

Discussion Paper Series – CRC TR 224

Discussion Paper No. 287
Project C 01

Rational vs. Irrational Beliefs in a Complex World

Gregor Boehl ¹
Cars Hommes ²

April 2021

¹ (corresponding author), gboehl@uni-bonn.de, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany

² ch.hommes@uva.nl, University of Amsterdam, Roetersstraat 11, 1001 NJ Amsterdam, The Netherlands and
Bank of Canada

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

Rational vs. Irrational Beliefs in a Complex World

Gregor Boehl^a, Cars Hommes^b,

^a*University of Bonn*

^b*University of Amsterdam and Bank of Canada*

April 14, 2021

Abstract

Can boundedly rational agents survive competition with fully rational agents? We develop a highly nonlinear heterogeneous agents model with rational forward looking versus boundedly rational backward looking agents and evolving market shares depending on their relative performance. Our novel numerical solution method detects equilibrium paths characterized by complex bubble and crash dynamics. Boundedly rational trend-extrapolators amplify small deviations from fundamentals, while rational agents anticipate market crashes after large bubbles and drive prices back close to fundamental value. Overall rational and non-rational beliefs co-evolve over time, with time-varying impact, and their interaction produces complex endogenous bubble and crashes, without any exogenous shocks.

Keywords: Heterogeneous agents, trend-extrapolation, bubbles, numerical solution method

JEL: C63, E03, E32, E44, E51

*Gregor Boehl (corresponding author): gboehl@uni-bonn.de, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany. Cars Hommes: ch.hommes@uva.nl, University of Amsterdam, Roetersstraat 11, 1001 NJ Amsterdam, The Netherlands. We are grateful to Herbert Dawid, Cees Diks, Reiner Franke, Thomas Hintermaier, Ken Judd, Isabelle Salle, Florian Wagener and participants of the 1st Behavioral Macroeconomics Workshop (Bamberg) for discussions and helpful comments on the content of this paper. Gregor Boehl gratefully acknowledges funding from the Bielefeld Graduate School in Economics and Management and the European Commission's EDEEM-Programme, financial support from the Alfred P. Sloan Foundation under the grant agreement G-2016-7176 for the Macroeconomic Model Comparison Initiative (MMCI) at the IMFS Frankfurt, and funding from the DFG under CRC-TR 224 (project C01) and under project number 441540692.

1 Introduction

“Given the uncertainty of the real world, the many actual and virtual traders will have many, perhaps equally many, forecast [...] If any group of traders was consistently better than average in forecasting stock prices, they would accumulate wealth and give their forecasts greater and greater weight. In this process, they would bring the present price closer to the true value.”

— Cootner, 1964

“. . . the rational expectations hypothesis is consistently rejected in individual forecast data. Individual forecast errors are systematically predictable from forecast revisions. [...] overreaction to information is the norm in individual forecast data, meaning that upward revisions are associated with realizations below forecasts.”

— Bordalo et al., 2020a

Do fully rational agents drive out irrational agents and bring asset prices back to fundamental value? This question has a long tradition in economics. Friedman (1953) has been an early advocate of the view that rational agents would outperform non-rational agents, who would lose money and be driven out of the market. The above quote from Cootner (1964) is supportive of this *rational view* and argues that not only will agents with more sophisticated forecasting rules drive out other rules, but in addition they will prevent “bubbles” and enforce prices to converge to fundamentals.

More recently the rational view has been challenged by behavioral approaches emphasizing the importance and empirical relevance of “irrational” extrapolative expectations in generating asset market bubbles. For example, Barberis et al. (2018) argue that bubbles may be triggered by shocks to fundamentals and amplified by extrapolative expectations. Bordalo et al. (2020b) show that expectations surveys have high explanatory power on stock price dynamics. Greenwood and Shleifer (2014) show that empirical expected returns are highly correlated with past returns and stock prices, while rational expectations would suggest the opposite. The second quote above is taken from Bordalo et al. (2020a), who document systematic overreaction of individual forecasters to news. The work of Hommes et al. (2005); Hommes (2021) shows the prevalence of coordination on extrapolative expectations in laboratory experiments with human subjects.

This paper is an attempt to reconcile the rational and the behavioral approach. Our main contribution lies in providing and solving a heterogeneous agents model in which “irrational agents” with extrapolative expectations and fully rational agents co-exist and their fractions co-evolve over time upon their relative success. The interactions between rational agents and extrapolative expectations create endogenous fluctuations characterized by irregular bubble and crash dynamics in line with the

empirical evidence above. Such heterogeneous agents models are highly nonlinear and therefore hard to solve even numerically. We provide a novel numerical method to compute rational expectations equilibria –i.e. solution paths along which rational agents have perfect foresight– characterized by bubble and crash patterns. Extrapolative expectations amplify small deviations from the fundamental into growing asset bubbles, while rational agents stabilize these bubbles when they become too large. These interactions between simple and sophisticated traders lead to complex bubble and crash patterns in asset prices.

A novel contribution of our paper is to develop a numerical method to solve for the global dynamics of a nonlinear system with endogenous cycles and chaotic equilibrium dynamics. Why is this necessary? As in Grandmont (1985, 1998) we approach the question of whether or not agents with non-rational expectations survive in a temporary equilibrium framework. The framework with competing heterogeneous beliefs is adapted from Brock and Hommes (1997, 1998). A temporary equilibrium map with heterogeneous expectations, including both forward looking and backward looking agents is of the form

$$x_t = F(E_t x_{t+1}, X_{t-1}), \tag{1}$$

where F is the law of motion of the equilibrium price x_t of a financial asset, and $X_{t-1} = \{x_{t-1}, x_{t-2}, \dots\}$ is the history of x_t .¹ Our particular interest lies in models where F is highly nonlinear and may exhibit complex equilibrium solutions, such as boom and bust cycles or chaotic and unpredictable bubble and crash dynamics. The law of motion (1) may, for example, represent a heterogeneous expectations model, where one type has simple backward looking expectations (e.g., some function of x_{t-1}) and another type has forward looking rational expectations (perfect foresight, hence the term $E_t x_{t+1}$). The dynamics of (1) is then only implicitly defined as the current state x_t depends on the past x_{t-1} , but also on the future x_{t+1} . When F is nonlinear such systems are hard to solve. The methodological contribution of our paper is to develop an explicit functional representation of the global solution to the nonlinear system (1). This representation can be understood as an iterative procedure over the sequence space, which allows us to simulate the global dynamics of the nonlinear system. The degree of complexity of a rational expectations solution is increasing tremendously when including boundedly rational agents in a model. To our best knowledge our paper is the first to successfully apply such techniques to detect cycles and chaotic temporary equilibrium dynamics.

The paper is organized as follows. After discussing some related literature in the rest of this section, the following Section 2 presents the temporary equilibrium asset pricing framework with heterogeneous expectations and Section 3 develops the numerical solution method. Section 4 dis-

¹ F may e.g. be an Euler equation derived from intertemporal utility maximization.

cusses some simple linear examples. In Section 5 we apply our numerical method to study the competition between rational and non-rational agents and cycles that may emerge. Section 6 discusses an example with chaotic equilibrium dynamics, while Section 7 concludes.

Related Literature

A large literature has addressed the question of whether fully rational agents drive out irrational agents. Rationality involves two important aspects: utility or profit maximization (Subjective Expected Utility, SEU) and expectations (Rational Expectations Hypothesis). Blume and Easley (2010) provide a survey and stimulating discussion of their and others' work on the *market selection hypothesis*, that is, the question: do markets redirect resources toward rational SEU decision makers? They conclude that utility maximization puts only few restrictions on asset prices and stress the importance to study differences in beliefs. Sandroni (2000) and Blume and Easley (2010) investigate how rational expectations fare against traders with non-rational beliefs. They show that if markets are dynamically complete, the economy is bounded, and traders have a common discount factor, then the market is dominated in the long run by those with correct expectations. Hence, if there is any trader with correct beliefs, in the long run asset prices converge to their rational expectations values. Finally, drawing on the literature on incomplete markets Cao (2018) find that over-optimistic agents not only survive but can even benefit from expectations other than rational, and can increase overall price volatility. None of these models generate realistic boom-bust cycles however, as in our nonlinear framework.

Another branch of literature in behavioral finance – the literature on noise traders – also studies models with rational and non-rational agents. “Noise traders”, a term due to Kyle (1985) and Black (1986), are investors whose changes in asset demand are not driven by news about economic fundamentals, but rather by non-fundamental considerations such as changes in expectations or market sentiment. Early and influential papers include work of De Long et al. (1990b,a), showing that noise traders can survive in a world with rational agents. The first paper shows that traders with incorrect beliefs can earn higher expected returns than those earned by traders with correct beliefs, because they take on extra risk. The second paper uses a 3-period model to show that in the presence of positive feedback traders, rational speculation can be destabilizing. This model thus explains overreaction to news about economic fundamentals, caused by rational informed speculators taking into account the presence of feedback traders. None of these models use a temporary equilibrium framework and bubble and crash cycles do not arise as in our model.

In an economy with heterogeneous agents, rational expectations are a strong assumption as rational agents must fully understand the economic environment, the behavior of all non-rational agents and their relative share among all traders, the latter being potentially time-varying. As

mentioned above, temporary equilibrium systems with rational versus non-rational agents are hard to study or even simulate. A less demanding assumption would be to replace the fully rational agents by agents who would behave rationally only in a world where all other agents behave rationally too. In a homogeneous rational world prices would be at fundamental value. Therefore, homogeneous rational agents are often called fundamentalists. There is a large literature on fundamentalists versus chartists models (see e.g. the survey in Hommes (2006)). These models are easier to handle as one would replace the future value x_{t+1} in (1) by a fundamental equilibrium value, which then typically leads to backward-looking learning models that are straightforward to simulate or can even be studied analytically. Such models however do not allow to analyse the interaction of heterogeneous agents if some of the agents are perfectly rational.

Our paper also relates to more recent heterogeneous agents models in behavioral macro economics, see e.g. the chapters in the *Handbook of Computational Economics, Vol. 4, Heterogeneous Agent Modeling* (Hommes and LeBaron, 2018), in particular the chapter of Branch and McGough (2018) on the micro-foundations of heterogeneous agents macro models and the recent survey on behavioral macroeconomics provided in Hommes (2021).

Our numerical method shows that the model with rational and trend-extrapolative agents can feature multiple equilibria, that is there may be multiple stable rational expectations equilibria. The type of dynamics then depends crucially on whether rational expectations are anchored around a fundamental or nonfundamental steady state. Prices are stationary –but larger than the fundamental value– if rational expectations are anchored around the nonfundamental steady state. In this equilibrium, the extrapolative drift away from the fundamental and the stabilizing effect of rational agents are in exact balance. The exact same model features boom-burst dynamics between the fundamental price and the nonfundamental steady state if rational expectations are anchored at the fundamental steady state. Such multiple equilibria that do not depend on fundamental variables are commonly referred to as *sunspots*.² This finding is important because expectations anchoring play a central role in macroeconomics, e.g. for the conduct of monetary economics (Woodford, 2005). The anchoring of long-term expectations is also considered as an explanation for the breakdown of the macroeconomic relationship between inflation and real activity (“the missing Deflation/Inflation puzzle”) by Coibion and Gorodnichenko (2015) and Ball and Mazumder (2018). Mertens and Ravn (2014) and Aruoba et al. (2018) discuss implications of expectations driven liquidity traps. In these models, a sunspot can trigger agents to coordinate on an (unstable) equilibrium in which the zero lower bound on nominal interest rate (ZLB) binds, which in turn causes unfavorable macroeconomic

²See e.g. Benhabib and Farmer (1999) for a survey on sunspots in macroeconomics.

conditions. The concept of sunspot equilibria has also been used to explain asset price bubbles (e.g. Gali, 2014). While to our best knowledge, of the sunspot equilibria treated in the literature at least one equilibrium is indetermined and hence dynamically unstable (e.g. the ZLB steady state or the sunspot asset price bubbles), in our model *both* steady states provide a stable equilibrium.

By exploring the chaotic dynamics that may result from the interaction of heterogeneous agents including rational traders, our paper also adds to the literature on the dynamics of financial crisis. A host of studies (see e.g. Schularick and Taylor, 2012; Greenwood and Hanson, 2013; López-Salido et al., 2017; Greenwood et al., 2020) documents that rapid asset price growth together with excessive credit growth are a good predictor of financial crisis and financial fragility events. Their work however can neither explain why asset prices start to rise in the first place, nor why they eventually collapse. We close this gap by showing that extrapolative dynamics can arise endogenously even if a sizable share of agents is fully rational. In our heterogeneous agents model such “bubbles” burst when deviations of asset prices from fundamental value become large and more and more agents switch to a rational forecasting rule endogenously bringing prices back close to fundamental value.

Lastly, by developing a numerical iterative procedure that allows solving systems with cycles and chaotic dynamics we also contribute to the literature on computational methods. Our method is related to the *extended path method* by Fair and Taylor (1983) as well as to the *policy function iteration* method proposed by Coleman (1990, 1991), which has been applied successfully in macroeconomic research.³ Instead of iterating on a numerical representation of the solution to F as in a policy function iteration algorithm, we directly iterate on the expected future trajectory (the “extended path” $\{E_t x_{t+s}\}_0^\infty$) given X_{t-1} . Other than with the extended path method, we do not ex-ante truncate this expected future trajectory and do not require any root findings methods, which may cause numerical instability. See Judd (1998), Miranda and Fackler (2004) and Ljungqvist and Sargent (2012) for comprehensive, general surveys of numerical methods in economics. Finally, our approach can be seen complementary to the literature on rational inattention (Sims, 2003, 2010) as we show that if the rational predictor is costly and alternative predictors are successful, a sizable fraction of agents has an incentive to opt for the non-rational prediction.

2 An asset pricing model with heterogeneous beliefs

In this section we recall the main dynamic equations of the asset pricing model with heterogeneous beliefs of Brock and Hommes (1998), henceforth BH98. This model allows for heterogeneous boundedly rational agents, who switch between different forecasting rules based upon their relative

³Recent examples include Gust et al. (2017); Atkinson et al. (2019); Richter and Throckmorton (2016).

performance. BH98 consider two broad classes of agents: fundamentalist who believe that the asset price is equal to its fundamental value – the discounted sum of expected future dividends– and chartists, who use simple forecasting strategies to extrapolate patterns in past prices. BH98 show that the heterogeneous expectations model exhibits complex, chaotic bubble and crash dynamics, with agents switching between stabilizing fundamentalist strategies and destabilizing trend-following strategies.

Notice that fundamentalists act as if they are rational agents in a homogeneous rational world. However, in a model that includes heterogeneous non-rational agents, the fundamentalists are not fully rational, because they do not take the behavior of non-rational agents into account. A novel feature of the current paper is that we consider fully rational agents in this heterogeneous agent setup, having perfectly rational expectations that take into account the behavior of other non-rational agents. Full rationality, however, makes the model hard to solve, even numerically.

Consider an asset market where agents can choose between a risk free asset paying a gross return $R = 1 + r$, where $r > 0$ is the interest rate, and a risky asset that pays an uncertain dividend y_t in each period. Agents are neither constrained in borrowing nor in short-selling. There are H different trader types with different expectations about the future price of the asset. Furthermore, each trader type is a myopic mean variance maximizer, which implies that trader type i 's demand $z_{i,t}$ for the risky asset is a linear function of his belief $x_{i,t+1}^e$ about the asset price in $t + 1$ as well as today's price. Let x_t be defined as the deviation of the asset price at time t from its fundamental value. The market clearing price then reads as a no-arbitrage condition of the form

$$Rx_t = \sum_{i=1}^H n_{i,t} x_{i,t+1}^e, \quad (2)$$

where $1/R$ represents the (time-invariant) discount rate and $n_{i,t}$ represents the fraction of trader type i with all fractions adding up to 1. This equation will also be called the *law of motion* (LOM) of the model. Notice that if all agents are fundamentalists, i.e., believe that the future deviation from fundamental value is zero, then the realized market price will be exactly at fundamental value. Hence, the fundamental value is a rational expectations equilibrium (REE) steady state of the model.

As a simple, but typical, example exhibiting cycles, we focus on a model with three types of agents of which two types will be symmetrical. Agents are either rational (type 1) or biased optimists (type 2) or biased pessimists (type 3). The biased forecasts are given by

$$x_{2,t+1}^e = +b \quad \text{and} \quad x_{3,t+1}^e = -b, \quad (3)$$

where the magnitude of the bias is denoted by $b > 0$. Rational agents have perfect information about the model, including the behaviour of all other agents, and the rational forecast is given by the expectation $E_t x_{t+1}$ conditional on the information available at time t .⁴

Market clearing with rational versus biased agents is then given by

$$Rx_t = n_{1,t}E_t x_{t+1} + (n_{2,t} - n_{3,t})b, \quad (4)$$

where $n_{h,t}$, $h = 1, 2, 3$ denotes the fractions of rational, optimistic and pessimistic agents at time t , with all fractions adding up to 1. In the case of a deterministic model, i.e. absent any exogenous disturbances, rational agents have perfect foresight and we can simply set $E_t x_{t+1} = x_{t+1}$ to obtain⁵

$$Rx_t = n_{1,t}x_{t+1} + (n_{2,t} - n_{3,t})b. \quad (5)$$

Following Brock and Hommes (1998), these fractions are updated according to the *performance measure* $\pi_{i,t}$ for each predictor i , given by *realized profits*:

$$\pi_{i,t} = (x_t - Rx_{t-1})(x_{i,t}^e - Rx_{t-1}) - \mathbb{1}_{RE}C, \quad (6)$$

where C is the information gathering cost for obtaining the rational expectations forecast and $\mathbb{1}_{RE}$ an indicator function that equals one if the agent is rational and zero otherwise. All other predictors are costless. The first bracket of the first term at the RHS of (6) denotes the actual resale value of the asset minus the opportunity costs for financing the purchase in the previous period. The term in the second brackets represents agent type i 's demand, resulting from the mean-variance maximization given the agent's past belief about the price. The choice of the performance measure is an essential ingredient of the model in determining the properties of the dynamic system. Realized profits from trading qualify as a natural choice for the fitness measure. Notice that the best forecaster does not necessarily earn the highest profit. In order to receive a positive profit it is sufficient to have made a correct choice on whether to go short or long. Likewise, any trader A that has a strong positive belief about next periods price will invest more money in the asset than another trader B with a relatively lower positive forecast. Even if B 's forecast was perfectly correct, trader A may still earn higher profits since he invested more. This feature, i.e. that profits are not proportional to

⁴In the rest of this work we are using the terms *rational* and *rational expectations* interchangeably. Like all agents in our model, rational agents are mean-variance maximizers, based on rational expectations of future prices.

⁵Notice that in this model, due to the absence of any short selling and borrowing constraints, a boom in asset prices is implicitly accompanied by an increase in leverage and credit growth. As mentioned earlier, this constitutes a feature of empirical relevance that we will further discuss in Section 6.

forecasting errors, is unique to financial markets and captured by Equation (6). Further, note that the asset pricing model with heterogeneous beliefs is a zero-sum game, that is, the profits averaged over all agent types add up to zero minus the information gathering costs for rational agents.⁶

The fraction of agents evolves endogenously depending on the relative performance. More precisely, the fraction of type $i = 1, \dots, H$ is determined by a *multinomial discrete choice model* depending on the past performance of the predictor:

$$n_{i,t} = \frac{e^{\beta\pi_{i,t-1}}}{\sum_{j=1}^H e^{\beta\pi_{j,t-1}}}. \quad (7)$$

If a predictor has been relatively more successful than others, it is more likely to be chosen and hence the fraction of agents using this predictor increases. β is called the *intensity of choice* which governs how quickly agents switch between predictors. If $\beta \rightarrow \infty$, all agents will immediately switch to the most successful predictor. This completes the specification of the asset pricing model with heterogeneous beliefs.

3 Numerical solution

We continue by describing our numerical solution method. Plugging the performance measure (6) into the fractions (7) and inserting the result into the market clearing equation (4), the model's state x_t at date t can be expressed as a (known) nonlinear function f :

$$x_t = f(x_{t+1}, x_{t-1}, x_{t-2}, x_{t-3}). \quad (8)$$

Notice again that x_t is a function of the *forward looking* rational forecast of next periods price x_{t+1} , as well as of *past values* $x_{t-1}, x_{t-2}, x_{t-3}$.⁷

We propose a novel numerical method to solve for the rational expectation equilibrium of a class of highly nonlinear problems where complicated dynamics may arise, and we are particularly interested to detect cycles or chaotic solutions. To our best knowledge we are the first in applying iterative numerical methods to solve for cyclic and chaotic rational expectations equilibria of nonlinear temporary equilibrium mappings.⁸ We are looking for a state-space solution of this model, i.e.

⁶More precisely, using the equilibrium pricing equation (2) a simple computation shows that $\sum_i n_{i,t}\pi_{i,t+1} = 0 - n_{R,t}C$, where $n_{R,t}$ is the fraction of rational agents in period t and C the cost for rational expectations.

⁷For the 3-type example with purely biased agents introduced in the previous section, the fractions $n_{i,t}$, through the performance measure, depend only on two lagged state variables x_{t-1}, x_{t-2} . In Section 6 we will consider another 2-type example, where one of the forecasting rules is a trend-following rule that depends on the lagged state variable x_{t-1} , so that the time-varying fractions, through the performance measure depend on an additional lagged variable x_{t-3} . Therefore our notation uses three lagged state variables to cover all examples.

⁸Grandmont (1985) shows the existence of cycles and chaos in simpler 1-dimensional overlapping generations

a functional representation that only depends on known states at time t . Since f is nonlinear and x_t depends on its future value x_{t+1} , we can not solve explicitly for a state-space solution in closed form and must rely on numerical solution techniques.

The intuition behind our numerical method goes as follows: the current price x_t depends on x_{t+1} , which again recursively depends on future prices. As the future is ex-ante unknown, any representation of the law of motion that depends on future prices is not useful and we require a solution to the system that depends only on past state variables. Assume that a steady state to f is known. We start by guessing that x_{t+1} is in the steady state and, given this guess, solve for x_t . We then update our guess on x_{t+1} by using our current computation for x_t and assuming that in x_{t+2} the system will be in the steady state. Using the new guess on x_{t+1} , we update the computation on x_t , which can then be used to re-evaluate the guess on x_{t+1} . That means, we go back and forth in time, thereby updating and extending guesses on the expected future trajectory of x . We stop once new guesses on values that are far in the future do not affect our approximation of x_{t+1} anymore. Such a (functional) recursive state-space representation that yields the current state of the system as a function of the history of past state state variables is commonly referred to as a *policy function*.

Thus, assume the existence of a policy function g for the model described by f in (8) for a given $X_{t-1} = \{x_{t-1}, x_{t-2}, x_{t-3}\}$:

$$x_t = g(x_{t-1}, x_{t-2}, x_{t-3}). \quad (9)$$

If a function g exists, using (8) it must satisfy

$$x_t = f(x_{t+1}, x_{t-1}, x_{t-2}, x_{t-3}) \quad (10)$$

$$= g(x_{t-1}, x_{t-2}, x_{t-3}). \quad (11)$$

By moving time forward one period and inserting g we get

$$x_{t+1} = g(x_t, x_{t-1}, x_{t-2}) \quad (12)$$

$$= g(g(x_{t-1}, x_{t-2}, x_{t-3}), x_{t-1}, x_{t-2}). \quad (13)$$

Then our problem boils down to finding a function g that satisfies

$$g(x_{t-1}, x_{t-2}, x_{t-3}) = f(g(g(x_{t-1}, x_{t-2}, x_{t-3}), x_{t-1}, x_{t-2}), x_{t-1}, x_{t-2}, x_{t-3}). \quad (14)$$

models of the form $x_t = f(x_{t+1})$. Our method computes cycles and chaotic rational expectations equilibria model for 2- and higher dimensional temporary equilibrium mappings.

If f is nonlinear we can not solve for the function g explicitly, but we can approximate it numerically. Define the approximation error ζ_t (also often called *Euler Equation error*, EEE) at time t as

$$\zeta_t = |g(X_{t-1}) - f(g(X_{t-1}), x_{t-1}, x_{t-2}), X_{t-1}|, \quad (15)$$

with $X_t = \{x_{t-1}, x_{t-2}, x_{t-3}\}$. The following proposition provides an algorithm to find a function g that approximates a solution to f :

Proposition 1. *An approximate state space solution $g(\cdot)$ to (8) with steady-state \bar{x} and a given X_{t-1} is, up to a precision of $\zeta_t < \epsilon$, given by*

$$x_t = g(X_{t-1}) = f(\hat{x}_{1,K}, X_{t-1}), \quad (16)$$

with

$$\hat{x}_{i,k} = \begin{cases} x_{t+i} & \text{for } i < 0, \\ f(\hat{x}_{i+1,k-1}, \hat{x}_{i-1,k}, \hat{x}_{i-2,k}, \hat{x}_{i-3,k}), & \text{for } 0 \leq i \leq k, \\ \bar{x} & \text{for } k < i, \end{cases} \quad (17)$$

where it holds for $K > 1$ that $|\hat{x}_{1,K} - \hat{x}_{1,K-1}| < \epsilon$. Note that $\hat{x}_{i,k}$ can be interpreted as the k^{th} guess on x_{t+i} .

In practice, this implies a simple iterative scheme over k and i . To clear notation, let $\hat{x}_{i,k}$ be the guess on x_{t+i} at iteration k and define $\hat{X}_{i,k} = \{\hat{x}_{i,k}, \hat{x}_{i-1,k}, \hat{x}_{i-2,k}\}$. Start the iteration with $k = i = 0$ and set $\hat{x}_{0,0} = f(\bar{x}, X_{t-1})$.⁹ Now increase $k \rightarrow 1$ and note that since $\hat{x}_{1,0} = \bar{x}$ (i.e. the guess on x_{t+1} for $k = 0$), it follows that $\hat{x}_{0,1} = \hat{x}_{0,0}$. For $k = i = 1$ it is that $\hat{x}_{1,1} = f(\bar{x}, \hat{X}_{0,1})$, where $\hat{X}_{0,1} = \{\hat{x}_{0,1}, x_{t-1}, x_{t-2}\}$. Increment k to $k = 2$ and recursively calculate $\hat{x}_{i,2}$ for $i \in \{0, 1, 2\}$. Continue to increase k , and for each k recursively calculate $\hat{x}_{i,k}$ for $i \in \{0, 1, \dots, k\}$ until it holds for $k = K$ that $|\hat{x}_{1,k} - \hat{x}_{1,k-1}| < \epsilon$. The iterative scheme is illustrated in Table 1.

Why does this work? By repeatedly increasing k and calculating $\hat{x}_{1,k}$, we get a better approximation of the expected future trajectory of x_t with every iteration. Incrementing k implies increasing the horizon for which this trajectory is calculated. For K , convergence of $|\hat{x}_{1,K} - \hat{x}_{1,K-1}|$ to a value below ϵ means that the iteration calculates expected values of this series that are far into the future and will not have a significant impact on the value of x_t . Note that the definition of ζ_t involves $g(X_t) = g(g(X_{t-1}, \cdot))$ inserted in f , i.e. the value of x_{t+1} as implied by g given X_{t-1} . For this

⁹In theory, many other guesses may also work where \bar{x} is not the steady state. However, as we are considering complicated maps with potential boom-burst dynamics, there may also be guesses for which the procedure may simply diverge due to infinite bubble solutions. We want to ex-ante rule out such solutions.

	$\hat{x}_{0,k}$	$\hat{x}_{1,k}$	$\hat{x}_{2,k}$	\cdots	$\hat{x}_{K-1,k}$	$\hat{x}_{K,k}$
$k = 0$	$f(\bar{x}, X_{t-1})$	\bar{x}	\bar{x}	\cdots	\bar{x}	\bar{x}
$k = 1$	$f(\bar{x}, X_{t-1})$	$f(\bar{x}, \hat{X}_{0,1})$	\bar{x}	\cdots	\bar{x}	\bar{x}
$k = 2$	$f(\hat{x}_{1,1}, X_{t-1})$	$f(\bar{x}, \hat{X}_{0,2})$	$f(\bar{x}, \hat{X}_{1,2})$	\cdots	\bar{x}	\bar{x}
$k = 3$	$f(\hat{x}_{1,2}, X_{t-1})$	$f(\hat{x}_{2,2}, \hat{X}_{0,3})$	$f(\bar{x}, \hat{X}_{1,3})$	\cdots	\bar{x}	\bar{x}
\vdots	\vdots	\vdots	\vdots		\vdots	\vdots
$k = K - 1$	$f(\hat{x}_{1,K-2}, X_{t-1})$	$f(\hat{x}_{2,K-2}, \hat{X}_{0,K-1})$	$f(\hat{x}_{3,K-2}, \hat{X}_{1,K-1})$	\cdots	$f(\bar{x}, \hat{X}_{K-2,K-1})$	\bar{x}
$k = K$	$f(\hat{x}_{1,K-1}, X_{t-1})$	$f(\hat{x}_{2,K-1}, \hat{X}_{0,K})$	$f(\hat{x}_{3,K-1}, \hat{X}_{1,K})$	\cdots	$f(\bar{x}, \hat{X}_{K-2,K})$	$f(\bar{x}, \hat{X}_{K-1,K})$

Table 1: Illustration of our numerical method. As above, $\hat{X}_{i,k} = \{\hat{x}_{i,k}, \hat{x}_{i-1,k}, \hat{x}_{i-2,k}\}$ and $x_{-s,k} = x_{t-s}$ for $s \in \{1, 2, 3\}$.

reason we require the convergence of $\hat{x}_{1,K}$ instead of $\hat{x}_{0,K}$, that is, the convergence of the guess on x_{t+1} given X_{t-1} . This ensures that the actual error on $g(X_{t-1})$ is also below ϵ at each time step t . The limiting function $g_K(X_{t-1})$ then approximates the true state-space representation g for a given X_{t-1} , and we can use g to simulate the global dynamics.

We are interested in the (long run) dynamics for the relevant domain given a vector of initial values $X_0 = (x_0, x_{-1}, x_{-2})$ and aim to find the solution given these initial values. Note that the limiting function $g_K(X_{t-1})$ is, while being an ϵ -precise approximation of g , not a global solution but holds only for X_{t-1} . That means we must reevaluate g for every point on the trajectory. Computationally, this can efficiently be implemented by re-using the series of $\hat{x}_{i,K}$ from $g_K(X_{t-2})$ for the evaluation of $g(X_{t-1})$ by just dropping $\hat{x}_{0,K}$. The evaluation of a time series for the example given in Figure 4 of length 1000 only takes about 210 milliseconds on a normal laptop. Further note that convergence of g_K for a given value of ϵ implies that the actual approximation error on g must be below ϵ .

Our method has a few significant advantages over alternative methods. First, it does neither rely on interpolation nor on root finding (as e.g. policy function iteration or the Fair-Taylor method). This maintains a high degree of accuracy and numerical stability. A crucial advantage of our method over methods that are based on a grid is that it is fully insensitive to the so called *curse of dimensionality*.¹⁰ Further, compared to grid-based methods our method does not require to ex-post specify of a bounded domain. For our case in which f is highly nonlinear this is of particular importance: for some regions in the immediate neighborhood of the trajectories discussed further below, there may only exist explosive solutions. As for these regions no rational expectations equilibria exist, global convergence of a grid-based policy function is hence impossible. Note that our method inherits its convergence properties, similar to other dynamic programming techniques, from the contraction mapping theorem. It is not strictly necessary that f is a contraction mapping as long as the contrac-

¹⁰The *curse of dimensionality* describes the problem that the computational complexity of standard grid-based methods grows exponentially in the number of states.

tion property holds for at least one implicit higher order mapping of f of the form $(f \circ f \circ \dots \circ f)$. Further note that our method can be seen as a policy function iteration solved for only one particular initial state X_{t-1} , where all grid points are chosen endogenously to solve for the anticipated path given X_{t-1} .

As a robustness exercise, we also adapt a policy function iteration algorithm to our highly nonlinear problem, thereby building on Coleman (1990, 1991). Details and simulations for the this alternative method can be found in Appendix A. Although approximation errors with the policy function iteration algorithm are several magnitudes larger, the method generally delivers the same results qualitatively and quantitatively, however at considerably larger computational costs.¹¹

4 Linear examples

In this section we provide two linear examples in order to get some intuition on how rational agents affect the dynamic of a heterogeneous agents asset pricing model. These examples also help to illustrate the workings behind our numerical method. Rational expectations solutions to linear systems are well understood and a solid body of literature provides robust and efficient solution methods (see e.g. Klein, 2000; Rendahl, 2017). Hence, if the mapping f in (8) is linear, our method can be expected to converge when the conditions of Blanchard and Kahn (1980) are satisfied. Note also that the state-space representation g is closely related to the Minimum State Variable (MSV) solution concept for rational expectations solutions of dynamic models in the adaptive learning literature (Evans and Honkapohja, 2003). The MSV solution technique is convenient in linear models, where it can be computed explicitly.

The first linear example illustrates how our method works in a homogeneous asset pricing model. Consider the simplest model

$$x_t = R^{-1}E_t x_{t+1} \tag{18}$$

where x_t and $E_t x_{t+1}$ are the deviation and expected deviation from fundamental price and $0 < R^{-1} < 1$ is the (constant) discount factor. Under perfect foresight the model becomes $x_t = x_{t+1}/R$. One can compute MSV solutions of the form $x_t = b x_{t-1}$. Substituting this solution in the model gives $x_t = (b^2/R)x_{t-1}$. This induces a so-called T-map: $T(b) = b^2/R$, a concept extensively used in the adaptive learning literature (Evans and Honkapohja, 2003). The fixed points of the T-map are $b = 0$ and $b = R$, corresponding to the two perfect foresight MSV solutions. To gain understanding of

¹¹An exception for this is the parameter region of the two-type model with sunspot equilibria that we discuss in Subsection 6.2. For the model specification discussed there approximation errors of the alternative method are large and simulation results misleading.

the stability properties of the system we can, similar to our iterative method from Section 3, look at a sequence of b -values under the T-map: $b_{k+1} = T(b_k)$.¹² A simple graphical analysis of the T-map shows that when the initial value $b_0 < R$, b_k converges to 0, while $b_0 > R$ leads to b_k diverging to ∞ . This simple example illustrates that (i) MSV solutions with perfect foresight exist, and (ii) iteration of the T-map shows that $b = 0$ is a stable MSV solutions, while $b = R$ is an unstable MSV solution. Iterations of the T-map are thus path-dependent on whether initial conditions are smaller or larger than the threshold $b_0 = R$. Iteration of the T-map is similar to the numerical approximation of the policy function g by a the limiting function g_k of the sequence g_0, g_1, \dots, g_k , etc. In fact, when applying our numerical algorithm with steady state function $\bar{x} = 0$, the sequence of numerical maps g_k exactly coincides with the sequence of T-maps, i.e. $g_k = T^k$.¹³

As a second linear example, consider a 3-type asset pricing model with rational agents versus trend-followers versus fundamentalists. To allow for complex eigenvalues the trend-extrapolating rule has an extra lag x_{t-2} , and is of the form $x_{2,t+1}^e = x_{t-1} + \gamma(x_{t-1} - x_{t-2})$. This generates oscillatory behaviour in the system. The fractions of the three types are n_{RE} , n_γ and $n_0 = 1 - n_{RE} - n_\gamma$ and they are – for now – fixed. The linear dynamic model is given by:

$$Rx_t = n_{RE}x_{t+1} + n_\gamma(x_{t-1} + \gamma(x_{t-1} - x_{t-2})) + n_0 \cdot 0, \quad (19)$$

where we set $R = \gamma = 1.1$.

Assuming $n_\gamma = 0.75$, Figure 1 illustrates four cases:

1. For $n_{RE} = 0$ trend-followers and fundamentalists share the market. As a result the price oscillates and converges quickly to its fundamental value.
2. Replacing a small fraction of the fundamentalist by rational agents, with $n_{RE} = 0.1$, leads to stable oscillations that converge slowly to the fundamental value (orange curve). This shows that rational agents are destabilising compared to fundamentalists, because rational agents are taking into account the oscillatory behavior caused by trend-followers.
3. The case $n_{RE} = .16$ has a larger share of rational agents and fewer fundamentalists and shows a perfect oscillation (green plot). With the trend-extrapolating coefficient $\gamma = R$ this leads to a system with complex eigenvalues on the unit circle. Comparing with the previous case, rational agents further amplify the persistent oscillations caused by trend-followers.
4. Any further increase of $n_{RE} > .16$ (not in the figure) at the cost of fewer fundamentalists

¹²In the learning literature, this concept is known as iterative E-stability.

¹³Note that because the model in (18) does not have any state variables, a limiting policy function g consistent with our exposition in Section 3 would simply be a constant: $g_k = c$ and independent of the history X_{t-1} .

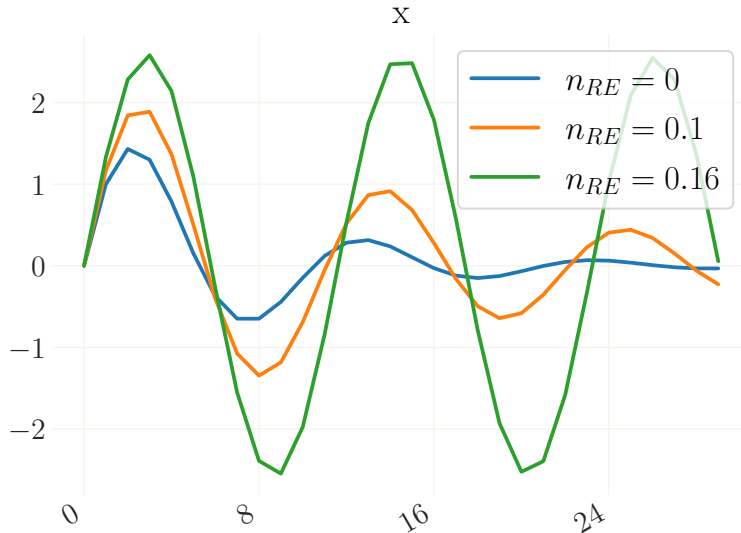


Figure 1: Time series of 3-type model with rational versus trend-extrapolating rule versus fundamentalists. Stable oscillations for $n_{RE} = 0$ and $n_{RE} = 0.1$ (blue and orange) and persistent oscillations for $n_{RE} = 0.16$ (green). When rational agents replace fundamentalists ($n_{RE} = 0.1$) slowly converging oscillations arise (orange), because rational agents anticipate the presence of trend-followers.

will cause the complex eigenvalue to cross the unit circle, leading to two eigenvalues with modulus larger than one. This violates the Blanchard-Kahn conditions and renders the rational expectations system indeterminate.

This linear example provides some intuition about the effect of rational agents in a heterogeneous (linear) world. In the presence of non-rational agents, rational agents are destabilizing compared to fundamentalists, because they take the boundedly rational behaviour into account. In the next section the fractions of these different types of agents become time varying, depending on the relative success of each strategy. With agents switching between different strategies, the system becomes highly nonlinear and to simulate the model we need to rely on our numerical solution technique.

5 Nonlinear simulations in the three-type model

In this section we apply the numerical solution method to the asset pricing model with heterogeneous expectations and use it to study the survival of (ir)rational beliefs. We consider two typical examples to illustrate the complex dynamical behavior in the model. The first example is the 3-type model with rational agents versus optimistic and pessimistic biased agents, as introduced in Section 2. This example exhibits periodic boom and bust cycles along which the fractions of rational and non-rational agents co-evolve over time. The second example is a 2-type model with rational agents

versus trend-extrapolators, which exhibits chaotic bubble and crash dynamics as discussed in Section 6. In both examples rational agents can not drive out irrational agents, but the fractions of rational and irrational beliefs rather co-evolve endogenously causing boom and bust cycles.

5.1 An example with biased agents and limit cycles

Consider the 3-type example with rational agents versus optimistic and pessimistic biased agents as given in (5), with time varying fractions given by (6) and (7). Since the fitness measure of the biased types includes two lags of x , this 3-type system is of the form

$$x_t = f(x_{t+1}, x_{t-1}, x_{t-2}). \tag{20}$$

There are three behavioral parameters: the intensity of choice β , the bias b and the costs C , whose benchmark values are given in Table 2. For our benchmark, the parameters β and b are normalized to 1 and the costs for rational expectations are set to zero. Independent of the parametrisation, there exists a trivial fundamental steady state $x = \bar{x} = 0$ where the steady state fractions $\bar{n}_2 = \bar{n}_3$ of both biased agents is equal. Throught this section, we only consider this fundamental steady state for our solution g .

R	β	b	C
0.99^{-1}	1	1	0

Table 2: Benchmark parametrisation

Let us first explore the potential dynamics of the system by varying the intensity of choice parameter β . The bifurcation diagram in Figure 2 shows the long-run dynamics of price deviations from fundamental as a function of the intensity of choice β .¹⁴ For low values of β , the simulations show that the fundamental steady state is stable and the solution converges to it. As β increases, the steady state becomes unstable and a limit cycle emerges after a Hopf-Bifurcation at $\beta \approx 1.2$. The amplitude of these cycles increases with β ; for higher values of β a stable 4-cycle arises. Figure D.26 in Appendix D illustrates the policy functions g for different values of β graphically. Phase diagrams can be found in Appendix C.

Figure 3 shows a similar bifurcation diagram for the bias parameter b : when the bias b increases, the system destabilizes through a Hopf-bifurcation and cycles of increasing amplitude arise as the bias increases. For large values of the bias b a stable 4-cycle arises.

¹⁴For each bifurcation diagram below and in Section 6 we draw a linear sequence of 1000 parameter values from the shown interval. For each draw, we use our method to simulate a time series of 2,000 periods, of which we discard the first 1,000 iterations as convergence period.





Figure 2: Bifurcations w.r.t. the intensity of choice β . All other parameters as in Table 2. Required minimal precision is $\epsilon = 1e - 14$. Figure 3: Bifurcations w.r.t. bias b . All other parameters as in Table 2. Required minimal precision is $\epsilon = 1e - 14$.

To understand the cycles in more detail, Figure 4 shows time series of the price deviations from the fundamental, the fractions of the three agent types, the profits of each of the three types and the realized profits (i.e. the profit of each type multiplied by the fraction of that type).¹⁵ The dynamics of prices and fractions is (almost) periodic and the fractions of the three types co-evolve over time. Moreover, the fractions of optimists and pessimists fluctuate significantly stronger than the fraction of rational agents.¹⁶ The reason is that the realized profits of rational agents always lie between the profits of the optimistic and the pessimistic traders: while the price forecast of rational agents is always perfect, their profit according to (6) – and hence their fraction – depends on the squared realized excess return $(x_t - Rx_{t-1})^2$. The realized profit of optimists or pessimists is given by the realized excess return times their asset demand. In a booming market optimistic traders will then earn higher profits than rational agents since they have overestimated the change in price and hence bought more of the risky asset than rational agents. During a bust, the opposite holds: pessimistic traders buy less (or short the risky asset more) and will earn higher profits than rational agents. As a result, the fraction of rational agents always lies between the fractions of optimists and pessimists.

For high values of the intensity of choice this effect becomes stronger, as illustrated in Figure 5 for $\beta = 500$. For such high intensity of choice, in each period the market is fully dominated by either optimists or pessimists, with the fraction of rational agents always being (almost) 0. In the

¹⁵The asset market is a zero sum game, that is, $\sum_h n_{ht} \pi_{h,t+1} = 0$. Hence, any positive profits by one type must be balanced by losses by another type. In case of information gathering costs for rational agents, aggregate net profits sum to $-n_R C$, i.e. minus the aggregate costs.

¹⁶Note that, due to the symmetric construction of optimistic and pessimistic traders, the dynamics of optimists and pessimists are in fact symmetric.

limiting case of $\beta = \infty$ in each period the fraction of optimists or pessimists is either 0 or 1, while the fraction of rational agents is 0. Asset prices then follow a 4-cycle subsequently dominated for two periods by optimists and the other two periods by pessimists.

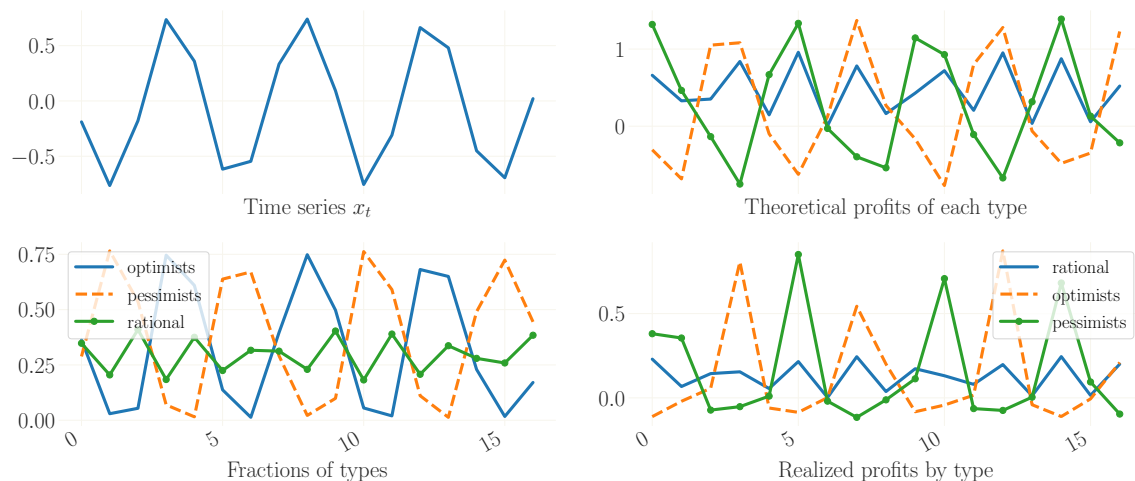


Figure 4: Time series for $\beta = 2.0$. Left: Prices and the evolution of the fraction of the different trader types over time. Right: Profits of the different trader types. Theoretical profits are as defined in Equation (6). Realized profits are theoretical profits multiplied by the fraction of each type respectively. All other parameters as in Table 2. Required minimal precision is $\epsilon = 1e - 14$.

This 3-type example illustrates that rational agents do not generally drive out boundedly rational agents from the market. If agents tend to switch to more successful strategy quickly, there is always a fraction of boundedly rational agents that is sufficiently successful to gather a share of the market. Any rational agent then adjusts his belief accordingly and by doing so amplifies the boundedly rational traders' beliefs. In the limit of an infinite intensity of choice, in fact the rational agents are driven out of the market and the price exhibits booms and busts in which optimistic and pessimistic traders subsequently dominate the market. These boom and bust cycles arise even in the absence of any information gathering costs for rational agents. Our numerical solution method easily picks up these boom and bust cycles with endogenously fluctuating rational and non-rational beliefs.

5.2 Comparing rational and fundamentalist traders

Fundamentalists traders are agents who believe that the price will always return to its fundamental value or, stated differently, who believe that the price deviation from fundamental will be 0. In a homogeneous world, fundamentalists are fully rational and the zero steady state would be the rational solution. In their asset pricing model with heterogeneous beliefs, BH98 focused on backward-looking expectations rules and considered many examples with fundamentalists versus chartists. However, BH98 did not consider the global dynamics of examples with forward-looking ra-

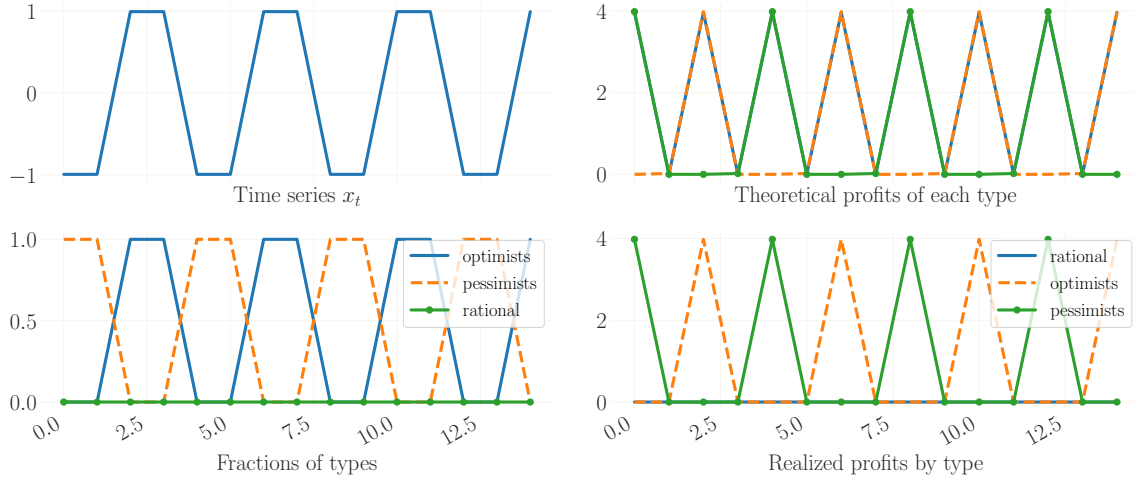


Figure 5: Time series for very high $\beta = 500$. Left: Time series of prices and the evolution of the fraction of the different trader types over time. Right: Profits of the different trader types. Theoretical profits are as defined in Equation (6). Realized profits are theoretical profits multiplied by the fraction of each type respectively. All other parameters as in Table 2. Required minimal precision is $\epsilon = 1e - 14$.

tional agents as we do here enabled by our numerical solution method. A natural question is whether in a heterogeneous agents model fully rational agents are stabilizing or destabilizing compared to fundamentalist agents.

To address the question, we compare the dynamics of two 3-type models. The first is our model with rational agents versus optimistic and pessimistic biased agents, as before. In the second model rational agents are replaced by fundamentalists, who always predict 0 deviation. The equilibrium pricing equation (5) then simplifies to

$$Rx_t = (n_{2,t} - n_{3,t})b, \quad (21)$$

with the modified fractions $n_{2,t}$ and $n_{3,t}$ (see BH98 for full details).

The bifurcation diagram in Figure 6 shows the long run behavior of both 3-type models as a function of the intensity of choice β . The diagram shows that the 3-type model with rational agents becomes unstable *earlier*, i.e., for lower value of the intensity of choice, than the 3-type model with fundamentalists. Hence, rational agents are destabilizing compared to fundamentalists in this 3-type world. The intuition is again that rational agents anticipate the presence of non-rational agents and thus amplify the effect of a positive or negative bias. In contrast, fundamentalists always predict the fundamental value and thus always add a stabilizing force in the presence of non-rational agents. In summary, the fact that rational agents take into account and anticipate the behavior of non-rational agents amplifies the destabilizing effect of these non-rational agents.




Figure 6: Bifurcations w.r.t. β . In cyan the same simulation with fundamentalists instead of rational agents. All other parameters as in Table 2. Required minimal precision is $\epsilon = 1e - 14$.

We have studied a 3-type example with rational agents versus optimists and pessimists, applying our numerical method to (20). The numerical approximation method works very well in this 3-type example and allows for solutions with a minimal required precision of $\epsilon = 1e - 14$. We cross-check our results with an alternative policy function iteration method in Appendix B. The alternative method provides quantitatively and qualitatively very similar results where average and maximum EEE are both very small. An likely reason for this is that the map f in (20) is bounded to the interval of biases, that is, $x_t \in [-b, b]$. Another reason for why the alternative method works well in the 3-type world is, that the trajectory does not describe any sudden jumps but is – over the respective domain – relatively smooth. In the next section we will consider a 2-type example with rational versus a trend-extrapolating rule. This environment poses more challenges to both numerical methods, because of the existence of explosive bubble solutions.

6 Chaotic solutions in a two-type model

We now consider another example with two trader types: rational agents vs. trend followers. Agents can choose between the rational, perfect foresight forecast at positive information gathering costs and a freely available trend-extrapolating forecast. Our numerical method detects chaotic bubble and crash dynamics. During the bubble phase trend-extrapolators dominate the market, thereby fueling the bubble. When the forecasting errors of the trend-extrapolators increase, the fraction of rational agents increases and eventually drives prices back close to its fundamental value, after which the story repeats and a new bubble forms.

The beliefs of trend followers are given by

$$x_{t+1}^e = \gamma x_{t-1}, \quad (22)$$

where γ denotes the degree of trend extrapolation. For $\gamma > 1$ trend-followers believe that the price deviation from the fundamental value will increase in the next period. The 2-type model with rational versus trend-following agents is given by

$$Rx_t = (1 - n_t)x_{t+1} + n_t\gamma x_{t-1} \quad (23)$$

$$\pi_{RE,t} = (x_t - Rx_{t-1})^2 - C \quad (24)$$

$$\pi_{BR,t} = (x_t - Rx_{t-1})(\gamma x_{t-2} - Rx_{t-1}) \quad (25)$$

$$n_t = \frac{e^{\beta\pi_{BR,t-1}}}{e^{\beta\pi_{BR,t-1}} + e^{\beta\pi_{RE,t-1}}}. \quad (26)$$

Here n_t is the fraction of trend-followers and $1 - n_t$ the fraction of rational agents, who incur information gathering costs $C > 0$.

Note that, since the last periods prediction of trend-followers now enters the profit equation, there is an extra time-lag and the function f is of the form (8) and the state-space from is $x_t = g(x_{t-1}, x_{t-2}, x_{t-3})$ as in (9), which poses additional computational challenges.

The following theoretical result about the fundamental steady state $x_t = 0$ and two additional non-fundamental steady states builds on Brock and Hommes (1998), who discuss this example briefly.

Lemma 1 (Existence of a steady state for the 2-type model). *Let $m^* = 1 - 2\frac{R-1}{\gamma-1}$ and x^* be the positive solution (if any) of*

$$x^* = \sqrt{z} \quad (27)$$

for

$$z = \frac{2 \operatorname{arctanh}(m^*)/\beta + C}{(\gamma - 1)(R - 1)}. \quad (28)$$

Let $E_1 = (0, \tanh(-\beta C/2))$, $E_2 = (x^*, m^*)$ and $E_3 = (-x^*, m^*)$.

1. For $\gamma < R$, E_1 is the unique steady state.
2. For $\gamma > 2R - 1$, there exist the three steady states E_1 , E_2 and E_3 .
3. For $R < \gamma < 2R - 1$, there are two possibilities:
 - (a) if $m^* < \tanh(-\beta C/2)$ then E_1 is the unique steady state.
 - (b) if $m^* > \tanh(-\beta C/2)$ then E_1 , E_2 , and E_3 are the steady states

Note that for the baseline simulations below, we initialize our policy function with the positive nonfundamental steady state whenever it exists, i.e. $\bar{x} \equiv \max\{0, x^*\}$. Throughout this section we use the parameters from Table 3. To eliminate approximation errors as much as possible, we set the required precision to $\epsilon = 1e - 14$.

R	β	γ	C
1.1	1.5	1.15	0.5

Table 3: Parametrisation for the two-type model.

6.1 Chaotic boom and bust cycles with rational traders

Figure 7 shows the bifurcation diagrams for the parameters C and γ , with $R = 1.1$, $\gamma = 1.15$ and $C = 0.5$. For low values of $\gamma < 1.15$ the solution converges to the fundamental steady state. As γ increases, a positive stable non-fundamental steady state arises, which becomes unstable for $\gamma \approx 1.172$. Cycles and chaos arise for higher values of the trend-extrapolation coefficient. Figure 8 shows the respective bifurcation diagram with respect to β . The figure suggests chaotic bubble and crash dynamics for values of β larger than 2.490., while Figure 7 (left) shows a similar bifurcation route to chaos as the cost parameter C increases.




Figure 7: Bifurcation diagrams w.r.t. cost parameter C (left) and γ (right). All other parameters as given in Table 3. Required minimal precision is $\epsilon = 1e - 14$.

The respective time series of prices, fractions and profits are displayed in Figure 9. The solution starts close to the fundamental value 0, with the fraction of trend followers larger than the fraction of rational agents because of the information gathering costs. With trend-followers dominating and sufficiently aggressive, prices increase exponentially. At some tipping point, however, the fraction of trend extrapolators drops to almost zero and all agents switch to the rational strategy. The asset

bubble bursts and prices drop back close to the fundamental steady state 0, from where the process repeats. Phase diagrams are provided in Appendix C.




Figure 8: Bifurcations w.r.t. β . All other parameters as given in Table 3. Required minimal precision is $\epsilon = 1e - 14$.

A graphical illustration of the policy function underlying the dynamics displayed in Figure 9 is shown in Figure 10. The mapping has three regions, independently of the value of x_{t-3} : for low values of x_{t-1} and high values of x_{t-2} the function maps to a low x_t (region 1). Likewise, high values of x_{t-1} and low values of x_{t-2} also map to low values of x_t (region 2). Along the diagonal of the (x_{t-1}, x_{t-2}) plane, the function maps to increasing values of x_t (region 3). In the bubble phase, the trajectory of Figure 9 travels along this diagonal. When x_{t-3} increases (as x_t rises), region 1 shrinks as region 2 grows and moves towards the diagonal (illustrated in the plane in the middle). When the trajectory is approaching its peaking point (right plane in Figure 10), region 1 disappeared and region 2 takes almost half of the image, thereby mapping x_t back close to the fundamental steady state. The figure also illustrates the complexity of the underlying problem, which our method is able to capture very well.

Intuitively it is clear that the fundamental steady state with $x_t = 0$ for all t , is unstable once $\gamma > R$ and $C > 0$ and β large enough. If this is the case, any increase of $C > 0$ or β will shrink the

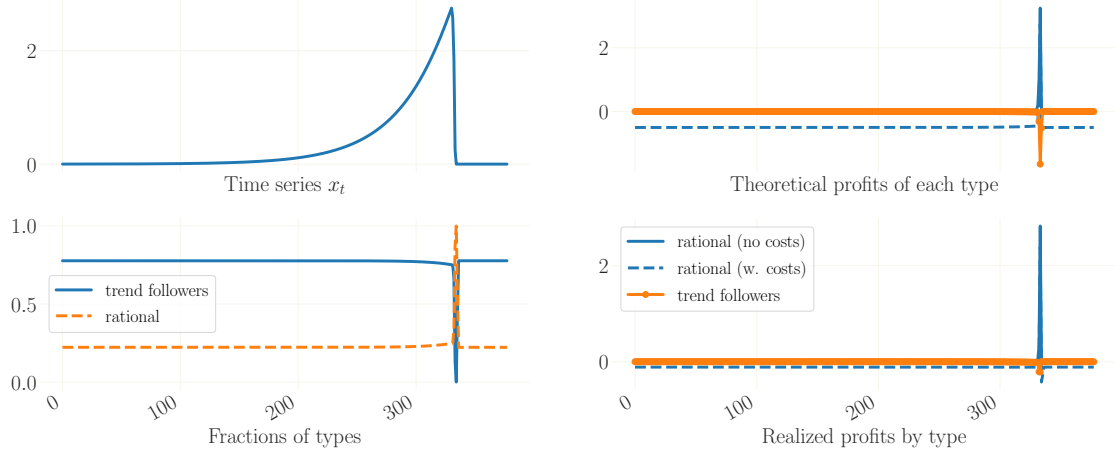


Figure 9: Time series for $\beta = 2.5$, $b = 0.0$, $\gamma = 1.15$, $R = 1.1$ and $C = 0.5$. Left: Time series of prices and the evolution of the fraction of the different trader types over time. Right: Profits of the different trader types. Theoretical profits are as defined in Equation (6). Realized profits are theoretical profits multiplied by the fraction of each type respectively. Required minimal precision is $\epsilon = 1e - 14$.

fraction of rational agents at the steady state and the extrapolative forces will therefore dominate around $x_t = 0$. This finding is confirmed by the time series and bifurcation diagrams.




Figure 10: Heatmap for slices of the policy function g for the chaotic time series in Figure 9. The mapping of f is nontrivial and not smooth on the relevant domain. For this figure we use the policy function iteration method outlined in Appendix A with a grid of $300 \times 300 \times 300$ points to obtain a high resolution of the image. Note that for the dynamics of the time series this increase in precision is not relevant.

To get some understanding of the complicated dynamics it is useful to consider the case $\beta = \infty$, that is, the case where all agents switching immediately to the better performing forecasting rule. At the fundamental steady state all agents are trend-followers, because of the costs for rational expectations. When $\gamma > R$, $C > 0$ and $\beta = \infty$ the fundamental steady state is *locally unstable*, but

globally stable, that is, after an unstable bubble phase solutions converge to the fundamental steady state. This implies existence of homoclinic orbits, that is, numerical solutions that converge to the fundamental steady state both forward and backwards in time. As stressed by Brock and Hommes (1997) the existence of homoclinic orbits implies complex, chaotic dynamics. The following theorem substantiates our numerical findings.

Theorem 1 (Homoclinic Orbit). *Let the intensity of choice $\beta = \infty$, costs $C > 0$, the trend-coefficient $\gamma > R$ and $\gamma \neq R^2$. Then the 2-type system has a homoclinic orbit.*

Proof. The profits of rational agents and trend-followers are given by

$$\pi_{1,t} = (x_t - Rx_{t-1})^2 - C \quad (29)$$

$$\pi_{2,t} = (x_t - Rx_{t-1})(\gamma x_{t-2} - Rx_{t-1}) \quad (30)$$

The switching depends on the sign of the profit difference:

$$\Delta\pi_t = \pi_{1,t} - \pi_{2,t} = (x_t - Rx_{t-1})(x_t - \gamma x_{t-2}) - C. \quad (31)$$

Since $\beta = \infty$, close to the fundamental steady state all agents will be trend-followers, because the first term of the RHS is close to zero and the information gathering costs for rationality dominate. When all agents are trend-followers the pricing equation simplifies to $x_t = (\gamma/R)x_{t-1}$. The fundamental steady state is unstable when $\gamma > R$, with the unstable eigenvector given by $((\gamma/R)^2, \gamma/R, 1)$. Take an initial state close to the fundamental steady state along the unstable eigenvector:

$$x_{-2} = \varepsilon, \quad x_{-1} = \varepsilon(\gamma/R), \quad x_t = \varepsilon(\gamma/R)^2. \quad (32)$$

As long as the fraction $n_{t,2} = 1$, the solution follows the unstable eigenvector and at time t , $x_t = \varepsilon(\gamma/R)^{t+2}$. To determine the point in time when agents switch to rational expectations, we need to look at the profit difference between rational expectations and the trend-following rule at date t :

$$\begin{aligned} \Delta\pi_t &= \left[\varepsilon \left(\frac{\gamma}{R}\right)^{t+2} - R\varepsilon \left(\frac{\gamma}{R}\right)^{t+1} \right] \left[\varepsilon \left(\frac{\gamma}{R}\right)^{t+2} - \gamma\varepsilon \left(\frac{\gamma}{R}\right)^t \right] - C \\ &= \varepsilon^2 \left(\frac{\gamma}{R}\right)^{t+1} \left(\frac{\gamma}{R} - R\right) \left(\frac{\gamma}{R}\right)^t \left[\left(\frac{\gamma}{R}\right)^2 - \gamma \right] - C \\ &= \varepsilon^2 \left(\frac{\gamma}{R}\right)^{2t+2} \left(\frac{\gamma - R^2}{R}\right)^2 - C. \end{aligned}$$

For $\gamma \neq R^2$ it holds that $\left(\frac{\gamma - R^2}{R}\right)^2 > 0$ while $\left(\frac{\gamma}{R}\right)^{2t+2}$ increases exponentially in t whenever $\gamma > R$.

This means that eventually $\Delta\pi_t > 0$ and, given $\beta = \infty$, all agents will then switch to rational expectations and the price will jump to the fundamental steady state. ■

Keeping Theorem 1 in mind, let us summarize the economic intuition behind the dynamics in Figures 8 - 9, thereby explicitly referring to the empirical findings of Greenwood and Shleifer (2014); Greenwood et al. (2020) and others. If price deviations from the fundamental are small, profits are low and the cost for the rational predictor is high relative to profits. This supports the rise of extrapolating forces, which are reinforced by the small fraction of rational agents. As prices are expected to rise, agents are willing to further indebt to finance asset purchases and credit will grow alongside with asset prices. As prices rise, the information costs (the costs of rationality) become negligible relative to profits. Rational agents however know that prices cannot continue to grow forever. Their foresight of the burst of the bubble in the future causes a small deviation of prices from the explosive trajectory. This however constitutes a relatively large loss in profits for extrapolating agents, which triggers a large fraction of agents to be willing to pay the costs of the rational predictor. An increase of the fraction of rational agents finally leads to the burst of the bubble.

6.2 Multiple equilibria and the role of steady-state anchoring

Let us have a closer look at values of β between 1.6 and 2.3, leaving all other parameters as before. For this range of parameters we detect multiple valid rational expectations equilibria with fundamentally different price dynamics. Equilibrium selection depends on whether rational expectations are anchored around the fundamental steady state ($x = 0$) or the non-fundamental steady state at x^* : if rational expectations are anchored around the fundamental steady state, we get similar dynamics as in subsection 6.1 with boom-burst dynamics in the intermediate region between 0 and x^* . When expectations are anchored at the nonfundamental steady state x^* , the nonfundamental steady state is stable and stationary.

In what follows we understand an *anchor* to long-term expectations as being the expected finite value of x when time goes to infinity. Technically, for our method \bar{x} is the initial guess at time t for *every* future state $x_{t+1}, x_{t+2}, \dots, x_{t+\infty}$. These guesses are alternated and corrected during the iterative procedure, and the iteration halts once adding $x_{t+K} = \bar{x}$ does not significantly change the evaluation of x_t (through x_{t+1}). That means the anchor \bar{x} is also the *finitely final* value of x when time approaches infinity, and all x_{t+s} , $s < K$ are chosen such that the trajectory leads in the neighborhood of \bar{x} .

Consider a value of $\beta = 2.14$ and note that for the given parameters it holds that $R < \gamma < (2R-1)$ and $m^* > \tanh(-\beta C/2)$. This implies that E_2 and E_3 are the two additional non-fundamental steady

states with $x^* \approx \pm 3.426$. In Figure 11 (left) we plot the time series of x_t for the chosen value of $\beta = 2.14$. The blue line depicts the dynamics from the bifurcation diagrams (Figures 7 and 8), for which the guesses on future expectations are initialized in the fundamental steady state, i.e. $\bar{x} = 0$. The price first approaches a value slightly below 2.4 and then shows erratic local boom-bust dynamics that are hard to predict. The orange line shows the same simulation exercise but initializing future expectations in the non-fundamental steady state x^* , i.e. $\bar{x} \equiv x^*$. Although this simulation is started from the same initial state X_0 , the trajectory converges to x^* , which is a locally stable steady state. Hence, the system is sensitive to the path of future expectations, which may not be unique. Anchoring future expectations around \bar{x} is a device of equilibrium selection which can heavily affect the global dynamics of f .




Figure 11: Left: time series for $\beta = 2.14$, all other parameters as in Table 3. Right: corresponding fractions of trend followers. For the blue line, the policy function is initialized in the unstable fundamental steady state. The policy function of the orange line is initialized at the locally stable non-fundamental steady state and exhibits no approximation errors. Both trajectories originate from the same initial conditions X_0 . Required minimal precision is $\epsilon = 1e - 8$.

To understand this feature, first note that both trajectories indeed constitute valid rational expectations equilibria. For both setups, the convergence criterion of the iterative procedure is reached and agents predict the future trajectory of prices with very high precision. After the time series is initialized with X_0 , prices grow with a growth rate $\phi = x_t/x_{t-1} \in (1, \gamma/R)$.¹⁷ At time $t = \bar{T}$, similarly to the chaotic example from the last subsection, rational agents believe that prices will not continue to rise beyond the threshold value around 2.4. This belief is self-fulfilling through the expectations feedback and prices drop. As a consequence, more agents will adapt the rational expectations forecast and prices will collapse further in period $t = \bar{T} + 1$, confirming the rational forecast in period $t = \bar{T}$. So far the dynamics were similar as in Subsection 6.1. The arising oscillating patterns are self-sustained and driven by the fact that the beliefs of rational agents are anchored below the threshold value close to 2.4.

¹⁷Note that it can be shown that for finite $\beta > 0$ and X_t , the growth rate of the boom-path (i.e., the growth rate of the trajectory that leads away from the fundamental steady-state) is not constant.

Consider now the case in which prices converge to x^* . As above, following X_0 the belief of rational agents that prices will continue to rise is also self-fulfilling. As rational expectations are anchored at x^* , the price will not collapse at $t = \bar{T}$ but increase further, finally converging to x^* . This non-fundamental steady state is locally stable because in x^* the weight n_t on the belief of trend-followers exactly cancels out extrapolating forces γ . It must hold that

$$\frac{1 - n^*}{R} + \frac{n^* \gamma}{R} = 1 \iff n^* = \frac{R - 1}{\gamma - 1}. \quad (33)$$

In summary, we find that not only are price beliefs self-fulfilling, but also the belief of the respective steady state is an important feature of a rational expectations equilibrium, and the anchoring of beliefs can have essential impact on the long-run dynamics of a system. Our exercise adds an important policy implication to the literature on expectation anchoring and multiple equilibria (see Section 1): if expectations are anchored around an unstable steady state, complicated dynamics can arise. If policy institutions can influence expectations, e.g. through forward guidance policies as conducted by the US Federal Reserve between 2008 and 2015, a successful re-anchoring of expectations to an alternative steady state has the power to stabilize the price dynamics.

7 Conclusion

In this work, we study the dynamics of a simple financial market model that is characterized by the coexistence of perfectly rational forward-looking and boundedly rational backward-looking agents. The relative share of each type of agent varies depending on their market performance, which gives rise to complicated endogenous bubble and crash dynamics. To find solutions to such a highly nonlinear heterogeneous agents model we develop and employ a novel and efficient iterative numerical method.

Rational agents anticipate the trend-extrapolating behavior of non-rational agents and therefore amplify instabilities caused by boundedly rational agents. Trend-extrapolators magnify small deviations from the fundamental price. The nontrivial interaction between rational forward-looking and non-rational backward-looking agents leads to bubble and crash dynamics with time-varying impact of rational and irrational behavior. As prices increase, rational agents anticipate and stabilize large bubbles, driving prices back to fundamental value.

Our heterogeneous agents model exhibits multiple equilibria which depend on the expected finite value when time goes to infinity. When future expectations are anchored on fundamental value asset price fluctuations exhibit bubble and crash patterns. In contrast, when finite expectations are anchored at a high non-fundamental steady state, this belief becomes self-fulfilling and prices are

larger than the fundamental price in the long run.

Our iterative method for numerically solving rational expectations equilibria in highly nonlinear heterogeneous agents models with rational and non-rational agents can be easily applied to other perfect foresight models in macroeconomics and finance.

References

- Aruoba, S.B., Cuba-Borda, P., Schorfheide, F., 2018. Macroeconomic dynamics near the zlb: A tale of two countries. *The Review of Economic Studies* 85, 87–118.
- Atkinson, T., Richter, A.W., Throckmorton, N.A., 2019. The zero lower bound and estimation accuracy. *Journal of Monetary Economics* .
- Ball, L., Mazumder, S., 2018. A Phillips Curve with Anchored Expectations and Short-Term Unemployment. *Journal of Money, Credit and Banking* 51, 111–137. doi:10.1111/jmcb.12502, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/jmcb.12502>.
- Barberis, N., Greenwood, R., Jin, L., Shleifer, A., 2018. Extrapolation and bubbles. *Journal of Financial Economics* 129, 203–227.
- Benhabib, J., Farmer, R.E., 1999. Indeterminacy and sunspots in macroeconomics. *Handbook of macroeconomics* 1, 387–448.
- Black, F., 1986. Noise. *The journal of finance* 41, 528–543.
- Blanchard, O.J., Kahn, C.M., 1980. The solution of linear difference models under rational expectations. *Econometrica: Journal of the Econometric Society* , 1305–1311.
- Blume, L., Easley, D., 2010. Heterogeneity, selection, and wealth dynamics. *Annu. Rev. Econ.* 2, 425–450.
- Bordalo, P., Gennaioli, N., Ma, Y., Shleifer, A., 2020a. Overreaction in macroeconomic expectations. *American Economic Review* 110, 2748–82.
- Bordalo, P., Gennaioli, N., Porta, R.L., Shleifer, A., 2020b. Expectations of fundamentals and stock market puzzles. Technical Report. National Bureau of Economic Research.
- Branch, W.A., McGough, B., 2018. Heterogeneous expectations and micro-foundations in macroeconomics, in: *Handbook of computational economics*. Elsevier. volume 4, pp. 3–62.
- Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. *Econometrica* , 1059–1095.

- Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control* 22, 1235–1274.
- Cao, D., 2018. Speculation and financial wealth distribution under belief heterogeneity. *The Economic Journal* 128, 2258–2281.
- Carroll, C.D., 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters* 91, 312–320.
- Coibion, O., Gorodnichenko, Y., 2015. Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation. *American Economic Journal: Macroeconomics* 7, 197–232. URL: <https://ideas.repec.org/a/aea/aejmac/v7y2015i1p197-232.html>.
- Coleman, W.J., 1990. Solving the stochastic growth model by policy-function iteration. *Journal of Business & Economic Statistics* 8, 27–29.
- Coleman, W.J., 1991. Equilibrium in a production economy with an income tax. *Econometrica: Journal of the Econometric Society* , 1091–1104.
- Cootner, P.H., 1964. The random character of stock market prices .
- De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990a. Noise trader risk in financial markets. *Journal of political Economy* 98, 703–738.
- De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990b. Positive feedback investment strategies and destabilizing rational speculation. *the Journal of Finance* 45, 379–395.
- Evans, G.W., Honkapohja, S., 2003. Expectations and the stability problem for optimal monetary policies. *The Review of Economic Studies* 70, 807–824.
- Fair, R.C., Taylor, J.B., 1983. Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models. *Econometrica: Journal of the Econometric Society* , 1169–1185.
- Friedman, M., 1953. *Essays in positive economics*. University of Chicago Press, Chicago. URL: <http://www.worldcat.org/oclc/38499088>.
- Gali, J., 2014. Monetary policy and rational asset price bubbles. *American Economic Review* 104, 721–52.
- Gerstner, T., Griebel, M., 1998. Numerical integration using sparse grids. *Numerical algorithms* 18, 209.

- Grandmont, J.M., 1985. On endogenous competitive business cycles. *Econometrica: Journal of the Econometric Society* , 995–1045.
- Grandmont, J.M., 1998. Expectations formation and stability of large socioeconomic systems. *Econometrica* , 741–781.
- Greenwood, R., Hanson, S.G., 2013. Issuer quality and corporate bond returns. *The Review of Financial Studies* 26, 1483–1525.
- Greenwood, R., Hanson, S.G., Shleifer, A., Sørensen, J.A., 2020. Predictable Financial Crises. Technical Report. Harvard Business School Working Paper.
- Greenwood, R., Shleifer, A., 2014. Expectations of returns and expected returns. *Review of Financial Studies* 27, 714–746.
- Gust, C., Herbst, E., López-Salido, D., Smith, M.E., 2017. The empirical implications of the interest-rate lower bound. *American Economic Review* 107, 1971–2006.
- Hommes, C., LeBaron, B., 2018. *Computational economics: heterogeneous agent modeling*. Elsevier.
- Hommes, C., Sonnemans, J., Tuinstra, J., Van de Velden, H., 2005. Coordination of expectations in asset pricing experiments. *The Review of Financial Studies* 18, 955–980.
- Hommes, C.H., 2006. Heterogeneous agent models in economics and finance. *Handbook of computational economics* 2, 1109–1186.
- Hommes, C.H., 2021. Behavioral & experimental macroeconomics and policy analysis: a complex systems approach. *Journal of economic literature* (forthcoming) .
- Judd, K.L., 1998. *Numerical methods in economics*. MIT press.
- Judd, K.L., Maliar, L., Maliar, S., Valero, R., 2014. Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain. *Journal of Economic Dynamics and Control* 44, 92–123.
- Klein, P., 2000. Using the generalized schur form to solve a multivariate linear rational expectations model. *Journal of economic dynamics and control* 24, 1405–1423.
- Kyle, A.S., 1985. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society* , 1315–1335.
- Ljungqvist, L., Sargent, T.J., 2012. *Recursive macroeconomic theory*. MIT press.

- López-Salido, D., Stein, J.C., Zakrajšek, E., 2017. Credit-market sentiment and the business cycle. *The Quarterly Journal of Economics* 132, 1373–1426.
- Mertens, K.R., Ravn, M.O., 2014. Fiscal policy in an expectations-driven liquidity trap. *The Review of Economic Studies* 81, 1637–1667.
- Miranda, M.J., Fackler, P.L., 2004. *Applied computational economics and finance*. MIT press.
- Rendahl, P., 2017. Linear time iteration. Technical Report. IHS Economics Series.
- Richter, A.W., Throckmorton, N.A., 2016. Are nonlinear methods necessary at the zero lower bound? .
- Sandroni, A., 2000. Do markets favor agents able to make accurate predictions? *Econometrica* 68, 1303–1341.
- Schularick, M., Taylor, A.M., 2012. Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008. *American Economic Review* 102, 1029–61.
- Sims, C.A., 2003. Implications of rational inattention. *Journal of monetary Economics* 50, 665–690.
- Sims, C.A., 2010. Rational inattention and monetary economics, in: *Handbook of monetary economics*. Elsevier. volume 3, pp. 155–181.
- Woodford, M., 2005. Central bank communication and policy effectiveness. Technical Report. National Bureau of Economic Research.

Appendix A Alternative solution method: policy function iteration

Additionally to the method outline in Section 3 we use a computational procedure that is based on the method of policy function iteration (Coleman, 1990, 1991) to solve for the rational expectation equilibrium.

Recall from Section 3 that our problem is to find a function g that satisfies

$$g(x_{t-1}, x_{t-2}, x_{t-3}) = f(g(x_{t-1}, x_{t-2}, x_{t-3}), x_{t-1}, x_{t-2}), x_{t-1}, x_{t-2}, x_{t-3}). \quad (\text{A.1})$$

Define $Y = (y_1, y_2, \dots, y_m)$ the vector of m grid points for each dimension of f and let $\mathfrak{Y} = Y \times Y \times Y$ be the grid on which g is defined. Hence, g resides on a cube in $\mathbb{R}^{\{m \times m \times m\}}$. Initialize the procedure with g_0 . Then, given the k th iterate g_k we can use (14) to find the $(k + 1)$ -iterate g_{k+1} :

$$\left\{ g_{k+1}(z_1, z_2, z_3) = f(g_k(g_k(z_1, z_2, z_3), z_1, z_2), z_1, z_2, z_3) \mid z_i \in Y \text{ for every } i \in \{1, 2, 3\} \right\}. \quad (\text{A.2})$$

If f is well behaved, i.e. f is continuous with bounded derivative on its domain, a reasonably chosen norm of $(g_{k+1} - g_k)$ converges to zero when k goes to infinity. Halt the iteration once $\|g_{k+1} - g_k\| < \varepsilon$ for some sufficiently small ε . The limiting function g_K then approximates the true state-space representation g and can be used to simulate the global dynamics. However, in this paper we consider highly nonlinear maps and the dynamics of f may be unstable for some regions in its domain \mathfrak{D} . E.g. for some vectors of initial states, f may exhibit only explosive solutions and the iterative procedure will diverge for these vectors. Such instabilities are indeed quite common for the functions we consider due to the possible existence of infinite bubble solutions. Divergence for some regions of the state space does however not rule out the existence of bounded solutions for other regions.

Given that an overall measure of convergence of the complete policy function may fail to display local convergence for some regions of the grid, we require a convergence measure that is insensitive to this problem. We are interested in the (long run) dynamics for the relevant domain given a vector of initial values $X_0 = (x_0, x_{-1}, x_{-2})$ and aim to find the solution given these initial values. Given (14), a solution for $x_1 = g(X_0)$ also implies finding a solution for $x_2 = g(X_1)$ as $g(X_0) = f(x_2, X_0)$. Recursively, finding a solution for $x_1 = g(X_0)$ hence implies solving for the complete *relevant* future path of x , given the initial state X_0 . That means that a sufficient *local* convergence criterion is

$$|g_{k+1}(X_0) - g_k(X_0)| \leq \varepsilon, \tag{A.3}$$

for a very small ε .¹⁸ We provide an assessment of accuracy and numerical stability based on average and maximum Euler equation errors.

A second important practical matter is the choice of grid points. A growing literature treats sophisticated methods for the choice of an appropriate grid (see e.g. Gerstner and Griebel, 1998; Judd et al., 2014). However, these methods are not particularly useful for our case as sometimes, some variables are highly cross-dependent, and sometimes they are not. This makes it hard to assign a higher density of grid points to some regions of the grid. It is also hard to ex-ante identify the necessary support of the grid because the amplitude of potential cycles may depend on the parameter choice. Endogenous grid methods (e.g. Carroll, 2006) are not particularly useful because the strong nonlinearities in our model impose a challenge on the convergence of the optimal choice of grid points themselves. Lastly, we can not use grids that put more weight to some regions of the state space

¹⁸For all numerical experiments in this paper we chose $\varepsilon = 1e - 10$ and $X_0 = (0.1, 0.2, 0)$.

than others for two reasons. First, because we can not identify these regions ex-ante.¹⁹ Second, even if we can ex-ante identify regions that are in particular relevant for the trajectory following X_0 , it might still be necessary to have a sufficiently dense grid in regions that are not actually close to this trajectory, as these regions may be important for convergence of the iterative procedure.

As a result, we only consider Cartesian grids. We initialize the grid with a comfortable support (y_0, y_m) , iterate on it for a small number of times and use the last iteration to simulate the model. We then shrink the grid to fit the bounds of the simulated time series plus some leeway.²⁰ We repeat this procedure until $(\min\{g_k(X_0)\}_0^T, \max\{g_k(X_0)\}_0^T) \approx (y_0, y_m)$. For the bifurcation diagrams in the following section we chose $m = 100$. This means for every parameter value considered in the bifurcation diagram we must find g which is represented on a grid of one million vectors. As all the examples we consider have a steady state at zero, we initialize the map $g_0 \equiv 0$ with zeroes. As we will discuss in Section 6, the choice to initialize g_0 at a certain steady state focal point can have implications for convergence and steady state selection.

Appendix B Comparison of the method from section 3 with policy function iteration

To check the accuracy of the policy function iteration algorithm, we use the average and maximum absolute Euler equation errors of the time series. Absolute Euler equation errors (EEE) are defined as the absolute value of the difference between the functional approximation g and the function f , where g is used to solve for the expectations x_{t+1}^e . This is equivalent to the calculation of approximation errors ζ_t in the main body. Hence:

$$EEE_{\text{average}} = \frac{1}{T} \sum_{t=0}^T |g(x_{t-1}, x_{t-2}, x_{t-3}) - f(g(\cdot), x_{t-1}, x_{t-2}, x_{t-3})| \quad (\text{B.1})$$

and

$$EEE_{\text{max}} = \max \left\{ \{|g(x_{t-1}, x_{t-2}, x_{t-3}) - f(g(\cdot), x_{t-1}, x_{t-2}, x_{t-3})|\}_{t=0}^T \right\} \quad (\text{B.2})$$

where T is the length of the simulated series after the convergence period.

For Figures 2 to 6 from the 3-type model with cycles, the average and maximum EEE are both very small. The time series presented in Figure 4 come with an average EEE of the magnitude of $1e - 5$ and a maximum EEE of $1e - 4$. This is documented in Figures B.12-B.14. The bifurcation maps and associated time series are qualitatively and quantitatively very similar to those presented

¹⁹For example, the grid that represents a limit cycle would need much detail close to the bounds of the limit cycle, less detail in the center.

²⁰Leeway is necessary because when approximated on a more fine-grained grid, g may again map beyond the bounds of the previous simulated series. This in fact happens quite frequently.

in Section 5, where we use the method we describe in Section 3.




Figure B.12: Average and maximum absolute Euler equation errors for the simulations in Figure 2 when using policy function iteration.




Figure B.13: Average and maximum absolute Euler equation errors for the simulations in Figure 3 when using policy function iteration.




Figure B.14: Average and maximum absolute Euler equation errors for the simulations in Figure 6 when using policy function iteration.

For the two-type example in (6), with more complex dynamics, non-negligible errors occur, which are documented in Figures B.16 and B.17. In fact, the policy function iteration method is unable to correctly solve for the dynamics of the intermediate region of β that we discuss in Subsection 6.2. Figure B.15 shows the same bifurcation diagram as Figure 8 from the main body, but using the policy function iteration algorithm.

Figure B.19 shows the errors over time for the time series from Figure 9, but using the policy

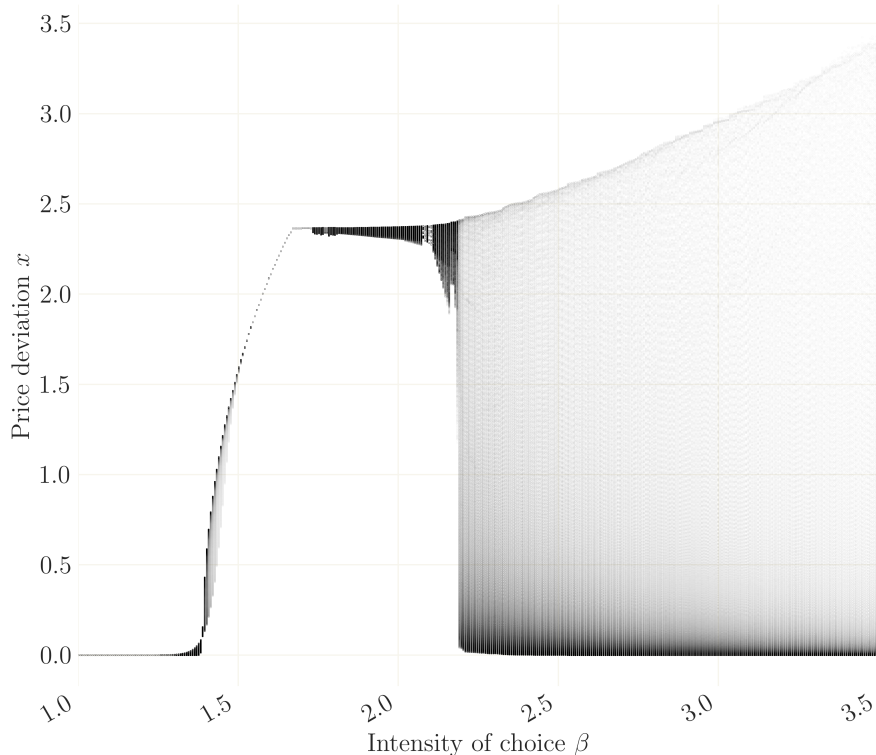


Figure B.15: Bifurcations w.r.t. β solved using policy function iteration. All other parameters as given in Table 3.




Figure B.16: Average and maximum absolute Euler equation errors for the simulations in Figure 8 when using policy function iteration. This plot shows that the numerical method is working well for the chaotic bubble and crash solutions. However, our numerical method is unable to well-approximate the policy function in the intermediate parameter region, where the system is close to the bifurcations of (stable) non-fundamental steady states.

function iteration method. Although the dynamics display complicated bubble-crash dynamics, errors are at a very small level and the time series match those from the main body. In contrast, Figure B.20 shows the simulations for $\beta = 2.14$ using policy function iteration to solve for the dynamics. Large errors occur and the policy function iteration algorithm is unable to solve for the boom-burst dynamics displayed in Figure B.20. The dashed red line draws $f(g(\cdot), X_{t-1})$, i.e the




Figure B.17: Average and maximum absolute Euler equation errors for the bifurcation w.r.t. the cost parameter C as in Figure 7 but when using policy function iteration.




Figure B.18: Average and maximum absolute Euler equation errors for the bifurcation w.r.t. the trend-following parameter γ as in Figure 7 but when using policy function iteration.

value of x_t as implied by f while using the policy function to approximate x_{t+1} . Indeed, while the policy function predicts the erratic waves quite well, the prices that are implied by using f directly is at times slightly lagging behind. The policy function iteration method correctly solves for the timing of a collapse, but the relatively large approximation error stems from the failure to predict the exact value of x during the course of the collapse. We explain this by the fact that, compared to the limit cycles in Section 5, the cycles only span over a relatively small domain of the whole grid.

Overall, in terms of approximation errors our method introduced in Section 3 performs several magnitudes better than policy function iteration.




Figure B.19: Top: time series for $\beta = 2.5$ solved using policy function iteration, all other parameters as in Table 3. The policy function is initialized with the unstable fundamental steady state $g_0 = 0$. Bottom: corresponding absolute Euler equation errors, which are the absolute value of the difference between the blue and the dashed red line.

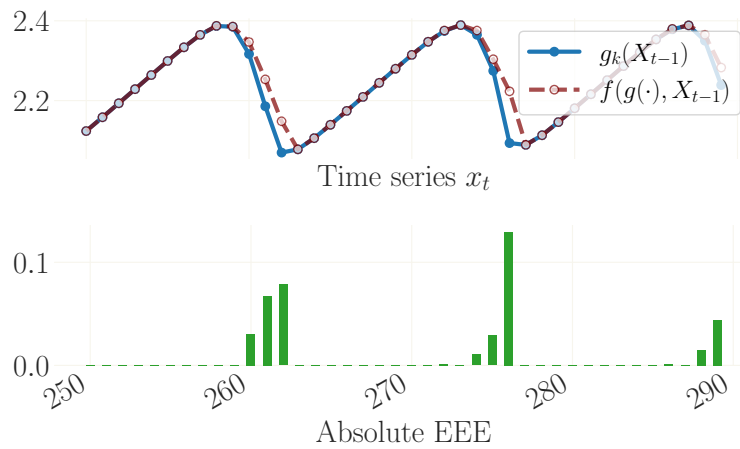


Figure B.20: Top: time series for $\beta = 2.14$ solved using policy function iteration, all other parameters as in Table 3. The policy function is initialized with the unstable fundamental steady state $g_0 = 0$. Bottom: corresponding absolute Euler equation errors, which are the absolute value of the difference between the blue and the dashed red line.

Appendix C Phase plots




Figure C.21: Phase plot in $(x_t, n_{1,t})$ -space for different values of β corresponding to the bifurcation diagram with limit cycles in Figure 3 (based on policy function iteration).




Figure C.22: Phase plot in $(x_t, n_{2,t})$ -space for different values of β corresponding to the bifurcation diagram with limit cycles in Figure 3 (based on policy function iteration).




Figure C.23: Phase plot in (x_{t-1}, x_t) -space for different values of β corresponding to the bifurcation diagram with limit cycles in Figure 3 (based on policy function iteration).




Figure C.24: Phase plot in (x_t, n_t) -space for different values of β corresponding to the bifurcation diagram with chaotic dynamics in Figure 8 (based on policy function iteration).




Figure C.25: Phase plot in (x_{t-1}, x_t) -space for different values of β corresponding to the bifurcation diagram with chaotic dynamics in Figure 8 (based on policy function iteration).

Appendix D Grapical illustrations of the policy functions

Figure D.26 represents slices of the policy function g from the 3-type example in Figure 3. The function is smooth for $\beta = 1.1$ (upper left panel) and, as reflected by an increase in contrast in the figure, becomes steeper as β increases. For $\beta = 7$ in the last diagram the function is very steep, in particular in the center. This coincides with the region for β in which the dynamics display a stable 4-cycle with values only in the periphery.




Figure D.26: Heatmap for the policy function g for $\beta = 1.1$ (ul), $\beta = 2.5$ (ur), $\beta = 4$ (ll) and $\beta = 7$ (lr), all other parameters as in Table 2. When β increases the function becomes steeper. At the diagonal a small change in x_{t-1} or x_{t-2} leads to a large change in $x_t = g(x_{t-1}, x_{t-2})$. The illustration is based on policy function iteration.

Figure 10 shows slices of the policy function for the chaotic example in Figure 9. The Figures D.27 and D.28 show the policy functions for Figure 11, initialized with different steady states.




Figure D.27: Heatmap for the policy function g for the time series in Figure 11 initialized in the unstable fundamental steady state ($g_0 = 0$). The illustration is based on policy function iteration.




Figure D.28: Heatmap for the policy function g for the time series in Figure 11 initialized in the locally stable non-fundamental steady state ($g_0 = x^*$). The illustration is based on policy function iteration.