{N_s}\left(A_{hx}^{g}g_{t_s}^{i}+\rho_{t_s}^{ihx}\right) \tag{4.44}$$

is 'close' to zero. I compute this average using estimates for the beliefs that follow directly from the difference between the survey and estimated empirical forecasts for all surveys (4.36). Note that by averaging the cross-sectional mean over time, I give less weight to beliefs in a survey with more respondents. By doing this, I avoid putting more (less) weight on surveys with more (less) respondents than average. It could be the case that beliefs are extreme at the time of a survey with many respondents, which would bias the overall mean. I treat the estimated mean as resulting from S observations, and therefore use a t-test with $S-1$ degrees of freedom to test whether the computed average is significantly different from zero.

Second, I test whether the variance of beliefs is bounded by the variance of the empirical distribution. For the empirical variance σ_{hx}^2, I use the estimate that follows from the estimation of the empirical transition functions (4.38) and (4.39) on the last vintage.[27] Again, I use the final belief estimates $A_{hx}^{g}g_{t_s}^{i}+\rho_{t_s}^{ihx}$ for all surveys s, horizons h, variables x and forecasters i. In line with the rationality axiom, I assume that the long-run average is zero, and estimate the variance as[28]

$$\frac{1}{S-1}\sum_{s=1}^{S}\frac{1}{N_s}\sum_{i=1}^{N_s}\left(A_{hx}^{g}g_{t_s}^{i}+\rho_{t_s}^{ihx}\right)^2. \tag{4.45}$$

As with calculating the mean, I put a lower weight on beliefs in a survey with more respondents to avoid a bias in the estimated variance. Again, I treat the estimate as resulting from S observations, and therefore use a (one-sided) χ^2-test with $S-1$

[27] Note that the variance estimate is different for each real-time vintage. I use the variance estimated on the final vintage, because this is the best estimate of the long-run empirical variance.

[28] As a robustness check, I compute the variance where instead of assuming zero mean, I use the time-averaged cross-sectional mean (4.44). The results are qualitatively similar.

degrees of freedom to test whether beliefs variance is significantly larger than the empirical variance.

I conduct both tests for all variables and horizons and exclude outliers.

4.2.6 Macroeconomic impact

I use two vector autoregressive (VAR) approaches to assess the role that market belief plays in macroeconomic fluctuations. The first captures the dynamics of mean market belief (Z_t), the output gap (y_t), inflation (π_t), and the interest rate (i_t). Here, subscript t indicates the quarter instead of month. I use my final estimate of mean market belief, measure inflation in terms of the chain-weighted GDP price level, and use the effective federal funds rate as the interest rate. I use four lags, which is standard in quarterly VARs.

After maximum likelihood estimation of the VAR, I rely on a Cholesky decomposition of the covariance matrix to identify structural shocks, where I use the following recursive ordering:

$$\begin{pmatrix} Z_t \\ y_t \\ \pi_t \\ i_t \end{pmatrix}. \tag{4.46}$$

I order market belief first for two reasons. First, mean market belief is measured in the middle of the quarter, while the other variables cover the whole quarter. Second, market belief already takes into account most publicly available information about changes to economic conditions within the quarter through the empirical forecast. The ordering of the other variables is standard.

For each dimension of mean market belief, I compute impulse responses to a belief shock to identify the impact on the macroeconomic variables. I also generate forecast error variance decompositions to quantify the importance of those shocks in macroeconomic fluctuations.

The second approach uses macroeconomic factors in addition to the three macrovariables used in the first approach, resulting in a factor-augmented VAR (FAVAR). I estimate the factors using the EM algorithm on the FRED-MD database. Before running the algorithm, I transform the variables in the panel to stationarity in a way that ensures that observations in the third month of a quarter correspond to a 3-month difference or growth rate in a quarterly average. I also remove outliers and standardize all variables to have zero mean and unit standard deviation. I then use the estimated factors for the third month of each quarter as the quarterly observations.

I use the following ordering to identify structural shocks:

$$\begin{pmatrix} Z_t \\ F_t \\ y_t \\ \pi_t \\ i_t \end{pmatrix} \tag{4.47}$$

I order the factors before the macroeconomic variables following Bernanke et al. (2005). Market belief is ordered first, with the same motivation as for the simpler

VAR. I construct impulse responses and forecast error variance decompositions to assess the impact on both the factors and the macroeconomic variables.

4.3 Results

From the real-time macroeconomic panel, I extract four factors ($r = 4$). Kurz and Motolese (2011) also use four factors to construct the empirical forecast. Together, these factors explain between 40% and 60% of the variation in the variables included in the panel, depending on the quarter in which the factors are calculated. From the fifth factor onwards, the marginal explained variance is lower than 5%.[29]

I use the extracted macroeconomic factors to compute the empirical forecasts in real-time. From 1978Q4 onwards, I compute differences between survey and empirical forecasts across variables and horizons and extract three belief factors ($q = 3$). These three factors explain roughly half of the variation in the forecast differences. The marginal contribution of each of these factors is more than 10%, with the contribution dropping below 10% for the fourth factor.[30]

The coefficients on mean market belief in the empirical forecasting equations (A^Z_{hx} in equations (4.38) and (4.39)) are significant for some variables and horizons, but not for all. For example, for the CPI inflation rate, the first and third belief factors are significant across horizons. On the other hand, only the first belief factor significantly contributes to forecasting real GDP, and only when forecasting the current quarter. Adding mean market belief does result in a larger adjusted R-squared for most variables and horizons.

Computing all empirical forecasts leads to a final (2019Q4) panel of forecast differences for 21 variables over 5 horizons, with 8,092 rows. In Figure 4.1 and Figure 4.2, I plot the average and standard deviation across survey respondents of the 1-quarter-ahead forecast differences for real output, inflation, and the interest rate over time. A positive value of real output belief means that on average, survey respondents are optimistic about economic growth. A negative value means that they are pessimistic. For inflation and interest rate beliefs, a positive (negative) value indicates that forecasters expect a value that is higher (lower) than the empirical forecast.

Market beliefs regarding these variables appear to be persistent. Output beliefs are 0.59 percentage points higher during recessions on average ($p < 10^{-3}$), indicating that forecasters expect a less negative impact or stronger recovery than under the empirical distribution. Inflation beliefs, on the other hand, tend to be negative during recessions: survey respondents expect inflation to be 0.72 percent lower during recessions on average ($p = 0.001$). This result is mostly driven by the large, negative inflation beliefs during the recession in the middle of the '70s. Interest rate beliefs are almost 75 basis points higher during recessions. Interestingly, this result is robust to excluding the zero-lower-bound period after 2008Q3.[31]

To gain more insight into the historical dynamics of beliefs, I zoom in on the stagflationary period from 1973 up to 1983 and on the period since the Volcker

[29] The PC_{p2} information criterion used by McCracken and Ng (2016) does not help to determine the optimal number of factors, as it does not reach a minimum for a number of factors that is lower than 15. I therefore use a small number of factors that still captures the majority of the variation in the panel.

[30] As with the macroeconomic panel, an information criterion does not help to determine the optimal number of factors.

[31] Mean beliefs about 3-month treasury bill rates are not significantly different during the ZLB period.

disinflation (from 1985 onwards). During this first period, inflation beliefs are 1.18 percentage points lower than during the rest of the sample ($p < 10^{-3}$). In the later period, they are 1.15 percentage points higher ($p < 10^{-3}$). Forecasters believed inflation to go down faster than justified by the data. After inflation got under control, they still expected it to be relatively high, apparently not convinced yet by the effectiveness of U.S. monetary policy. Kurz and Motolese (2011) also find positive inflation beliefs on average after the Volcker disinflation. Output beliefs are 26 basis points lower after the disinflationary period ($p = 0.035$). I cannot distinguish these periods for interest rate beliefs, because data is only available from 1981Q3 onwards.

Disagreement across forecasters, as measured by the standard deviation of the forecast differences, increases during recessions.

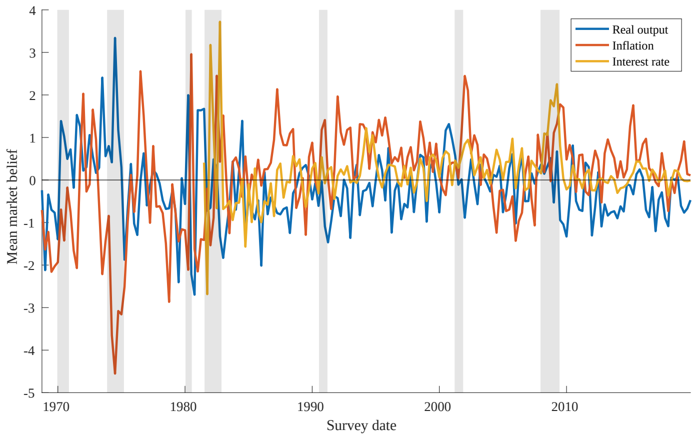

Figure 4.1: Average across survey respondents of the difference between survey and empirical 1-quarter ahead forecasts for real GDP, the inflation rate in terms of the chain-weighted GDP price index, and the 3-month treasury bill rate. Real output differences are logged and reported in percentages. Inflation is measured in annualized quarterly rates. Grey bars indicate NBER recessions.

4.3.1 Belief factors

Based on the final estimate of beliefs, I find that the first belief factor explains 21% of the variance in the forecast differences, the second 18%, and the third 12%. Together, they explain more than half of the variation. This means that the variable-horizon-specific beliefs account for almost half of the fluctuationsin sentiment.

To interpret the role of the factors in the fluctuations of beliefs about each variable-horizon separately, I regress all forecast differences on the three belief factors for each variable-horizon combination. This amounts to estimating equation (4.36), which expresses the forecast differences for variable x over horizon h as the corresponding factor loadings A^{g}_{hx} times the belief factors g^{i}_{t} plus the variable-horizon-specific beliefs ρ^{ihx}_{t}. The ρ^{ihx}_{t} correspond to the idiosyncratic errors. I standardize the forecast differences and belief factors to have zero mean and unit variance prior to regression, so that the loadings measure the factors' impact in terms of standard deviations.

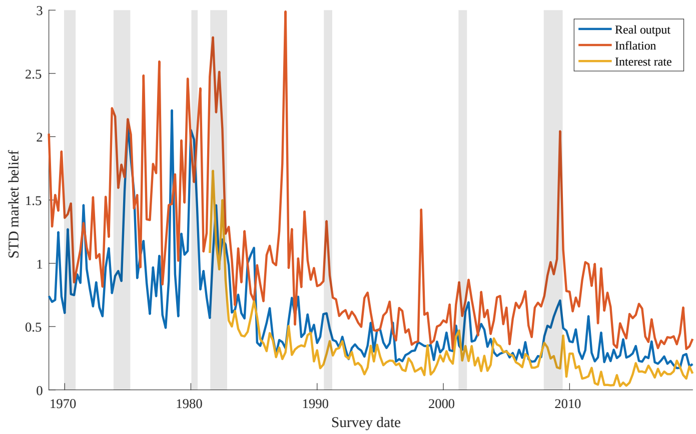

Figure 4.2: Standard deviation across survey respondents of the difference between survey and empirical 1-quarter ahead forecasts for real GDP, the inflation rate in terms of the chain-weighted GDP price index, and the 3-month treasury bill rate. Real output differences are logged and reported in percentages. Inflation is measured in annualized quarterly rates. Grey bars indicate NBER recessions.

In Figure 4.3, I present the R-squared for each regression, measuring the fraction of variation in the variable-horizons that is explained by the belief factors. The full names of each of the variables can be found in Appendix C.1. The figure shows large dispersions between variables. Some variables, like government consumption and investment, are poorly explained by the belief factors, while others, like output and inflation, are well explained by them. Overall, most variance is explained of the three middle horizons, almost 80% for some variables.

In Figure 4.4 and Figure 4.5, I present the factor loadings A_{hx}^{g} for each variable-horizon combination. A positive (negative) loading indicates that the forecast difference increases (decreases) in response to an increase in the respective belief factor. The figure also indicates whether the loading significantly differs from zero. For example, a one-standard-deviation increase in the first factor significantly increases one-year-ahead core CPI beliefs by about one standard deviation (see left panel of Figure 4.4).

Ceteris paribus, the first factor mostly moves mid- to long-term forecast differences of inflation and interest rates, with large negative impacts on residential investment, housing, and corporate profits. Because it moves inflation and real activity (as measured by real consumption and real GDP) in opposite directions, this seems to refer to supply-side beliefs. The fact that housing beliefs negatively load on the first factor might have to do with the positive relation with interest rate beliefs. The increase in interest rate beliefs is presumably due to an expected tightening of monetary policy to battle inflation. Two deviations from the supply-side interpretation are increases in employment and decreases in unemployment beliefs. These deviations mostly concern short-term beliefs however, while the inflation and real activity loadings mostly concern long-term beliefs. Herbst and Winkler (2020) find a similar supply-side factor in their analysis of disagreement using the SPF. Opposed

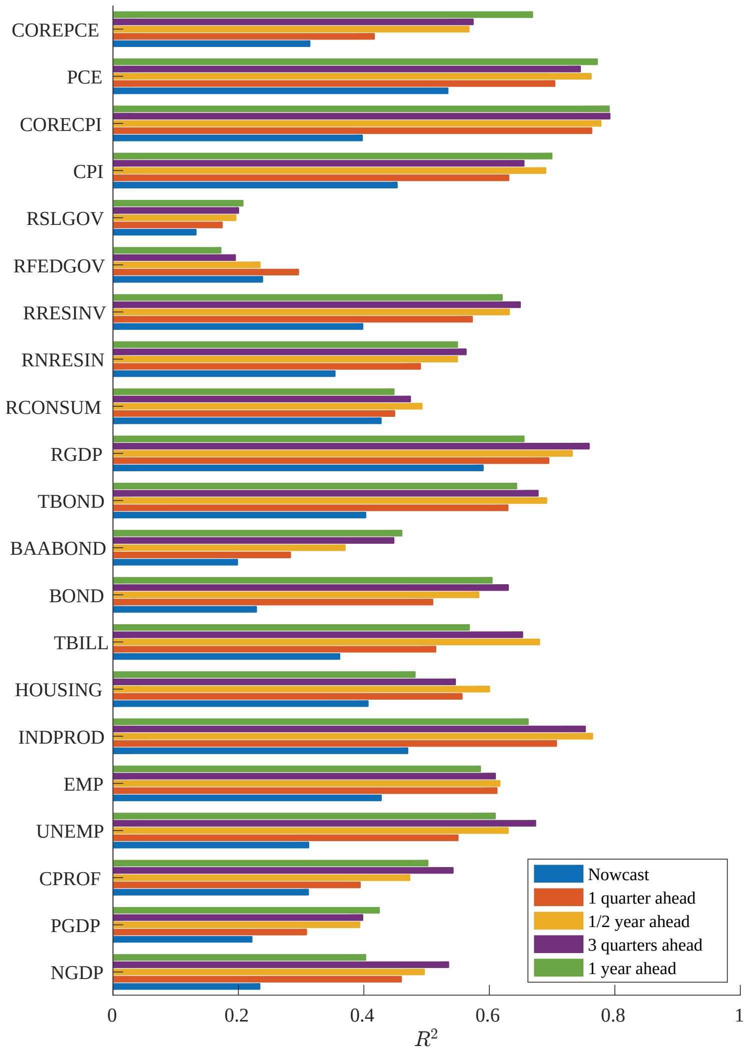

Figure 4.3: R-squared values for regressions of each variable-horizon forecast difference on all three dimensions of belief.

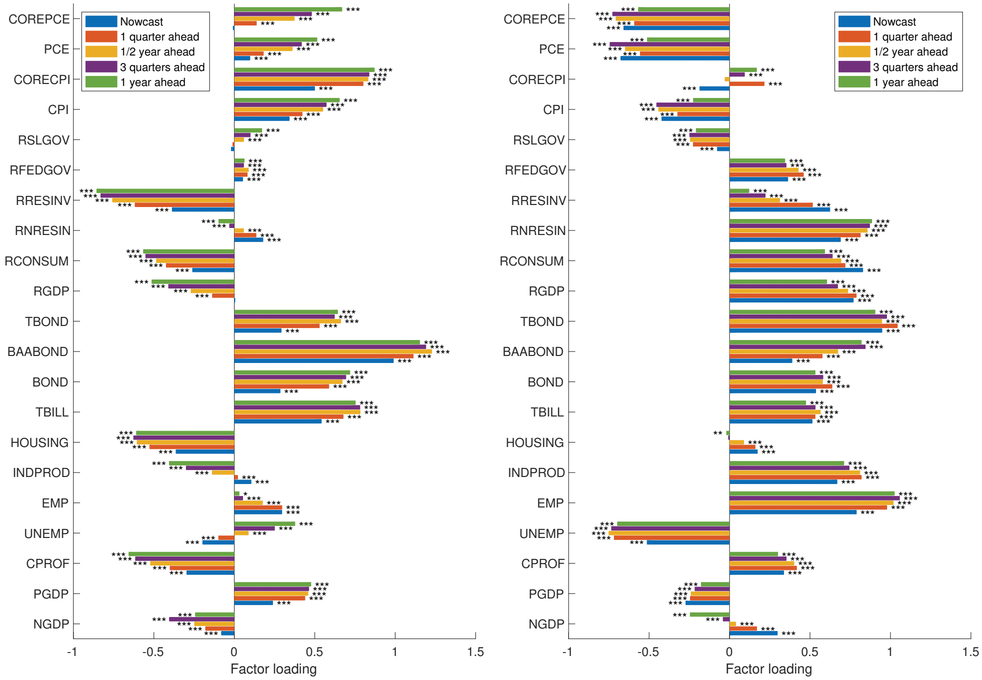

Figure 4.4: Factor loadings of forecast differences on the first (*left*) and second (*right*) belief factor for all variable-horizon combinations. Both forecast differences and belief factors are standardized to have unit variance, so the loadings are measured in terms of standard deviations. Asterisks indicate significance: *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

to my approach, they also consider long-term forecasts (over a 10-year horizon) and find that long-term GDP and productivity forecasts load negatively on this factor. In the context of their semi-structural model, they interpret this pattern as disagreement about permanent productivity shocks. Because of the negative loading of corporate profits, they disregard an interpretation as disagreement about mark-up shocks.

The second belief factor also moves inflation and real activity in opposite directions, and is therefore indicative of supply-side beliefs. However, its explanatory power is focused on the shorter horizons. To distinguish between the first and second belief factor, I name them long-term and short-term supply-side factor, respectively. Another important difference is that the second factor moves interest rate and real activity beliefs in the same direction. Given the opposite impact the factor has on inflation, these interest rate loadings do not seem in line with conventional monetary policy.

Both inflation and real activity load positively on the third factor, indicating that it represents demand-side beliefs. The negative loadings of (short-term) interest rates suggest that these beliefs are related to monetary policy. Housing and residential investment load positively on the third factor.

I plot the evolution of the three dimensions of mean market belief Z in Figure 4.6.

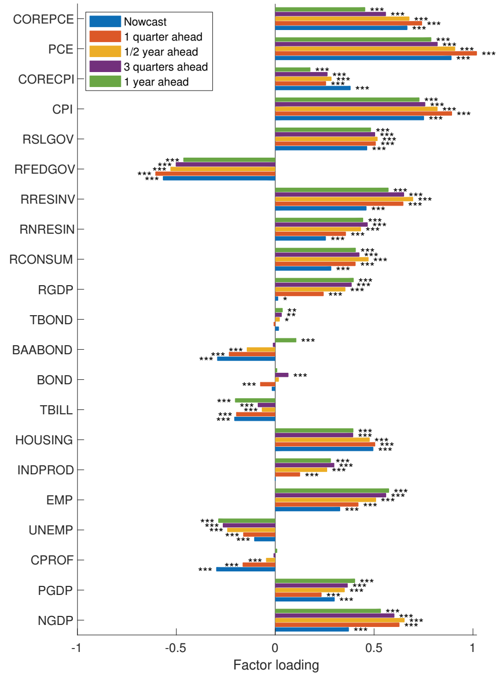

Figure 4.5: Factor loadings of forecast differences on the third belief factor for all variable-horizon combinations. Both forecast differences and belief factors are standardized to have unit variance, so the loadings are measured in terms of standard deviations. Asterisks indicate significance: *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

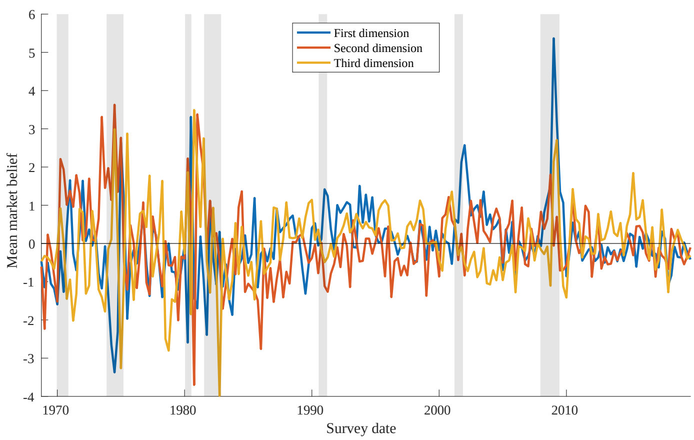

Figure 4.6: The three dimensions of mean market belief. Grey bars indicate NBER recessions.

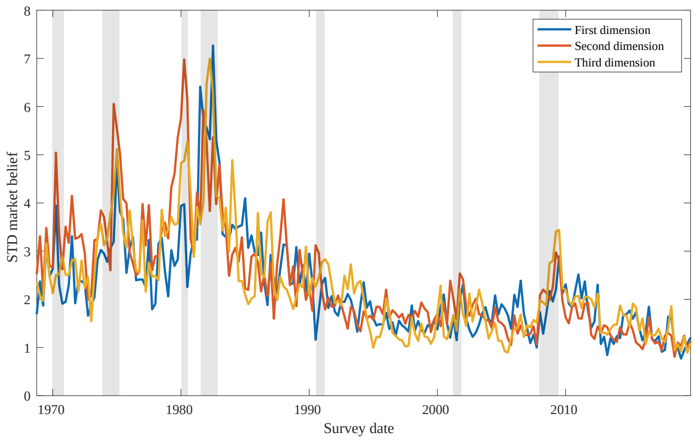

Figure 4.7: Cross-sectional standard deviation of three dimensions of market belief. Grey bars indicate NBER recessions.

I have standardized the time series to have zero mean and unit variance. With some exceptions, the first factor tends to be below average in the ’70s and early ’80s. In fact, it is 0.71 standard deviations lower in this period ($p < 10^{-3}$) compared to the rest of the sample. This is an interesting finding, because the supply-side inflation scenario that the first dimension represents fits this stagflationary period. It suggests that it took the forecasters a while to incorporate this scenario in their forecasts. The first factor keeps playing an important role after the Volcker disinflation, with peaks in each of the three recessions since the ’90s. On average, it is 0.74 standard devations higher after 1985 than before ($p < 10^{-3}$).

The second factor exhibits an opposite pattern: it is 0.70 standard deviations higher during the stagflationary period ($p = 0.006$), and 0.46 lower after the Volcker disinflation ($p = 0.002$). However, this short-term supply-side factor moves inflation and real activity in opposite directions to the long-term supply-side factor, so the long- and short-term beliefs are consistent. Interestingly, the second factor is significantly higher during recessions: about 0.9 standard deviations ($p < 10^{-3}$). This suggests that survey respondents are more optimistic about real activity during recessions, and they expect inflation to be lower than the empirical forecast. Their interest rate beliefs are more positive during recessions.

The third dimension of beliefs is slightly higher since the Volcker disinflation (0.35 standard deviations, $p = 0.019$), indicating that monetary-policy induced demand-side scenario gets more weight compared to empirical distribution.

All three dimensions of mean market belief are most volatile in the first part of the sample, and least volatile during the Great Moderation. Figure 4.7 shows the cross-sectional standard deviations over time. Disagreement along the three belief dimensions appear to move together. Again, forecaster disagreement increases during recessions. Moreover, there seems to be a long-term trend towards more consensus among the forecasters.

4.3.2 Belief dynamics

I present the estimates for the belief transition function (4.43) in Table 4.1. Slightly abusing previously introduced notation, I here denote by g_t^j the j-th dimension of belief at time t, not the belief of agent j. The first three columns, which correspond to $\bar{A}_Z$, show that the beliefs are persistent and that the three belief dimensions are interrelated. For example, having more positive beliefs along the first dimension at the time of survey s (putting more weight on the supply-side inflation scenario than the empirical distribution) is associated with a less optimistic economic outlook at the time of the next survey (dimension 2 for $s+1$). The effect size is of 0.14 is measured in terms of standard deviations of the belief dimensions.

Table 4.1: Belief dynamics

	$g^1_{t_{s-1}}$	$g^2_{t_{s-1}}$	$g^3_{t_{s-1}}$	$\widehat{u}^{F^1}_{t_s}$	$\widehat{u}^{F^2}_{t_s}$	$\widehat{u}^{F^3}_{t_s}$	$\widehat{u}^{F^4}_{t_s}$
$g^1_{t_s}$	0.432	0.071	0.028	−0.064	0.153	0.193	0.262
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.010)	(0.011)
$g^2_{t_s}$	−0.142	0.380	0.091	−0.305	−0.217	0.108	0.025
	(0.010)	(0.010)	(0.010)	(0.011)	(0.010)	(0.010)	(0.010)
$g^3_{t_s}$	0.131	−0.080	0.408	0.130	−0.176	0.224	−0.080
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)

Note. The first three columns together represent $\bar{A}_Z$ in the empirical transition function of beliefs (4.43). Here, $g^j_{t_s}$ denotes the j-th dimension of belief at the time of survey s. The last four columns represent $\bar{A}^F_Z$ in (4.43), where $\widehat{u}^{F^j}_{t_s}$ denotes the factor j residual resulting from estimation of the real-time factor transition function (4.42). Both beliefs and residuals are standardized to have unit variance before estimation of (4.43), so that the coefficients measure impact in terms of standard deviations. Standard errors are in parentheses.

The last four columns of Table 4.1 correspond to $\bar{A}^F_Z$, which determines the role of learning in the dynamics of beliefs. Most of its entries are highly significant and sizeable, indicating that learning plays an important role. Moreover, including the learning terms increases the average adjusted R-squared by 13.2%, from 18.4% without learning to 31.6% with learning. In line with the belief factors' distinct economic interpretations, each macroeconomic factor affects each belief factor differently.

4.3.3 Correlations

I now consider the correlations of the belief innovations ϵ_t^{ig}. Looking at the factors extracted from the panel of residuals resulting from the estimation of the belief transition function, I find that together, they explain 75% of the variance of the residuals. The individual components contribute 27%, 27%, and 21%. This supports the assumption that beliefs are correlated across agents, which plays a central role in Rational Beliefs theory.

4.3.4 Rationality

In Figure 4.8, I present the results for the rationality tests that I lay out in section 4.2.5. On the left-hand side, I plot the time-averaged cross-sectional mean of the difference between the survey and empirical forecasts for each variable-horizon combination. The figure also indicates whether the average is significantly different from zero. For example, the forecasters are pessimistic about real output on average, and more pessimistic over longer horizons. Over a 1-year horizon, the time-averaged log-difference between the average survey and empirical forecasts is about half a percentage point. I find larger averages (in absolute sense) for longer horizons for many of the variables.

The figure also shows that a zero average is the exception rather than the rule. This result is in line with earlier studies, and only means that the forecasters do not satisfy the rationality axiom over the time period that I analyse, at least for the individual variable-horizons. Rational Beliefs theory requires beliefs to have a zero time average in the long run.

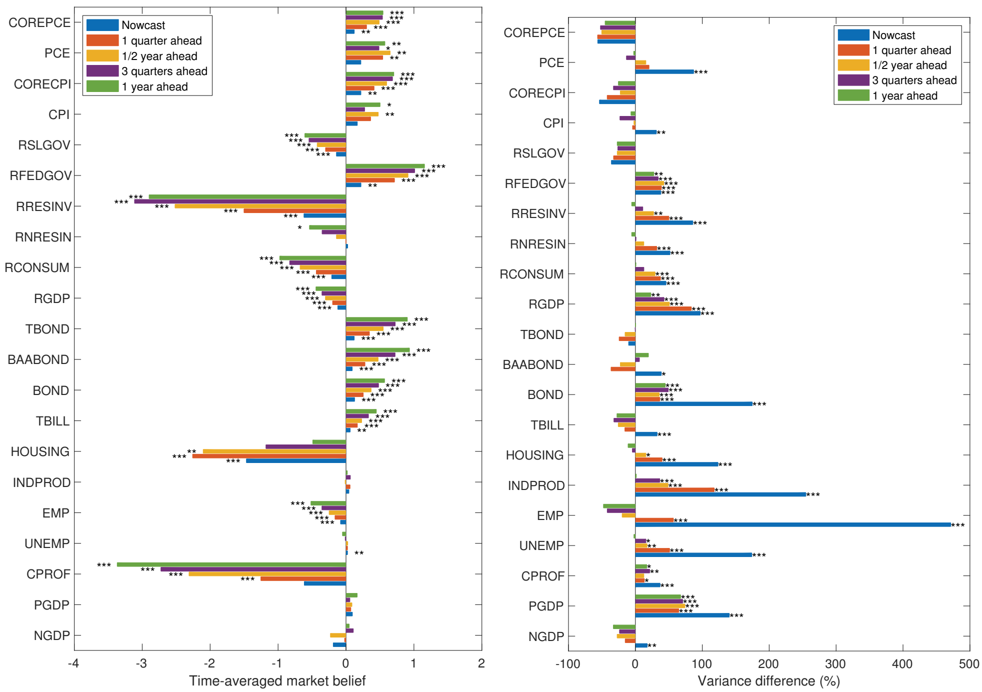

Figure 4.8: *Left*: The average difference between the survey and empirical forecasts as defined in (4.44). *Right*: The percentage-difference between belief variance as defined in (4.45) and the estimated empirical variance σ^2_{hx}. Asterisks indicate p-values for a two-sided t-test with null that time-averaged mean market belief is zero, and for a one-sided χ^2-test with null that belief variance is smaller than, or equal to the empirical variance: $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

On the right-hand side of Figure 4.8, I report the results for the variance test. The first result that stands out is that the belief variances for the nowcasts are much larger than the empirical variances for many of the variables. The largest difference is for the employment rate (EMP): almost 500%. It is important to note here that employment forecasts have only been included in the SPF since 2003Q4, which means that relatively few belief observations are available to estimate the variance. However, industrial production beliefs, which can be computed over the full 1968Q4 – 2019Q4 sample, also show a large variability over shorter horizons. This either suggests that irrational animal spirits play a larger role over shorter horizons, or that my estimation of the empirical distribution misses some aspects that are more important over shorter horizons. Both explanations seem to be at odds with the fact that on average, short-horizon survey forecasts are actually closest to the empirical forecast for most variables. Nevertheless, it might still be the case that the variability of beliefs is larger over short horizons, even though on average they are closer. Overall, I reject the null hypothesis that the variance is smaller than the empirical variance for 38 of the 105 variable-horizon combinations (at the one-percent level). The majority of beliefs appear to satisfy the bound on variance implied by the rationality axiom.

4.3.5 Macroeconomic impact

I present impulse response functions for all three dimensions of market belief in the VAR specification in Figure 4.9, Figure 4.10, and Figure 4.11. Here, Z^j indicates dimension j of mean market belief. The responses of market belief are in terms of standard deviations. The first dimension does not elicit any significant responses at the 5% level (Figure 4.9). The results do hint at a funds rate increase of about 25 basis points after a one-standard-deviation increase along the first dimension.[32]

A positive shock to the second belief dimension is associated with negative responses of both the output gap and the funds rate (Figure 4.10). A one-standard-deviation shock leads to a decrease of the output gap of about 0.3 percentage points, and to a decrease in the federal funds rate of about 0.5 percentage points, or two standard policy moves. Both responses are persistent, lasting for more than a year. Given the fact that such a shock is associated with an increase in optimism about the economic outlook, this is an interesting result. Kurz et al. (2018) provide a possible explanation for the negative output response. They point out that two opposing forces are at work. On the one hand, agents expect higher wages in the future, and therefore want to increase consumption today. On the other hand, they prefer to work less today, and work more tomorrow for a higher wage. Which of the two effects dominates, depends on beliefs and the monetary policy regime. The data suggests that the substitution effect dominates, leading to a drop in the output gap. There is no straightforward interpretation of the negative response of the interest rate, as Kurz et al. (2013) and Kurz et al. (2018) both show theoretically that monetary policy is complex in an economy with diverse beliefs. Interestingly, the interest rate increase is at odds with the increased interest rate beliefs due to the second factor shock.

The third dimension of belief is associated with a brief, small increase in the output gap, followed by a return to equilibrium (Figure 4.11). The federal funds rate responds with an increase of about one standard policy move (25 basis points). The increase becomes insignificant after half a year. In this case, optimism about real activity results in an increase in economic growth, suggesting that the income effect

[32] As I show below, the response is significant and almost twice as large in the FAVAR specification.

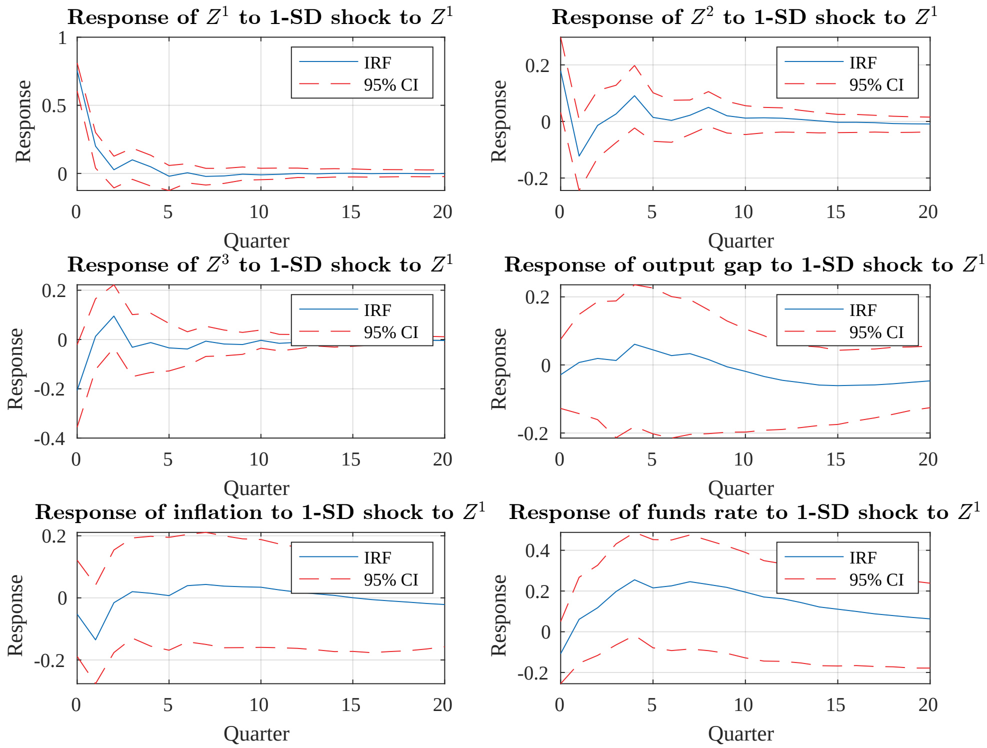

Figure 4.9: VAR impulse responses to a one-standard deviation positive shock to the first dimension of mean market belief. Dimension j of mean market belief is denoted by Z^j. The factors' responses are measured in terms of standard deviations. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

dominates, increasing consumption. Again, the funds rate impact is opposite to the impact on interest rate beliefs.

In Figure 4.12, I plot the sum of the forecast error variance contributed by each of the three belief factors for each variable included in the VAR. The variable that is most affected by the belief factors is the federal funds rate, about 40% of whose long-term fluctuations can be attributed to belief shocks. As can be seen in the individual FEVDs in Appendix C.3, this result is mostly driven by the second (real activity) dimension of beliefs. The contribution of beliefs to fluctuations of the output gap is smaller, but still substantial at 15%. The variance contribution for inflation is small, at about 6%.

For the FAVAR specification, I first use the FRED-MD panel to compute three macroeconomic factors with the EM algorithm. I use three factors here, because the marginal contribution to the average explained variance drops below 5% for the fourth factor.[33] The first factor accounts for 23% of the panel's variance, and is a real activity and employment factor. The second factor contributes a marginal 11% to the

[33] Again, using an information criterion does not help to find an optimal (small) number of factors.

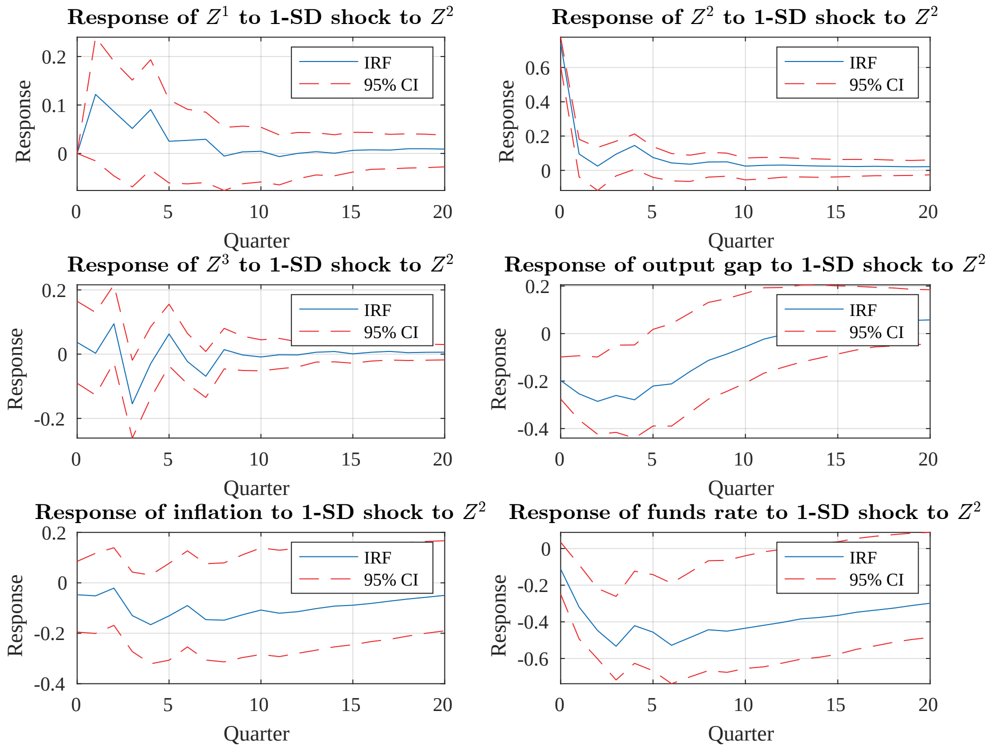

Figure 4.10: VAR impulse responses to a one-standard deviation positive shock to the second dimension of mean market belief. Dimension j of mean market belief is denoted by Z^j. The factors' responses are measured in terms of standard deviations. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

explained variance, and mostly drives housing and interest-rate(-spread)s. The third factor is an inflation factor that accounts for about 9% of the variance. Together, the three factors explain 43% of the panel's variance.

The FAVAR impulse responses of output, inflation, and the interest rate are similar to those in the VAR specification (see Appendix C.3). Two exceptions stand out. First, a one-standard-deviation increase in the first belief dimension is associated with a funds rate increase of about 0.4 percentage points. Second, the negative response of the output gap to a shock in the second belief factor is no longer significant at the 5% level. The macroeconomic factors themselves respond to belief shocks, but the responses are generally small (smaller than a 0.2 standard-deviations in absolute value).

The FEVD sums are similar for the VAR variables when the factors are included (Figure 4.13), but it is interesting to note that the belief shocks also play a substantial role in the fluctuations of the factors themselves. Most notably the first, real activity factor (15%) and second, interest rate factor (20%). This is in line with the VAR results.

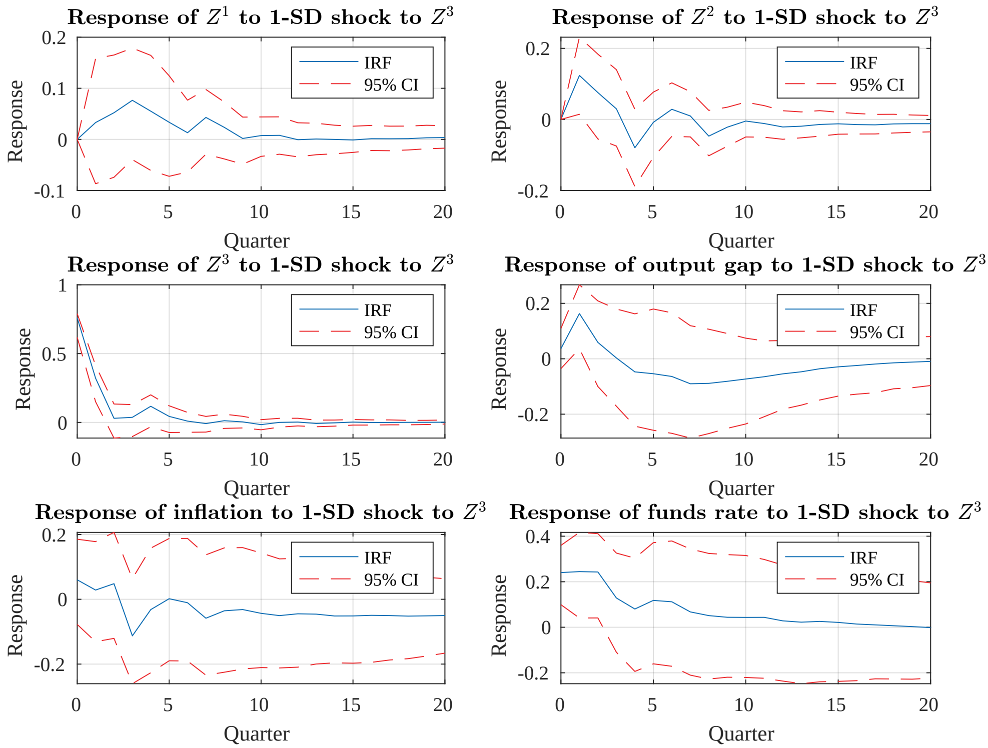

Figure 4.11: VAR impulse responses to a one-standard deviation positive shock to the third dimension of mean market belief. Dimension j of mean market belief is denoted by Z^j. The factors' responses are measured in terms of standard deviations. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

4.3.6 Robustness

The main conclusions of this chapter are robust to changing the construction of the empirical forecast, and to excluding certain parts of the sample from the analysis. In these deviations from the baseline specification, sentiment is still multidimensional, with distinct macroeconomic impacts along each dimension.

First, I discuss the results when leaving market beliefs out of the empirical transition function, and only including a lag of the observable and the real-time macroeconomic factors as predictors. The belief factor loadings are very similar, as well as the variance explained by the three belief factors together. The rationality tests give the same results. The average forecast differences are somewhat smaller, but there are still a large number of variable-horizons for which beliefs do not average to zero over time. The same is true for belief variances: they are slightly smaller, but significantly larger than the empirical variance for a similar number of variable-horizons.

The belief dynamics are also similar, both qualitatively and quantitatively. The belief factors are interrelated and learning explains a substantial part of their dynamics. Their relation to macroeconomic fluctuations, as identified by the (FA)VAR

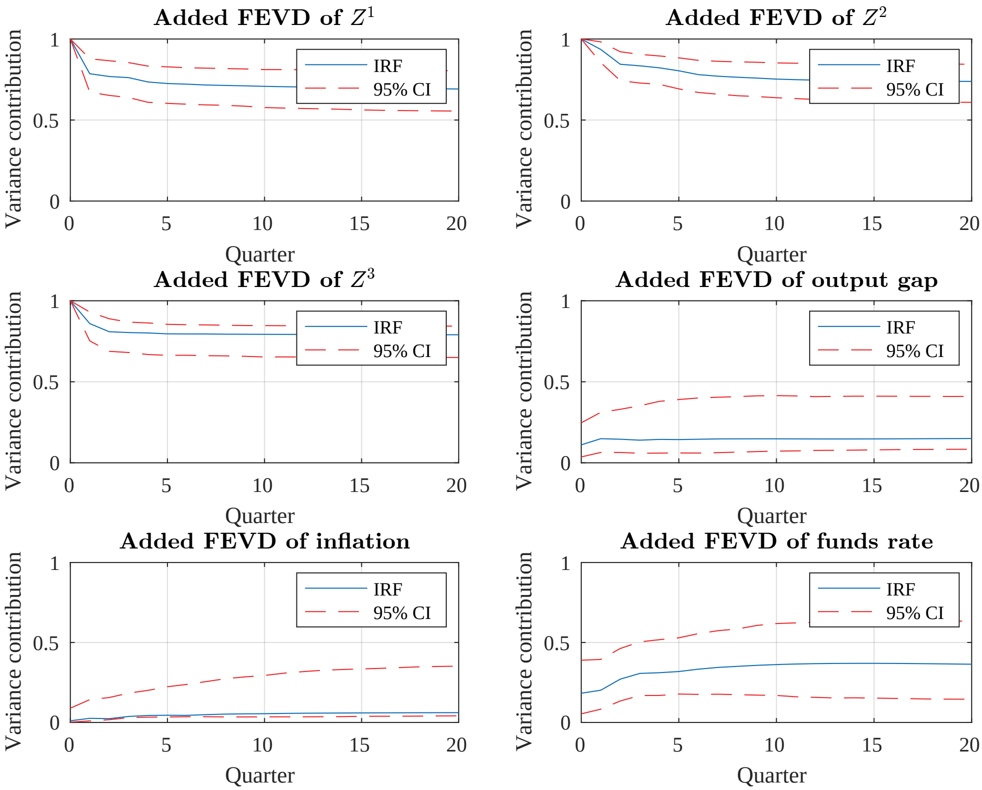

Figure 4.12: Sum of VAR FEVDs when beliefs are shocked. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

analysis, is also comparable. I highlight three differences. First, the first belief factor also has a significant impact on the funds rate in the VAR specification (of about 30 basis points), and its impact in the FAVAR specification is even larger than in my baseline results. Second, the impulse responses of the FRED factors are more significant, but of similar size and form. Third, the significant interest rate impact of the second belief factor in the baseline results for the FAVAR disappears in this alternative specification. However, beliefs still play an important role in fluctuations of the funds rate. Overall, the contribution of beliefs is somewhat smaller compared to the baseline, but still significant.

As a second robustness exercise, I only use surveys for which the survey deadline is known. This excludes surveys conducted before 1990Q3. I do not include mean market belief in the empirical forecast, because there is a limited history of surveys in this case. Moreover, the first robustness test indicates that it leads to similar conclusions. The three extracted belief factors together explain about half of the variance in all the forecast differences. The individual contributions of the factors are also similar to the baseline sample. There are differences in the factor loadings, however. The first factor is a demand-side factor, with positive loadings of real activity,

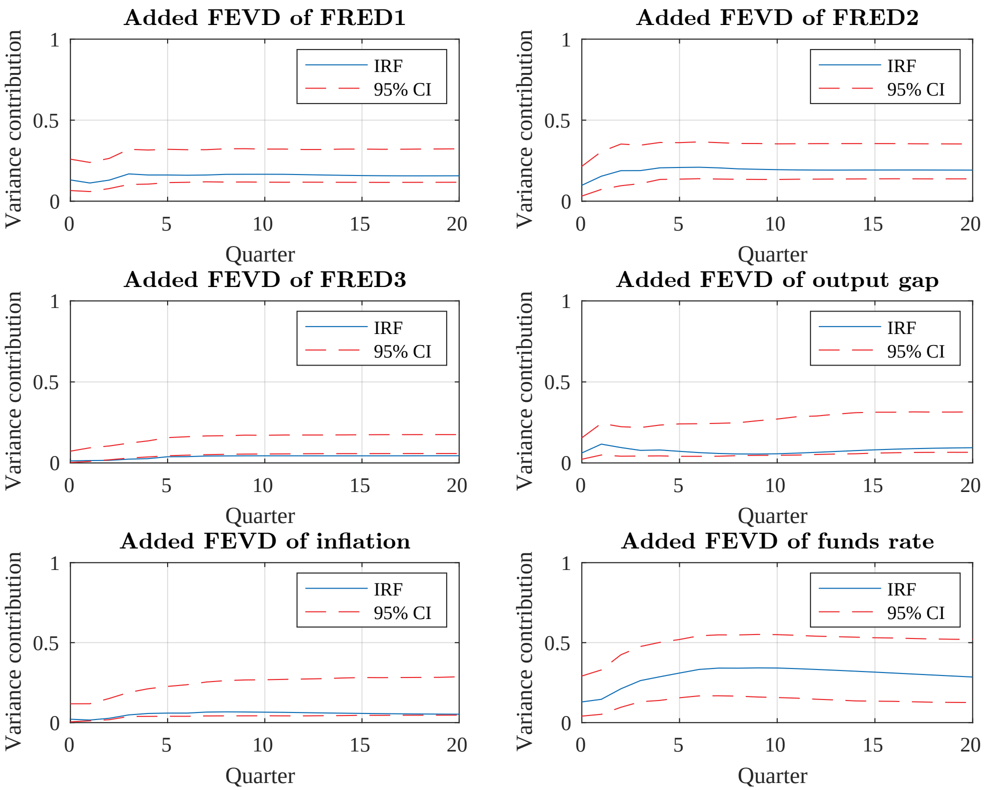

Figure 4.13: Sum of FAVAR FEVDs when beliefs are shocked. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications. FRED*i* denotes factor *i* of the FRED-MD panel.

inflation, and interest rate beliefs. It moves forecast differences regarding profits and housing in the opposite direction. The second factor is similar to the short-term supply-side factor in the baseline sample, moving output and the interest rate beliefs in the same direction, and beliefs about inflation in the opposite direction. The third factor is similar to the baseline's demand-side factor.

The factors' macroeconomic impact is much larger in the VAR specification, contributing 50% to the forecast error variance of the output gap, more than 20% to that of inflation, and 60% to that of the interest rate. The contributions are smaller in the FAVAR specification, but still sizable, at 25%, 15%, and 40%, respectively. These contributions are mostly driven by the second dimension of market belief, like in the baseline case. It causes a persistent decrease in the output gap, with a maximum impact of minus 40 basis points after six quarters. The funds rate impact is even larger: minus 50 basis points after six quarters.

For the last robustness exercise, I exclude data after 2008Q3 to avoid the ZLB. I exclude five SPF variables that are only available shortly before this cut-off point, namely, employment, Moody's baa corporate bond yield, core CPI inflation, PCE inflation, and core PCE inflation. Again, I extract three belief factors that together explain about 50% of the variance of the forecast differences. The first factor is similar to the second factor in the baseline case. It is difficult to categorize the other two factors, because they move similar variables in different directions. In terms of macroeconomic impact, however, the results are comparable: the total FEVD contribution of the belief factors is about 20% for the output gap, 30% for the interest rate, and 6% for inflation.

4.4 Conclusion

I study beliefs in the context of Rational Beliefs theory. For a broad selection of economic variables, I compute the difference between survey forecasts and non-judgmental forecasts based on data available at the time of the corresponding surveys. I study the factor structure underlying these differences, which constitute the beliefs. A one-dimensional belief measure can only capture about a fifth of all the variation in the forecast differences. I conclude that beliefs are multidimensional, where each dimension has a distinct economic interpretation: (1) a long-term supply-side dimension, (2) a short-term supply-side dimension, and (3) a demand-side dimension. This multidimensionality has important implications for the dynamics and macroeconomic impact of beliefs. The different dimensions are orthogonal in the long-run, but their dynamics are interrelated over shorter periods. They have distinct impacts on the macroeconomic system. Overall, the three belief dimensions that I study play an important role in macroeconomic fluctuations, with the largest role played in the dynamics of the funds rate, followed by the contribution to output gap fluctuations. I do not find evidence for a role of beliefs in the evolution of inflation.

Regarding the characteristics of the beliefs, I can again make a distinction between the different dimensions. I find that learning plays a vital role in forming beliefs, but it affects each dimension differently. Furthermore, I find substantive evidence that the forecasters' beliefs exhibit irrationality. This evidence depends on two conditions. First, that my specification of the empirical distribution is reasonable. Future research could investigate whether alternative specifications lead to different results.[34] Second, it is conditional on my 50-year sample being sufficiently long. Strictly speaking, the rationality principle of Rational Beliefs is only guaranteed to hold in the limit $t \to \infty$. It could be the case that forecasters would pass the rationality tests over a longer time period. However, my results for a period spanning decades, with multiple generations of forecasters and diverse economic conditions, suggests that they would not. For practical purposes at least, like understanding and forecasting economic systems in the short-run (shorter than 50 years), one cannot rely on beliefs averaging out over time, and it is important to better understand their drivers.

My macroeconomic impact analysis indicates that the interest rate is most responsive to beliefs.[35] This result deserves more attention in future empirical research, zooming in specifically on the role of beliefs in monetary policy. Because I use revised data to estimate the macroeconomic VAR, I do not take into account the real-time information set of policymakers. By using real-time data to estimate a monetary policy

[34] Including non-linearities in the empirical transition functions might be one avenue to explore, where a trade-off will have to be made between model complexity and parameter stability.

[35] The role of beliefs in monetary theory is also an important theme in the theoretical Rational Beliefs literature (e.g., Kurz et al., 2005, 2013, 2018).

reaction function that includes the estimated belief factors, one could more clearly disentangle the effect of the different belief factors.

Another interesting avenue for future research would be to try and explain specific episodes where large shocks to the beliefs occur. Such an exercise could shed more light on the drivers of beliefs and on their macroeconomic impact.

Chapter 5

Conclusion

In this dissertation, I have investigated the intertwined roles of uncertainty and beliefs in economics. In chapter 2, I have modeled FOMC members' expectations using Bayesian learning. To analyze the impact of the uncertainty surrounding their macroeconomic forecasts, I have used real-time data about the economy and interest rate decisions at FOMC meetings. My analysis indicates that this uncertainty plays a role in monetary policy that is separate from the roles played by official (Greenbook) projections and uncertainty regarding the stock market. Specifically, it suggests that FOMC members set lower interest rates in a high-uncertainty regime compared to a low-uncertainty one.

In the third chapter, I have investigated the role of uncertainty in the belief-formation process. I have shown how different risk preferences and varying degrees of uncertainty surrounding the performance of forecasting rules can lead to heterogeneous expectations. Using empirical estimates for the distribution of risk aversion in a simple asset pricing model, I have shown that this belief-formation mechanism can lead to chaotic dynamics, including asset price bubbles.

In chapter 4, I have investigated the multidimensionality of sentiment in the context of Rational Beliefs theory. In this theory, uncertainty about the economy's true laws of motion leaves room for temporary waves of optimism and pessimism. My results suggest that sentiment is indeed multidimensional, with each dimension having a different macroeconomic impact.

I have made several suggestions for further research in the conclusions of each chapter. In what follows, I describe additional avenues for future research to build on this dissertation's results and address some of its limitations. I focus on potential syntheses between the chapters.

In chapter 2, I have used "a relatively small set of key relationships" (Greenspan, 2004, p. 37) to capture policymakers' beliefs. It would be insightful to use a more data-rich environment to derive an uncertainty measure. The Bayesian time-varying parameter VAR of Koop and Korobilis (2013) that I use in this chapter is especially well-suited for such an environment. One could use the comprehensive real-time dataset that I have compiled for the fourth chapter. The analysis in chapter 4 offers another suggestion for future research. It finds a prominent role for sentiment in interest rate dynamics, so it makes sense to include it in the policy reaction function. One might even use the dispersion along the various dimensions of sentiment as measures of uncertainty, where each one can have a distinct influence on monetary policy.

Rational Beliefs theory offers a derivation of sentiment's law of motion based on Bayesian learning and a combination of quantitative and qualitative information (Kurz et al., 2013). However, chapter 4 has shown that beliefs might not always satisfy the rationality principles of Rational Beliefs theory (at least on time scales shorter than 50 years). One might try to explain the deviations using a heuristic switching model like the one I have studied in the third chapter. Here, agents would use heuristics not to forecast economic variables themselves but their deviations from the empirical forecast. Such an analysis might provide further insight into the drivers of the different dimensions of beliefs.

Appendix A

Appendix to chapter 2

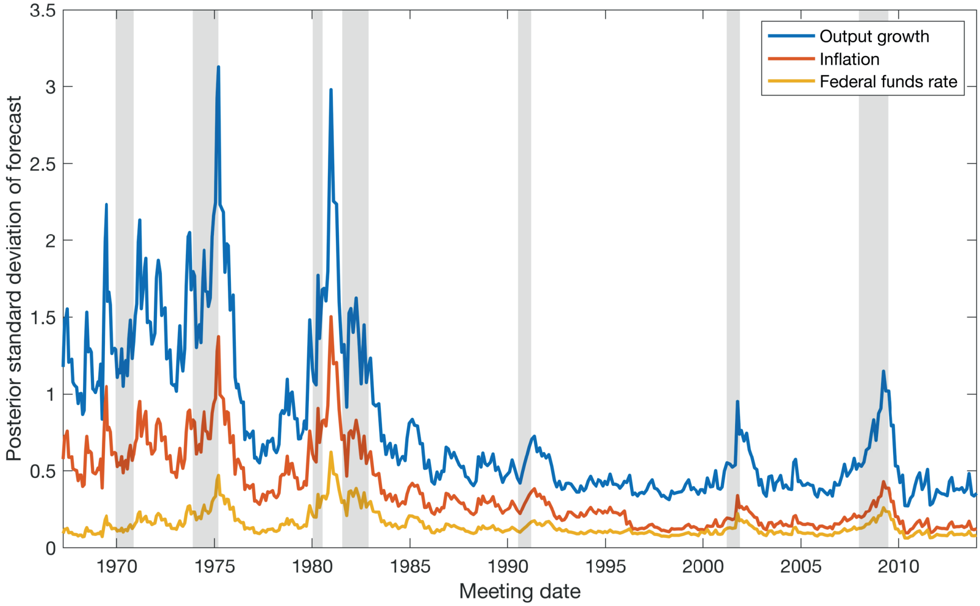

Figure A.1: Uncertainty related to one-quarter-ahead forecasts of output growth (U_{ym}), inflation ($U_{\pi m}$), and the federal funds rate (U_{im}). The measures are computed as the square root of the corresponding diagonal element in the covariance matrix of the posterior density of forecasts ($\widetilde{R}_{m,T_m+1}$). The grey bars indicate NBER recessions.

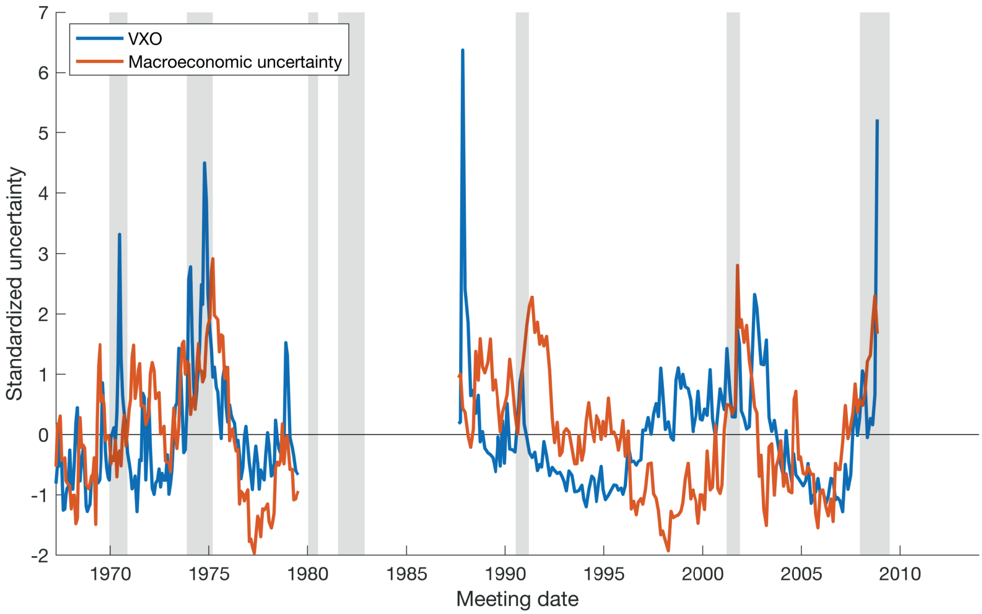

Figure A.2: Financial uncertainty proxy (VXO index) and entropy measure of macroeconomic uncertainty on Martins-Burns-Miller (1969 – 1979) and Greenspan-Bernanke (1987 – 2008) samples. Both series are standardized to have zero mean and unit standard deviation on the respective subsamples. The grey bars indicate NBER recessions.

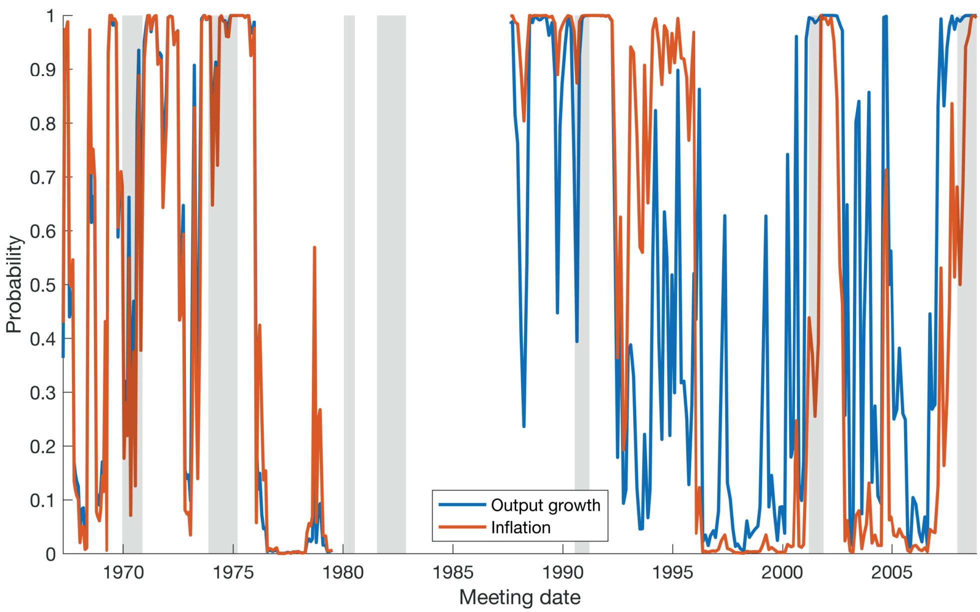

Figure A.3: Probability of being in a high-uncertainty regime, as defined in terms of output growth uncertainty (blue) and inflation uncertainty (red). The probabilities are calculated for the Martins-Burns-Miller (1969 – 1979) and Greenspan-Bernanke (1987 – 2008) samples separately, that is, using the subsample values for the median and standard deviation in (2.9). The grey bars indicate NBER recessions.

Appendix B

Appendix to chapter 3

B.1 Does a representative agent exist?

In this appendix, we show that while a representative agent does exist for the asset pricing model by Brock and Hommes (1998), no such aggregation is possible for our model.

First, a rational agent could not represent a heterogeneous economy, since an asset market with a rational agent would never deviate from the fundamental price. Second, an agent who would use the same heuristic forecasting method as the agents in the heterogeneous economy would only choose one rule at a time. She could therefore not represent the heterogeneous economy in which both heuristics can be used at the same time. We conclude that a representative agent would have to diversify across the heuristics instead of switching between them.

We start by deriving the pricing equation in the representative agent market, as well as the weights that she assigns to both rules in terms of the mean and variance of their performance. The mean-variance utility (3.1) does not have a maximum if $\gamma \leq 0$, so we assume that the representative agent has risk aversion $\gamma_{R,t} \in (0, \infty)$ for all t. In this case, aggregate demand is equal to individual demand, given by the representative agent version of equation (3.4):

$$\frac{E_{R,t}\left[R^e_{t+1}\right]}{\gamma_{R,t}\, V_{R,t}\left[R^e_{t+1}\right]}, \tag{B.1}$$

where $E_{R,t}$ and $V_{R,t}$ denote the representative agent's beliefs about expectation and variance, conditional on information set $\mathbb{I}_t$.

Like the heterogeneous agents, the representative agent is assumed to be familiar with the fundamental price and dividend processes: $E_{R,t}\left[p^*_{t+1} + d_{t+1}\right] = E_t\left[p^*_{t+1} + d_{t+1}\right]$. However, instead of choosing only one heuristic, the agent combines the fundamentalist and momentum rules in a weighted average to predict the price deviation from the fundamental:

$$E_{R,t}\left[x_{t+1}\right] = w_{F,t}\, h_F\left(x_{t-1}\right) + w_{M,t}\, h_M\left(x_{t-1}, x_{t-2}\right), \tag{B.2}$$

where $w_{F,t}$ and $w_{M,t}$ are the period t weights assigned to the fundamentalist and momentum rules, respectively. Note that we use the same two heuristics as in our model. However, the derivations and conclusions in this appendix hold for any two heuristics, including those explored by Brock and Hommes (1998).

The price deviations from the fundamental are given by

$$x_t = \frac{1}{R^f} \left\{ w_{F,t} \, f \, x_{t-1} + (1 - w_{F,t}) \left[x_{t-1} + m \left(x_{t-1} - x_{t-2} \right) \right] \right\}. \tag{B.3}$$

The weights for period t are determined by maximizing the following mean-variance performance measure with respect to $w_{F,t}$ and $w_{M,t}$:

$$\Phi_t = \begin{pmatrix} w_{F,t} & w_{M,t} \end{pmatrix} \begin{pmatrix} \langle u_F \rangle_{t-1} \\ \langle u_M \rangle_{t-1} \end{pmatrix} - \frac{\gamma_{R,t}}{2} \begin{pmatrix} w_{F,t} & w_{M,t} \end{pmatrix} \Sigma_{t-1} \begin{pmatrix} w_{F,t} \\ w_{M,t} \end{pmatrix}, \tag{B.4}$$

such that $w_{M,t+1} + w_{F,t+1} = 1$ and $0 \leq w_{F,t+1} \leq 1$. Here, Σ_t is the covariance matrix given by

$$\Sigma_t = \begin{pmatrix} \widetilde{\sigma}^2_{F,t} & \widetilde{\sigma}_{FM,t} \\ \widetilde{\sigma}_{FM,t} & \widetilde{\sigma}^2_{M,t} \end{pmatrix}, \tag{B.5}$$

with $\widetilde{\sigma}_{FM,t}$ a measure of the covariance between the utilities of the two heuristics in period t.[1] Note that period t risk aversion enters the performance measure, while mean and (co)variance measures are from period $t-1$. The timing is similar to that in the heterogeneous economy; period t mean and variance are not yet known when determining the weights for that period.

By maximizing the mean-variance performance measure (B.4), the representative agent chooses the weights given in the following proposition:

Proposition 6 (Representative agent weights)**.** *The weights chosen by the representative agent with relative risk aversion coefficient $\gamma_{R,t}$ are given by*

$$w_{F,t+1} = \min \left[\max \left(w^*_{F,t+1}, 0 \right), 1 \right],$$

where $w^*_{F,t+1}$ is the interior solution given by

$$w^*_{F,t+1} = \frac{\gamma^{-1}_{R,t+1} \left(\langle u_F \rangle_t - \langle u_M \rangle_t \right) - \widetilde{\sigma}_{FM,t} + \widetilde{\sigma}^2_{M,t}}{\widetilde{\sigma}^2_{F,t} - 2\, \widetilde{\sigma}_{FM,t} + \widetilde{\sigma}^2_{M,t}},$$

and

$$w_{M,t+1} = 1 - w_{F,t+1}.$$

Proof. See Appendix B.3. □

We now consider the representative agent for the model of Brock and Hommes (1998), who model the fractions of agents that use the heuristics with the multinomial

[1] We do not need to specify this measure here, but a definition in line with (3.14) would be $\widetilde{\sigma}_{FM,t} = \langle u_F \, u_M \rangle_t - \langle u_F \rangle_t \langle u_M \rangle_t$, with $\langle u_F \, u_M \rangle_t = \eta \, \langle u_F \, u_M \rangle_{t-1} + (1 - \eta) \, u_{F,t} \, u_{M,t}$.

logit (MNL).[2] The MNL can be derived from a stochastic utility model:[3] The utility $\widetilde{U}_{s,i,t}$ derived from heuristic $s \in \{F, M\}$ by agent i is expressed as

$$\widetilde{U}_{s,i,t} = U_{s,t} + \epsilon_{s,i,t}. \tag{B.6}$$

Here, $U_{s,t}$ is a performance measure for period t, which is observable at the aggregate level and can be equal to the weighted average of squared forecasting errors that we use in our model, or any other measure of expected performance. The stochastic term $\epsilon_{s,i,t}$ represents the stochastic differences between agents. It reflects all utility determinants that are not captured by $U_{s,t}$ as well as measurement errors and errors resulting from a potential misspecification of the functional form of $U_{s,t}$.

The representative agent is a mean-variance optimizer, who therefore wants to maximize the individual utilities, while penalizing their spread. The spread of individual utilities around the performance measure is measured by the variance of the stochastic term. The size of the penalty is determined by the representative agent's risk aversion. The MNL arises when the $\epsilon_{s,i,t}$ are assumed to be i.i.d across agents and heuristics according to the double exponential distribution. The variance of this distribution is given by $\pi^2/(6\beta^2)$, with $\beta > 0$ the intensity of choice parameter. To be able to compare the original model with our representative agent aggregation, we make the same assumption regarding the distribution of the stochastic utility term. To translate the results of Proposition 6 to this setting, we have

$$\langle u_F \rangle_t = U_{F,t}, \quad \langle u_M \rangle_t = U_{M,t}, \quad \widetilde{\sigma}_{FM} = 0, \quad \widetilde{\sigma}_F^2 = \widetilde{\sigma}_M^2 = \frac{\pi^2}{6\beta^2}.$$

With our fundamentalist and momentum heuristics, the Brock and Hommes (1998) deviations from the fundamental price are given by

$$x_t = \frac{1}{R^f} \left\{ n_{F,t}\, f\, x_{t-1} + (1 - n_{F,t}) \left[x_{t-1} + m\,(x_{t-1} - x_{t-2}) \right] \right\}, \tag{B.7}$$

with $n_{F,t}$ the fraction of agents using the fundamentalist rule, given by the MNL:

$$n_{F,t+1} = \frac{\exp(\beta\, U_{F,t})}{\exp(\beta\, U_{F,t}) + \exp(\beta\, U_{M,t})}. \tag{B.8}$$

Comparing with equation (B.3), we see that the representative agent model is equivalent to the Brock and Hommes model precisely when $w_{F,t+1} = n_{F,t+1}$ for arbitrary $U_{F,t}$ and $U_{M,t}$. In the following proposition, we show that this is possible when the representative agent has a specific, time-varying risk aversion.

Proposition 7 (Representative agent for multinomial logit switching)**.** *Suppose that the representative agent believes that the performance measures of the two forecasting heuristics are independent ($\widetilde{\sigma}_{FM,t} = 0$ for all t), and have equal, constant variance, given by*

$$\widetilde{\sigma}_F^2 = \widetilde{\sigma}_M^2 = \frac{\pi^2}{6\beta^2},$$

[2] Most heuristic switching studies use the MNL, following Brock and Hommes (1997). The MNL is well-documented in the discrete choice literature (e.g., Anderson et al., 1992).

[3] Most studies in the heuristic switching literature that refer to a specific derivation of the MNL (e.g., Hommes, 2013), refer to the interpretation of the stochastic utility model that we use here.

where β is the intensity of choice in the MNL. Furthermore suppose that she has a time-varying risk aversion given by

$$\gamma_{R,t+1} = \begin{cases} \dfrac{6\beta^2(U_{F,t} - U_{M,t})}{\pi^2} \dfrac{\exp(\beta\, U_{F,t}) + \exp(\beta\, U_{M,t})}{\exp(\beta\, U_{F,t}) - \exp(\beta\, U_{M,t})} & \text{if } U_{F,t} \neq U_{M,t}; \\ \dfrac{12\beta}{\pi^2} & \text{if } U_{F,t} = U_{M,t}. \end{cases}$$

Then the representative agent market is equivalent to a heterogeneous market with agents switching between the two heuristics, where the fractions of agents following the two rules are governed by the MNL.

Proof. See Appendix B.3. □

The agent's risk aversion $\gamma_{R,t}$ can be set to any value when $U_{F,t} = U_{M,t}$. However, we have made the particular choice presented in the proposition to ensure continuity of risk aversion as a function of $U_{F,t}$ and $U_{M,t}$. We also note that more general formulations of the proposition can be proven. The only requirements are that the variances of the two heuristics are equal and non-zero. An aggregation can also be achieved when variances are time-varying, or when the performance measures are correlated.[4]

We now turn to the representative agent for our market consisting of mean-variance optimizing agents with heterogeneous risk preferences. Comparing the representative agent price dynamics (B.3) with the heterogeneous system (3.20), we see that the representative agent economy is identical to it if and only if the weight assigned by the representative agent to the fundamentalist heuristic is equal to the risk-tolerance fraction represented by that heuristic in the heterogeneous economy in every period, that is, if and only if $w_{F,t} = \Theta_{F,t}$ for all $t \in \mathbb{N}_0$. Now suppose that for some period t, we have that $\langle u_F \rangle_t = \langle u_M \rangle_t$ and $\widetilde{\sigma}_{FM,t} < \widetilde{\sigma}^2_{F,t} < \widetilde{\sigma}^2_{M,t}$. It follows from equation (3.19) that $\Theta_{F,t+1} = 1$, while Proposition 6 tells us that $w_{F,t+1} < 1$, independent of the agents' risk aversion. It follows that $w_{F,t+1} \neq \Theta_{F,t+1}$, implying that our heterogeneous market of switching agents cannot be captured by a representative agent. The same argument holds if the representative agent does not take into account the covariance (i.e., $\widetilde{\sigma}_{FM,t} = 0$).[5]

B.2 Introducing noise traders

In this appendix, we derive the pricing equation for our model when it includes noise traders in addition to heuristic switching agents. The noise traders buy a random amount of the risky asset in each period. We denote by α_N the fraction of noise traders in the market, which leaves a fraction $1 - \alpha_N$ of heuristic switching agents. Let e_t be the stochastic demand of the noise traders in period t (in our simulations, this demand is drawn from a normal distribution). Equating supply and demand

[4] A proof is available upon request.

[5] One can prove that this result also holds when allowing the representative agent to divide her wealth over a finite set of (time-varying) risk aversions (proof available upon request).

now gives

$$0 = (1-\alpha_N)\int_0^\infty \frac{\widetilde{E}_{\gamma,t}\left[R^e_{t+1}\right]}{\gamma\,\widetilde{V}_{\gamma,t}\left[R^e_{t+1}\right]}\,g(\gamma)\,d\gamma + \alpha_N\,e_t$$
$$p_t = \frac{1}{R^f\,\Theta}\int_0^\infty \frac{g(\gamma)}{\gamma}\,\widetilde{E}_{\gamma,t}\left[p_{t+1}+d_{t+1}\right]\,d\gamma + \frac{1}{R^f\,\Theta}\,\frac{\alpha_N}{1-\alpha_N}\,e_t.$$

In the case of fundamentalists and chartists, this gives the following equation for the deviations from the fundamental price:

$$x_t = \frac{1}{R^f}\left\{\Theta_{F,t}\,f\,x_{t-1} + (1-\Theta_{F,t})\left[x_{t-1} + m\,(x_{t-1}-x_{t-2})\right]\right\} + \epsilon_t,$$

where we have defined the resulting price shock in terms of the noise traders' demand:

$$\epsilon_t = \frac{1}{R^f\,\Theta}\,\frac{\alpha_N}{1-\alpha_N}\,e_t.$$

B.3 Proofs

Proof of Proposition 1. We write $\left(x^*, \langle u_F\rangle^*, \langle u_M\rangle^*, \left\langle u_F^2\right\rangle^*, \left\langle u_M^2\right\rangle^*\right)$ to denote the steady state. It follows from the definition of the system (3.20) that

$$\langle u_F\rangle^* = \eta\,\langle u_F\rangle^* - (1-\eta)\left(x^* - f\,x^*\right)^2$$
$$\langle u_F\rangle^* = -\,(x^*)^2\left(1-f\right)^2.$$

Similarly, we have that $\left\langle u_F^2\right\rangle^* = (x^*)^4\left(1-f\right)^4$. Furthermore, it follows that

$$\langle u_M\rangle^* = \eta\,\langle u_M\rangle^*,$$

so that $\langle u_M\rangle^* = 0$ (since $0 < \eta < 1$). Similarly, $\left\langle u_M^2\right\rangle^* = 0$.

It now follows from (3.19) that

$$\Theta_F^* = \begin{cases} 0 & \text{if } x^* \neq 0; \\ \frac{1}{2} & \text{if } x^* = 0, \end{cases}$$

Note that if $\Theta_F^* = 0$, we have that

$$x^* = \frac{x^*}{R^f},$$

which means that $x^* = 0$. We conclude that $(0, 0, 0, 0, 0)$ is the only steady state. □

Proof of Proposition 2. We write the steady state as (v^*, x^*). Because $v_t = x_{t-1}$, we have that $v^* = x^*$. It then follows from the definition of A in (3.25) that

$$\frac{1}{R^f}\left[\Theta_F\,f\,x^* + (1-\Theta_F)\,x^*\right] = x^*$$
$$x^*\left[R^f - 1 + \Theta_F(1-f)\right] = 0$$
$$x^* = 0,$$

where in the last step we have used that $R^f - 1 + \Theta_F(1-f) > 0$, since $R^f > 1$, $0 \le \Theta_F \le 1$, and $f < 1$. We conclude that $(0,0)^T$ is the only steady state. $\square$

Proof of Proposition 3. If $\Theta_F = 1$, the map A simplifies to

$$A = \begin{pmatrix} 0 & 1 \\ 0 & \dfrac{f}{R^f} \end{pmatrix},$$

with eigenvalues 0 and f/R^f. Because $f < 1$ and $R^f > 1$, both eigenvalues have absolute value smaller than 1, which implies that the steady state is stable.

If $\Theta_F = 0$, we have

$$A = \begin{pmatrix} 0 & 1 \\ -\dfrac{m}{R^f} & \dfrac{1+m}{R^f} \end{pmatrix},$$

with characteristic equation

$$R^f \lambda^2 - (1+m)\lambda + m = 0,$$

and eigenvalues

$$\lambda_\pm = \frac{1 + m \pm \sqrt{(1+m)^2 - 4mR^f}}{2R^f}.$$

The eigenvalues are complex when

$$\begin{aligned} (1+m)^2 - 4mR^f &< 0 \\ m^2 + \left(2 - 4R^f\right) m + 1 &< 0. \end{aligned}$$

The zeros of the quadratic function on the left-hand side lie at

$$\begin{aligned} m_\pm &= \frac{4R^f - 2 \pm \sqrt{(2-4R^f)^2 - 4}}{2} \\ &= 2\left(R^f \pm \sqrt{R^f \left(R^f - 1\right)}\right) - 1. \end{aligned}$$

Note that $m_\pm \in \mathbb{R}$ because $R^f > 1$. It follows that $\lambda_\pm \in \mathbb{C}$ when $m_- < m < m_+$, and $\lambda_\pm \in \mathbb{R}$ if $m \le m_-$ or $m \ge m_+$.

Because m_- is strictly decreasing in R^f[6] and

$$\lim_{R^f \to \infty} R^f - \sqrt{R^f \left(R^f - 1\right)} = \frac{1}{2},^{7}$$

we have that

$$m_- > 2\left(\frac{1}{2}\right) - 1 = 0.$$

Because on the other hand

$$\sqrt{R^f(R^f - 1)} > \sqrt{(R^f - 1)^2} = R^f - 1,$$

it follows that

$$m_- < 2\left(R^f - R^f + 1\right) - 1 = 1.$$

We also have that

$$m_+ = R^f + 2\sqrt{R^f \left(R^f - 1\right)} + R^f - 1 > R^f.$$

We investigate $|\lambda_\pm|$ to derive the results about stability. The steady state is stable if the eigenvalues lie within the unit circle in the complex plane ($|\lambda_\pm| < 1$), neutral when they lie on the unit circle ($|\lambda_\pm| = 1$), and unstable if they lie outside the unit circle($|\lambda_\pm| > 1$). If $m_- < m < m_+$, we have $\lambda_\pm \in \mathbb{C}$, and

$$|\lambda_\pm|^2 = \frac{1}{4\left(R^f\right)^2}\left[(1+m)^2 + 4mR^f - (1+m)^2\right] = \frac{m}{R^f}.$$

---

[6] To prove this, consider m_- a function of $R^f \in (1, \infty)$, and take the derivative with respect to R^f to get

$$\frac{dm_-}{dR^f} = 2\left(1 - \frac{2R^f - 1}{2\sqrt{(2-4R^f)^2 - 4}}\right).$$

For $R^f = 2$, the derivative is equal to

$$\frac{dm_-}{dR^f} = 2\left(1 - \frac{3}{2\sqrt{2}}\right) < 0.$$

The derivative is equal to 0 if and only if

$$2R^f - 1 = 2\sqrt{(2-4R^f)^2 - 4}.$$

Squaring both sides of the equation gives $1 = 0$, a contradiction. It follows that the derivative of m_- with respect to R^f never vanishes, and is negative for $R^f = 2$. Because in addition this derivative is continuous on $(1, \infty)$, it is negative everywhere on this interval. We conclude that m_- is strictly decreasing in R^f.

[7] To see this, note that for $x \in (1, \infty)$:

$$\begin{aligned} x - \sqrt{x(x-1)} &= \frac{\left(x - \sqrt{x^2 - x}\right)\left(x + \sqrt{x^2 - x}\right)}{x + \sqrt{x^2 - x}} \\ &= \frac{x}{x + \sqrt{x^2 - x}} \\ &= \frac{1}{1 + \sqrt{1 - \frac{1}{x}}} \to \frac{1}{2} \quad (x \to \infty). \end{aligned}$$

Note that $m_- < R^f < m_+$, so that $\lambda_\pm \in \mathbb{C}$ and $|\lambda_\pm| = 1$ if $m = R^f$, implying a neutral steady state. It furthermore follows that the steady state is stable for $m_- < m < R^f$, and unstable for $R^f < m < m_+$.

If $\lambda_\pm \in \mathbb{R}$ (i.e., $m \leq m_-$ or $m \geq m_+$), we have that

$$(1+m)^2 - 4mR^f \geq 0,$$

and

$$\sqrt{(1+m)^2 - 4mR^f} \leq \sqrt{(1+m)^2} = 1+m,$$

(using $m \geq 0$) which implies that

$$1 + m - \sqrt{(1+m)^2 - 4mR^f} \geq 0.$$

In this case, it follows that

$$|\lambda_\pm| = \lambda_\pm = \frac{1 + m \pm \sqrt{(1+m)^2 - 4mR^f}}{2R^f},$$

so that

$$\lambda_- = |\lambda_-| \leq |\lambda_+| = \lambda_+. \tag{B.9}$$

We now prove that the steady state is stable for $0 \leq m \leq m_-$. Given inequality (B.9), it is enough to prove that $\lambda_+ < 1$. Let us write λ_+ as a function of m:

$$\lambda_+(m) = \frac{1 + m + \sqrt{(1+m)^2 - 4mR^f}}{2R^f}.$$

First note that

$$\lambda_+(m_-) = \frac{1 + m_- + 0}{2R^f} < \frac{1}{R^f} < 1.$$

Next, we consider $m \in [0, m_-)$. Because $R^f > 1$, it follows that

$$\lambda_+(0) = \frac{1}{R^f} < 1.$$

We continue by proving that $\lambda_+(m)$ is strictly decreasing on $[0, m_-)$, which implies that $\lambda_+(m) < 1$ for all $m \in [0, m_-)$, since it is smaller than 1 at $m = 0$. For $m \in [0, m_-)$, we have that

$$\frac{d\lambda_+}{dm}(m) = \frac{1}{2R^f}\left[1 + \frac{1 + m - 2R^f}{\sqrt{(1+m)^2 - 4mR^f}}\right],$$

implying that

$$\frac{d\lambda_+}{dm}(0) = \frac{1 - R^f}{R^f} < 0.$$

Because the derivative is continuous on $[0, m_-)$, it would have to vanish for some $m' \in (0, m_-)$ in order to change sign. The derivative vanishes at m' if and only if

$$1 + m' - 2R^f = -\sqrt{(1+m')^2 - 4m'R^f}.$$

Squaring both sides and rearranging leads to

$$\left(R^f\right)^2 = R^f.$$

This equality has no solutions for $R^f > 1$, and because the squares are not equal, no $m' \in (0, m_-)$ exists for which the derivative of λ_+ vanishes. We conclude that the derivative is negative on $[0, m_-)$, implying that λ_+ is strictly decreasing on this interval. As argued above, this means that $\lambda_+(m) < 1$ for all $m \in [0, m_-)$. We already showed that $\lambda_+(m) < 1$ for $m = m_-$, so we can conclude that the steady state is stable for $m \leq m_-$.

To prove that the steady state is unstable for $m \geq m_+$, it suffices to show that $\lambda_- > 1$, because of (B.9). Let us now write λ_- as a function of m:

$$\lambda_-(m) = \frac{1 + m - \sqrt{(1+m)^2 - 4mR^f}}{2R^f}.$$

We have that

$$\lambda_-(m_+) = \frac{1 + m_+ - 0}{2R^f} = \frac{R^f + \sqrt{R^f(R^f - 1)}}{R^f} > 1.$$

For $m \in (m_+, \infty)$ the derivative is given by

$$\frac{d\lambda_-}{dm}(m) = \frac{1}{2R^f}\left[1 - \frac{1 + m - 2R^f}{\sqrt{(1+m)^2 - 4mR^f}}\right].$$

The sign of the derivative is determined by the function

$$\zeta : (m_+, \infty) \longrightarrow \mathbb{R}, \quad m \mapsto 1 + m - 2R^f - \sqrt{(1+m)^2 - 4mR^f}.$$

If ζ is positive, the derivative is negative, if $\zeta = 0$, the derivative vanishes, and if it is negative, the derivative has positive sign. We can extend the domain of ζ to include m_+, and find that

$$\zeta(m_+) = \sqrt{R^f(R^f - 1)} > 0.$$

Since ζ is continuous and nowhere vanishing,[8] it is positive on the whole interval (m_+, ∞), which implies that the derivative of λ_- with respect to m is negative on this interval. We conclude hat λ_- is strictly decreasing as a function of $m \in (m_+, \infty)$. We

[8]We did a similar derivation above: $\zeta = 0$ if and only if

$$1 + m - 2R^f = -\sqrt{(1+m)^2 - 4mR^f}.$$

Squaring both sides and rearranging gives us

$$\left(R^f\right)^2 = R^f,$$

which has no solutions for $R^f > 1$. We conclude that ζ cannot vanish when $R^f > 1$.

finalize the proof that $\lambda_- > 1$ for $m \in (m_+, \infty)$ by showing that $\lambda_-(m) \to 1$ as $m \to \infty$:

$$\begin{aligned}
\lim_{m\to\infty} \lambda_-(m) &= \lim_{m\to\infty} \frac{1}{2R^f}\left[1+m-\sqrt{(1+m)^2-4mR^f}\right] \\
&= \lim_{m\to\infty} \frac{1}{2R^f}\left[1+m+\sqrt{(1+m)^2-4mR^f}\right]^{-1} \\
&\qquad \times \left[1+m-\sqrt{(1+m)^2-4mR^f}\right] \\
&\qquad \times \left[1+m+\sqrt{(1+m)^2-4mR^f}\right] \\
&= \lim_{m\to\infty} \frac{1}{2R^f}\left[\frac{4mR^f}{1+m+\sqrt{(1+m)^2-4mR^f}}\right] \\
&= \lim_{m\to\infty} \frac{1}{2R^f}\left[\frac{4R^f}{\frac{1}{m}+1+\sqrt{\frac{1}{m^2}+\frac{2}{m}+1-\frac{4R^f}{m}}}\right] \\
&= 1.
\end{aligned}$$

We have shown that λ_- strictly decreases from a value above 1 at $m = m_+$ to 1 as $m \to \infty$. This implies that $\lambda_-(m) > 1$ for all $m \in [m_+, \infty)$. We conclude that the steady state is unstable for $m \geq m_+$.

Combining the stability results for complex eigenvalues with those for real eigenvalues, we can conclude that the steady state is stable for $0 < m < R^f$, and unstable for $m > R^f$. This concludes the proof. □

Proof of Proposition 4. The map A, as defined in (3.25) has characteristic equation

$$R^f \lambda^2 - \left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right] \lambda + (1-\Theta_F)\, m.$$

The eigenvalues are given by

$$\begin{aligned}
\lambda_\pm &= \frac{1}{2R^f}\Bigg\{\Theta_F\, f + (1-\Theta_F)\,(1+m) \\
&\qquad \pm \sqrt{\left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2 - 4R^f\,(1-\Theta_F)\,m}\Bigg\} \\
&= \frac{\Theta_F\, f + (1-\Theta_F)\,(1+m) \pm \sqrt{\xi\left(\Theta_F; R^f, f, m\right)}}{2R^f},
\end{aligned}$$

where we have defined $\xi : [0,1] \to \mathbb{R}$,

$$\begin{aligned}
\Theta_F \mapsto &\left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2 - 4R^f\,(1-\Theta_F)\,m \\
&= (f-1-m)^2\,\Theta_F^2 + \left[2(1+m)(f-1-m) + 4R^f m\right]\Theta_F \\
&\quad + (1+m)^2 - 4R^f m,
\end{aligned}$$

with parameters $R^f \in (1, \infty)$, $f \in [0, 1)$, and $m \in (0, \infty)$.

Complexity of the eigenvalues The eigenvalues of A are real if $\xi \geq 0$ and complex if $\xi < 0$. The zeros of ξ are given by

$$\Theta_{F,\pm} = \frac{(1+m-f)(1+m) - 2mR^f \pm 2\sqrt{m^2 \left(R^f\right)^2 - f\, m\, R^f (1+m-f)}}{(1+m-f)^2}.$$

The eigenvalues are complex when these zeros are real and $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$. They are real when the zeros of ξ are complex, or when they are real and $\Theta_F \notin (\Theta_{F,-}, \Theta_{F,+})$. The expression in the square root in the definition for $\Theta_{F,\pm}$ can be rewritten as

$$mR^f \left[f^2 - (1+m)\, f + mR^f\right].$$

Since $m > 0$, the sign of this expression is determined by the function

$$\chi : [0,1) \longrightarrow \mathbb{R}, \quad f \mapsto f^2 - (1+m)\, f + mR^f.$$

The zeros $\Theta_{F,\pm}$ are complex when $\chi < 0$, and real when $\chi \geq 0$. Also note that if $\chi = 0$, we have that $\Theta_{F,-} = \Theta_{F,+}$, which means that $\xi \geq 0$ and that the eigenvalues are real. The zeros of χ are given by

$$f_\pm = \frac{1 + m \pm \sqrt{(1+m)^2 - 4mR^f}}{2}.$$

We conclude that the eigenvalues of A are real if $f_- \leq f \leq f_+$ or $\Theta_F \leq \Theta_{F,-}$ or $\Theta_F \geq \Theta_{F,+}$, and complex otherwise.

The modulus of the eigenvalues Note that

$$\sqrt{\xi(\Theta_F; R^f, f, m)} \leq \sqrt{\left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2} = \Theta_F\, f + (1-\Theta_F)\,(1+m)^9,$$

which implies that

$$\Theta_F\, f + (1-\Theta_F)\,(1+m) - \sqrt{\xi\left(\Theta_F; R^f, f, m\right)} \geq 0.$$

If the eigenvalues are real ($\xi \geq 0$), this means that

$$\lambda_- = |\lambda_-| \leq |\lambda_+| = \lambda_+. \tag{B.10}$$

If the eigenvalues are complex, we have that

$$\lambda_\pm = \frac{1}{2R^f}\Bigg\{\Theta_F\, f + (1-\Theta_F)\,(1+m) \\ \pm i\sqrt{4R^f\,(1-\Theta_F)\, m - \left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2}\Bigg\},$$

[9] Note that $\sqrt{\xi(\Theta_F; R^f, f, m)} \neq \sqrt{\left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2}$ if $\Theta_F < 1$ and $m > 0$.

so that

$$
\begin{aligned}
|\lambda_\pm|^2 &= \frac{1}{4\left(R^f\right)^2}\Big\{ \left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2 \\
&\qquad + 4R^f\,(1-\Theta_F)\,m - \left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2 \Big\} \\
&= \frac{(1-\Theta_F)\,m}{R^f}.
\end{aligned}
$$

It follows that $|\lambda_\pm| = 1$ at a threshold value given by

$$
\overline{\Theta}_F = 1 - \frac{R^f}{m}. \tag{B.11}
$$

If $\Theta_F < \overline{\Theta}_F$, we have that $|\lambda_\pm| > 1$, and if $\Theta_F > \overline{\Theta}_F$, it follows that $|\lambda_\pm| < 1$.

Now assume that $m > 1$. We are going to show that $\Theta_{F,-} < \overline{\Theta}_F < \Theta_{F,+}$. Consider $\Theta_{F,\pm}$ as functions of f:[10] $\Theta_{F,\pm} : [0,1) \to \mathbb{R}$,

$$
f \mapsto \frac{(1+m-f)(1+m) - 2m\,R^f \pm 2\sqrt{m^2\,\left(R^f\right)^2 - f\,m\,R^f(1+m-f)}}{(1+m-f)^2}.
$$

It follows that

$$
\Theta_{F,\pm}(0) = 1 + \frac{\pm 2mR^f - 2mR^f}{(1+m)^2}.
$$

We find that $\Theta_{F,+}(0) = 1 > \overline{\Theta}_F$, while $\Theta_{F,-}(0) < \overline{\Theta}_F$ if and only if

$$
\begin{aligned}
\frac{4mR^f}{(1+m)^2} &> \frac{R^f}{m} \\
4 &> \frac{(1+m)^2}{m^2} \\
4 &> 1 + \frac{2}{m} + \frac{1}{m^2} \\
m &> 1,
\end{aligned}
$$

where we have used the fact that $m > 0$. Hence it follows that $\Theta_{F,-}(0) < \overline{\Theta}_F < \Theta_{F,+}(0)$. Since $\Theta_{F,\pm}$ are continuous functions of f and $\overline{\Theta}_F$ is independent of f, we can only have $\overline{\Theta}_F \notin (\Theta_{F,-}, \Theta_{F,+})$ if $\Theta_{F,-}(f) = \overline{\Theta}_F$ or $\Theta_{F,+}(f) = \overline{\Theta}_F$ for some $f \in [0,1)$. However, note that $\Theta_{F,\pm} = \overline{\Theta}_F$ if and only if

$$
\begin{aligned}
\frac{(1+m-f)(1+m) - 2m\,R^f \pm 2\sqrt{m^2\,\left(R^f\right)^2 - f\,m\,R^f(1+m-f)}}{(1+m-f)^2} &= 1 - \frac{R^f}{m} \\
\pm 2\sqrt{m^2\,\left(R^f\right)^2 - f\,m\,R^f(1+m-f)} &= R^f\left[2(f-1) + m - \frac{(1-f)^2}{m}\right] \\
&\quad - f(1+m)
\end{aligned}
$$

[10] Later in this proof, we show that $m > m_-$, which includes the case of $m > 1$, implies that $\Theta_{F,\pm} \in \mathbb{R}$ and $\Theta_{F,-} < \Theta_{F,+}$. We do this by noting that the $f_\pm$ are complex if $m_- < m < m_+$ and showing that $f_\pm > 1$ if $m \geq m_+$.

Squaring both sides and solving for R^f leads to

$$R^f = \frac{f\,m}{f-1-3m} \text{ or } R^f = \frac{f\,m}{f-1+m}.$$

These solutions contradict $R^f > 1$, because we have that

$$\frac{f\,m}{f-1-3m} < 0,$$

and since $m > 1$, it follows that

$$\begin{aligned} f &< 1 \\ (m-1)f &< m-1 \\ f\,m &< f-1+m \\ \frac{f\,m}{f-1+m} &< 1, \end{aligned}$$

so $\Theta_{F,\pm} \neq \overline{\Theta}_F$. We conclude that $\Theta_{F,-} < \overline{\Theta}_F < \Theta_{F,+}$ if $m > 1$.

Before we continue with the proof, we formulate a lemma about the the range of the eigenvalues when they are real. Its proof can be found at the end of this appendix.

Lemma 1. *The functions defined by*

$$\lambda_\pm : D \to \mathbb{R}, \quad \Theta_F \mapsto \frac{\Theta_F\, f + (1-\Theta_F)(1+m) \pm \sqrt{\xi\left(\Theta_F; R^f, f, m\right)}}{2R^f},$$

with $0 \leq f < 1$, $m > 0$, $R^f > 1$ *, and where the domain* D *is chosen to ensure that* $D \subset [0,1]$ *and that the function is real:*

$$D = \begin{cases} [0,1] & \text{if } f_- \leq f \leq f_+ \text{ or } \Theta_{F,\pm} \leq 0 \text{ or } \Theta_{F,\pm} \geq 1; \\ [0,\Theta_{F,-}] \cup [\Theta_{F,+},1] & \text{if } 0 \leq \Theta_{F,-} < \Theta_{F,+} \leq 1; \\ [0,\Theta_{F,-}] & \text{if } 0 \leq \Theta_{F,-} < 1 < \Theta_{F,+}; \\ [\Theta_{F,+},1] & \text{if } \Theta_{F,-} < 0 < \Theta_{F,+} \leq 1; \\ \emptyset & \text{if } \Theta_{F,-} < 0 \text{ and } \Theta_{F,+} > 1; \end{cases}$$

satisfy

$$\lambda_\pm \neq 1.$$

The case $0 < m \leq m_-$ Suppose that $0 < m \leq m_-$. Because $m \leq m_-$, it follows that $(1+m)^2 - 4mR^f \geq 0$ (see the proof of Proposition 3), and the zeros $f_\pm$ of χ are real. Also note that

$$\sqrt{(1+m)^2 - 4mR^f} < \sqrt{(1+m)^2} = 1+m,$$

so that $f_- > 0$. We also have that

$$f_- \leq \frac{1+m_-}{2} < 1,$$

because $m_- < 1$ (Proposition 3). This means that there exist $f \in (0,1)$ for which $f_- \leq f \leq f_+$. For these values of f, the eigenvalues of A are real. We conclude that the eigenvalues are real when $0 < m \leq m_-$ and $f_- \leq f \leq f_+$. If $0 < m \leq m_-$ and $f \notin [f_-, f_+]$, the zeros $\Theta_{F,\pm}$ of ξ are real and $\Theta_{F,-} < \Theta_{F,+}$. In this case, the eigenvalues are complex if $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$, and real otherwise.

Let us now consider the stability of the steady state when $0 < m \leq m_-$. In the following, we show that the steady state is stable if the eigenvalues are real. Given inequality (B.10), it suffices to show that $\lambda_+ < 1$. First assume that $f_- \leq f \leq f_+$ or $\Theta_{F,\pm} \leq 0$ or $\Theta_{F,\pm} \geq 1$, so that we can define λ_+ as a continuous function of $\Theta_F \in [0,1]$ as in Lemma 1. From Proposition 3, we know that $\lambda_+(1) < 1$. Because λ_+ is continuous and Lemma 1 gives us that $\lambda_+ \neq 1$, it follows that $\lambda_+ < 1$.

Now assume that $\Theta_{F,\pm} \in \mathbb{R}$. If $0 < \Theta_{F,-} < 1$ and $\Theta_F \leq \Theta_{F,-}$, we can define λ_+ as a function of $\Theta_F \in [0, \Theta_{F,-}]$. Proposition 3 implies that $\lambda_+(0) < 1$. Because λ_+ is continuous on $[0, \Theta_{F,-}]$ and Lemma 1 tells us that $\lambda_+ \neq 1$, it follows that $\lambda_+ < 1$. If $0 < \Theta_{F,+} < 1$ and $\Theta_F \geq \Theta_{F,+}$, we can define λ_+ as a continuous function on $[\Theta_{F,+}, 1]$. Combining this with the fact that $\lambda_+(1) < 1$ and Lemma 1, we conclude that $\lambda_+ < 1$.

In the case that $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$, the eigenvalues of A are complex. We have already shown that the modulus of the eigenvalues depends on a threshold fraction of fundamentalists $\overline{\Theta}_F$, as defined in equation (B.11). If $m \leq m_- < 1$, we have that $\overline{\Theta}_F < 0$, and $\Theta_F > \overline{\Theta}_F$. This means that $|\lambda_\pm| < 1$ and the steady state is stable.

The case $m_- < m < m_+$ We start by showing that $\Theta_{F,-} < 0 < \Theta_{F,+} \leq 1$. If $m_- < m < m_+$, the zeros $f_\pm$ of χ are complex, and $\Theta_{F,\pm} \in \mathbb{R}$. We know from Proposition 3 that if $m_- < m < m_+$, the eigenvalues of A are complex at $\Theta_F = 0$, which means that $\xi(0) < 0$. This implies for the zeros of ξ that $\Theta_{F,-} < 0$ and $\Theta_{F,+} > 0$. From the definition of ξ, we know that

$$\xi(1) = f^2 \geq 0,$$

where $\xi(1) = 0$ if and only if $f = 0$. This means that either $\Theta_{F,+} = 1$ (when $f = 0$), or $\Theta_{F,+} \in (0,1)$ (when $f > 0$). We conclude that $\Theta_{F,-} < 0 < \Theta_{F,+} \leq 1$, where $\Theta_{F,+} = 1$ if and only if $f = 0$.

Now $\Theta_F < \Theta_{F,+}$ implies that $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$, since $\Theta_{F,-} < 0 \leq \Theta_F$. This means that the eigenvalues of A are complex, and the stability of the steady state depends on the threshold value $\overline{\Theta}_F$ like before.

Lastly, we show that the steady state is stable if $\Theta_F \geq \Theta_{F,+}$. In this case, it follows that the eigenvalues are real, and inequality (B.10) implies that it suffices to show that $\lambda_+ < 1$. Defining λ_+ as a function of Θ_F on $[\Theta_{F,+}, 1]$, Lemma 1 implies that it suffices to show that $\lambda_+(\Theta_F) < 1$ for some $\Theta_F \in [\Theta_{F,+}, 1]$. We already know from Proposition 3 that $\lambda_+(1) < 1$, so we can conclude that the steady state is stable if $\Theta_F \geq \Theta_{F,+}$.

The case $m \geq m_+$ If $m \geq m_+$, the zeros $f_\pm$ of χ are real, but they satisfy $f_\pm > 1$, which means that $\Theta_{F,\pm} \in \mathbb{R}$ and $\Theta_{F,-} < \Theta_{F,+}$ (because $f < 1$). To see that $f_\pm > 1$, note that χ has its minimum at f^*, which satisfies

$$f^* = \frac{1+m}{2} \geq \frac{1+m_+}{2} > \frac{1+R^f}{2} > 1.$$

If $m = m_+$, it follows that $f_\pm = f^* > 1$. If $m > m_+$, we have that $f_\pm \in \mathbb{R}$ and $\chi(f^*) < 0$. On the other hand, we know that

$$\chi(1) = m\left(R^f - 1\right) > 0.$$

Because χ is continuous, it follows that $f_- \in (1, f^*)$, and

$$1 < f_- < f_+.$$

We continue by showing that $0 \leq \Theta_{F,-} < \Theta_{F,+} \leq 1$. Note that ξ takes its minimal value at

$$\Theta_F^* = \frac{(1+m-f)(1+m) - 2mR^f}{(1+m-f)^2}.$$

We have that

$$m \geq m_+ = 2\left(R^f + \sqrt{R^f\,\left(R^f - 1\right)}\right) - 1 > 2R^f - 1,$$

so that

$$\begin{aligned} m + 1 - 2R^f &> 0 \\ m^2 + m(1 - 2R^f) &> 0 \\ m(1+m) - 2mR^f &> 0. \end{aligned}$$

This implies that

$$(1+m-f)(1+m) - 2mR^f > m(1+m) - 2mR^f > 0,$$

and $\Theta_F^* > 0$.

Moreover, it follows that $\Theta_F^* < 1$ if and only if

$$\begin{aligned} (1+m-f)(1+m) - 2mR^f &< (1+m-f)^2 \\ f^2 - (1+m)f + 2mR^f &> 0. \end{aligned} \tag{B.12}$$

Because $m \geq m_+ > R^f > 1$, we have that

$$f^2 - (1+m)f + 2mR^f > -(1+m) + 2mR^f = m\left(2R^f - 1\right) - 1 > 0,$$

and inequality (B.12) holds, so that $\Theta_F^* < 1$. Because in addition $f < f_-$, as we have shown earlier, we conclude that

$$\xi(\Theta_F^*) < 0, \quad 0 < \Theta_F^* < 1. \tag{B.13}$$

We furthermore have that

$$\xi(0) = (1+m)^2 - 4mR^f \geq 0,$$

because $m \geq m_+$, with $\xi(0) = 0$ if and only if $m = m_+$. In combination with (B.13), this means that $\Theta_{F,-} \geq 0$ and $\Theta_{F,-} = 0$ if and only if $m = m_+$. We also know that

$\xi(1) \geq 0$, with $\xi(1) = 0$ if and only if $f = 0$ (as mentioned earlier). Combining this with (B.13) implies that $\Theta_{F,+} \leq 1$ and that $\Theta_{F,+} = 1$ if and only if $f = 0$.

We now consider the stability of the steady state of A. If $\Theta_F \leq \Theta_{F,-}$, the eigenvalues of A are real. In the following, we show that the steady state is unstable in this case, meaning $|\lambda_\pm| > 1$. Inequality (B.10) implies that showing $\lambda_- > 1$ is sufficient. We can define λ_- as a real-valued function of $\Theta_F \in [0, \Theta_{F,-}]$. Proposition 3 implies that $\lambda_-(0) > 1$, since $m \geq m_+ > R^f$. Because λ_- is continuous and $\lambda_- \neq 1$ (Lemma 1), it follows that $\lambda_- > 1$ for all $\Theta_F \in [0, \Theta_{F,-}]$. We conclude that the steady state is unstable for $\Theta_F \leq \Theta_{F,-}$.

If $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$, the eigenvalues of A are complex, and the stability of the steady state is again determined by the threshold $\overline{\Theta}_F$. The proof that the steady state is stable if $\Theta_F \geq \Theta_{F,+}$ (implying real eigenvalues) carries over from the case of $m \in (m_-, m_+)$.

Taking all cases together, the steady state is unstable if $\Theta_F < \overline{\Theta}_F$, and stable if $\Theta_F > \overline{\Theta}_F$. To see this, first note that we have shown that the steady state is stable if $m \leq m_-$, while in this case indeed $\Theta_F > \overline{\Theta}_F$, because $\overline{\Theta}_F < 0$. Second, in the case of $m_- < m < m_+$, we have shown that the stability of the steady state depends on $\overline{\Theta}_F$ as described above if $\Theta_F < \Theta_{F,+}$. If $\Theta_F \geq \Theta_{F,+}$, the steady state is stable, while it indeed also holds that $\Theta_F > \overline{\Theta}_F$, since $\Theta_{F,+} > \overline{\Theta}_F$. Third, note that if $m \geq m_+$, we have that the stability of the steady state is determined by $\overline{\Theta}_F$ if $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$. If $\Theta_F < \Theta_{F,-}$ ($\Theta_F > \Theta_{F,-}$), the steady state is unstable (stable), which is consistent with the claim above, since $\Theta_{F,-} < \overline{\Theta}_F < \Theta_{F,+}$.

Furthermore, we have that the complexity of the eigenvalues is determined by Θ_F relative to the $\Theta_{F,\pm}$ as described above, except for the case in which the $\Theta_{F,\pm}$ are complex: If $m \leq m_-$ and $f_- \leq f \leq f_+$, the $\Theta_{F,\pm}$ are complex, and the eigenvalues of A are real. □

Proof of Proposition 6. The representative agent maximizes Φ_t given by equation (B.4). Omitting time subscripts for clarity and using that $w_M = 1 - w_F$, this reduces to

$$\begin{aligned}\Phi &= w_F\left(\langle u_F\rangle - \langle u_M\rangle\right) + \langle u_M\rangle \\ &\quad - \frac{\gamma_R}{2}\left[w_F^2\,\widetilde{\sigma}_F^2 + 2\,w_F\,(1-w_F)\,\widetilde{\sigma}_{FM} + (1-w_F)^2\,\widetilde{\sigma}_M^2\right].\end{aligned}$$

The weights are constrained to the $[0, 1]$ interval:

$$-w_F \leq 0; \tag{B.14a}$$

$$w_F - 1 \leq 0. \tag{B.14b}$$

The first and second derivatives of the objective function with respect to w_F are given by

$$\begin{aligned}\frac{\partial\Phi}{\partial w_F} &= \langle u_F\rangle - \langle u_M\rangle - \gamma_R\left[w_F\,\widetilde{\sigma}_F^2 + (1-2\,w_F)\,\widetilde{\sigma}_{FM} - (1-w_F)\,\widetilde{\sigma}_M^2\right]; \\ \frac{\partial^2\Phi}{\partial w_F^2} &= -\gamma_R\left(\widetilde{\sigma}_F^2 - 2\,\widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2\right).\end{aligned}$$

The parenthetical expression in the second derivative is just the variance of the difference between heuristics' performance ($\widetilde{\sigma}_F^2 - 2\,\widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2 = V[u_F - u_M]$), which is non-negative. This implies that the objective function is concave (since $\gamma > 0$). Because we also have that the constraints are convex and continuously differentiable with respect to w_F, the Karush-Kuhn-Tucker conditions are sufficient. Denoting the optimum by w_F^* and the multipliers corresponding to constraints (B.14a) and (B.14b) by $\mu_{1,2}$, these conditions are given by

$$\left.\frac{\partial \Phi}{\partial w_F}\right|_{w_F^*} = \mu_2 - \mu_1; \tag{B.15a}$$

$$-w_F^* \leq 0; \tag{B.15b}$$

$$w_F^* - 1 \leq 0; \tag{B.15c}$$

$$\mu_{1,2} \geq 0; \tag{B.15d}$$

$$-\mu_1 \, w_F^* = 0; \tag{B.15e}$$

$$\mu_2 \, (w_F^* - 1) = 0. \tag{B.15f}$$

We start with the interior solution: When none of the constraints are binding, we have that $\mu_1 = \mu_2 = 0$, and condition (B.15a) gives that

$$\gamma_R^{-1}(\langle u_F \rangle - \langle u_M \rangle) = w_F^* \, \widetilde{\sigma}_F^2 + (1 - 2\, w_F^*)\, \widetilde{\sigma}_{FM} - (1 - w_F^*)\, \widetilde{\sigma}_M^2$$

$$w_F^* = \frac{\gamma_R^{-1}\,(\langle u_F \rangle - \langle u_M \rangle) - \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2}{\widetilde{\sigma}_F^2 - 2\,\widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2}. \tag{B.16}$$

The two constraints cannot be binding at the same time. If (B.14a) is binding, we have $w_F^* = 0$, $\mu_1 > 0$, $\mu_2 = 0$, and condition (B.15a) implies that

$$\gamma_R^{-1}\,(\langle u_F \rangle - \langle u_M \rangle) - \widetilde{\sigma}_{FM} - \widetilde{\sigma}_M^2 < 0$$

Comparing with equation (B.16), we see that this corresponds to the interior solution being negative. Analogously, if constraint (B.14b) is binding, it follows that $w_F^* = 1$, $\mu_1 = 0$, $\mu_2 > 0$, and condition (B.15a) implies that

$$\gamma_R^{-1}\,(\langle u_F \rangle - \langle u_M \rangle) - \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2 > \widetilde{\sigma}_F^2 - 2\,\widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2,$$

which corresponds to the interior solution being strictly larger than 1. We conclude that the optimal weight put on the fundamentalist rule is

$$w_F^* = \min\left[\max\left(\frac{\gamma_R^{-1}\,(\langle u_F \rangle - \langle u_M \rangle) - \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2}{\widetilde{\sigma}_F^2 - 2\,\widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2}, 0\right), 1\right],$$

with the optimal weight for the momentum rule given by $w_M^* = 1 - w_F^*$. □

Proof of Proposition 7. For simplicity, we write

$$\sigma^2 = \widetilde{\sigma}_F^2 = \widetilde{\sigma}_M^2 = \frac{\pi^2}{6\beta^2}$$

and rewrite the MNL fractions as

$$n_{F,t+1} = \frac{\exp(\beta \Delta_t)}{\exp(\beta \Delta_t) + 1},$$

where we have defined $\Delta_t = U_{F,t} - U_{M,t}$. To prove the theorem, we have to show that the representative agent weights satisfy $w_{F,t+1} = n_{F,t+1}$ for arbitrary Δ_t, which would imply that the pricing equations (B.3) and (B.7) for the representative agent market and the MNL switching market are the same.

Because we have that $\widetilde{\sigma}_{FM,t} = 0$, it follows from Proposition 6 that the interior solution to the representative agent maximization problem is given by

$$w^*_{F,t+1} = \frac{\Delta_t}{2\gamma_{R,t+1}\,\sigma^2} + \frac{1}{2}.$$

First note that when $\Delta_t = 0$, it follows that

$$w_{F,t+1} = w^*_{F,t+1} = \frac{1}{2} = n_{F,t+1}.$$

Now suppose that $\Delta_t \neq 0$. We can rewrite the representative agent's risk aversion as defined in the proposition in terms of Δ_t and σ^2:

$$\gamma_{R,t+1} = \frac{\Delta_t}{\sigma^2}\,\frac{\exp(\beta\Delta_t) + 1}{\exp(\beta\Delta_t) - 1},$$

which gives that

$$\begin{aligned} w^*_{F,t+1} &= \frac{1}{2\frac{\exp(\beta\Delta_t)+1}{\exp(\beta\Delta_t)-1}} + \frac{1}{2}. \\ &= \frac{\exp(\beta\Delta_t)}{\exp(\beta\Delta_t) + 1} \\ &= n_{F,t+1}. \end{aligned}$$

Since $0 < n_{F,t+1} < 1$, we have that $w_{F,t+1} = w^*_{F,t+1} = n_{F,t+1}$. We conclude that $w_{F,t+1} = n_{F,t+1}$ for arbitrary Δ_t. □

Proof of Lemma 1. We have that $\lambda_\pm = 1$ if and only if

$$\begin{aligned} &\Theta_F\, f + (1-\Theta_F)(1+m) \pm \sqrt{\xi\left(\Theta_F; R^f, f, m\right)} = 2R^f \\ &\pm \sqrt{\left[\Theta_F\, f + (1-\Theta_F)(1+m)\right]^2 - 4R^f(1-\Theta_F)\, m} = 2R^f \\ &\qquad\qquad - \Theta_F\, f - (1-\Theta_F)(1+m). \end{aligned} \tag{B.17}$$

Squaring both sides and rearranging gives

$$R^f = 1 + \Theta_F(f - 1 - m) \leq 1.$$

This is a contradiction, since $R^f > 1$. Since the squares are not equal, equality (B.17) cannot hold. We conclude that $\lambda_\pm \neq 1$. □

B.4 Cobweb application

B.4.1 The linear cobweb model with diverse risk preferences

In this appendix, we present the implications of our expectation-formation framework in a linear cobweb model. The version of the cobweb model we present here is inspired by Hommes (2013). We use the linear version to focus on the non-linearities arising from the expectation-formation process. The cobweb model represents a market for non-storable goods (e.g., corn), which are produced in the period before they are sold in the market. In each period, suppliers have to decide how much they are going to produce. This decision crucially depends on the price for which they can expect to sell their produced goods in the next period. This means that in period t they have to form expectations about the price in period $t + 1$.

Market mechanism

The supply function of an individual producer is linear and given by $s\,p_t^e$, with $s > 0$ and p_t^e the forecast for the period t price. This supply function is consistent with profit maximization under a quadratic cost function $c(q) = q^2/2s$ (Brock and Hommes, 1997).[11] Now assume that we have a continuum of producers on the unit interval $(0, 1)$. Denoting by $p_t^e(i)$ the price forecast of producer $i \in (0, 1)$, we find that aggregate supply is given by

$$\begin{aligned} S_t &= \int_0^1 s\,p_t^e(i)\,di \\ &= s\,\overline{p}_t^e, \end{aligned}$$

expressed in terms of the average of all producers' forecasts: $\overline{p}_t^e = \int_0^1 p_t^e(i)\,di$.[12]

Aggregate demand is given by $D(p_t) = a - d\,p_t$, where $a, d > 0$. By adding the market clearing condition, we get the system that governs the dynamics of the model:

$$D(p_t) = a - d\,p_t, \quad a, d > 0, \tag{B.18a}$$

$$S(\overline{p}_t^e) = s\,\overline{p}_t^e, \quad s > 0, \tag{B.18b}$$

$$D(p_t) = S(\overline{p}_t^e). \tag{B.18c}$$

Beliefs

We investigate the dynamics when producers can choose between two strategies: rational and naive. The rational strategy requires complete information about the model and all the other agents' expectations, and perfectly predicts the price in the next period. The naive strategy predicts that the price in the next period will be the same as the last observed price. Using the notation $p_{R,t+1}^e$ for the rational and $p_{N,t+1}^e$

[11] Expected profit in period t is given by $\Pi_t^e(q) = q\,p_t^e - q^2/2s$. We have FOC $p_t^e - q^*/s = 0$, which gives $q^* = s\,p_t^e$. Because $\Pi_t^e(q)$ is strictly concave, q^* is a maximum.

[12] This identity readily extends to the discrete case, with the integral replaced by a summation.

for the naive forecast, we have that

$$p^e_{R,t+1} = p_{t+1}; \tag{B.19a}$$

$$p^e_{N,t+1} = p_t. \tag{B.19b}$$

The choice of strategy is governed by the same framework as in the asset pricing model: Suppliers use the mean-variance performance measure defined in equation (3.15) to find the strategy that best balances performance and risk. Different from the asset pricing model, a strategy's performance measure does not only depend on the squared forecasting errors, but also on the cost of using the strategy. We assume that use of the naive rule is free, but that the rational rule has a cost C_R. The interpretation is that resources are needed to gather information about the structure of the model and the expectations of the other suppliers. The specification of these strategies, including the information costs of the rational rule is taken from Brock and Hommes (1997).

Note that the rational rule always has zero forecasting errors, meaning that its utility is always equal to minus its cost: $u_{R,t} = u_R = -C_R$. The utility for the naive rule is given by

$$u_{N,t} = -\left(p_t - p_{t-1}\right)^2.$$

The average utility, average squared utility and variance are calculated in the same way as in the asset pricing model (equations 3.14a, 3.14b, and 3.14c). Note that the rational rule has zero variance, since its utility is constant.

The following result establishes a link between squared forecasting errors and profits, providing support for our choice of utility:

Proposition 8. *In the the linear cobweb model with cost function $c(q) = q^2/2s$, the squared forecasting error is equal to the difference between maximum profit and realized profit, up to a constant. They are equal for $s = 2$.*

Proof. Maximum profit is given by

$$\begin{aligned}\Pi^*_t &= s\, p_t^2 - \frac{(s\, p_t)^2}{2s} \\ &= \frac{s}{2}\, p_t^2,\end{aligned}$$

while realized profit is given by

$$\begin{aligned}\Pi_t &= s\, p^e_t\, p_t - \frac{(s\, p^e_t)^2}{2s} \\ &= \frac{s}{2}\left[2\, p^e_t\, p_t - (p^e_t)^2\right].\end{aligned}$$

It follows that

$$\begin{aligned}\Pi^*_t - \Pi_t &= \frac{s}{2}\left[p_t^2 - 2\, p^e_t\, p_t + (p^e_t)^2\right] \\ &= \frac{s}{2}\left(p_t - p^e_t\right)^2,\end{aligned}$$

which is $s/2$, a constant, times the squared forecasting error. The constant equals 1 for $s = 2$. □

We are now in the position to determine the period t fractions of producers that use the forecasting strategies. Since they have to form one-period ahead forecasts, these fractions are determined in the previous period $(t-1)$ and based on information from that period. Similar to the asset pricing model, we can determine the risk aversion coefficient for which an agent would be indifferent between the two rules in period $t-1$, again denoted by $\overline{\gamma}_{t-1}$. Because the rational rule is riskless, suppliers with risk aversion below the threshold $\overline{\gamma}_{t-1}$ will choose the naive rule, and those with risk aversion above the threshold will choose the rational rule. If $\widetilde{\sigma}_{N,t-1} \neq 0$, it follows that

$$\overline{\gamma}_{t-1} = 2\,\frac{\langle u_N \rangle_{t-1} + C_R}{\widetilde{\sigma}^2_{N,t-1}}.$$

The fraction of producers using the naive rule in period t is then given by

$$n_{N,t} = G\left(\overline{\gamma}_{t-1}\right),$$

with G the empirical cumulative distribution function of risk aversion. The fraction of rational producers is given by

$$n_{R,t} = 1 - G\left(\overline{\gamma}_{t-1}\right).$$

If $\widetilde{\sigma}_{N,t-1} = 0$, the fractions are only determined by $\langle u_N \rangle_{t-1}$: If $\langle u_N \rangle_{t-1} < -C_R$, we have $n_{N,t} = 0$, if $\langle u_N \rangle_{t-1} > -C_R$, we have $n_{N,t} = 1$, and if $\langle u_N \rangle_{t-1} = -C_R$, we have $n_{N,t} = n_{R,t} = 1/2$.

The average forecast in period t is given by

$$\overline{p}^e_t = n_{R,t}\, p_t + n_{N,t}\, p_{t-1}. \tag{B.20}$$

Combining equations (B.18a) – (B.20), we find that the clearing price is given by

$$p_t = \frac{a - n_{N,t}\, s\, p_{t-1}}{n_{R,t}\, s + d}.$$

The steady state price is given by

$$p^* = \frac{a}{s+d}, \tag{B.21}$$

where we have used that $n_{R,t} + n_{N,t} = 1$. Note that this steady state price is constant, contrary to the fundamental price that we derived in the asset pricing model.

For $a = 0$, the steady state is 0 and the model reduces to the deviations-from-steady-state model of Muth (1961). In some of our simulations, we include an aggregate supply shock ϵ_t, due to weather conditions for example. In that case, aggregate supply (B.18b) is replaced by

$$S(\overline{p}^e_t) = s\,\overline{p}^e_t + \epsilon_t,$$

so that the market clearing price becomes

$$p_t = \frac{a - n_{N,t}\, s\, p_{t-1} - \epsilon_t}{n_{R,t}\, s + d}.$$

The steady state is the same if the supply shock has zero mean.

We denote the phase space by $Z = \mathbb{R} \times \mathbb{R}_{\leq 0} \times \mathbb{R}_{\geq 0}$. The evolution of the resulting 3-dimensional discrete dynamical system is governed by the one-step map

$$\psi : Z \to Z, \quad \begin{pmatrix} p_t \\ \langle u_N \rangle_t \\ \langle u_N^2 \rangle_t \end{pmatrix} \mapsto \begin{pmatrix} p_{t+1} \\ \eta \langle u_N \rangle_t - (1-\eta)\left(p_{t+1} - p_t\right)^2 \\ \eta \langle u_N^2 \rangle_t + (1-\eta)\left(p_{t+1} - p_t\right)^4 \end{pmatrix}. \tag{B.22}$$

If $\widetilde{\sigma}_{N,t} \neq 0$, next period's price p_{t+1} is given by

$$p_{t+1} = \frac{a - G\left(2\,\dfrac{\langle u_N \rangle_t + C_R}{\langle u_N^2 \rangle_t - \langle u_N \rangle_t^2}\right) s\, p_t - \epsilon_t}{\left[1 - G\left(2\,\dfrac{\langle u_N \rangle_t + C_R}{\langle u_N^2 \rangle_t - \langle u_N \rangle_t^2}\right)\right] s + d}.$$

If $\widetilde{\sigma}_{N,t} = 0$, it is given by

$$p_{t+1} = \begin{cases} \dfrac{a - \epsilon_t}{s + d} & \text{if } \langle u_N \rangle_t < -C_R; \\[2ex] \dfrac{a - \dfrac{s\,p_t}{2} - \epsilon_t}{\dfrac{s}{2} + d} & \text{if } \langle u_N \rangle_t = -C_R; \\[2ex] \dfrac{a - s\,p_t - \epsilon_t}{d} & \text{if } \langle u_N \rangle_t > -C_R. \end{cases}$$

B.4.2 Dynamics

Dynamics with constant fractions

Before introducing endogenous switching between the two forecasting strategies into the cobweb model, we study its properties when the fractions of rational and naive agents are constant. This will later help us understand the dynamics of the full model.

Proposition 9. *When the fractions of rational and naive producers are constant, the clearing price in the linear cobweb model evolves according to*

$$p_t = \left(-\frac{n_N\, s}{n_R\, s + d}\right)^t (p_0 - p^*) + p^*,$$

with p_0 the initial price and p^* the steady state price given by equation (B.21)*. Stability of the steady state depends on the fraction of rational producers relative to a threshold given by*

$$\overline{n}_R = \frac{s - d}{2s}.$$

When $n_R > \overline{n}_R$, the steady state is globally asymptotically stable, meaning that it is stable and everywhere attracting. For $n_R = \overline{n}_R$, the steady state is neutral and the system follows a two-cycle. When $n_R < \overline{n}_R$, the steady state is unstable and all solutions are unbounded, except for the steady state.

Proof. The clearing price in period t is given by

$$p_t = \frac{a}{n_R\, s + d} - \frac{n_N\, s}{n_R\, s + d}\, p_{t-1}.$$

The homogeneous problem

$$p_t = -\frac{n_N\, s}{n_R\, s + d}\, p_{t-1}$$

has general solution

$$p_t = \left(-\frac{n_N\, s}{n_R\, s + d}\right)^t c,$$

with c some constant. Knowing that the steady state p^* is a particular solution for the full, inhomogeneous problem, we can add it to the general solution for the homogeneous problem and solve for c to get

$$p_t = \left(-\frac{n_N\, s}{n_R\, s + d}\right)^t (p_0 - p^*) + p^*.$$

The stability of the steady state changes at the threshold value $\overline{n}_R$ given by

$$\begin{aligned} \frac{\overline{n}_N\, s}{\overline{n}_R\, s + d} &= 1 \\ (1 - \overline{n}_R)s &= \overline{n}_R\, s + d \\ \overline{n}_R &= \frac{s - d}{2s}. \end{aligned}$$

For $n_R > \overline{n}_R$, the steady state is stable and globally attracting. When $n_R = \overline{n}_R$, the clearing price jumps back and forth between p_0 and $2p^* - p_0$ in a two-cycle, indicating metastability. For $n_R < \overline{n}_R$, lastly, $|p_t - p^*|$ increases indefinitely when $p_0 \neq p^*$. □

In the special case that all producers use the rational strategy at all times ($n_{R,t} = n_R = 1$), the system is always at the steady state. Even when the initial price p_0 is far from the steady state price, the system immediately converges to the steady state.

When all producers are using the naive rule ($n_{N,t} = n_N = 1$), the model's dynamics depends on the relation between supply and demand parameters s and d. In the case that $s < d$, $\overline{n}_R < 0$ and $n_R > \overline{n}_R$: the system converges to the steady state, while price deviations alternate between positive and negative. The intuition is that after an initial price above the steady state, suppliers expect the same deviation in the next period, and produce more than steady state supply $s\, p^*$. This oversupply then leads to a price below steady state in the following period. Because $s < d$ however, the increase in supply is smaller than needed to reach the same absolute deviation from steady state, and the new price is closer to the steady state price. Next, the suppliers expect the same low price level, and decide to produce less than steady state supply. This again leads to an above-steady-state price that is yet closer to the steady state. This process of ever smaller deviations from steady-state supply continues until the steady state is reached.

When $s = d$, the clearing price jumps back and forth between p_0 and $2p^* - p_0$ in a two-cycle: The over or undersupply compared to steady state is precisely enough to reach the same price deviation in the next period, but on the other side of the steady state. Lastly, in the case $s > d$, the steady state is unstable, and deviations from the steady state become arbitrarily large.

Note that in general the steady state is globally asymptotically stable when $d > s$, since in that case $\overline{n}_R < 0$, while $n_R \geq 0$.

Dynamics with endogenous fractions

We now allow for endogenous switching between the two forecasting strategies. We set the cost of using the rational strategy to 1 (i.e., $C_R = 1$). We use numerical methods to study the systems' dynamics. We provide intuitive explanations in terms of the theoretical results presented in Proposition 9. We present the results on the price dynamics in terms of deviations from the steady state, denoted by x_t:

$$x_t = p_t - p^*.$$

In all our simulations, we set $d = 1$ and $s = 2$. This means that the squared forecasting error is equal to the profit deviations from the maximum (Proposition 8). Moreover, it implies that $s > d$: The steady state is unstable under the naive rule, and the threshold fraction of rational producers is given by $\overline{n}_R = 0.25$. The interaction between the stabilizing and destabilizing forces of the rational and naive suppliers will give rise to rich dynamics: stable in some parameter ranges and chaotic in others.

We start with the deterministic model, with zero supply shock (i.e., $\epsilon_t = 0$). When producers have substantial memory ($\eta = 0.9$), chaotic dynamics arise, as can be seen in the two top-left panels of Figure B.1. The system starts close to the steady state, with the fraction of rational producers above the threshold value $\overline{n}_R = 0.25$, so that the steady state is stable and attracting (Proposition 9). The fraction of rational agents quickly drops below the threshold, leading to an unstable steady state and increasing volatility in the clearing price. When the forecasting errors of the naive rule become too large, producers begin to switch back to the rational rule, starting with those who have the largest rate of risk aversion. As soon as the fraction of rational producers exceeds the threshold of 0.25, the system starts to stabilize, and the producers gradually switch back to the naive rule to save costs (starting with those with lowest γ).

A key characteristic of this chaotic solution is that not all producers switch to the rational rule during volatile periods: Both forecasting strategies are used at all times. We can identify two reasons for this behaviour. First, there are producers left with low enough risk aversion to stick with the naive rule. Second, there is significant persistence in the performance measures of the naive rule (because $\eta = 0.9$), meaning that the larger volatility only gradually affects the mean-variance utility of that rule. The result is that the system never reaches the steady state. When the fraction of rational producers drops below 0.25 after a volatile episode, the steady state becomes unstable, and price deviations from the steady state start increasing again, giving rise to a new volatile episode: Stable periods sow the seeds of volatile periods. This cycle repeats itself indefinitely, each time in a slightly different way.

For the specification used in the top-right panel of Figure B.1, we find the following Lyapunov exponents: $\lambda_1 = 0.0749$, $\lambda_2 = -0.1049$, $\lambda_3 = -0.1342$ (Benettin et al., 1980). The positive maximal exponent indicates that the dynamics is indeed chaotic. This system is more sensitive to initial conditions than the asset pricing system that

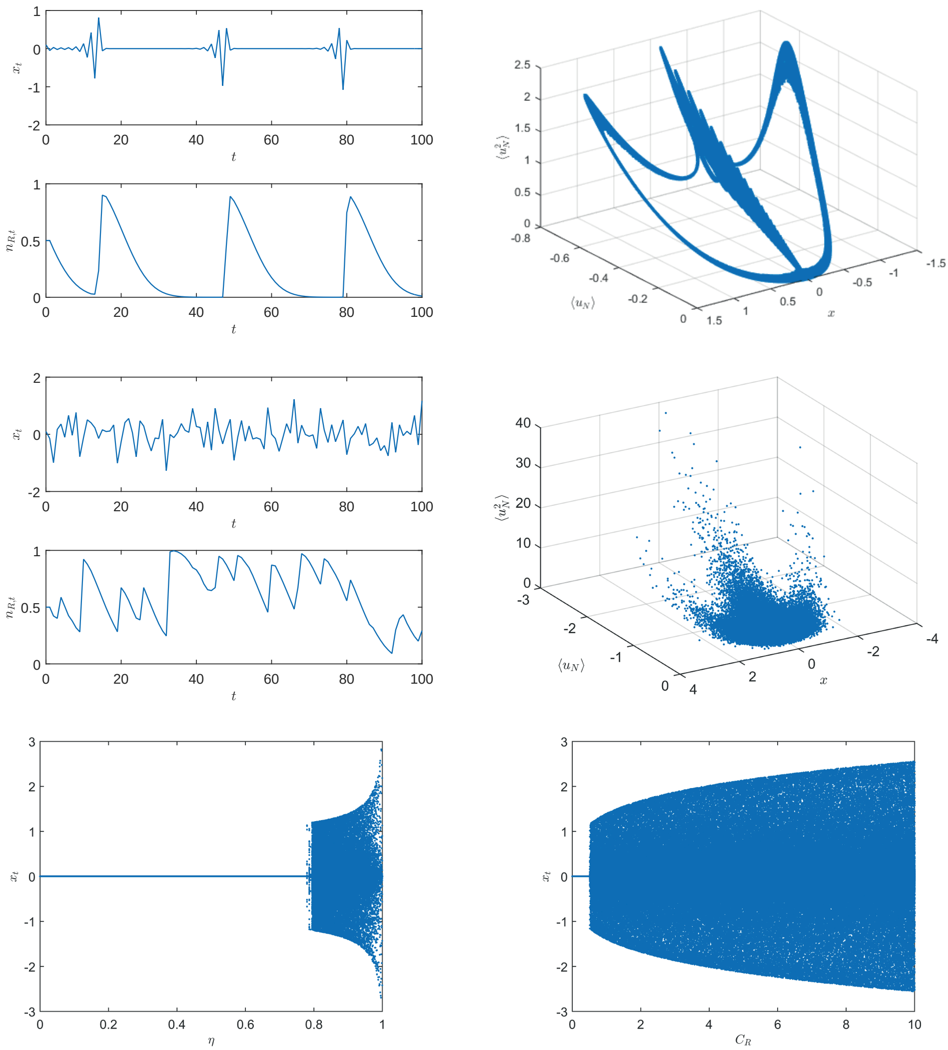

Figure B.1: *Two top-left panels:* Price deviations from the steady state (top) and fraction following rational strategy (bottom) for $\eta = 0.9$, $C_R = 1$, and $\epsilon_t = 0$. *Top-right panel:* Strange attractor for $\eta = 0.9$, $C_R = 1$, and $\epsilon_t = 0$. *Two mid-left panels:* Price deviations from the steady state (top) and fraction following rational strategy (bottom) for $\eta = 0.9$, $C_R = 1$, and with i.i.d. aggregate supply shock $\epsilon_t \sim \mathcal{N}(0,1)$. *Mid-right panel:* Noisy attractor for $\eta = 0.9$, $C_R = 1$, and with i.i.d. aggregate supply shock $\epsilon_t \sim \mathcal{N}(0,1)$. *Bottom-left panel:* Long-term behaviour of the price deviations from the steady state for different values of the memory parameter η, keeping C_R at 1. *Bottom-right panel:* Long-term behaviour of the price deviations from the steady state for different values of the cost of the rational rule C_R, keeping η at 0.9. *Parameter values:* $a = 0$, $d = 1$, $s = 2$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $p_0 = 0.10$, $\langle u_N \rangle_0 = -0.50$, $\langle u_N^2 \rangle_0 = 0.57$ (implying $n_{R,1} = 0.50$).

we studied in the main body of this chapter, which had a maximal exponent of $\lambda_1 = 0.0035$. We present the strange attractor for this solution in the top-right panel of Figure B.1. It has a box-counting dimension of approximately 1.6.

If we include a standard normal, i.i.d. aggregate supply shock $\epsilon_t \sim \mathcal{N}(0, 1)$, price deviations from the steady state become larger, and the rational rule becomes more important because of the added volatility (see the mid-left panels of Figure B.1). The strange attractor of the deterministic system is still discernible in the attractor of this noisy solution, which is larger and less well-defined, as we show in the bottom-right panel of Figure B.1.

To explore how sensitive the dynamics is to changes in the model parameters, we plot the long-term behaviour of the price deviations from the steady state for different values of η and C_R in the bottom panels of Figure B.1. For each value of the parameter, we plot the last 1,000 price deviations of 2,000 period simulation. For $\eta = 0.5$, the system stays at the steady state after the 1,000 initialization periods. For a value of η slightly above 0.8, the system becomes chaotic, and the long-term behaviour has no clear structure. The more producers remember, the larger the deviations from the steady state become, as is reflected in the increasing bandwidth for higher values of η. This can be explained by the fact that a larger η means that the utility of the naive rule is affected more gradually in times of high volatility, meaning that it takes longer for producers to switch to the rational rule and that there is therefore more time for the system to move away from the steady state.

We see a similar threshold value for the cost of the rational rule (C_R). As long as the cost is low enough, the naive rule cannot compete with the rational rule, and the system converges to the steady state eventually. However, at a cost slightly below 1, the dynamics becomes chaotic. As the cost becomes higher, the naive rule becomes relatively more attractive, and the price deviations become larger.

Appendix C

Appendix to chapter 4

C.1 Data

C.1.1 Survey of Professional Forecasters

Table C.1 describes the variables that are forecast by the SPF respondents and included in my analysis. It also reports the date of the first survey that includes the variable.

Transformations The transformation codes in Table C.1 refer to the following transformations:

II. $\widetilde{x}_t = x_t$;

V. $\widetilde{x}_t = 100 \log x_t$;

VI. $\widetilde{x}_t = 100 \left[\left(\frac{x_t}{x_{t-3}} \right)^4 - 1 \right]$.

Here, $\widetilde{x}_t$ corresponds to the transformation introduced in the theoretical framework (section 4.2.1) and x_t indicates the original series, and the 3-month running average in the case of monthly variables. As in the theoretical framework, t is measured in months, so that transformation VI denotes the annualized quarterly growth rate.

Table C.1: SPF variables included in the analysis

Variable name	First included	TCODE	Description
NGDP	1968Q4	VI	Forecasts for the quarterly and annual level of nominal GDP. Seasonally adjusted, annual rate, billions $. Prior to 1992, these are forecasts for nominal GNP.
PGDP	1968Q4	VI	Forecasts for the quarterly and annual level of the chain-weighted GDP price index. Seasonally adjusted, index, base year varies. 1992 - 1995, GDP implicit deflator. Prior to 1992, GNP implicit deflator.

Table C.1: SPF variables included in the analysis (continued)

Variable name	First included	TCODE	Description
CPROF	1968Q4	V	Forecasts for the quarterly and annual level of nominal corporate profits after tax excluding IVA and CCAdj. Seasonally adjusted, annual rate, billions $. Beginning with the survey of 2006:Q1, this variable includes IVA and CCAdj.
UNEMP	1968Q4	II	Forecasts for the quarterly average and annual average unemployment rate. Seasonally adjusted, percentage points. Quarterly forecasts are for the quarterly average of the underlying monthly levels.
EMP	2003Q4	V	Forecasts for the quarterly average and annual average level of nonfarm payroll employment. Seasonally adjusted, thousands of jobs. Quarterly forecasts are for the quarterly average of the underlying monthly levels.
INDPROD	1968Q4	V	Forecasts for the quarterly average and annual average level of the index of industrial production. Seasonally adjusted, index, base year varies. Quarterly forecasts are for the quarterly average of the underlying monthly levels.
HOUSING	1968Q4	V	Forecasts for the quarterly average and annual average level of housing starts. Seasonally adjusted, annual rate, millions. Quarterly forecasts are for the quarterly average of the underlying monthly levels.
TBILL	1981Q3	II	Forecasts for the quarterly average and annual average three-month Treasury bill rate. Percentage points. Quarterly forecasts are for the quarterly average of the underlying daily levels.
BOND	1981Q3	II	Forecasts for the quarterly average and annual average level of Moody's Aaa corporate bond yield. Percentage points. Prior to 1990:Q4, this is the new, high-grade corporate bond yield (Business Conditions Digest variable 116). Quarterly forecasts are for the quarterly average of the underlying daily levels.
BAABOND	2010Q1	II	Forecasts for the quarterly average and annual average level of Moody's Baa corporate bond yield. Percentage points. Quarterly forecasts are for the quarterly average of the underlying daily levels.
TBOND	1992Q1	II	Forecasts for the quarterly average and annual average 10-year Treasury bond rate. Percentage points. Quarterly forecasts are for the quarterly average of the underlying daily levels.

Table C.1: SPF variables included in the analysis (continued)

Variable name	First included	TCODE	Description
RGDP	1968Q4	V	Forecasts for the quarterly and annual level of chain-weighted real GDP. Seasonally adjusted, annual rate, base year varies. 1992 - 1995, fixed-weighted real GDP. Prior to 1992, fixed-weighted real GNP. Prior to 1981Q3, RGDP is computed by using the formula NGDP / PGDP * 100.
RCONSUM	1981Q3	V	Forecasts for the quarterly and annual level of chain-weighted real personal consumption expenditures. Seasonally adjusted, annual rate, base year varies. Prior to 1996, fixed-weighted real personal consumption expenditures.
RNRESIN	1981Q3	V	Forecasts for the quarterly and annual level of chain-weighted real nonresidential fixed investment. Also known as business fixed investment. Seasonally adjusted, annual rate, base year varies. Prior to 1996, fixed-weighted real nonresidential fixed investment.
RRESINV	1981Q3	V	Forecasts for the quarterly and annual level of chain-weighted real residential fixed investment. Seasonally adjusted, annual rate, base year varies. Prior to 1996, fixed-weighted real residential fixed investment.
RFEDGOV	1981Q3	V	Forecasts for the quarterly and annual level of chain-weighted real federal government consumption and gross investment. Seasonally adjusted, annual rate, base year varies. Prior to 1996, real fixed-weight federal government purchases of goods and services.
RSLGOV	1981Q3	V	Forecasts for the quarterly and annual level of chain-weighted real state and local government consumption and gross investment. Seasonally adjusted, annual rate, base year varies. Prior to 1996, real fixed-weighted state and local government purchases of goods and services.
CPI	1981Q3[a]	II	Forecasts for the headline CPI inflation rate. Seasonally adjusted, annual rate, percentage points. Quarterly forecasts are annualized quarter-over-quarter percent changes of the quarterly average price index level. The quarterly price index level is the quarterly average of the underlying monthly price index levels.
CORECPI	2007Q1	II	Forecasts for the core CPI inflation rate. Seasonally adjusted, annual rate, percentage points. Quarterly forecasts are annualized quarter-over-quarter percent changes of the quarterly average price index level. The quarterly price index level is the quarterly average of the underlying monthly price index levels.

Table C.1: SPF variables included in the analysis (continued)

Variable name	First included	TCODE	Description
PCE	2007Q1	II	Forecasts for the headline chain-weighted PCE inflation rate. Seasonally adjusted, annual rate, percentage points. Quarterly forecasts are annualized quarter-over-quarter percent changes of the quarterly average price index level. The quarterly price index level is the quarterly average of the underlying monthly price index levels.
COREPCE	2007Q1	II	Forecasts for the core chain-weighted PCE inflation rate. Seasonally adjusted, annual rate, percentage points. Quarterly forecasts are annualized quarter-over-quarter percent changes of the quarterly average price index level. The quarterly price index level is the quarterly average of the underlying monthly price index levels.

The second column (First included) gives the date of the first survey that includes forecasts for the variable. Transformation codes (TCODE) are explained above. Variable descriptions are taken from SPF documentation version 6 May 2020 (see `https://www.philadelphiafed.org/-/media/frbp/assets/surveys-and-data/survey-of-professional-forecasters/spf-documentation.pdf`).
[a] Real-time seasonally adjusted CPI data is only available for vintages from 1994 onwards. I therefore only include CPI forecasts in the analysis from 1994 onwards.

C.1.2 Real-time macroeconomic panel

In Table C.2, I describe all variables that are included in the real-time panel that I use to extract macroeconomic factors. It includes all variables that are included in the SPF (Table C.1), with monthly observations where possible. When the data is quarterly, I mention it in the 'Note' column of the table. The panel contains many additional variables in the groups (a) output and income, (b) the labour market, (c) housing, (d) consumption, investment, orders, and inventories, (e) money and credit, (f) interest and exchange rates, (g) prices, (h) the stock market. For each variable, I report the date of the first survey for which it is included in the real-time panel. In Figure C.2, I plot the panel width over time. For some variables, real-time data becomes unavailable for a limited time after the first time it is included in the panel. This causes incidental drops in the panel width. Additionally, I report the earliest available observation in the real-time panel associated with the 2019Q4 survey.

Transformations The transformation codes are different for the real-time panel than used in Table C.1. I denote the original (monthly) series by x_t, and the transformed series by X_t:

1. $X_t = x_t$;
2. $X_t = x_t - x_{t-3}$;
5. $X_t = 100(\log x_t - \log x_{t-3})$;
6. $X_t = 100(\log x_t - 2\log x_{t-3} + \log x_{t-6})$;

7. $X_t = 100\left(\frac{x_t}{x_{t-3}} - \frac{x_{t-3}}{x_{t-6}}\right)$.

The transformations are based on those from the FRED-MD panel (McCracken and Ng, 2016). I use 3-month differences instead of monthly differences, because I use the extracted factors to forecast quarterly quantities. Furthermore, I apply different transformations to some of the variables based on stationarity tests and in line with the transformations used by Jurado et al. (2015) for their diffusion index forecasts.

Seasonal adjustment For the Manufacturing Business Outlook Survey (BOS) and consumer price index, only non-seasonally adjusted data is available in real-time. (In the case of CPI, real-time seasonally adjusted data is available from 1994.) As these are important elements of the macroeconomic panel, I calculate a real-time seasonal adjustment using a simplified[1] version of the X-11-ARIMA seasonal adjustment program (Dagum, 1980). X-11-ARIMA is based on the X-11 program of the U.S. Census Bureau that was introduced in 1965 (Findley et al., 1998). It was used by the U.S. Bureau of Labor Statistics until the introduction of its successor X-12-ARIMA. I compute back- and forecasts using a data-determined ARIMA model[2] and follow the procedure outlined in Appendix A of Findley et al. (1998). For BOS data, I use an additive decomposition into trend, seasonal, and irregular components, and only compute the seasonal adjustment when at least 5 years of data is available.[3] For CPI, I use a multiplicative decomposition and use the logged series to compute back- and forecasts. I compare my own adjustment to the official one for the most recent vintage (2019Q4). For the BOS, the root-mean-squared difference between the official seasonally adjusted data and my own adjusted data is 2.32, down from 10.72 before the seasonal adjustment. For CPI, the root-mean-squared difference is 0.15, compared to 0.38 before the adjustment. In Figure C.1, I compare the official seasonally adjusted data to my own adjustment. It shows that the simplified X-11-ARIMA program does a good job of approximating the official adjustment. Note that I do not use my own seasonal adjustment to calculate empirical CPI forecasts before 1994, not to introduce forecast differences that are due to the seasonal adjustment.

[1] The simplification mostly lies in the fact that I do not use an extreme value correction, and that I use a 13-term Henderson filter instead of a data-determined filter length.

[2] The ARIMA specification-selection algorithm is based on Hyndman and Khandakar (2008). I only carry out the first step for computational efficiency.

[3] For the first 5 years of the BOS, I include data that is not seasonally adjusted.

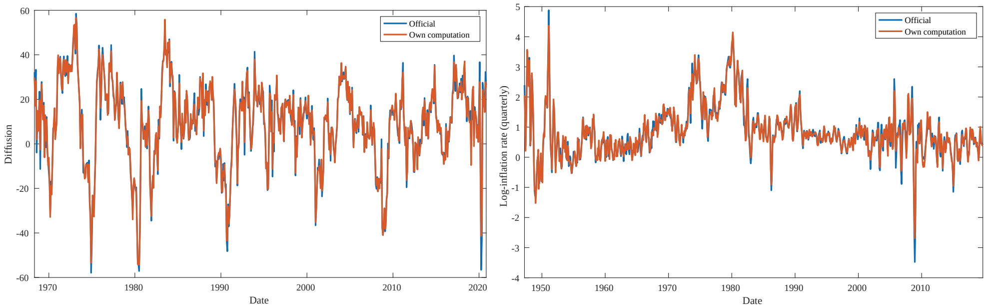

Figure C.1: Comparison of own seasonal adjustment to official seasonal adjustment for the general activity index from the BOS (*left*) and the quarterly log-inflation rate (*right*).

Table C.2: Variables included in the real-time panel

Description	First incl.	First obs.	TCODE	Note
		Output and income		
Real GDP	1968Q4	1959:03	5	Quarterly
Nominal GDP	1968Q4	1959:03	6	Quarterly
Real personal income	1968Q4	1959:04	5	Deflated by CPI
Corporate profits	1968Q4	1959:03	5	Quarterly
Industrial production index	1968Q4	1959:01	5	
IP index: manufacturing	1968Q4	1959:01	5	
BOS current activity index	1969Q1	1968:08	2	Own seasonal adjustment
BOS future activity index	1969Q1	1968:08	2	Own seasonal adjustment
Capacity utilization: manufacturing	1979Q4	1959:01	2	
Real exports	1968Q4	1959:03	5	Quarterly
Real imports	1968Q4	1959:03	5	Quarterly
		Labour market		
Civilian labour force	1968Q4	1959:01	5	
Civilian employment	1968Q4	1959:01	5	
Civilian unemployment rate	1968Q4	1959:01	2	
Average duration of unemployment	1972Q1	1959:01	2	In weeks
Civilians unemployed < 5 weeks	1968Q4	1959:01	5	
Civilians unemployed $5 - 14$ weeks	1968Q4	1959:01	5	
Civilians unemployed 15+ weeks	1968Q4	1959:01	5	
Civilians unemployed $15 - 26$ weeks	1982Q1	1959:01	5	
Civilians unemployed 27+ weeks	1982Q1	1959:01	5	
All employees: total nonfarm	1968Q4	1959:01	5	
All employees: goods-producing	1971Q4	1959:01	5	
All employees: construction	1968Q4	1959:01	5	
All employees: manufacturing	1968Q4	1959:01	5	
All employees: durable goods	1968Q4	1959:01	5	
All employees: nondurable goods	1968Q4	1959:01	5	
All employees: service-producing	1971Q4	1959:01	5	

Table C.2: Variables included in the real-time panel (continued)

Description	First incl.	First obs.	TCODE	Note
All employees: trade, transp. & util.	1968Q4	1959:01	5	
All employees: wholesale trade	1968Q4	1959:01	5	
All employees: retail trade	1968Q4	1959:01	5	
All employees: financial activities	1968Q4	1959:01	5	
All employees: government	1968Q4	1959:01	5	
Avg weekly hours: manufacturing	1968Q4	1959:01	2	
Avg weekly overtime hours: man.	1968Q4	1959:01	2	
BOS current employment	1969Q1	1968:08	2	Own seasonal adjustment
BOS current average workweek	1969Q1	1968:08	2	Own seasonal adjustment
Housing				
Housing starts	1968Q4	1959:04	5	
Housing starts: 1-unit structures	1972Q2	1959:04	5	
Housing starts: northeast	1973Q2	1959:04	5	
Housing starts: midwest	1973Q2	1959:04	5	
Housing starts: south	1973Q2	1959:04	5	
Housing starts: west	1973Q2	1959:04	5	
New private housing permits	1999Q4	1960:04	5	
New 1-family houses sold	1999Q3	1963:04	5	
Real residential fixed investment	1968Q4	1959:03	5	Quarterly
Consumption, investment, orders, and inventories				
Real pers. consumption expenditures	1968Q4	1959:04	5	Quarterly before 1980Q1[a]
BOS current new orders	1969Q1	1968:08	2	Own seasonal adjustment
BOS current shipments	1969Q1	1968:08	2	Own seasonal adjustment
BOS current unfilled orders	1969Q1	1968:08	2	Own seasonal adjustment
BOS current delivery time	1969Q1	1968:08	2	Own seasonal adjustment
BOS current inventories	1969Q1	1968:08	2	Own seasonal adjustment
BOS future capital expenditures	1969Q1	1968:08	2	Own seasonal adjustment
Consumer sentiment index	1968Q4	1960:02	2	Monthly from Jan 1978[b]

Table C.2: Variables included in the real-time panel (continued)

Description	First incl.	First obs.	TCODE	Note
Real nonresidential fixed investment	1968Q4	1959:03	5	Quarterly
Real federal gov. cons. & inv.	1968Q4	1959:03	5	Quarterly
Real state and local gov. cons. & inv.	1968Q4	1959:03	5	Quarterly
		Money and credit		
M1 money stock	1968Q4	1959:07	6	
M2 money stock	1971Q2	1959:07	6	
Real M2 money stock	1971Q2	1959:04	5	Deflated by CPI
Monetary base	1980Q2	1959:07	6	Missing 1982Q4 – 1985Q2, and 1992Q1 – 1992Q2
Total reserves of depository inst.	1968Q4	1959:07	6	Missing 1969Q4 – 1980Q1, 1982Q4 – 1985Q2, and 1992Q1 – 1992Q2
Nonborrowed reserves of dep. inst.	1968Q4	1959:07	7	Missing 1969Q4 – 1980Q1, 1982Q4 – 1985Q2, and 1992Q1 – 1992Q2
Commercial and industrial loans	1997Q1	1959:01	6	
Real estate loans	1997Q1	1959:01	6	
Total nonrevolving credit	1999Q3	1959:01	6	
Nonrevolving credit to personal income	1999Q3	1959:04	2	Own calculation
Total consumer credit	1997Q1	1959:01	6	Owned and securitized
Money zero maturity (MZM) stock	1998Q4	1959:07	6	
Securities in bank credit	1998Q4	1959:01	6	
		Interest and exchange rates		
Effective federal funds rate	1968Q4	1959:04	2	
3-Month aa fin. commercial paper	1968Q4	1959:04	2	
3-Month treasury bill	1968Q4	1959:01	2	
6-Month treasury bill	1968Q4	1959:04	2	
1-Year treasury rate	1968Q4	1959:04	2	
5-Year treasury rate	1968Q4	1959:04	2	
10-Year treasury rate	1968Q4	1959:01	2	
Moody's aaa corporate bond yield	1968Q4	1959:01	2	

Table C.2: Variables included in the real-time panel (continued)

Description	First incl.	First obs.	TCODE	Note
Moody's baa corporate bond yield	1968Q4	1959:01	2	
3-Month comm. paper – funds rate	1968Q4	1959:01	1	
3-Month treasury – funds rate	1968Q4	1959:01	1	
6-Month treasury – funds rate	1968Q4	1959:01	1	
1-Year treasury – funds rate	1968Q4	1959:01	1	
5-Year treasury – funds rate	1968Q4	1959:01	1	
10-Year treasury – funds rate	1968Q4	1959:01	1	
Moody's aaa bond yield – funds rate	1968Q4	1959:01	1	
Moody's baa bond yield – funds rate	1968Q4	1959:01	1	
Switzerland / U.S. exchange rate	1968Q4	1959:04	5	
Japan / U.S. exchange rate	1968Q4	1959:04	5	
U.S. / U.K. exchange rate	1968Q4	1959:04	5	
Canada / U.S. exchange rate	1968Q4	1959:04	5	
		Prices		
GDP price index	1968Q4	1959:03	6	Quarterly
Consumer price index: all items	1968Q4	1959:01	6	Own s.a. before 1994[c]
Core consumer price index	1997Q1	1959:01	6	Excl. food and energy
Personal cons. exp.: price index	1968Q4	1959:07	6	Quarterly before 2000Q3[a]
Core personal cons. exp.: price index	1996Q1	1959:07	6	Excl. food and energy; quarterly before 2000Q3[a]
Producer price index: all commodities	1997Q1	1959:01	6	
Crude oil price	1985Q2	1959:07	6	
BOS current prices paid	1969Q1	1968:08	2	Own seasonal adjustment
BOS current prices received	1969Q1	1968:08	2	Own seasonal adjustment
		Stock market		
S&P 500: composite	1968Q4	1959:04	5	
S&P 500: industrials	1968Q4	1959:04	5	
S&P500: dividend yield	1968Q4	1959:04	2	
S&P500: price-earnings ratio	1968Q4	1959:04	5	

Table C.2: Variables included in the real-time panel (continued)

Description	First incl.	First obs.	TCODE	Note
VXO	1968Q4	1962:07	1	

The second column (First incl.) gives the date of the first survey whose associated real-time panel included the variable. The third column (First obs.) indicates the month of the first observation for the variable in the panel associated with the 2019Q4 survey. Transformation codes (TCODE) are explained above. Own seasonal adjustment refers to the adjustment described above. [a] Quarterly before YYYYQQ means that real-time panels associated with surveys before the YYYYQQ survey contain a quarterly series of the variable, whereas later panels contain a monthly series. [b] This consumer sentiment index is based on the University of Michigan Surveys of Consumers. This survey has been conducted monthly only from January 1978 onwards, but quarterly observations are available from February 1960 until that time.
[c] For vintages before 1994, only non-seasonally adjusted data is available. From 1994 onwards, seasonally adjusted data is available for the full length of the sample.

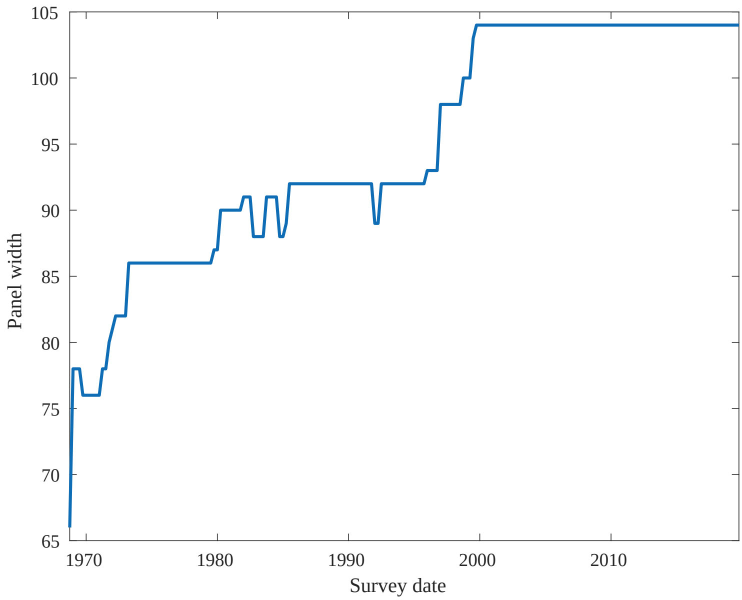

Figure C.2: Width of the real-time panel for all surveys.

C.2 Proof of Proposition 5

C.2.1 Empirical distribution

I start with the empirical distribution. I first prove that time $t+h$ distributions of mean market belief, factors, and changes in observables can be expressed in terms of time t values.

Lemma 2. *The empirical distributions of mean market belief, factors, and changes in observables at time $t+h$ ($h \geq 1$) can be written*

$$Z^m_{t+h} = A_{hZ} Z_t + \epsilon^{hZ}_{t+h}, \qquad \epsilon^{hZ}_{t+h} \sim \mathcal{N}(0, \Sigma_{hZ}), \tag{C.1}$$
$$F^m_{t+h} = A_{hF} F_t + A^Z_{hF} Z_t + \epsilon^{hF}_{t+h}, \qquad \epsilon^{hF}_{t+h} \sim \mathcal{N}(0, \Sigma_{hF}), \tag{C.2}$$
$$\widehat{x}^m_{t+h} = \bar{c}_{hx} + \bar{\lambda}_{hx} \widehat{x}_t + \bar{A}^F_{hx} F_t + \bar{A}^Z_{hx} Z_t + \epsilon^{hx}_{t+h}, \qquad \epsilon^{hx}_{t+h} \sim \mathcal{N}(0, \bar{\sigma}^2_{hx}). \tag{C.3}$$

Proof. I start with mean market belief (C.1). The base case $h = 1$ follows directly from equation (4.17). Simply define

$$A_{1Z} = A_Z, \quad \epsilon^{1Z}_{t+1} = \bar{\epsilon}^Z_{t+1}.$$

Now assume that equation (C.1) holds for $h = \eta \geq 1$. It follows that

$$\begin{aligned}
Z^m_{t+\eta+1} &= A_{1Z} Z^m_{t+\eta} + \epsilon^{1Z}_{t+\eta+1} \\
&= A_{1Z} \left(A_{\eta Z} Z_t + \epsilon^{\eta Z}_{t+\eta} \right) + \epsilon^{1Z}_{t+\eta+1} \\
&= A_{\eta+1,Z} Z_t + \epsilon^{\eta+1,Z}_{t+\eta+1},
\end{aligned}$$

where

$$A_{\eta+1,Z} = A_{1Z} A_{\eta Z}, \quad \epsilon^{\eta+1,Z}_{t+\eta+1} \sim \mathcal{N}(0, \Sigma_{\eta+1,Z}), \quad \Sigma_{\eta+1,Z} = A_{1Z} \Sigma_{\eta Z} A^T_{1Z} + \Sigma_{1Z}.$$

This concludes the proof for mean market belief by induction. I can use this result to make a similar argument for F^m_{t+h}. Again, the case $h = 1$ is trivial, following directly from equation (4.16). Assuming that (C.2) holds for $h = \eta \geq 1$, it follows that

$$\begin{aligned}
F^m_{t+\eta+1} &= A_{1F} F^m_{t+\eta} + A^Z_{1F} Z^m_{t+\eta} + \epsilon^{1F}_{t+\eta+1} \\
&= A_{1F} \left(A_{\eta F} F_t + \epsilon^{\eta F}_{t+\eta} \right) + A^Z_{1F} \left(A_{\eta Z} Z_t + \epsilon^{\eta Z}_{t+\eta} \right) + \epsilon^{1F}_{t+\eta+1} \\
&= A_{\eta+1,F} F_t + A^Z_{\eta+1,F} Z_t + \epsilon^{\eta+1,F}_{t+\eta+1}.
\end{aligned}$$

Here, I define

$$A_{\eta+1,F} = A_{1F} A_{\eta F}, \quad A^Z_{\eta+1,F} = A^Z_{1F} A_{\eta Z},$$

and it follows that

$$\epsilon^{\eta+1,F}_{t+\eta+1} \sim \mathcal{N}(0, \Sigma_{\eta+1,F}), \quad \Sigma_{\eta+1,F} = A_{1F} \Sigma_{\eta F} A^T_{1F} + A^Z_{1F} \Sigma_{\eta Z} \left(A^Z_{1F} \right)^T + \Sigma_{1F}.$$

Finally, the proof for the stationary change in observables follows the same reasoning. The case $h = 1$ follows from (4.18). Now suppose that the result is valid for some case

$h = \eta \geq 1$. I have that

$$
\begin{aligned}
\widehat{x}^m_{t+\eta+1} &= \bar{c}_{1x} + \bar{\lambda}_{1x}\widehat{x}^m_{t+\eta} + \bar{A}^F_{1x}F^m_{t+\eta} + \bar{A}^Z_{1x}Z^m_{t+\eta} + \epsilon^{1x}_{t+\eta+1} \\
&= \bar{c}_{1x} + \bar{\lambda}_{1x}\left(\bar{c}_{\eta x} + \bar{\lambda}_{\eta x}\widehat{x}_t + \bar{A}^F_{\eta x}F_t + \bar{A}^Z_{\eta x}Z_t + \epsilon^{\eta x}_{t+\eta}\right) + \\
&\qquad \bar{A}^F_{1x}\left(A_{\eta F}F_t + A^Z_{\eta F}Z_t + \epsilon^{\eta F}_{t+\eta}\right) + \bar{A}^Z_{1x}\left(A_{\eta Z}Z_t + \epsilon^{\eta Z}_{t+\eta}\right) + \epsilon^{1x}_{t+\eta+1} \\
&= \bar{c}_{\eta+1,x} + \bar{\lambda}_{\eta+1,x}\widehat{x}_t + \bar{A}^F_{\eta+1,x}F_t + \bar{A}^Z_{\eta+1,x}Z_t + \epsilon^{\eta+1,x}_{t+\eta+1},
\end{aligned}
$$

where

$$
\begin{aligned}
&\bar{c}_{\eta+1,x} = \bar{c}_{1x} + \bar{\lambda}_{1x}\bar{c}_{\eta x}, \quad \bar{\lambda}_{\eta+1,x} = \bar{\lambda}_{1x}\bar{\lambda}_{\eta x}, \quad \bar{A}^F_{\eta+1,x} = \bar{\lambda}_{1x}\bar{A}^F_{\eta x} + \bar{A}^F_{1x}A_{\eta F}, \\
&\bar{A}^Z_{\eta+1,x} = \bar{\lambda}_{1x}\bar{A}^Z_{\eta x} + \bar{A}^F_{1x}A^Z_{\eta F} + \bar{A}^Z_{1x}A_{\eta Z}, \quad \epsilon^{\eta+1,x}_{t+\eta+1} \sim \mathcal{N}(0, \bar{\sigma}^2_{\eta+1,x}), \\
&\bar{\sigma}^2_{\eta+1,x} = \bar{\lambda}^2_{1x}\bar{\sigma}^2_{\eta x} + \bar{A}^F_{1x}\,\Sigma_{\eta F}\left(\bar{A}^F_{1x}\right)^T + \bar{A}^Z_{1x}\,\Sigma_{\eta Z}\left(\bar{A}^Z_{1x}\right)^T + \bar{\sigma}^2_{1x}.
\end{aligned}
$$

This concludes the proof by induction. □

I am now in a position to prove the empirical distribution in Proposition 5.

Proof of Proposition 5 (empirical distribution). This is another proof by induction. The base case $h = 1$ follows directly from the fact that $[\widetilde{x}_{t+1} - \widetilde{x}_t]^m = \widehat{x}^m_{t+1}$. Now suppose that the proposition holds for $h = \eta \geq 1$. It follows that

$$
\begin{aligned}
[\widetilde{x}_{t+\eta+1} - \widetilde{x}_t]^m &= [\widetilde{x}_{t+\eta+1} - \widetilde{x}_{t+\eta}]^m + [\widetilde{x}_{t+\eta} - \widetilde{x}_t]^m \\
&= \widehat{x}^m_{t+\eta+1} + c_{\eta x} + \lambda_{\eta x}\widehat{x}_t + A^F_{\eta x}F_t + A^Z_{\eta x}Z_t + u^{\eta x}_{t+\eta} \\
&= \bar{c}_{\eta+1,x} + \bar{\lambda}_{\eta+1,x}\widehat{x}_t + \bar{A}^F_{\eta+1,x}F_t + \bar{A}^Z_{\eta+1,x}Z_t + \epsilon^{\eta+1,x}_{t+\eta+1} + \\
&\qquad c_{\eta x} + \lambda_{\eta x}\widehat{x}_t + A^F_{\eta x}F_t + A^Z_{\eta x}Z_t + u^{\eta x}_{t+\eta} \\
&= c_{\eta+1,x} + \lambda_{\eta+1,x}\widehat{x}_t + A^F_{\eta+1,x}F_t + A^Z_{\eta+1,x}Z_t + u^{\eta+1,x}_{t+\eta+1},
\end{aligned}
$$

with

$$
\begin{aligned}
&c_{\eta+1,x} = \bar{c}_{\eta+1,x} + c_{\eta x}, \quad \lambda_{\eta+1,x} = \bar{\lambda}_{\eta+1,x} + \lambda_{\eta x}, \quad A^F_{\eta+1,x} = \bar{A}^F_{\eta+1,x} + A^F_{\eta x}, \\
&A^Z_{\eta+1,x} = \bar{A}^Z_{\eta+1,x} + A^Z_{\eta x}, \quad u^{\eta+1,x}_{t+\eta+1} \sim \mathcal{N}(0, \sigma^2_{\eta+1,x}), \quad \sigma^2_{\eta+1,x} = \bar{\sigma}^2_{\eta+1,x} + \sigma^2_{\eta x}.
\end{aligned}
$$

This concludes the proof. □

C.2.2 Perceived distribution

The proof for the perceived distribution is similar. The main difference is that I have to take into account individual beliefs. Again, I first need to establish expressions for the time $t + h$ distributions.

Lemma 3. *Agent i's perceived distributions of her individual beliefs, mean market belief, factors, and changes in observables at time $t+h$ ($h \geq 1$) can be written*

$$g^i_{t+h} = A_{hg} g^i_t + \epsilon^{ihg}_{t+h}, \qquad \epsilon^{ihg}_{t+h} \sim \mathcal{N}(0, \Sigma_{hg}), \tag{C.4}$$

$$Z^i_{t+h} = A_{hZ} Z_t + A^g_{hZ} g^i_t + \epsilon^{ihZ}_{t+h}, \qquad \epsilon^{ihZ}_{t+h} \sim \mathcal{N}(0, \check{\Sigma}_{hZ}), \tag{C.5}$$

$$F^i_{t+h} = A_{hF} F_t + A^Z_{hF} Z_t + A^g_{hF} g^i_t + \epsilon^{ihF}_{t+h}, \qquad \epsilon^{ihF}_{t+h} \sim \mathcal{N}(0, \check{\Sigma}_{hF}), \tag{C.6}$$

$$\widehat{x}^i_{t+h} = \bar{c}_{hx} + \bar{\lambda}_{hx} \widehat{x}_t + \bar{A}^F_{hx} F_t + \bar{A}^Z_{hx} Z_t + \tag{C.7}$$

$$\bar{A}^g_{hx} g^i_t + \bar{\rho}^{ihx}_t + \epsilon^{ihx}_{t+h}, \qquad \epsilon^{ihx}_{t+h} \sim \mathcal{N}(0, \check{\sigma}^2_{hx}). \tag{C.8}$$

Proof. I start with the base case $h = 1$ for the individual beliefs distribution. It follows from the perceived distribution of the factors (4.19) that

$$F^i_{t+1} - A_F F_t - A^Z_F A_Z Z_t = \bar{A}^g_F g^i_t + \bar{\epsilon}^{iF}_{t+1}.$$

Plugging this into the transition function for beliefs (4.5) gives

$$\begin{aligned} g^i_{t+1} &= A_Z g^i_t + A^F_Z \left(\bar{A}^g_F g^i_t + \bar{\epsilon}^{iF}_{t+1} \right) + \epsilon^{ig}_{t+1}. \\ &= A_{1g} g^i_t + \epsilon^{i1g}_{t+1}, \end{aligned}$$

where

$$A_{1g} = A_Z + A^F_Z \bar{A}^g_F, \quad \epsilon^{i1g}_{t+1} \sim \mathcal{N}(0, \Sigma_{1g}), \quad \Sigma_{1g} = A^F_Z \check{\Sigma}_F \left(A^F_Z\right)^T + \Sigma_g.$$

Now assume that equation (C.4) holds for $h = \eta \geq 1$. It follows that

$$\begin{aligned} g^i_{t+\eta+1} &= A_{1g} g^i_{t+\eta} + \epsilon^{i1g}_{t+\eta+1} \\ &= A_{1g} \left(A_{\eta g} g^i_t + \epsilon^{i\eta g}_{t+\eta} \right) + \epsilon^{i1g}_{t+\eta+1} \\ &= A_{\eta+1,g} g^i_t + \epsilon^{i,\eta+1,g}_{t+\eta+1}, \end{aligned}$$

where

$$A_{\eta+1,g} = A_{1g} A_{\eta g}, \quad \epsilon^{i,\eta+1,g}_{t+\eta+1} \sim \mathcal{N}(0, \Sigma_{\eta+1,g}), \quad \Sigma_{\eta+1,g} = A_{1g} \Sigma_{\eta g} A^T_{1g} + \Sigma_{1g}.$$

This concludes the proof for individual beliefs by induction. Now for mean market belief. The base case ($h = 1$) again follows directly. In particular, note that A_{1Z} is the same as the coefficient matrix appearing in Lemma 2 (both being equal to A_Z). Now assume that (C.5) is valid for some $\eta \geq 1$. It follows that

$$\begin{aligned} Z^i_{t+\eta+1} &= A_{1Z} Z^i_{t+\eta} + A^g_{1Z} g^i_{t+\eta} + \epsilon^{i1Z}_{t+\eta+1} \\ &= A_{1Z} \left(A_{\eta Z} Z_t + A^g_{\eta Z} g^i_t + \epsilon^{i\eta Z}_{t+\eta} \right) + A^g_{1Z} \left(A_{\eta g} g^i_t + \epsilon^{i\eta g}_{t+\eta} \right) + \epsilon^{i1Z}_{t+\eta+1} \\ &= A_{\eta+1,Z} Z_t + A^g_{\eta+1,Z} g^i_t + \epsilon^{i,\eta+1,Z}_{t+\eta+1}, \end{aligned}$$

where $A_{\eta+1,Z}$ is the same as in Lemma 2 and

$$\begin{aligned} A^g_{\eta+1,Z} &= A_{1Z} A^g_{\eta Z} + A^g_{1Z} A_{\eta g}, \quad \epsilon^{i,\eta+1,Z}_{t+\eta+1} \sim \mathcal{N}(0, \check{\Sigma}_{\eta+1,Z}) \\ \check{\Sigma}_{\eta+1,Z} &= A_{1Z} \check{\Sigma}_{\eta Z} A^T_{1Z} + A^g_{1Z} \Sigma_{\eta g} \left(A^g_{1Z}\right)^T + \check{\Sigma}_{1Z}. \end{aligned}$$

This proves the case of mean market belief by induction. The proof for (C.6) combines (C.4) and (C.5) through induction. I omit it here, because it is very similar to the previous proof. Now for the stationary change in observables. I use the same induction recipe: I build on the trivial $h = 1$ case and assume that the result holds for some horizon $\eta \geq 1$, which gives

$$\begin{aligned}
\widehat{x}^i_{t+\eta+1} &= \bar{c}_{1x} + \bar{\lambda}_{1x}\widehat{x}^i_{t+\eta} + \bar{A}^F_{1x}F^i_{t+\eta} + \bar{A}^Z_{1x}Z^i_{t+\eta} + \bar{A}^g_{1x}g^i_{t+\eta} + \bar{\rho}^{i1x}_{t+\eta} + \epsilon^{i1x}_{t+\eta+1} \\
&= \bar{c}_{1x} + \bar{\lambda}_{1x}\left(\bar{c}_{\eta x} + \bar{\lambda}_{\eta x}\widehat{x}_t + \bar{A}^F_{\eta x}F_t + \bar{A}^Z_{\eta x}Z_t + \bar{A}^g_{\eta x}g^i_t + \bar{\rho}^{i\eta x}_t + \epsilon^{i\eta x}_{t+\eta}\right) + \\
&\qquad \bar{A}^F_{1x}\left(A_{\eta F}F_t + A^Z_{\eta F}Z_t + A^g_{\eta F}g^i_t + \epsilon^{i\eta F}_{t+\eta}\right) + \bar{A}^Z_{1x}\left(A_{\eta Z}Z_t + A^g_{\eta Z}g^i_t + \epsilon^{i\eta Z}_{t+\eta}\right) + \\
&\qquad \bar{A}^g_{1x}\left(A_{\eta g}g^i_t + \epsilon^{i\eta g}_{t+\eta}\right) + \bar{\rho}^{i1x}_{t+\eta} + \epsilon^{i1x}_{t+\eta+1} \\
&= \bar{c}_{\eta+1,x} + \bar{\lambda}_{\eta+1,x}\widehat{x}_t + \bar{A}^F_{\eta+1,x}F_t + \bar{A}^Z_{\eta+1,x}Z_t + \bar{A}^g_{\eta+1,x}g^i_t + \bar{\rho}^{i,\eta+1,x}_t + \epsilon^{i,\eta+1,x}_{t+\eta+1}.
\end{aligned}$$

Here, all coefficients that do not refer to individual beliefs are the same as in Lemma 2. The other coefficients are given by

$$\begin{aligned}
&\bar{A}^g_{\eta+1,x} = \bar{\lambda}_{1x}\bar{A}^g_{\eta x} + \bar{A}^F_{1x}A^g_{\eta F} + \bar{A}^Z_{1x}A^g_{\eta Z} + \bar{A}^g_{1x}A_{\eta g}, \quad \bar{\rho}^{i,\eta+1,x}_t = \bar{\lambda}_{1x}\bar{\rho}^{i\eta x}_t, \\
&\epsilon^{i,\eta+1,x}_{t+\eta+1}\mathcal{N}(0, \check{\sigma}^2_{\eta+1,x}), \quad \check{\sigma}^2_{\eta+1,x} = \bar{\lambda}^2_{1x}\check{\sigma}^2_{\eta x} + \bar{A}^F_{1x}\,\check{\Sigma}_{\eta F}\left(\bar{A}^F_{1x}\right)^T + \bar{A}^Z_{1x}\,\check{\Sigma}_{\eta Z}\left(\bar{A}^Z_{1x}\right)^T + \\
&\bar{A}^g_{1x}\,\Sigma_{\eta g}\left(\bar{A}^g_{1x}\right)^T + \sigma^2_{x\rho} + \check{\sigma}^2_{1x},
\end{aligned}$$

where I use the fact that $\bar{\rho}^{i1x}_{t+\eta} = \rho^{ix}_{t+\eta} \sim \mathcal{N}(0, \sigma^2_{x\rho})$. This concludes the proof by induction. □

The perceived distribution of h-period change to observables follows readily from Lemma 3.

Proof of Proposition 5 (perceived distribution). This proof is analogous to the case of the empirical distribution, with a few extra terms for individual beliefs. □

C.3 Macroeconomic impact of market belief: Additional results

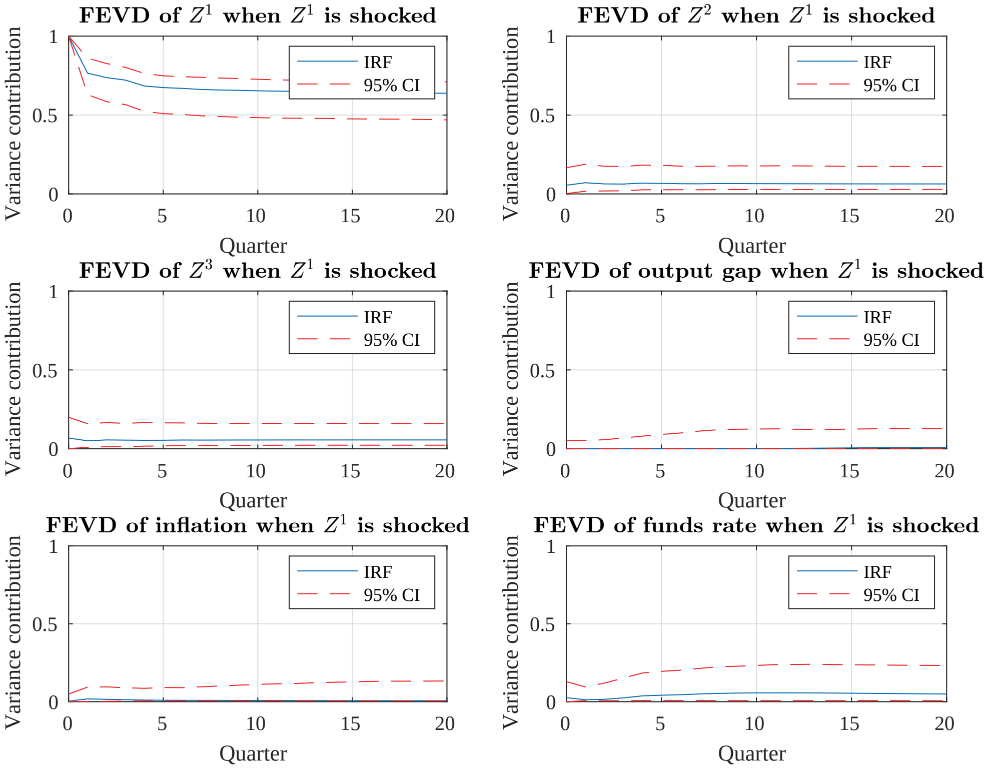

Figure C.3: VAR FEVD when the first dimension of mean market belief is shocked. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

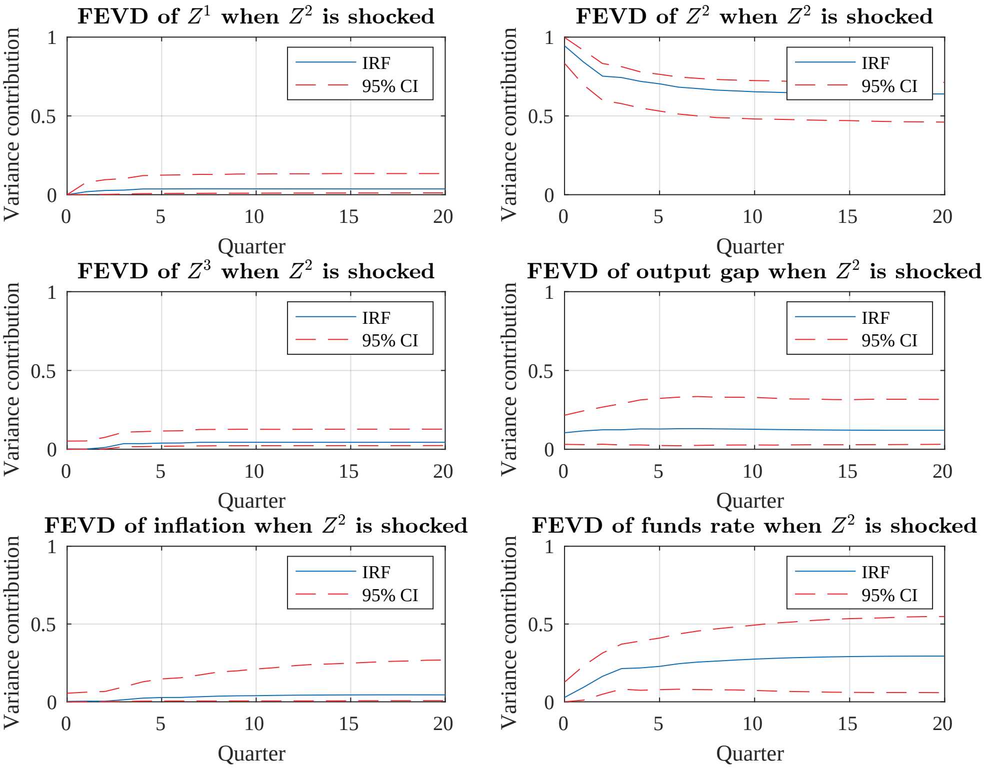

Figure C.4: VAR FEVD when the second dimension of mean market belief is shocked. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

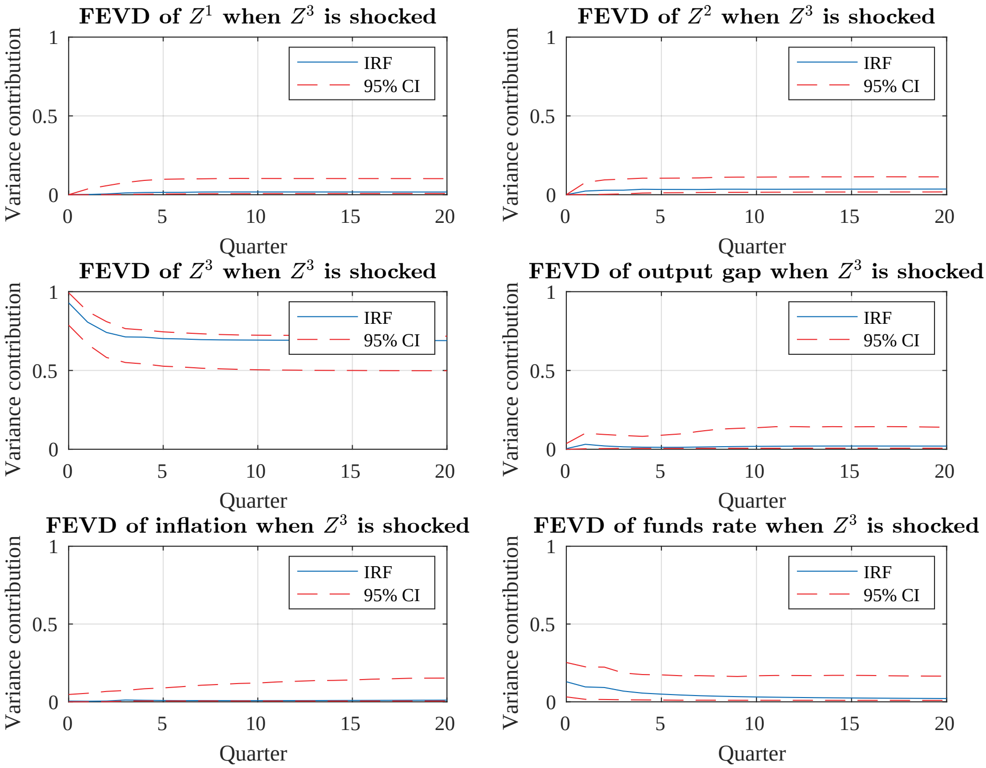

Figure C.5: VAR FEVD when the third dimension of mean market belief is shocked. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

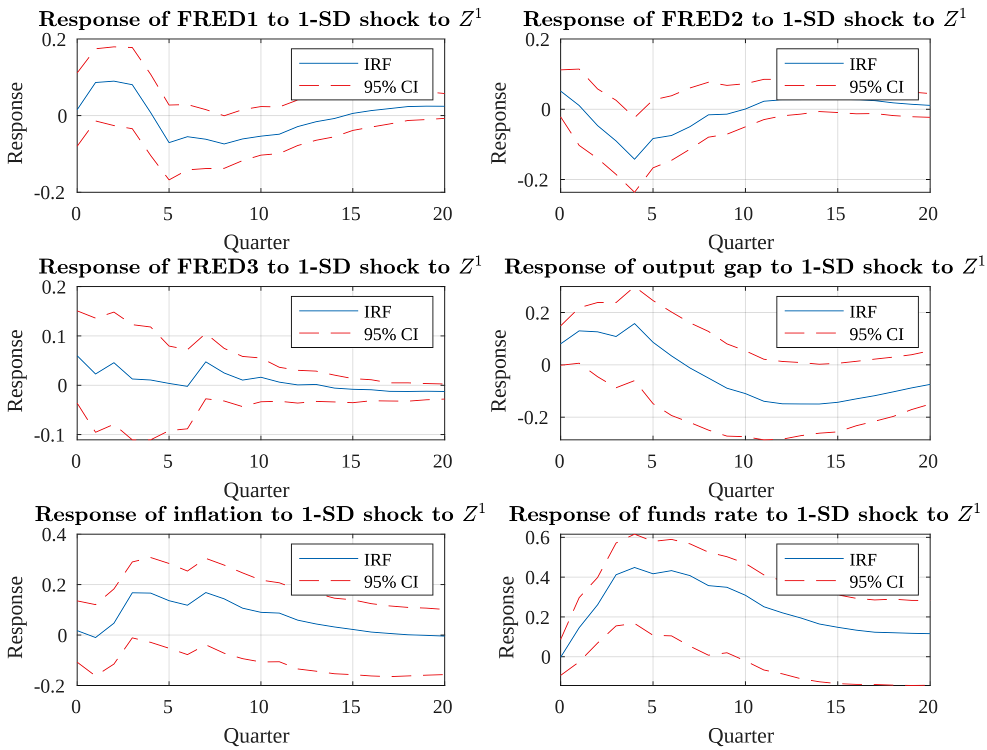

Figure C.6: FAVAR impulse responses to a one-standard deviation positive shock to the first dimension of mean market belief. FREDi denotes factor i of the FRED-MD panel. The factors' responses are measured in terms of standard deviations. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

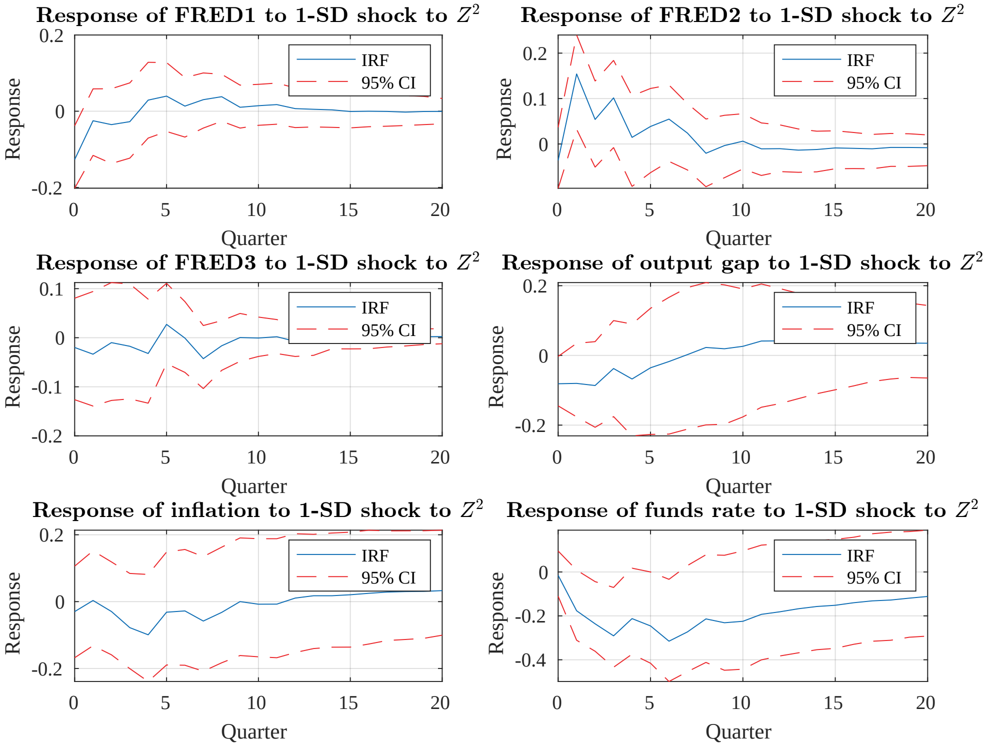

Figure C.7: FAVAR impulse responses to a one-standard deviation positive shock to the second dimension of mean market belief. FREDi denotes factor i of the FRED-MD panel. The factors' responses are measured in terms of standard deviations. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

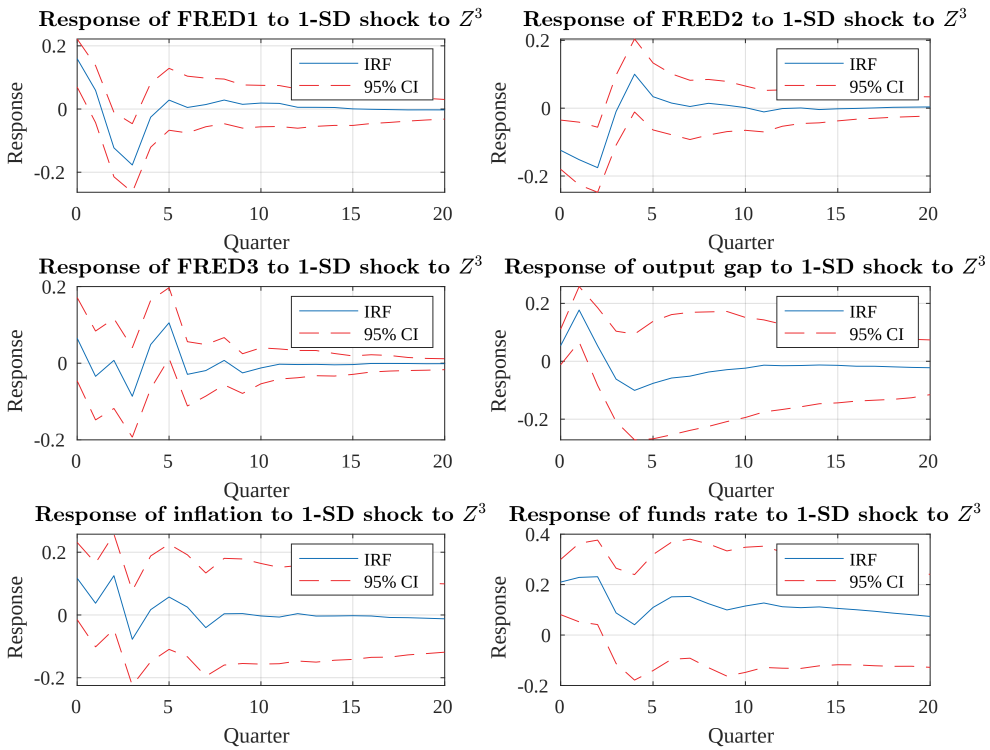

Figure C.8: FAVAR impulse responses to a one-standard deviation positive shock to the third dimension of mean market belief. FREDi denotes factor i of the FRED-MD panel. The factors' responses are measured in terms of standard deviations. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

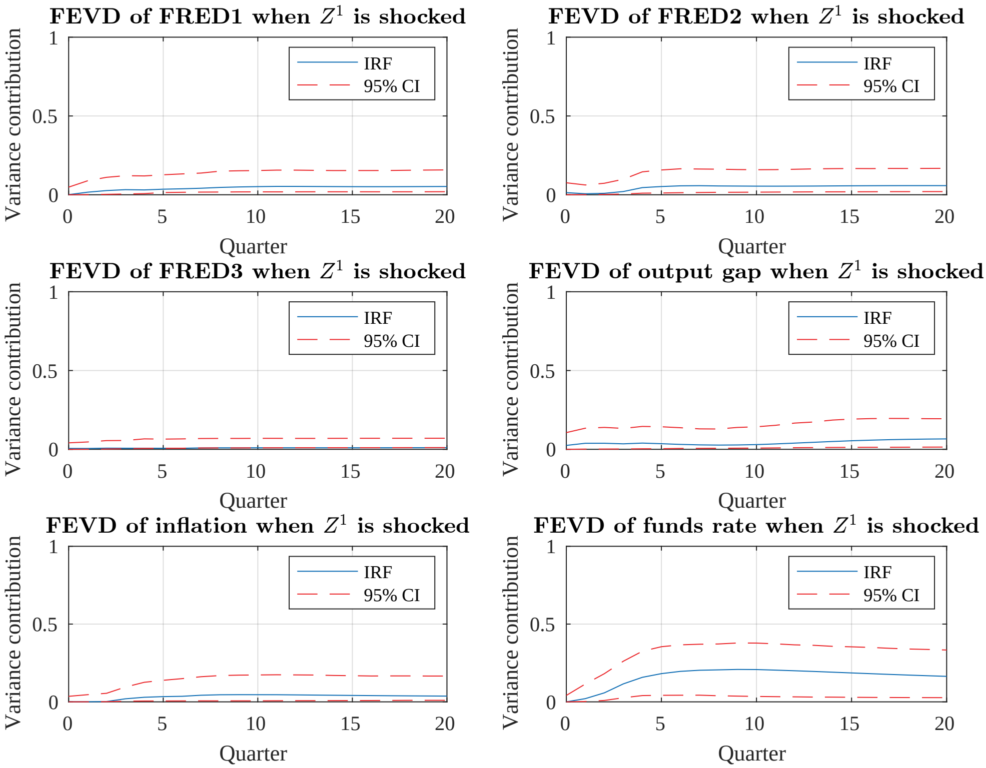

Figure C.9: FAVAR FEVD when the first dimension of mean market belief is shocked. FREDi denotes factor i of the FRED-MD panel. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

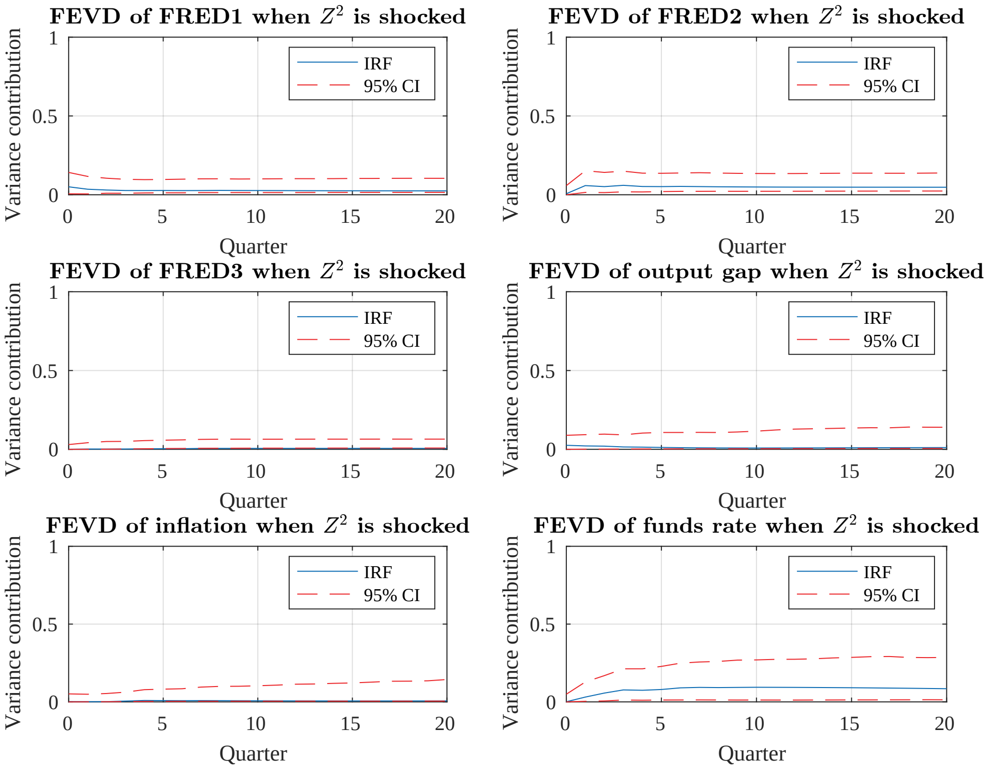

Figure C.10: FAVAR FEVD when the second dimension of mean market belief is shocked. FRED*i* denotes factor *i* of the FRED-MD panel. Dimension *j* of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

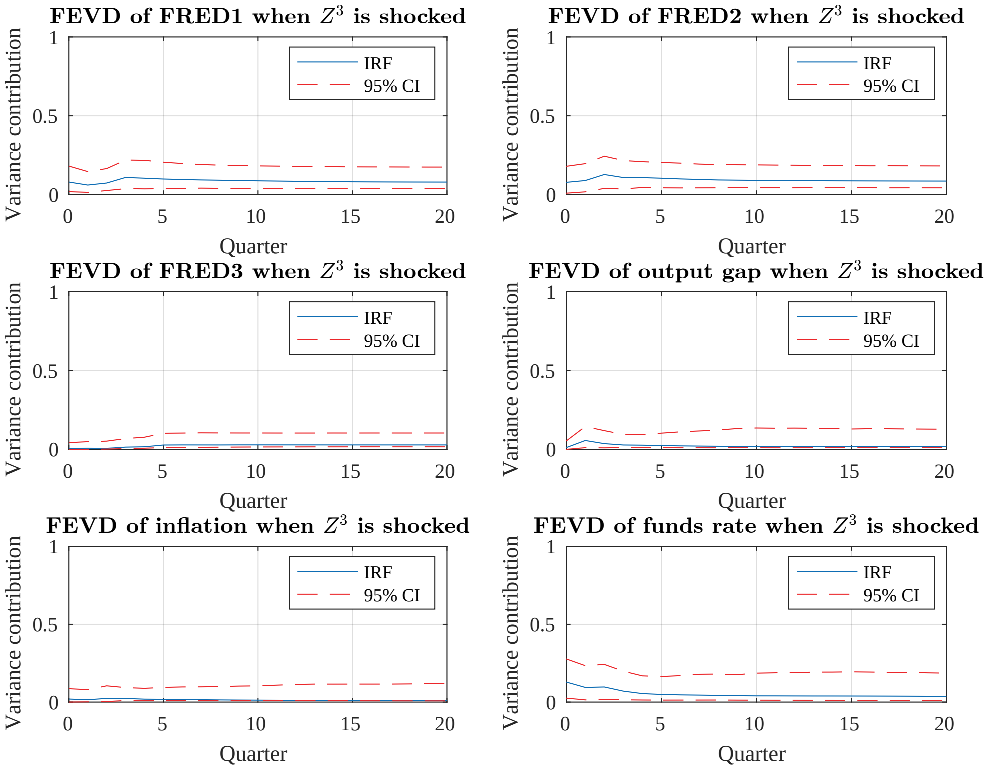

Figure C.11: FAVAR FEVD when the third dimension of mean market belief is shocked. FREDi denotes factor i of the FRED-MD panel. Dimension j of mean market belief is denoted by Z^j. The dashed lines indicate bootstrapped 95% confidence intervals based on 1,000 replications.

References

Aarbu, K. O. and Schroyen, F. (2009). Mapping risk aversion in Norway using hypothetical income gambles. Working Paper, Norwegian School of Economics and Business Administration, Bergen.

Ahmed, N. A. and Gokhale, D. V. (1989). Entropy expressions and their estimators for multivariate distributions. *IEEE Transactions on Information Theory*, 35(3):688–692.

Akerlof, G. A. and Shiller, R. J. (2009). *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*. Princeton University Press, Princeton, NJ.

Anderson, S. P., De Palma, A., and Thisse, J. F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.

Andrews, D. W. K. (1991). Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica*, 59(3):817–858.

Ang, A. (2014). *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press, Oxford.

Anufriev, M. and Hommes, C. (2012). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(4):35–64.

Anufriev, M. and Panchenko, V. (2009). Asset prices, traders' behavior and market design. *Journal of Economic Dynamics and Control*, 33(5):1073–1090.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Batra, R. N. and Ullah, A. (1974). Competitive Firm and the Theory of Input Demand under Price Uncertainty. *Journal of Political Economy*, 82(3):537–548.

Beaudry, P. and Willems, T. (2018). On the Macroeconomic Consequences of Over-Optimism. Working Paper 24685, National Bureau of Economic Research.

Benettin, G., Galgani, L., Giorgilli, A., and Strelcyn, J.-M. (1980). Lyapunov characteristic exponents for smooth dynamical systems and for hamiltonian systems; A method for computing all of them. Part 2: Numerical application. *Meccanica*, 15(1):21–30.

Berger, J. O. and Smith, L. A. (2019). On the Statistical Formalism of Uncertainty Quantification. *Annual Review of Statistics and Its Application*, 6(1):433–460.

Bernanke, B. S. (1983). Irreversibility, Uncertainty, and Cyclical Investment. *The Quarterly Journal of Economics*, 98(1):85–106.

Bernanke, B. S., Boivin, J., and Eliasz, P. (2005). Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach. *The Quarterly Journal of Economics*, 120(1):387–422.

Bhandari, A., Borovicka, J., and Ho, P. (2019). Survey data and subjective beliefs in business cycle models. Working Paper 19-14, Federal Reserve Bank of Richmond.

Biernacki, C., Celeux, G., and Govaert, G. (2003). Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate Gaussian mixture models. *Computational Statistics & Data Analysis*, 41(3):561–575.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2018). Really Uncertain Business Cycles. *Econometrica*, 86(3):1031–1065.

Boccara, N. (2010). *Modeling Complex Systems*. Graduate Texts in Physics. Springer, New York, NY.

Born, B., Breuer, S., and Elstner, S. (2018). Uncertainty and the Great Recession. *Oxford Bulletin of Economics and Statistics*, 80(5):951–971.

Branch, W. A. (2004). The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations. *The Economic Journal*, 114(497):592–621.

Brock, W. A. and Hommes, C. (1997). A rational route to randomness. *Econometrica*, 65(5):1059–1095.

Brock, W. A. and Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8–9):1235–1274.

Caggiano, G., Castelnuovo, E., and Nodari, G. (2018). Risk management-driven policy rate gap. *Economics Letters*, 171:235–238.

Caldara, D., Fuentes-Albero, C., Gilchrist, S., and Zakrajšek, E. (2016). The macroeconomic impact of financial and uncertainty shocks. *European Economic Review*, 88:185–207.

Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Chahrour, R. and Jurado, K. (2018). News or Noise? the Missing Link. *American Economic Review*, 108(7):1702–1736.

Chan, Y. L. and Kogan, L. (2002). Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices. *Journal of Political Economy*, 110(6):1255–1285.

Chiappori, P.-A., Gandhi, A., Salanié, B., and Salanié, F. (2009). Identifying Preferences under Risk from Discrete Choices. *American Economic Review*, 99(2):356–362.

Chiarella, C. and He, X.-Z. (2002). Heterogeneous beliefs, risk and learning in a simple asset pricing model. *Computational Economics*, 19(1):95–132.

Chiarella, C. and He, X.-Z. (2003). Heterogeneous beliefs, risk, and learning in a simple asset-pricing model with a market maker. *Macroeconomic Dynamics*, 7(4):503–536.

Chiarella, C., He, X.-Z., and Zwinkels, R. C. J. (2014). Heterogeneous expectations in asset pricing: Empirical evidence from the S&P500. *Journal of Economic Behavior & Organization*, 105:1–16.

Choi, S., Fisman, R., Gale, D., and Kariv, S. (2007). Consistency and heterogeneity of individual behavior under uncertainty. *American Economic Review*, 97(5):1921–1938.

Cohen, A. and Einav, L. (2007). Estimating risk preferences from deductible choice. *American Economic Review*, 97(3):745–788.

Coibion, O. (2012). Are the Effects of Monetary Policy Shocks Big or Small? *American Economic Journal: Macroeconomics*, 4(2):1–32.

Coibion, O. and Gorodnichenko, Y. (2011). Monetary Policy, Trend Inflation, and the Great Moderation: An Alternative Interpretation. *American Economic Review*, 101(1):341–370.

Conlisk, J. (1996). Why bounded rationality? *Journal of Economic Literature*, 34(2):669–700.

Croushore, D. and Stark, T. (2001). A real-time data set for macroeconomists. *Journal of Econometrics*, 105(1):111–130.

Croushore, D. D. (1993). Introducing: The survey of professional forecasters. *Business Review - Federal Reserve Bank of Philadelphia*, 6:3.

Dagum, E. B. (1980). The X-11-ARIMA seasonal adjustment method. *Statistics Canada*, (12–564E):677–796.

De Grauwe, P. (2012). *Lectures on Behavioral Macroeconomics*. Princeton University Press, Princeton, NJ.

De Grauwe, P., Foresti, P., and Ji, Y. (2019). Fiscal Policies in Booms and Busts. CEPR Discussion Paper 13740, Centre for Economic Policy Research, London.

De Grauwe, P. and Macchiarelli, C. (2015). Animal spirits and credit cycles. *Journal of Economic Dynamics and Control*, 59:95–117.

Doan, T., Litterman, R., and Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, 3(1):1–100.

Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2012). The intergenerational transmission of risk and trust attitudes. *The Review of Economic Studies*, 79(2):645–677.

Dovern, J. (2015). A multivariate analysis of forecast disagreement: Confronting models of disagreement with survey data. *European Economic Review*, 80:16–35.

Durbin, J. and Koopman, S. J. (2012). Filtering, smoothing and forecasting. In *Time Series Analysis by State Space Methods: Second Edition*. Oxford University Press.

Elliott, G. and Timmermann, A. (2008). Economic Forecasting. *Journal of Economic Literature*, 46(1):3–56.

Enders, Z., Hünnekes, F., and Müller, G. J. (2019). Firm expectations and economic activity. CESifo Working Paper 7623.

Evans, C., Fisher, J., Gourio, F., and Krane, S. (2016). Risk Management for Monetary Policy Near the Zero Lower Bound. *Brookings Papers on Economic Activity*, 2015(1):141–219.

Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.

Falck, E., Hoffmann, M., and Hürtgen, P. (2019). Disagreement about inflation expectations and monetary policy transmission. *Journal of Monetary Economics*.

Falconer, K. (2004). *Fractal Geometry: Mathematical Foundations and Applications*. John Wiley & Sons.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., and Sunde, U. (2018). Global evidence on economic preferences. *The Quarterly Journal of Economics*, 133(4):1645–1692.

Fève, P. and Guay, A. (2019). Sentiments in SVARs. *The Economic Journal*, 129(618):877–896.

Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C., and Chen, B.-C. (1998). New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program. *Journal of Business & Economic Statistics*, 16(2):127–152.

Gaunersdorfer, A., Hommes, C. H., and Wagener, F. O. O. (2008). Bifurcation routes to volatility clustering under evolutionary learning. *Journal of Economic Behavior & Organization*, 67(1):27–47.

Gnabo, J.-Y. and Moccero, D. N. (2015). Risk management, nonlinearity and aggressiveness in monetary policy: The case of the US Fed. *Journal of Banking & Finance*, 55:281–294.

Greenspan, A. (2004). Risk and Uncertainty in Monetary Policy. *American Economic Review*, 94(2):33–40.

Gros, C. (2015). *Complex and Adaptive Dynamical Systems*. Springer International Publishing, Cham.

Herbst, E. and Winkler, F. (2020). The Factor Structure of Disagreement. Technical report.

Hommes, C. (2001). Financial markets as nonlinear adaptive evolutionary systems. *Quantitative Finance*, 1(1):149–167.

Hommes, C. (2013). *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press, Cambridge.

Hommes, C. (2021). Behavioral and Experimental Macroeconomics and Policy Analysis: A Complex Systems Approach. *Journal of Economic Literature*, 59(1):149–219.

Hommes, C., Massaro, D., and Weber, M. (2019). Monetary policy under behavioral expectations: Theory and experiment. *European Economic Review*, 118:193–212.

Hommes, C., Sonnemans, J., Tuinstra, J., and van de Velden, H. (2005). Coordination of expectations in asset pricing experiments. *The Review of Financial Studies*, 18(3):955–980.

Hommes, C., Sonnemans, J., Tuinstra, J., and van de Velden, H. (2008). Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization*, 67(1):116–133.

Hyndman, R. J. and Khandakar, Y. (2008). Automatic Time Series Forecasting: The forecast Package for R. *Journal of Statistical Software*, 27(1):1–22.

Jaeger, D. A., Dohmen, T., Falk, A., Huffman, D., Sunde, U., and Bonin, H. (2010). Direct evidence on risk attitudes and migration. *The Review of Economics and Statistics*, 92(3):684–689.

Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216.

Kalai, E. and Lehrer, E. (1993). Rational learning leads to Nash equilibrium. *Econometrica*, 61(5):1019–1045.

Kimball, M. S., Sahm, C. R., and Shapiro, M. D. (2008). Imputing risk tolerance from survey responses. *Journal of the American Statistical Association*, 103(483):1028–1038.

Kirman, A. P. (1992). Whom or what does the representative individual represent? *The Journal of Economic Perspectives*, 6(2):117–136.

Kocherlakota, N. (2010). Modern macroeconomic models as tools for economic policy. *The Region*, 24(1):5–21.

Koop, G. and Korobilis, D. (2013). Large time-varying parameter VARs. *Journal of Econometrics*, 177(2):185–198.

Kording, K. P. (2014). Bayesian statistics: Relevant for the brain? *Current Opinion in Neurobiology*, 25:130–133.

Kurz, M. (1994). On the structure and diversity of rational beliefs. *Economic Theory*, 4(6):877–900.

Kurz, M., Jin, H., and Motolese, M. (2005). The role of expectations in economic fluctuations and the efficacy of monetary policy. *Journal of Economic Dynamics and Control*, 29(11):2017–2065.

Kurz, M. and Motolese, M. (2011). Diverse beliefs and time variability of risk premia. *Economic Theory*, 47(2-3):293–335.

Kurz, M., Motolese, M., Piccillo, G., and Wu, H. (2018). Monetary policy with diverse private expectations. Working Paper.

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.

Leeper, E. M. and Zha, T. (2003). Modest policy interventions. *Journal of Monetary Economics*, 50(8):1673–1700.

Lei, V., Noussair, C. N., and Plott, C. R. (2001). Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality. *Econometrica*, pages 831–859.

Leland, H. E. (1968). Saving and Uncertainty: The Precautionary Demand for Saving. *The Quarterly Journal of Economics*, 82(3):465–473.

Liebermann, J. (2014). Real-Time Nowcasting of GDP: A Factor Model vs. Professional Forecasters. *Oxford Bulletin of Economics and Statistics*, 76(6):783–811.

Litterman, R. B. (1986). Forecasting with Bayesian Vector Autoregressions: Five Years of Experience. *Journal of Business & Economic Statistics*, 4(1):25–38.

Lucas, R. (1976). Econometric policy evaluation: A critique. In Brunner, K. and Meltzer, A. H., editors, *The Phillips Curve and Labor Markets*, volume 1 of *Carnegie-Rochester Conference Series on Public Policy*, pages 19–46. American Elsevier, New York.

Ludvigson, S. C., Ma, S., and Ng, S. (2015). Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response? Working Paper 21803, National Bureau of Economic Research.

Mankiw, N. G., Reis, R., and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER Macroeconomics Annual*, 18:209–248.

Marcellino, M., Stock, J. H., and Watson, M. W. (2006). A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of Econometrics*, 135(1):499–526.

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, 7(1):77–91.

McCracken, M. W. and Ng, S. (2016). FRED-MD: A Monthly Database for Macroeconomic Research. *Journal of Business & Economic Statistics*, 34(4):574–589.

Milani, F. (2011). Expectation shocks and learning as drivers of the business cycle. *The Economic Journal*, 121(552):379–401.

Motolese, M. and Nielsen, C. K. (2007). Rational Beliefs Theory: A review. *Rivista Internazionale di Scienze Sociali*, page 35.

Mumtaz, H. and Theodoridis, K. (2019). Dynamic effects of monetary policy shocks on macroeconomic volatility. *Journal of Monetary Economics*.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29(3):315–335.

Nelson, D. B. (1992). Filtering and forecasting with misspecified ARCH models I: Getting the right variance with the wrong model. *Journal of Econometrics*, 52(1):61–90.

Orlik, A. and Veldkamp, L. (2014). Understanding Uncertainty Shocks and the Role of Black Swans. Working Paper 20445, National Bureau of Economic Research.

Orphanides, A. (2001). Monetary Policy Rules Based on Real-Time Data. *American Economic Review*, 91(4):964–985.

Panetta, F. (2020). The price of uncertainty and uncertainty about prices: Monetary policy in the post-COVID-19 economy.

Paravisini, D., Rappoport, V., and Ravina, E. (2016). Risk aversion and wealth: Evidence from person-to-person lending portfolios. *Management Science*, 63(2):279–297.

Park, B.-J. (2014). Time-varying, heterogeneous risk aversion and dynamics of asset prices among boundedly rational agents. *Journal of Banking & Finance*, 43:150–159.

Pfajfar, D. (2013). Formation of rationally heterogeneous expectations. *Journal of Economic Dynamics and Control*, 37(8):1434–1452.

Pigou, A. (1927). *Industrial Fluctuations*. MacMillan, London.

Rich, R. and Tracy, J. (2010). The Relationships among Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts. *The Review of Economics and Statistics*, 92(1):200–207.

Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review*, 94(4):1055–1084.

Roth, C. and Wohlfart, J. (2018). How do expectations about the macroeconomy affect personal expectations and behavior? CESifo Working Paper 7154.

Schmitt, N. (2020). Heterogeneous expectations and asset price dynamics. *Macroeconomic Dynamics*, page forthcoming.

Schumacher, C. and Breitung, J. (2008). Real-time forecasting of German GDP based on a large factor model with monthly and quarterly data. *International Journal of Forecasting*, 24(3):386–398.

Simon, H. A. (1957). *Models of Man; Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. Wiley, Oxford.

Smets, F. and Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606.

Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica: Journal of the Econometric Society*, pages 1119–1151.

Stark, T. (2010). Realistic Evaluation of Real-Time Forecasts in the Survey of Professional Forecasters. Technical report.

Stark, T. (2013). SPF panelists' forecasting methods: A note on the aggregate results of a November 2009 special survey. Technical report.

Stock, J. H. and Watson, M. W. (2002). Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business & Economic Statistics*, 20(2):147–162.

Tinbergen, J. (1937). *An Econometric Approach to Business Cycle Problems*. Hermann & Compagnie, Paris.

Tinbergen, J. (1939). *Statistical Testing of Business-Cycle Theories*. League of Nations, Geneva.

Viscusi, W. K. (1985). Are Individuals Bayesian Decision Makers? *American Economic Review*, 75(2):381–385.

von Gaudecker, H.-M., van Soest, A., and Wengström, E. (2011). Heterogeneity in risky choice behavior in a broad population. *American Economic Review*, 101(2):664–694.

Wieland, J. F. and Yang, M.-J. (2020). Financial Dampening. *Journal of Money, Credit and Banking*, 52(1):79–113.

Wieland, V. and Wolters, M. H. (2011). The diversity of forecasts from macroeconomic models of the US economy. *Economic Theory*, 47(2-3):247–292.

Winkler, F. (2019). The role of learning for asset prices and business cycles. *Journal of Monetary Economics*.

Xiouros, C. and Zapatero, F. (2010). The representative agent of an economy with external habit formation and heterogeneous risk aversion. *The Review of Financial Studies*, 23(8):3017–3047.

Yu, A. J. (2007). Adaptive Behavior: Humans Act as Bayesian Learners. *Current Biology*, 17(22):R977–R980.

Nederlandse samenvatting

In dit proefschrift onderzoek ik de rollen van onzekerheid en verwachtingen in de economie.

Er zijn verschillende bronnen van onzekerheid. Op een fundamenteel niveau weten we niet met zekerheid wat de huidige staat van de economie is. Veel macro-economische indicatoren, zoals economische groei en inflatie, worden herzien nadat ze gepubliceerd zijn. Deze herzieningen kunnen substantieel zijn, zeker in turbulente perioden als de COVID-19 pandemie of de financiële crisis van 2008. Zelfs met volledige kennis over de huidige staat van de economie zouden we geen perfecte voorspellingen kunnen doen, omdat we de bewegingswetten van de economie niet kennen.

Onzekerheid dwingt ons om verwachtingen te vormen over de huidige en toekomstige staat van de economie. Neem een scenario waarin iemand moet beslissen om nu een huis te kopen, of te wachten. Cruciale factoren in haar besluit zijn de huidige staat van de huizenmarkt en haar toekomstige inkomen. Omdat beide factoren een mate van onzekerheid met zich meebrengen, moet ze er verwachtingen over vormen om een beslissing te kunnen nemen. Als ze bijvoorbeeld verwacht dat er een economische crisis aankomt, wat haar kans op werkloosheid vergroot, zou ze het kopen van een huis misschien willen uitstellen.

Waar onzekerheid leidt tot het vormen van verwachtingen, brengen verwachtingen zelf ook een mate van onzekerheid met zich mee. Ten eerste bestaan goede voorspellingen uit een verzameling mogelijke uitkomsten en een kansverdeling over die uitkomsten. Ten tweede kennen we niet de verwachtingen van alle andere economische actoren, wat zorgt voor meer onzekerheid over het gehele economische systeem.

Hoofdstuk 2

Er is geen consensus over de directe economische impact van onzekerheid, maar de verwachtingen van economische actoren over deze impact kunnen de economie ook beïnvloeden. Zeker wanneer onzekerheid de beslissingen van beleidsmakers beïnvloedt, kan de impact groot zijn. In het tweede hoofdstuk van dit proefschrift, gebaseerd op gezamenlijk werk met Giulia Piccillo, onderzoek ik of macro-economische onzekerheid invloed heeft op het monetaire beleid in de VS. Acht keer per jaar komt de Federal Open Market Committee (FOMC) bijeen om het monetaire beleid te beoordelen op basis van de huidige economische situatie. Ik veronderstel dat de FOMC-leden een standaard macro-economisch model gebruiken om een oordeel te vormen over de economische situatie. Dit model beschrijft de relaties tussen economische groei, inflatie, en de rente. Ik veronderstel ook dat de beleidsmakers Bayesiaanse statistiek gebruiken om hun ideeën te updaten: als er nieuwe data beschikbaar is, passen ze hun overtuigingen over de parameters van het model aan. Deze overtuigingen worden weergegeven door een kansverdeling. Uit de spreiding leid ik een maat van macro-economische onzekerheid af. Bij het opstellen van deze maat van onzekerheid, gebruik ik de macro-economische data zoals beschikbaar op het moment dat de betreffende FOMC-bijeenkomsten werden gehouden.

Ik schat de impact van deze real-time Bayesiaanse maat van macro-economische onzekerheid op rentebesluiten van de FOMC. Ik houd hierbij rekening met de impact van de economische voorspellingen die voor iedere bijeenkomst voorbereid worden door de Federal Reserve Board. Ik neem ook een maat van financiële onzekerheid mee om de rollen van financiële en macro-economische onzekerheid te kunnen onderscheiden. Ik concludeer dat beleidsmakers besluiten tot een significant lagere rente in tijden van grotere macro-economische onzekerheid.

Hoofdstuk 3

Mensen hebben verschillende houdingen ten aanzien van onzekerheid. Sommigen houden van risico, terwijl anderen juist risico-avers zijn. In het derde hoofdstuk van dit proefschrift, gebaseerd op gezamenlijk werk met Giulia Piccillo, onderzoek ik of deze houdingen invloed hebben op verwachtingen. Ik gebruik een bestaand verwachtingsmodel waarin economische actoren kiezen tussen simpele voorspellingsregels om verwachtingen te vormen. Zij baseren deze keuze op hun prestaties in het verleden. In deze keuze introduceer ik een rol voor risico-aversie: actoren baseren hun keuze zowel op de prestaties van de regels als de variabiliteit van deze prestaties. Omdat ze verschillende houdingen hebben ten aanzien van risico, kiezen ze verschillende regels. Dit leidt tot heterogene verwachtingen. Om het model empirisch te valideren, trek ik de risico-aversie van de actoren uit een verdeling gebaseerd op enquête data.

Ik combineer dit verwachtingsmodel met een simpele financiële markt met twee investeringsmogelijkheden, de één risicovol en de ander risicoloos. Actoren in het model vormen verwachtingen over het rendement van beide investeringen om te bepalen hoe ze hun geld verdelen. Simulaties laten zien dat de wisselwerking tussen verschillende verwachtingen kan zorgen voor onvoorspelbare marktbubbels en -crashes. Het toevoegen van kleine stochastische prijsschokken leidt tot grotere bubbels en kan markten destabiliseren.

Hoofdstuk 4

Eerdere studies naar de rol van 'animal spirits' (sentiment, emoties) in de economie trekken geen eenduidige conclusie. In hoofdstuk 4 stel ik hiervoor een verklaring voor: deze studies meten verschillende dimensies van sentiment die verschillende macro-economische invloeden hebben. Om mijn hypothese te toetsen, gebruik ik de theorie van Rational Beliefs. Deze theorie definieert animal spirits als tijdelijke afwijkingen van verwachtingen over economische grootheden van de verwachtingen die worden geïmpliceerd door de data die op dat moment beschikbaar is (de zogenaamde empirische verwachtingen). Deze definitie suggereert ook dat de animal spirits gemeten kunnen worden als het verschil tussen geobserveerde voorspellingen en de empirische voorspellingen. Enquête data over de voorspellingen van professionele voorspellers dekt de eerste helft van de vergelijking. Ik benader de empirische voorspellingen door een groot panel van real-time data te verzamelen dat alle relevante aspecten van de economie dekt, en hierop een statistisch model toe te passen om voorspellingen te produceren.

Ik gebruik de volledige 50 jaar aan beschikbare enquête data. Iedere enquête beschrijft de verwachtingen van ongeveer 40 voorspellers voor diverse economische variabelen (bijvoorbeeld productie, prijzen, rente, huisvesting), en over meerdere termijnen. Wanneer ik de corresponderende empirische voorspellingen daarvan aftrek, ontstaat een panel van sentiment voor alle enquêtes en alle voorspellers, over alle combinaties van variabelen en termijnen. Ik gebruik vervolgens een statistische procedure om drie dimensies te identificeren die samen ongeveer 50% van de animal spirits van de voorspellers verklaren. Ik concludeer dat sentiment inderdaad multidimensionaal is: de eerste dimensie verklaart slechts een vijfde van het sentiment van de voorspellers. Verder concludeer ik dat iedere dimensie een eigen macro-economische impact heeft, wat mijn hypothese ondersteunt.

Om de empirische analyse te ondersteunen, ontwikkel ik een theoretisch raamwerk gebaseerd op de theorie van Rational Beliefs. Dit stelt mij in staat om toetsbare implicaties van de theorie af te leiden. Daardoor dient dit hoofdstuk ook als een toets van de empirische relevantie van de theorie. Mijn resultaten ondersteunen sommige implicaties van de theorie, maar ondermijnen andere.

Curriculum vitae

Thomas Gomez was born on 12 November 1991 in Alkmaar, the Netherlands. He received cum laude BSc degrees in Mathematics and Physics from Utrecht University in 2014. He completed a Research Master in Multidisciplinary Economics at Utrecht University School of Economics, for which he received a cum laude MSc degree in 2016. In September 2016, Thomas started working as a PhD candidate at Utrecht University. He is currently working as a Quantitative Analyst at an algorithmic trading firm.

U.S.E. Dissertation Series

USE 001 **Bastian Westbrock** (2010): *Inter-firm networks: economic and sociological perspectives.*

USE 002 **Yi Zhang** (2011): *Institutions and International Investments: Evidence from China and Other Emerging Markets.*

USE 003 **Ryan van Lamoen** (2011): *The Relationship between Competition and Innovation Measuring Innovation and Causality.*

USE 004 **Martijn Dröes** (2011): *House Price Uncertainty in the Dutch Owner-Occupied Housing Market.*

USE 005 **Thomas van Huizen** (2012): *Behavioural Assumptions in Labour Economics: Analysing Social Security Reforms and Labour Market Transitions.*

USE 006 **Martijn Boermans** (2012): *International Entrepreneurship and Enterprise Development.*

USE 007 **Joras Ferwerda** (2012): *The Multidisciplinary Economics of Money Laundering.*

USE 008 **Federico D'Onofrio** (2013): *Observing the country: a history of Italian agricultural economics, 1900-1930.*

USE 009 **Saraï Sapulete** (2013): *Works Council Effectiveness: Determinants and Outcomes.*

USE 010 **Britta Hoyer** (2013): *Network Formation under the Threat of Disruption.*

USE 011 **Coen Rigtering** (2013): *Entrepreneurial Orientation: Multilevel Analysis and Consequences.*

USE 012 **Beate Cesinger** (2013): *Context and Complexity of International Entrepreneurship as a Field of Research.*

USE 013 **Jan de Dreu** (2013): *Empirical essays on the governance of financial institutions.*

USE 014 **Lu Zhang** (2013): *Industrial Specialization: Determinants, Processes and Consequences.*

USE 015 **Matthias Filser** (2013): *Strategic Issues in Entrepreneurship and Family Business Research.*

USE 016 **Mikko Pohjola** (2013): *A Compilation of Studies on Innovation in Firms: Capabilities, Strategies, and Performance.*

USE 017 **Han-Hsin Chang** (2013): *Heterogeneity in Development.*

USE 018 **Suzanne Heijnen** (2014): *Analyses of sickness absence.*

USE 019 **Mark Kattenberg** (2014): *The Economics of Social Housing: Implications for Welfare, Consumption, and Labor Market Composition.*

USE 020 **Daniel Possenriede** (2014): *The Economics of Temporal and Locational Flexibility of Work.*

USE 021 **Dirk Gerritsen** (2014): *The Relevance of Security Analyst Opinions for Investment Decisions.*

USE 022 **Shiwei Hu** (2014): *Development in China and Africa.*

USE 023 **Saara Tamminen** (2014): *Heterogeneous Firms, Mark-Ups, and Income Inequality.*

USE 024 **Marcel van den Berg** (2014): *Does Internationalization Foster Firm Performance?*

USE 025 **Emre Akgündüz** (2014): *Analyzing maternal employment and child care quality.*

USE 026 **Jasper Lukkezen** (2014): *From Debt Crisis to Sovereign Risk.*

USE 027 **Vesile Kutlu** (2015): *Essays on Subjective Survival Probabilities, Consumption, and Retirement Decisions.*

USE 028 **Brigitte Crooijmans** (2015): *Leiden fusies tot efficiëntere woningcorporaties? Een exploratieve studie naar schaalvoordelen in de sociale huisvesting.*

USE 029 **Andrej Svorenčík** (2015): *The Experimental Turn in Economics: a History of Experimental Economics.*

USE 030 **Secil Danakol** (2015): *Foreign Direct Investment, Foreign Aid and Domestic Entrepreneurship.*

USE 031 **Ioana Deleanu** (2015): *Anti-Money Laundering Efforts: Failures, Fixes and the Future.*

USE 032 **Jaap Oude Mulders** (2016): *Organizations, managers, and the employment of older workers after retirement.*

USE 033 **Malka de Castro Campos** (2016): *Private Consumption-Savings Behavior and Macroeconomic Imbalances.*

USE 034 **Tahereh Rezaei Khavas** (2016): *Fairness concerns and cooperation in context.*

USE 035 **Joyce Delnoy** (2016): *Auctions with Competing Sellers and Behavioral Bidders.*

USE 036 **Krista Bruns** (2017): *Emergence and Diffusion of Institutions and their Effect on Economic Growth.*

USE 037 **Daan van der Linde** (2017): *Democracies under Rising Inequality: New Tests of the Redistributive Thesis.*

USE 038 **Swantje Falcke** (2017): *On the move: Analyzing immigration determinants and immigrant outcomes.*

USE 039 **Joep Steegmans** (2017): *House Prices and Household Mobility in The Netherlands: Empirical Analyses of Financial Characteristics of the Household.*

USE 040 **Najmeh Rezaei Khavas** (2017): *Essays in Information Economics.*

USE 041 **Maryam Imanpour** (2017): *The Role of Social Networks for Combating Money Laundering.*

USE 042 **Ye Li** (2018): *Hydrogen Infrastructure Decisions through a Real Option Lens.*

USE 043 **Li Lin** (2018): *Leadership across cultural contexts.*

USE 044 **Werner Liebregts** (2018): *Hidden entrepreneurship: Multilevel analyses of the determinants and consequences of entrepreneurial employee activity.*

USE 045 **Ian Koetsier** (2018): *Government debt: The economic consequences of natural disasters and pension funds' herding.*

USE 046 **Jordy Meekes** (2019): *Local Labour Markets, Job Displacement And Agglomeration Economies.*

USE 047 **Timur Pasch** (2019): *Essays On The Design Of The Management Accounting System: Determinants, Components And Effects.*

USE 048 **Jeroen Content** (2019): *The role of relatedness and entrepreneurship in regional economic development.*

USE 049 **Franziska Heinicke** (2019): *Essays on self-image and preferences for honesty.*

USE 050 **Rebean Al-silefanee** (2019): *Entrepreneurship and Private Sector Development: The Case of Kurdistan Region of Iraq.*

USE 051 **Markus Meinzer** (2019): *Countering cross-border tax evasion and avoidance: An assessment of OECD policy design from 2008 to 2018.*

USE 052 **Zornitza Kambourova** (2019): *Women's Adverse Health Events and Labor Market Participation.*

USE 053 **Tim van der Valk** (2019): *Household finance in France and the Netherlands 1960-2000: An evolutionary approach.*

USE 054 **Milena Dinkova** (2019): *Brace yourselves, Pension is coming: Consumption, financial literacy and tailored pension communication.*

USE 055 **Lisa Dumhs** (2019): *Finding the right job: School-to-work transitions of vocational students in the Netherlands.*

USE 056 **Dea Tusha** (2020): *FDI spillovers in developing countries: channels, conditions, challenges.*

USE 057 **Jingyang Liu** (2020): *Money and credit dynamics in the euro area.*

USE 058 **An Duong** (2020): *Financial integration, trade, and productivity.*

USE 059 **Katharina Weddige-Haaf** (2021): *Real and Financial Asymmetries in the Euro Area.*

USE 060 **Peter Gerbrands** (2021): *Tax Dynamics and Money Laundering. Simulating Policy Reforms in a Complex System.*

USE 061 **Timo Verlaat** (2022): *Carrot and Stick: Experiments With Social Welfare Policies.*

USE 062 **Lucia Rossel Flores** (2022): *A multidisciplinary analysis of tax reform: from politics to human behavior.*

USE 063 **Wanxiang Cai** (2022): *Social capital and crowdfunding. A multilevel perspective.*

USE 064 **Vincent Schippers** (2022): *The local economic impacts of natural disasters. A view from outer space.*

USE 065 **Peter D van der Meer** (2022): *Job insecurity and mental health. Essays on* *the effect of job insecurity on mental health and the moderating effect of religiousness and psychological factors.*

In this dissertation, I examine the roles of uncertainty and beliefs in the economy. In chapter 2, I investigate whether macroeconomic uncertainty affects monetary-policy decisions in the US. I create a measure of macroeconomic uncertainty as felt by the policymakers and analyze its impact on their interest rate decisions. I find that policymakers set a significantly lower interest rate in times of higher macroeconomic uncertainty. In the third chapter, I investigate the role of risk in belief formation. I use a model in which economic agents choose between simple forecasting rules to form beliefs. They base their choice on the rules' historical performance and the variability (risk) of that performance. Agents have different risk preferences, and therefore choose different rules. I analyze the implications of this belief-formation model in a stylized financial market and show that the interactions between different beliefs can drive unpredictable booms and busts. In chapter 4, I study the role of sentiment in the macroeconomy. I use 50 years of survey data on the expectations of professional economic forecasters. I measure sentiment as the difference between observed forecasts and forecasts implied by the data available at the time. I conclude that sentiment is multidimensional, where every dimension has a distinct macroeconomic impact. My results furthermore indicate that the survey forecasts are not always rational.

ISBN 978-94-91870-52-1