VU Research Portal

Comparing Behavioural Heterogeneity Across Asset Classes

ter Ellen, Saskia; Hommes, Cars; Zwinkels, Remco

published in

Journal of Economic Behavior and Organization
2021

DOI (link to publisher)

[10.2139/ssrn.2944022](https://doi.org/10.2139/ssrn.2944022)
[10.1016/j.jebo.2019.12.013](https://doi.org/10.1016/j.jebo.2019.12.013)

document version

Publisher's PDF, also known as Version of record

document license

Article 25fa Dutch Copyright Act

[Link to publication in VU Research Portal](#)

citation for published version (APA)

ter Ellen, S., Hommes, C., & Zwinkels, R. (2021). Comparing Behavioural Heterogeneity Across Asset Classes. *Journal of Economic Behavior and Organization*, 185, 747-769. <https://doi.org/10.2139/ssrn.2944022>, <https://doi.org/10.1016/j.jebo.2019.12.013>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Comparing behavioural heterogeneity across asset classes[☆]

Saskia ter Ellen^{a,*}, Cars H. Hommes^{b,c}, Remco C.J. Zwinkels^d

^a Norges Bank, Bankplassen 2, Oslo 0151, Norway

^b University of Amsterdam, Roetersstraat 11, Amsterdam 1018 WB, the Netherlands

^c Bank of Canada, 234 Wellington St W, Ottawa, ON K1A 0G9, Canada

^d Vrije Universiteit Amsterdam, De Boelelaan 1105, Amsterdam 1081 HV, the Netherlands

ARTICLE INFO

Article history:

Received 13 May 2019

Revised 13 December 2019

Accepted 14 December 2019

Available online 5 February 2020

JEL classification:

C24

E31

G12

G15

Keywords:

Financial markets

Heterogeneous expectations

Market stability

ABSTRACT

We estimate an endowment-based asset pricing model in which agents have heterogeneous and time-varying beliefs about the future price on a range of asset classes. This gives insight into the extent behaviour differs across assets, and what this implies for market stability. We find evidence for behavioural heterogeneity for all asset classes but equity. Heterogeneity is especially large and persistent in asset classes for which limits to arbitrage are more binding. In less constrained (financial) markets, agents update their beliefs more frequently. Consequently, the probability of behavioural bubbles and crashes is substantially higher in macroeconomic asset classes than in financial asset classes.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

There is ample empirical evidence showing that individual beliefs are not consistent with rational expectation representative agent type models (e.g. Greenwood and Shleifer, 2014).¹ Although it has often been argued that such inconsistencies average out at the market level, there are numerous theoretical studies suggesting that this is not necessarily the case, especially when there is a sufficiently large group of non-fundamental traders (e.g. Cutler et al., 1990; DeLong et al., 1990; Lux, 1995; Barberis et al., 1998; Brock and Hommes, 1997; 1998).² Limitations to the rational expectations hypothesis are also suggested by the fact that market prices display dynamics such as volatility clustering, heavy tails, and bubbles and

[☆] This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We thank participants of the Research in Behavioral Finance Conference (RBFC), Society for Nonlinear Dynamics and Econometrics (SNDE), Computing in Economics and Finance (CEF), Workshop on Economic Heterogeneous Interaction Agents (WEHIA), and the International Symposium on Computational Economics and Finance, as well as seminar participants at Norges Bank.

* Corresponding author.

E-mail addresses: saskia.ter-ellen@norges-bank.no (S. ter Ellen), c.h.hommes@uva.nl (C.H. Hommes), r.zwinkels@vu.nl (R.C.J. Zwinkels).

¹ Greenwood and Shleifer (2014) study six different surveys on stock price expectations. Similar conclusions have been drawn based on, e.g. inflation surveys (Branch, 2004) and stock, bond, and foreign exchange surveys (MacDonald, 2000).

² Cutler et al. (1990) propose a model with fundamental traders alongside feedback traders, and show that the latter group of traders affects prices. DeLong et al. (1990) show that a group of noise traders of sufficient size in an otherwise rational market will affect prices and survive in the long run. Lux (1995) models a group of traders that are not fully informed about the fundamental, and copy the behaviour of other traders, so-called herding behaviour. Barberis et al. (1998) set up a model in which over- and underreaction to news co-exist in time-varying proportions, and form a possible explanation for

crashes, often referred to as the ‘stylised facts’ of financial markets. Traditional asset pricing models based on a rational representative agent can generally not explain such dynamics.

Several alternative explanations have been brought forward, often relaxing various assumptions underlying these traditional models, such as (in)complete information, (no) limits to arbitrage, and regarding risk perception, preferences, or expectation formation, to name a few. Many of these explanations require one to deviate from the assumption of full rationality, creating scope for heterogeneity between agents. A number of studies show that expectation formation can be summarised by certain rules of thumb; a large part of the experimental literature in this field is surveyed by [Assenza et al. \(2014\)](#). [Bloomfield and Hales \(2002\)](#) show that participants of an experiment switch between a trend following and a mean reverting rule when forecasting firm earnings, conditional on the recent price realisations. [MacDonald \(2000\)](#) summarises the main findings from studying quantitative surveys for stock, bond, and foreign exchange markets, and finds that most survey participants use forecasting rules with short term trend extrapolation and long term mean reversion. [Greenwood and Shleifer \(2014\)](#) find evidence for trend extrapolation in stock market expectations.

A stream of research, triggered by papers by [Brock and Hommes \(1997, 1998\)](#), have been quite successful at replicating the stylised facts of financial markets (including the aforementioned volatility clustering and heavy tails) using an approach with heterogeneity in expectations, taking into account the experimental and survey-based evidence of expectation formation. [Lof \(2014\)](#) shows that considering such ‘puzzling’ market dynamics as emergent properties of a system composed of heterogeneous agents indeed generates more realistic outcomes than allowing for time-variation in the discount factor. At the same time, these models also incorporate more realistic expectation formation processes. The empirical literature on heterogeneous agent models such as those developed in [Brock and Hommes \(1997, 1998\)](#) and [Lof \(2014\)](#), has by now established ample evidence on the importance of behavioural heterogeneity for explaining financial market dynamics.³

In this paper, we use a model in which agents have heterogeneous expectations about future prices, and estimate it on a broad set of assets to get a better understanding of differences in agent behaviour and corresponding market stability across asset classes. Various versions of the model, with ranging numbers and types of agents, different profit and/or switching functions, have been estimated successfully on a large number of individual asset classes before. So far, the literature has concluded that allowing for heterogeneity and switching in the expectation formation process is important when describing the dynamics in individual markets. Due to the variation in the functional forms of expectation formation rules and the switching functions, however, the results have, to the best of our knowledge, never been directly compared with each other across many different assets.⁴ We therefore adopt a generalised version of the expectation formation rules and switching function allowing us to compare results across assets. Specifically, as in [Brock and Hommes \(1997, 1998\)](#) and [De Grauwe and Grimaldi \(2006\)](#) we take a model with mean-variance utility and two types of agents with heterogeneous price expectations, fundamentalists and chartists. Fundamentalists expect the market price to mean-revert towards the (perceived) fundamental value, and chartists expect the price to move away from its fundamental value. In addition, agents are allowed to switch between types conditional on recent (relative) performance.⁵ Due to the self-referential nature of asset markets, the law of motion of the market price changes with the switching behaviour of the agents.

We estimate the model on a data set consisting of 220 quarterly observations from 1960 to 2015 of equity prices, foreign exchange rates, commodity prices, real estate prices, and inflation rates.⁶ The results can be compared across assets based on three main parameters. The intensity of choice parameter describes how fast agents update their expectation formation model, i.e., how quickly agents switch between the fundamentalist and chartist group conditional on past performance. The other two coefficients of interest are the fundamentalist mean reversion and the chartist extrapolation coefficients. The difference between the fundamentalists and chartists governs the heterogeneity between market participants. Combined, the three coefficients determine the stability of the system, and therefore the sensitivity of the market to behavioural bubbles and crashes. Both heterogeneity and switching contribute to the instability of markets. We find that switching between types is more prevalent in (financial) markets that are populated by sophisticated investors and in which there are few limits to arbitrage, such as foreign exchange markets. Heterogeneity, however, is especially large and persistent for asset classes with more binding limits to arbitrage such as house prices, causing those to be especially prone to behavioural bubbles and crashes.

Our paper makes an economic as well as methodological contribution. With the analysis conducted in this paper, we get a better understanding of which type of markets and market characteristics are more prone to bubbles, what the market characteristics are that drive that bubble sensitivity, and how investor behaviour contributes to that. In such a way, the paper contributes to the academic literature on heterogeneous agent models, while at the same time offering a different

the momentum and mean-reversion patterns in financial markets. [Brock and Hommes \(1997\)](#) introduce the concept of adaptively rational expectations, in which agents rationally choose from a set of expectation functions in a cobweb type model. In an example with two types, rational and naive, the authors show that under certain conditions highly irregular equilibrium prices converge to a strange attractor.

³ For an overview of the empirical literature on heterogeneous agent models, see [Chen et al. \(2012\)](#), [ter Ellen and Verschoor \(2018\)](#) and [Lux and Zwinkels \(2018\)](#).

⁴ It should be noted though that [Westerhoff and Franke \(2012\)](#) compare results for stock and foreign exchange markets.

⁵ This distinction is in line with the experimental results of, among others, [Bloomfield and Hales \(2002\)](#), who show that participants switch between a trend following and a mean reverting rule, conditional on the recent price realisations.

⁶ For notational convenience, we will also refer to inflation as an ‘asset class’ in the remainder of the paper, even though we use data on the consumer price index rather than tradeable inflation derivatives, such as inflation swaps or inflation protected bonds. Inflation, though, is also typically modelled as the result of underlying inflation expectations.

perspective on certain policy relevant issues. In this respect, our finding that the macroeconomic assets are especially prone to bubbles are interesting, because the focus of policy makers tends to be on the (excessively) volatile financial markets.

In order to be able to make an economically sensible comparison across markets, this paper also makes a methodological contribution. The function governing the switching between types is akin to the smooth transition autoregressive (STAR) models in financial econometrics; see Teräsvirta (1994) and Dijk et al. (2002). In STAR models, it is assumed that the behaviour of the variable is conditional on the value of a transition variable. Comparison across transition variables, however, is hard because the variable is typically not unit free, and as such, neither are the estimated coefficients. We propose a normalised transition variable, that is unit free. As a result, the estimated coefficients also become unit free and therefore comparable across space and time. This does not only allow us to compare the switching intensity for different asset markets, but is also helpful for the estimation procedure, as it makes estimation results less sensitive to outliers or specific episodes.

The remainder of this paper is organised as follows. Section 2 introduces the generic heterogeneous agent model and studies the stability of the model and the estimation procedure. Section 3 covers a description of the data we use and the methodology we apply to estimate the model. Results for various specifications are presented and discussed in Section 4. Section 5 concludes.

2. Heterogeneous agent model and market stability

The papers by Brock and Hommes (1997, 1998) have triggered a stream of research attempting to replicate the ‘stylised facts’ of financial markets (including the aforementioned volatility clustering and heavy tails) using the heterogeneity approach. For example, De Grauwe and Grimaldi (2006) systematically demonstrate how a model with heterogeneous expectations is able to replicate all stylised facts of foreign exchange markets. Furthermore, the heterogeneity approach has been proposed as an explanation for several puzzles in (financial) economics, such as the UIP-puzzle in exchange rate economics; see e.g. Spronk et al. (2013). Heterogeneous agent models have been extensively confronted with all sorts of data. Especially stock markets (Boswijk et al., 2007; Hommes and in ‘t Veld, 2017; Chiarella et al., 2014; Lof, 2014) and foreign exchange markets (Frankel and Froot, 1990; De Jong et al., 2010) have been extensively analysed, but the model has also showed to be useful to explain the price dynamics in for instance housing markets (Kouwenberg and Zwinkels, 2014; Bolt et al., 2019), option markets (Frijns et al., 2010), commodity markets (ter Ellen and Zwinkels, 2010; Baur and Glover, 2014; Westerhoff and Reitz, 2005), credit markets (Chiarella et al., 2015) and collectibles such as wine (Fernandez-Perez et al., 2019). Also the process of macroeconomic variables, such as inflation, can be described by a heterogeneous agent model, as in Cornea-Madeira et al. (2019). Although there is plenty of empirical evidence for a plethora of different markets, it is not straightforward to compare the results of these studies with one another. The main reason for the limited comparability of behaviour across papers, is the variation in functional forms of the expectation formation rules. Consequently, the parameter governing the switching between expectation formation rules also cannot be compared across papers since it is not unit free. In this section we introduce a heterogeneous agent model that is applicable to a broad range of asset classes, and contains a unit free switching parameter.

2.1. The model

We will now introduce our model, which is based on Brock and Hommes (1997, 1998, hereafter BH98). Assume an economy with a single risky asset with price p_t that pays a stochastic dividend y_t .⁷ Wealth then evolves according to:

$$W_{t+1} = RW_t + (p_{t+1} + y_{t+1} - Rp_t)z_t, \quad (1)$$

in which R is 1 plus the risk free rate and z_t the demand for the risky asset.

Investors are mean-variance optimisers such that their demand for the risky asset solves:

$$\text{Max}_z \{E_{ht} W_{t+1} - (a/2)V_{ht}(W_{t+1})\}, \quad (2)$$

in which a is the risk aversion parameter, which is the same for all agents, and V_{ht} the variance of wealth. Solving (2) yields an optimal demand for the risky asset z equal to:

$$z_{ht} = E_{ht}(p_{t+1} + y_{t+1} - Rp_t)/a\sigma^2. \quad (3)$$

Assume that there are H groups of investors, who have heterogeneous expectations about the future price p_{t+1} but homogeneous expectations about the future dividend y_{t+1} .⁸ Now assume that n_{ht} is the fraction of type h investors in period t , with $\sum_{h=1}^H n_h = 1$ and $n_h \in (0, 1)$. Total demand for the risky asset is then given by:

$$\sum_{h=1}^H n_{ht} \{E_{ht}(p_{t+1} + y_{t+1} - Rp_t)/a\sigma^2\}. \quad (4)$$

⁷ Note that y_t may be equal to zero. For (consumption) assets that do not pay out dividends (such as commodities) y_t can be thought of as a convenience yield, i.e. the benefit of physically holding the asset, for example in times of scarcity. In the case of housing, y_t may be seen as rents (as discussed in Dieci and Westerhoff, 2016; Schmitt and Westerhoff, 2019), or the interest rate on the foreign bond in case of foreign currency.

⁸ This is motivated by the fact that prices are much harder to forecast than dividends; companies have the tendency to keep dividends as stable as possible, see e.g. Leary and Michaely (2011).

Without loss of generality, we can put the outside supply of the risky asset to zero, such that:⁹

$$\sum_{h=1}^H n_{ht} \{E_{ht}(p_{t+1} + y_{t+1} - Rp_t) / a\sigma^2\} = 0, \quad (5)$$

and

$$Rp_t = \sum_{h=1}^H n_{ht} E_{ht}(p_{t+1} + y_{t+1}). \quad (6)$$

The fundamental price as perceived by the agents in the model is given by p_t^* . In case of a dividend paying asset, the fundamental price can be thought of as the discounted cashflow $p^* = \frac{\bar{y}}{(R-1)}$ when the dividend process is i.i.d. with mean $\bar{y}$. It is then convenient to write the model in terms of deviations from the fundamental, $x_t = p_t - p_t^*$.

Under the assumption that all beliefs of the groups in H are of the form:

$$E_{ht}(p_{t+1} + y_{t+1}) = E_t(p_{t+1}^* + y_{t+1}) + f_h(x_{t-1}, \dots, x_{t-L}). \quad (7)$$

Eq. (6) can then be written as the weighted average expectation of the heterogeneous agents in the market:

$$Rx_t = \sum_{h=1}^H n_{ht} f_h(x_{t-1}, \dots, x_{t-L}). \quad (8)$$

Consistent with the literature on heterogeneous agents and motivated by the experimental literature, we assume two types of agents, fundamentalists and chartists. The fundamentalists expect the market price to converge to the fundamental value, and thus x_t to converge to zero. Hence, $f_F = \phi_F x_{t-1}$ with $0 < \phi_F < 1$. As such, fundamentalists exhibit stabilising behaviour as they bring the market price closer to the fundamental value. Chartists, on the other hand, exhibit destabilising behaviour by expecting the deviation between price and fundamental to increase. Hence, $f_C = \phi_C x_{t-1}$ with $\phi_C > 1$. Given that both groups use the same source of information when forming expectations (i.e., x_{t-1}), heterogeneity is measured by the difference in coefficients $(\phi_C - \phi_F)$.¹⁰

The full pricing equation is then given by:

$$Rx_t = n_{Ft} \phi_F x_{t-1} + n_{Ct} \phi_C x_{t-1}. \quad (9)$$

Given that both groups use the same information set, agents are able to switch between groups conditional on the relative performance of the groups at zero information costs. The higher the performance of one group is compared to the other, the more likely it is that more agents will switch to this rule. This assumption is again consistent with the experimental results of, e.g., Bloomfield and Hales (2002). We first define the performance measure π_{ht} . We assume that agents base their choice on the relative ability of the groups to forecast x_t over the previous I periods. Specifically:

$$\pi_{ht} = \sum_{i=1}^I |x_{t-i} - \phi_h x_{t-i-2}|, \quad (10)$$

in which I is the memory parameter. In the benchmark setting, we set $I = 1$, such that agents only consider the most recent forecast error, but we study the sensitivity of the empirical results to this choice in Section 4.1. Note that in this definition, π_{ht} is a negative function of performance.

It might be that some agents are faster than other in switching between groups. Therefore, as in BH98, n_{ht} is endogenously determined by means of a multinomial logit function given by:

$$n_{ht} = \exp\left(-\beta \frac{\pi_{ht}}{\pi_{Ft} + \pi_{Ct}}\right) / Z_t \quad (11)$$

$$Z_t = \sum_{F,C} \exp\left(-\beta \frac{\pi_{ht}}{\pi_{Ft} + \pi_{Ct}}\right), \quad (12)$$

which simplifies to:

$$n_{Ft} = \left(1 + \exp\left(\beta \frac{\pi_{Ft} - \pi_{Ct}}{\pi_{Ct} + \pi_{Ft}}\right)\right)^{-1} \quad (13)$$

⁹ This may not be the case for all asset classes that we consider in this paper. However, the resulting model is observationally equivalent to assuming risk-neutral investors, as in Schmitt and Westerhoff (2019), in the case of time-varying risk aversion a , risk σ , or outside supply, which lead to time-varying risk premia. In the situation these elements are constant, the expectation formation function E_{ht} can be interpreted as a price impact function since the (constant) a and σ will only be scaling the expectations.

¹⁰ Note that when estimating the model, it is possible that the parameters are different from each other but both smaller than one. In that case, both groups of agents have mean reverting beliefs, but one group expects the reversion to be faster than the other group. In such a case, we refer to the group with the relatively fast (slow) expected mean reversion as fundamentalists (chartists).

(a) Absolute prediction errors

(b) Relative prediction errors

Fig. 1. Absolute versus relative prediction errors. *Notes:* This figure shows how smooth the transition between two types of agents is for different values of β and a range of performance function values. In both graphs, the x-axis denotes the value of the performance function and the y-axis the proportion of agents of type 1. The lines represent different values for β , ranging from 0 (the horizontal line at 0.5) to 100 (the steepest line) with steps of 1. The left graph shows the transition for different values of the absolute performance function, the right graph for the relative performance function.

$$n_{Ct} = 1 - n_{Ft} = \left(1 + \exp\left(\beta \frac{\pi_{Ct} - \pi_{Ft}}{\pi_{Ft} + \pi_{Ct}}\right)\right)^{-1}, \quad (14)$$

in which π_{ht} is the performance of group h in period t .¹¹ By construction, $n_{ht} \in (0, 1)$ and $n_{Ft} + n_{Ct} = 1$. Coefficient β determines the sensitivity of agents to differences in performance between the two groups and is known as the intensity of choice parameter. A positive β implies that agents switch towards the group with the better performance. With $\beta = 0$, agents are not sensitive and remain in their group; as a result, $n_{Ft} = n_{Ct} = 0.5 \forall t$. The higher β , the quicker agents will decide to switch between groups conditional on the relative performance difference. With β sufficiently large, all agents will immediately switch from one group to the other such that n_{Ct} and n_{Ft} are either equal to zero or one. This parameter is important in determining market stability, as we will illustrate later.

Note that this switching function differs from BH98. Specifically, whereas Brock and Hommes (1997, 1998) use absolute prediction errors $\pi_{Ft} - \pi_{Ct}$, we use relative prediction errors $\frac{\pi_{Ft} - \pi_{Ct}}{\pi_{Ft} + \pi_{Ct}}$. We motivate this choice and discuss the implications in Section 2.2.

2.2. Market stability

Whereas the stability properties of the BH98 model have already been studied extensively, we briefly illustrate the stability of our model in this subsection as we have adjusted the original model somewhat to our purpose of cross-market comparison. The stability conditions of our model we find here will also help us later when judging the stability of the asset markets in the empirical part of the paper.

The switching function in the original BH98 model is different from how we defined it in the previous section. Specifically, we allow agents to switch based on relative prediction errors, $\frac{\pi_{Ft} - \pi_{Ct}}{\pi_{Ft} + \pi_{Ct}}$, whereas the original model uses absolute prediction errors, $\pi_{Ft} - \pi_{Ct}$. This was first applied in ter Ellen and Zwinkels (2010). There are three main reasons why we implement relative rather than absolute prediction errors. First, it makes the intensity of choice parameter β comparable across time and space. Because β is not unit free, its magnitude is conditional on the definition of performance as well as the (time-varying) variance of this performance. By normalising the prediction errors to a number between -1 and 1 , we will be able to compare the estimated β across all asset classes we consider. Second, the normalisation is helpful for the estimation procedure. By normalising, the characteristics of the statistical properties of the prediction errors will be more stable over time and will not have extreme values, such that estimation results are less likely to be driven by specific episodes. Third and finally, also intuitively it is safe to assume that agents judge performances relatively rather than absolutely. A 10% difference in performance between strategy A and B is more informative than a 10-unit difference. In the Appendix to this paper, we compare the estimation properties of the relative versus the absolute prediction errors.

Fig. 1 displays the relation between performance differences, both absolute and relative, and fundamentalist weights for a range of β values from zero to 100. It illustrates the effect of different levels of agent sensitivity to prediction errors in

¹¹ Note that the switching function with relative performance measures requires that $\pi_{ht} > 0$. With, for example, $\pi_{Ft} > 0$ and $\pi_{Ct} = 0$, $n_{Ft} = (1 + \exp(\beta))^{-1}$, and thereby independent of π_{Ft} . In case both $\pi_{Ft} = 0$ and $\pi_{Ct} = 0$, n_{ht} is indeterminate. Our performance measure (10), however, ensures that $\pi_{ht} \geq 0$. Empirically, we never observe that the performance measure reaches zero.

Fig. 2. Bifurcation analysis. Notes: This figure shows bifurcation plots for the absolute versus relative prediction errors. Parameter values are set to $\phi_F = 0.8$, $\phi_C = 1.1$, and $I = 1$. The x-axis denotes β and the y-axis represents the period 1000 value.

deciding about which forecasting model to use. The figures show the same general tendency that the S-shape of the logit switching function becomes more pronounced for higher values of β . On one extreme, $\beta = 0$ results in a horizontal line for both configurations, indicating no sensitivity to performance differences. On the other extreme, $\beta \rightarrow 100$ results in a step-wise function for both configurations, indicating infinite sensitivity to performance differences such that all agents are either fundamentalist or chartist. The differences lie in the non-extreme cases. The figure confirms that the absolute performance difference can take any value; as a result, the weights can also take any value between zero and unity. In the case of relative differences, however, the relative performance difference remains within the -1 to 1 range. As a result, the range over which the weights move, increases with β . For example, for $\beta = 1$, the minimum weight is 0.26 and the maximum weight is 0.73.

The stability of the asset market described by our heterogeneous agent model is determined by the coefficient set consisting of β , ϕ_F , and ϕ_C . As shown in the original (Brock and Hommes, 1997) paper, models consisting of switching fundamentalists and chartists do not necessarily converge to the fundamental equilibrium in which price is equal to the fundamental value ($x = 0$). Instead, complex dynamics can emerge. Here we illustrate the effect of an increasing β given a set of ϕ_F and ϕ_C using a bifurcation diagram. Due to the different properties of the switching function, the price dynamics generated by the model with relative performance differences will also be different. To study the differences in dynamics, we create bifurcation plots.

In a bifurcation plot one can see how the steady-state properties of the model changes when one changes certain parameters. In these simulations, we are interested in the stability of the steady state for different values of β . We set the parameter values to $\phi_F = 0.8$, $\phi_C = 1.1$, and $I = 1$. Fig. 2 presents bifurcation diagrams for the model with both absolute and relative performance differences in which we vary β from zero to ten with steps of 0.01.

Fig. 2 illustrates that both configurations of the model produce stable fundamental equilibria of $\bar{x} = 0$ up to a certain level of β , after which a bifurcation occurs and the equilibrium becomes unstable. The bifurcation points of the two configurations, though, lie at different values of β . For the absolute prediction errors, the bifurcation point lies at $\beta = 6.95$. For the relative difference, this is $\beta = 2.05$. In other words, the stability region is smaller for the relative differences than for the absolute differences.

These figures illustrate that market stability is directly related to the behaviour of individual agents in the market. In Fig. 2 we study the effect of β , but the stability properties of the model are also determined by ϕ_C and ϕ_F . When the parameter estimates of a certain asset class in the empirical section of this paper result in a stable equilibrium point, it implies that this market is stable and not very prone to bubbles. When the parameter set does not yield a stable equilibrium, however, it implies that the deterministic skeleton of the model is unstable and the market is sensitive to the endogenous creation of price bubbles. Fig. 2 suggests that the switching model exhibits a pitchfork bifurcation, where the fundamental steady state becomes unstable and two stable non-fundamental steady states are created, one above and one below the fundamental value.¹²

¹² Bolt et al. (2019) estimate a very similar two-type switching model on housing prices and show analytically that a pitchfork bifurcation occurs.

Table 1
Descriptive statistics.

	Equity		Currencies		Commodities		Macro	
	S&P500	NASDAQ	USDJPY	USDUKP	Gold	Oil	CPI	House
Mean	0.1368	0.2088	−0.0654	−0.0306	0.1550	0.1157	0.0930	0.1238
Median	0.1451	0.2333	−0.0341	−0.0179	0.0966	0.1166	0.0750	0.1344
Max	0.5189	0.9686	0.2728	0.2512	1.0004	0.8994	0.2409	0.3071
Min	−0.5023	−0.5912	−0.4933	−0.4172	−0.3420	−0.8743	0.0261	−0.1433
Std. dev.	0.1823	0.2530	0.1407	0.1193	0.2925	0.3433	0.0522	0.1016
Skew	−0.5251	−0.8105	−0.4085	−0.5267	0.8329	−0.1817	1.2070	−0.7892
Kurt	3.7341	4.4633	3.3328	3.4314	3.2899	2.9781	3.6215	3.7421
Obs.	203	151	203	203	171	163	203	202

Notes: This table shows descriptive statistics of x_t per asset class, defined as $\ln(p_t) - \ln(p_t^*)$, where $p_t^* = \sum_{i=1}^{20} p_{t-i+1}$.

3. Data and methods

Considering we want to compare the parameters of the model across asset classes, we need to employ data that is as comparable as possible in terms of frequency, sample period, and geography. After all, we want to be able to say that the differences in parameters result from differences in asset classes, rather than from differences in sample period or market composition. As such, the choices we make in the benchmark setup are based on maximum comparability for as many assets as possible. The asset classes that we consider, are equity (the S&P500 index and the Nasdaq), foreign exchange (UK Pound/US Dollar and Japanese Yen/US Dollar), commodities (WTI crude oil, gold), and two macroeconomic variables (the Case–Shiller house price index and the US Consumer Price Index). For the main results, we rely on quarterly data as this frequency allows for the inclusion of macroeconomic variables that are not available at a higher frequency. We execute robustness checks with data on monthly frequency, excluding the macro-variables.

Even though the notion of mean reversion towards some fundamental value is intuitively appealing and consistent with experimental findings, empirically implementing it is challenging because it is unknown what the fundamental value should be; see Fama (1991) on the dual hypothesis problem. An important point to realise, though, is that the fundamental value in our heterogeneous agent model does not have to be the asset's actual underlying value. Instead, it should be a proxy that boundedly rational agents might perceive as a fundamental anchor in their expectation formation process. As such, it should be a fundamental value based on a well-known model that is relatively easy to calculate using publicly available information. Furthermore, in our case we also require a fundamental value that is methodologically comparable across assets such that the results are comparable. To achieve this goal, we take two approaches. First, in the benchmark case we take a moving average of the market price as a fundamental value. The advantage of this approach is its simplicity, its applicability to all asset classes, and the fact that the approach is exactly the same across assets. Furthermore, the moving-average has a number of characteristics one would expect from a fundamental proxy. First, by construction prices mean-revert to their fundamental value. Furthermore, this approach embeds the result of Shiller (1981) that market prices exhibit excess volatility relative to their fundamental value. We take a moving average of 20 quarters, or 5 years. Our second approach to calculating a fundamental value proxy is more sophisticated and based on discounted cashflows. This results in slightly more advanced fundamental proxies, but also creates differences between the asset classes and excludes the assets that do not provide a cash-flow to the agent, such as commodities.

All data is retrieved from Thomson Reuters through Datastream, apart from the S&P500 and housing data, which are retrieved from the website of Robert Shiller.¹³ The sample period covers 1960Q1 to 2015Q2 totalling 222 observations, as far as data availability allows it. Gold, oil, and Nasdaq have slightly shorter sample periods starting in 1968, 1970, and 1978, respectively. Table 1 presents the descriptive statistics of our assets; specifically, we present the statistics of the log-price deviation from fundamental $x_t = \log(p_t) - \log(p_t^*)$ using the fundamental proxy based on the moving-average proxy. We take log-deviations such that x_t represents a percentage price deviation, which is again directly comparable across asset classes.

Table 1 shows that the mean and median values of x_t tend to be positive, illustrating the increasing trend in asset prices over time. The currencies are the exception to this rule, with a negative mean, suggesting that the U.S. dollar has depreciated vis-a-vis the U.K. pound and Japanese yen over our sample period. The minimum–maximum range and the standard deviation give an indication about the market volatility. The commodities are especially volatile, followed by equity, currencies, and the macroeconomic variables housing and CPI.

Fig. 3 displays the evolution of x_t , the deviation between the current price and its fundamental value for each asset. Mispricing is much more persistent for housing and CPI, as can be seen by the smooth movement and persistently positive level of x_t . This might be explained by the presence of chartists in the market, who drive the price further away from its fundamental by extrapolating the deviations further in the future. In contrast, deviations from the fundamental are relatively short-lived in financial(ised) markets, and typically move around zero.

¹³ See <http://www.econ.yale.edu/shiller/data.htm>.

Fig. 3. Price divergence from fundamental X . Notes: This figure presents x_t per asset class, defined as $\ln(p_t) - \ln(p_t^*)$, with $p_t^* = \sum_{i=1}^{20} p_{t-i+1}$.

Table 2
Benchmark estimation results.

	Equity		Currencies		Commodities		Macro	
	S&P500	Nasdaq	USDJPY	USDUKP	Gold	Oil	CPI	House
Static								
$\phi/2$	0.920*** (40.042)	0.888*** (32.077)	0.933*** (36.115)	0.917*** (35.498)	0.945*** (48.046)	0.838*** (17.263)	0.995*** (114.19)	0.992*** (111.25)
c	0.012 (1.811)	0.021* (1.651)	−0.003 (−0.705)	−0.002 (−0.539)	0.002 (0.297)	0.010 (0.601)	0.000 (0.267)	0.001 (0.701)
LL	224.71	99.21	277.28	305.79	174.73	38.47	678.48	551.82
Obs	190	138	190	190	158	150	190	189
Switching								
ϕ_C	1.276* (1.834)	1.142*** (13.365)	0.939*** (37.011)	1.194*** (11.744)	0.985*** (55.613)	0.990*** (6.596)	1.086*** (60.625)	1.414*** (62.994)
ϕ_F	0.552 (0.790)	0.716*** (7.087)	0.697*** (13.361)	0.566*** (4.361)	0.575*** (12.318)	0.439** (2.336)	0.889*** (28.235)	0.667*** (23.880)
β	0.142 (0.326)	1.117 (1.143)	30.949 (0.431)	1.034** (2.199)	19.074 (0.800)	2.730 (0.840)	1.300*** (6.073)	2.125*** (11.125)
c	0.011* (1.727)	0.019 (0.013)	−0.007 (−1.588)	−0.002 (−0.420)	0.003 (0.477)	0.007 (0.485)	0.000 (0.261)	−0.001 (−0.904)
LL	224.82	101.18	279.11	310.13	185.12	41.08	704.08	630.54
Obs	190	138	190	190	158	150	190	189
$P_{\phi_F=\phi_C}$	0.604	0.013	0.000	0.003	0.000	0.083	0.000	0.000
$\bar{\phi}$	0.914	0.929	0.818	0.880	0.780	0.715	0.988	1.041

Notes: This table shows the results of estimating Eq. (15) on quarterly data based on the moving-average fundamental proxy. T-statistics are in parentheses; *, **, *** represent significance at the 10%, 5%, and 1% level, respectively. $P_{\phi_F=\phi_C}$ denotes the P-value of the Wald test on equality of parameters ϕ_F and ϕ_C .

The empirical model based on Eq. (9) we take to the data is given by:

$$x_t = c + n_{Ft}\phi_F x_{t-1} + n_{Ct}\phi_C x_{t-1} + \varepsilon_t, \quad (15)$$

in which ε_t is the residual and c is an intercept we include to ensure that $E(\varepsilon_t) = 0$. Furthermore, in this benchmark model we assume that the risk-free rate is equal to zero, or $R = 1$. We test the sensitivity to this choice in Section 4.1. The weights n_{Ft} and n_{Ct} are given by Eqs. (13) and (10).

Estimation is done using (quasi) maximum likelihood as is common in the literature for these reduced-form models. We set the memory parameter l equal to 1, and test the robustness of the results to this choice in Section 4.1. The starting values for ϕ_F , ϕ_C , and β are 0.8, 1.1, and 1, respectively. The intensity of choice parameter β is restricted to positive values in the estimation procedure.

4. Results

Table 2 presents the estimation results of the benchmark model, with quarterly data, a moving-average fundamental based on $M = 20$ quarters, and a memory length $l = 1$ quarter.

The top panel labelled ‘Static’ presents the results with $\beta = 0$, such that $n_{Ft} = n_{Ct} = 0.5\forall t$. As a result, the size of the estimated coefficients is $\phi/2$. First of all, we observe that the persistence in x_t is rather different across asset classes, from very persistent for the macro variables, i.e., 0.995 for CPI, to more mean reverting for the more financialised assets, such as 0.838 for oil. The mean reversion is indicative for the efficiency of the asset, as it illustrates how quickly the market price reverts to its fundamental value. Given that we calculate x_t as the log-difference between price and fundamental, an average ϕ of 0.935 implies that the market expects prices to mean-revert with 6.5% per period ($1 - 0.935 = 0.065$).

The bottom panel of Table 2 labelled ‘Switching’ presents the results of the full switching model in which the switching parameter β is estimated as a free parameter. For all but one asset, the S&P500, we observe that allowing agents to switch between strategies adds to the explanatory power of the model. The Wald-test of equality of parameters, $\phi_F = \phi_C$ is only accepted for the S&P500 at the 10% significance level.¹⁴ In other words, for seven out of eight asset classes we find evidence for behavioural heterogeneity. The degree of heterogeneity is economically relevant; the average difference between ϕ_F and ϕ_C is 0.463, ranging from 0.747 for the house price index to 0.198 for CPI. This implies that fundamentalists, on average, expect mean-reversion to occur 46% per period faster than chartists. For five out of eight assets, we observe that the chartist coefficient ϕ_C is larger than unity; this is significant in four cases. This implies that chartists in these markets expect $|x_t|$ to increase, so the market price to move away from the fundamental value. This has implications for market stability because

¹⁴ Because ϕ_F and ϕ_C are not identified under the null hypothesis of no switching, a standard likelihood ratio is not informative regarding the added value of switching. Teräsvirta (1994) shows, however, that a significant difference between the two auto-regressive parameters is a sufficient condition.

the market price will move away from the fundamental in periods of chartist domination. As such, it might be the case that prices do not converge to the fundamental equilibrium $x_t = 0$ with such a parameter set. $\bar{\phi}$ represents the average or equilibrium value of ϕ for the switching model. With $\bar{\phi} > 1$, the model is locally unstable and price does not converge to its fundamental value in equilibrium. The results indicate that the housing market is indeed locally unstable. The other macro-variable, CPI, shows the second highest $\bar{\phi}$, but in this case it is slightly smaller than unity and therefore locally stable.

The estimated switching parameter β is positive, implying that agents switch towards the group with the smaller forecast error in the previous period. In other words, the switching function functions as a positive-feedback rule. The average β equals 9.535, ranging from 1.034 for the U.S. dollar–U.K. pound currency pair to 30.949 for the U.S. dollar–Japanese yen currency pair.¹⁵ Overall, it appears that β is somewhat lower for the macro-variables, and higher for the highly liquid financial markets. In other words, agents are more sensitive to performance difference in financial markets than in more macroeconomic assets, and therefore switch more between strategies in financial assets.

Fig. 4 shows the relative performance from the fundamentalist perspective, $\frac{\pi_{Ft} - \pi_{Ct}}{\pi_{Ft} + \pi_{Ct}}$, for all assets. Fig. 4 indicates that the relative forecast errors hover around zero for all assets. The amplitude is bounded between -1 and $+1$; values that are also attained for most assets. The volatility in the relative errors is somewhat lower for the housing market and, to a lesser extent, CPI. This might explain the substantially higher T -statistics for the estimated β 's for housing and CPI in Table 2.

Fig. 5 presents the estimated fundamentalist weights n_{Ft} for all assets, and shows that the weights move around the average of 0.50. This is by construction of the switching function given by Eq. (13). The variability and amplitude of the weights is related to the estimated β coefficients. For example, the weights for the dollar–yen exchange rate jump between zero and one very frequently ($\beta = 30.9$), whereas the weights for the S&P500 are static at 0.5 ($\beta = 0$). The fluctuations in the fundamentalist weights can be explained with the movement in the price deviation x_t as shown in Fig. 3. For example, the fundamentalist weight for gold is zero between 2003 and 2012 with the exception of a single spike in 2005. This corresponds with the run-up in gold prices as illustrated by a prolonged increase in x_t in Fig. 3. Furthermore, we can recognise the crash in the US housing market in 2006 by the corresponding spike in the fundamentalist weight. This switches to a spike in chartist weight in 2009 as the price undershoots the fundamental value.

Fig. 6 presents the estimated market sentiment, defined by $n_{Ft}\phi_F + n_{Ct}\phi_C$, for all assets. When this parameter exceeds one, the market expects an increase of $|x_t|$. When this is the case, price is not converging towards the fundamental equilibrium point of $x_t = 0$ at that point in time, but rather the market is unstable and exhibits a temporary bubble.

Fig. 6 shows that the market sentiment follows the inverse pattern of the fundamentalist weights, which is by construction. This implies that when the weight on chartists is high, the aggregate market impact is more likely to be destabilising. For a number of assets, we observe that sentiment is above one in certain periods (Nasdaq, gold, CPI, and housing). In such episodes, agent behaviour is destabilising, driving the price further away from its fundamental value and the market exhibits a temporary bubble. Especially the housing markets shows periods of strong instability with peaks to 1.3, implying that market participants expect mispricing to grow by 30% over the next quarter.

4.1. Alternative specifications

The choices we made for the benchmark configuration presented in the previous section were based on maximising the number of included assets and maximising the comparability across assets. The question is whether these choices affect the empirical findings. Therefore, in this section we estimate the model under different configurations to see whether the main results about agent behaviour and corresponding market stability continue to hold. In each test we run, the number of assets or the cross-market comparability will be affected, but it will allow us to draw inference on the sensitivity of the results.

4.1.1. Monthly data

First, we estimate the model using monthly data rather than quarterly data. This has the consequence that the housing market drops out because monthly data is not available. It might help, though, in finding more reliable coefficient estimates as can be seen in the Appendix that the estimation is sensitive to small-sample issues. We use the same moving-average length in calculating the fundamental value (5 years) and memory length in the switching function ($I = 1$ period). Table 3 presents the results.

The estimation results using monthly data in Table 3 are highly comparable to those using quarterly data in Table 2. Although the exact coefficient estimates are somewhat different, the rankings in terms of persistence and switching are equivalent. The evidence on heterogeneity is somewhat less, as we no longer find heterogeneity for the Nasdaq and the oil market. The stability properties of the estimated models are comparable, apart from the housing market. Specifically, all assets converge to the fundamental equilibrium.

¹⁵ Standard t -tests do not apply for judging the significance of β due to the nonlinear structure of the switching function. Because ϕ_F and ϕ_C are not identified under the null hypothesis of no switching, β is a nuisance parameter. The significance can again be judged by the test of equality of ϕ_F and ϕ_C ; see Teräsvirta (1994).

Fig. 4. Relative performance. Notes: This figure presents the relative performance $\frac{\pi_{F1} - \pi_{CI}}{\pi_{F1} + \pi_{CI}}$ from the fundamentalist perspective.

Fig. 5. Fundamental weights. Notes: This figure presents estimated fundamental weights n_{Ft} .

Fig. 6. Market sentiment. Notes: This figure presents the estimated market sentiment, defined as $n_F \phi_F + n_C \phi_C$.

Table 3
Estimation results monthly.

	Equity		Currencies		Commodities		Macro
	S&P500	Nasdaq	USDJPY	USDUKP	Gold	Oil	CPI
Static							
$\phi/2$	0.982*** (150.4)	0.971*** (113.5)	0.983*** (134.9)	0.975*** (129.1)	0.979*** (146.6)	0.963*** (96.26)	0.999*** (514.2)
c	0.003 (1.562)	0.007* (1.772)	−0.001 (−0.519)	−0.001 (−0.684)	0.001 (0.371)	0.002 (0.475)	0.000 (0.411)
LL	1133.0	589.2	1244.0	1287.2	725.4	450.0	2642.6
Obs	594	437	594	594	498	474	594
Switching							
ϕ_C	1.010*** (35.53)	1.252*** (4.413)	0.996*** (122.0)	1.009*** (33.66)	1.015*** (92.14)	0.989*** (8.122)	1.039*** (304.9)
ϕ_F	0.946*** (31.66)	0.680*** (2.361)	0.949*** (79.23)	0.911*** (29.17)	0.961*** (105.0)	0.939*** (7.599)	0.792*** (43.26)
β	1.276 (0.654)	0.163 (1.346)	777.9 (0.030)	3.093 (0.729)	9.009 (0.889)	1.787 (0.101)	2.269*** (32.44)
c	0.003* (1.735)	0.007* (1.738)	−0.001 (−0.835)	−0.001 (−0.606)	0.000 (−0.066)	0.002 (0.458)	0.001 (4.084)
LL	1133.9	589.86	1247.9	1292.6	729.5	450.1	2742.3
Obs	594	437	594	594	498	474	594
$P_{\phi_F=\phi_C}$	0.261	0.317	0.001	0.095	0.001	0.838	0.000
$\hat{\phi}$	0.978	0.966	0.973	0.960	0.988	0.964	0.916

Notes: This table shows the results of estimating Eq. (15) on monthly data using the moving-average fundamental proxy. T -statistics are in parentheses; *, **, *** represent significance at the 10%, 5%, and 1% level, respectively. $P_{\phi_F=\phi_C}$ denotes the P -value of the Wald test on equality of parameters ϕ_F and ϕ_C .

4.1.2. Exogenous fundamental

In the benchmark case we used a moving average of the price level as a proxy for the fundamental value. By doing so, we implicitly assumed that market prices mean-revert to their fundamental value in the long run, which is reasonable. To test the sensitivity of the results to this choice, though, we re-estimate the model using an exogenous fundamental value for the asset classes for which this is possible. We do this for the asset classes for which there is an economic model available, that can also reasonably be assumed to be used by the (boundedly rational) agents in our model.¹⁶ We obtain fundamentals for the assets under consideration in the following way:

- **Equity:** Dividend discounted model, as in Boswijk et al. (2007), based on $p_t^* = Y_t(1+g)/(r-g)$, in which Y_t is dividend, g is the growth rate of dividends, and r the discount rate.
- **Foreign exchange:** Purchasing Power Parity (PPP) model, as in ter Ellen et al. (2013). We take the PPP exchange rates as calculated by Datastream.
- **Housing:** Dividend discount model, using rents as dividends, as in Kouwenberg and Zwinkels (2014) and similar to Bolt et al. (2019).

Fig. 7 plots the deviation of the market prices from their corresponding fundamental values for the variables for which we have an exogenous fundamental value. Even though it is impossible to say that the fundamental values are correct, the figures illustrate that the chosen fundamental proxies behave as one expects from a fundamental value. Specifically, we observe that the market price oscillates around the fundamental proxies, or that x_t oscillates around one in all five cases¹⁷, and that the fundamental follows a more stable path than the market prices (Shiller, 1981). Note that this also indicates that the moving-average in the benchmark configuration represents a reasonable fundamental proxy.

Table 4 presents the estimation results using the exogenous fundamental values. The estimation results with exogenous fundamental are again highly comparable to the benchmark results, especially for the USDJPY currency pair and the housing market. The intensity of choice β is consistently around 30 for the USDJPY and around 2 for the housing market. The estimates are slightly different for the equity markets and the USDUKP. The S&P500 now shows significant heterogeneity, whereas the Nasdaq does not. The USDUKP results now indicate somewhat more intense switching than before.

¹⁶ The commodities and CPI drop out because they do not yield a cashflow making it challenging to obtain an exogenous fundamental value.

¹⁷ Johansen cointegration tests indicate that the price and fundamental series are cointegrated for all assets. Results available on request.

Fig. 7. Exogenous fundamentals. Notes: This figure presents the fundamental value proxies based on exogenous models and data combined with the actual market values.

4.1.3. The effect of memory

In the benchmark configuration we assumed that agents only consider the most recent period in comparing the relative performance of groups when deciding on which forecasting rule to use in the next period, i.e., $l = 1$. In this subsection we study the effect of this choice by increasing l to 2, 4, and 8, i.e., six months, one year, and two years. Table 5 presents the results.

Table 5 presents a clear trend regarding the effect of increasing the memory parameter. A higher l is consistently accompanied by a decrease in the added value of switching. These results suggest that agents in the markets we consider do not consider more than one quarter of past performance when deciding to switch between groups. This is to be expected given the relatively low frequency of our data.

Table 4
Estimation results with exogenous fundamental.

	Equity		Currencies		Macro
	S&P500	Nasdaq	USDJPY	USDUKP	House
Static					
$\phi/2$	0.980*** (71.50)	0.964*** (63.23)	0.961*** (44.12)	0.958*** (43.45)	0.994*** (174.8)
c	0.006 (1.033)	0.028 (1.472)	−0.004 (−0.737)	0.007 (1.388)	0.001 (0.970)
LL	235.1	91.42	305.5	319.3	616.0
Obs	207	157	209	197	207
Switching					
ϕ_C	0.984*** (71.16)	1.128*** (6.670)	0.978*** (41.46)	0.959*** (50.73)	1.586*** (14.08)
ϕ_F	0.628*** (10.95)	0.798*** (4.774)	0.920*** (25.37)	0.660*** (32.31)	0.247* (1.797)
β	52954 (0.000)	0.267 (0.322)	33.68 (0.204)	42.26 (0.370)	1.802*** (23.04)
c	0.003 (0.582)	0.025 (1.204)	−0.006 (−0.973)	0.010** (2.399)	0.000 (0.181)
LL	236.6	91.90	306.8	323.5	748.0
Obs	207	157	209	197	207
$P_{\phi_C=\phi_F}$	0.000	0.321	0.114	0.000	0.000
ϕ	0.806	0.963	0.949	0.810	0.917

Notes: This table shows the results of estimating Eq. (15) on quarterly data with fundamental values based on exogenous data. T-statistics are in parentheses; *, **, *** represent significance at the 10%, 5%, and 1% level, respectively. $P_{\phi_F=\phi_C}$ denotes the P-value of the Wald test on equality of parameters ϕ_F and ϕ_C .

Table 5
The effect of memory.

	Equity		Currencies		Commodities		Macro	
	S&P500	Nasdaq	USDJPY	USDUKP	Gold	Oil	CPI	House
$I = 2$								
ϕ_C	0.942*** (25.72)	1.066*** (5.412)	0.932*** (37.01)	1.078*** (9.952)	1.026*** (31.49)	0.904*** (16.70)	1.075*** (64.49)	1.344*** (39.74)
ϕ_F	0.900*** (15.03)	0.627** (2.274)	0.935*** (13.36)	0.665*** (4.649)	0.643*** (9.461)	0.439** (2.581)	0.907*** (27.03)	0.735*** (24.32)
β	36.68 (0.146)	1.277 (0.662)	0.000 (0.431)	1.931 (0.924)	7.670 (1.329)	51.35 (0.330)	1.671*** (4.739)	2.166*** (9.236)
$P_{\phi_F=\phi_C}$	0.612	0.347	NA	0.078	0.000	0.012	0.000	0.000
$I = 4$								
ϕ_C	1.147** (2.256)	0.909*** (15.48)	1.847** (2.279)	1.267*** (3.499)	2.043** (2.484)	1.156 (1.629)	1.135*** (13.15)	1.275*** (29.98)
ϕ_F	0.652 (1.268)	0.833*** (7.186)	−0.088 (−0.111)	0.505 (1.370)	−0.404 (−0.508)	0.439 (0.588)	0.772*** (7.653)	0.815*** (24.12)
β	0.547 (0.461)	24.70 (0.214)	−0.575 (−1.472)	0.658 (1.131)	−1.061*** (−4.863)	−0.707 (−0.336)	−1.394*** (−5.225)	−2.400*** (−6.768)
$P_{\phi_F=\phi_C}$	0.628	0.598	0.226	0.293	0.129	0.620	0.040	0.000
$I = 8$								
ϕ_C	1.172 (1.238)	0.925*** (2.815)	2.064 (1.386)	1.173 (0.331)	1.611 (1.195)	1.133 (1.231)	1.288*** (3.270)	1.176*** (24.82)
ϕ_F	0.642 (0.673)	0.833*** (2.570)	−0.306 (−0.212)	0.656 (0.186)	0.156 (0.123)	0.439 (0.459)	0.625 (1.597)	0.905*** (25.97)
β	−0.371 (−0.278)	9.216 (0.063)	0.601 (1.387)	−0.129 (−0.067)	−0.772* (−1.815)	1.224 (0.273)	1.096*** (4.257)	3.090*** (3.421)
$P_{\phi_F=\phi_C}$	0.780	0.887	0.419	0.942	0.578	0.711	0.398	0.001

Notes: This table shows the results of estimating Eq. (15) on quarterly data using the moving-average fundamental proxy while varying the memory parameter I . T-statistics are in parentheses; *, **, *** represent significance at the 10%, 5%, and 1% level, respectively. $P_{\phi_F=\phi_C}$ denotes the P-value of the Wald test on equality of parameters ϕ_F and ϕ_C .

Table 6

The effect of the risk-free rate.

	Equity		Currencies		Commodities		Macro	
	S&P500	Nasdaq	USDJPY	USDUKP	Gold	Oil	CPI	House
Static								
$\phi/2$	0.930*** (39.99)	0.897*** (32.09)	0.945*** (36.13)	0.931*** (35.79)	0.970*** (47.46)	0.850*** (17.39)	1.025*** (112.87)	1.007*** (112.0)
c	0.012* (1.887)	0.023* (1.725)	−0.004 (−0.726)	−0.002 (−0.522)	0.002 (0.230)	0.010 (0.607)	−0.001 (−0.799)	0.001 (0.902)
LL	222.52	97.977	275.41	303.06	172.85	36.61	669.25	552.61
Obs	190	138	190	190	158	150	190	189
Switching								
ϕ_C	1.072*** (6.622)	1.144*** (13.54)	1.401*** (5.854)	1.197*** (12.51)	1.159*** (12.65)	0.996*** (6.746)	1.161*** (27.72)	1.421*** (57.70)
ϕ_F	0.787*** (22.98)	0.723*** (7.184)	0.412 (1.611)	0.588*** (4.567)	0.701*** (4.856)	0.439** (2.435)	0.906*** (31.49)	0.654*** (21.39)
β	0.000 (0.207)	1.148 (1.134)	0.489 (1.231)	1.140*** (2.267)	1.046 (1.616)	2.857 (0.841)	1.224*** (6.490)	1.913*** (10.80)
c	0.012*** (3.853)	0.021 (1.630)	−0.005 (−0.978)	−0.002 (−0.431)	0.003 (0.419)	0.008 (0.525)	0.000 (−0.308)	0.001 (1.413)
LL	222.52	100.07	277.37	307.82	176.10	39.13	692.14	627.30
Obs	190	138	190	190	158	150	190	189
$P_{\phi_F=\phi_C}$	NA	0.013	0.043	0.003	0.046	0.069	0.000	0.000
ϕ	0.930	0.934	0.906	0.892	0.930	0.717	1.033	1.037

Notes: This table shows the results of estimating Eq. (15) on quarterly data using the moving-average fundamental proxy with the memory parameter $l = 1$. T-statistics are in parentheses; *, **, *** represent significance at the 10%, 5%, and 1% level, respectively. $P_{\phi_F=\phi_C}$ denotes the P-value of the Wald test on equality of parameters ϕ_F and ϕ_C .

4.1.4. Taking the risk-free rate into account

In the benchmark specification, we put the risk-free rate equal to zero such that $R = 1 + r_f = 1$ in Eq. (9). In this robustness test, we do include a risk-free rate. Specifically, we take the 1-month T-bill rate from the website of Ken French. Table 6 presents the estimation results.

The estimation results in Table 6 indicate that including the risk-free rate does not change the result qualitatively. In terms of heterogeneity, the difference between the fundamentalist and chartist coefficients ϕ_F and ϕ_C , the results are highly similar. The same holds for the added value of switching: we find significant switching for all asset classes except the S&P500. The order of magnitude of the intensity of choice parameter β is somewhat different for two assets, USDJPY and gold, compared to the benchmark case. In the benchmark case, these two assets displayed very high values of β . In the analysis including the risk-free rate, the significance of these two assets is marginal, so the differences in the estimated β 's might be due to identification issues.

4.1.5. Heterogeneous beliefs about returns

The model in Section 2 is written in terms of deviations from the fundamental price, x_t . The question is whether this choice is important for the estimation results. Therefore, we also estimate the model in terms of returns (as in, for example, ter Ellen and Zwinkels, 2010).¹⁸

To illustrate the difference between the models estimated in deviations from the fundamental and in returns, we also estimate BH98 on returns. Considering this model has slightly different behavioural rules and properties we will now briefly introduce the BH98 model in returns.

The price change of the asset is a weighted average of the expectations of fundamentalists and chartists:

$$\Delta p_t = n_{Ft} E_{F,t-1}(\Delta p_t) + n_{Ct} E_{C,t-1}(\Delta p_t). \quad (16)$$

Fundamentalists expect the price to revert back to its fundamental value:

$$E_{F,t-1}(\Delta p_t) = \phi_F(p_{t-1}^* - p_{t-1}). \quad (17)$$

Chartists extrapolate past returns:

$$E_{C,t-1}(\Delta p_t) = \phi_C(p_{t-1} - p_{t-2}). \quad (18)$$

¹⁸ The underlying reason of this difference lies in the implicit assumption about market clearing. Specifically, the model in deviations assumes that markets clear based on a Walrasian auctioneer. The model in returns is disequilibrium model in which a market maker adjust prices consistent with excess demand. See Hommes (2006) for an in-depth discussion about the micro-structure in agent-based models.

Table 7

Estimation results for model in returns.

	Equity		Currencies		Commodities		Macro	
	S&P500	Nasdaq	USDJPY	USDUKP	Gold	Oil	CPI	House
Static								
ϕ_F	−0.053 (−0.941)	−0.075 (−1.012)	−0.054 (−1.008)	−0.099* (−1.809)	0.056 (1.296)	−0.225** (−2.080)	0.116*** (5.562)	−0.028*** (−4.433)
ϕ_C	0.299* (1.962)	−0.035 (−0.203)	0.204 (1.255)	0.378** (2.611)	0.258* (1.916)	0.151 (1.105)	0.689*** (6.569)	2.026*** (38.93)
c	0.017** (2.618)	0.031** (2.217)	−0.007 (−1.446)	−0.003 (−0.858)	0.008 (1.002)	0.019 (1.065)	0.001 (1.035)	0.002*** (4.216)
LL	217.0	91.43	273.4	304.5	168.2	30.5	686.7	742.0
Obs	190	138	190	190	158	150	190	189
Switching								
ϕ_F	−0.053 (−0.930)	−0.071** (−2.285)	−0.094* (−1.844)	−0.099 (−1.611)	0.049 (1.573)	−0.665*** (−5.581)	0.032** (2.191)	−0.014* (−1.866)
ϕ_C	0.299** (1.961)	0.046 (0.394)	0.168 (0.939)	0.378** (2.609)	0.271** (2.635)	−0.056 (−0.626)	0.751*** (13.41)	1.321*** (14.46)
β	0.000 (0.000)	92.19 (0.089)	7.049 (0.430)	0.000 (0.000)	18.28 (0.431)	76.29 (0.393)	7.807 (1.286)	1.743*** (8.810)
c	0.017** (2.616)	0.035** (2.484)	−0.007 (−1.629)	−0.003 (−0.842)	0.009 (1.272)	0.013 (0.796)	0.003*** (3.744)	0.001*** (3.102)
LL	217.0	92.22	274.0	304.5	169.8	37.3	694.6	751.1
Obs	190	138	190	190	158	150	190	189
LR	0.00	1.58	1.30	0.00	3.09	13.66	15.83	18.07

Notes: This table shows the results of estimating the model in terms of returns Δp_t rather than deviations from the fundamental value $p_t - p_t^*$ using quarterly data. T-statistics are in parentheses; *, **, *** represent significance at the 10%, 5%, and 1% level, respectively.

The weights of the two groups are determined in the same way as in Eq. (13), and the performance depends on the forecasting abilities of each strategy:

$$\pi_{ht} = - \sum_{i=1}^I |\Delta p_{t-i} - E_{h,t-i-1} \Delta p_{t-i}|. \quad (19)$$

Table 7 presents the estimation results for the model in returns, again using quarterly data, $I = 1$, and a fundamental proxy based on a moving-average of the price level. The results show a clear presence of both fundamentalists and chartists: For most assets, we find a negative ϕ_F suggesting the presence of mean-reverting fundamentalists, and a positive ϕ_C suggesting the presence of trend-chasing chartists. CPI forms the exception, as fundamentalists expect the price level to deviate further from the fundamental. This finding could be driven by the strong upward trend in CPI combined with the moving average fundamental that is by construction lagging behind the current price level. The switching results indicate that the dynamic model adds explanatory power to the model for the commodities and the macro-variables, but not for the equity markets and the currencies. Within these assets, the ranking of speed of switching is comparable to the other configurations above. Striking is again the consistency for the housing market: again we find a β roughly equal to 2.

4.2. Discussion

Estimating (various specifications of) a generic heterogeneous agent model for various asset classes creates scope to compare the results across these asset classes, as well as to previous empirical studies. Comparing across asset classes gives us the opportunity to draw some conclusions on market characteristics and behavioural heterogeneity. At the same time, comparing the results to previous evidence also emphasises the importance of choosing the model specifications.

Given that limits to arbitrage (Shleifer and Vishny, 1997) are largest in macroeconomic assets such as housing markets, it can be expected that these markets are most prone to behavioural bubbles. Shleifer and Vishny (1997) argue that there are three limits to arbitrage that might prevent markets from being fully efficient: implementation costs, fundamental risk, and noise traders. All three might play a role in explaining our results. First, implementation costs are relatively high for real estate, as these markets are not very liquid compared to financial markets. As such, it is costly to do transactions. Second, fundamental risks are high, because houses are highly heterogeneous. Combined with the high transaction costs, it makes it hard to hedge the fundamental risks embedded in real estate transactions. Finally, noise trader risk might also be relatively important in real estate markets due to the wide range of different market participants. The market consists of professional investors as well as retail investors who also have a consumption motive for being active in the real estate market. Furthermore, most retail investors only make a few transactions in the real estate market in a lifetime, such that the

probability of sub-rational behaviour is increased.¹⁹ Moreover, [Hong and Stein \(2003\)](#) show that in markets where there are short-sale constraints, speculative bubbles may arise because bullish investors do not enter the market, or leave the market earlier. Consequently, in unconstrained markets, both bullish and bearish investors are represented, whereas in markets with short-sale constraints, there is an overrepresentation of investors with optimistic price expectations. This makes such markets more prone to bubbles. A prime example of a market without such constraints is the foreign exchange market, while the housing market is a market for which short-sale constraints are very binding.

This is confirmed by the empirical results: the heterogeneity between fundamentalists and chartists is relatively large for the macroeconomic variables house prices and inflation. Specifically, for these variables we find that the chartists are destabilising, meaning that they expect the price deviation from the fundamental value to increase over time. Combined with relatively slow switching, it is conceivable that the chartists dominate the market for a relatively long period, causing large and persistent mispricing. As such, we find that the stability of markets is least for the macroeconomic variables of inflation and house prices. At the same time, we find that the intensity of choice parameter is relatively high for financial markets, implying that agents switch between expectation functions relatively quickly. Hence, it might be that volatility is high in financial markets, also compared to the fundamental value, but due to the rapid switching between groups there is a lower probability for bubbles to develop. The main differences between companies in the S&P500 and the Nasdaq is that the companies in the Nasdaq are more prone to limits to arbitrage as they are more difficult to value.²⁰ This is what we see reflected in our results: Overall, we find slightly more heterogeneity and switching for the Nasdaq than for the S&P500. Our results are robust for the exact choice of functional form, definition of fundamental value, data frequency, and model configuration.

It is quite puzzling that, no matter the exact estimation specification, there does not seem to be significant heterogeneity and switching in the S&P500 equity market. This is in contrast to earlier studies, both those using survey data ([Greenwood and Shleifer, 2014](#)) and using price data (e.g. [Boswijk et al., 2007](#); [Amilon, 2008](#); [Chiarella et al., 2014](#); [Lof, 2014](#); [Hommes and in 't Veld, 2017](#); [Schmitt, 2020](#)). Although some of these differences may be explained by differences in estimation sample, frequency, or model specification, even the study that is closest to one of our specifications ([Hommes and in 't Veld, 2017](#), quarterly data for S&P500, 1950–2016, with a Gordon growth-based fundamental) is able to detect behavioural heterogeneity. The main difference may be that their model includes a memory parameter that takes into account a longer history of performance. Indeed, we do see some evidence in [Table 5](#) that increasing our memory parameter to $l = 4$ results in a $\phi_C > 1$ and $a\phi_F < 1$.²¹ The results of [Frijns et al. \(2018\)](#), who examine heterogeneous beliefs among retail and institutional investors, may shed some more light on the overall lack of switching for equity in our results. First of all, they show that there are indeed large differences in the behaviour of retail versus institutional investors, as alluded to above. However, contrary to what one may expect, retail investors in their study behave more like fundamentalists on average, whereas their institutional investors are more likely to use chartist-type strategies. Although they find strong switching between strategies within each group, the behaviour of the two groups often seems to cancel out on aggregate, which is why we may not pick up significant switching in our aggregate price series.

As all the empirical evidence on behavioural heterogeneity in the foreign exchange market is based on data with higher frequencies and exogenous fundamentals, it is not that straightforward to compare our estimates with earlier findings. Likewise, estimation for commodity prices has often been conducted on monthly data. However, for these markets it is very common to use a moving-average fundamental proxy. In contrast to our findings, [Baur and Glover \(2014\)](#) find that for a model specification based on returns, monthly data, and a performance-based switching function, chartists in the gold market are contrarian, whereas fundamentalists are destabilising. Comparing estimation results for the oil market with those of [ter Ellen and Zwinkels \(2010\)](#), for a model specification very similar to [Baur and Glover \(2014\)](#), we find more similarities. Whereas we do not find significant chartism in our estimation results for a model in returns, this may be explained by the fact that we do not take asymmetry into account, as [ter Ellen and Zwinkels \(2010\)](#) do.

Finally, for our macroeconomic series, housing and inflation, the studies by [Bolt et al. \(2019\)](#) and [Cornea-Madeira et al. \(2019\)](#) come closest. Although our results are qualitatively in line with results in [Bolt et al. \(2019\)](#) for the U.S. housing market (presence of both chartists and fundamentalists and evidence of switching), our parameter estimates are more extreme. In our specification using an exogenous fundamental, chartists are very extrapolative, expecting the price deviation from the fundamental to increase with 59% next quarter. Fundamentalists, on the other hand, are much more stabilising than in [Bolt et al. \(2019\)](#), expecting misalignment to decrease by about 75% over the next quarter. This is also in contrast to [Kouwenberg and Zwinkels \(2014\)](#), who find that fundamentalists' price reversion is slow with 1.4% per quarter on average. To the best of our knowledge, [Cornea-Madeira et al. \(2019\)](#) is the only other paper that estimates a heterogeneous agent model on inflation data. Their model specification is based on an exogenous fundamental (based on marginal costs), a memory parameter of four quarters, and a naive chartist rule comparable to the one in our model on returns. As such, their model specification is so different from ours that it is infeasible to compare parameter estimates. A feature that is present in their results as well as ours, is that the weights of the two groups is relatively persistent, i.e. agents do not switch easily between groups.

¹⁹ See also the work of [Salzman and Zwinkels \(2017\)](#), who show that many behavioural biases are stronger in real estate than in other markets.

²⁰ They are younger, have higher market-to-book values, and pay out less dividends ([Baker and Wurgler, 2006](#)).

²¹ However, note that in [Table 5](#) we estimate a model with the moving-average fundamental proxy rather than the Gordon growth-based fundamental.

There remain important differences in the ways the models in these papers are specified. This emphasises the importance of keeping the model specification constant when comparing behavioural heterogeneity over various asset classes. When model specifications vary, it is already difficult to directly compare parameter estimates for the same asset class, let alone across asset classes. This paper makes an important contribution in that respect.

5. Conclusion

We have estimated a generic heterogeneous agent model on various asset classes, ranging from macroeconomic variables such as CPI and house prices, to financial markets such as the foreign exchange market. We find that whereas switching is more intense in financial assets, there is more heterogeneity in the macro-assets. As a result, the latter are more instable in the sense that behavioural bubbles have a higher probability of occurring. This might be explained by the fact that the limits to arbitrage are relatively large in these markets. Our findings are qualitatively robust to the choice of fundamental, data frequency, and model configuration.

Declaration of Competing Interest

There are no financial or personal relationships with other people or organisations that could inappropriately influence (bias) this work.

Appendix

One of the reasons why we use the normalised performance measure is ease of estimation. To illustrate this point, we run a simulation and estimation exercise. Specifically, we run stochastic simulations of the model, adding a noise term $\epsilon_t \sim (0; 0.1)$, with $\phi_F = 0.8$ and $\phi_C = 1.1$. We set $I = 1$ and vary β in the simulation from zero to two with steps of 0.04. Subsequently, we estimate the model on the simulated data using maximum likelihood. We do this 1,000 times for each level of β , with 220²² or 1000 observations per run to study the effect of small samples. The distribution of estimated values of $\hat{\beta}$ will allow us to draw inference about the accuracy and efficiency of the estimation procedure. Fig. 8 presents the average as well as the standard error of the 1000 estimated $\hat{\beta}$'s for each level of β for $t = 220$.

Panel (a) of Fig. 8 displays the results for the model with absolute prediction errors. The upper left figure shows the average estimated $\hat{\beta}$. Whereas the actual β ranges from zero to two, the estimated $\hat{\beta}$ hovers around 1000 with a high level of variation. The rough coefficient estimate, therefore, is not informative about the underlying switching mechanism. We therefore winsorise the coefficient estimates at the 10% level.²³ The resulting figure, in the lower left panel, shows an upward sloping line from a little below zero to two, reflecting the underlying β values.

Panel (b) of Fig. 8 displays results for the model with relative prediction errors. The upper left figure shows a noisy, but clearly upward sloping line. Whereas the actual values range from zero to two, the estimates appear to range from zero to approximately 100. The standard error of the estimated $\hat{\beta}$ s is substantially lower (approx. 20 times) than for the model with absolute prediction errors. After winsorising the coefficient estimates at the 10% level, we observe a very smooth line exactly reflecting the actual underlying β 's in the lower left panel. The standard errors are also very low, and appear to decrease as the underlying β increases. Overall, it is clear that the estimation produces much more accurate and efficient estimates for the model with relative prediction errors.

To formalise the findings in Fig. 8, we estimate the following equation $\hat{\beta}_i = c_1 + c_2\beta_i + \epsilon_i$. An unbiased $\hat{\beta}_i$ should give $c_1 = 0$ and $c_2 = 1$. Table 8 presents the results.

The upper half of Panel (a) in Table 8 shows results for the absolute prediction errors. Consistent with Fig. 8, we observe that the estimated $\hat{\beta}$ is not informative regarding the actual β . After winsorisation, we do observe a positive and significant c_2 , although it is significantly larger than unity. The winsorised standard error decreases somewhat as β increases. The lower half of Panel (a) in Table 8 shows the results for the relative prediction errors. The raw $\hat{\beta}$ estimates show a positive correlation with β , although the c_2 coefficient is far too high. After winsorisation, the estimates are very well behaved, with $c_1 = 0.151$ and $c_2 = 0.825$.

Panel (b) in Table 8 presents the results for the simulation of $t = 1000$ periods. Overall, the results are very similar to those with $t = 220$ indicating that the differences between the models with absolute and relative prediction errors are not driven by small sample issue. We do observe that the non-winsorised results are somewhat stronger for the longer simulation paths. This implies that the longer simulation paths eliminate some of the extreme values in the β estimates.

²² This is roughly equal to the number of observations we have in the empirical part of the paper.

²³ Specifically, we remove the 10% highest values of $|\hat{\beta}|$.

(a) Absolute prediction errors

(b) Relative prediction errors

Fig. 8. Estimation. *Notes:* This figure shows estimated β coefficients for the absolute versus relative prediction errors based on $t = 220$. The y-axis of the left column denotes the estimated β value, whereas the x-axis denotes the 'true' β of the underlying data generating process. The y-axis of the right column denotes the estimated standard errors for different values of β .

Table 8
Estimation bias.

	Coeff.	St. err.	Win. coeff	Win. st. err.
Panel (a): $t = 220$				
Absolute prediction errors				
c_1	682.285*** (4.778)	20.101*** (8.067)	−0.743*** (−7.082)	0.009*** (−27.816)
c_2	230.504 (1.492)	2.621 (1.054)	1.350*** (18.433)	−0.001*** (−2.072)
R^2	0.01	−0.01	0.89	0.06
Relative prediction errors				
c_1	−41.915*** (−3.313)	0.792*** (2.872)	0.151*** (5.318)	0.001*** (14.520)
c_2	87.813*** (7.102)	0.544*** (2.217)	0.825*** (37.912)	−0.000*** (−5.352)
R^2	0.48	0.07	0.99	0.73
Panel (b): $t = 1000$				
Absolute prediction errors				
c_1	−18.451 (−0.181)	5.807* (1.746)	−0.020* (−1.692)	0.002*** (47.62)
c_2	399.84*** (3.025)	6.913* (1.878)	1.009*** (54.72)	0.000** (−2.601)
R^2	0.11	0.03	0.98	0.11
Relative prediction errors				
c_1	−0.019 (−0.117)	0.006 (1.389)	0.097*** (8.897)	0.000*** (15.21)
c_2	1.318*** (5.864)	0.003 (0.503)	0.906*** (108.23)	0.000*** (−3.557)
R^2	0.50	−0.02	1.00	0.54

Notes: This table shows the results of estimating $\hat{\beta}_t = c_1 + c_2\beta_t + \epsilon_t$ using $\hat{\beta}$ estimates from simulated data. An unbiased $\hat{\beta}_t$ should give $c_1 = 0$ and $c_2 = 1$. *, **, *** represent significance at the 10%, 5%, and 1% level, respectively.

References

- Amilon, H., 2008. Estimation of an adaptive stock market model with heterogeneous agents. *J. Empir. Financ.* 15 (2), 342–362.
- Assenza, T., Bao, T., Hommes, C., Massaro, D., 2014. Experiments on expectations in macroeconomics and finance. In: Duffy, J. (Ed.), *Experiments in Macroeconomics (Research in Experimental Economics: A Research Annual)*, 17. Emerald Insight, pp. 11–70. doi:10.1108/S0193-230620140000017002.
- Baker, M., Wurgler, J., 2006. Investor sentiment and the cross-section of stock returns. *J. Financ.* 61 (4), 1645–1680.
- Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. *J. Financ. Econ.* 49 (3), 307–343.
- Baur, D.G., Glover, K.J., 2014. Heterogeneous expectations in the gold market: specification and estimation. *J. Econ. Dyn. Control* 40 (1), 116–133. doi:10.1016/j.jedc.2014.01.001.
- Bloomfield, R., Hales, J., et al., 2002. Predicting the next step of a random walk: experimental evidence of regime-shifting beliefs. *J. Financ. Econ.* 65 (3), 397–414. doi:10.1016/S0304-405X(02)00147-2.
- Bolt, W., Demertzis, M., Diks, C., Hommes, C., Van Der Leij, M., 2019. Identifying booms and busts in house prices under heterogeneous expectations. *J. Econ. Dyn. Control* 103, 234–259.
- Boswijk, H.P., Hommes, C.H., Manzan, S., 2007. Behavioral heterogeneity in stock prices. *J. Econ. Dyn. Control* 31 (6), 1938–1970. doi:10.1016/j.jedc.2007.01.001.
- Branch, W.A., 2004. The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations. *Econ. J.* 114 (497), 592–621. doi:10.1111/j.1468-0297.2004.00233.x.
- Brock, W.A., Hommes, C.H., 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. *J. Econ. Dyn. Control* 22 (8–9), 1235–1274. doi:10.1016/S0165-1889(98)00011-6.
- Chen, H., Joslin, S., Tran, N.K., 2012. Rare disasters and risk sharing with heterogeneous beliefs. *Rev. Financ. Stud.* 25 (7), 2189–2224. doi:10.1093/rfs/hhs064.
- Chiarella, C., ter Ellen, S., He, X.-Z., Wu, E., et al., 2015. Fear or fundamentals? Heterogeneous beliefs in the European sovereign CDS market. *J. Empir. Financ.* 32 (1), 19–34. doi:10.1016/j.jempfin.2014.11.003.
- Chiarella, C., He, X.-Z., Zwinkels, R.C.J., 2014. Heterogeneous expectations in asset pricing: empirical evidence from the S&P500. *J. Econ. Behav. Organ.* 105 (1), 1–16. doi:10.1016/j.jebo.2014.03.003. <http://linkinghub.elsevier.com/retrieve/pii/S0167268114000729>
- Cornea-Madeira, A., Hommes, C.H., Massaro, D., 2019. Behavioral heterogeneity in U.S. inflation dynamics. *J. Bus. Econ. Stat.* 37 (2), 288–300.
- Cutler, D., Poterba, J., Summers, L., 1990. Speculative dynamics and the role of feedback traders. *Am. Econ. Rev.* 80 (2), 63–68. <http://EconPapers.repec.org/RePEc:aea:aecrev:v:80:y:1990:i:2:p:63-68>
- De Grauwe, P., Grimaldi, M., 2006. Exchange rate puzzles: a tale of switching attractors. *Eur. Econ. Rev.* 50, 1–33. doi:10.1016/j.euroecorev.2004.07.005.
- De Jong, E., Verschoor, W.F.C., Zwinkels, R.C.J., 2010. Heterogeneity of agents and exchange rate dynamics: evidence from the EMS. *J. Int. Money Financ.* 29 (8), 1652–1669. doi:10.1016/j.jimonfin.2010.05.007.
- DeLong, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. *J. Polit. Econ.* 98 (4), 703–738. Reprinted in R. H. Thaler (Ed.), *Advances in Behavioral Finance*, Russell Sage Foundation, 1993.
- Dieci, R., Westerhoff, F., 2016. Heterogeneous expectations, boom-bust housing cycles, and supply conditions: a nonlinear economic dynamics approach. *J. Econ. Dyn. Control* 71, 21–44.

- Dijk, D.V., Teräsvirta, T., Franses, P.H., 2002. Smooth transition autoregressive models a survey of recent developments. *Econom. Rev.* 21 (911221489), 1–47.
- Fama, E.F., 1991. Efficient capital markets: II. *J. Financ.* 46 (5), 1575–1617. doi:10.1111/j.1540-6261.1991.tb04636.x.
- Fernandez-Perez, A., Frijns, B., Tourani-Rad, A., Weisskopf, J.-P., 2019. Behavioural heterogeneity in wine investments. *Appl. Econ.* 51 (30), 3236–3255.
- Frankel, J.A., Froot, K.A., 1990. Chartists, fundamentalists, and trading in the foreign exchange market. *Am. Econ. Rev.* 80 (2), 181–185.
- Frijns, B., Huynh, T., & Tourani-Rad, A., Westerholm, J. (2018). Heterogeneous Beliefs Among Retail and Institutional Investors. Auckland University of Technology. Working paper.
- Frijns, B., Lehnert, T., Zwinkels, R.C.J., 2010. Behavioral heterogeneity in the option market. *J. Econ. Dyn. Control* 34 (11), 2273–2287.
- Greenwood, R., Shleifer, A., 2014. Expectations of returns and expected returns. *Rev. Financ. Stud.* 27 (3), 714–746. doi:10.1093/rfs/hht082.
- Hommes, C.H., 2006. Chapter 23 – Heterogeneous agents models in economics and finance. In: *Handbook of Computational Economics*, Vol. 2. Elsevier, pp. 1109–1186. doi:10.1016/S1574-0021(05)02023-X.
- Hommes, C.H., in 't Veld, D., 2017. Booms, busts and behavioural heterogeneity in stock prices. *J. Econ. Dyn. Control* 80 (1), 1533–1597. <https://ideas.repec.org/p/tin/wpaper/20150088.html>
- Hong, H., Stein, J.C., 2003. Differences of opinion, short-sales constraints, and market crashes. *Rev. Financ. Stud.* 16 (2), 487–525.
- Kouwenberg, R., Zwinkels, R., 2014. Forecasting the US housing market. *Int. J. Forecast.* 30 (3), 415–425. doi:10.1016/j.ijforecast.2013.12.010.
- Leary, M.T., Michaely, R., 2011. Determinants of dividend smoothing: empirical evidence. *Rev. Financ. Stud.* 24 (10), 3197–3249.
- Lof, M., 2014. Rational speculators, contrarians and excess volatility. *Manag. Sci.* 61 (8), 1889–1901. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2194792
- Lux, T., 1995. Herd behaviour, bubbles and crashes. *Econ. J.* 105 (431), 881–896.
- Lux, T., Zwinkels, R.C.J., 2018. Empirical validation of agent-based models. In: LeBaron, B., Hommes, C.H. (Eds.), *Handbook of Computational Economics*, 4. Elsevier, pp. 437–488.
- MacDonald, R., 2000. Expectations formation and risk in three financial markets: surveying what the surveys say. *J. Econ. Surv.* 14 (1), 69–100. doi:10.1111/1467-6419.00105.
- Salzman, D., Zwinkels, R.C.J., 2017. Behavioural real estate. *J. Real Estate Lit.* 25 (1), 77–106.
- Schmitt, N., 2020. Heterogeneous expectations and asset price dynamics. *Macroecon. Dyn.* forthcoming
- Schmitt, N., Westerhoff, F., 2019. Short-run momentum, long-run mean reversion and excess volatility: an elementary housing model. *Econ. Lett.* 176, 43–46.
- Shiller, R.J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *Am. Econ. Rev.* 71 (3), 421–435.
- Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. *J. Financ.* 52 (1), 35–55.
- Spronk, R., Verschoor, W.F.C., Zwinkels, R.C.J., 2013. Carry trade and foreign exchange rate puzzles. *Eur. Econ. Rev.* 60 (1), 17–31.
- ter Ellen, S., Verschoor, W.F.C., 2018. Heterogeneous beliefs and asset price dynamics: a survey of recent evidence 24, 53–79.
- ter Ellen, S., Verschoor, W.F.C., Zwinkels, R.C.J., 2013. Dynamic expectation formation in the foreign exchange market. *J. Int. Money Financ.* 37 (1), 75–97. doi:10.1016/j.jimonfin.2013.06.001.
- ter Ellen, S., Zwinkels, R.C.J., 2010. Oil price dynamics: a behavioral finance approach with heterogeneous agents. *Energy Econ.* 32 (6), 1427–1434. doi:10.1016/j.eneco.2010.03.003.
- Teräsvirta, T., 1994. Specification, estimation, and evaluation of smooth transition autoregressive models. *J. Am. Stat. Assoc.* 89 (25), 208–218.
- Westerhoff, F., Franke, R., 2012. Converse trading strategies, intrinsic noise and the stylized facts of financial markets. *Quant. Financ.* 12, 425–436.
- Westerhoff, F., Reitz, S., 2005. Commodity price dynamics and the nonlinear market impact of technical traders: empirical evidence for the US corn market. *Phys. A: Stat. Mech. Appl.* 349 (3–4), 641–648. doi:10.1016/j.physa.2004.11.015.