WORKING PAPER

Comparing behavioural heterogeneity across asset classes

NORGES BANK RESEARCH

12 | 2017

SASKIA TER ELLEN
CARS H. HOMMES
AND
REMCO C.J. ZWINKELS

NORGES BANK

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:
FacilityServices@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail
FacilityServices@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-819-0 (online)
ISBN 978-82-7553-990-6 (online)

Comparing Behavioural Heterogeneity Across Asset Classes*

Saskia ter Ellen† Cars H. Hommes‡ Remco C.J. Zwinkels§

June 2017

Abstract

We estimate a generic agent-based model in which agents have heterogeneous beliefs about the future price to see to what extent behaviour differs across assets, and what this implies for market stability. We find evidence for behavioural heterogeneity for all asset classes, except for equities. Heterogeneity is especially pronounced for macro-economic variables. Agents update their beliefs frequently in financial markets, and only gradually in the case of macro-economic variables. Consequently, we find that the probability of behavioural bubbles is substantially higher for the macro-economic variables than for financial assets.

Keywords: financial markets, heterogeneous expectations, market stability

JEL classification codes: E31, G12, G15.

*This Working 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 of Nonlinear Dynamics and Econometrics (SNDE), Computational Economics and Finance (CEF), and the International Symposium on Computational Economics and Finance, as well as seminar participants at Norges Bank.

†Norges Bank; E: saskia.ter-ellen@norges-bank.no.

‡University of Amsterdam and Tinbergen Institute; E: c.h.hommes@uva.nl

§Vrije Universiteit Amsterdam and Tinbergen Institute; E: r.zwinkels@vu.nl

1 Introduction

Market prices show dynamics that cannot always be explained by models based on the assumption of representative agents, such as volatility clustering, heavy tails, and bubbles and crashes. One way to better understand these dynamics is to see them as emergent properties of a system composed of heterogeneous agents. In this paper, we use a model in which agents have heterogeneous expectations about future prices, and estimate the model on a broad set of assets to get a better understanding of differences in agent behavior and corresponding market stability across asset classes. Specifically, as in Brock and Hommes (1997, 1998) the model has two types of agents, fundamentalists and chartists, who are able to switch between types conditional on relative performance. We find that switching between types is more prevalent in financial markets such as foreign exchange markets. Heterogeneity, however, is especially pronounced in macro-economic price series such as house prices and CPI, causing these markets to be especially prone to behavioural bubbles.

Despite its theoretical elegance and convenience, the rational expectations hypothesis (REH) by Muth (1961) does not appear to hold up empirically at the individual level. Greenwood and Shleifer (2014) conclude based on six different surveys that stated beliefs are not consistent with rational expectation representative agent type models. Although Branch (2004) finds that there is heterogeneity in inflation surveys as well, he concludes that it can be rational to stick to a predictor function that is not the most accurate when the alternatives are costly. Experimental studies show similar results. As early as 1976, Schmalensee used experimental methods to study the belief formation process and finds evidence against rationality. More recently, Anufriev and Hommes (2012) show that evolutionary selection among heterogeneous expectation rules explains experimental outcomes and aggregate price behaviour much better than homogeneous benchmarks. Bloomfield and Hales (2002) run an experiment in which they explicitly tell participants that they are forecasting a random variable; yet, forecasts display all sorts of structures.

The question that arises is how expectations are formed if they are not rational; stepping away from rationality creates an infinite number of degrees of freedom. Furthermore, stepping away from rationality also creates scope for heterogeneity between agents. A number of studies show that expectation formation can be summarized 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 switch between a trend-following and a mean-reverting rule, conditional on the recent price realizations. MacDonald (2000) summarizes the main findings from studying quantitative surveys, and finds that most survey participants use forecasting rules with short-term trend extrapolation and long-term mean reversion. Greenwood and Shleifer (2014) also find evidence for trend extrapolation.

The subsequent question is whether the boundedly rational and heterogeneous expectations at the individual level affect asset prices due to the self-referential nature of asset markets, or whether heterogeneity averages out in the limit and markets converge to efficiency. A number of theoretical studies have looked into this issue. Cutler et al. (1990) propose a model with fundamental traders alongside feedback traders who form expectations based on past returns, and show that the latter group of traders does affect prices. Delong et al. (1993) theoretically show that a group of noise traders of sufficient size in an otherwise rational market will affect prices and survive in the long run. Likewise, Lux (1995) models a group of traders that are not fully informed about the fundamentals, 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 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 of functions, rational and naive, the authors show that under certain conditions highly irregular equilibrium prices converge to a strange attractor. In a follow-up paper, Brock and Hommes (1998) apply the same concept, but with two boundedly rational expectation functions in an asset market.

The papers by Brock and Hommes (1997, 1998) have triggered a stream of research attempting to replicate the 'stylized facts' of financial markets, including the aforementioned volatility clustering and heavy tails using the heterogeneity approach; see e.g. Hommes (2006) for an overview. The empirical literature on such heterogeneous agent models has by now established ample evidence on the importance of behavioural heterogeneity for explaining financial market dynamics, on both the micro and the macro level; for an overview of the empirical literature on heterogeneous agent models, see Chen et al. (2012) and Lux and Zwinkels (2017). Heterogeneous agent models have been extensively tested using all sorts of data. Various versions of the model, with varying numbers and types of agents, different profit and/or switching functions, and varying results, have been estimated successfully on a large number of asset classes. 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; Spronk et al., 2013) have been extensively analyzed, but the model has also showed itself useful in explaining the price dynamics in, for instance, housing markets (Kouwenberg and Zwinkels, 2014; Bolt et al., 2014), option markets (Frijns et al., 2010), commodity markets (ter Ellen and Zwinkels, 2010; Baur and Glover, 2014; Westerhoff and Reitz, 2005), and credit markets (Chiarella et al., 2015). Even macro-economic variables such as inflation can be described by a heterogeneous agent model, as in Cornea-Madeira et al. (2017)[1]. So far, the literature has concluded that allowing for heterogeneity and switching is important when describing the dynamics of individual markets. Due to the variation in the estimated models, however, the results have never been directly compared with each other.

In this paper we estimate a generic heterogeneous agent model with fundamentalists, chartists, and switching on a broad set of assets in order to see whether behaviour differs across markets, and what that means for the stability of these markets. The limited comparability between the existing papers is due to the different functional forms in the set of expectation formation rules, but mainly due to the fact that the main coefficient in the function governing the switching between rules is unit free, such that its magnitude cannot be compared across space and time. Therefore, in this paper we adopt a generalized version of the switching function that allows us to compare results across assets. Subsequently, 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. The results can be compared based on three main parameters. The intensity of choice parameter describes how fast agents update their beliefs when they get new information. The other coefficients of interest are the fundamentalist mean reversion and the chartist extrapolation coefficients. Combined, the three coefficients determine the stability of the system, and therefore the sensitivity of the model to behavioural bubbles.

[1]Although inflation is not the result of supply and demand in some sort of exchange, it is typically modeled as a self-referential process in which the realization is a function of the ex-ante expectations.

Our results suggest that the intensity of choice parameter is relatively high for financial markets, implying that agents switch between expectation functions relatively quickly. At the same time, we find that the heterogeneity between fundamentalists and chartists is relatively pronounced for the macro-economic variables house prices and inflation. Both contribute to the instability of markets, but the mechanism is different. We find that the stability of markets is least for the macro-economic variables of inflation and house prices. Our results are robust for the exact choice of model, definition of fundamental value, data frequency, and model configuration.

With the analysis conducted in this paper, we aim to get a better understanding of which types of markets are more prone to bubbles, what the market characteristics are that drive that bubble sensitivity, and how investor behaviour contributes to that. Ultimately, we want to contribute to a measure of 'bubble sensitivity' with this framework, and evaluate the risk of certain policies that might drive prices away from their fundamental values. In such a way, the paper contributes to the academic literature on heterogeneous agent models, while at the same time offering a different perspective on certain policy-relevant issues. In this respect, our finding that the macro-economic variables are especially prone to bubbles is interesting, because the focus of policy makers tends to be on the (excessively) volatile financial markets.

The remainder of this paper is organized as follows. Section 2 sets out 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. The results are discussed in Section 4. Finally, we show the results of several robustness checks in Section 5, and Section 6 concludes.

2 Heterogeneous Agent Model and Market Stability

In order to evaluate and compare the characteristics of investor behaviour across several asset classes, we estimate a generic stylized model with heterogeneous agents that is appropriate for each asset class. This heterogeneous agent model will be able to provide estimated coefficients that have a straightforward interpretation and can be easily compared across these different asset classes. We use a slightly modified version of the model of Brock and Hommes (1998, hereinafter BH98) as the basic model on which we build our empirical analysis. Roughly speaking, there are two different versions of this model. The first version, which is closest to the BH98 model, models the price in deviations from its fundamental value. Agents have heterogeneous expectations about the convergence of the price to the fundamental value, and thus about the size of this deviation. In this setup, one group of agents has mean-reverting beliefs, meaning they expect the price of the asset to revert back to its fundamental value (fundamentalists). The other group of agents has extrapolative beliefs, which means that they expect the deviation to increase in the future (chartists). The second version models changes in the price (returns), and agents form beliefs about future returns. Again, fundamentalists in this model expect that the price of an asset reverts to its fundamental value. However, chartists are considered to be trend chasers. They base their expectations of future prices solely on past price movements. In both versions of the BH98 model, agents can switch between a chartist and a fundamentalist strategy, often based on the (relative) profitability of these strategies in the past.

2.1 The Model

We will now introduce our model which is based on Brock and Hommes (1997, 1998). Assume an economy with a single risky asset with price p_t that pays a stochastic dividend y_t[2]. Wealth then evolves according to

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

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

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

$$Max_z \left\{ E_{ht}W_{t+1} - (a/2)V_{ht}(W_{t+1}) \right\} \tag{2}$$

in which a is the risk aversion parameter and V_{ht} the variance of wealth. Solving yields an optimal demand for the risky asset z equal to

[2]For assets that do not pay out dividends such as, for example, commodities, $y_t = 0 \forall t$

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

Assume that there are H groups of investors and that n_{ht} is the fraction of type h investors in period t. Total demand for the risky asset is then given by

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

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

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

and

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

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

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

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

the pricing Equation (6) can then be written as

$$Rx_t = \sum n_{ht} f_h(x_{t-1}, ..., x_{t-L}) \tag{8}$$

Consistent with the literature on heterogeneous agents, we assume two types of traders, fundamentalists and chartists. The fundamentalists expect the price level to converge to the fundamental value, and thus x to converge to zero. Hence, $f_F = \phi_F x_{t-1}$ with $\phi_F < 1$. Chartists, on the other hand, are destabilizing and expect the deviation between price and fundamental value to increase. Hence, $f_C = \phi_C x_{t-1}$ with $\phi_C > 1$[3].

[3] Note that when estimating the model, it is possible to find heterogeneity, but with both parameters less than one. In that case, both groups of agents have mean-reverting beliefs, but one group expects

The full pricing equation is then given by

$$Rx_t = n_{Ft}\phi_F x_{t-1} + n_{Ct}\phi_C x_{t-1} \tag{9}$$

Agents are able to switch between groups conditional on the relative performance of the groups. The smaller the forecast error of one group is compared to the other, the more likely it is that more agents will switch to this rule. In other words, n_{ht} is endogenously determined by means of

$$n_{ht} = \exp\left(\beta\frac{\pi_{ht}}{\pi_{Ft}+\pi_{Ct}}\right)/Z_t \tag{10}$$

$$Z_t = \sum \exp\left(\beta\frac{\pi_{ht}}{\pi_{Ft}+\pi_{Ct}}\right) \tag{11}$$

which simplifies to

$$n_{Ft} = \left(1+\exp\left(\beta\frac{\pi_{Ft}-\pi_{Ct}}{\pi_{Ft}+\pi_{Ct}}\right)\right)^{-1} \tag{12}$$

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

in which π_{ht} is the performance of group h. Coefficient β determines the sensitivity of agents to differences in performance between the two groups. With $\beta = 0$, agents are not sensitive and remain in their group; as result, $n_{Ft} = n_{Ct} = 0.5\forall t$. The higher β is, the quicker agents will decide to switch between groups conditional on the relative performance difference. This parameter is important in determining market stability, as we will show later.

Note that this switching function differs from the one in the Brock and Hommes reference model. Specifically, whereas Brock and Hommes (1997) use absolute profit differences $\pi_{Ft} - \pi_{Ct}$, we use relative profit differences $\frac{\pi_{Ft}-\pi_{Ct}}{\pi_{Ft}+\pi_{Ct}}$. We discuss why and what the implications are in Section 2.2.

To complete the model, we need to 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

the reversion to be faster than the other group.

previous I periods. Specifically,

$$\pi_{ht} = \sum_{i=1}^{I} |x_{t-i} - \phi_h x_{t-i-2}| \tag{14}$$

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 5.

2.2 Market Stability

Whereas the stability properties of the BH98 model have already been studied extensively, we briefly show 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 profit differences, $\frac{\pi_{Ft}-\pi_{Ct}}{\pi_{Ft}+\pi_{Ct}}$, whereas the original model uses absolute profit differences, $\pi_{Ft}-\pi_{Ct}$[4]. This was first applied in ter Ellen and Zwinkels (2010). There are two main reasons why we implement relative rather than absolute profit differences. First, it makes the intensity of choice parameter β comparable across time and markets. Because β is unit-free, its magnitude cannot be interpreted, as it is conditional on the definition of performance as well as the (time-varying) variance of this performance. By normalizing the profit difference to a number between -1 and 1, we will be able to compare the estimated β across all asset classes we consider. Second, the normalization is helpful for the estimation procedure. By normalizing, the characteristics of the statistical properties of the profit difference 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. In the Appendix to this paper, we compare the estimation properties of the relative versus the absolute profit differences.

Figure 1 shows the relation between performance differences, both absolute and relative, and fundamentalist weights for a range of β values from zero to 100. It illustrates

[4]This transformation functions correctly only when performance $\pi_{ht} >> 0$. This is indeed the case for our definition of performance.

Figure 1: Absolute versus Relative Profit Differences

(a) Absolute profit difference (b) Relative profit difference

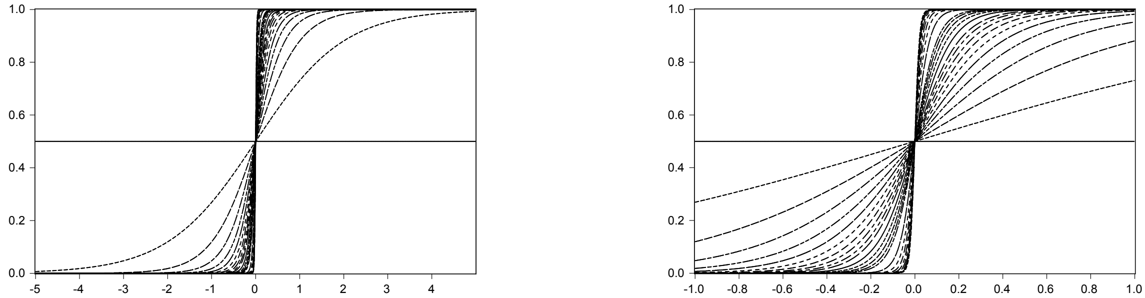

Notes: This figure shows how smooth the transition between two types of agents is for different values of β and a range of profit function values. In both graphs, the x-axis denotes the value of the profit function and the y-axis the proportion of agents of type 1. The lines represent different values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

the effect of different levels of an agent's sensitivity to profit differences in deciding on 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. At one extreme, $\beta = 0$ results in a horizontal line for both configurations, indicating no sensitivity to performance differences. At the other extreme, $\beta \to \infty$ results in a step-wise function for both configurations, indicating infinite sensitivity to performance differences such that all agents are either fundamentalists or chartists. 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 already 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. Second, we illustrate the effect of varying ϕ_F and ϕ_C given a certain level of β

Figure 2: Bifurcation analysis

(a) Absolute profit difference (b) Relative profit difference

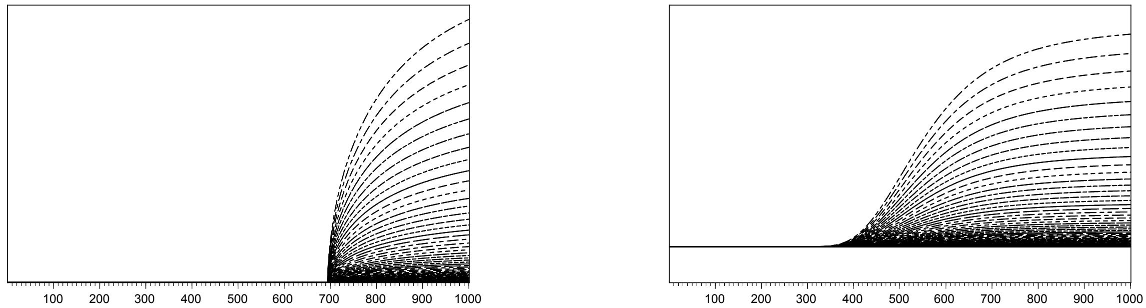

Notes: This figure shows bifurcation plots for the absolute versus relative profit difference. Parameter values are set to $\phi_F = 0.8$, $\phi_C = 1.1$, and $I = 1$. The x-axis denotes $\beta * 100$ and the lines represent various equilibria.

by means of basins of attraction. 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 and we plot the parameter basins of attraction for both specifications.

In a bifurcation plot one can see how the dynamics 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 $I = 1$. Figure 2 presents bifurcation diagrams for the model with both absolute and relative performance differences in which we vary β between zero and ten with steps of 0.01.

Figure 2 illustrates that both configurations of the model produce stable fundamental equilibria of $x = 0$ 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 profit difference, 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.

To further compare the stability properties of the two configurations, we construct basins of attraction for different combinations of ϕ_F and ϕ_C given a certain level of β. We leave $I = 1$ and set $\beta = 2$ and $\beta = 5$. Figure 3 presents the results.

Figure 3 shows the stability of the model for different values of ϕ_F and ϕ_C. When the combination of certain values of values of ϕ_F and ϕ_C lies in the red area, the model con-

Figure 3: Basin of attraction

(a) Absolute profit difference with $\beta = 2$ (b) Absolute profit difference with $\beta = 5$

(c) Relative profit difference with $\beta = 2$ (d) Relative profit difference with $\beta = 5$

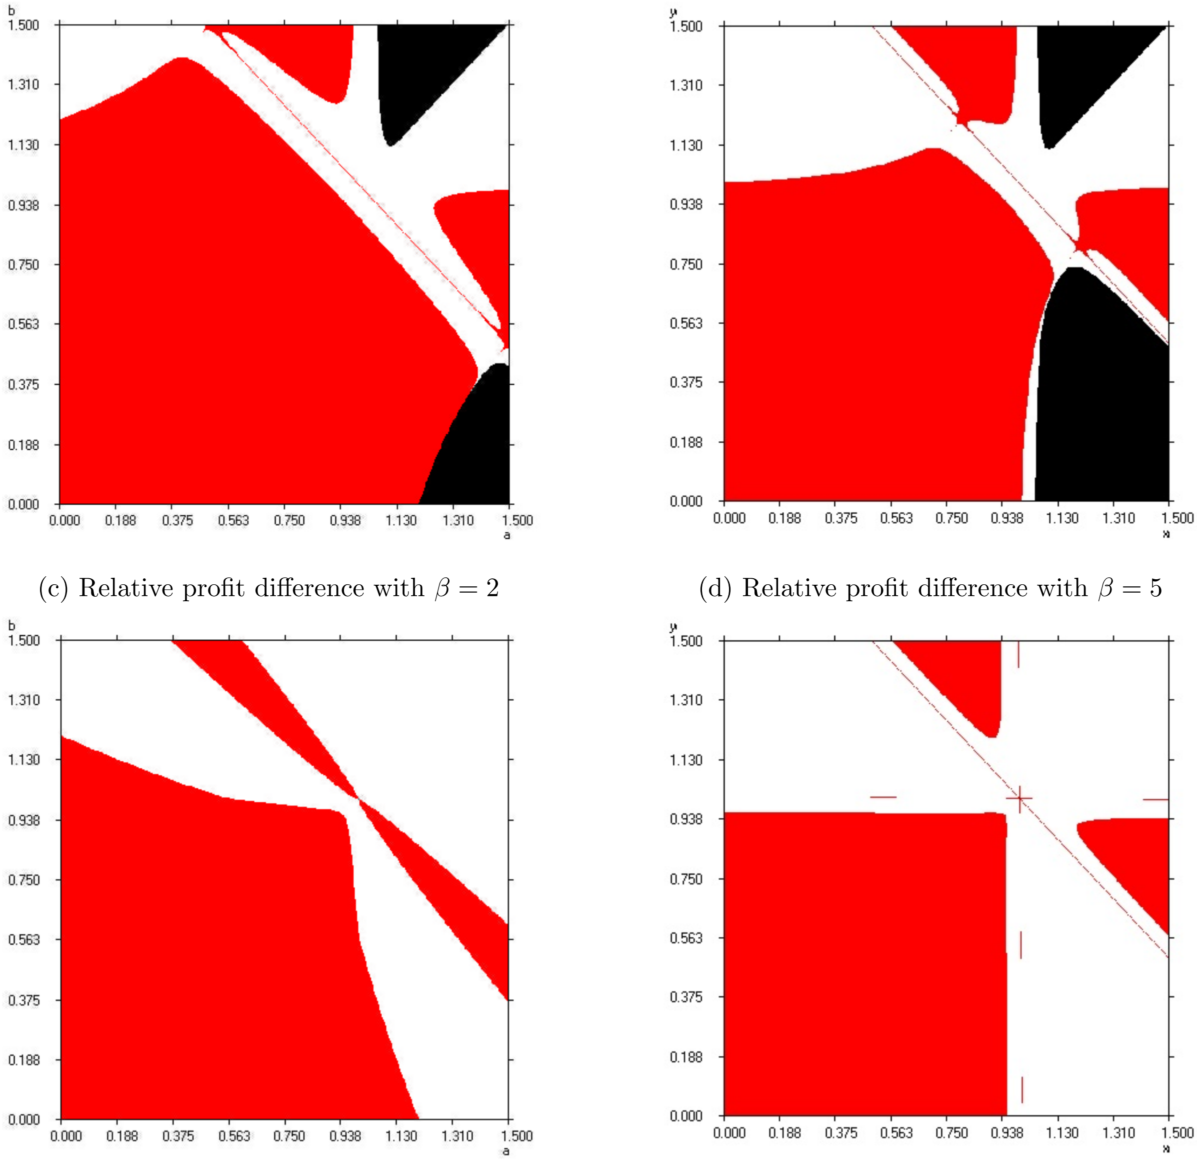

Notes: This figure shows basins of attraction for the absolute versus relative profit difference and for different combinations of ϕ_F, ϕ_C on the x- and y-axis. Furthermore, $I = 1$ and we compare $\beta = 2$ and $\beta = 5$. Red areas represent convergence, black areas represent divergence, and white ares represent non-convergence.

verges to the fundamental equilibrium $x = 0$. The white area indicates non-convergence, and black indicates divergence (i.e., an explosive price-path). Whereas the convergence region appears to be larger for the configuration with absolute profit differences, the configuration with relative profit difference does not generate divergence regions. When looking at the case of $\beta = 5$, we observe that the stability region for absolute profit differences shrinks somewhat, but that the regions for relative differences change shape substantially.

These figures show that market stability is directly related to the behaviour of individual agents in the market. When the parameter estimates of a certain asset class in the empirical section of this paper lie in the convergence area, it implies that this market is stable and not very prone to bubbles. When the parameter set does not belong to the convergence area, however, it implies that the market is sensitive to the endogenous creation of price bubbles.

3 Data and Methods

Considering that we want to compare the behavioural 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 (culture). 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 equities (S&P500 index), foreign exchange (UK pound / US dollar and Japanese yen / US dollar), commodities (WTI crude oil, gold), and two macro-economic 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 macro-economic 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, 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 realize, 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 fundamental value proxy has a number of characteristics one would expect. First, by construction this proxy assures that 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 is slightly more sophisticated and based on (exogenous) models and data. 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 are obtained from Thomson Reuters through Datastream, apart from the equity and housing data, which are obtained from the website of Robert Shiller[5]. The sample period covers 1960Q1 to 2015Q2 totalling 222 observations, as far as data availability allows it. Gold and oil have slightly shorter sample periods starting in 1968 and 1970, respectively. Table 1 presents the descriptive statistics of our assets; specifically, we present the statistics of the log-price deviation from fundamental value $x_t = log(p_t) - log(p_t^*)$ using the fundamental value 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. Currencies are the exception, with a negative mean, suggesting that the US dollar has depreciated vis-a-vis the UK pound and Japanese yen. The minimum-maximum range and the standard deviation give an

[5]See http://www.econ.yale.edu/˜shiller/data.htm.

Table 1: Descriptive Statistics

	Equity	Currencies		Commodities		Macro	
	S&P500	USDJPY	USDUKP	Gold	Oil	CPI	House
Mean	0.1368	-0.0654	-0.0306	0.1550	0.1157	0.0930	0.1238
Median	0.1451	-0.0341	-0.0179	0.0966	0.1166	0.0750	0.1344
Max	0.5189	0.2728	0.2512	1.0004	0.8994	0.2409	0.3071
Min	-0.5023	-0.4933	-0.4172	-0.3420	-0.8743	0.0261	-0.1433
Std.Dev.	0.1823	0.1407	0.1193	0.2925	0.3433	0.0522	0.1016
Skew	-0.5251	-0.4085	-0.5267	0.8329	-0.1817	1.2070	-0.7892
Kurt	3.7341	3.3328	3.4314	3.2899	2.9781	3.6215	3.7421
Obs.	203	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 $p_t^* = \sum_{i=1}^{20} p_{t-i+1}$.

indication about market volatility. The commodities are especially volatile, followed by equity, currencies, and the macro-economic variables.

Figure 4 shows 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 the CPI, as can be seen by the smooth movement and persistently positive level of x_t. This can be explained by the presence of chartists in the market, who drive the price further away from its fundamental value by extrapolating the deviations further in the future. In contrast, deviations from the fundamental value are relatively short-lived in financial(ized) markets, and often move around zero.

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

$$x_t = c + n_{Ft}\phi_F x_{t-1} + n_{Ct}\phi_C x_{t-1} + \varepsilon_t \tag{15}$$

in which ε_t is the residual and c is an intercept we include to ensure that $E(\varepsilon_t) = 0$. The weights n_{Ft} and n_{Ct} are given by Equations (12) and (14).

Estimation is done using (quasi) maximum likelihood as is common in the literature for these reduced-form models. We set the memory parameter I equal to 1, and test the robustness of the results to this choice in Section 5. 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.

Figure 4: Price Divergence from Fundamental Value X

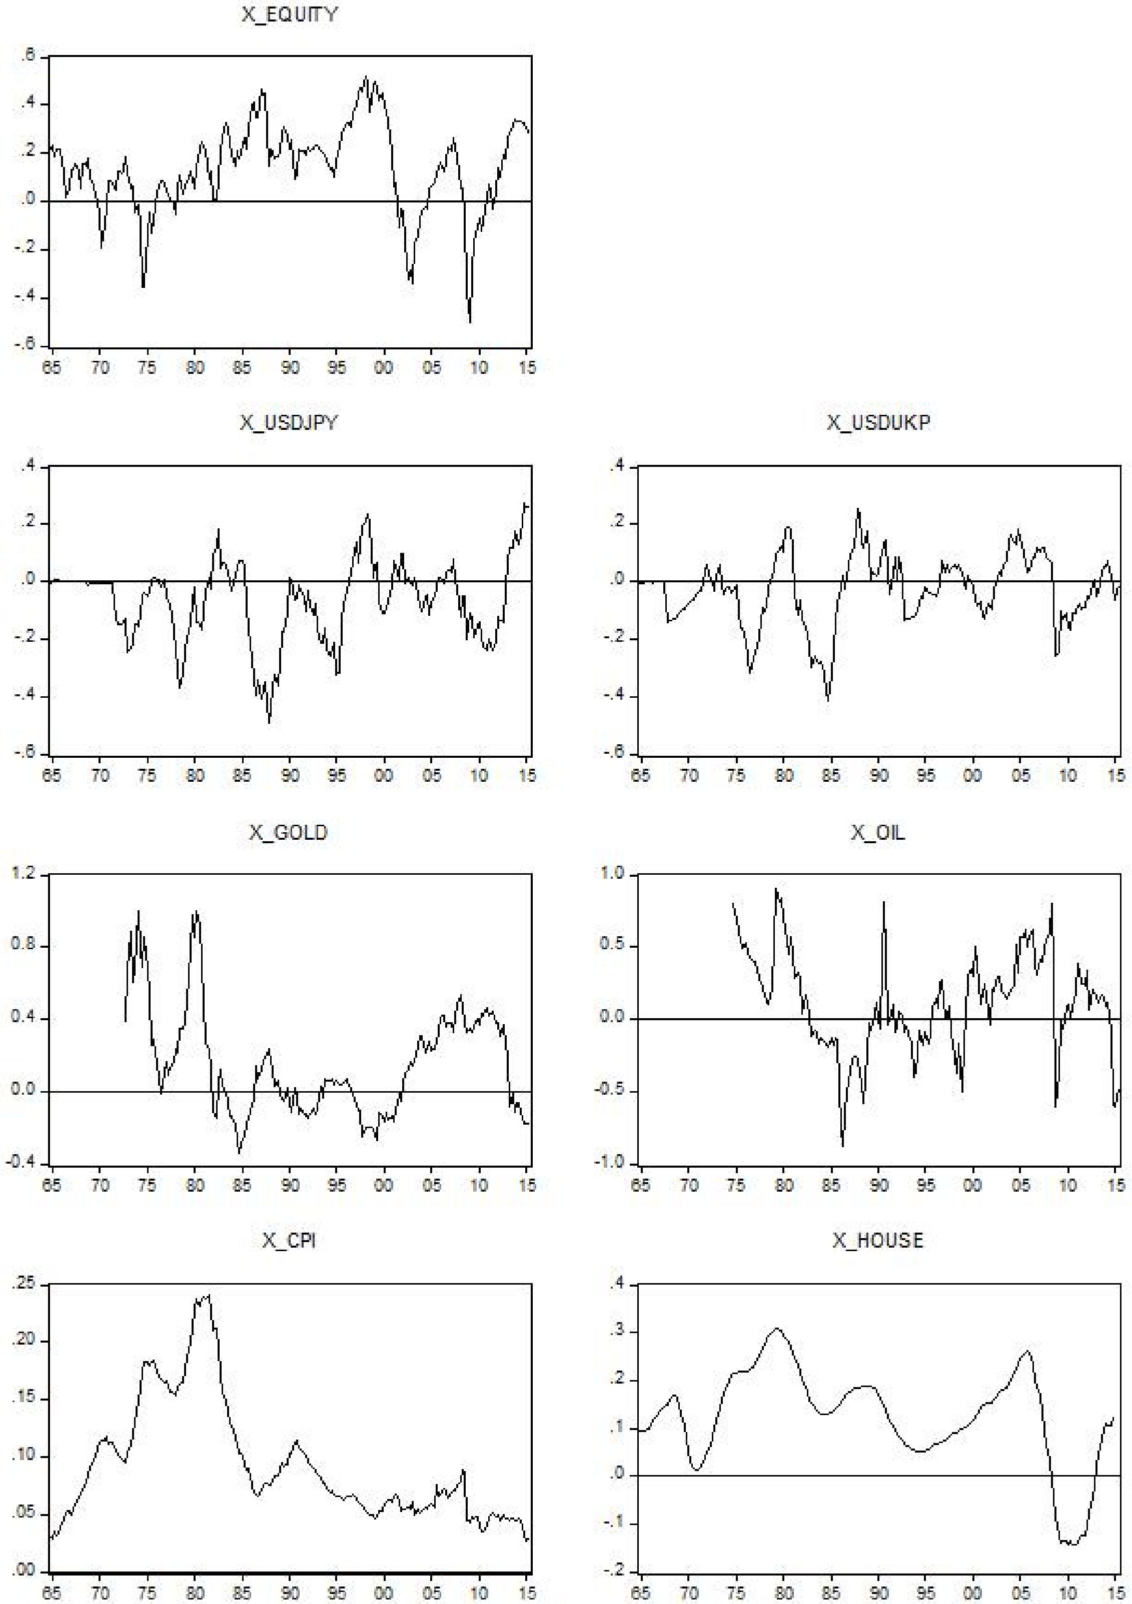

Notes: This figure presents x_t per asset class over time, defined as $ln(p_t) - ln(p_t^*)$, with $p_t^* = \sum_{i=1}^{20} p_{t-i+1}$.

Table 2: Benchmark Estimation Results

	Equity	Currencies		Commodities		Macro	
	S&P500	USDJPY	USDUKP	Gold	Oil	CPI	House
	Static						
ϕ	1.839***	1.866***	1.834***	1.908***	1.676***	1.989***	1.984***
	(40.042)	(36.115)	(35.498)	(48.046)	(17.263)	(114.19)	(111.25)
c	0.012	-0.003	-0.002	0.002	0.010	0.000	0.001
	(1.811)	(-0.705)	(-0.539)	(0.297)	(0.601)	(0.267)	(0.701)
LL	224.71	277.28	305.79	174.73	38.47	678.48	551.82
Obs	190	190	190	158	150	190	189
	Switching						
ϕ_C	1.276*	0.939***	1.194***	0.985***	0.990***	1.086***	1.414***
	(1.834)	(37.011)	(11.744)	(55.613)	(6.596)	(60.625)	(62.994)
ϕ_F	0.552	0.697***	0.566***	0.575***	0.439**	0.889***	0.667***
	(0.790)	(13.361)	(4.361)	(12.318)	(2.336)	(28.235)	(23.880)
β	0.142	30.949	1.034**	19.074	2.730	1.300***	2.125***
	(0.326)	(0.431)	(2.199)	(0.800)	(0.840)	(6.073)	(11.125)
c	0.011*	-0.007	-0.002	0.003	0.007	0.000	-0.001
	(1.727)	(-1.588)	(-0.420)	(0.477)	(0.485)	(0.261)	(-0.904)
LL	224.82	279.11	310.13	185.12	41.08	704.08	630.54
Obs	190	190	190	158	150	190	189
$P_{\phi_F=\phi_C}$	0.604	0.000	0.003	0.000	0.083	0.000	0.000

Notes: This table shows the results of estimating Equation (15) on quarterly data based on the moving-average fundamental value proxy. *, **, *** 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 Main Results

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

The top panel labeled 'Static' presents the results with $\beta = 0$, such that $n_{Ft} = n_{Ct} = 0.5 \forall t$. As a result, the effect size of the estimated coefficients is $0.5 * \phi$. 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., $0.5 * 1.989 = 0.9945$ for CPI, to more mean-reverting for the more financialized assets, such as $0.5 * 1.676 = 0.838$ 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 value, an average ϕ of 1.871 implies that the market expects prices to mean-revert with 6.45% per period ($1 - 0.5 * 1.871 = 0.0645$).

The bottom panel of Table 2 labeled '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 class, equities, 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 equities at the 10% significance level[6]. In other words, for six out of seven 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 the CPI. This implies that fundamentalists, on average, expect mean-reversion to occur 46% per period faster than chartists. For four out of seven assets, we observe that the chartist coefficient ϕ_C is larger than unity. 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 the market price will move away from the fundamental value 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. We further study the stability of the deterministic skeleton of the model using the estimated coefficient sets below.

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 US Dollar - UK Pound currency pair to 30.949 for the US dollar - Japanese yen currency pair[7]. 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 macro-economic variables, and therefore switch more between strategies in financial assets.

Figure 5 shows the estimated fundamentalist weights n_{Ft} for the assets for which we

[6]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.

[7]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)

find significant switching.

Figure 5 shows that the weights move around the average of 0.50. This is by construction of the switching function given by Equation (12). The variability and amplitude of the weights is related to the estimated β coefficients. For example, the weights for the dollar - yen exchange rate continuously jump between zero and one ($\beta = 30.9$) whereas the weights for CPI move slowly approximately between 0.2 and 0.8 ($\beta = 1.3$). The fluctuations in the fundamentalist weights can be explained by the movement in the price deviation x_t as shown in Figure 4. 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 Figure 4. Furthermore, we can recognize 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.

Figure 6 presents the estimated market sentiment, defined by $n_{Ft}\phi_F + n_{Ct}\phi_C$, for the assets for which we find significant switching. When this parameter exceeds one, the market expects an increase of $|x_t|$. When this is the case, we can speak of a locally unstable system because at that point in time price is not converging towards the fundamental equilibrium point of $x_t = 0$; further below we study the global stability properties of the estimated asset classes.

Figure 6 shows that 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 destabilizing. For a number of assets, we observe that sentiment is above one in certain periods (housing, CPI, and dollar-pound). In such episodes, agent behaviour is destabilizing, driving the price further away from its fundamental value. 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.

Whereas Figure 6 shows that for certain asset classes there are periods of market instability, we cannot conclude anything about the global stability of the estimated models, as studied in Section 2.2 above. To study the global stability of the model using the set of estimated coefficients per asset class, we take a simulation approach. Specifically, we take the set of estimated coefficients and run deterministic simulations to examine whether

Figure 5: Fundamentalist Weights

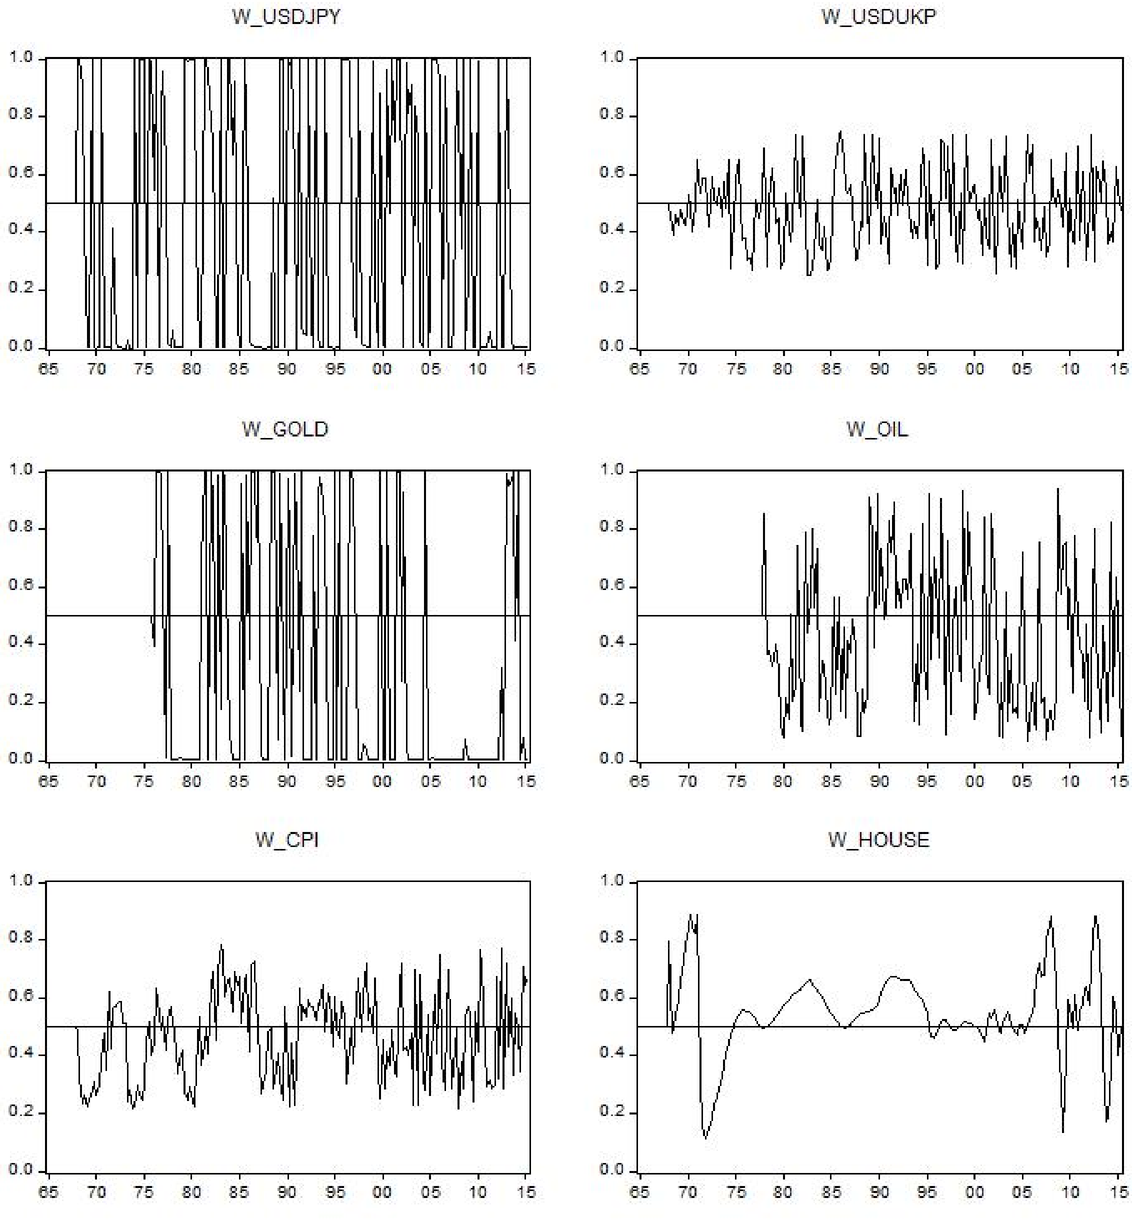

Notes: This figure presents estimated fundamentalist weights n_{Ft} for each asset class over time.

Figure 6: Market Sentiment

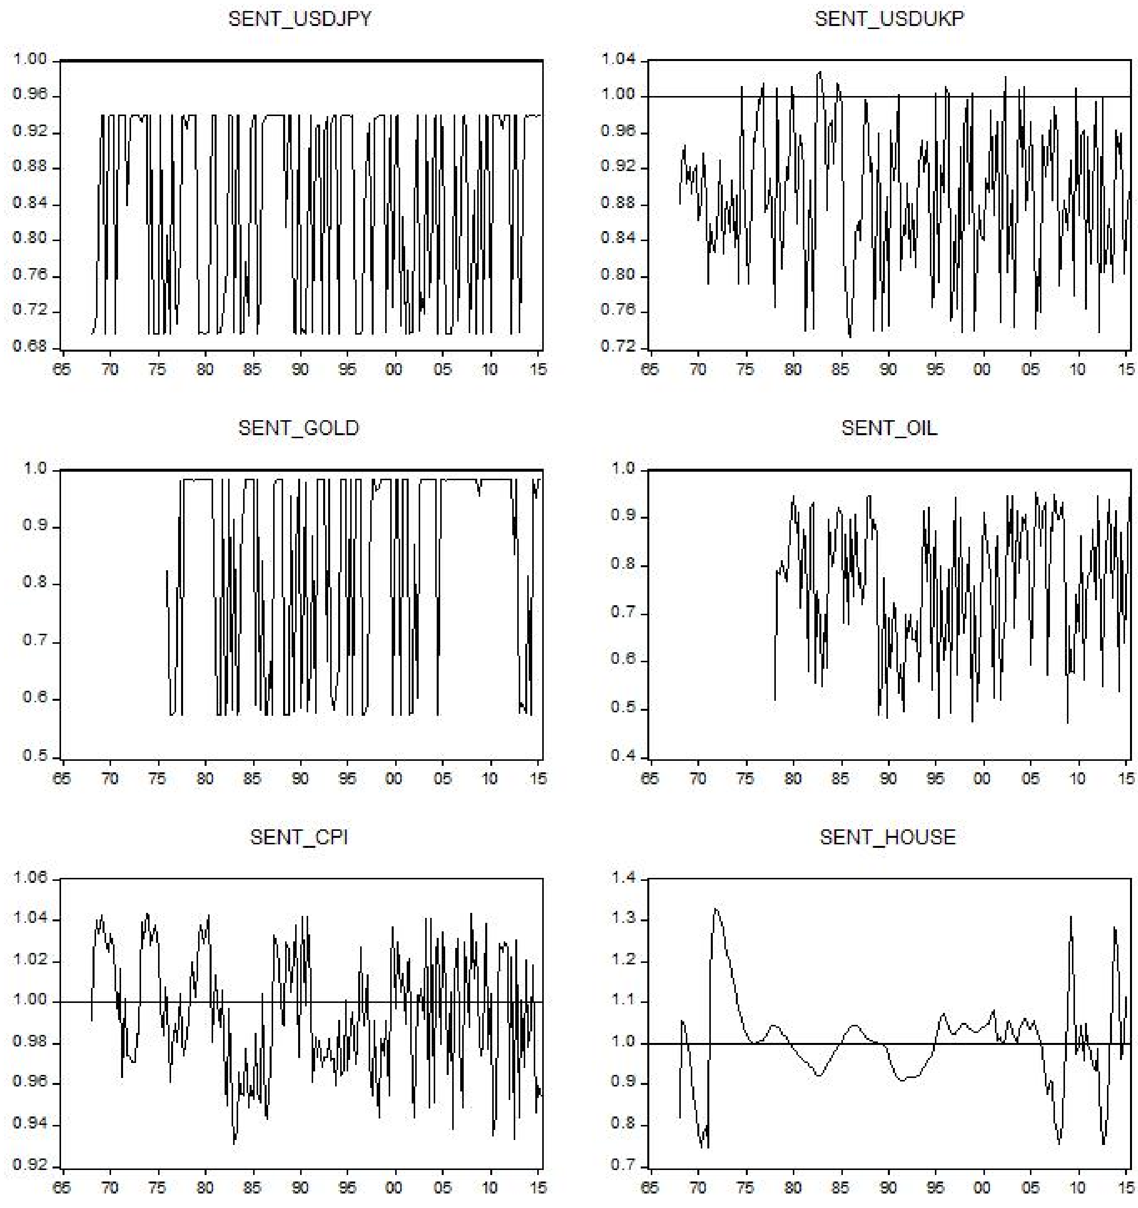

Notes: This figure presents the estimated market sentiment over time for the assets for which we find significant switching, defined as $n_{Ft}\phi_F + n_{Ct}\phi_C$.

the market price converges to its fundamental equilibrium $x_t = 0$. We set the starting values equal to $x_0 = 1$ and $n_{Ft} = n_{Ct} = 0.5$. Figure 7 presents the simulated price paths.

The top six panels of Figure 7 present the first 100 periods of the simulation for six asset classes; prices are kept constant for the first ten periods[8]. For the currencies and the commodities, we observe that the price falls to the fundamental value, i.e., x_t converges to zero. Although there is some variation in the speed of mean reversion, which is somewhat higher for the commodities than for the currencies, the market dynamics are globally stable for all four asset classes. Figure 6 showed that the dollar-pound market is locally unstable at points, but this does not translate to global instability although we do observe that mean-reversion is slowest for the dollar-pound currency within the group of currencies and commodities.

The results are somewhat different for the macro-economic assets. For the CPI, we observe that the price does not converge to the fundamental value within the first 100 periods of the observations; $x_{100} = 0.312$. For the house price index, we observe an explosive price path over the first 100 observations with $x_{100} = 34.702$. To see what happens after the first 100 periods, the bottom two panels of Figure 7 depict the first 700 periods of the simulation paths for the macro-assets. The CPI is highly persistent, but it does converge towards the fundamental equilibrium. The simulated house price, however, peaks at $3.35 \cdot 10^9$ in period 583 after which it drops and converges to the fundamental equilibrium $x_t = 0$.

Whereas Figure 7 presents the simulated price paths, Figure 8 presents the simulated fundamentalist weights n_{Ft} for the first 700 periods. Whereas the currencies and commodities all converge to the fundamental equilibrium within a reasonable number of periods, Figure 8 shows that the out-of-equilibrium dynamics are quite different across assets. For the dollar-yen currency pair and gold we observe that all agents immediately switch towards fundamentalism, causing x_t to converge to zero relatively quickly. For dollar-pound, the fraction of fundamentalists remains constant at its starting value of 0.5. This reflects the effect of the relatively low estimated β. For the oil market, weights also remain constant at 0.5 until period 120, after which it jumps to 0.9. The macro-economic variables show again a different picture. For the CPI, we observe a constant wiggle in the fraction of fundamentalists, although the amplitude is extremely small. The

[8] We do not include equity as we did not find significant evidence for behavioural heterogeneity. As a result, the model is stable by definition as $\phi < 2$ for equity in the top half of Table 2.

Figure 7: Deterministic Skeleton: Simulated X

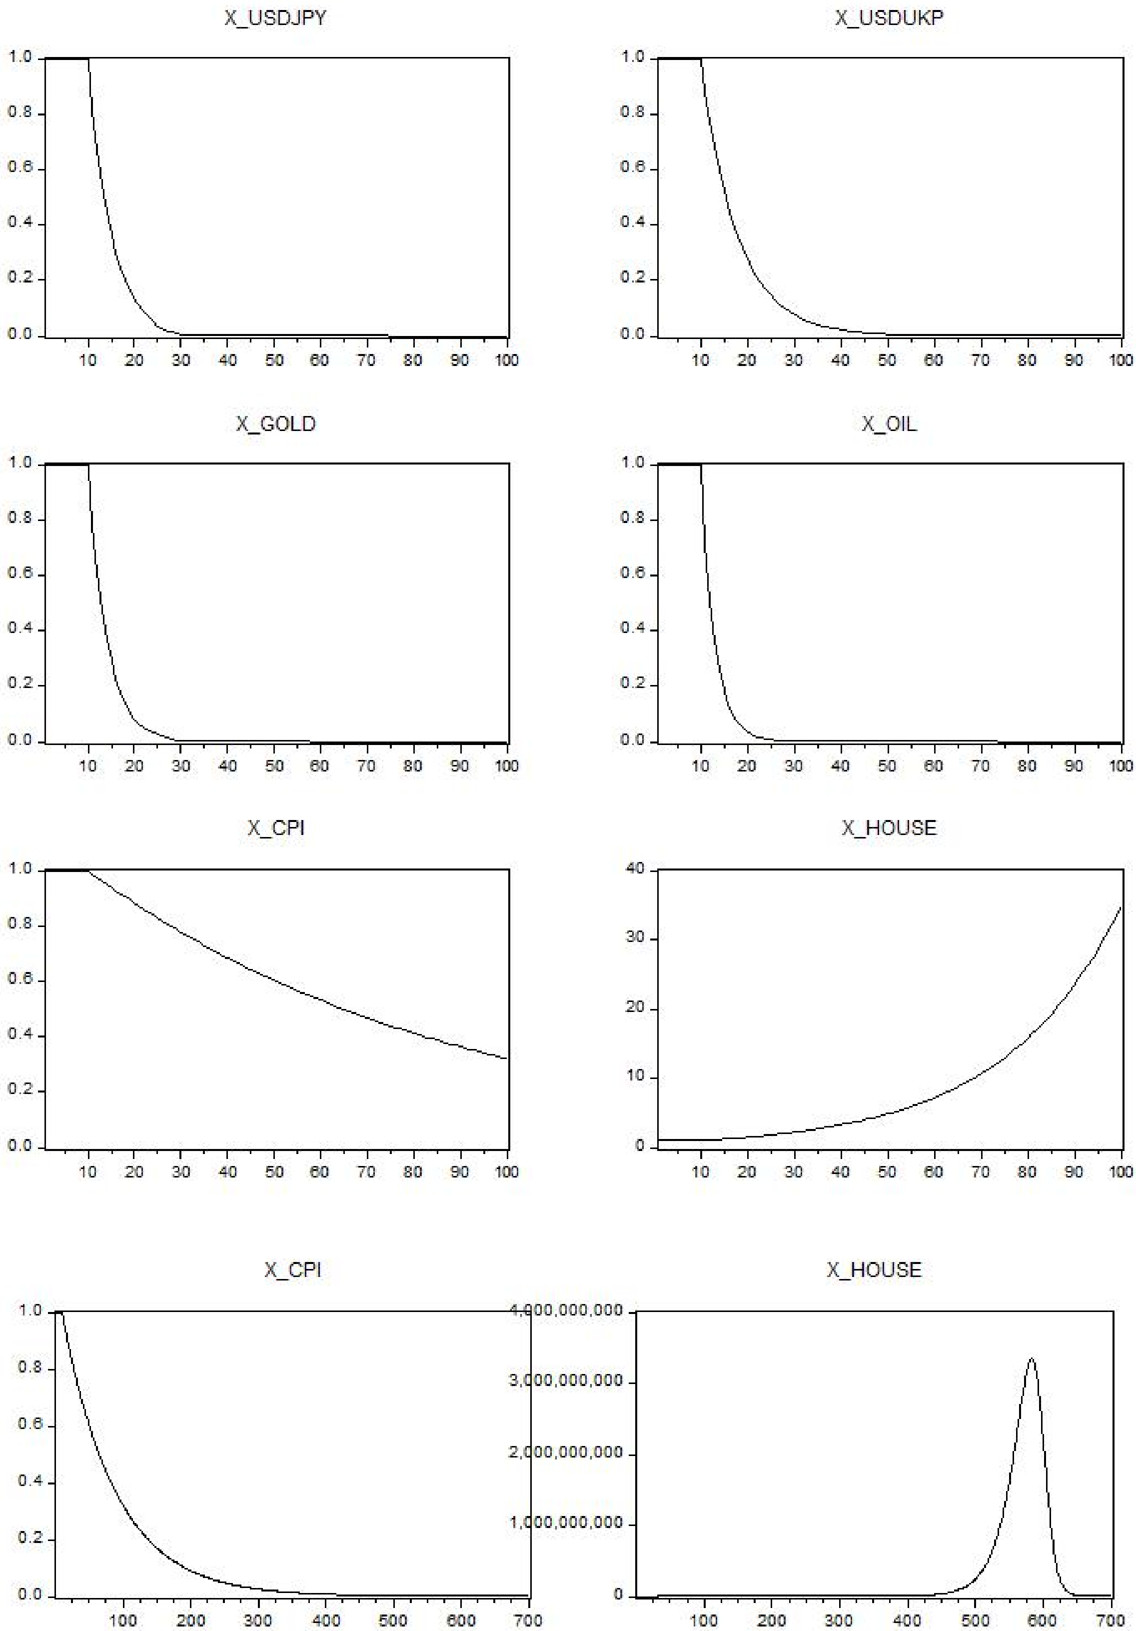

Notes: This figure presents the price divergence from fundamental value x_t obtained from simulations using the estimated coefficient sets.

Figure 8: Deterministic Skeleton: Simulated n_{Ft}

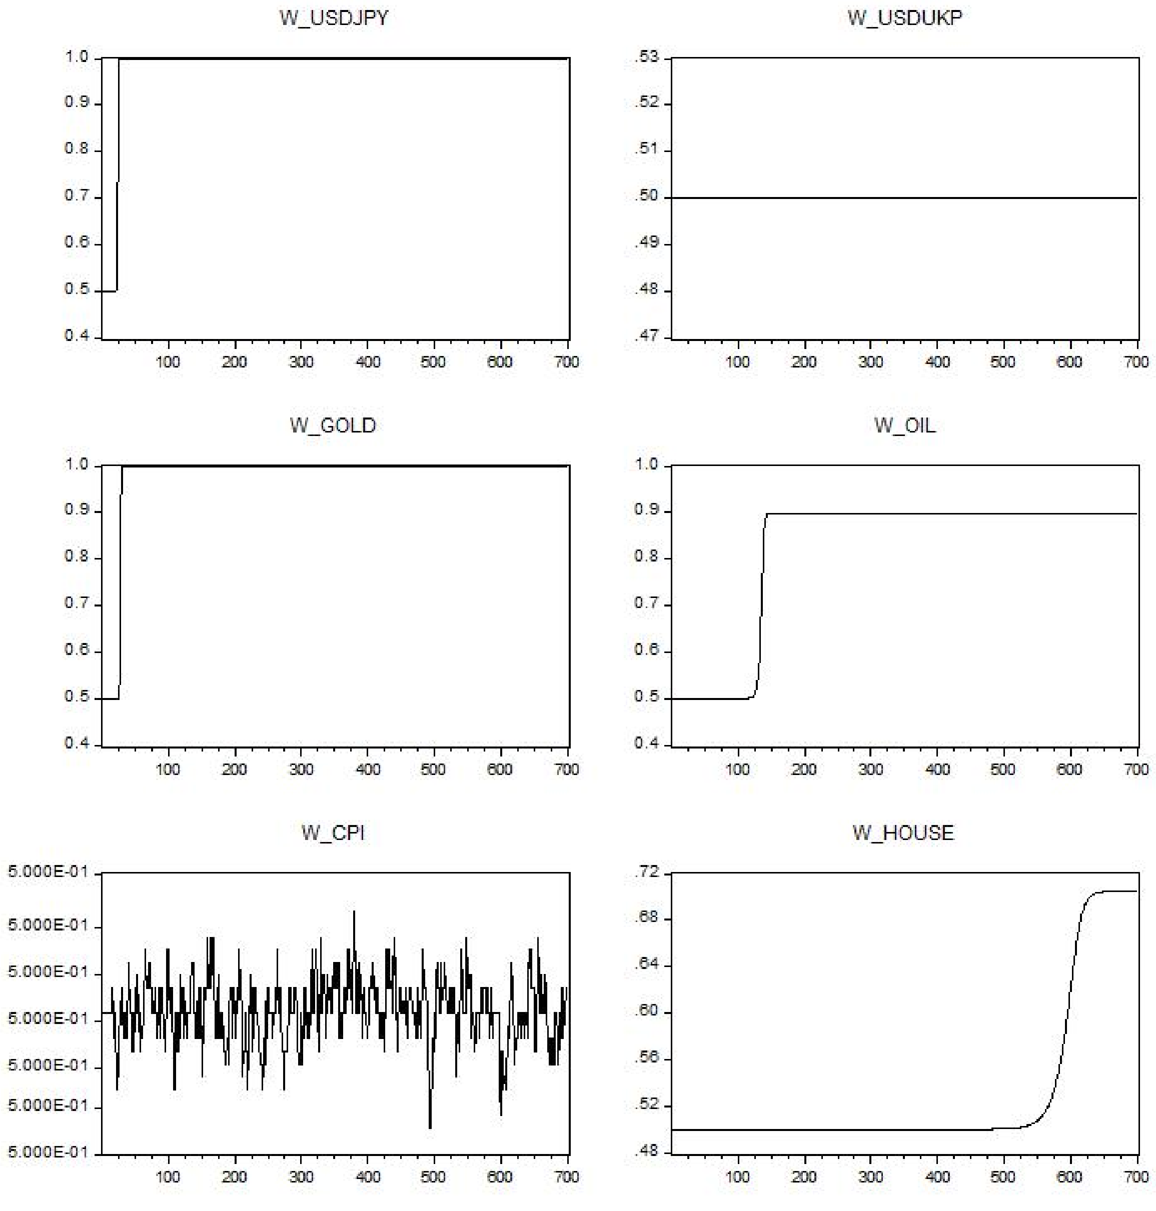

Notes: This figure presents the fundamentalist weights n_{Ft} obtained from simulations using the estimated coefficient sets.

simulated fraction of fundamentalists in the housing market remains constant at 0.5 for the first 500 periods. With $n_{Ft} = 0.5$, market sentiment $n_{Ft}\phi_F + n_{Ct}\phi_C = 1.04$ causing the explosive price path. At approximately $t = 600$, n_{Ft} jumps to 0.7, causing sentiment to drop below unity and thereby x_t to converge to zero.

5 Robustness

The choices we made for the benchmark configuration presented in the previous section were based on maximizing the number of assets and maximizing 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.

5.1 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 are 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 with 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 stability properties of the estimated models are comparable. Specifically, all assets converge to the fundamental equilibrium apart from the US dollar - UK pound exchange rate. This was not the case in the quarterly data, although we did find local instability for that market.

5.2 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

Table 3: Estimation Results Monthly

	Equity	Currencies		Commodities		Macro
	S&P500	USDJPY	USDUKP	Gold	Oil	CPI
			Static			
ϕ	1.964***	1.965***	1.950***	1.959***	1.925***	1.998***
	(150.4)	(134.9)	(129.1)	(146.6)	(96.26)	(514.2)
c	0.003	-0.001	-0.001	0.001	0.002	0.000
	(1.562)	(-0.519)	(-0.684)	(0.371)	(0.475)	(0.411)
LL	1133.0	1244.0	1287.2	725.4	450.0	2642.6
Obs	594	594	594	498	474	594
			Switching			
ϕ_C	1.010***	0.996***	1.009***	1.015***	0.989***	1.039***
	(35.53)	(122.0)	(33.66)	(92.14)	(8.122)	(304.9)
ϕ_F	0.946***	0.949***	0.911***	0.961***	0.939***	0.792***
	(31.66)	(79.23)	(29.17)	(105.0)	(7.599)	(43.26)
β	1.276	777.9	3.093	9.009	1.787	2.269***
	(0.654)	(0.030)	(0.729)	(0.889)	(0.101)	(32.44)
c	0.003*	-0.001	-0.001	0.000	0.002	0.001
	(1.735)	(-0.835)	(-0.606)	(-0.066)	(0.458)	(4.084)
LL	1133.9	1247.9	1292.6	729.5	450.1	2742.3
Obs	594	594	594	498	474	594
$P_{\phi_F=\phi_C}$	0.261	0.001	0.095	0.001	0.838	0.000
$Stability$	FP	FP	EX	FP	FP	FP

Notes: This table shows the results of estimating Equation (15) on monthly data using the moving-average fundamental value proxy. *, **, *** 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. FP denotes fixed point; EX denotes explosive.

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[9]. We obtain fundamentals for the assets under consideration in the following way:

- ***Equities***: 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).

Figure 9 plots the market prices together with 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 value proxies behave as one expects from a fundamental value. Specifically, we observe that the market price oscillates around the fundamental value proxies in all four cases[10], and that the fundamental value 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 value proxy.

Table 4 presents the estimation results using the exogenous fundamental values. The estimation results with exogenous fundamental values are again highly comparable with the benchmark results, especially for the dollar-pound currency pair and the housing market. The intensity of choice β is consistently around 30 for the dollar-yen and around 2 for the housing market. The estimates are slightly different for the equity market and

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

[10]Johanssen cointegration tests indicate that the price and fundamental value series are cointegrated for all four assets. Results available on request.

Figure 9: Exogenous Fundamentals

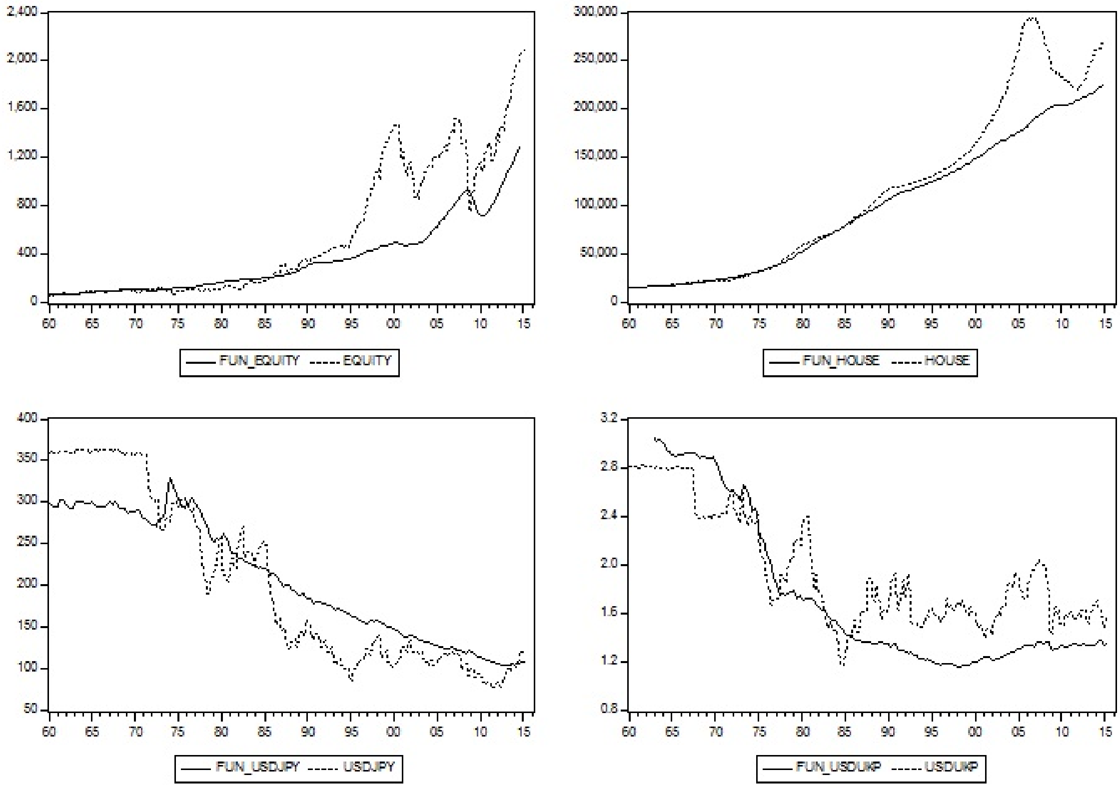

Notes: This figure presents the fundamental value proxies based on exogenous models and data combined with the actual market values.

Table 4: Estimation Results with Exogenous Fundamental

	Equity	Currencies		Macro
	S&P500	USDJPY	USDUKP	House
	Static			
ϕ	1.960***	1.922***	1.916***	1.989***
	(71.50)	(44.12)	(43.45)	(174.8)
c	0.006	-0.004	0.007	0.001
	(1.033)	(-0.737)	(1.388)	(0.970)
LL	235.1	305.5	319.3	616.0
Obs	207	209	197	207
	Switching			
ϕ_C	0.984***	0.978***	0.959***	1.586***
	(71.16)	(41.46)	(50.73)	(14.08)
ϕ_F	0.628***	0.920***	0.660***	0.247*
	(10.95)	(25.37)	(32.31)	(1.797)
β	52954	33.68	42.26	1.802***
	(0.000)	(0.204)	(0.370)	(23.04)
c	0.003	-0.006	0.010**	0.000
	(0.582)	(-0.973)	(2.399)	(0.181)
LL	236.6	306.8	323.5	748.0
Obs	207	209	197	207
$P_{\phi_C=\phi_F}$	0.000	0.114	0.000	0.000
Stability	FP	FP	FP	FP

Notes: This table shows the results of estimating Equation (15) on quarterly data with fundamental values based on exogenous data. *, **, *** 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. FP denotes fixed point; EX denotes explosive.

the dollar-pound. Results for the equity market are somewhat mixed. The likelihood ratio test still indicates that switching has no added value, but the Wald test now indicates that $\phi_F \neq \phi_C$. The dollar-pound results now indicate somewhat more intense switching than before.

5.3 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., $I = 1$. In this subsection we study the effect of this choice by increasing I 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 I is consistently accompanied by a decrease in the added value of switching. These results suggests that agents in the markets we consider do not consider more than one quarter of past performance when deciding to switch between groups.

5.4 Heterogeneous Beliefs about Returns

The model in Section 2 is written in terms of deviations from the fundamental value, 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))[11].

To illustrate the difference between the models estimated in deviations from the fundamental value and in returns, we will 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.

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

[11] 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 adjusts prices consistent with excess demand. See Hommes (2006) for an in-depth discussion about the micro-structure of agent-based models

Table 5: The Effect of Memory

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

Notes: This table shows the results of estimating Equation (15) on quarterly data using the moving-average fundamental proxy while varying the memory parameter I. *, **, *** 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: Estimation Results for Model in Returns

	Equity	Currencies		Commodities		Macro	
	S&P500	USDJPY	USDUKP	Gold	Oil	CPI	House
				Static			
ϕ_F	-0.053	-0.054	-0.099*	0.056	-0.225**	0.116***	-0.028***
	(-0.941)	(-1.008)	-(1.809)	(1.296)	(-2.080)	(5.562)	(-4.433)
ϕ_C	0.299*	0.204	0.378**	0.258*	0.151	0.689***	2.026***
	(1.962)	(1.255)	(2.611)	(1.916)	(1.105)	(6.569)	(38.93)
c	0.017**	-0.007	-0.003	0.008	0.019	0.001	0.002***
	(2.618)	(-1.446)	(-0.858)	(1.002)	(1.065)	(1.035)	(4.216)
LL	217.0	273.4	304.5	168.2	30.5	686.7	742.0
Obs	190	190	190	158	150	190	189
				Switching			
ϕ_F	-0.053	-0.094*	-0.099	0.049	-0.665***	0.032**	-0.014*
	(-0.930)	(-1.844)	(-1.611)	(1.573)	(-5.581)	(2.191)	(-1.866)
ϕ_C	0.299**	0.168	0.378**	0.271**	-0.056	0.751***	1.321***
	(1.961)	(0.939)	(2.609)	(2.635)	(-0.626)	(13.41)	(14.46)
β	0.000	7.049	0.000	18.28	76.29	7.807	1.743***
	(0.000)	(0.430)	(0.000)	(0.431)	(0.393)	(1.286)	(8.810)
c	0.017**	-0.007	-0.003	0.009	0.013	0.003***	0.001***
	(2.616)	(-1.629)	(-0.842)	(1.272)	(0.796)	(3.744)	(3.102)
LL	217.0	274.0	304.5	169.8	37.3	694.6	751.1
Obs	190	190	190	158	150	190	189
LR	0.00	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. *, **, *** represent significance at the 10%, 5%, and 1% level, respectively.

Fundamentalists expect the price to revert to its fundamental value:

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

Chartists extrapolate past returns:

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

The weights of the two groups are determined in the same way as in Equation (12), and the performance depends on the forecasting abilities of each strategy:

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

Table 6 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 estimation results for the model in returns shows a clear presence of both fundamentalists and chartists: For most assets, we find a negative ϕ_F suggesting the presence of fundamentalists, and a positive ϕ_C suggesting the presence of chartists. The CPI constitutes the exception, as fundamentalists expect the price level to deviate further from the fundamental value. This finding could be driven by the strong upward trend in the 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 market and the currencies. Within these assets, the ranking of speed of switching is comparable to the other configurations above. The consistency for the housing market is striking: again we find a β roughly equal to 2.

6 Conclusion

We have estimated a generic heterogeneous agent model on various asset classes, ranging from macro-economic variables such as CPI and house prices, to fast moving financial markets such as the foreign exchange market. We find that whereas switching is more intense in financial assets, the macro-economic variables are more unstable in the sense that behavioural bubbles have a higher probability of occurring. Our findings are qualitatively robust to the choice of fundamental, data frequency, and model configuration. The results have important implications for policymakers because whereas the focus has typically been on the risk of financial market volatility, we find that slow-moving macro variables pose a bigger threat to financial-economic stability.

References

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

Assenza, T., T. Bao, C. Hommes, and D. Massaro (2014). Experiments on expectations in macroeconomics and finance. *In J. Duffy (Ed.), Experiments in Macroeconomics*, 11–70.

Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of Financial Economics 49*(3), 307–343.

Baur, D. G. and K. J. Glover (2014). Heterogeneous expectations in the gold market: Specification and estimation. *Journal of Economic Dynamics and Control 40*(1), 116–133.

Bloomfield, R. and J. Hales (2002). Predicting the next step of a random walk: Experimental evidence of regime-shifting beliefs. *Journal of Financial Economics 65*(3), 397–414.

Bolt, W., M. Demertzis, C. Diks, C. Hommes, and M. Van Der Leij (2014). Identifying booms and busts in house prices under heterogeneous expectations. *De Nederlandsche Bank Working Paper No. 450*.

Boswijk, H. P., C. H. Hommes, and S. Manzan (2007). Behavioral heterogeneity in stock prices. *Journal of Economic Dynamics and Control 31*(6), 1938–1970.

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

Brock, W. A. and C. H. Hommes (1997). A rational route to randomness. *Econometrica 65*(5), 1059–1095.

Brock, W. A. and C. H. Hommes (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control 22*(8-9), 1235–1274.

Chen, H., S. Joslin, and N. K. Tran (2012). Rare disasters and risk sharing with heterogeneous beliefs. *Review of Financial Studies 25*(7), 2189–2224.

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

Chiarella, C., S. ter Ellen, X. Z. He, and E. Wu (2015). Fear or fundamentals? Heterogeneous beliefs in the European sovereign CDS market. *Journal of Empirical Finance 32*(1), 19–34.

Cornea-Madeira, A., C. H. Hommes, and D. Massaro (2017). Behavioral heterogeneity in U.S. inflation dynamics. *Journal of Business and Economic Statistics forthcoming.*

Cutler, D., J. Poterba, and L. Summers (1990). Speculative dynamics and the role of feedback traders. *American Economic Review 80*(2), 63–68.

De Jong, E., W. F. C. Verschoor, and R. C. J. Zwinkels (2010). Heterogeneity of agents and exchange rate dynamics: Evidence from the EMS. *Journal of International Money and Finance 29*(8), 1652–1669.

Delong, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann (1993). Noise trader risk in financial markets. *Journal of Political Economy 98*(4). Reprinted in R. H. Thaler (Ed.), Advances in Behavioral Finance, Russell Sage Foundation, 1993.

Fama, E. F. (1991). Efficient capital markets: II. *The Journal of Finance 46*(5), 1575–1617.

Frankel, J. A. and K. A. Froot (1990). Chartists, fundamentalists, and trading in the foreign exchange market. *American Economic Review 80*(2), 181–185.

Frijns, B., T. Lehnert, and R. C. J. Zwinkels (2010). Behavioral heterogeneity in the option market. *Journal of Economic Dynamics and Control 34*(11), 2273–2287.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. *Review of Financial Studies 27*(3), 714–746.

Hommes, C. H. (2006). Chapter 23 - Heterogeneous agents models in economics and finance. In *Handbook of Computational Economics*, Volume 2, pp. 1109–1186.

Hommes, C. H. and D. in 't Veld (2017). Booms, busts and behavioural heterogeneity in stock prices. *Journal of Economic Dynamics and Control 80*(1), 1533–1597.

Kouwenberg, R. and R. Zwinkels (2014). Forecasting the US housing market. *International Journal of Forecasting 30*(3), 415–425.

Lof, M. (2014). Rational speculators, contrarians and excess volatility. *Management Science 61*(8), 1889–1901.

Lux, T. (1995). Herd behaviour, bubbles and crashes. *The Economic Journal 105*(431), 881–896.

Lux, T. and R. C. Zwinkels (2017). Empirical validation of agent-based models, volume 4. *In B. LeBaron and C.H. Hommes (Eds.), Handbook of Computational Economics, Heterogeneous Agent Models forthcoming.*

MacDonald, R. (2000). Expectations formation and risk in three financial markets: Surveying what the surveys say. *Journal of Economic Surveys 14*(1), 69–100.

Muth, J. (1961). Rational expectations and the theory of price movements. *Econometrica 29*(3), 315–335.

Schmalensee, R. (1976). An experimental study of expectation formation. *Econometrica 44*(1), 17–41.

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review 71*(3), 421 – 435.

Spronk, R., W. F. C. Verschoor, and R. C. J. Zwinkels (2013). Carry trade and foreign exchange rate puzzles. *European Economic Review 60*(1), 17–31.

ter Ellen, S., W. F. C. Verschoor, and R. C. J. Zwinkels (2013). Dynamic expectation formation in the foreign exchange market. *Journal of International Money and Finance 37*(1), 75–97.

ter Ellen, S. and R. C. J. Zwinkels (2010). Oil price dynamics: A behavioral finance approach with heterogeneous agents. *Energy Economics 32*(6), 1427–1434.

Teräsvirta, T. (1994). Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association 89*(25), 208–218.

Westerhoff, F. and S. Reitz (2005). Commodity price dynamics and the nonlinear market impact of technical traders: Empirical evidence for the US corn market. *Physica A: Statistical Mechanics and its Applications 349*(3-4), 641–648.

Appendix

One of the reasons why we use the normalized profit 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 220[12] or $1,000$ observations per run to study the effect of small samples. The distribution of estimated values of $\widehat{\beta}$ will allow us to draw inference about the accuracy and efficiency of the estimation procedure. Figure 10 presents the average as well as the standard error of the $1,000$ estimated $\widehat{\beta}$'s for each level of β for $t = 220$.

Panel (a) of Figure 10 displays the results for the model with absolute profit differences. The upper left figure shows the average estimated $\widehat{\beta}$. Whereas the actual β ranges from zero to two, the estimated $\widehat{\beta}$ hovers around $1,000$ with a high level of variation. The rough coefficient estimate, therefore, is not informative about the underlying switching mechanism. We therefore winsorize the coefficient estimates at the 10% level[13]. 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 Figure 10 displays results for the model with relative profit differences. 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 $\widehat{\beta}$s is substantially lower (approx. 20 times) than for the model with absolute profit differences. After winsorizing 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 profit differences.

To formalize the findings in Figure 10, we estimate the following equation $\widehat{\beta}_i = c_1 + c_2\beta_i + \epsilon_i$. An unbiased $\widehat{\beta}_i$ should give $c_1 = 0$ and $c_2 = 1$. Table 7 presents the results.

The upper half of Panel (a) in Table 7 shows results for the absolute profit differences.

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

[13]Specifically, we remove the 10% highest values of $\left|\widehat{\beta}\right|$.

Figure 10: Estimation

(a) Absolute profit difference

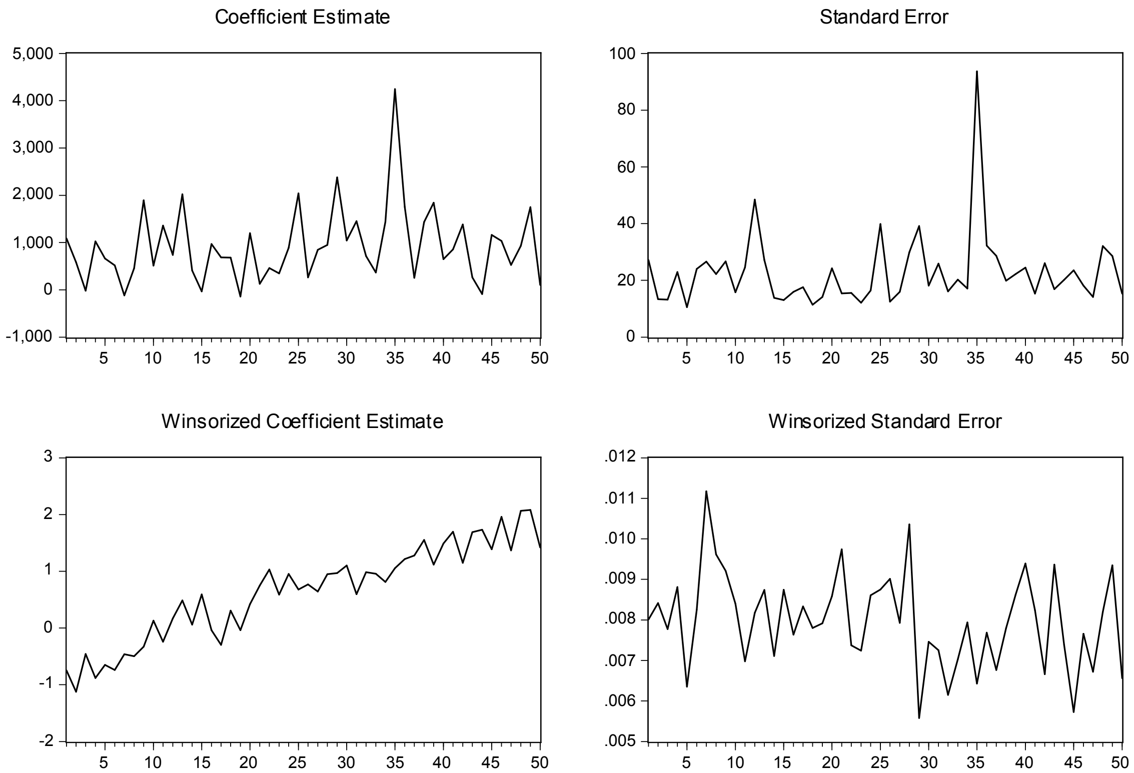

(b) Relative profit difference

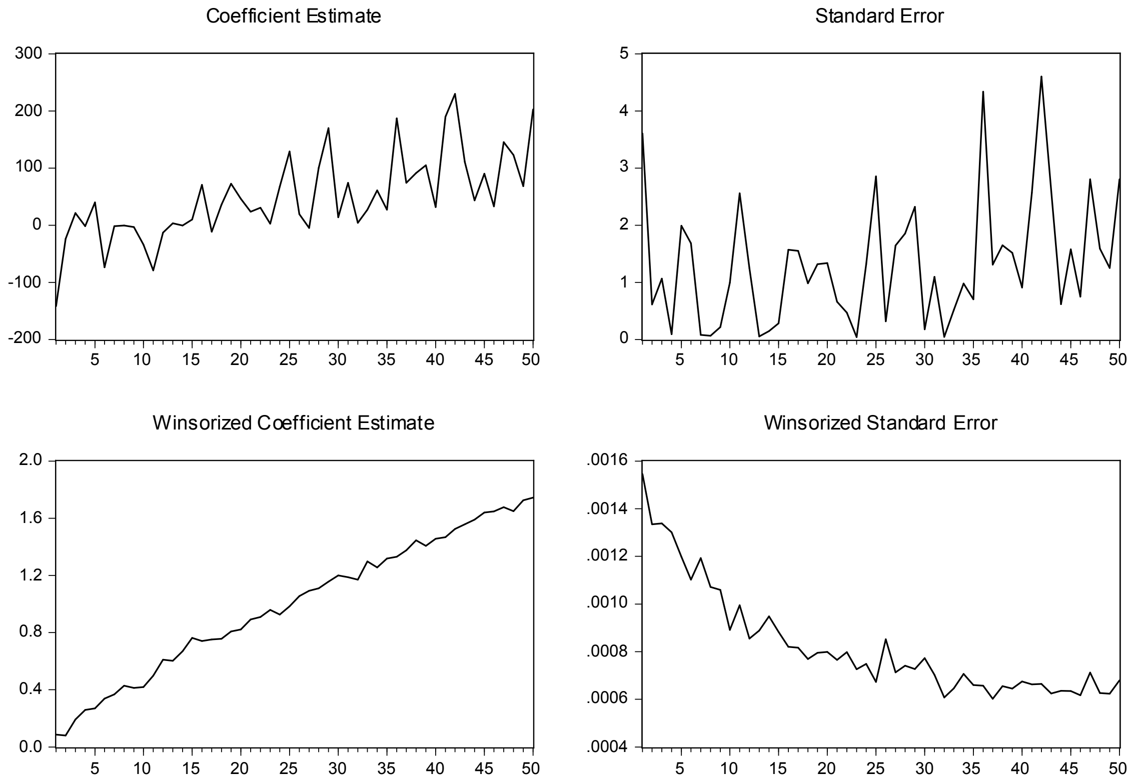

Notes: This figure shows estimated β coefficients for the absolute versus relative profit difference based on $t = 220$. The y-axis of the left column denotes the estimated β value, whereas the x-axis denotes the 'true' β of the underlying DGP, multiplied by 25. The y-axis of the right column denotes the estimated standard errors for different values of β.

Table 7: Estimation Bias

	Coeff.	St.Err.	Win. Coeff	Win. St.Err.
	Panel (a): t = 220			
	Absolute profit difference			
c_1	682.285***	20.101***	-0.743***	0.009***
	(4.778)	(8.067)	(-7.082)	(-27.816)
c_2	230.504	2.621	1.350***	-0.001***
	(1.492)	(1.054)	(18.433)	(-2.072)
R^2	0.01	-0.01	0.89	0.06
	Relative profit difference			
c_1	-41.915***	0.792***	0.151***	0.001***
	(-3.313)	(2.872)	(5.318)	(14.520)
c_2	87.813***	0.544***	0.825***	-0.000***
	(7.102)	(2.217)	(37.912)	(-5.352)
R^2	0.48	0.07	0.99	0.73
	Panel (b): t = 1,000			
	Absolute profit difference			
c_1	-18.451	5.807*	-0.020*	0.002***
	(-0.181)	(1.746)	(-1.692)	(47.62)
c_2	399.84***	6.913*	1.009***	0.000**
	(3.025)	(1.878)	(54.72)	(-2.601)
R^2	0.11	0.03	0.98	0.11
	Relative profit difference			
c_1	-0.019	0.006	0.097***	0.000***
	(-0.117)	(1.389)	(8.897)	(15.21)
c_2	1.318***	0.003	0.906***	0.000***
	(5.864)	(0.503)	(108.23)	(-3.557)
R^2	0.50	-0.02	1.00	0.54

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

Consistent with Figure 10, we observe that the estimated $\widehat{\beta}$ is not informative regarding the actual β. After winsorization, we do observe a positive and significant c_2, although it is significantly larger than unity. The winsorized standard error decreases somewhat as β increases. The lower half of Panel (a) in Table 7 shows the results for the relative profit differences. The raw $\widehat{\beta}$ estimates show a positive correlation with β, although the c_2 coefficient is far too high. After winsorization, the estimates are very well behaved, with $c_1 = 0.151$ and $c_2 = 0.825$.

Panel (b) in Table 7 presents the results for the simulation of $t = 1,000$ periods. Overall, the results are very similar to those with $t = 220$ indicating that the differences between the models with absolute and relative profit differences are not driven by the small sample issue. We do observe that the non-winsorized 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.