Gomez, Thomas; Piccillo, Giulia

Working Paper

Diverse Risk Preferences and Heterogeneous Expectations in an Asset Pricing Model

CESifo Working Paper, No. 8003

Provided in Cooperation with:
Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Gomez, Thomas; Piccillo, Giulia (2019) : Diverse Risk Preferences and Heterogeneous Expectations in an Asset Pricing Model, CESifo Working Paper, No. 8003, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:
https://hdl.handle.net/10419/215005

CESifo WORKING PAPERS

8003
2019

December 2019

Diverse Risk Preferences and Heterogeneous Expectations in an Asset Pricing Model

Thomas Gomez, Giulia Piccillo

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 8003*

Diverse Risk Preferences and Heterogeneous Expectations in an Asset Pricing Model

Abstract

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

JEL-Codes: D810, D840, G110, G120, G410.

Keywords: heterogeneous risk aversion, bounded rationality, heterogeneous expectations, heuristic switching, asset pricing.

Thomas Gomez
Utrecht University / The Netherlands
t.gomez@uu.nl

Giulia Piccillo
Maastricht University / The Netherlands
g.piccillo@maastrichtuniversity.nl

8 October 2019
This work was supported by the Complex Systems Fund, with special thanks to Peter Koeze.

1. Introduction

From historical examples like the Tulip Mania to more recent ones like the US housing bubble, it seems that bubbles are a recurring characteristic of asset markets. While it is hard to identify bubbles in real markets because fundamental values cannot be observed, experimental asset markets, with known fundamentals, can also exhibit price bubbles and crashes (e.g., Smith et al., 1988). It is important to understand and model the mechanisms underlying these price dynamics, in order to provide policymakers with the necessary tools to promote stability in real financial markets.

Asset price bubbles cannot always be explained by rational, speculative motives. Lei et al. (2001) show that bubbles can emerge even when speculation is not possible, suggesting that the behaviour is inherently irrational. Hommes et al. (2005, 2008) specifically look at irrationality in expectations. In their learning-to-forecast experiments, subjects are rewarded for accurately forecasting future price realizations of a risky asset. At the individual level, their results indicate that the subjects use different, simple rules to predict future asset prices. At the aggregate level, they find that the interaction between the forecasting rules used by the subjects can lead to different asset price dynamics in the same experimental setting. These dynamics include slow convergence to, and significant deviations from the fundamental.

Anufriev and Hommes (2012) show that these different price dynamics can be explained by modelling expectations with the heuristic switching framework introduced by Brock and Hommes (1997). In line with the learning-to-forecast results, this framework assumes that individuals choose between a number of simple forecasting rules (heuristics) to form their expectations. They base their choice on the past performance of those rules. As the performance changes over time, the individuals update their choice, switching between heuristics. The heuristic switching framework is in line with a large literature on boundedly rational, heterogeneous expectations, where heterogeneity changes over time. Mankiw et al. (2003), for example, show that consumers as well as professional forecasters have diverse inflation expectations, and Wieland and Wolters (2010) find time-varying

heterogeneity in model forecasts. Heuristic switching models are able to explain certain stylized facts of financial markets, like excess volatility and volatility clustering (Anufriev and Panchenko, 2009; Gaunersdorfer et al., 2008). These models have also been applied in macroeconomics, investigating the effects of monetary policy (e.g., Hommes et al., 2019), fiscal policy (e.g., De Grauwe et al., 2019), and the financial sector (De Grauwe and Macchiarelli, 2015). Branch (2004) estimates a heuristic switching model using inflation forecasting data. His results indicate that consumers do indeed switch between forecasting strategies based on past performance.

In the context of this heuristic switching literature, our main contribution is to study the natural consequence of the standard asset pricing assumption that agents are mean-variance optimizers.[1] Specifically, they make a trade-off between forecasting performance and variability of that performance when choosing a heuristic. This means that risk aversion plays a key role in expectation formation: Agents with different risk preferences might choose a different heuristic, because their risk aversion determines the importance of performance variability for their choice. As the heuristics' performance changes over time, agents reconsider their choice. This can lead them to switch between forecasting rules. In this way, diverse risk preferences can lead to time-varying heterogeneity of expectations. We model this diversity by using risk preference estimates for the general population provided by Kimball et al. (2008) and Aarbu and Schroyen (2009).[2]

We argue that diverse risk preferences are a natural source of heterogeneous expectations for two reasons.[3] First, risk preferences play a central role in economic theory and have been linked empirically to a variety of outcomes related to, for example, career choice, financial decision making, and migration (Dohmen et al., 2012; Falk et al., 2018; Jaeger et al., 2010). Second, risk preferences are indeed heterogeneous, as shown

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

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

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

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

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

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

price shocks. We find that small price shocks can be amplified by the chaotic dynamics, turning deterministically stable solutions into explosive ones.

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

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

Second, Chiarella and He (2002, 2003) study a heuristic switching model where risk aversion is again linked to the chosen heuristic, but the risk aversion parameter for a given heuristic is constant. They study this model, which also incorporates learning, in two institutional settings: Chiarella and He (2002) use a Walrasian scenario, while Chiarella and He (2003) use a market-maker. They find that the diversity of risk preferences

matters for the asset pricing dynamics. Like Park (2014), they do not take into account the implications of diverse risk preferences for expectation formation.

2. The model

2.1. An asset market with heuristic switching mean-variance optimizers

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

Agents differ only in their constant relative risk aversion (CRRA) parameter γ, drawn from a distribution defined on $(0, \infty)$. We assume a population size normalized to 1 and denote the density function by g, with cumulative distribution function G. Focusing on relative risk aversion as it appears in the CRRA utility function allows us to use the empirical estimates of its distribution that can be found in the literature (see subsection 2.3).

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

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

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

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

conditions.[4]

We have that[5]

$$\widetilde{E}_{\gamma,t}[W_{\gamma,t+1}] = z_{\gamma,t}\,\widetilde{E}_{\gamma,t}[R^e_{t+1}] + R^f\,W_{\gamma,t} \tag{2a}$$

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

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

$$z_{\gamma,t} = \frac{\widetilde{E}_{\gamma,t}\left[R^e_{t+1}\right]}{\gamma\,\widetilde{V}_{\gamma,t}\left[R^e_{t+1}\right]}. \tag{3}$$

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

$$Z_t = \int_0^\infty \frac{\widetilde{E}_{\gamma,t}\left[R^e_{t+1}\right]}{\gamma\,\widetilde{V}_{\gamma,t}\left[R^e_{t+1}\right]}\,g(\gamma)\,d\gamma. \tag{4}$$

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

Equating aggregate demand and supply, $Z_t = 0$, gives an expression for the risky

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

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

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

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

asset price in period t:

$$p_t = \frac{1}{R^f \, \Theta} \int_0^\infty \frac{g(\gamma)}{\gamma} \, \widetilde{E}_{\gamma,t} \left[p_{t+1} + d_{t+1} \right] \, d\gamma, \tag{5}$$

where we have defined the aggregate risk tolerance

$$\Theta = \int_0^\infty \frac{g(\gamma)}{\gamma} \, d\gamma. \tag{6}$$

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

The fundamental price for the risky asset, denoted by p_t^*, is the price that would arise in a homogeneous, rational market. We denote expectation and variance, conditional on information set $\mathbb{I}_t$, by E_t and V_t. Equation (5) implies that

$$p_t^* = \frac{E_t \left[p_{t+1}^* + d_{t+1} \right]}{R^f}. \tag{7}$$

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

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

$$\widetilde{E}_{\gamma,t} \left[p_{t+1}^* + d_{t+1} \right] = E_t \left[p_{t+1}^* + d_{t+1} \right]. \tag{8}$$

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

deviations from the fundamental.

Second, the agents form beliefs about future price deviations from the fundamental by using a forecasting heuristic. These heuristics are simple rules that predict price deviations based on a specific number of observed previous deviations. Formally, the heuristics are indexed by the finite set $\mathcal{I}$, and for $s \in \mathcal{I}$, we define a heuristic as a function $h_s : \mathbb{R}^L \to \mathbb{R}$, $(x_{t-1}, \ldots, x_{t-L}) \mapsto h_s(x_{t-1}, \ldots, x_{t-L})$, where $L \in \mathbb{N}$ indicates the number of lags of price deviations that are taken into account.[6] In period t, agents of type γ use heuristic $s_{\gamma,t} \in \mathcal{I}$, and their beliefs about price deviations from the fundamental are given by

$$\widetilde{E}_{\gamma,t}\left[x_{t+1}\right] = h_{s_{\gamma,t}}\left(x_{t-1}, \ldots, x_{t-L}\right). \tag{9}$$

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

Together, equations (8) and (9) pin down the agents' beliefs about the gross return per risky asset: $\widetilde{E}_{\gamma,t}\left[p_{t+1} + d_{t+1}\right] = \widetilde{E}_{\gamma,t}\left[p^*_{t+1} + d_{t+1}\right] + \widetilde{E}_{\gamma,t}\left[x_{t+1}\right]$. This means that we can solve for the price deviations from the fundamental using equations (5) and (7):

$$x_t = \frac{1}{R^f\,\Theta} \int_0^\infty \frac{g(\gamma)}{\gamma}\, h_{s_{\gamma,t}}\left(x_{t-1}, \ldots, x_{t-L}\right)\, d\gamma. \tag{10}$$

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

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

between risky asset demand and risk aversion (see equation 3).

We define $\Gamma_{s,t}$ as the set of risk aversion rates for which agents use heuristic s in period t, and $\Theta_{s,t}$ as the fraction of aggregate risk tolerance (defined in equation 6) represented by agents that use heuristic s in period t:

$$\Theta_{s,t} = \Theta^{-1} \int_{\Gamma_{s,t}} \frac{g(\gamma)}{\gamma}\, d\gamma. \tag{11}$$

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

Now we can rewrite equation (10) as

$$x_t = \frac{1}{R^f} \sum_{s\in\mathcal{I}} \Theta_{s,t}\, h_s\left(x_{t-1}, \ldots, x_{t-L}\right). \tag{12}$$

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

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

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

$$u_{s,t} = -\left[x_t - h_s(x_{t-2}, \ldots, x_{t-L-1})\right]^2. \tag{13}$$

They calculate a weighted average $\langle u_s \rangle_t$, a weighted squared average $\langle u_s^2 \rangle_t$, and a weighted

variance $\widetilde{\sigma}^2_{s,t}$ of the observed performance of heuristic s, where past performance is weighted by a memory parameter η that satisfies $0 < \eta < 1$:[7]

$$\langle u_s \rangle_t = \eta \langle u_s \rangle_{t-1} + (1-\eta)\, u_{s,t} \tag{14a}$$

$$\left\langle u_s^2 \right\rangle_t = \eta \left\langle u_s^2 \right\rangle_{t-1} + (1-\eta)\, u_{s,t}^2 \tag{14b}$$

$$\widetilde{\sigma}^2_{s,t} = \left\langle u_s^2 \right\rangle_t - \langle u_s \rangle_t^2. \tag{14c}$$

The weighted variance satisfies $\widetilde{\sigma}^2_{s,t} \geq 0$ if $\widetilde{\sigma}^2_{s,0} \geq 0$.[8] The average and variance are combined in a mean-variance performance measure, defined as

$$\Psi_{\gamma,s,t} = \langle u_s \rangle_t - \frac{\gamma}{2} \widetilde{\sigma}^2_{s,t}. \tag{15}$$

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

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

2.2. Fundamentalists and momentum traders

Following previous literature (Anufriev and Panchenko, 2009; Gaunersdorfer et al., 2008), and in line with empirical evidence (e.g., Chiarella et al., 2014), we consider the

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

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

case in which agents choose between a fundamentalist and a trend-following (chartist) rule. We focus on just two rules to keep the model simple, but one could extend it to include more rules.

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

$$h_F(x_{t-1}) = f\, x_{t-1}. \tag{16}$$

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

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

$$h_M(x_{t-1}, x_{t-2}) = x_{t-1} + m\,(x_{t-1} - x_{t-2})\,. \tag{17}$$

The larger m, the stronger the past momentum will be extrapolated into future deviations from the fundamental.[9]

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

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

Looking at the definition of the mean-variance performance measure for choosing between the heuristics (15), we see that the choice in period t is independent of risk aversion if the variances of the rules' performance are equal ($\widetilde{\sigma}_{F,t} = \widetilde{\sigma}_{M,t}$). This means that all agents use the same rule in period $t+1$: If the fundamentalist rule has performed better on average ($\langle u_F \rangle_t > \langle u_M \rangle_t$), everyone uses the fundamentalist rule in period $t+1$ and $\Theta_{F,t+1} = 1$. On the other hand, if the momentum rule has performed better on average, we have that $\Theta_{F,t+1} = 0$. In the case that the heuristics have performed equally well on average, the agents are evenly divided between the two heuristics in period $t+1$, and $\Theta_{F,t+1} = \frac{1}{2}$.

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

$$\overline{\gamma}_t = 2\,\frac{\langle u_F \rangle_t - \langle u_M \rangle_t}{\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t}}. \tag{18}$$

If $\widetilde{\sigma}_{F,t} > \widetilde{\sigma}_{M,t}$, agents of type $\gamma < \overline{\gamma}_t$ choose the fundamentalist heuristic for period $t+1$. Only agents with a low enough risk aversion are willing to accept the larger variance in order to profit from the better average performance. When $\langle u_F \rangle_t < \langle u_M \rangle_t$ in this case, $\overline{\gamma}_t < 0$, and all agents use the momentum rule in period $t+1$: When the risk is bigger and average performance worse, none of the agents use the fundamentalist heuristic. If $\widetilde{\sigma}_{F,t-1} > \widetilde{\sigma}_{M,t-1}$ and $\langle u_F \rangle_{t-1} \geq \langle u_M \rangle_{t-1}$, $\Theta_{F,t}$ is given by equation (11), with $s = F$ and $\Gamma_{F,t} = (0, \overline{\gamma}_{t-1})$.

The opposite holds when $\widetilde{\sigma}_{F,t-1} < \widetilde{\sigma}_{M,t-1}$: Agents of type $\gamma > \overline{\gamma}_{t-1}$ choose the fundamentalist heuristic for period t, so that $\Theta_{F,t} = 1$ when $\langle u_F \rangle_{t-1} > \langle u_M \rangle_{t-1}$, and $\Theta_{F,t}$ is given by equation (11), with $s = F$ and $\Gamma_{F,t} = (\overline{\gamma}_{t-1}, 0)$ if $\langle u_F \rangle_{t-1} \leq \langle u_M \rangle_{t-1}$.

We now have that

$$\Theta_{F,t+1} = \begin{cases} \Theta^{-1} \int_0^{\overline{\gamma}_t} \frac{g(\gamma)}{\gamma}\, d\gamma & \text{if } \widetilde{\sigma}_{F,t} > \widetilde{\sigma}_{M,t} \text{ and } \langle u_F \rangle_t \geq \langle u_M \rangle_t; \\ \Theta^{-1} \int_{\overline{\gamma}_t}^{\infty} \frac{g(\gamma)}{\gamma}\, d\gamma & \text{if } \widetilde{\sigma}_{F,t} < \widetilde{\sigma}_{M,t} \text{ and } \langle u_F \rangle_t \leq \langle u_M \rangle_t; \\ 1 & \text{if } \widetilde{\sigma}_{F,t} \leq \widetilde{\sigma}_{M,t} \text{ and } \langle u_F \rangle_t > \langle u_M \rangle_t; \\ 0 & \text{if } \widetilde{\sigma}_{F,t} \geq \widetilde{\sigma}_{M,t} \text{ and } \langle u_F \rangle_t < \langle u_M \rangle_t; \\ \frac{1}{2} & \text{if } \widetilde{\sigma}_{F,t} = \widetilde{\sigma}_{M,t} \text{ and } \langle u_F \rangle_t = \langle u_M \rangle_t. \end{cases} \tag{19}$$

The function $\Theta_{F,t+1}$ is discontinuous at points where $\widetilde{\sigma}_{F,t} = \widetilde{\sigma}_{M,t}$ and $\langle u_F \rangle_t = \langle u_M \rangle_t$: If we start in the $\langle u_F \rangle_t > \langle u_M \rangle_t$ area for example, and move towards the $\langle u_F \rangle_t < \langle u_M \rangle_t$ area while keeping $\widetilde{\sigma}_{F,t} = \widetilde{\sigma}_{M,t}$, $\Theta_{F,t+1}$ jumps from 1 to 1/2 to 0 as we pass the $\langle u_F \rangle_t = \langle u_M \rangle_t$ point. This discontinuity has consequences for the analysis of the dynamics that we discuss in subsection 3.2.

Combining equations (12) and (14), the dynamical system with fundamentalists and momentum traders is given by

$$\begin{cases} x_t = \dfrac{1}{R^f} \left\{ \Theta_{F,t}\, f\, x_{t-1} + (1 - \Theta_{F,t}) \left[x_{t-1} + m \left(x_{t-1} - x_{t-2} \right) \right] \right\} \\ \langle u_F \rangle_t = \eta \langle u_F \rangle_{t-1} - (1 - \eta) \left(x_t - f\, x_{t-2} \right)^2 \\ \langle u_M \rangle_t = \eta \langle u_M \rangle_{t-1} - (1 - \eta) \left[x_t - x_{t-2} - m \left(x_{t-2} - x_{t-3} \right) \right]^2 \\ \left\langle u_F^2 \right\rangle_t = \eta \left\langle u_F^2 \right\rangle_{t-1} + (1 - \eta) \left(x_t - f\, x_{t-2} \right)^4 \\ \left\langle u_M^2 \right\rangle_t = \eta \left\langle u_M^2 \right\rangle_{t-1} + (1 - \eta) \left[x_t - x_{t-2} - m \left(x_{t-2} - x_{t-3} \right) \right]^4, \end{cases} \tag{20}$$

where the risk-tolerance fractions $\Theta_{F,t}$ are given by equation (19).

2.3. Empirical grounding

To empirically ground our model, we draw from the rich literature on risk aversion, which includes studies on experiments, surveys, and decision-making in the field. Kimball et al. (2008) estimate the CRRA parameter for risk tolerance θ. They determine the distribution of risk tolerance by using hypothetical income gambles by respondents to the US Health and Retirement Study.[10] Kimball et al. find that a log-normal distribution fits the individual-level risk tolerance data well, because it has a fat right tail and imposes non-negative risk aversion. An additional advantage is that this distribution is computationally efficient. They estimate that $\log \theta \sim \mathcal{N}(-1.84, 0.73)$. This implies that $\log \gamma \sim \mathcal{N}(1.84, 0.73)$ and that the relative rate of risk aversion has mean 8.22 and standard deviation 6.90.

A limitation of Kimball et al.'s estimated distribution is that their respondents are all between 51 and 61 years of age. Aarbu and Schroyen (2009) take a similar approach, but use Norwegian survey data that covers ages 18 – 74. They find a lower average rate of risk aversion of 3.92, which is in line with earlier studies, with standard deviation 2.95. Figure 1 shows the density functions for both estimates. We find similar results with both distributions, but we present the results using Aarbu and Schroyen's estimate, since they cover a broader age group.

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

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

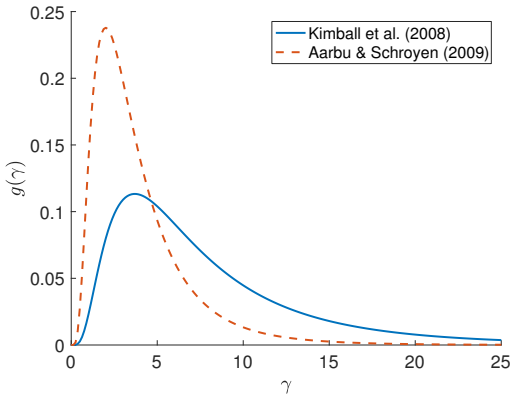

Figure 1: Distribution of relative rate of risk aversion, as estimated by Kimball et al. (2008): $\log \gamma \sim \mathcal{N}(1.84, 0.73)$, and Aarbu and Schroyen (2009): $\log \gamma \sim \mathcal{N}(1.14, 0.67)$.

3. Analytical results

3.1. Does a representative agent exist?

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

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

In Appendix A, we show that our market with heterogeneous agents cannot be captured by a representative agent with time-varying risk aversion. We furthermore prove that a representative agent aggregation does exist for the asset pricing model by Brock

and Hommes (1998). Specifically, we show that their heterogeneous market with heuristic switching agents is equivalent to one with a diversifying representative agent whose risk aversion is given by

$$\gamma_{R,t+1} = \begin{cases} \dfrac{6\beta^2(U_{F,t} - U_{M,t})}{\pi^2} \dfrac{\exp(\beta U_{F,t}) + \exp(\beta U_{M,t})}{\exp(\beta U_{F,t}) - \exp(\beta U_{M,t})} & \text{if } U_{F,t} \neq U_{M,t}; \\ \dfrac{12\beta}{\pi^2} & \text{if } U_{F,t} = U_{M,t}. \end{cases} \tag{21}$$

Here, $\beta > 0$ is the intensity of choice parameter. It measures how quickly agents switch between heuristics, and plays a key role in previous heuristic switching models. The heuristic performance measures $U_{F,t}$ and $U_{M,t}$ for period t can be equal to the weighted average of squared forecasting errors that we use in our model, or any other measure of performance.[11]

3.2. Dynamics

In this section, we present analytical results on the dynamics of the asset pricing model as defined by the dynamical system (20). These results give an idea of the type of behaviour that the model can exhibit, and help us understand the simulations that we present in section 4. First, there is only one steady state, with the risky asset price at the fundamental and the weighted averages and variances at zero:

Lemma 1 (Existence of steady state)**.** *The system* (20) *has precisely one steady state, namely* $(0, 0, 0, 0, 0)$.

Proof. See Appendix C. □

We cannot use standard linearization techniques to study the stability of the steady state, because the risk-tolerance fraction $\Theta_{F,t}$ has a discontinuity at this point. However, we can investigate its stability under the assumption that agents stick to one of the two

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

forecasting rules, so that the risk-tolerance fractions are constant and equal to an exogenously given value Θ_F. This assumption simplifies the analysis even further, because we can ignore the part of the system that governs the switching, namely the weighted averages and variances of the heuristics' performance. In what follows, we derive three lemmas that characterize the behaviour under the constant, exogenous fraction assumption. After this analysis, we relax the assumption to return to the original model, and use these lemmas to understand its behaviour.

3.2.1. Exogenous, constant fractions

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

$$\begin{pmatrix} \nu_t \\ x_t \end{pmatrix} = A \begin{pmatrix} \nu_{t-1} \\ x_{t-1} \end{pmatrix} \tag{22a}$$

$$A = \begin{pmatrix} 0 & 1 \\ -\frac{1}{R^f}\left(1-\Theta_F\right) m & \frac{1}{R^f}\left[\Theta_F\, f + \left(1-\Theta_F\right)\left(1+m\right)\right] \end{pmatrix}. \tag{22b}$$

Lemma 2 (Existence of steady state with constant fractions)**.** *The system* (22) *has* $(0,0)^T$ *as its only steady state.*

Proof. See Appendix C. □

This steady state corresponds to the risky asset price being at the fundamental, like that of the unrestricted system (20).

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

Lemma 3 (Dynamics in homogeneous market)**.** *If* $\Theta_F = 1$*, the steady state of* (22) *is stable, and the eigenvalues of* A *are real. Let*

$$m_\pm = 2\left(R^f \pm \sqrt{R^f\,\left(R^f-1\right)}\right) - 1, \tag{23}$$

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

1. *If* $0 < m < R^f$*, the steady state is stable.*
2. *If* $m > R^f$*, the steady state is unstable.*
3. *If* $m \leq m_-$ *or* $m \geq m_+$*, the eigenvalues of* A *are real.*
4. *If* $m_- < m < m_+$*, the eigenvalues of* A *are complex.*[12]

Proof. See Appendix C. □

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

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

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

If the market consists of a constant mix of fundamentalists and momentum traders, the dynamics is described by the following lemma, which is represented graphically in Figure 2:

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

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

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

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

and

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

1. *The steady state of (22) is unstable if $\Theta_F < \overline{\Theta}_F$.*
2. *The steady state of (22) is stable if $\Theta_F > \overline{\Theta}_F$.*
3. *If $m \leq m_-$ and $f_- \leq f \leq f_+$, the eigenvalues of A are real. Otherwise, we have:*
   (a) *The eigenvalues are real if $\Theta_F \leq \Theta_{F,-}$ or $\Theta_F \geq \Theta_{F,+}$.*
   (b) *The eigenvalues are complex if $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$.*
4. *If $m > 1$, it follows that $\Theta_{F,-} < \overline{\Theta}_F < \Theta_{F,+}$.*
5. *If $m > m_-$, it follows that $0 < \Theta_{F,+} \leq 1$, where $\Theta_{F,+} = 1$ if and only if $f = 0$.*
6. *If $m_- < m < m_+$, it follows that $\Theta_{F,-} < 0$, and if $m \geq m_+$, we have that $0 \leq \Theta_{F,-} < \Theta_{F,+}$, where $\Theta_{F,-} = 0$ if and only if $m = m_+$.*

Proof. See Appendix C. □

This lemma tells us that in a heterogeneous market with constant fractions, the stability of the steady state depends on the risk-tolerance fraction of fundamentalists, relative to a threshold $\overline{\Theta}_F$. If the fundamentalist fraction is relatively low ($\Theta_F < \overline{\Theta}_F$), the steady state is unstable, while a relatively high fraction ($\Theta_F > \overline{\Theta}_F$) implies a stable steady state. The size of the momentum parameter m also plays an important role. First, note that if the momentum parameter is small enough, namely $m < R^f$, the steady state is stable, independent of the fraction of fundamentalists: In this case we have that $\overline{\Theta}_F < 0$, which means that $\Theta_F > \overline{\Theta}_F$, since Θ_F is positive. Second, a larger m

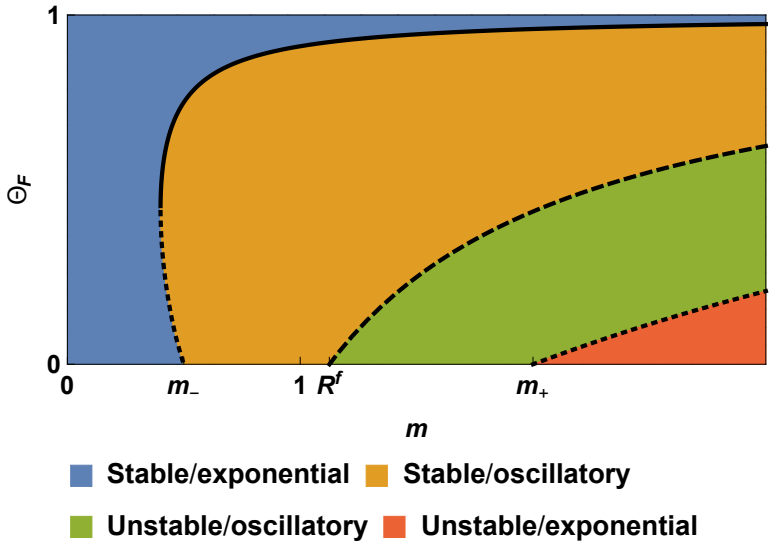

Figure 2: Graphical representation of Lemma 4, showing risky asset price dynamics with fixed fractions of fundamentalists and momentum traders, for different values of the momentum parameter m and the risk-tolerance fraction of fundamentalists Θ_F. Each of the four regions is characterized by two properties: the stability of the steady state and the type of dynamics (exponential or oscillatory). The solid and dotted lines represent $\Theta_{F,+}$ and $\Theta_{F,-}$, respectively, while the dashed line represents $\overline{\Theta}_F$. The shapes of these lines depend on the values of R^f and f, but the four regions remain. As R^f gets larger (smaller), the blue and red areas become smaller (larger). Changes in f have the opposite effect. The boundaries between the different regions cross the m-axis at m_-, R^f, and m_+, regardless of the parameter values.

implies a larger $\overline{\Theta}_F$, meaning that a larger risk-tolerance fraction of fundamentalists is required for the steady state to be stable. These results are consistent with the intuition behind Lemma 3 for the homogeneous case: The fundamentalist traders bring stability in the market, while the momentum rule is destabilizing, where a larger momentum parameter implies a stronger destabilizing effect.

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

We can use above results to describe the behaviour of the risky asset price for specific ranges of the momentum parameter and the risk-tolerance fraction of fundamentalists. As shown in Figure 2, we can distinguish between four regions. First, if Θ_F is large

or m small, the asset price exponentially converges towards the fundamental (the blue region). Second, the system exhibits converging oscillations when m is relatively small or Θ_F relatively large (the orange region). Third, if $m > R^f$, the asset price can exhibit oscillations diverging from the fundamental, depending on the value of Θ_F (the green region). Fourth, if $m \geq m_+$, the price exponentially diverges from the fundamental when Θ_F is small enough (the red region).

3.2.2. Endogenous, time-varying fractions

We return to the unrestricted system (20), in which agents are allowed to switch between forecasting strategies. Lemma 3 and Lemma 4 help us characterize its behaviour, because we can see the constant-fractions system (22) as part of the switching system. In each period t, $\Theta_{F,t}$ and the lagged asset price deviations x_{t-1} and x_{t-2} determine the current deviation x_t through (22). Given x_t, the weighted average and variances are pinned down by (20), which in turn determine $\Theta_{F,t+1}$ through (19). This new fraction then serves as the input for the constant-fractions system in period $t+1$, and the cycle repeats.

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

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

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

4. Numerical results

We now turn to numerical methods to study the dynamics of the asset pricing model with switching. The interesting dynamics occurs when the fundamentalist heuristic is attractive even when all agents use the momentum rule. This is the case when $m_- < m < m_+$, so that the asset price oscillates around the fundamental when $\Theta_F < \Theta_{F,+}$: During the oscillations, the fundamentalist rule becomes more attractive as the asset price moves towards the fundamental. For certain parameter combinations, the fundamentalist rule becomes attractive enough to make agents switch from the momentum rule. We use $R^f = 1.01$, $\eta = 0.2$, $f = 0.6$, and $m = 1.1$ as the baseline parameter values in our simulations. The distribution of relative risk aversion follows the empirically validated density $g(\gamma)$ that was introduced in subsection 2.3.

These parameter values correspond to $m_- \approx 0.82$, $m_+ \approx 1.22$, $\Theta_{F,-} \approx -0.02$, $\Theta_{F,+} \approx 0.84$, and $\overline{\Theta}_F \approx 0.08$. It follows that indeed $m_- < m < m_+$, so that $\Theta_{F,-} < 0$. Using the insights from subsection 3.2, we know that the asset price oscillates away from the fundamental if the risk-tolerance fraction of fundamentalists is smaller than 8%, oscillates towards it if the fundamentalist fraction lies between 8% and 84%, and exponentially converges towards the fundamental if the fraction is larger than 84%.

4.1. Chaotic dynamics

In the two left panels of Figure 3, we show the evolution of the price deviations from the fundamental, and of the risk-tolerance fraction of fundamentalists for the baseline

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

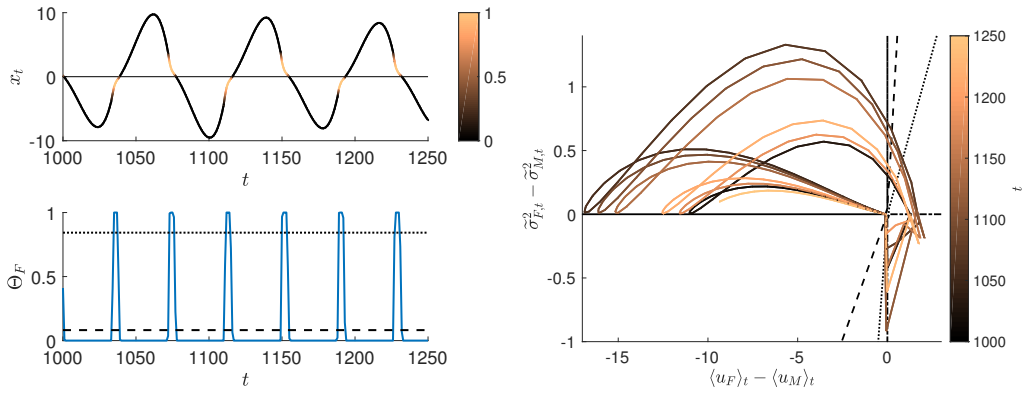

Figure 3: *Top left:* Price deviations from the fundamental for periods $1000 - 1250$, where the colour gradient represents the value of Θ_F. *Bottom left:* Risk-tolerance fraction choosing the fundamentalist strategy for periods $1000 - 1250$. The dashed and dotted lines indicate $\overline{\Theta}_F$ and $\Theta_{F,+}$, respectively. *Right:* Evolution of the dynamical system in the $(\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ plane for periods $1000 - 1250$. The solid, dashed, dotted, and dot-dashed black lines indicate the 0, $\overline{\Theta}_F$, $\Theta_{F,+}$, and 1 contour lines of $\Theta_{F,t+1}$, respectively. *Parameter values:* $R^f = 1.01$, $\eta = 0.2$, $f = 0.6$, $m = 1.1$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $x_{-2} = x_{-1} = x_0 = 0.1$, $\langle u_F \rangle_0 = -1$, $\langle u_M \rangle_0 = -0.90$, $\langle u_F^2 \rangle_0 = 2$, $\langle u_M^2 \rangle_0 = 1.91$ (implying $\Theta_{F,1} = 0.5$).

The right panel of Figure 3 shows the behaviour of the system over the same period in the $(\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ plane. Its position in this plane determines the risk-tolerance fraction $\Theta_{F,t+1}$ in the next period, through equation (19). We have added several contour lines for $\Theta_{F,t+1}$ to the figure, in particular those corresponding to $\Theta_{F,t+1} = 0$ (solid) and $\Theta_{F,t+1} = 1$ (dot-dashed): In the second quadrant, corresponding to $\langle u_F \rangle_t - \langle u_M \rangle_t < 0$ and $\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t} > 0$, we have that $\Theta_{F,t+1} = 0$, and in the fourth quadrant, corresponding to $\langle u_F \rangle_t - \langle u_M \rangle_t > 0$ and $\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t} < 0$, we have that $\Theta_{F,t+1} = 1$. Note that all contour lines come together in the discontinuity at the origin. Each cycle, the system stays in the second quadrant relatively long, which means that all agents follow the momentum rule. As the fundamentalist rule becomes more attractive, it moves towards the first quadrant, where the variance of the fundamentalist rule is still larger than that of the momentum rule. This means that the most risk-tolerant investors are the first to switch to the fundamentalist rule. When the variance of the fundamentalist rule drops below that of the momentum rule, the system moves towards the fourth quadrant, where all agents are fundamentalists. As the momentum rule starts picking up on the asset price movement towards the fundamental and becomes more attractive, the system briefly visits the third quadrant. It moves back to the second quadrant, with only momentum traders in the market, when the variance of the momentum rule becomes smaller than that of the fundamentalist rule. This cycle repeats itself indefinitely, each time in a slightly different way.

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

and policy-making.

The dynamics of a chaotic system can be characterized by a spectrum of Lyapunov exponents that indicate how nearby solutions converge or diverge over time. A negative exponent indicates convergence, while a positive exponent indicates divergence. The largest, or maximal, Lyapunov exponent, denoted by λ_1, determines whether the dynamics is chaotic (see, e.g., Gros, 2015). The average factor by which the distance between two neighbouring points increases or decreases in each time-step is given by e^{λ_1} (Boccara, 2010, ch. 5). Using the algorithm proposed by Benettin et al. (1980), we find for the asset pricing model that $\lambda_1 = 0.0035$.[13] This positive exponent implies exponential divergence: Solutions that start close to each other drift apart over time. This sensitive dependence on initial conditions is the defining characteristic of a chaotic system. Given the value $\lambda_1 = 0.0035$, the distance between neighbouring points doubles after approximately 200 time-steps.

Figure 4 shows the long-term behaviour of the asset pricing system. It plots the points of a solution consisting of one million time-steps (after an initialization period of 1000 steps) in the $(\langle u_F\rangle_t - \langle u_M\rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ plane (left panel) and in $x_t - (\langle u_F\rangle_t - \langle u_M\rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ space (right panel). In the left panel, we also show a contour plot for the fraction of fundamentalists Θ_F. It shows that the behaviour in Figure 3 is representative of the long-term dynamics. The system often comes close to the discontinuity at the origin, and it spends most time in the second quadrant, where all traders use the momentum rule. The right panel shows that these points correspond to large asset price deviations, while those points where x_t is small correspond to larger fractions of fundamentalists.

Both panels show that the solution is confined to a specific set of points: the strange attractor. Strange attractors arise in chaotic dynamical systems when trajectories are attracted to a subset of their phase space, but have chaotic dynamics within this subset:

[13] The other Lyapunov exponents are $\lambda_2 = 0.0010$, $\lambda_3 = -1.4917$, $\lambda_4 = -1.6059$, $\lambda_5 = -1.6094$, $\lambda_6 = -2.4687$, $\lambda_7 = -23.8493$. To calculate the exponents, we have used parameter values equal to those used in Figure 3, and an initialization period of 2000 time-steps.

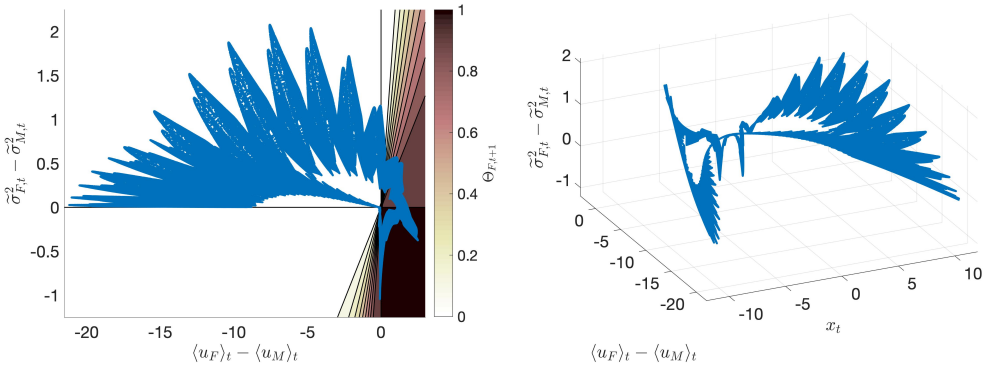

Figure 4: *Left:* Attractor of the asset pricing system (blue) and contour plot of $\Theta_{F,t+1}$ (colour gradient) in the $(\langle u_F\rangle_t - \langle u_M\rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ plane. *Right:* Attractor of the asset pricing system in $x_t - (\langle u_F\rangle_t - \langle u_M\rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ space. *Parameter values:* $R^f = 1.01$, $\eta = 0.2$, $f = 0.6$, $m = 1.1$, $\log\gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $x_{-2} = x_{-1} = x_0 = 0.1$, $\langle u_F\rangle_0 = -1$, $\langle u_M\rangle_0 = -0.90$, $\langle u_F^2\rangle_0 = 2$, $\langle u_M^2\rangle_0 = 1.91$ (implying $\Theta_{F,1} = 0.5$).

Nearby solutions that enter the attractor drift apart over time, while staying within the confines of the attractor. The dimension of a strange attractor is a useful tool in studying chaotic systems. A strange attractor's dimension quantifies the complexity of the system by representing its effective number of degrees of freedom. It can also be compared with other theoretical studies and empirically tested. The strange attractor in our model has a box-counting (fractal) dimension (Falconer, 2004, ch. 3) of about 1.8. This value is close to the dimensions reported by Brock and Hommes (1998).

The dynamics of the asset pricing model crucially depends on the parameter values. One example is the threshold $\overline{\Theta}_F$, which is determined by the ratio between R^f and m. Here we focus on the memory parameter η. Figure 5 shows long-term asset price deviations in the deterministic system for different values of this parameter: For each value of η, this diagram plots 1000 asset price deviations after an initialization period of 1000 time-steps. Solutions with deviations that are larger than 1000 in absolute value are omitted.

For η larger than 0.63, the system is unstable, and the amplitude of the asset price oscillations increases indefinitely. In this case, the momentum traders take over the market, and $\Theta_{F,t}$ stays constant at 0. For some values between 0.40 and 0.63, the solution

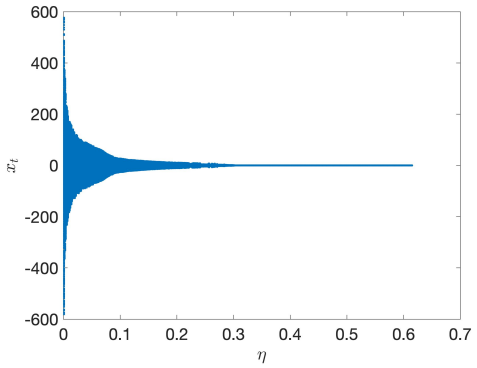

Figure 5: Long-term behaviour of the asset price deviations from the fundamental for different values of the memory parameter η. Deviations that are larger than 1000 in absolute value are omitted. *Parameter values:* $R^f = 1.01$, $f = 0.6$, $m = 1.1$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $x_{-2} = x_{-1} = x_0 = 0.1$, $\langle u_F \rangle_0 = -1$, $\langle u_M \rangle_0 = -0.90$, $\langle u_F^2 \rangle_0 = 2$, $\langle u_M^2 \rangle_0 = 1.91$ (implying $\Theta_{F,1} = 0.5$).

converges to the steady state, with the asset price at the fundamental. For other values between 0.40 and 0.63, the dynamics is chaotic, with small asset price deviations from the fundamental. The dynamics is also chaotic if η is smaller than 0.40, with larger deviations for smaller η.

Note on the wealth distribution. While the risk aversion distribution does not change in the asset pricing model, the wealth distribution can shift over time. More risk-tolerant traders invest a larger share of their wealth in the risky asset, which means that they can earn bigger returns. They also make bigger losses, but without a bankruptcy mechanism in place, they stay in the market. Shifts of wealth between traders does not affect the asset price however, because their demand for the risky asset is independent of their wealth (see equation 3). We have chosen not to include a bankruptcy mechanism to keep the asset pricing model as simple as possible, and focus on the dynamics resulting from the heterogeneous expectations.

Cobweb application. To investigate to what extent the richness of the dynamics resulting from our expectation-formation framework depends on the choice of model, we also study its dynamics in the cobweb model (see Appendix D). The cobweb model is used by Muth (1961) to introduce the rational expectations hypothesis, and by Brock and Hommes

(1997) to introduce their heuristic switching model. Contrary to the asset pricing model, this market exhibits negative expectations feedback: A higher expected price will lead to a lower realized price. Like in Brock and Hommes (1997), agents can choose between a stable, rational strategy, and an unstable, naive strategy, where the former is costlier. In this different context, our mean-variance switching model gives rise to a dynamics that is just as rich as that observed in the asset pricing model, including deterministic chaos.

4.2. Resilience to price shocks

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

$$x_{t+1} = \frac{1}{R^f} \left\{ \Theta_{F,t} \, f \, x_t + (1 - \Theta_{F,t}) \left[x_t + m \left(x_t - x_{t-1} \right) \right] \right\} + \epsilon_t. \tag{24}$$

We simulate this stochastic version of the model with an i.i.d. price shock $\epsilon_t \sim \mathcal{N}(0, 0.05)$, which is small compared to the range of price deviations that occur in the deterministic system. The results are presented in the top panels of Figure 6.

We do not just observe noise around the deterministic solution of Figure 3, but the dynamics is different. The heterogeneous beliefs, especially the trend-following heuristic, amplify the small noise, leading to larger and less regular asset price deviations. Periods where the fundamentalists take over the market also last longer. The changed dynamics is reflected in the long-term behaviour, which is presented in the top-right panel of Figure 6: The shape of the attractor is preserved, but it is bigger and less well-defined.

In the following, we investigate the stochastic system in different parameter ranges. When $\eta = 0.4$, the deterministic system converges to the steady state.[14] However, when we introduce a small normally distributed price shock with a standard deviation of 0.01,

[14] This is true for a wide range of initial values, as long as the initial asset price deviations do not become too large. The convergence is not exponential and relatively slow: Starting from the initial conditions used in Figure 3, the solution oscillates towards the steady state in approximately 2500 time-steps.

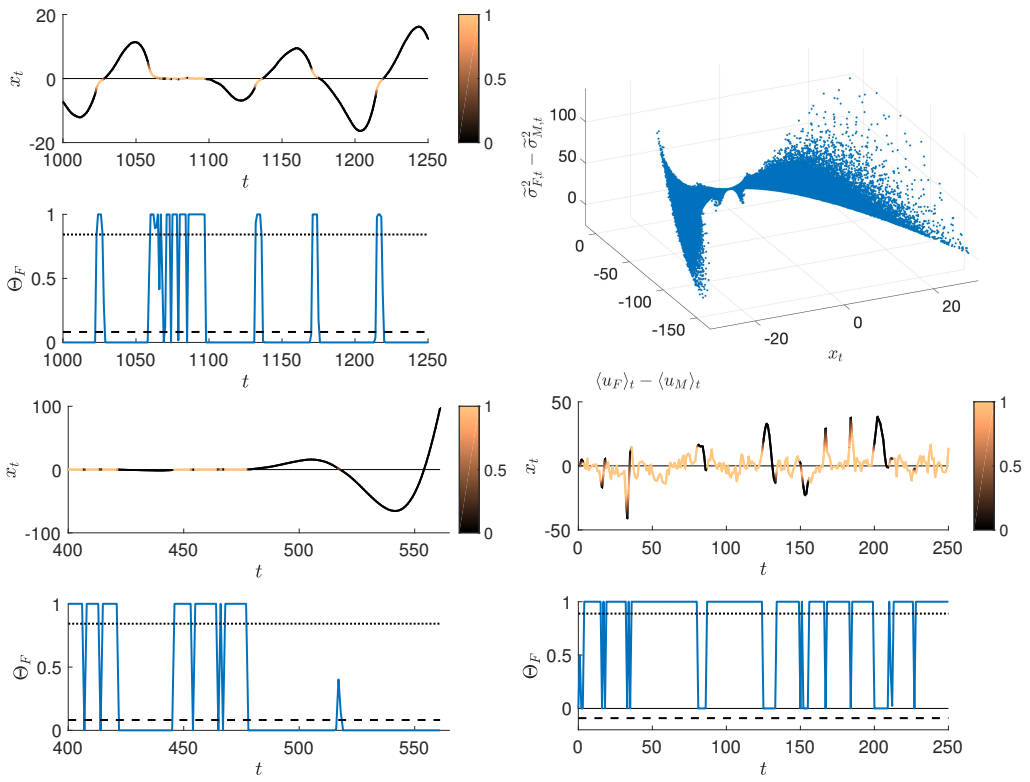

Figure 6: *Two top-left panels:* Price deviations from the fundamental (top) and risk-tolerance fraction choosing the fundamentalist strategy (bottom) for periods $1000 - 1250$, with $R^f = 1.01$, $\eta = 0.2$, and x_{t+1} given by (24) with an i.i.d. price shock $\epsilon_t \sim \mathcal{N}(0, 0.05)$. *Top-right panel:* Attractor in $x_t - (\langle u_F \rangle_t - \langle u_M \rangle_t) - (\widetilde{\sigma}^2_{F,t} - \widetilde{\sigma}^2_{M,t})$ space of the noisy asset pricing system depicted in the top-left panels. *Two bottom-left panels:* Price deviations from the fundamental (top) and risk-tolerance fraction choosing the fundamentalist strategy (bottom) for periods $400 - 561$, with $R^f = 1.01$, $\eta = 0.4$, and x_{t+1} given by (24) with an i.i.d. price shock $\epsilon_t \sim \mathcal{N}(0, 0.05)$. *Two bottom-right panels:* Price deviations from the fundamental (top) and risk-tolerance fraction choosing the fundamentalist strategy (bottom) for periods $0 - 250$, with $R^f = 1.2$, $\eta = 0.4$, and x_{t+1} given by (24) with an i.i.d. price shock $\epsilon_t \sim \mathcal{N}(0, 5)$. *Note:* A colour gradient represents the value of Θ_F, and dashed and dotted lines indicate $\overline{\Theta}_F$ and $\Theta_{F,+}$, respectively. *Parameter values:* $f = 0.6$, $m = 1.1$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $x_{-2} = x_{-1} = x_0 = 0.1$, $\langle u_F \rangle_0 = -1$, $\langle u_M \rangle_0 = -0.90$, $\langle u_F^2 \rangle_0 = 2$, $\langle u_M^2 \rangle_0 = 1.91$ (implying $\Theta_{F,1} = 0.5$).

the solution is pushed out of equilibrium, and exhibits behaviour similar to that shown in the top panels of Figure 6.

Increasing the standard deviation of the shock to 0.05, the solution becomes explosive. In the bottom-left panels of Figure 6, we show its behaviour for periods 400-561. It starts close to the fundamental, with long periods of only fundamentalists in the market. After about 80 time-steps, the momentum traders take over, and the asset price starts oscillating around the fundamental with ever-increasing amplitude. Although the deterministic system is stable, it is less resilient than the chaotic $\eta = 0.2$ system, because it is less able to absorb small price shocks.

Following Lemma 4, we showed that the deterministic asset pricing model always converges to the steady state when $m < R^f$, because in this case $\Theta_{F,t} \geq 0 > \overline{\Theta}_F$. If we adjust the $\eta = 0.4$ system by setting $R^f = 1.2 > m = 1.1$, it is much more resilient. The bottom-right panels of Figure 6 show the solution with large price shocks $\epsilon_t \sim \mathcal{N}(0, 5)$. Note that indeed, the dashed line, corresponding to $\overline{\Theta}_F$, lies below zero, and hence $\Theta_{F,t}$ never comes below it. Even though the price shocks are bigger, the solution stays closer to the fundamental than in the case where $R^f < m$. The market is dominated by fundamentalists, which stabilize it. Momentum traders only take over from time to time to give rise to short-lived booms and busts.

5. Conclusion

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

We have proven that this model cannot be captured by a representative agent. Furthermore, we have shown that a steady state exists, and have characterized the asset

price dynamics under the assumption of constant fractions: Stability of the steady state depends on the risk-tolerance fraction of fundamentalists relative to a threshold. This threshold depends on the strength of the momentum rule relative to the risk-free rate. If the risk-free rate is larger than the strength of the momentum rule ($R^f > m$), the steady state is always stable. A lower interest rate implies a larger threshold, which means that a larger fundamentalist fraction is needed for the steady state to be stable. This result is particularly interesting given the current low-interest-rate climate, and suggests that monetary stimulus can lead to less stable financial markets.

Numerically, we have shown that allowing for time-varying fractions, our asset pricing model can exhibit chaotic dynamics with a strange attractor. Bubbles and crashes emerge endogenously and unpredictably. By including stochastic asset price shocks, we have shown that heterogeneous beliefs and the resulting non-linearities have implications for resilience. Even when the endogenous dynamics (without stochastic shocks) is minimal or non-existent, small price shocks can cause large asset price bubbles.

We highlight three promising avenues for further research. First, agents could be modelled to form expectations by combining heuristics in a weighted average, in the same way that the representative agent does in subsection 3.1. While exploring the implications of this assumption, we have found that a representative agent does exist when the heterogeneous agents diversify across heuristics in this way, given that she has a time-varying risk aversion.[15] This could be linked to the time-varying risk aversion that is used in the more traditional finance literature that uses representative agents (e.g., Campbell and Cochrane, 1999).

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

[15] A proof is available upon request.

agents with different risk preferences (Chan and Kogan, 2002; Xiouros and Zapatero, 2010).

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

References

Aarbu, K. O., Schroyen, F., 2009. Mapping risk aversion in Norway using hypothetical income gambles. Working Paper, Norwegian School of Economics and Business Administration, Bergen.

Anderson, S. P., De Palma, A., Thisse, J. F., 1992. Discrete Choice Theory of Product Differentiation. MIT Press, Cambridge, MA.

Ang, A., 2014. Asset Management: A Systematic Approach to Factor Investing. Oxford University Press, Oxford.

Anufriev, M., Hommes, C., 2012. Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. American Economic Journal: Microeconomics 4 (4), 35–64.

Anufriev, M., Panchenko, V., 2009. Asset prices, traders' behavior and market design. Journal of Economic Dynamics and Control 33 (5), 1073–1090.

Benettin, G., Galgani, L., Giorgilli, A., Strelcyn, J.-M., 1980. Lyapunov characteristic exponents for smooth dynamical systems and for hamiltonian systems; A method for computing all of them. Part 2: Numerical application. Meccanica 15 (1), 21–30.

Boccara, N., 2010. Modeling Complex Systems. Graduate Texts in Physics. Springer, New York, NY.

Branch, W. A., 2004. The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations. The Economic Journal 114 (497), 592–621.

Brock, W. A., Hommes, C., 1997. A rational route to randomness. Econometrica 65 (5), 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 (8–9), 1235–1274.

Campbell, J. Y., Cochrane, J. H., 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. Journal of Political Economy 107 (2), 205–251.

Chan, Y. L., Kogan, L., 2002. Catching up with the Joneses: Heterogeneous preferences and the dynamics of asset prices. Journal of Political Economy 110 (6), 1255–1285.

Chiarella, C., He, X.-Z., 2002. Heterogeneous beliefs, risk and learning in a simple asset pricing model. Computational Economics 19 (1), 95–132.

Chiarella, C., He, X.-Z., 2003. Heterogeneous beliefs, risk, and learning in a simple asset-pricing model with a market maker. Macroeconomic Dynamics 7 (4), 503–536.

Chiarella, C., He, X.-Z., Zwinkels, R. C. J., 2014. Heterogeneous expectations in asset pricing: Empirical evidence from the S&P500. Journal of Economic Behavior & Organization 105, 1–16.

Choi, S., Fisman, R., Gale, D., Kariv, S., 2007. Consistency and heterogeneity of individual behavior under uncertainty. American Economic Review 97 (5), 1921–1938.

Cohen, A., Einav, L., 2007. Estimating risk preferences from deductible choice. American Economic Review 97 (3), 745–788.

De Grauwe, P., Foresti, P., Ji, Y., 2019. Fiscal Policies in Booms and Busts. CEPR Discussion Paper 13740, Centre for Economic Policy Research, London.

De Grauwe, P., Macchiarelli, C., 2015. Animal spirits and credit cycles. Journal of Economic Dynamics and Control 59, 95–117.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2012. The intergenerational transmission of risk and trust attitudes. The Review of Economic Studies 79 (2), 645–677.

Falconer, K., 2004. Fractal Geometry: Mathematical Foundations and Applications. John Wiley & Sons.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., Sunde, U., 2018. Global evidence on economic preferences. The Quarterly Journal of Economics 133 (4), 1645–1692.

Gaunersdorfer, A., Hommes, C. H., Wagener, F. O. O., 2008. Bifurcation routes to volatility clustering under evolutionary learning. Journal of Economic Behavior & Organization 67 (1), 27–47.

Gros, C., 2015. Complex and Adaptive Dynamical Systems. Springer International Publishing, Cham.

Hommes, C., 2013. Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems. Cambridge University Press, Cambridge.

Hommes, C., Massaro, D., Weber, M., 2019. Monetary policy under behavioral expectations: Theory and experiment. European Economic Review 118, 193–212.

Hommes, C., Sonnemans, J., Tuinstra, J., van de Velden, H., 2005. Coordination of expectations in asset pricing experiments. The Review of Financial Studies 18 (3), 955–980.

Hommes, C., Sonnemans, J., Tuinstra, J., van de Velden, H., 2008. Expectations and bubbles in asset pricing experiments. Journal of Economic Behavior & Organization 67 (1), 116–133.

Jaeger, D. A., Dohmen, T., Falk, A., Huffman, D., Sunde, U., Bonin, H., 2010. Direct evidence on risk attitudes and migration. The Review of Economics and Statistics 92 (3), 684–689.

Kimball, M. S., Sahm, C. R., Shapiro, M. D., 2008. Imputing risk tolerance from survey responses. Journal of the American Statistical Association 103 (483), 1028–1038.

Lei, V., Noussair, C. N., Plott, C. R., 2001. Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality. Econometrica, 831–859.
Mankiw, N. G., Reis, R., Wolfers, J., 2003. Disagreement about inflation expectations. NBER Macroeconomics Annual 18, 209–248.
Markowitz, H., 1952. Portfolio Selection. The Journal of Finance 7 (1), 77–91.
Muth, J. F., 1961. Rational expectations and the theory of price movements. Econometrica 29 (3), 315–335.
Nelson, D. B., 1992. Filtering and forecasting with misspecified ARCH models I: Getting the right variance with the wrong model. Journal of Econometrics 52 (1), 61–90.
Paravisini, D., Rappoport, V., Ravina, E., 2016. Risk aversion and wealth: Evidence from person-to-person lending portfolios. Management Science 63 (2), 279–297.
Park, B.-J., 2014. Time-varying, heterogeneous risk aversion and dynamics of asset prices among boundedly rational agents. Journal of Banking & Finance 43, 150–159.
Pfajfar, D., 2013. Formation of rationally heterogeneous expectations. Journal of Economic Dynamics and Control 37 (8), 1434–1452.
Smith, V. L., Suchanek, G. L., Williams, A. W., 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. Econometrica: Journal of the Econometric Society, 1119–1151.
von Gaudecker, H.-M., van Soest, A., Wengström, E., 2011. Heterogeneity in risky choice behavior in a broad population. American Economic Review 101 (2), 664–694.
Wieland, V., Wolters, M. H., 2010. The diversity of forecasts from macroeconomic models of the US economy. Economic Theory 47 (2-3), 247–292.
Xiouros, C., Zapatero, F., 2010. The representative agent of an economy with external habit formation and heterogeneous risk aversion. The Review of Financial Studies 23 (8), 3017–3047.

Appendix A. Does a representative agent exist?

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

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

We start by deriving the pricing equation in the representative agent market, as well as the weights that she assigns to both rules in terms of the mean and variance of their performance. The mean-variance utility (1) does not have a maximum if $\gamma \leq 0$, so we assume that the representative agent has risk aversion $\gamma_{R,t} \in (0, \infty)$ for all t. In this case, aggregate demand is equal to individual demand, given by the representative agent version of equation (4):

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

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

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

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

where $w_{F,t}$ and $w_{M,t}$ are the period t weights assigned to the fundamentalist and mo-

mentum rules, respectively. Note that we use the same two heuristics as in our model. However, the derivations and conclusions in this appendix hold for any two heuristics, including those explored by Brock and Hommes (1998).

The price deviations from the fundamental are given by

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

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

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

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

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

with $\widetilde{\sigma}_{FM,t}$ a measure of the covariance between the utilities of the two heuristics in period t.[16] Note that period t risk aversion enters the performance measure, while mean and (co)variance measures are from period $t-1$. The timing is similar to that in the heterogeneous economy; period t mean and variance are not yet known when determining the weights for that period.

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

Lemma 5 (Representative agent weights)**.** *The weights chosen by the representative*

[16] We do not need to specify this measure here, but a definition in line with (14) would be $\widetilde{\sigma}_{FM,t} = \langle u_F\, u_M \rangle_t - \langle u_F \rangle_t \langle u_M \rangle_t$, with $\langle u_F\, u_M \rangle_t = \eta\, \langle u_F\, u_M \rangle_{t-1} + (1-\eta)\, u_{F,t}\, u_{M,t}$.

agent with relative risk aversion coefficient $\gamma_{R,t}$ *are given by*

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

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

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

and

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

Proof. See Appendix C. □

We now consider the representative agent for the model of Brock and Hommes (1998), who model the fractions of agents that use the heuristics with the multinomial logit (MNL).[17] The MNL can be derived from a stochastic utility model:[18] The utility $\widetilde{U}_{s,i,t}$ derived from heuristic $s \in \{F, M\}$ by agent i is expressed as

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

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

The representative agent is a mean-variance optimizer, who therefore wants to maximize the individual utilities, while penalizing their spread. The spread of individual

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

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

utilities around the performance measure is measured by the variance of the stochastic term. The size of the penalty is determined by the representative agent's risk aversion. The MNL arises when the $\epsilon_{s,i,t}$ are assumed to be i.i.d across agents and heuristics according to the double exponential distribution. The variance of this distribution is given by $\pi^2/(6\beta^2)$, with $\beta > 0$ the intensity of choice parameter. To be able to compare the original model with our representative agent aggregation, we make the same assumption regarding the distribution of the stochastic utility term. To translate the results of Lemma 5 to this setting, we have

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

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

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

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

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

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

Lemma 6 (Representative agent for multinomial logit switching)**.** *Suppose that the representative agent believes that the performance measures of the two forecasting heuristics are independent ($\widetilde{\sigma}_{FM,t} = 0$ for all t), and have equal, constant variance, given by*

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

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

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

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

Proof. See Appendix C. □

The agent's risk aversion $\gamma_{R,t}$ can be set to any value when $U_{F,t} = U_{M,t}$. However, we have made the particular choice presented in the lemma to ensure continuity of risk aversion as a function of $U_{F,t}$ and $U_{M,t}$. We also note that more general formulations of the lemma can be proven. The only requirements are that the variances of the two heuristics are equal and non-zero. An aggregation can also be achieved when variances are time-varying, or when the performance measures are correlated.[19]

We now turn to the representative agent for our market consisting of mean-variance optimizing agents with heterogeneous risk preferences. Comparing the representative agent price dynamics (A.3) with the heterogeneous system (20), we see that the representative agent economy is identical to it if and only if the weight assigned by the representative agent to the fundamentalist heuristic is equal to the risk-tolerance fraction represented by that heuristic in the heterogeneous economy in every period, that is, if and only if $w_{F,t} = \Theta_{F,t}$ for all $t \in \mathbb{N}_0$. Now suppose that for some period t, we have that $\langle u_F \rangle_t = \langle u_M \rangle_t$ and $\widetilde{\sigma}_{FM,t} < \widetilde{\sigma}^2_{F,t} < \widetilde{\sigma}^2_{M,t}$. It follows from equation (19) that $\Theta_{F,t+1} = 1$, while Lemma 5 tells us that $w_{F,t+1} < 1$, independent of the agents' risk aversion. It follows that $w_{F,t+1} \neq \Theta_{F,t+1}$, implying that our heterogeneous market of

[19] A proof is available upon request.

switching agents cannot be captured by a representative agent. The same argument holds if the representative agent does not take into account the covariance (i.e., $\widetilde{\sigma}_{FM,t} = 0$).[20]

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

Appendix B. Introducing noise traders

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

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

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

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

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

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

Appendix C. Proofs

Proof of Lemma 1. We write $\left(x^*, \langle u_F\rangle^*, \langle u_M\rangle^*, \left\langle u_F^2\right\rangle^*, \left\langle u_M^2\right\rangle^*\right)$ to denote the steady state. It follows from the definition of the system (20) that

$$\begin{aligned}\langle u_F\rangle^* &= \eta\,\langle u_F\rangle^* - (1-\eta)\,(x^* - f\,x^*)^2\\ \langle u_F\rangle^* &= -\,(x^*)^2\,(1-f)^2\,.\end{aligned}$$

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

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

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

It now follows from (19) that

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

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

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

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

Proof of Lemma 2. We write the steady state as (ν^*, x^*). Because $\nu_t = x_{t-1}$, we have that $\nu^* = x^*$. It then follows from the definition of A in (22) that

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

where in the last step we have used that $R^f - 1 + \Theta_F(1-f) > 0$, since $R^f > 1$,

$0 \leq \Theta_F \leq 1$, and $f < 1$. We conclude that $(0,0)^T$ is the only steady state. □

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

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

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

If $\Theta_F = 0$, we have

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

with characteristic equation

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

and eigenvalues

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

The eigenvalues are complex when

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

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

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

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

$\lambda_\pm \in \mathbb{R}$ if $m \leq m_-$ or $m \geq m_+$.

Because m_- is strictly decreasing in R^f[21] and

$$\lim_{R^f \to \infty} R^f - \sqrt{R^f\ (R^f - 1)} = \frac{1}{2},$$[22]

we have that

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

Because on the other hand

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

it follows that

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

---

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

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

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

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

The derivative is equal to 0 if and only if

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

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

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

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

We also have that

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

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

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

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

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

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

and

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

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

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

In this case, it follows that

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

so that

$$\lambda_- = |\lambda_-| \leq |\lambda_+| = \lambda_+. \tag{C.1}$$

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

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

First note that

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

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

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

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

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

implying that

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

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

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

Squaring both sides and rearranging leads to

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

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

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

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

We have that

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

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

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

The sign of the derivative is determined by the function

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

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

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

Since ζ is continuous and nowhere vanishing,[23] it is positive on the whole interval

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

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

(m_+, ∞), which implies that the derivative of λ_- with respect to m is negative on this interval. We conclude hat λ_- is strictly decreasing as a function of $m \in (m_+, \infty)$. We finalize the proof that $\lambda_- > 1$ for $m \in (m_+, \infty)$ by showing that $\lambda_-(m) \to 1$ as $m \to \infty$:

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

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

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

Proof of Lemma 4. The map A, as defined in (22) has characteristic equation

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

Squaring both sides and rearranging gives us

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

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

The eigenvalues are given by

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

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

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

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

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

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

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

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

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

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

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

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

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

The modulus of the eigenvalues. Note that

$$\sqrt{\xi(\Theta_F; R^f, f, m)} \leq \sqrt{\left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2} = \Theta_F\, f + (1-\Theta_F)\,(1+m)^{24},$$

which implies that

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

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

$$\lambda_- = |\lambda_-| \leq |\lambda_+| = \lambda_+. \tag{C.2}$$

If the eigenvalues are complex, we have that

$$\lambda_{\pm} = \frac{1}{2R^f}\Big\{\Theta_F\, f + (1-\Theta_F)\,(1+m) \\ \pm i\sqrt{4R^f\,(1-\Theta_F)\,m - \left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2}\Big\},$$

[24] Note that $\sqrt{\xi(\Theta_F; R^f, f, m)} \neq \sqrt{\left[\Theta_F\, f + (1-\Theta_F)\,(1+m)\right]^2}$ if $\Theta_F < 1$ and $m > 0$.

so that

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

It follows that $|\lambda_{\pm}| = 1$ at a threshold value given by

$$
\overline{\Theta}_F = 1 - \frac{R^f}{m}. \tag{C.3}
$$

If $\Theta_F < \overline{\Theta}_F$, we have that $|\lambda_{\pm}| > 1$, and if $\Theta_F > \overline{\Theta}_F$, it follows that $|\lambda_{\pm}| < 1$.

Now assume that $m > 1$. We are going to show that $\Theta_{F,-} < \overline{\Theta}_F < \Theta_{F,+}$. Consider $\Theta_{F,\pm}$ as functions of f:[25] $\Theta_{F,\pm} : [0,1) \to \mathbb{R}$,

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

It follows that

$$
\Theta_{F,\pm}(0) = 1 + \frac{\pm 2mR^f - 2mR^f}{(1+m)^2}.
$$

We find that $\Theta_{F,+}(0) = 1 > \overline{\Theta}_F$, while $\Theta_{F,-}(0) < \overline{\Theta}_F$ if and only if

$$
\begin{aligned}
\frac{4mR^f}{(1+m)^2} &> \frac{R^f}{m} \\
4 &> \frac{(1+m)^2}{m^2} \\
4 &> 1 + \frac{2}{m} + \frac{1}{m^2} \\
m &> 1,
\end{aligned}
$$

[25] Later in this proof, we show that $m > m_-$, which includes the case of $m > 1$, implies that $\Theta_{F,\pm} \in \mathbb{R}$ and $\Theta_{F,-} < \Theta_{F,+}$. We do this by noting that the $f_{\pm}$ are complex if $m_- < m < m_+$ and showing that $f_{\pm} > 1$ if $m \geq m_+$.

where we have used the fact that $m > 0$. Hence it follows that $\Theta_{F,-}(0) < \overline{\Theta}_F < \Theta_{F,+}(0)$. Since $\Theta_{F,\pm}$ are continuous functions of f and $\overline{\Theta}_F$ is independent of f, we can only have $\overline{\Theta}_F \notin (\Theta_{F,-}, \Theta_{F,+})$ if $\Theta_{F,-}(f) = \overline{\Theta}_F$ or $\Theta_{F,+}(f) = \overline{\Theta}_F$ for some $f \in [0, 1)$. However, note that $\Theta_{F,\pm} = \overline{\Theta}_F$ if and only if

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

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

Squaring both sides and solving for R^f leads to

$$R^f = \frac{f\,m}{f-1-3m} \text{ or } R^f = \frac{f\,m}{f-1+m}.$$

These solutions contradict $R^f > 1$, because we have that

$$\frac{f\,m}{f-1-3m} < 0,$$

and since $m > 1$, it follows that

$$\begin{aligned} f &< 1 \\ (m-1)f &< m-1 \\ f\,m &< f-1+m \\ \frac{f\,m}{f-1+m} &< 1, \end{aligned}$$

so $\Theta_{F,\pm} \neq \overline{\Theta}_F$. We conclude that $\Theta_{F,-} < \overline{\Theta}_F < \Theta_{F,+}$ if $m > 1$.

Before we continue with the proof, we formulate a lemma about the the range of the eigenvalues when they are real. Its proof can be found at the end of this appendix.

Lemma 7. *The functions defined by*

$$\lambda_{\pm} : D \to \mathbb{R}, \quad \Theta_F \mapsto \frac{\Theta_F\, f + (1-\Theta_F)\,(1+m) \pm \sqrt{\xi\,(\Theta_F; R^f, f, m)}}{2R^f},$$

with $0 \leq f < 1$, $m > 0$, $R^f > 1$, *and where the domain* D *is chosen to ensure that* $D \subset [0,1]$ *and that the function is real:*

$$D = \begin{cases} [0,1] & \text{if } f_- \leq f \leq f_+ \text{ or } \Theta_{F,\pm} \leq 0 \text{ or } \Theta_{F,\pm} \geq 1; \\ [0,\Theta_{F,-}] \cup [\Theta_{F,+},1] & \text{if } 0 \leq \Theta_{F,-} < \Theta_{F,+} \leq 1; \\ [0,\Theta_{F,-}] & \text{if } 0 \leq \Theta_{F,-} < 1 < \Theta_{F,+}; \\ [\Theta_{F,+},1] & \text{if } \Theta_{F,-} < 0 < \Theta_{F,+} \leq 1; \\ \emptyset & \text{if } \Theta_{F,-} < 0 \text{ and } \Theta_{F,+} > 1; \end{cases}$$

satisfy

$$\lambda_{\pm} \neq 1.$$

The case $0 < m \leq m_-$. Suppose that $0 < m \leq m_-$. Because $m \leq m_-$, it follows that $(1+m)^2 - 4mR^f \geq 0$ (see the proof of Lemma 3), and the zeros $f_{\pm}$ of χ are real. Also note that

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

so that $f_- > 0$. We also have that

$$f_- \leq \frac{1+m_-}{2} < 1,$$

because $m_- < 1$ (Lemma 3). This means that there exist $f \in (0,1)$ for which $f_- \leq f \leq f_+$. For these values of f, the eigenvalues of A are real. We conclude that the eigenvalues are real when $0 < m \leq m_-$ and $f_- \leq f \leq f_+$. If $0 < m \leq m_-$ and $f \notin [f_-, f_+]$, the zeros $\Theta_{F,\pm}$ of ξ are real and $\Theta_{F,-} < \Theta_{F,+}$. In this case, the eigenvalues are complex if $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$, and real otherwise.

Let us now consider the stability of the steady state when $0 < m \leq m_-$. In the following, we show that the steady state is stable if the eigenvalues are real. Given inequality (C.2), it suffices to show that $\lambda_+ < 1$. First assume that $f_- \leq f \leq f_+$ or $\Theta_{F,\pm} \leq 0$ or $\Theta_{F,\pm} \geq 1$, so that we can define λ_+ as a continuous function of $\Theta_F \in [0,1]$ as in Lemma 7. From Lemma 3, we know that $\lambda_+(1) < 1$. Because λ_+ is continuous and Lemma 7 gives us that $\lambda_+ \neq 1$, it follows that $\lambda_+ < 1$.

Now assume that $\Theta_{F,\pm} \in \mathbb{R}$. If $0 < \Theta_{F,-} < 1$ and $\Theta_F \leq \Theta_{F,-}$, we can define λ_+ as a function of $\Theta_F \in [0, \Theta_{F,-}]$. Lemma 3 implies that $\lambda_+(0) < 1$. Because λ_+ is continuous on $[0, \Theta_{F,-}]$ and Lemma 7 tells us that $\lambda_+ \neq 1$, it follows that $\lambda_+ < 1$. If $0 < \Theta_{F,+} < 1$ and $\Theta_F \geq \Theta_{F,+}$, we can define λ_+ as a continuous function on $[\Theta_{F,+}, 1]$. Combining this with the fact that $\lambda_+(1) < 1$ and Lemma 7, we conclude that $\lambda_+ < 1$.

In the case that $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$, the eigenvalues of A are complex. We have already shown that the modulus of the eigenvalues depends on a threshold fraction of fundamentalists $\overline{\Theta}_F$, as defined in equation (C.3). If $m \leq m_- < 1$, we have that $\overline{\Theta}_F < 0$, and $\Theta_F > \overline{\Theta}_F$. This means that $|\lambda_\pm| < 1$ and the steady state is stable.

The case $m_- < m < m_+$. We start by showing that $\Theta_{F,-} < 0 < \Theta_{F,+} \leq 1$. If $m_- < m < m_+$, the zeros $f_\pm$ of χ are complex, and $\Theta_{F,\pm} \in \mathbb{R}$. We know from Lemma 3 that if $m_- < m < m_+$, the eigenvalues of A are complex at $\Theta_F = 0$, which means that $\xi(0) < 0$. This implies for the zeros of ξ that $\Theta_{F,-} < 0$ and $\Theta_{F,+} > 0$. From the definition of ξ, we know that

$$\xi(1) = f^2 \geq 0,$$

where $\xi(1) = 0$ if and only if $f = 0$. This means that either $\Theta_{F,+} = 1$ (when $f = 0$), or $\Theta_{F,+} \in (0,1)$ (when $f > 0$). We conclude that $\Theta_{F,-} < 0 < \Theta_{F,+} \leq 1$, where $\Theta_{F,+} = 1$ if and only if $f = 0$.

Now $\Theta_F < \Theta_{F,+}$ implies that $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$, since $\Theta_{F,-} < 0 \leq \Theta_F$. This means that the eigenvalues of A are complex, and the stability of the steady state depends on the threshold value $\overline{\Theta}_F$ like before.

Lastly, we show that the steady state is stable if $\Theta_F \geq \Theta_{F,+}$. In this case, it follows that the eigenvalues are real, and inequality (C.2) implies that it suffices to show that $\lambda_+ < 1$. Defining λ_+ as a function of Θ_F on $[\Theta_{F,+}, 1]$, Lemma 7 implies that it suffices to show that $\lambda_+(\Theta_F) < 1$ for some $\Theta_F \in [\Theta_{F,+}, 1]$. We already know from Lemma 3 that $\lambda_+(1) < 1$, so we can conclude that the steady state is stable if $\Theta_F \geq \Theta_{F,+}$.

The case $m \geq m_+$. If $m \geq m_+$, the zeros $f_\pm$ of χ are real, but they satisfy $f_\pm > 1$, which means that $\Theta_{F,\pm} \in \mathbb{R}$ and $\Theta_{F,-} < \Theta_{F,+}$ (because $f < 1$). To see that $f_\pm > 1$, note that χ has its minimum at f^*, which satisfies

$$f^* = \frac{1+m}{2} \geq \frac{1+m_+}{2} > \frac{1+R^f}{2} > 1.$$

If $m = m_+$, it follows that $f_\pm = f^* > 1$. If $m > m_+$, we have that $f_\pm \in \mathbb{R}$ and $\chi(f^*) < 0$. On the other hand, we know that

$$\chi(1) = m\left(R^f - 1\right) > 0.$$

Because χ is continuous, it follows that $f_- \in (1, f^*)$, and

$$1 < f_- < f_+.$$

We continue by showing that $0 \leq \Theta_{F,-} < \Theta_{F,+} \leq 1$. Note that ξ takes its minimal value at

$$\Theta_F^* = \frac{(1+m-f)(1+m) - 2mR^f}{(1+m-f)^2}.$$

We have that

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

so that

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

This implies that

$$(1 + m - f)(1 + m) - 2mR^f > m(1 + m) - 2mR^f > 0,$$

and $\Theta_F^* > 0$.

Moreover, it follows that $\Theta_F^* < 1$ if and only if

$$\begin{aligned} (1 + m - f)(1 + m) - 2mR^f &< (1 + m - f)^2 \\ f^2 - (1 + m)f + 2mR^f &> 0. \end{aligned} \tag{C.4}$$

Because $m \geq m_+ > R^f > 1$, we have that

$$f^2 - (1 + m)f + 2mR^f > -(1 + m) + 2mR^f = m\left(2R^f - 1\right) - 1 > 0,$$

and inequality (C.4) holds, so that $\Theta_F^* < 1$. Because in addition $f < f_-$, as we have shown earlier, we conclude that

$$\xi(\Theta_F^*) < 0, \quad 0 < \Theta_F^* < 1. \tag{C.5}$$

We furthermore have that

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

because $m \geq m_+$, with $\xi(0) = 0$ if and only if $m = m_+$. In combination with (C.5), this

means that $\Theta_{F,-} \geq 0$ and $\Theta_{F,-} = 0$ if and only if $m = m_+$. We also know that $\xi(1) \geq 0$, with $\xi(1) = 0$ if and only if $f = 0$ (as mentioned earlier). Combining this with (C.5) implies that $\Theta_{F,+} \leq 1$ and that $\Theta_{F,+} = 1$ if and only if $f = 0$.

We now consider the stability of the steady state of A. If $\Theta_F \leq \Theta_{F,-}$, the eigenvalues of A are real. In the following, we show that the steady state is unstable in this case, meaning $|\lambda_\pm| > 1$. Inequality (C.2) implies that showing $\lambda_- > 1$ is sufficient. We can define λ_- as a real-valued function of $\Theta_F \in [0, \Theta_{F,-}]$. Lemma 3 implies that $\lambda_-(0) > 1$, since $m \geq m_+ > R^f$. Because λ_- is continuous and $\lambda_- \neq 1$ (Lemma 7), it follows that $\lambda_- > 1$ for all $\Theta_F \in [0, \Theta_{F,-}]$. We conclude that the steady state is unstable for $\Theta_F \leq \Theta_{F,-}$.

If $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$, the eigenvalues of A are complex, and the stability of the steady state is again determined by the threshold $\overline{\Theta}_F$. The proof that the steady state is stable if $\Theta_F \geq \Theta_{F,+}$ (implying real eigenvalues) carries over from the case of $m \in (m_-, m_+)$.

Taking all cases together, the steady state is unstable if $\Theta_F < \overline{\Theta}_F$, and stable if $\Theta_F > \overline{\Theta}_F$. To see this, first note that we have shown that the steady state is stable if $m \leq m_-$, while in this case indeed $\Theta_F > \overline{\Theta}_F$, because $\overline{\Theta}_F < 0$. Second, in the case of $m_- < m < m_+$, we have shown that the stability of the steady state depends on $\overline{\Theta}_F$ as described above if $\Theta_F < \Theta_{F,+}$. If $\Theta_F \geq \Theta_{F,+}$, the steady state is stable, while it indeed also holds that $\Theta_F > \overline{\Theta}_F$, since $\Theta_{F,+} > \overline{\Theta}_F$. Third, note that if $m \geq m_+$, we have that the stability of the steady state is determined by $\overline{\Theta}_F$ if $\Theta_{F,-} < \Theta_F < \Theta_{F,+}$. If $\Theta_F < \Theta_{F,-}$ ($\Theta_F > \Theta_{F,-}$), the steady state is unstable (stable), which is consistent with the claim above, since $\Theta_{F,-} < \overline{\Theta}_F < \Theta_{F,+}$.

Furthermore, we have that the complexity of the eigenvalues is determined by Θ_F relative to the $\Theta_{F,\pm}$ as described above, except for the case in which the $\Theta_{F,\pm}$ are complex: If $m \leq m_-$ and $f_- \leq f \leq f_+$, the $\Theta_{F,\pm}$ are complex, and the eigenvalues of A are real. □

Proof of Lemma 5. The representative agent maximizes Φ_t given by equation (A.4).

Omitting time subscripts for clarity and using that $w_M = 1 - w_F$, this reduces to

$$\begin{aligned}\Phi = {} & w_F \left(\langle u_F \rangle - \langle u_M \rangle\right) + \langle u_M \rangle \\ & - \frac{\gamma_R}{2} \left[w_F^2 \, \widetilde{\sigma}_F^2 + 2\, w_F \left(1 - w_F\right) \widetilde{\sigma}_{FM} + \left(1 - w_F\right)^2 \widetilde{\sigma}_M^2\right].\end{aligned}$$

The weights are constrained to the $[0, 1]$ interval:

$$-w_F \leq 0; \tag{C.6a}$$

$$w_F - 1 \leq 0. \tag{C.6b}$$

The first and second derivatives of the objective function with respect to w_F are given by

$$\begin{aligned}\frac{\partial \Phi}{\partial w_F} &= \langle u_F \rangle - \langle u_M \rangle - \gamma_R \left[w_F \, \widetilde{\sigma}_F^2 + \left(1 - 2\, w_F\right) \widetilde{\sigma}_{FM} - \left(1 - w_F\right) \widetilde{\sigma}_M^2\right]; \\ \frac{\partial^2 \Phi}{\partial w_F^2} &= -\gamma_R \left(\widetilde{\sigma}_F^2 - 2\, \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2\right).\end{aligned}$$

The parenthetical expression in the second derivative is just the variance of the difference between heuristics' performance ($\widetilde{\sigma}_F^2 - 2\,\widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2 = V[u_F - u_M]$), which is non-negative. This implies that the objective function is concave (since $\gamma > 0$). Because we also have that the constraints are convex and continuously differentiable with respect to w_F, the Karush-Kuhn-Tucker conditions are sufficient. Denoting the optimum by w_F^* and the multipliers corresponding to constraints (C.6a) and (C.6b) by $\mu_{1,2}$, these conditions are

given by

$$\left.\frac{\partial \Phi}{\partial w_F}\right|_{w_F^*} = \mu_2 - \mu_1; \tag{C.7a}$$

$$-w_F^* \leq 0; \tag{C.7b}$$

$$w_F^* - 1 \leq 0; \tag{C.7c}$$

$$\mu_{1,2} \geq 0; \tag{C.7d}$$

$$-\mu_1 \, w_F^* = 0; \tag{C.7e}$$

$$\mu_2 \, (w_F^* - 1) = 0. \tag{C.7f}$$

We start with the interior solution: When none of the constraints are binding, we have that $\mu_1 = \mu_2 = 0$, and condition (C.7a) gives that

$$\begin{aligned} \gamma_R^{-1}(\langle u_F \rangle - \langle u_M \rangle) &= w_F^* \, \widetilde{\sigma}_F^2 + (1 - 2\, w_F^*) \, \widetilde{\sigma}_{FM} - (1 - w_F^*) \, \widetilde{\sigma}_M^2 \\ w_F^* &= \frac{\gamma_R^{-1} \, (\langle u_F \rangle - \langle u_M \rangle) - \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2}{\widetilde{\sigma}_F^2 - 2\, \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2}. \end{aligned} \tag{C.8}$$

The two constraints cannot be binding at the same time. If (C.6a) is binding, we have $w_F^* = 0$, $\mu_1 > 0$, $\mu_2 = 0$, and condition (C.7a) implies that

$$\gamma_R^{-1} \, (\langle u_F \rangle - \langle u_M \rangle) - \widetilde{\sigma}_{FM} - \widetilde{\sigma}_M^2 < 0$$

Comparing with equation (C.8), we see that this corresponds to the interior solution being negative. Analogously, if constraint (C.6b) is binding, it follows that $w_F^* = 1$, $\mu_1 = 0$, $\mu_2 > 0$, and condition (C.7a) implies that

$$\gamma_R^{-1} \, (\langle u_F \rangle - \langle u_M \rangle) - \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2 > \widetilde{\sigma}_F^2 - 2\, \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2,$$

which corresponds to the interior solution being strictly larger than 1. We conclude that

the optimal weight put on the fundamentalist rule is

$$w_F^* = \min\left[\max\left(\frac{\gamma_R^{-1}\left(\langle u_F\rangle - \langle u_M\rangle\right) - \widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2}{\widetilde{\sigma}_F^2 - 2\,\widetilde{\sigma}_{FM} + \widetilde{\sigma}_M^2}, 0\right), 1\right],$$

with the optimal weight for the momentum rule given by $w_M^* = 1 - w_F^*$. □

Proof of Lemma 6. For simplicity, we write

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

and rewrite the MNL fractions as

$$n_{F,t+1} = \frac{\exp\left(\beta\Delta_t\right)}{\exp\left(\beta\Delta_t\right) + 1},$$

where we have defined $\Delta_t = U_{F,t} - U_{M,t}$. To prove the theorem, we have to show that the representative agent weights satisfy $w_{F,t+1} = n_{F,t+1}$ for arbitrary Δ_t, which would imply that the pricing equations (A.3) and (A.7) for the representative agent market and the MNL switching market are the same.

Because we have that $\widetilde{\sigma}_{FM,t} = 0$, it follows from Lemma 5 that the interior solution to the representative agent maximization problem is given by

$$w_{F,t+1}^* = \frac{\Delta_t}{2\gamma_{R,t+1}\,\sigma^2} + \frac{1}{2}.$$

First note that when $\Delta_t = 0$, it follows that

$$w_{F,t+1} = w_{F,t+1}^* = \frac{1}{2} = n_{F,t+1}.$$

Now suppose that $\Delta_t \neq 0$. We can rewrite the representative agent's risk aversion as

defined in the lemma in terms of Δ_t and σ^2:

$$\gamma_{R,t+1} = \frac{\Delta_t}{\sigma^2} \frac{\exp(\beta\Delta_t) + 1}{\exp(\beta\Delta_t) - 1},$$

which gives that

$$\begin{aligned} w^*_{F,t+1} &= \frac{1}{2\frac{\exp(\beta\Delta_t)+1}{\exp(\beta\Delta_t)-1}} + \frac{1}{2}. \\ &= \frac{\exp(\beta\Delta_t)}{\exp(\beta\Delta_t) + 1} \\ &= n_{F,t+1}. \end{aligned}$$

Since $0 < n_{F,t+1} < 1$, we have that $w_{F,t+1} = w^*_{F,t+1} = n_{F,t+1}$. We conclude that $w_{F,t+1} = n_{F,t+1}$ for arbitrary Δ_t. □

Proof of Lemma 7. We have that $\lambda_\pm = 1$ if and only if

$$\begin{aligned} &\Theta_F f + (1 - \Theta_F)(1 + m) \pm \sqrt{\xi(\Theta_F; R^f, f, m)} = 2R^f \\ &\pm \sqrt{[\Theta_F f + (1 - \Theta_F)(1 + m)]^2 - 4R^f (1 - \Theta_F) m} = 2R^f \\ &\qquad\qquad - \Theta_F f - (1 - \Theta_F)(1 + m). \end{aligned} \tag{C.9}$$

Squaring both sides and rearranging gives

$$R^f = 1 + \Theta_F(f - 1 - m) \leq 1.$$

This is a contradiction, since $R^f > 1$. Since the squares are not equal, equality (C.9) cannot hold. We conclude that $\lambda_\pm \neq 1$. □

Appendix D. Cobweb application

Appendix D.1. The linear cobweb model with diverse risk preferences

In this appendix, we present the implications of our expectation-formation framework in a linear cobweb model. The version of the cobweb model we present here is inspired by Hommes (2013). We use the linear version to focus on the non-linearities arising from the expectation-formation process. The cobweb model represents a market for non-storable goods (e.g., corn), which are produced in the period before they are sold in the market. In each period, suppliers have to decide how much they are going to produce. This decision crucially depends on the price for which they can expect to sell their produced goods in the next period. This means that in period t they have to form expectations about the price in period $t+1$.

Appendix D.1.1. Market mechanism

The supply function of an individual producer is linear and given by $s\,p_t^e$, with $s > 0$ and p_t^e the forecast for the period t price. This supply function is consistent with profit maximization under a quadratic cost function $c(q) = q^2/2s$ (Brock and Hommes, 1997).[26] Now assume that we have a continuum of producers on the unit interval $(0, 1)$. Denoting by $p_t^e(i)$ the price forecast of producer $i \in (0, 1)$, we find that aggregate supply is given by

$$\begin{aligned} S_t &= \int_0^1 s\,p_t^e(i)\,di \\ &= s\,\overline{p}_t^e, \end{aligned}$$

expressed in terms of the average of all producers' forecasts: $\overline{p}_t^e = \int_0^1 p_t^e(i)\,di$.[27]

Aggregate demand is given by $D(p_t) = a - d\,p_t$, where $a, d > 0$. By adding the market

[26] Expected profit in period t is given by $\Pi_t^e(q) = q\,p_t^e - q^2/2s$. We have FOC $p_t^e - q^*/s = 0$, which gives $q^* = s\,p_t^e$. Because $\Pi_t^e(q)$ is strictly concave, q^* is a maximum.

[27] This identity readily extends to the discrete case, with the integral replaced by a summation.

clearing condition, we get the system that governs the dynamics of the model:

$$D(p_t) = a - d\,p_t, \quad a, d > 0, \tag{D.1a}$$

$$S(\overline{p}_t^e) = s\,\overline{p}_t^e, \quad s > 0, \tag{D.1b}$$

$$D(p_t) = S(\overline{p}_t^e). \tag{D.1c}$$

Appendix D.1.2. Beliefs

We investigate the dynamics when producers can choose between two strategies: rational and naive. The rational strategy requires complete information about the model and all the other agents' expectations, and perfectly predicts the price in the next period. The naive strategy predicts that the price in the next period will be the same as the last observed price. Using the notation $p^e_{R,t+1}$ for the rational and $p^e_{N,t+1}$ for the naive forecast, we have that

$$p^e_{R,t+1} = p_{t+1}; \tag{D.2a}$$

$$p^e_{N,t+1} = p_t. \tag{D.2b}$$

The choice of strategy is governed by the same framework as in the asset pricing model: Suppliers use the mean-variance performance measure defined in equation (15) to find the strategy that best balances performance and risk. Different from the asset pricing model, a strategy's performance measure does not only depend on the squared forecasting errors, but also on the cost of using the strategy. We assume that use of the naive rule is free, but that the rational rule has a cost C_R. The interpretation is that resources are needed to gather information about the structure of the model and the expectations of the other suppliers. The specification of these strategies, including the information costs of the rational rule is taken from Brock and Hommes (1997).

Note that the rational rule always has zero forecasting errors, meaning that its utility is always equal to minus its cost: $u_{R,t} = u_R = -C_R$. The utility for the naive rule is

given by

$$u_{N,t} = -\left(p_t - p_{t-1}\right)^2.$$

The average utility, average squared utility and variance are calculated in the same way as in the asset pricing model (equations 14a, 14b, and 14c). Note that the rational rule has zero variance, since its utility is constant.

The following result establishes a link between squared forecasting errors and profits, providing support for our choice of utility:

Lemma 8. *In the the linear cobweb model with cost function $c(q) = q^2/2s$, the squared forecasting error is equal to the difference between maximum profit and realized profit, up to a constant. They are equal for $s = 2$.*

Proof. Maximum profit is given by

$$\begin{aligned}\Pi_t^* &= s\,p_t^2 - \frac{(s\,p_t)^2}{2s} \\ &= \frac{s}{2}\,p_t^2,\end{aligned}$$

while realized profit is given by

$$\begin{aligned}\Pi_t &= s\,p_t^e\,p_t - \frac{(s\,p_t^e)^2}{2s} \\ &= \frac{s}{2}\left[2\,p_t^e\,p_t - (p_t^e)^2\right].\end{aligned}$$

It follows that

$$\begin{aligned}\Pi_t^* - \Pi_t &= \frac{s}{2}\left[p_t^2 - 2\,p_t^e\,p_t + (p_t^e)^2\right] \\ &= \frac{s}{2}\left(p_t - p_t^e\right)^2,\end{aligned}$$

which is $s/2$, a constant, times the squared forecasting error. The constant equals 1 for $s = 2$. □

We are now in the position to determine the period t fractions of producers that use the forecasting strategies. Since they have to form one-period ahead forecasts, these fractions are determined in the previous period $(t-1)$ and based on information from that period. Similar to the asset pricing model, we can determine the risk aversion coefficient for which an agent would be indifferent between the two rules in period $t-1$, again denoted by $\overline{\gamma}_{t-1}$. Because the rational rule is riskless, suppliers with risk aversion below the threshold $\overline{\gamma}_{t-1}$ will choose the naive rule, and those with risk aversion above the threshold will choose the rational rule. If $\widetilde{\sigma}_{N,t-1} \neq 0$, it follows that

$$\overline{\gamma}_{t-1} = 2\,\frac{\langle u_N \rangle_{t-1} + C_R}{\widetilde{\sigma}^2_{N,t-1}}.$$

The fraction of producers using the naive rule in period t is then given by

$$n_{N,t} = G\left(\overline{\gamma}_{t-1}\right),$$

with G the empirical cumulative distribution function of risk aversion. The fraction of rational producers is given by

$$n_{R,t} = 1 - G\left(\overline{\gamma}_{t-1}\right).$$

If $\widetilde{\sigma}_{N,t-1} = 0$, the fractions are only determined by $\langle u_N \rangle_{t-1}$: If $\langle u_N \rangle_{t-1} < -C_R$, we have $n_{N,t} = 0$, if $\langle u_N \rangle_{t-1} > -C_R$, we have $n_{N,t} = 1$, and if $\langle u_N \rangle_{t-1} = -C_R$, we have $n_{N,t} = n_{R,t} = 1/2$.

The average forecast in period t is given by

$$\overline{p}^e_t = n_{R,t}\, p_t + n_{N,t}\, p_{t-1}. \tag{D.3}$$

Combining equations (D.1a) – (D.3), we find that the clearing price is given by

$$p_t = \frac{a - n_{N,t}\, s\, p_{t-1}}{n_{R,t}\, s + d}.$$

The steady state price is given by

$$p^* = \frac{a}{s + d}, \tag{D.4}$$

where we have used that $n_{R,t} + n_{N,t} = 1$. Note that this steady state price is constant, contrary to the fundamental price that we derived in the asset pricing model.

For $a = 0$, the steady state is 0 and the model reduces to the deviations-from-steady-state model of Muth (1961). In some of our simulations, we include an aggregate supply shock ϵ_t, due to weather conditions for example. In that case, aggregate supply (D.1b) is replaced by

$$S(\overline{p}_t^e) = s\, \overline{p}_t^e + \epsilon_t,$$

so that the market clearing price becomes

$$p_t = \frac{a - n_{N,t}\, s\, p_{t-1} - \epsilon_t}{n_{R,t}\, s + d}.$$

The steady state is the same if the supply shock has zero mean.

We denote the phase space by $Z = \mathbb{R} \times \mathbb{R}_{\leq 0} \times \mathbb{R}_{\geq 0}$. The evolution of the resulting 3-dimensional discrete dynamical system is governed by the one-step map

$$\psi : Z \to Z, \quad \begin{pmatrix} p_t \\ \langle u_N \rangle_t \\ \langle u_N^2 \rangle_t \end{pmatrix} \mapsto \begin{pmatrix} p_{t+1} \\ \eta \langle u_N \rangle_t - (1-\eta)\,(p_{t+1} - p_t)^2 \\ \eta \langle u_N^2 \rangle_t + (1-\eta)\,(p_{t+1} - p_t)^4 \end{pmatrix}. \tag{D.5}$$

If $\widetilde{\sigma}_{N,t} \neq 0$, next period's price p_{t+1} is given by

$$p_{t+1} = \frac{a - G\left(2\,\dfrac{\langle u_N\rangle_t + C_R}{\langle u_N^2\rangle_t - \langle u_N\rangle_t^2}\right) s\, p_t - \epsilon_t}{\left[1 - G\left(2\,\dfrac{\langle u_N\rangle_t + C_R}{\langle u_N^2\rangle_t - \langle u_N\rangle_t^2}\right)\right] s + d}.$$

If $\widetilde{\sigma}_{N,t} = 0$, it is given by

$$p_{t+1} = \begin{cases} \dfrac{a - \epsilon_t}{s + d} & \text{if } \langle u_N\rangle_t < -C_R; \\ \dfrac{a - \dfrac{s\, p_t}{2} - \epsilon_t}{\dfrac{s}{2} + d} & \text{if } \langle u_N\rangle_t = -C_R; \\ \dfrac{a - s\, p_t - \epsilon_t}{d} & \text{if } \langle u_N\rangle_t > -C_R. \end{cases}$$

Appendix D.2. Dynamics

Appendix D.2.1. Dynamics with constant fractions

Before introducing endogenous switching between the two forecasting strategies into the Cobweb model, we study its properties when the fractions of rational and naive agents are constant. This will later help us understand the dynamics of the full model.

Lemma 9. *When the fractions of rational and naive producers are constant, the clearing price in the linear cobweb model evolves according to*

$$p_t = \left(-\frac{n_N\, s}{n_R\, s + d}\right)^t (p_0 - p^*) + p^*,$$

with p_0 the initial price and p^* the steady state price given by equation (D.4). Stability of the steady state depends on the fraction of rational producers relative to a threshold given by

$$\overline{n}_R = \frac{s - d}{2s}.$$

When $n_R > \overline{n}_R$, the steady state is globally asymptotically stable, meaning that it is

stable and everywhere attracting. For $n_R = \overline{n}_R$, the steady state is neutral and the system follows a two-cycle. When $n_R < \overline{n}_R$, the steady state is unstable and all solutions are unbounded, except for the steady state.

Proof. The clearing price in period t is given by

$$p_t = \frac{a}{n_R\, s + d} - \frac{n_N\, s}{n_R\, s + d}\, p_{t-1}.$$

The homogeneous problem

$$p_t = -\frac{n_N\, s}{n_R\, s + d}\, p_{t-1}$$

has general solution

$$p_t = \left(-\frac{n_N\, s}{n_R\, s + d}\right)^t c,$$

with c some constant. Knowing that the steady state p^* is a particular solution for the full, inhomogeneous problem, we can add it to the general solution for the homogeneous problem and solve for c to get

$$p_t = \left(-\frac{n_N\, s}{n_R\, s + d}\right)^t (p_0 - p^*) + p^*.$$

The stability of the steady state changes at the threshold value $\overline{n}_R$ given by

$$\begin{aligned} \frac{\overline{n}_N\, s}{\overline{n}_R\, s + d} &= 1 \\ (1 - \overline{n}_R)s &= \overline{n}_R\, s + d \\ \overline{n}_R &= \frac{s - d}{2s}. \end{aligned}$$

For $n_R > \overline{n}_R$, the steady state is stable and globally attracting. When $n_R = \overline{n}_R$, the clearing price jumps back and forth between p_0 and $2p^* - p_0$ in a two-cycle, indicating metastability. For $n_R < \overline{n}_R$, lastly, $|p_t - p^*|$ increases indefinitely when $p_0 \neq p^*$. □

In the special case that all producers use the rational strategy at all times ($n_{R,t} =$

$n_R = 1$), the system is always at the steady state. Even when the initial price p_0 is far from the steady state price, the system immediately converges to the steady state.

When all producers are using the naive rule ($n_{N,t} = n_N = 1$), the model's dynamics depends on the relation between supply and demand parameters s and d. In the case that $s < d$, $\overline{n}_R < 0$ and $n_R > \overline{n}_R$: the system converges to the steady state, while price deviations alternate between positive and negative. The intuition is that after an initial price above the steady state, suppliers expect the same deviation in the next period, and produce more than steady state supply $s\,p^*$. This oversupply then leads to a price below steady state in the following period. Because $s < d$ however, the increase in supply is smaller than needed to reach the same absolute deviation from steady state, and the new price is closer to the steady state price. Next, the suppliers expect the same low price level, and decide to produce less than steady state supply. This again leads to an above-steady-state price that is yet closer to the steady state. This process of ever smaller deviations from steady-state supply continues until the steady state is reached.

When $s = d$, the clearing price jumps back and forth between p_0 and $2p^* - p_0$ in a two-cycle: The over or undersupply compared to steady state is precisely enough to reach the same price deviation in the next period, but on the other side of the steady state. Lastly, in the case $s > d$, the steady state is unstable, and deviations from the steady state become arbitrarily large.

Note that in general the steady state is globally asymptotically stable when $d > s$, since in that case $\overline{n}_R < 0$, while $n_R \geq 0$.

Appendix D.2.2. Dynamics with endogenous fractions

We now allow for endogenous switching between the two forecasting strategies. We set the cost of using the rational strategy to 1 (i.e., $C_R = 1$). We use numerical methods to study the systems' dynamics. We provide intuitive explanations in terms of the theoretical results presented in Lemma 9. We present the results on the price dynamics

in terms of deviations from the steady state, denoted by x_t:

$$x_t = p_t - p^*.$$

In all our simulations, we set $d = 1$ and $s = 2$. This means that the squared forecasting error is equal to the profit deviations from the maximum (Lemma 8). Moreover, it implies that $s > d$: The steady state is unstable under the naive rule, and the threshold fraction of rational producers is given by $\overline{n}_R = 0.25$. The interaction between the stabilizing and destabilizing forces of the rational and naive suppliers will give rise to rich dynamics: stable in some parameter ranges and chaotic in others.

We start with the deterministic model, with zero supply shock (i.e., $\epsilon_t = 0$). When producers have substantial memory ($\eta = 0.9$), chaotic dynamics arise, as can be seen in the two top-left panels of Figure D.7. The system starts close to the steady state, with the fraction of rational producers above the threshold value $\overline{n}_R = 0.25$, so that the steady state is stable and attracting (Lemma 9). The fraction of rational agents quickly drops below the threshold, leading to an unstable steady state and increasing volatility in the clearing price. When the forecasting errors of the naive rule become too large, producers begin to switch back to the rational rule, starting with those who have the largest rate of risk aversion. As soon as the fraction of rational producers exceeds the threshold of 0.25, the system starts to stabilize, and the producers gradually switch back to the naive rule to save costs (starting with those with lowest γ).

A key characteristic of this chaotic solution is that not all producers switch to the rational rule during volatile periods: Both forecasting strategies are used at all times. We can identify two reasons for this behaviour. First, there are producers left with low enough risk aversion to stick with the naive rule. Second, there is significant persistence in the performance measures of the naive rule (because $\eta = 0.9$), meaning that the larger volatility only gradually affects the mean-variance utility of that rule. The result is that the system never reaches the steady state. When the fraction of rational producers drops

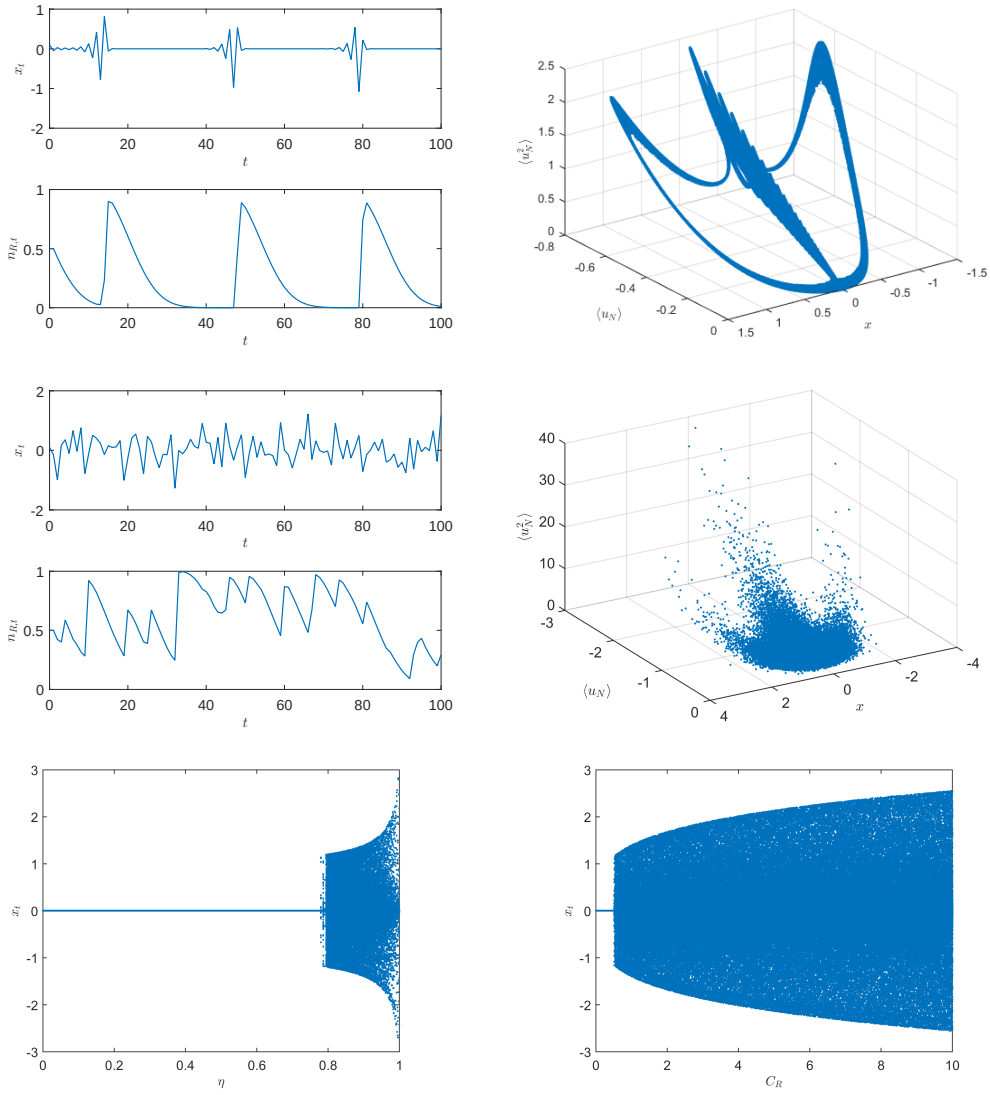

Figure D.7: *Two top-left panels:* Price deviations from the steady state (top) and fraction following rational strategy (bottom) for $\eta = 0.9$, $C_R = 1$, and $\epsilon_t = 0$. *Top-right panel:* Strange attractor for $\eta = 0.9$, $C_R = 1$, and $\epsilon_t = 0$. *Two mid-left panels:* Price deviations from the steady state (top) and fraction following rational strategy (bottom) for $\eta = 0.9$, $C_R = 1$, and with i.i.d. aggregate supply shock $\epsilon_t \sim \mathcal{N}(0, 1)$. *Mid-right panel:* Noisy attractor for $\eta = 0.9$, $C_R = 1$, and with i.i.d. aggregate supply shock $\epsilon_t \sim \mathcal{N}(0, 1)$. *Bottom-left panel:* Long-term behaviour of the price deviations from the steady state for different values of the memory parameter η, keeping C_R at 1. *Bottom-right panel:* Long-term behaviour of the price deviations from the steady state for different values of the cost of the rational rule C_R, keeping η at 0.9. *Parameter values:* $a = 0$, $d = 1$, $s = 2$, $\log \gamma \sim \mathcal{N}(1.14, 0.67)$. *Initial values:* $p_0 = 0.10$, $\langle u_N \rangle_0 = -0.50$, $\langle u_N^2 \rangle_0 = 0.57$ (implying $n_{R,1} = 0.50$).

below 0.25 after a volatile episode, the steady state becomes unstable, and price deviations from the steady state start increasing again, giving rise to a new volatile episode: Stable periods sow the seeds of volatile periods. This cycle repeats itself indefinitely, each time in a slightly different way.

For the specification used in the top-right panel of Figure D.7, we find the following Lyapunov exponents: $\lambda_1 = 0.0749$, $\lambda_2 = -0.1049$, $\lambda_3 = -0.1342$ (Benettin et al., 1980). The positive maximal exponent indicates that the dynamics is indeed chaotic. This system is more sensitive to initial conditions than the asset pricing system that we studied in the main body of this paper, which had a maximal exponent of $\lambda_1 = 0.0035$. We present the strange attractor for this solution in the top-right panel of Figure D.7. It has a box-counting dimension of approximately 1.6.

If we include a standard normal, i.i.d. aggregate supply shock $\epsilon_t \sim \mathcal{N}(0, 1)$, price deviations from the steady state become larger, and the rational rule becomes more important because of the added volatility (see the mid-left panels of Figure D.7). The strange attractor of the deterministic system is still discernible in the attractor of this noisy solution, which is larger and less well-defined, as we show in the bottom-right panel of Figure D.7.

To explore how sensitive the dynamics is to changes in the model parameters, we plot the long-term behaviour of the price deviations from the steady state for different values of η and C_R in the bottom panels of Figure D.7. For each value of the parameter, we plot the last 1,000 price deviations of 2,000 period simulation. For $\eta = 0.5$, the system stays at the steady state after the 1,000 initialization periods. For a value of η slightly above 0.8, the system becomes chaotic, and the long-term behaviour has no clear structure. The more producers remember, the larger the deviations from the steady state become, as is reflected in the increasing bandwidth for higher values of η. This can be explained by the fact that a larger η means that the utility of the naive rule is affected more gradually in times of high volatility, meaning that it takes longer for producers to switch to the rational rule and that there is therefore more time for the system to move away from the

steady state.

We see a similar threshold value for the cost of the rational rule (C_R). As long as the cost is low enough, the naive rule cannot compete with the rational rule, and the system converges to the steady state eventually. However, at a cost slightly below 1, the dynamics becomes chaotic. As the cost becomes higher, the naive rule becomes relatively more attractive, and the price deviations become larger.