

TI 2015-088/II

Tinbergen Institute Discussion Paper



Booms, Busts and Behavioural Heterogeneity in Stock Prices

Cars Hommes

Daan in 't Veld

Faculty of Economics and Business, University of Amsterdam, and Tinbergen Institute, the Netherlands.

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Booms, busts and behavioural heterogeneity in stock prices

Cars Hommes*, Daan in 't Veld**

This draft: July 2015

Abstract

We empirically evaluate a behavioural model with boundedly rational traders who disagree about the persistence of deviations from the fundamental stock price. Fundamentalist traders believe in mean-reversion, while chartists extrapolate trends. Agents gradually switch between the two rules, based upon their relative performance, leading to self-reinforcing regimes of mean-reversion and trend-following. For the fundamental price we use well-known models of Gordon (1962) and Campbell and Cochrane (1999). We estimate the two-type switching model using U.S. stock prices until 2012Q4 and find significant behavioural heterogeneity. Our model suggests that behavioural regime switching strongly amplifies booms and busts in stock prices.

Keywords: behavioural finance, bounded rationality, heterogeneous expectations, stock prices, financial crisis

JEL classifications: G12; C22; G01.

[☆]The authors would like to thank Cees Diks and Marco van der Leij for sharing programming code and helpful discussions, and Jules van Binsbergen, Peter Boswijk, William Branch, Peter Broer (discussant), Lennart Dek, George Evans, Bochen Li (discussant), Sebastiano Manzan, Simon van Norden and Remco Zwinkels for useful comments. We also thank participants at the Computing in Economics and Finance conference CEF2012 (Prague), the INEXC conference 2013 “Expectational Coordination and the Stock Market” (Paris), the Midwest Finance Association Annual Conference 2014 (Orlando), the Research in Behavioural Finance Conference 2014 (Rotterdam) and seminars at the University of Amsterdam, Tinbergen Institute, UC Irvine, the University of Oregon, the University of Pisa, the CPB Netherlands Bureau for Economic Policy Analysis, and SEO Economic Research. The paper has been accepted for the European Finance Association conference in Vienna, 19-22 August 2015. The research has been supported by the Netherlands Organisation for Scientific Research (NWO) under the Complexity program *Understanding financial instability through complex systems* (grant no. 645.000.015).

*Amsterdam School of Economics and CeNDEF, University of Amsterdam and Tinbergen Institute, C.H.Hommes@uva.nl

**Corresponding author: Amsterdam School of Economics and CeNDEF, University of Amsterdam and Tinbergen Institute, Valckeniersstraat 65-67, 1018 XE Amsterdam, the Netherlands, D.intVeld@uva.nl, tel.: ++31 20 525 4335, fax.: ++31 20 525 4349.

1. Introduction

Economic reality shows the limitations of standard asset pricing models with a representative rational agent only concerned with economic fundamentals. In 2008 the S&P500 stock index, the financial bellwether of the dominant market economy of the U.S., and many other stock indices, lost around one half of their total value. While the bankruptcy of Lehman Brothers amounted to a clear fundamental shock to the economy, it is difficult to explain all of this loss as a rational re-evaluation of fundamentals. Other behavioural explanations need to be considered. In this paper we present evidence from S&P500 data that market sentiment switches between different behavioural regimes, which amplified shocks such as the Lehman bankruptcy, and more generally amplifies booms and busts of the economy.

We first apply the idea of switching market sentiment to a basic framework that provides a fundamental value of the price-dividend ratio: the standard Gordon solution based on a constant risk premium. Within this framework we introduce a simple behavioural model with some agents believing in mean-reversion of stock prices (called fundamentalists) and others (called chartists) who expect a continuation of the trend. Agents gradually switch between the two rules, based upon their relative performance, so they learn and adapt their behaviour if the market situation changes and the losses of their strategy become too large. Because of the positive expectations feedback in asset markets, self-reinforcing behavioural regimes of mean-reversion and trend-following arise endogenously in the model, explaining large and persistent deviations of the S&P500 from the Gordon fundamental value.

A convenient feature of our model is that it is formulated in deviations from a fundamental price, so that it can be tested against any suitable fundamental benchmark. Behavioural heterogeneity can therefore complement the mainstream financial literature on stock market fluctuations by providing an *amplification mechanism* to explain excess volatility (Shiller, 1981). To this end we combine our model with the consumption-habit asset pricing model of Campbell and Cochrane (1999). They argue in a standard representative-agent framework that booms and busts in asset prices are driven by countercyclical variation in risk premia, which in turn are inversely related to consumption relative to a slow-moving habit level. We show that even if part of the variation in the price-dividend ratio can be explained by consumption-driven variation in risk premia, our model still gives significant parameter estimates and adds explanatory power due to behavioural heterogeneity. Overall, we argue that there is strong evidence for heterogeneous beliefs amplifying booms and busts in the stock market.

Standard asset pricing models do not take heterogeneity into account as these models assume the expectations of individual investors are rational and can be described by a representative agent. Asset prices should in this view equal the *fundamental value* of expected discounted sum of future cashflows, or more specifically dividend payoffs (Campbell and Shiller, 1988a). Various

reasons have been proposed why this fundamental value could change over time, as in Campbell and Cochrane (1999). Bansal and Yaron (2004) argue for the effects of long-run economic uncertainty on asset prices. Pástor and Veronesi (2006) and Ofek and Richardson (2003) give particular (but very different) explanations for the high valuations of technology firms in the late 1990s. Nevertheless, these explanations may not be sufficient to fully explain stock market fluctuations. More specifically, we show that for the consumption-habit model of Campbell and Cochrane (1999), behavioural heterogeneity is a significant amplification mechanism.

With the contention that the financial crisis cannot be sufficiently explained by economic fundamentals, our paper fits within the *behavioural* finance literature. Departing from the strongest form of rationality opens up the alternative view that stock prices may have been overpriced. The behavioural finance literature is surveyed in e.g. Hirshleifer (2001) and Barberis and Thaler (2003). Barberis and Thaler (2003) stress the finding that traders with flawed expectations can not always be driven away from the market. As these traders distort supply and demand based on fundamentals, assets can be partly mispriced. In their words: “One of the biggest successes of behavioral finance is a series of theoretical papers showing that in an economy where rational and irrational traders interact, irrationality *can* have a substantial and long-lived impact on prices.” (Barberis and Thaler, 2003, p. 1053, their emphasis).

Barberis and Thaler (2003) also state that careful empirical analysis remains the main challenge for behavioural models. As one recent example, Branch and Evans (2010) develop a framework with agents learning the parameters of their underparameterised forecasting models and reproduce regime-switching returns and volatilities in monthly U.S. stock data. Adam and Marcet (2011) and Adam et al. (2013) provide another example where investors’ subjective beliefs are shown to drive booms and busts in the S&P 500’s price-dividend ratio. In these examples however the model is calibrated to replicate certain characteristics in the data. Moreover, these models assume learning by a homogeneous representative agent: see Pástor and Veronesi (2009) for a stimulating survey. Our simple behavioural model assumes heterogeneous agents and contains few parameters that can be estimated directly.

We will model our boundedly rational traders within the *heterogeneous agents* asset pricing framework of Brock and Hommes (1997, 1998)¹. The literature on heterogeneous agents models (HAMs) has been growing in the last decades and is extensively reviewed in e.g. Hommes (2006), LeBaron (2006) and Lux (2009). For example, HAMs have been applied to stock prices empirically in Boswijk et al. (2007), Chiarella et al. (2014) and Lof (2012, 2014). Switching models with heterogeneous agents have also been applied to other financial markets, in particular

¹Other related early heterogeneous agents models include the noise trader models of DeLong et al. (1990a,b), the model with ‘newswatchers’ versus momentum traders of Hong and Stein (1999) and the evolutionary belief models of Jouini and Napp (2010) and Jouini et al. (2012). These models also assume bounded rationality of (at least one type of) agents, but do not allow for switching between different strategies.

exchange rates (Kirman and Teyssière, 2002; Westerhoff and Reitz, 2003; Alfarano et al., 2005; de Jong et al., 2010), but also for example to option prices (Frijns et al., 2010) and oil prices (ter Ellen and Zwinkels, 2010). This empirical literature is growing fast, see e.g. Chen et al. (2012) for an overview.

Our paper makes four contributions to the empirical literature on behavioural asset pricing. Most importantly, we generalise the asset pricing model with heterogeneous agents and test it against two benchmark fundamentals, the Gordon model and the Campbell-Cochrane consumption-habit model. A second novelty in the literature is that we introduce agents' memory of earlier realised excess returns. This will lead to gradual (rather than instant) switching and makes the model applicable to quarterly data with a simple economic interpretation. A third, methodological contribution is to run Monte Carlo simulations to clarify two difficulties in estimating HAMs: the stationarity of the time series and the significance of the switching intensity. Finally, we look in greater detail at the price dynamics in the recent turbulent years in terms of fundamentals and amplification mechanisms, as our time series includes both the dot-com bubble and the global financial crisis. For example, for the Campbell-Cochrane consumption-habit fundamental, our model explains the financial crisis as being triggered by an exogenous shock (the Lehman Brothers bankruptcy) and strongly amplified by coordination on trend-following behaviour.

Many factors have contributed to the rising interest in behavioural heterogeneity. First, laboratory experiments with human subjects have been performed to study individual expectations and aggregate outcomes, e.g. Hommes et al. (2005, 2008). Experimental studies have the benefit that the underlying asset market fundamentals can be fully controlled; for an overview of the use of laboratory experiments to test for heterogeneous expectations, see Hommes (2011). Anufriev and Hommes (2012) find in experimental asset pricing data that subjects switch between different forecasting rules, consistent with the theoretical model of Brock and Hommes (1998). An interesting finding from these laboratory experiments is that under *positive expectations feedback* coordination on trend-following strategies amplifies asset market fluctuations (Heemeijer et al., 2009).

Second, empirical evidence has shown that switching based on past performance is relevant for real financial markets. For example, Ippolito (1989), Chevalier and Ellison (1997), Sirri and Tufano (1998) and Karceski (2002) found in mutual funds data that money flows out of past poor performers into good performers. Pension funds also switch away from bad performers (Del Guercio and Tkac, 2002). Investors in the stock market can be expected to display similar switching behaviour when choosing between different strategies.

Third, there is growing interest in survey data on expectations of financial specialists, which can be traced back to Frankel and Froot (1987). Comparing six different data sources, Green-

wood and Shleifer (2013) show that surveys of stock market investors are highly positively correlated with each other, supporting the idea that they do reflect actual beliefs. The heterogeneity in price expectations also changes over time, as shown in Shiller (1987, 2000) and Vissing-Jorgensen (2004). All three forms of microlevel evidence of actual traders shifting between simple behavioural rules motivate our aggregate model.

We emphasise that the behaviourally heterogeneous expectations of our investors are not model-consistent as in the traditional rational expectations framework. Yet agents are boundedly rational in the sense that they switch to better performing rules, which then become almost self-fulfilling. We show that the data supports self-reinforcing temporary coordination on either mean-reversion or trend-following. Still, on which type of behaviour agents will coordinate is difficult to foresee in advance: the market is unpredictable in the short run. Fundamentals play a complementary role in explaining mean-reversion in stock market fluctuations and make prices predictable in the long run. Overall, strategy switching serves as an amplification mechanism for booms and busts.

The paper is organised as follows. Section 2 develops the general asset pricing model with heterogeneous agents. In Section 3 we present our main estimation results under the standard Gordon fundamental value. Section 4 provides Monte Carlo simulations to test the robustness of our results, as well as simulated time series generated by our model that illustrate the endogenous behavioural regimes. In Section 5 we combine our model with the consumption-habit model of Campbell and Cochrane (1999) that has a time-varying risk premium. Section 6 concludes. Further robustness checks are provided in the Appendix.

2. Model description

We derive a stylised asset pricing model with heterogeneous agents, generalising Brock and Hommes (1998) and Boswijk et al. (2007) to a model that allows for time-variation in dividends and discount rates.² We assume that investors have perfect knowledge of the underlying fundamental process, and are therefore able to calculate the ‘fundamental value’, which in this section will be derived in general terms. The general form of our model to be estimated is

$$x_t = \frac{1}{R^*} \sum_{h=1}^H n_{h,t} E_{h,t}[x_{t+1}], \quad (1)$$

²In this paper we focus on the price-dividend ratio as it is the classical way to model stock prices. There is also a large literature of modeling asset prices based on book values or earnings, e.g. Campbell and Shiller (2001). Boswijk et al. (2007) estimate an asset pricing model with behavioural heterogeneity both for fundamental valuation based on dividends and earnings, and find robust results.

where x_t is the price-dividend ratio in *deviations from the fundamental value*, $n_{h,t}$ the fractions of agents having belief $E_{h,t}$, and $1/R^*$ is the expected effective discount factor, to be specified below. An advantage of this general form is that the model can be estimated around any benchmark fundamental valuation model. In Section 3, we will specify the present value model with a constant risk premium based on Gordon (1962) for the fundamental value. In Section 5, we consider another benchmark fundamental value with a time-varying risk premium, the consumption-habit model of Campbell and Cochrane (1999).

Even though agents know the fundamental value, they have different beliefs about how the price of the asset deviates from its fundamental. The changes of agents' beliefs will lead to fluctuating market sentiment. For both benchmark fundamental value models in Sections 3 and 5, we will show that there is significant evidence for behavioural heterogeneity in the data.

An important idea of our model is to separate behavioural factors influencing prices from fundamental factors. Our main assumptions aim to model heterogeneous beliefs of investors on top of an asset pricing framework that is as general and flexible as possible. This two-step approach allows us to estimate the model with heterogeneous beliefs in *deviations from any* specification of the underlying fundamental process. In particular, our general heterogeneous agents model can be directly applied to the Campbell-Cochrane fundamental asset pricing benchmark. In Section 2.1 we discuss the fundamental value, and in Section 2.2 the behavioural expectation rules of the agents. In Section 2.3 we present the econometric form of our model.

2.1. Fundamental value

Consider a risky financial asset that pays a random dividend payoff D_t at time t . The opportunity cost for investing in the risky asset is captured by the discount rate R_{t+1} which may in general vary over time. The standard pricing equation (cf. e.g. Cochrane, 2001, p. 10) is

$$P_t = E_t \left[\frac{P_{t+1} + D_{t+1}}{R_{t+1}} \right]. \quad (2)$$

As Cochrane (2001, p. 37) emphasises, this equation does *not* presupposes a representative agent; it rather applies to each individual investor. Today's price is the expected discounted sum of tomorrow's price and of tomorrow's dividend payoff. Notice that R_{t+1} here refers to an objective *ex ante* discount rate about which agents have identical expectations, as specified below. The heterogeneity will apply to expectations about the future price P_{t+1} .

We focus on possible belief disagreement about future prices, but assume agreement about fundamentals. This approach reflects the idea of investors that prices are determined endogenously and partly depend on expectations about the next period's price, while the fundamentals follow an exogenous stochastic process. Thus, all agents have identical beliefs about the dividend

$D_{t+1} = (1 + g_{t+1})D_t$ and its discounted value:

$$E_t \left[\frac{D_{t+1}}{R_{t+1}} \right] = E_t \left[\frac{1 + g_{t+1}}{R_{t+1}} \right] D_t. \quad (3)$$

We can therefore rewrite the pricing equation (2) in terms of the price-dividend ratio (PD) ratio $\delta_t \equiv P_t/D_t$ as

$$\begin{aligned} \delta_t &= E_t \left[\frac{1}{R_{t+1}} \frac{D_{t+1}}{D_t} (\delta_{t+1} + 1) \right] \\ &= E_t \left[\frac{1 + g_{t+1}}{R_{t+1}} (\delta_{t+1} + 1) \right]. \end{aligned} \quad (4)$$

The *fundamental value* P_t^* is obtained under rational expectations from the present value of all future cash flows (see e.g. Boswijk et al., 2007, p. 1965). Substituting the basic pricing equation (2) forward under rational expectations, applying the law of iterated expectations, and imposing the transversality condition leads to

$$\begin{aligned} P_t^* &= E_t \left[\sum_{j=1}^{\infty} \left(\prod_{k=1}^j \frac{1}{R_{t+k}} \right) D_{t+j} \right] \\ &= E_t \left[\sum_{j=1}^{\infty} \left(\prod_{k=1}^j \frac{1 + g_{t+k}}{R_{t+k}} \right) D_t \right], \end{aligned} \quad (5)$$

and the fundamental PD ratio equals

$$\delta_t^* = E_t \left[\sum_{j=1}^{\infty} \left(\prod_{k=1}^j \frac{1 + g_{t+k}}{R_{t+k}} \right) \right]. \quad (6)$$

The fundamental values of the price and PD ratio are presented here in the most general form, but will simplify in subsequent sections to the special cases of Gordon (1962) and Campbell and Cochrane (1999) by assumptions on g_{t+1} and R_{t+1} .

We assume that all agents know the fundamental price P_t^* , but disagree about how the next period's price will deviate from P_{t+1}^* , e.g. because some investors may use non-fundamental trading decisions. Therefore agent type h tries to predict the next period's P_{t+1} by its subjective expectations $E_{h,t}[P_{t+1}]$ which may differ from P_{t+1}^* . To allow for a stationary underlying process, we focus on the subjective expectation about the price-dividend ratio $E_{h,t}[\delta_{t+1}]$, which consequently also possibly differ from δ_{t+1}^* .

2.2. Behavioural heterogeneous beliefs

To model heterogeneity we consider H types of investors using different expectation rules. The fractions or weights of agents using a particular belief $E_{h,t}$ are denoted by $n_{h,t}$. We assume that the PD pricing equation (4) holds at the aggregate level, averaging over all agents' expectations, i.e.

$$\delta_t = \sum_{h=1}^H n_{h,t} E_{h,t} \left[\frac{1+g_{t+1}}{R_{t+1}} (\delta_{t+1} + 1) \right]. \quad (7)$$

Equation (7), with the market price reflecting average beliefs, is typically derived in an underlying model with market clearing and an appropriate utility function.³ The fractions $n_{h,t}$ of beliefs are endogenous and will be modelled below.

We further specify the pricing process by separating behavioural heterogeneity from fundamental factors. More precisely, we assume common beliefs on fundamental factors such as growth rates g_{t+1} and discount rates R_{t+1} . We also assume that agents' behavioural beliefs $E_{h,t}[\delta_{t+1}]$ are independent of objective expectations about fundamental factors, i.e.

$$E_{h,t} \left[\frac{1+g_{t+1}}{R_{t+1}} (\delta_{t+1} + 1) \right] = E_t \left[\frac{1+g_{t+1}}{R_{t+1}} \right] E_{h,t} [\delta_{t+1} + 1]. \quad (8)$$

In order to simplify the behavioural beliefs around the fundamental δ_{t+1}^* , we define the *unconditional* expectation $1/R^* \equiv E \left[E_t \left[\frac{1+g_{t+1}}{R_{t+1}} \right] \right]$ as the *expected effective discount factor* for pricing stocks in terms of the PD ratio. Hence, our behavioural assumption is that the expected effective discount factor is constant, while the rational-expectations fundamental δ_t^* may be time-varying. Equation (7) becomes

$$\delta_t = \frac{1}{R^*} \sum_{h=1}^H n_{h,t} E_{h,t} [\delta_{t+1} + 1]. \quad (9)$$

It will be convenient to formulate the model in *deviations* from the fundamental value $x_t \equiv \delta_t - \delta_t^*$. We assume that all agents have common and rational beliefs about the fundamental value:

$$E_{h,t}[\delta_{t+1}^*] = E_t[\delta_{t+1}^*] = R^* \delta_t^* - 1. \quad (10)$$

³For example, Boswijk et al. (2007) derive equation (7) from a CARA utility function for a fixed discount rate $R_{t+1} = 1 + r$, as in the standard Gordon model. In Section 5 we directly apply equation (7) to the Campbell-Cochrane benchmark fundamental.

Agents' behavioural beliefs about the next period's PD ratio can be formulated as

$$\begin{aligned} E_{h,t}[\delta_{t+1}] &= E_t[\delta_{t+1}^*] + E_{h,t}[x_{t+1}] \\ &= E_t[\delta_{t+1}^*] + f_h(x_{t-1}, \dots, x_{t-L}), \end{aligned} \quad (11)$$

where $E_{h,t}[x_{t+1}]$ represents the expected deviation of the PD ratio from the fundamental value, expressed as a function $f_h(\cdot)$ of the L last observed deviations.⁴

Under these assumptions about heterogeneous expectations, price deviations from the fundamental can be simplified as

$$\begin{aligned} x_t = \delta_t - \delta_t^* &= \frac{1}{R^*} \sum_{h=1}^H n_{h,t} E_{h,t}[\delta_{t+1} + 1] - \delta_t^* \\ &= \frac{1}{R^*} \sum_{h=1}^H n_{h,t} E_{h,t}[x_{t+1}]. \end{aligned} \quad (12)$$

The standard asset pricing model based on future stock prices and dividends (2) has now been reformulated as a dynamic HAM (12) in which price deviations from the fundamental value depend only on discounted expected future price deviations.

We stress that as the model is formulated in deviations from a fundamental PD ratio, it can be used with different benchmark fundamentals. The crucial assumption that has been made is that agents have common beliefs about fundamental factors, but have heterogeneous beliefs about deviations from fundamental. Also note that the fundamental benchmark with a rational representative agent is nested as a special case of our model when all agent types believe $E_{h,t}[x_{t+1}] = 0$, for all $h = 1, \dots, H$. The model for price fluctuations around the fundamental value (12) holds for any choice of agent types and for any choice of the fundamental value δ_t^* . This setup is convenient to test empirically whether any deviations from a benchmark fundamental are significant.

For the empirical estimation, we consider the simplest form of heterogeneity in belief types which are linear in the last observation:

$$E_{h,t}[x_{t+1}] = \phi_h x_{t-1}. \quad (13)$$

Choosing $H = 2$ types is sufficient to capture an essential difference between agents. Some agents (called fundamentalists) believe in *mean-reversion* of the stock price to its fundamental value and have a parameter $(0 <) \phi_1 < 1$. Other agents (called chartists) believe that the price (in

⁴Note that agents at time t do not observe the contemporaneous price and react to past realised prices only. This assumption is common in the literature and for example used by Hong and Stein (1999) to model momentum traders.

the short run) will move away from the fundamental value and have $\phi_2 > 1$. Chartists expect a continuation of the trend and will be a destabilising factor in the model when their impact becomes large.

The behavioural finance literature has a long tradition of models with fundamentalists and chartists; see Hommes (2006) and LeBaron (2006) for extensive surveys. In a recent overview of the empirical HAM literature, Chen et al. (2012) classify the broad variety of agent-based economic papers and underline that the simple fundamentalist-chartist opposition is often sufficient to explain stylised facts from asset price data that seem ‘puzzles’ in a rational representative agent framework. Aoki (2002) argues with a theoretical model that the behaviour of many different market participants can often be clustered in just two groups. Another recent example is Lof (2014), who applies the VAR approach to a heterogeneous asset pricing model with fundamentalists and contrarians.⁵

Agents use simple rules to predict future prices, but switch to other strategies if their predictions become too far-off from actual prices: investors learn from their mistakes. Our agents learn through reinforcement learning or evolutionary selection based upon the relative performance of their forecasting rule. The fractions of agents belonging to one of the two types are updated with a multinomial logit model as in Brock and Hommes (1997) with *intensity of choice* β^6 :

$$n_{h,t+1} = \frac{e^{\beta U_{h,t}}}{\sum_{j=1}^H e^{\beta U_{j,t}}}. \quad (14)$$

In order to specify the performance measure of belief type h , $U_{h,t}$, we need to consider the profits of agent types. Following Brock and Hommes (1998) and Boswijk et al. (2007), we consider the following profit function in price deviations x_t :

$$\pi_{h,t+1} = (E_{h,t}[x_{t+1}] - R^* x_t)(x_{t+1} - R^* x_t). \quad (15)$$

This expression for realised profits has the intuitive property that it is proportional to agents’ demand (depending on the expectations $E_{h,t}[x_{t+1}] - R^* x_t$) times the realised excess return (depending on realisations $x_{t+1} - R^* x_t$).⁷

⁵Others have proposed models with two types that have more advanced, time-varying adaptive learning beliefs, for example Branch and Evans (2010).

⁶Again, our choice for fluctuating fractions is founded in earlier behavioural finance literature. Chen et al. (2012) state that “*evolving fractions* have been considered to be a [main] cause of many stylised facts” (p. 15, their emphasis).

⁷In a technical appendix accompanying this paper we show that this profit function is consistent with a myopic mean-variance demand function, from which Boswijk et al. (2007) derive the market clearing equation (7) for the Gordon fundamental value. Strictly speaking, the profit function in Boswijk et al. (2007) is proportional to the expression in (15) by a constant factor C which is captured in the estimate $\beta^* = \beta C$.

It will turn out that for our application of the model to quarterly data, realised profits of more than one period in the past should be accounted for in the performance measure. To this end we introduce in the performance measure a *memory* parameter ω :

$$U_{h,t} = (1 - \omega)\pi_{h,t} + \omega U_{h,t-1}, \quad (16)$$

so that the most recent observed profit receives weight $(1 - \omega)$. The relative weight of the j -th lag of realised profits is thus $\omega^j(1 - \omega)$ and decreases in j .

We now have a complete specification of the fluctuating fractions of the $H = 2$ belief types $n_{1,t}$ and $n_{2,t}$. As profits in equation (15) for a certain belief h increase and its performance measure exceeds that of the other belief, more agents will choose this belief, according to the multinomial logit model (14). Thus there is a positive relation between realised profits and fractions of the agents' belief types.

2.3. Econometric form

The HAM in equation (12) with the additional assumptions about expectation formation can be written as an econometric AR(1) model with a time-varying coefficient after adding an error term:

$$\begin{aligned} x_t &= \frac{1}{R^*}(n_{1,t}\phi_1 + n_{2,t}\phi_2)x_{t-1} + \varepsilon_t, \\ &\equiv \varphi_t x_{t-1} + \varepsilon_t. \end{aligned} \quad (17)$$

The error terms are assumed to be independently and identically distributed: $\varepsilon_t \sim IID(0, \sigma^2)$. Economically speaking, the error term ε_t captures exogenous fundamental shocks to underlying dividends and discount rates, which affect prices but are unobserved to investors when making expectations $E_{h,t}[x_{t+1}]$.

The time-varying coefficient φ_t replaces the constant parameter in a regular AR(1)-model and is interpreted here as the *average market sentiment*. As market sentiment rises, prices stay for a longer number of periods away from the fundamental value. Combining equations (14),

(15) and (16), fractions depend nonlinearly on the four parameter and all past realisations:⁸

$$n_{1,t} = f_{n1}(\phi_1, \phi_2, \beta, \omega; x_{t-1}, x_{t-2}, \dots, x_1), \quad (18)$$

$$n_{2,t} = 1 - n_{1,t}. \quad (19)$$

Equations (17), (18) and (19) summarise our heterogeneous agents model for price deviations from any fundamental value. The key idea of the model is *positive expectations feedback*. Initially there is some distribution of fundamentalists and chartists. As shocks are feeded into realised prices, one of the two strategies may receive a higher payoff and attracts more followers given an intensity of choice $\beta > 0$. Because the price is determined by market clearing, these fractions affect the next period price: if overall market sentiment is higher, the next period's price will also be higher. This leads to almost self-fulfilling expectations. For example, when chartists (with $\phi_2 > 1$) dominate, temporary bubbles may arise triggered by fundamental shocks and amplified by trend-following expectations.

The goal of estimating our model is to quantify the effect of positive feedback and switching regimes of market sentiment. This obviously depends on the estimations of the parameters ϕ_1 and ϕ_2 . If these parameters are closer to each other, the effects of switching decreases. We will show that the difference between ϕ_1 and ϕ_2 is statistically significant, and also economically significant, in the sense that the heterogenous agents model produces substantially different market predictions than a representative agent model.

3. Estimation results for a simple behavioural model

The estimation follows a two-step procedure in line with the model description above. First we estimate the fundamental value of the Gordon model based on dividends and a constant risk premium. Second, we estimate the heterogeneous agents model summarised by equations (17), (18) and (19) with nonlinear least squares. In Section 3.3 we introduce a linear representative agent benchmark model to compare our estimation with.

⁸The exact mathematical formula for the fraction of fundamentalists (not important for our line of argument) equals

$$n_{1,t} = (1 + \exp[\beta(\phi_1 - \phi_2) \sum_{j=0}^{t-4} [\omega^j(1 - \omega)x_{t-3-j}(x_{t-1-j} - R^*x_{t-2-j})]])^{-1},$$

In this equation the index $j = 1, 2, \dots$ corresponds to the j -th lag of realised profits that enters the performance measure through memory in (16). At $j = t - 4$, the first observation x_1 is used in determining the fraction n_t , which puts an upper bound on the memory of realised profits. It should be clear that this formula can only be used for $t \geq 4$.

3.1. Estimating the fundamental value of the Gordon model

In this section we specify the fundamental values of stock prices P_t^* and PD ratios δ_t^* using the standard model based on Gordon (1962). The textbook Gordon solution for the fundamental PD ratio under discrete time is constant and equal to:

$$\delta^* = \frac{1+g}{r-g}, \quad (20)$$

where g is the expected growth rate of dividends, and $r = i + RP$ is the sum of the expected risk free rate i and the risk premium on stocks RP , both assumed to be constant. This follows immediately from substituting $g_{t+1} = g$ and $R_{t+1} = 1 + r$ in equation (6).

We follow Boswijk et al. (2007) in using the *dynamic* Gordon model instead of the standard (static) Gordon model. In the dynamic Gordon model, agents can extract possible changes in the future parameters g_{t+1} and $r_{t+1} = i_{t+1} + RP$ from data on dividend growth rates and interest rates available at time t . This approach is more flexible and allows for time variation in the fundamental PD ratio around δ^* . We will show, however, that the time variation in the fundamental PD ratio of the dynamic Gordon model is relatively small. Notice that the dynamic Gordon model presupposes a fixed risk premium.

Agents use a simple AR(1) rule to update their beliefs with the last observation in the risk free rate and growth rate:

$$E_t[r_{t+j}] = r + \rho^j(r_t - r)E_t[g_{t+j}] = g + \tau^j(g_t - g) \quad (21)$$

Boswijk et al. (2007) show, using the approach of Poterba and Summers (1988), that the time-varying fundamental PD ratio is to a first-order Taylor approximation given by:

$$\delta_t^* = \frac{1+g}{r-g} + \frac{\rho(1+g)}{(r-g)(1+r-\rho(1+g))}(r_t - r) + \frac{\tau(1+r)}{(r-g)(1+r-\tau(1+g))}(g_t - g). \quad (22)$$

To estimate the static part δ^* in equation (22), we use updated data on the S&P500 prices and dividends originally provided by Shiller (2005), with $T = 252$ end-of-quarter observations from 1950Q1 until 2012Q4. For an easier interpretation, we will focus on the *yearly* price dividend ratio even though it is based on quarterly observations. We also estimate the parameters using yearly data for comparison. See Table 1 for the results.

Our findings are close to previous estimates on the postwar period, such as a real yearly dividend yield of around 3.5% and a yearly risk premium of around 2.5%. We find some small differences when we estimate the dynamic Gordon model directly on yearly data, because more data points are used in the quarterly estimation. Ignoring deviations in the dividend growth

Table 1: Estimation of the fundamental value.

Frequency	π	d/p	g	r	i	RP	R^*	δ^*
Quarterly	0.91	0.84	0.40	1.24	0.54	0.70	1.008	119.5
Yearly equivalent	3.71	3.40	1.62	5.06	2.19	2.87	1.034	29.9
Yearly	3.71	3.37	1.31	4.69	2.25	2.43	1.033	30.1

Values used for estimating the static Gordon solution $\delta^* = (1+g)/(r-g)$: π is the average inflation rate, d/p is the average dividend yield D_t/P_{t-1} , g is the average dividend growth rate, $r = d/p + g$ equals the risk free rate plus the required risk premium on stocks, i is the average real return on T-notes with a 10-year maturity, $RP = r - i$ is the risk premium and $R^* = (1+r)/(1+g)$ is the expected effective discount rate. We use the CPI index to deflate the nominal variables. All numbers except R^* and δ^* are multiplied by 100. The estimation is done using both quarterly and yearly data from 1951Q1-2012Q4; yearly equivalent estimates based on the quarterly estimation are presented using geometric progression.


and risk free rates, we find a yearly constant value of $\delta^* = 29.9$ for the quarterly estimation.

Next, we estimate the AR(1) rules used by the agents to update the fundamental value. We find quarterly values of $\rho = 0.40$ for the persistence in the risk free rate and $\tau = 0.50$ for the growth rate, and we calculate the fundamental value δ_t^* according to equation (22). Figure 1 plots the fundamental value based on the price-dividend ratio δ_t^* and dividends D_t next to the observed value of the S&P500 index and observed PD ratios. As displayed in the bottom panel of Figure 1, δ_t^* fluctuates in the interval of $[20.3, 40.8]$, but for most periods it stays close to the constant average value of $\delta^* = 29.9$.

The S&P500 clearly exhibits excess volatility, that is, it fluctuates much more than its underlying fundamentals; an important point already made by Shiller (1981). Until 1990, the S&P500 is seen to fluctuate relatively quietly around the value that was to be expected from future dividends. After 1995 this changes: stock prices, most notably of firms in the internet and information technology sector, rose much more than was justified by the dividend pay-outs. At the top of the dot-com bubble in 2000, when the PD ratio reached almost 90 compared to the fundamental value of around 30, the S&P500 was by a factor three overpriced relative to the dividend-based fundamental.

One may argue that we find excess volatility and overpricing in Figure 1 because we do not assume a time-varying risk premium. In Section 5, we follow the more standard approach that at least part of the variation in asset prices is due to variation in risk premia, using the consumption-habit asset pricing model of Campbell and Cochrane (1999) as fundamental benchmark. Irrespective of the underlying prices, all agents in our behavioural model are aware that prices differ from fundamentals, but do not believe that they can use this knowledge to gain higher profits. This behavioural element is supported by survey data, in particular for the period around the turn of the millennium. Both Shiller (2000) and Vissing-Jorgensen (2004) found that the majority of respondent investors in 2000 were aware of the overvaluation of stock

Figure 1: The S&P500 index with its fundamental value, corrected for inflation (top panel), and the realised and fundamental yearly PD ratio (bottom panel).



prices, but did not expect that the mispricing would be corrected within a period of a year.

Another observation is that, in our model with a fixed risk premium, the financial crisis of 2008 is of a quite different nature than the burst of the dot-com bubble. After 2000 prices went down for three years, but stayed above the dynamic Gordon fundamental value. Partly driven by the securitisation activities of large investment banks, dividends rose steadily from 2003 to 2008, which in turn drove prices up again, perhaps more than justified by fundamentals. After the bankruptcy of Lehman Brothers on 15 September 2008, the stock market crashed and prices returned very closely to the fundamental value. Only afterwards, when the market already started to recover, dividends started to fall.

3.2. *Estimation of the heterogeneous agents model*

This section is devoted to the estimation of the heterogeneous agents model using the time series $x_t = \delta_t - \delta_t^*$ of price deviations from the Gordon benchmark. From Figure 1 it is clear that there is quite some structure in x_t . It is highly persistent and interrupted by phases of mean-reversion. We perform nonlinear least squares to estimate the heterogeneous agents model and interpret the asset price fluctuations by different behavioural regimes.

Table 2 shows the estimation results for our four parameters ϕ_1 , ϕ_2 , β and ω , under four different model specifications (A), (B), (C) and (D). For model specification (A), we fix the intensity of choice $\beta = 1$ for reasons discussed below, and estimate the remaining three parameters. For model (B), we estimate all four parameters simultaneously, while model (C) fixes the intensity of choice at a different, high value of $\beta = 10$. Model (D) fixes $\beta = 1$ as in model (A) and has no memory, i.e. the memory parameter $\omega = 0$.

In all estimations, the belief parameters ϕ_1 and ϕ_2 are significantly different from zero, emphasising that the fundamental value is not very informative in the short run. In the model specifications A, B and C with memory, the belief coefficients furthermore show the essential difference between fundamentalism and chartism: $\phi_1 < 1$ and $\phi_2 > 1$. If the regression includes memory the estimated difference $\Delta\phi \equiv \phi_2 - \phi_1$ is significant and around 0.07 to 0.09: the hypothesis $\Delta\phi = 0$ is rejected at the 5% level. This is our main evidence for behavioural heterogeneity in the S&P500 data.

The estimated intensity of choice β is not found to be significant in model version (B), where all four parameters are estimated simultaneously. This is a typical finding in nonlinear “smooth transition” AR models, as changes in β have often very little effect on the fit of the model. For a detailed general discussion of this issue, see Teräsvirta (1994). In Section 4.1, we will use Monte Carlo (MC) simulations and show that the insignificance should not be a concern, as the t -test for β simply lacks power given the size of our sample. These MC simulations show that for a sample size of 252, even if our estimated model (A) with a $\beta = 1 > 0$ would be the true

Table 2: Estimation of the belief coefficients ϕ_1 and ϕ_2 , the intensity of choice β and the memory parameter ω in the heterogeneous agents model $x_t = \frac{1}{R^*}(n_{1,t}\phi_1 + n_{2,t}\phi_2)x_{t-1} + \varepsilon_t$, with $R^* = 1.008$ and the fractions $n_{1,t}$ and $n_{2,t}$ updated according to (18) and (19). All specifications are estimated with nonlinear least squares except for specification (D), which is found by a grid search.

	(A)	(B)	(C)	(D)
ϕ_1	0.936*** (0.027)	0.947*** (0.023)	0.940*** (0.020)	0.981 (-)
ϕ_2	1.026*** (0.021)	1.017*** (0.021)	1.026*** (0.020)	0.981 (-)
β	1	2.443 (7.268)	10	1
ω	0.824*** (0.154)	0.800*** (0.152)	0.852*** (0.190)	0
$\Delta\phi$	0.090** (0.039)	0.070** (0.033)	0.087** (0.029)	0 (-)
T	252	252	252	252
s^2	14.12	14.09	13.87	14.08
R^2	0.952	0.953	0.954	0.951
AIC	2.579	2.583	2.552	2.592
BIC	2.649	2.667	2.623	2.698

*, **, *** denote significance at the 10%, 5% and 1% level, respectively.

Standard errors are denoted in brackets.

The R^2 denotes the proportion variation in δ_t explained by the model.

underlying data generating process, we would not reject the null hypothesis of $\beta = 0$. The MC simulations also show that for sufficiently large sample sizes, this null hypothesis is correctly rejected. Stated differently, the test for no switching ($\beta = 0$) has low power against the HAM model in small samples. This result for nonlinear smooth transition models is similar in spirit as for example the low power of the Dickey-Fuller unit root test for linear near unit-root AR(1) processes, a well-known empirical problem for small samples.

A solution for the insignificant β is to fix it at, e.g. at $\beta = 1$. In the robustness checks in Appendix A, we show that the explanatory power of the model (as measured by the R^2) is much less sensitive to the value of β than to the values of the other parameters (see in particular Figure A.7). Notice also that with $\beta = 10$ fixed at a higher value, as in model (C), the remaining estimated parameters hardly change. Because a non-zero β is necessary to identify different regimes as well as the level of memory ω , the model (A) with a value of 1 is our main focus in the rest of this paper. The model with only the three relevant parameters ϕ_1 , ϕ_2 and ω makes economic sense and is also preferred over model (B) on the basis of Akaike's information criterion (AIC) and the Bayesian information criterion (BIC). Taking $\beta = 0$ reduces the model to a linear AR(1) model and seriously reduces the fit of the model, but any positive value from say $\beta \geq 1$ can be used in the HAM and yields similar estimation results.

The memory parameter ω is strongly significant. This means that shocks that are observed


more than one period ago are also taken into account in the switching between beliefs. While the value around 0.8 in model (A) might seem high, it implies that more than half of the information is extracted from observations in the last year ($1 - 0.824^4 \approx 54\%$). So while memory is important in estimating HAMs on higher frequency data, it does not require unrealistic processing abilities from the agents. The model remains consistent with the behavioural background of bounded rationality. We also estimate model (D) without memory ($\omega = 0$). In this case the two estimated regimes become identical ($\phi_1 = \phi_2$) and the nonlinear least squares estimation is no longer identified; the results for model (D) in Table 2 were found by a grid search. Because of the equal belief parameters, the estimation of model (D) is identical to that of a linear AR(1) model, as we will show in Section 3.3. This underlines our statement that memory is important to explain the regime shifts in the given quarterly dataset of the stock index.

From the estimation of the HAM we can infer the estimated fraction of fundamentalists $n_{1,t}$ over time (Figure 2, top panel). This plot points at a structural break in 1995.⁹ After some initial large shocks in the beginning of the sample, the fraction remained within the interval $[0.4, 0.7]$ for most of the periods. Starting from 1995, however, two successive regimes of trend-following and mean-reversion are evident. Trend-following dominated in the 1990s, amplifying the stock price run-up during the dot-com bubble. The mean-reversion regime continued until prices came just below the Gordon fundamental value in 2009, but fundamentalists remained to dominate the markets for most of the periods in recent years. The inclusion of a memory parameter is needed to distinguish these transitions from the relatively high noise levels.¹⁰

⁹In the estimation of the HAM in Section 5 we allow for a structural break in 1995Q1, a jump in the risk premium, in the Campbell and Cochrane (1999) fundamental benchmark. Allowing for a structural break in the dynamic Gordon model as is done in Appendix A yields similar results and in particular significant behavioural heterogeneity.

¹⁰Boswijk et al. (2007) estimate the model on *annual* stock price deviations from the static Gordon fundamental value between 1871 and 2003. For their lower frequency data, they do find significant behavioural heterogeneity without memory. Their estimated fractions fluctuate heavily over the whole period, sometimes with large swings from close to 0 to close to 1.

Figure 2: Estimated fraction of fundamentalists (top panel) and the corresponding market sentiment (bottom panel) for the HAM (A) under the Gordon model.



The fractions can be translated directly into the estimated market sentiment φ_t over time (Figure 2, bottom panel). The system is locally stable, as for the market sentiment with equal distribution of beliefs it holds that $\frac{\phi_1 + \phi_2}{2R^*} = 0.973 < 1$. However, the plot shows that temporary destabilisation with explosive market sentiment, i.e. $\varphi_t > 1$, is possible when the market is dominated by chartists. This happened for a consecutive number of quarters during 1995-2000, and chartists strongly amplified the magnitude of the dot-com bubble. After the bubble burst, market sentiment remained low for a relatively long period and slowly recovered, at which point the financial crisis hit in 2008. The fact that the model generates these genuinely different and intuitive regimes makes it economically of interest.

Under the assumption of a constant risk premium, the financial crisis was merely a correction back to fundamentals. The relatively small fluctuations in market sentiment since 2001Q1 indicate two points at which some investors moved away from the fundamentalists belief. First, in 2006 and 2007 the recovery of the stock market increased the fraction of chartists to almost 50% and the market sentiment up to 0.970; but already by 2008Q2, *before* the bankruptcy of Lehman Brothers, fundamentalists constituted already almost 100% of the market. According to our HAM around the Gordon fundamental benchmark, the financial crisis therefore has been strongly amplified as a correction back to fundamentals. In 2010 a second temporary upheaval can be observed, but again most agents turned back to mean-reversion.

3.3. A representative agent benchmark model

The variable ϕ_t in equation (17) is a time-varying AR(1)-parameter within the interval $[\frac{\phi_1}{R^*}, \frac{\phi_2}{R^*}]$. As an obvious linear benchmark, we also estimate an AR(1) model with a constant parameter, see Table 3. Notice that this is *not* the rational representative agent benchmark, which would coincide with the fundamental benchmark, but rather a representative agent believe in constant, linear mean reversion. The estimated coefficient of the linear model is 0.973, which is close to 1. This points to the possibility of non-stationarity, implying that the estimated standard errors should be interpreted with care. The Dickey-Fuller test shows that the null hypothesis of a unit root (i.e. $\phi = 1$) can not be rejected. In Section 4.1 we address the result of this test in more detail by Monte Carlo simulations, and find that it is most probably caused by the Dickey-Fuller test having low power in our relatively small sample of $T = 252$.

Table 3: Estimation of the AR(1) model $x_t = \phi x_{t-1} + \varepsilon_t$

Variable	Coefficient	(Std. Err.)
ϕ	0.973***	(0.013)
<i>DF</i>	-2.414	<i>p</i> -value=0.402
<i>T</i>	252	
s^2	14.08	
R^2	0.951	
<i>AIC</i>	2.584	
<i>BIC</i>	2.626	

*, **, *** denote significance at the 10%, 5% and 1% level, respectively.
The R^2 denotes the proportion variation in δ_t explained by the model.

From comparing the R^2 , the heterogeneous agents models with memory are seen to be an improvement over the AR model. The result of a partial F-test shows that the improvement of model (A) over the AR model is significant: the F-statistic is 4.90 and exceeds the 5%-critical value of 3.84. Also, the heterogeneous agents model (A) is preferred over the AR model with a lower *AIC*, although the HAM has a lower *BIC* than the AR model. Admittedly, in absolute terms the explanatory power is roughly 95% for both models. The mild improvement of the explanatory power should be considered in relation with the unpredictability of stocks. The estimations on the shorter horizon suffer from considerable noise levels and different models are therefore inevitably more alike. Campbell and Shiller (1988b) already noted that stock prices become less predictable when they are measured over intervals of less than a year rather than over intervals of several years.

We conclude that the HAM statistically outperforms the AR model. Another interesting feature of our heterogeneous agents model is that it gives an intuitive economic interpretation of medium-run bubbles that is lacking in representative agent benchmark models.¹¹ The next

¹¹It is not straightforward to improve the (extremely) simple AR(1) model in a way other than we propose. For

section will discuss the differences between the linear and the nonlinear switching models in more depth by Monte Carlo simulations. It will become clear that the two models are economically very different.

4. Monte Carlo simulations

To gain understanding about the properties of the HAM and evaluate differences with the simple representative agent AR(1) model, we use the estimated equations as Data-Generating Processes (DGPs) in our Monte Carlo simulations. For these two DGPs we draw shocks from a normal distribution with zero mean and variance equal to the sample variance of the errors s^2 . For the HAM version (A) the DGP is

$$x_t = \phi_t x_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma^2 = 14.12), \quad (23)$$

where $\phi_t \in [\frac{\phi_1}{R^*}, \frac{\phi_2}{R^*}] = [0.928, 1.017]$ is updated according to (17), (18) and (19) with $\phi_1 = 0.936$, $\phi_2 = 1.026$, $\beta = 1$ and $\omega = 0.824$. For the model with a fixed AR(1)-coefficient the DGP is

$$x_t = 0.973x_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma^2 = 14.08). \quad (24)$$

Remember that the simple AR(1) model is a straightforward benchmark model to test for homogeneity. The HAM is essentially an AR(1) model with a time-varying coefficient.

We first evaluate the power and size of the tests we have used in the estimation in Section 3, by running these two DGPs for some number of periods T and considering the outcomes of these tests. In Section 4.2 we use the DGPs to generate time series starting at three observed points in time, in order to illustrate the potential differences between the nonlinear HAM and linear benchmark models in predicting PD ratios.

4.1. Evaluating the power of the main tests

Our tests of interest are the test for homogeneity ($H_0 : \phi_1 = \phi_2$) and for no switching ($H_0 : \beta = 0$) for the HAM, and the Dickey-Fuller test for a unit root ($H_0 : \phi = 1$) for the AR model. Since the true DGP is unknown, we perform these three tests under both DGPs, resulting in six combinations. For example, we estimate a HAM (including a free parameter β) on data generated by a simple AR(1) model, and investigate whether we find significant switching or heterogeneity.

example, in estimations of more general AR(p) model the higher order autocorrelation terms are typically not significant. Another possibility is to allow for time-varying volatility using GARCH-errors (Bollerslev, 1986), a method that is successful in explaining daily stock returns. We estimated an AR(1)-GARCH(1,1) model, in which the dynamics after 1995 can be interpreted with high clustered volatility. This interpretation, however, misses the different regimes that seem present in the data. We find that the HAM also dominates the AR-GARCH model in explaining the data.

In this case it is unlikely that we reject the null hypotheses of no switching and homogeneity, but by pure coincidence of the error realisations, it is possible.

For each of the six combinations of DGP and test, we make B simulation runs of sample size T , make for every run estimations on the simulated time series and check whether the p -value of the particular test is below the nominal significance level α . If the null hypothesis is false, e.g. when homogeneity is tested on HAM-generated data, the proportion of times we reject it measures the *power* of the test. If the null hypothesis is true, e.g. when homogeneity is tested on AR-generated data, the rejection probability measures the *size* of the test. In Tables 4 and 5 the Monte Carlo-estimations of the power are shown, and between brackets the size.

To check the asymptotic properties of the tests, we take a large sample size of $T = 5,000$ and $B = 1,000$ simulation runs (see Table 4). The DF test is asymptotically working correctly: in large samples it successfully identifies the true underlying processes to be stationary, and always rejects the null hypothesis of a unit root. Similarly the power of the homogeneity test is 100% for a large sample size. However, the test for no switching fails to reject the null hypothesis in 12% of the cases, even for this large sample size. This already indicates that rejecting the null of no switching may be difficult even when the true DGP is the HAM.

Table 4: Rejection probabilities: nominal significance level $\alpha = 5\%$, $T = 5,000$, $B = 1,000$, parameters for HAM version (A) in Table 2 and for the AR model as in Table 3.

True model:	DGP \	Test: H_0	unit root	homogeneity	no switching
			$\varphi = 1$	$\phi_1 = \phi_2$	$\beta = 0$
heterogeneous	HAM		100%	100%	88%
representative	AR		100%	(2%)	(0%)

We repeat the Monte Carlo evaluation for the actual sample size of the data ($T = 252$) and $B = 10,000$ simulation runs. Table 5 shows the results. The DF test performs poorly in small samples. While the low small-sample power of the DF test for near unit-root AR processes is well-known in the literature, our results show that the power decreases even further if the true model is a nonlinear HAM. In particular, for our estimated HAM parameters the power is reduced to almost half of the power if the true model an AR process. Therefore, the fact that we failed to reject a unit root in the data is not surprising.

Table 5: Rejection probabilities: nominal significance level $\alpha = 5\%$, $T = 252$, $B = 10,000$, parameters for HAM version (A) in Table 2 and for the AR model as in Table 3.

True model:	DGP \	Test: H_0	unit root	homogeneity	no switching
			$\varphi = 1$	$\phi_1 = \phi_2$	$\beta = 0$
heterogeneous	HAM		6.6%	46.2%	0.0%
representative	AR		11.1%	(4.2%)	(0.0%)

The test for no switching is utterly uninformative in small samples, as it does not reject in any of the replications when the true DGP is a HAM. In other words, the reason that we do not find a significant β in the data is because the corresponding t -test lacks power. The test for homogeneity on the other hand remains to have a relatively high power (46% rejections for the HAM versus 4% for the AR model). This result underlines that the detected heterogeneity in the data is much more likely to be driven by real heterogeneity in the underlying process than by pure chance in a representative agent world.¹²

4.2. *What are the economic effects of different behavioural regimes?*


To understand the economic mechanism that is at play in our heterogeneous agents model, we examine some simulated time series in greater depth. As an example, we simulate the model after three periods: 1997Q3 ($t = 191$), 2002Q2 ($t = 210$) and 2004Q4 ($t = 220$). At each of these points in time the PD ratio was close to 60, relatively far above the fundamental value. Using the DGP given in (23), we calculate different quantiles of the simulated distribution over a rolling horizon. To start the simulations from the initial market situation at $t = 191$, $t = 210$ and $t = 220$, the observed fractions $n_{h,t}$ and performances $U_{h,t}$ of both rules are fed into the simulation. In the left panels of Figure 3, the median prediction and the 5%-, 30%-, 70%- and 95% quantiles for the heterogeneous agents model are presented.¹³

It is striking that the heterogeneous agents model allows for the possibility of a large bubble after 1997Q3, but at the same time generates strong mean-reversion after 2002Q2. After 2004Q4 the median prediction is also decreasing quite quickly to the fundamental value. These differences can be explained by the key mechanism in the model: positive expectations feedback. In 1997Q3, the estimated fraction of fundamentalists is low ($n_{1,t} = 0.04$, see Figure 2) and the market sentiment parameter exceeds 1 ($\phi_t = 1.01$). With partly self-fulfilling expectations and many investors believing in a trend, stock prices typically move further away from fundamentals, which also happened during the dot-com bubble. Our model suggests that the bubble could have been even more pronounced: the top of the 95%-quantile is 106.6 and is reached after 26 quarters in 2004Q1. Note though that bubbles end endogenously: the increasingly high expectations of chartists are bound to overshoot the realised prices, leading to a fundamentalist mean-reverting regime. Our model thus explains the stock market boom in the late 1990s as a temporary bubble triggered by fundamentals and strongly amplified by trend-following behaviour.

¹²Under the Data-Generated Process by the AR(1) model, the tests for homogeneity and for no switching reject the null hypothesis in less than 5% of the simulations, indicating that the size of these tests is not controlled at the nominal level. This effect, probably due to the nonlinearity of the estimated model, does not affect our conclusions.

¹³For simplicity we ignore here small deviations in the fundamental value of the dynamic Gordon model and consider the static Gordon solution $\delta^* = 29.9$.

Figure 3: Realised and simulated PD ratio after 1997Q3 (top panels), after 2002Q2 (middle panels) and after 2004Q4 (bottom panels). The left panels show the simulated distributions for the heterogeneous agents model, feeded with last observed fractions and performances, and the right panels show the simulated distributions for the AR model.



For comparison, we also simulate the representative agent AR model (right panels of Figure 3) starting from the same periods. Note that the linear AR model gives almost identical predictions in 1997Q3, 2002Q2 and 2004Q4, and completely misses the dot-com bubble. The AR simulations are symmetrically distributed around the median and return slowly to the fundamental value with a constant coefficient $\phi < 1$. The HAM is *nonlinear*: it can allow for a bubble in the short run, but also generates a faster return to the fundamental value if fundamentalists become dominant. This rapid decline, as for example in the simulations after 2002Q2, occurs if after some negative shocks the fundamentalists belief keeps attracting more followers, which decreases market sentiment and consequently prices.

The heterogeneous agents model, built upon positive expectations feedback, generates simulated time series of prices that are economically quite different from linear representative agent models. Our model is suitable for making medium-run projections of future prices when predictions of rational representative agent models are unreliable. The financial crisis is within a representative agent world typically perceived as an extreme event. The PD ratio of 27.8 in 2009Q1 is below its 5% quantile of the AR model simulated after 2004Q4, more than four years before the crisis (bottom right panel of Figure 3). Because of the high estimated fraction of fundamentalists, the heterogeneous agents model predicts lower prices after 2004Q4 than the AR model, and its upper 95% confidence interval (above the 5% quantile) does contain the possibility of the large drop in stock prices during the financial crisis.

5. Estimation results using a consumption-habit time-varying risk premium

We have seen that under a constant risk premium stock prices exhibit considerable deviations from their fundamentals and in particular there is large overpricing after 1990. The heterogeneous agents model explains this overpricing as being triggered by fundamental shocks and strongly amplified by a long regime of trend-following behaviour up to the end of 2000. A different explanation from mainstream finance is that the discount rate changed to very low values, such that the same expected future payoffs were valued higher. In Section 3 we allowed the discount rate to vary, but only with predictable variation in the risk free rate; the risk premium was assumed to be constant. In this section we relax the assumption of a constant risk premium in order to study the robustness of our results. We will show that, even after introducing considerable time-variation in the risk premium, significant evidence for behavioural heterogeneity in beliefs remains.

We follow the consumption-habit model for stock price fluctuations by Campbell and Cochrane (1999). This approach is well-known and recently summarised and advocated in Cochrane (2011). The main idea of their model is that investors demand a higher risk premium as consumption decreases during a recession, and conversely become less risk averse when consumption goes up during an economic boom. In order to translate continuing rising consumption levels to a

stationary level of risk aversion, Campbell and Cochrane (1999) define a slow moving “habit” or moving average consumption level, and consider the relative distance between consumption and this habit, called “surplus consumption”. Using this surplus consumption, the model predicts the evolution of price-dividend ratios over time.¹⁴

Our two-step methodology in search for evidence of heterogeneous agents can be applied to any benchmark model for the fundamental value, and here we use the consumption-habit model as the fundamental stock index value. In the first step, we will use the specification of Campbell and Cochrane (1999) for surplus consumption and fit their model on actual PD ratios. Inspired by the line of thought in Cochrane (2011), we allow for one structural break in surplus consumption to capture high asset price values after the 1990s. In the second step, we will estimate the heterogeneous agents model on deviations from the fitted price-dividend ratios, and test whether time-varying risk premia make a significant difference.

5.1. Estimating the Campbell-Cochrane model

Below we give a short summary of the habit consumption asset pricing model. We start by recalling the two-period equation for the PD ratio (4) from Section 2:

$$\delta_t = E_t \left[\frac{1}{R_{t+1}} \frac{D_{t+1}}{D_t} (\delta_{t+1} + 1) \right]. \quad (25)$$

To specify the discount rate R_{t+1} , in this case stochastic, Campbell and Cochrane (1999) stipulate that agents maximise the utility function:

$$U = E_t \sum_{j=1}^{\infty} \kappa^j u(C_{t+j}, H_{t+j}) = E_t \sum_{j=1}^{\infty} \kappa^j \frac{(C_{t+j} - H_{t+j})^{1-\gamma} - 1}{1-\gamma}, \quad (26)$$

where C , H , κ and γ are respectively consumption, the habit level, the subjective time discount factor and the utility curvature parameter. The *surplus consumption ratio* S_t is defined as the relative difference between consumption and the habit level:

$$S_t = \frac{C_t - H_t}{C_t}. \quad (27)$$

From the first-order conditions of (26), the discount rate R_{t+1} equals the inverse of the intertemporal marginal rate of substitution M_{t+1} :

$$\frac{1}{R_{t+1}} = M_{t+1} = \kappa \frac{u_c(C_{t+1}, H_{t+1})}{u_c(C_t, H_t)} = \kappa \left(\frac{S_{t+1}}{S_t} \frac{C_{t+1}}{C_t} \right)^{-\gamma}. \quad (28)$$

¹⁴The consumption-habit model has been challenged by Brunnermeier and Nagel (2008), who analyse microdata on how households allocate their wealth between risky and riskless assets.

Combining equations (25) and (28), Campbell and Cochrane (1999) show that the fundamental PD ratio can be expressed as a function of surplus consumption ratio S_t only, as it is the only state variable.

We now focus on the calculation of the surplus consumption S_t . Campbell and Cochrane (1999) assume that the log surplus consumption ratio $s_t \equiv \log S_t$ evolves as a heteroskedastic AR(1) process:

$$s_{t+1} = (1 - \phi^s)\bar{s} + \phi^s s_t + \lambda(s_t)(c_{t+1} - c_t - \mu_c). \quad (29)$$

The symbols ϕ^s , μ_c and $\bar{s}$ denote parameters for the function of the surplus consumption ratio and $\lambda(s_t)$ is the sensitivity function of s_{t+1} to the deviation of consumption growth from its long run average. Consumption growth is modelled as an i.i.d. lognormal process on $c_t \equiv \log C_t$:

$$c_{t+1} = c_t + \mu_c + v_{t+1}^c, \quad v_{t+1}^c \sim IID(0, \sigma_c^2). \quad (30)$$

Campbell and Cochrane (1999) find the steady-state surplus consumption ratio $\bar{S} \equiv \exp(\bar{s})$:

$$\bar{S} = \sigma_c \sqrt{\frac{\gamma}{1 - \phi}}, \quad (31)$$

where γ is a parameter of utility curvature. They define the sensitivity function as

$$\lambda(s_t) = \begin{cases} (\bar{S})^{-1} \sqrt{1 - 2(s_t - \bar{s})} - 1 & \text{if } s_t \leq s_{max} \\ 0 & \text{if } s_t > s_{max}, \end{cases} \quad (32)$$

where $s_{max} \equiv \log S_{max}$ is the value of s_t at which the upper expression (32) becomes zero:

$$s_{max} = \bar{s} + \frac{1}{2}(1 - \bar{S}^2). \quad (33)$$

Using the expressions above, we can relate observed consumption levels C_t to surplus consumption ratios S_t by four free parameters: μ_c , σ_c , ϕ^s and γ . We use updated data on real per capita U.S. consumption of nondurable goods and services originally provided by Chen and Ludvigson (2009) with $T = 243$ observations from 1952Q1 to 2012Q3.¹⁵ We match the mean and standard deviation of log consumption growth μ_c and σ_c to the data. The parameters ϕ^s and γ are taken identical to Campbell and Cochrane (1999).

Table 6 presents the estimates of the parameters μ_c and σ_c with values of other parameters that are either assumed or implied by the model. Our steady-state surplus consumption ratio $\bar{S}$ is around 0.037 and somewhat lower than the value of Campbell and Cochrane (1999) of 0.057,

¹⁵Note that no data is available for the period 1950Q1-1951Q4, which reduces our sample for the PD ratio by 8 observations. At the time of writing also the last observation 2012Q4 was not available.

because we find smaller variation in consumption growth ($\sigma_c = 0.94$ rather than 1.50, on yearly basis). There are also some small differences between quarterly and yearly surplus consumption ratios; we will use the quarterly values in the further analysis.

Table 6: Parameter values for the habit consumption model.

Parameter	Variable	Value (quarterly)	Value (yearly)
<i>Matched:</i>			
Mean consumption growth (%)	μ_c	0.46	1.84
S.d. of consumption growth (%)	σ_c	0.47	0.94
<i>Assumed:</i>			
Persistence coefficient	ϕ^s	0.97	0.87
Utility curvature	γ	2.00	2.00
<i>Implied:</i>			
Steady-state surplus consumption	$\bar{S}$	0.038	0.037
Maximum surplus consumption	S_{max}	0.063	0.061

Figure 4 shows consumption in logs and relative to consumption at the start of the sample, and the habit level implied by the model under the assumption that the surplus consumption ratio starts at the steady state at the beginning of the sample 1952Q1. The lower panel of Figure 4 shows the surplus consumption ratio. Similar to the estimations of Campbell and Cochrane (1999), the estimated surplus consumption ratio tracks the macroeconomic trends, such as the consumption boom in the 1960s (though with a lag) and the boom in the 1980s. The most recent part of the time series shows the ongoing Great Recession during which consumption dropped down abruptly and persistently to a level very close to the habit.

In the Campbell-Cochrane model the price-dividend ratio is a nearly log-linear function of the surplus consumption ratio. We will estimate the log-linear relationship between the surplus consumption ratio and the PD ratio δ_t as follows:


$$\delta_t = b(S_t)^p + u_t, \quad (34)$$

where $b > 0$ and $p > 0$ are the parameters specifying the log-linear relationship, and u_t is the error term. Instead, Campbell and Cochrane (1999) search for a numerical solution of the PD ratio (25) by plugging in discount rates (28) and using a numerical integrator to evaluate the conditional expectation over the normally distributed consumption shocks v_{t+1}^c . Direct estimation is perhaps less precise but much simpler and leads to a graphical representation similar to Cochrane (2011, p. 1073).

It turns out that the statistical fit of this relationship is rather poor, as seen in Table 7, with an R^2 of 4%. The estimate of p is below 0 (pointing to a negative relationship) but insignificant.¹⁶

¹⁶In fact, the fitted PD ratios of the consumption-habit model without structural break are almost constant

Figure 4: Log consumption per capita of nondurable goods and services and the habit level (top panel), and surplus consumption ratio (bottom panel), under the assumption that the surplus consumption ratio starts at the steady state in 1952Q1.



The main reason for the low fit is that the model does not capture the large stock market boom in the 1990s. Cochrane (2011) circumvents this problem by focusing on PD ratios after 1990 only. As Campbell and Cochrane (1999) note: “Growth in consumption of nondurables and services was surprisingly low in the early 1990s, so our model predicts a fall in price/dividend ratios rather than the increase we see in the data.” They list (exogenous) reasons, such as shifts in corporate financial policy and shifts in consumption due to rising income inequality or demographic effects, why the model based on the time series C_t might underestimate the PD ratio in this period.

Table 7: Estimation of the log-linear relation between the surplus consumption ratio and PD ratio.

Variable	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
b	28.29	(589.3)	303.9***	(78.59)
p	-0.074	(1.18)	0.830***	(0.149)
S_{break}	0		0.093***	(0.007)
T	243		243	
s^2	260.5		58.43	
R^2	0.040		0.785	

*, **, *** denote significance at the 10%, 5% and 1% level, respectively.

To capture these “shifts”, we consider as a fundamental benchmark a consumption-habit model with one structural break in the 1990s. More precisely, we allow for a structural break in surplus consumption ratio in quarter 1995Q1 (corresponding to $t = 173$) by estimating

$$\delta_t = \begin{cases} b(S_t)^p + u_t & \text{if } t < 173 \\ b(S_t + S_{break})^p + u_t & \text{if } t \geq 173, \end{cases} \quad (35)$$


The last two columns of Table 7 present the results of this log-linear regression with one structural break. All three coefficients are significant and of the expected sign. We observe that a very large break in surplus consumption ratio of 0.093 is required, almost one and a half times the maximum value $\bar{S}$. Figure 5 plots the fitted PD ratios resulting from this model, δ_t^{CC} . Given the large break, the model does track some of the variation in PD ratios, also after 1995, as argued by Cochrane (2011). Allowing for a structural break in the consumption-habit fundamental value improves the R^2 from 4% to 79%.

The inclusion of one structural break seems, in principal, reasonable because we have a relatively long time series.¹⁷ For example, Pástor and Stambaugh (2001) establish multiple structural breaks in the equity premium of the CSRP NYSE value-weighted portfolio from 1840

and very close to those of the (dynamic) Gordon model. Estimation of the HAM using these fundamental PD ratios leads therefore to similar results as in Section 3.

¹⁷In Appendix A, we also estimate our model using a Gordon fundamental value as in Section 3 with one structural break in 1995Q1. These estimation results are very similar and support our main result of significant behavioural heterogeneity.

Figure 5: Realised PD ratio and its fitted values based on the Campbell-Cochrane model with a structural break in 1995Q1.



to 1999 and identify the sharpest drop in the 1990s. Lettau and Van Nieuwerburgh (2008) consider various econometric techniques to detect shifts in the mean of the price-dividend ratio and estimate a similar timing of the largest break, which is also consistent with our finding of a large and significant structural break in 1995.¹⁸

Although a structural break improves the fit of the model, for the estimates $b = 303.9$ and $p = 0.830$ that are consistent with the whole sample, large deviations from the fundamental value remain. A lower risk premium due to high surplus consumption ratios cannot fully explain the dot-com bubble as can be seen from Figure 5. The consumption-habit model also fails to explain why the stock market fell so deep in 2008Q3 and recovered so quickly afterwards, because the surplus consumption ratio stayed low for the whole period after 2008Q3. So even after allowing for one structural break, the fit of the Campbell-Cochrane model is far from perfect.

5.2. Estimation of the HAM under the consumption-habit fundamental value

The Campbell-Cochrane model, despite being able to reproduce some general patterns in asset prices by a time-varying risk premium, does not fully explain asset price movements. Even if we disregard the large unexplained structural break in the 1990s, the model fails to account sufficiently for the two biggest events after 1990, namely the dot-com bubble and the financial crisis. We therefore extend the model with behavioural heterogeneity between agents as we did

¹⁸Our results do not depend qualitatively on the exact timing of the structural break. For example, changing the structural break to 1990Q1 ($t = 153$) leads to similar parameter estimates: $b = 269.7^{***}$, $p = 0.773^{***}$, $S_{break} = 0.091^{***}$; $R^2 = 0.647$.

for the Gordon model.¹⁹ In other words, we re-estimate the model summarised by equations (17), (18) and (19) using the time series of PD ratios in deviation from the consumption-habit ‘fundamental value’:

$$x_t^{CC} \equiv \delta_t - \delta_t^{CC}. \quad (36)$$

The results for the HAM estimated on x_t^{CC} are presented in Table 8.²⁰

Table 8: Estimation of the belief coefficients ϕ_1 and ϕ_2 , the intensity of choice β and the memory parameter ω in the HAM $x_t^{CC} = \frac{1}{R^*}(n_{1,t}\phi_1 + n_{2,t}\phi_2)x_{t-1}^{CC} + \varepsilon_t$, with $R^* = 1.008$ and the fractions $n_{1,t}$ and $n_{2,t}$ updated according to (18) and (19), using PD ratios δ_t^{CC} fitted to the Campbell-Cochrane model as in the last two columns of Table 7. All specifications are estimated with nonlinear least squares.

Variable	(A)	(B)	(C)	(D)
ϕ_1	0.699*** (0.071)	0.733*** (0.058)	0.759*** (0.052)	0.796*** (0.100)
ϕ_2	1.017*** (0.068)	0.985*** (0.052)	0.980*** (0.050)	0.933*** (0.098)
β	1	3.768 (6.035)	10	1
ω	0.686*** (0.129)	0.719*** (0.118)	0.817*** (0.016)	0
$\Delta\phi$	0.318*** (0.121)	0.253*** (0.085)	0.221*** (0.078)	0.137 (0.186)
T	243	243	243	243
s^2	13.94	13.72	13.89	14.40
R^2	0.948	0.949	0.948	0.946
AIC	2.664	2.654	2.657	2.687
BIC	2.750	2.755	2.744	2.759

*, **, *** denote significance at the 10%, 5% and 1% level, respectively.

Standard errors are denoted in brackets.

The R^2 denotes the proportion variation in δ_t explained by the model.

The parameter estimates are, perhaps surprisingly, in many ways similar to those in Section 3 under a constant risk premium. Most importantly, we find that the belief parameters ϕ_1 and ϕ_2 are significantly different from each other. Comparing the estimations of model version (A) with three parameters on x_t^{CC} and x_t , we find that the estimate of $\Delta\phi$ is larger - 0.32 instead of 0.08 - and highly significant. Even when taking time-varying risk premia into account, there is significant evidence for different behavioural regimes in the data.

¹⁹By extending the Campbell-Cochrane model with heterogeneous beliefs, we complement theoretical work by Bhamra and Uppal (2014). They generalise a framework in which agents have “catching up with the Joneses” utility functions to heterogeneous priors and heterogeneous preferences, and obtain a closed-form solution.

²⁰For simplicity, we maintain the constant value of the expected effective discount rate at $R^* = 1.008$, obtained from the Gordon model in Section 3.1. The estimation results are robust to reasonable variations in R^* .

There is also significant memory for switching between the two behavioural rules. Setting $\omega = 0$ as is done in model version (D) leads to a lower fit and an insignificant $\Delta\phi$. The estimated intensity of choice is again not significant, because the model fit is very insensitive to the exact value of $\beta > 0$ and its t -test has low power; see Section 5.3 for a more detailed discussion. Model version (A), in which the value is fixed at $\beta = 1$, has a higher *AIC* than model (B), but is preferred over (B) based on the lower *BIC*. We focus on model specification (A), which raises the fit for the consumption-habit model with one structural break from an R^2 of 79% to 95% by allowing for behavioural heterogeneity.²¹


The main differences with the predictions of the HAM estimation in Section 3 lie in the values of the two beliefs, and make sense if we consider the different fundamental benchmark we have used. The value of ϕ_1 is much lower, because the fundamental value is generally closer to the actual PD ratio. For example, after the burst of the dot-com bubble, prices came within a few years back to the fundamental price based on a large surplus consumption ratio and a low risk premium. Also the parameter ϕ_2 is lower and in fact very close to 1, less than one standard deviation away; slightly above 1 in (A) and slightly below 1 in (B) and (C). Hence, while the ‘chartist’ belief is still significantly different from the mean-reverting fundamentalist belief, it is not significantly different from predicting the last observed price, that is, ‘naive expectations’, consistent with a belief that the PD deviation from the fundamental benchmark follows a random walk.

Figure 6 plots the estimated fraction of fundamentalists $n_{1,t}$ and the market sentiment ϕ_t with the time-varying risk premium. The most striking difference is that the fraction of fundamentalists is fluctuating heavily during the entire sample. Before 1995, the proportion of chartists rarely becomes very high, and asset prices stay relatively close to fundamentals. In the 1990s and around the year 2000 market sentiment comes close to 1 (and sometimes exceeds it), destabilising the market and amplifying the dot-com stock boom.

The heterogeneous agents model with a fundamental benchmark with time-varying risk premium based on surplus consumption gives an interesting explanation for the depth of the financial crisis. The model predicts that the fundamental price around 2008 is close to 50, based on the higher surplus consumption ratio after the structural break in 1995 and the lower implied risk premium (see Figure 5). The bankruptcy of Lehman Brothers lead to a large negative shock at a time when a majority of investors was following a chartist strategy, as the fundamentalist belief was unrewarding in the past periods. This behavioural overreaction of investors to the

²¹Note that the HAM using a time-varying risk premium has a slightly lower R^2 (94.8% instead of 95.2%) while having more parameters. The reason lies in the two-step estimation procedure of first the fundamental value and then the heterogeneous agents part on price deviations from this fundamental. Estimating all parameters simultaneously is possible and gives a higher R^2 for the HAM using the consumption-habit time-varying risk premium.

Figure 6: Estimated fraction of fundamentalists (top panel) and the corresponding market sentiment (bottom panel) for the HAM (A) using the Campbell-Cochrane model as fundamental benchmark.



Lehman Brothers shock, with the market sentiment exceeding 1, strongly amplified the financial crisis. Only starting from 2009Q3 price-dividend ratios returned to the fundamentals. During this period dividends increased, so prices started increasing again.

5.3. Evaluating the tests for the model under a consumption-habit fundamental value

As a final robustness check for the model with a time-varying risk premium, we run Monte Carlo simulations using the estimated HAM and a benchmark AR model as Data-Generating Processes.²² We first estimate the linear benchmark model, see Table 9. The constant value φ is as expected halfway the beliefs of the two heterogeneous groups ϕ_1 and ϕ_2 : $0.863 \approx (0.698 + 1.107)/(2 * 1.008)$. The Dickey-Fuller test rejects the null hypothesis of a unit root in PD ratios in deviation from the consumption-habit fundamental value with one structural break.

Using Monte Carlo simulations, we evaluate the Dickey-Fuller test for a unit root ($\varphi = 1$), the test for homogeneity ($\phi_1 = \phi_2$) and the test for no switching ($\beta = 0$) given the estimated parameters in the model with a time-varying risk premium (similar as in Section 4.1 for the

²²It is not straightforward to make out-of-sample predictions for PD ratios with these DGPs as in Section 4.2, because we would then have to make predictions about future consumption growth. This is outside the scope of the current paper.

Table 9: Estimation of the AR(1) model $x_t^{CC} = \phi x_{t-1}^{CC} + \varepsilon_t$, using PD ratios fitted to the Campbell-Cochrane model as in the last two columns of Table 7.

Variable	Coefficient	(Std. Err.)
ϕ	0.863***	(0.032)
DF	-4.284***	p -value<0.01
T	243	
s^2	14.22	
R^2	0.946	
AIC	2.684	
BIC	2.742	

*, **, *** denote significance at the 10%, 5% and 1% level, respectively.

The R^2 denotes the proportion variation in δ_t explained by the model.

model with a constant risk premium). In Table 10 this is done for a large sample size of $T = 5000$ and in Table 11 for the actual sample size of 243. In both cases we generate a large number of B simulated time series under the different DGPs and check the outcome of the three tests for each simulation.

The results are qualitatively the same as in the case with a constant risk premium. The Dickey-Fuller test is asymptotically correct in rejecting the unit root, but has lower power for a smaller sample size of $T = 243$. Because the actual AR(1) parameter is here well below 1 ($\phi = 0.863$) the power is still 98%; much higher than the low power of 11% found in Table 3 (for $\phi = 0.973$). If the actual underlying DGP is a nonlinear HAM, the power of the DF-test decreases to 80%.

Table 10: Rejection probabilities: nominal significance level $\alpha = 5\%$, $T = 5,000$, $B = 1,000$, parameters for HAM version (A) in Table 8 and for the AR model as in Table 9.

True model:	DGP \	Test:	unit root	homogeneity	no switching
		H_0	$\phi = 1$	$\phi_1 = \phi_2$	$\beta = 0$
heterogeneous	HAM		100%	100%	100%
representative	AR		100%	(2%)	(0%)

Table 11: Rejection probabilities: nominal significance level $\alpha = 5\%$, $T = 243$, $B = 10,000$, parameters for HAM version (A) in Table 8 and for the AR model as in Table 9.

True model:	DGP \	Test:	unit root	homogeneity	no switching
		H_0	$\phi = 1$	$\phi_1 = \phi_2$	$\beta = 0$
heterogeneous	HAM		80.3%	84.0%	0.0%
representative	AR		98.2%	(3.2%)	(0.0%)

When estimating a heterogeneous agents model one should be looking for significantly different beliefs (i.e. a significant $\Delta\phi$), and not necessarily for significant switching (a significant β). Both tests work asymptotically correct (having a high power and low size), but the test for no switching hardly ever rejects (0.0%) in small sample and is thus uninformative. Again, this

explains why the estimated intensity of choice β is not significant. The test for homogeneity has large power (84%) even in small sample. The detected heterogeneity in the S&P500 data is therefore most probably not driven by random shocks in a representative agent model (as the attributed p-value is 3.2%).

6. Conclusion

In this paper we investigate the value added of explaining asset pricing movements using a heterogeneous beliefs model. Motivated by the global financial crisis, our focus is on booms and busts that take place in the medium run, and we estimate our model on quarterly data. A convenient feature of our HAM is that it can be formulated around any benchmark fundamental. We use two well-known models as fundamental benchmark values for the price-dividend ratio, the dynamic Gordon model with a constant risk premium and the Campbell-Cochrane consumption-habit model with a time-varying risk premium. Using deviations from these fundamental benchmarks, we find a statistically significant improvement of the heterogeneous agents model over these standard representative agent models, and significant evidence for behavioural heterogeneity.

The global financial crisis displayed events that were sometimes inconceivable in a representative agent world. Our heterogeneous agents approach can shed light upon unexpected swings in asset prices driven by positive expectations feedback. Investors switch between strategies based upon their relative performance. When the group of fundamentalist traders gains momentum, prices return more quickly to the fundamental value. Temporary bubbles, triggered by small fundamental shocks, are strongly amplified when a majority of investors coordinates on chartist beliefs. The endogenous transitions between the regimes driven by relative profitability makes the model attractive in helping to explain booms and busts in asset pricing data.

In our simple heterogeneous agents model with fundamentalists and chartists, we find differences in expectations – measured quarterly and in deviations from the fundamental value – of 10 to 30 percentage points depending on the model that is used for the fundamental value. We show (using Monte Carlo simulations) that this large and significant effect cannot be expected to arise from standard representative agent models. Our model with different behavioural regimes of market sentiment is economically significant in the sense that it gives better predictions than representative agent models in periods before the dot-com bubble and the financial crisis.

A limitation of our methodology is that we only use aggregate data to estimate behavioural rules of individual investors. The simple linear behavioural rules of our model are also found laboratory experimental markets (e.g. Hommes et al., 2005) and describe individual forecasting behaviour quite well. A recent and promising line of research combines stock market data with increasingly reliable survey measures of investors' expectations; e.g. Adam et al. (2013) present

a model where agents are ‘internally rational’ but hold subjective beliefs about stock prices and calibrate it on data from the S&P500 and the UBS Gallup Survey, showing robustness to other survey data. These general patterns in survey data seem to be consistent with the main conclusions in our paper. Our model predicts that during the dot-com bubble around 2000 investors were aware that prices were too high compared to measures of true fundamental value, and surveys as in e.g. Shiller (2000) support this claim. An analysis of behavioural heterogeneity combining aggregate price data and survey measures is an interesting possibility for future research.

We pay specific attention to price dynamics around the financial crisis of 2008. Clearly, the interpretation of this event generally depends on which benchmark fundamental is used to estimate the model. If we consider the Gordon fundamental value, the stock market index has been overpriced since 1995 (see Figure 1). Under this Gordon fundamental, the stock market crash in 2008 is a correction towards the fundamental, while after 2008 a new bubble has formed. If we instead use the Campbell-Cochrane consumption-habit fundamental model with one structural break, the S&P500 seems to be fairly priced in recent years, although systematic deviations remain (see Figure 5). The financial downturn in 2008 is here interpreted as a strong, temporary overreaction to the bankruptcy of Lehman Brothers (see also Figure 6, bottom panel, for the estimated market sentiment). These differences illustrate that policy makers should have a good sense of the underlying fundamental value before they can consider measures to stabilise financial markets. More research is required to investigate which policy measures should be used to contain the amplifying effects of behavioural heterogeneity.

Our analysis adds to evidence that behavioural heterogeneity and strategy switching plays an important role in asset price dynamics. Boundedly rational traders can be expected to survive in financial markets and amplify booms and busts. In particular, we show that agents switched between fundamentalist and chartist beliefs, and strongly reinforced the decline in asset prices after the shock of the Lehman Brothers bankruptcy. Uncertainty due to behavioural heterogeneity has important implications for risk management. Policy makers should therefore not focus exclusively on rational representative agent models, but should take behavioural heterogeneity into account in assessing the systemic risk in financial markets.

References

- Adam, K., Beutel, J., Marcet, A., 2013. Stock price booms and expected capital gains. Mimeo, University of Mannheim.
- Adam, K., Marcet, A., 2011. Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory* 146 (3), 1224–1252.
- Alfarano, S., Lux, T., Wagner, F., 2005. Estimation of agent-based models: The case of an

- asymmetric herding model. *Computational Economics* 26 (1), 19–49.
- Anufriev, M., Hommes, C., 2012. Evolution of market heuristics. *Knowledge Engineering Review* 27 (2), 255–271.
- Aoki, M., 2002. Open models of share markets with two dominant types of participants. *Journal of Economic Behavior & Organization* 49 (2), 199–216.
- Bansal, R., Yaron, A., 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59 (4), 1481–1509.
- Barberis, N., Thaler, R., 2003. A survey of behavioral finance. In: Constantinides, G., Harris, M., Stulz, R. M. (Eds.), *Handbook of the Economics of Finance*. Vol. 1. North-Holland: Elsevier, pp. 1053–1128.
- Bhamra, H. S., Uppal, R., 2014. Asset prices with heterogeneity in preferences and beliefs. *Review of Financial Studies* 27 (2), 519–580.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31 (3), 307–327.
- Boswijk, H. P., Hommes, C. H., Manzan, S., 2007. Behavioral heterogeneity in stock prices. *Journal of Economic Dynamics and Control* 31 (6), 1938–1970.
- Branch, W., Evans, G., 2010. Asset return dynamics and learning. *Review of Financial Studies* 23 (4), 1651–1680.
- Brock, W. A., Hommes, C. H., 1997. A rational route to randomness. *Econometrica* 65 (5), 1059–1096.
- Brock, W. A., Hommes, C. H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22 (8-9), 1235–1274.
- Brunnermeier, M. K., Nagel, S., 2008. Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals’ asset allocation. *American Economic Review* 98 (3), 713–736.
- Campbell, J. Y., Cochrane, J. H., 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107 (2), 205–251.
- Campbell, J. Y., Shiller, R. J., 1988a. The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies* 1 (3), 195–228.
- Campbell, J. Y., Shiller, R. J., 1988b. Stock prices, earnings, and expected dividends. *Journal of Finance* 43 (3), 661–76.

- Campbell, J. Y., Shiller, R. J., 2001. Valuation ratios and the long-run stock market outlook: An update. NBER Working Paper 8221, National Bureau of Economic Research, Inc.
- Chen, S.-H., Chang, C.-L., Du, Y.-R., 2012. Agent-based economic models and econometrics. *The Knowledge Engineering Review* 27 (2), 187–219.
- Chen, X., Ludvigson, S. C., 2009. Land of addicts? An empirical investigation of habit-based asset pricing models. *Journal of Applied Econometrics* 24 (7), 1057–1093.
- Chevalier, J., Ellison, G., 1997. Risk taking by mutual funds as a response to incentives. *Journal of Political Economy* 105 (6), 1167–1200.
- Chiarella, C., He, X.-Z., Zwinkels, R. C., 2014. Heterogeneous expectations in asset pricing: Empirical evidence from the S&P500. *Journal of Economic Behavior and Organization* 105, 1–16.
- Cochrane, J. H., 2001. *Asset Pricing*. Princeton University Press.
- Cochrane, J. H., 2011. Presidential address: Discount rates. *Journal of Finance* 66 (4), 1047–1108.
- de Jong, E., Verschoor, W., Zwinkels, R., 2010. Heterogeneity of agents and exchange rate dynamics: Evidence from the EMS. *Journal of International Money and Finance* 29 (8), 1652–1669.
- Del Guercio, D., Tkac, P. A., 2002. The determinants of the flow of funds of managed portfolios: Mutual funds vs. pension funds. *Journal of Financial and Quantitative Analysis* 37 (4), 523–558.
- DeLong, J., Shleifer, A., Summers, L., Waldmann, R., 1990a. Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance* 45 (2), 379–395.
- DeLong, J., Shleifer, A., Summers, L., Waldmann, R., 1990b. Noise trader risk in financial markets. *Journal of Political Economy* 98 (4), 703–738.
- Frankel, J. A., Froot, K. A., 1987. Using survey data to test standard propositions regarding exchange rate expectations. *American Economic Review* 77 (1), 133–53.
- Frijns, B., Lehnert, T., Zwinkels, R. C., 2010. Behavioral heterogeneity in the option market. *Journal of Economic Dynamics and Control* 34 (11), 2273–2287.
- Gordon, M. J., 1962. *The Investment, Financing, and Valuation of the Corporation*. Irwin.
- Greenwood, R., Shleifer, A., 2013. Expectations of returns and expected returns. NBER Working Paper 18686, National Bureau of Economic Research.

- Heemeijer, P., Hommes, C., Sonnemans, J., Tuinstra, J., 2009. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control* 33 (5), 1052–1072.
- Hirshleifer, D., 2001. Investor psychology and asset pricing. *Journal of Finance* 56 (4), 1533–1597.
- Hommes, C., 2006. Heterogeneous agent models in economics and finance. In: Tesfatsion, L., Judd, K. L. (Eds.), *Handbook of Computational Economics*. Vol. 2. North-Holland: Elsevier, pp. 1109–1186.
- Hommes, C., 2011. The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control* 35 (1), 1–24.
- Hommes, C., Sonnemans, J., Tuinstra, J., van de Velden, H., 2005. Coordination of expectations in asset pricing experiments. *Review of Financial Studies* 18 (3), 955–980.
- Hommes, C., Sonnemans, J., Tuinstra, J., van de Velden, H., 2008. Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization* 67 (1), 116–133.
- Hong, H., Stein, J. C., 1999. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance* 54 (6), 2143–2184.
- Ippolito, R. A., 1989. Efficiency with costly information: A study of mutual fund performance, 1965–1984. *The Quarterly Journal of Economics* 104 (1), 1–23.
- Jouini, E., Napp, C., 2010. Unbiased disagreement in financial markets, waves of pessimism and the risk-return trade-off. *Review of Finance*, 1–27.
- Jouini, E., Napp, C., Viossat, Y., 2012. Evolutionary beliefs and financial markets. *Review of Finance*, 727–766.
- Karceski, J., 2002. Returns-chasing behavior, mutual funds, and beta’s death. *Journal of Financial and Quantitative Analysis* 37 (4), 559–594.
- Kirman, A., Teyssière, G., 2002. Microeconomic models for long memory in the volatility of financial time series. *Studies in Nonlinear Dynamics & Econometrics* 5 (4), 3.
- LeBaron, B., 2006. Agent-based computational finance. In: Tesfatsion, L., Judd, K. L. (Eds.), *Handbook of computational economics*. Vol. 2. North-Holland: Elsevier, pp. 1187–1233.
- Lettau, M., Van Nieuwerburgh, S., 2008. Reconciling the return predictability evidence. *Review of Financial Studies* 21 (4), 1607–1652.
- Lof, M., 2012. Heterogeneity in stock prices: A star model with multivariate transition function. *Journal of Economic Dynamics and Control* 36 (12), 1845–1854.

- Lof, M., 2014. Rational speculators, contrarians and excess volatility. *Management Science* (forthcoming).
- Lux, T., 2009. Stochastic behavioral asset pricing models and the stylized facts. In: Hens IV, T., Schenk-Hoppé, K. R. (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. Vol. 1. North-Holland Amsterdam, p. 161.
- Ofek, E., Richardson, M., 2003. Dotcom mania: The rise and fall of internet stock prices. *Journal of Finance* 58 (3), 1113–1138.
- Pástor, L., Stambaugh, R. F., 2001. The equity premium and structural breaks. *Journal of Finance* 56 (4), 1207–1239.
- Pástor, L., Veronesi, P., 2006. Was there a Nasdaq bubble in the late 1990s? *Journal of Financial Economics* 81 (1), 61–100.
- Pástor, L., Veronesi, P., 2009. Learning in financial markets. *Annual Review of Financial Economics* 1 (1), 361–381.
- Poterba, J. M., Summers, L. H., 1988. Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics* 22 (1), 27–59.
- Shiller, R. J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71 (3), 421–36.
- Shiller, R. J., 1987. Investor behavior in the October 1987 stock market crash: Survey evidence. NBER Working Paper 2446, National Bureau of Economic Research.
- Shiller, R. J., 2000. Measuring bubble expectations and investor confidence. *Journal of Psychology and Financial Markets* 1 (1), 49–60.
- Shiller, R. J., 2005. *Irrational Exuberance*, 2nd Edition. Princeton University Press.
- Sirri, E. R., Tufano, P., 1998. Costly search and mutual fund flows. *Journal of Finance* 53 (5), 1589–1622.
- ter Ellen, S., Zwinkels, R. C., 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 (425), 208–218.
- Vissing-Jorgensen, A., 2004. Perspectives on behavioral finance: Does “irrationality” disappear with wealth? Evidence from expectations and actions. In: *NBER Macroeconomics Annual 2003*, Volume 18. The MIT Press, pp. 139–208.

Westerhoff, F. H., Reitz, S., 2003. Nonlinearities and cyclical behavior: The role of chartists and fundamentalists. *Studies in Nonlinear Dynamics & Econometrics* 7 (4).

Appendix A. Robustness analysis

In this appendix we analyse the robustness of our estimation results of the heterogeneous agents model of the PD ratios in deviations from the Gordon benchmark model.

In Figure A.7, we show that the explanatory power of the model (as measured by the R^2) is much less sensitive to the value of β than to that of the other parameters ϕ_1 , ϕ_2 and ω . For $\beta \geq 1$ the R^2 is essentially flat. This is the reason why we often find a nonsignificant β coefficient. Because a non-zero β is necessary to identify different regimes as well as the level of memory ω , we have fixed β to the (arbitrary) value of 1 in most of our estimations. Fixing β to some larger value yields very similar results.

Next, we try to improve the fundamental value of the Gordon model by adding a structural break in the risk premium starting from 1995Q1, just as we did for the Campbell Cochrane consumption-habit fundamental model in Section 5. We find that before 1995Q1, the risk premium was high and around 0.85% per quarter; from 1995Q1 the risk premium was much lower at around 0.32% per quarter. This leads to higher fundamental prices after 1995: $\delta_{pre95}^* = 25.3$ versus $\delta_{post95Q1}^* = 54.2$. We consider for the Gordon model with one structural break again both the static and the dynamic model. For the dynamic Gordon model, there is a higher impact of deviations from the average values g and r . See Figure A.8.

We reestimate model versions (A) with $\beta = 1$, (B) with all four parameters and (D) with $\beta = 1$ and $\omega = 0$, in Table A.12 for the static Gordon and in Table A.13 for the dynamic Gordon model, both with a break in 1995Q1. The main result is robust for the structural break: if the model includes memory, the two beliefs are significantly different from each other. The estimated value of $\Delta\phi$ is between the estimate for the Gordon model without break (around 0.09) and the estimate for the Campbell Cochrane consumption-habit model (up to 0.30), namely 0.15 to 0.20. The estimates of ϕ_2 are very close to 1 (naive expectations), slightly above 1 for the static and slightly below 1 for the dynamic Gordon model. In general, our results are robust and our conclusions do not depend on our choice of the benchmark fundamental value.

Figure A.7: Sensitivity of R^2 in the HAM. The dot denotes the choice of parameters as in model (A) giving $R^2 = 0.952$. In each of the diagrams, one of the four parameters (ϕ_1 , ϕ_2 , ω or β) is altered while retaining the values of the other three. The dotted line indicates the fit of the AR(1) model ($R^2 = 0.951$).




Figure A.8: Realised PD ratio and its fundamental values based on the static and dynamic Gordon model with one structural break in 1995Q1.




Table A.12: Estimation of the belief coefficients ϕ_1 and ϕ_2 , the intensity of choice β and the memory parameter ω in the HAM $x_t^{SG+break} = \frac{1}{R^*}(n_{1,t}\phi_1 + n_{2,t}\phi_2)x_{t-1}^{SG+break} + \varepsilon_t$, with $R^* = 1.008$ and the fractions $n_{1,t}$ and $n_{2,t}$ updated according to (18) and (19), under a static Gordon model with one structural break in 1995Q1. All specifications are estimated with nonlinear least squares.

	(A)	(B)	(D)
ϕ_1	0.817*** (0.053)	0.827*** (0.048)	0.930*** (0.150)
ϕ_2	1.019*** (0.036)	1.010*** (0.035)	0.945*** (0.164)
β	1	1.957 (3.106)	1
ω	0.640** (0.276)	0.613** (0.271)	0
$\Delta\phi$	0.203*** (0.073)	0.184** (0.067)	0.016 (0.309)
T	252	252	252
s^2	10.90	10.88	11.39
R^2	0.955	0.955	0.953
AIC	2.525	2.531	2.556
BIC	2.581	2.601	2.598

The R^2 denotes the proportion variation in δ_t explained by the model.

Table A.13: Estimation of the belief coefficients ϕ_1 and ϕ_2 , the intensity of choice β and the memory parameter ω in the HAM $x_t^{DG+break} = \frac{1}{R^*}(n_{1,t}\phi_1 + n_{2,t}\phi_2)x_{t-1}^{DG+break} + \epsilon_t$, with $R^* = 1.008$ and the fractions $n_{1,t}$ and $n_{2,t}$ updated according to (18) and (19), under a dynamic Gordon model with one structural break in 1995Q1. All specifications are estimated with nonlinear least squares except for specification (D), which is found by a grid search.

	(A)	(B)	(D)
ϕ_1	0.816*** (0.054)	0.829*** (0.052)	0.914 (-)
ϕ_2	0.989*** (0.045)	0.977*** (0.043)	0.914 (-)
β	1	2.359 (5.080)	1
ω	0.752*** (0.176)	0.759*** (0.123)	-
$\Delta\phi$	0.174** (0.082)	0.148** (0.075)	0 (-)
T	252	252	252
s^2	17.32	17.26	-
R^2	0.945	0.945	0.944
AIC	2.728	2.731	2.738
BIC	2.813	2.829	2.808

The R^2 denotes the proportion variation in δ_t explained by the model.