

TI 2013-204/II
Tinbergen Institute Discussion Paper



Behaviorally Rational Expectations and almost Self-Fulfilling Equilibria

Cars Hommes

CeNDEF, 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

Behaviorally Rational Expectations and Almost Self-Fulfilling Equilibria

Cars Hommes^a

^a*CeNDEF, University of Amsterdam and Tinbergen Institute
Valckenierstraat 65-67, 1018XE Amsterdam, NL*

August 19, 2013

Abstract

Rational expectations assumes perfect, model consistency between beliefs and market realizations. Here we discuss behaviorally rational expectations, characterized by an observable, parsimonious and intuitive form of consistency between beliefs and realizations. We discuss three case-studies. Firstly, a New Keynesian macro model with a representative agent learning an optimal, but misspecified, AR(1) rule to forecast inflation consistent with observed sample mean and first-order autocorrelations. Secondly, an asset pricing model with heterogeneous expectations and agents switching between a mean-reverting fundamental rule and a trend-following rule, based upon their past performance. The third example concerns learning-to-forecast laboratory experiments, where under positive feedback individuals coordinate expectations on non-rational, almost self-fulfilling equilibria with persistent price fluctuations very different from rational equilibria.

JEL Classification: D84, D83, E32, C92

Keywords: Expectation feedback, self-fulfilling beliefs, heuristic switching model, experimental economics

1 Introduction

Traditional economics is built on the rationality paradigm. All agents are assumed to be perfectly rational, using all available information and rationality is common knowledge. In a perfect rational world, the population of agents is represented by a perfectly rational representative agent, markets are in equilibrium and as a result markets are efficient and prices reflect economic fundamentals. Agents have rational expectations (RE), that is, expectations are model consistent, so that beliefs and market realizations, on average, coincide optimally. The rational expectation hypothesis (*REH*, Muth, 1961, Lucas, 1972) is the cornerstone of new classical economic and financial analysis.

Behavioral economics and finance are gaining ground however and many systematic deviations from rationality –behavioral biases– have been documented, e.g., hyperbolic discounting, money illusion, ambiguity effect, loss aversion, status quo bias, survivorship bias, etc. Some possible biases in expectation formation that have been discussed in the literature include belief bias, anchoring, bandwagon effect, (Bayesian) conservatism, exaggerated expectations, optimistic or pessimistic beliefs and overconfidence. In this paper we discuss a behavioral theory of expectations, replacing perfect model consistency by observable, more intuitive and parsimonious consistency requirements between beliefs and realizations.

Much work has been done on boundedly rational expectations and adaptive learning (Sargent, 1993). Boundedly rational agents act as econometricians or statisticians using an econometric forecasting model and updating the parameters over time as additional observations become available; see, e.g., Evans and Honkapohja (2001, 2013) and Bullard (2006) for extensive surveys and references. Part of this literature involves studying under which conditions learning will converge to the rational expectations equilibrium (REE). Another part of the literature studies non-rational learning equilibria, explaining excess volatility, as e.g. in Bullard (1994), Hommes and Sorger (1998), Hommes and Rosser (2001) and Adam and Marcet (2011).

This paper puts forth the notion of *behaviorally rational expectations*. While adaptive learning has stressed a statistical approach to expectations and forecasts, we stress simplicity, parsimony and an intuitive behavioral interpretation. Agents base their expectations on simple forecasting heuristics, with the class of forecasting rules disciplined by behavioral consistency between beliefs and market realizations. Hence, we back off from perfect rational expectations and model consistency, but impose a reasonable degree of observational consistencies between beliefs and realizations. Aggregate behavior is then characterized by *almost self-fulfilling* behavioral learning equilibria. A rational expectations equilibrium may arise as a special case in which the equilibrium is perfectly self-fulfilling, but typically behavioral learning equilibria exhibit excess volatility and may persistently deviate from the rational benchmark. Our approach is similar in spirit to the concepts of *natural expectations* (Fuster et al., 2010; Beshears et al., 2013) and *smart heuristics* and the *adaptive toolbox* (Gigerenzer et al., 1999; Gigerenzer and Selten, 2001). A key difference however is that we impose plausible consistency requirements between beliefs and realizations.

We discuss the concept of behavioral rationality through three case studies; more examples may be found in Hommes (2013). The first is in macroeconomics, a New Keynesian Philips curve (NKPC) with a representative agent learning the simplest, but *misspecified*, univariate AR(1) rule to forecast future inflation in an economy that is too complex to fully understand (Hommes and Zhu, 2013). The parameters of the AR(1) rule are pinned down by simple and intuitive consistency requirements: the mean and the first-order autocorrelation of the AR(1) forecasting rule coincide with the realizations. The representative agent does not fully understand the complex economy, but makes a mistake believing that the true economy follows an AR(1) process. This mistake becomes self-fulfilling (Grandmont, 1998) –in terms of the sample average and first order sample autocorrelation– when the agent learns the optimal parameters of the simple univariate rule.

The second case study is an asset pricing model with heterogeneous expectations

and agents choosing between a mean-reverting fundamental rule and a trend-following rule (Boswijk et al., 2007; Hommes and in't Veld, 2013). Switching between strategies is endogenous, based upon the relative past performance of the rules. Hence, agents believe that the market may switch between a mean-reverting and an explosive phase and this belief becomes self-fulfilling when switching behavior is driven by profitability.

The third example concerns learning-to-forecast laboratory experiments in positive and negative feedback systems (Bao et al., 2012). Negative feedback systems are rather stable and individuals quite easily learn to coordinate on the perfectly self-fulfilling REE. In contrast, under positive feedback *almost self-fulfilling equilibria* arise in the laboratory markets.

A common finding is that in positive feedback markets aggregate behavior is not well described by perfectly self-fulfilling rational equilibrium. Instead, individuals coordinate their expectations on almost self-fulfilling equilibria, very different from exact self-fulfilling equilibria, with excess volatility and persistent price fluctuations deviating significantly from perfect rationality.

The paper is organized as follows. Section 2 discusses the New Keynesian Philips curve with a representative behaviorally rational agent. Section 3 discusses an asset pricing model with fundamentalists versus trend-extrapolators. Section 4 discusses individual and aggregate behavior in laboratory experiments under positive and negative feedback. Finally, Section 5 concludes.

2 A New Keynesian Philips curve

We consider the New Keynesian Philips curve (NKPC) with inflation driven by an exogenous AR(1) process y_t for the firm's real marginal cost or the output gap. Inflation and the output gap (real marginal cost) evolve according to (e.g. Woodford,

2003; Lansing, 2009)

$$\begin{cases} \pi_t = \delta\pi_{t+1}^e + \gamma y_t + u_t, \\ y_t = a + \rho y_{t-1} + \varepsilon_t, \end{cases} \quad (1)$$

where π_t is the inflation at time t , π_{t+1}^e is the subjective expected inflation at date $t+1$ and y_t is the output gap or real marginal cost, $\delta \in [0, 1)$ is the representative agent's subjective time discount factor, $\gamma > 0$ is related to the degree of price stickiness in the economy and $\rho \in [0, 1)$ describes the persistence of the AR(1) driving process. u_t and ε_t are i.i.d. stochastic disturbances with zero mean and finite absolute moments with variances σ_u^2 and σ_ε^2 , respectively. We refer to u_t as a markup shock, which is often motivated by the presence of an uncertain, variable tax rate and to ε_t as a demand shock, that is uncorrelated with the markup shock.

Agents are boundedly rational and do not know the exact form of the actual law of motion (1). As in Hommes and Zhu (2013), we assume that agents do not recognize that inflation π_t is driven by an exogenous stochastic process y_t for marginal costs or output gap. Although agents may observe y_t they do not realize or they do not believe that the state of the economy is determined by the exogenous process y_t . Instead agents only use past observations $\pi_{t-1}, \pi_{t-2}, \dots$, and a simple linear model to forecast the future state of the economy (i.e. inflation), whose structure and complexity they do not fully understand. We focus on the simplest case where agents' perceived law of motion (PLM) is a simple, parsimonious univariate AR(1) process. Hommes and Zhu (2013) follow the idea of a *self-fulfilling mistake* (Grandmont, 1998) imposing consistency requirements on the parameters of the AR(1) rule, as explained below, so that the mistake becomes self-fulfilling.

If agents believe that inflation π_t follows a univariate AR(1) process, and try to learn its parameters α and β , the implied actual law of motion (ALM) becomes

$$\begin{cases} \pi_t = \delta[\alpha_t + \beta_t^2(\pi_{t-1} - \alpha_t)] + \gamma y_t + u_t, \\ y_t = a + \rho y_{t-1} + \varepsilon_t. \end{cases} \quad (2)$$

We assume that agents use sample autocorrelation learning (SAC-learning; Hommes

and Sorger, 1998) to learn the parameters α and β . That is, for any finite set of observations $\{\pi_0, \pi_1, \dots, x_t\}$, the sample average is given by

$$\alpha_t = \frac{1}{t+1} \sum_{i=0}^t \pi_i, \quad (3)$$

and the first-order sample autocorrelation coefficient is given by

$$\beta_t = \frac{\sum_{i=0}^{t-1} (\pi_i - \alpha_t)(\pi_{i+1} - \alpha_t)}{\sum_{i=0}^t (\pi_i - \alpha_t)^2}. \quad (4)$$

Hence α_t and β_t are updated over time as new information arrives. SAC-learning has a simple behavioral interpretation, since the sample average and first order autocorrelation coefficient –the persistence of a time series– may be “guestimated” from an observed time series.

A first-order *Stochastic Consistent Expectations Equilibrium* (SCEE) (α^*, β^*) is an equilibrium for which the mean and first-order autocorrelation of the implied actual law of motion (2) exactly coincide with α^* and β^* respectively. Hommes and Zhu (2013) prove the following proposition:

Proposition. In the case that $0 < \rho < 1$ and $0 \leq \delta < 1$, there exists at least one nonzero first-order stochastic consistent expectations equilibrium (SCEE) (α^*, β^*) for the New Keynesian Philips curve (2) with $\alpha^* = \frac{\gamma^a}{(1-\delta)(1-\rho)} = \bar{\pi}^*$.

The sample mean $\alpha^* = \frac{\gamma^a}{(1-\delta)(1-\rho)} = \bar{\pi}^*$ exactly coincides with the RE sample mean, and therefore along a SCEE inflation is unbiased in the long run. The inflation persistence, as measured by the first-order autocorrelation β^* is typically much larger than the RE persistence (which equals ρ , the autocorrelation of the exogenous driving force). Moreover, Hommes and Zhu (2013) show that for the New Keynesian Philips curve (1) multiple SCEE may coexist. Coexistence of multiple SCEE is illustrated in Figure 1. The first-order autocorrelation coefficient $F(\beta)$ of the implied ALM (2) is given by

$$F(\beta) = \delta\beta^2 + \frac{\gamma^2\rho(1 - \delta^2\beta^4)}{\gamma^2(\delta\beta^2\rho + 1) + (1 - \rho^2)(1 - \delta\beta^2\rho) \cdot \frac{\sigma_u^2}{\sigma_\varepsilon^2}}. \quad (5)$$

Figure 1a illustrates that, for empirically plausible parameter values, $F(\beta)$ has three fixed points $\beta_1^* \approx 0.3066$, $\beta_2^* \approx 0.7417$ and $\beta_3^* \approx 0.9961$. Hence, we have co-existence of three first-order SCEE (α^*, β^*) . Figures 1b and 1c illustrate the time series of inflation along the coexisting SCEE. Inflation has low persistence along the SCEE (α^*, β_1^*) , but very high persistence along the SCEE (α^*, β_3^*) . The time series of inflation along the high persistence SCEE in Figure 1c has in fact similar persistence characteristics and amplitude of fluctuation as in empirical U.S. inflation data. Furthermore, Figure 1c illustrates that inflation in the high persistence SCEE has much stronger persistence than REE inflation, where the first-order autocorrelation coefficient of REE inflation equals $\rho = 0.9$, significantly less than $\beta_3^* = 0.9961$.

Figure 2 illustrates the inflation dynamics in the Keynesian Philips curve under constant gain SAC-learning¹. In the case of multiple equilibria, time-varying parameter estimates make it possible that a sequence of shocks moves the economy from one equilibrium to another. As shown in Branch and Evans (2010), a real-time constant gain learning formulation in an asset pricing model is capable of reproducing regime-switching asset returns and volatilities. Figure 2 illustrates these switches between different basins of attraction, where other parameters are taken as in Figure 1. The inflation π_t clearly exhibits regime switching, with the autocorrelation coefficient β_t switching between phases of high persistence with near unit root behaviour and phases of low persistence with close to 0 autocorrelation and the variance estimator

$$R_t = \frac{1}{t+1} \sum_{i=0}^t (\pi_t - \alpha_t)^2, \quad (6)$$

switching between phases of high and low volatility.

This simple example illustrates the occurrence of almost self-fulfilling equilibria, with behaviorally rational expectations, in a representative agent framework. The

¹Hommes and Zhu (2013) derive the recursive constant gain version of SAC-learning. Under SAC-learning, the “gain”, i.e., the change in the updating of the parameter values α_t and β_t converges to 0 because of the $1/t$ terms averaging over the entire past. In the constant gain version of SAC-learning the term $1/t$ is replaced by a parameter κ representing the constant gain.




Figure 1: (a) The first-order autocorrelation β^* of the SCEE correspond to the three intersection points of $F(\beta)$ in (5) (bold curve) with the perceived first-order autocorrelation β (dotted line); (b) time series of inflation in low-persistence SCEE $(\alpha^*, \beta_1^*) = (0.03, 0.3066)$; (c) times series of inflation in high-persistence SCEE $(\alpha^*, \beta_3^*) = (0.03, 0.9961)$ (bold curve) and time series of REE inflation (dotted curve). Empirically plausible parameter values are: $\delta = 0.99$, $\gamma = 0.075$, $a = 0.0004$, $\rho = 0.9$, $\sigma_\varepsilon = 0.01$ [$\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$], and $\sigma_u = 0.003162$ [$u_t \sim N(0, \sigma_u^2)$], so that $\frac{\sigma_u^2}{\sigma_\varepsilon^2} = 0.1$.




Figure 2: New Keynesian Philips curve under constant gain learning with gain parameter $\kappa = 0.008$. Top panel: inflation π_t and sample average α_t ; Middle panel: first-order autocorrelation coefficient β_t , switching between high and low persistence; Bottom panel: variance estimator R_t switching between high and low volatility.

representative agent makes a mistake and does not fully recognize the complex structure of the economy. Instead, the agent uses the simple univariate AR(1) forecasting rule. By learning, the mistake becomes self-fulfilling in terms of the two most important observable statistics, the sample average and first-order autocorrelation, which may also be “guestimated” from observed time series.

3 A behavioral asset pricing model

We consider a stylized yet microfounded asset pricing model with heterogeneous beliefs and the estimation of a 2-type model on S&P500 data. A detailed derivation may be found in Brock and Hommes (1998). Investors can choose between a risk free asset paying a fixed return r and a risky asset (say a stock) paying stochastic dividends. We assume that investors have perfect knowledge of the exogenous cash flow process, and thus know the ‘fundamental value’ of the risky asset, but differ in their beliefs about the price of the asset in the short run. Denote Y_t as the dividend payoff, P_t as the asset price and $z_{h,t}$ as the number of shares bought by investor i .

Wealth in period $t + 1$ for trader type h is given by

$$W_{h,t+1} = (1 + r)W_{h,t} + (P_{t+1} + Y_{t+1} - (1 + r)P_t)z_{h,t}, \quad (7)$$

Let $E_{h,t}$ and $V_{h,t}$ denote the “beliefs” or forecasts of trader type h about conditional expectation and conditional variance. Agents are assumed to be myopic mean-variance maximizers so that the demand $z_{h,t}$ of type h is

$$z_{h,t} = \frac{E_{h,t}[P_{t+1} + Y_{t+1} - (1 + r)P_t]}{aV_{h,t}[P_{t+1} + Y_{t+1} - (1 + r)P_t]} = \frac{E_{h,t}[P_{t+1} + Y_{t+1} - (1 + r)P_t]}{a\sigma^2}, \quad (8)$$

where a is the risk aversion and, for simplicity, the conditional variance $V_{h,t} = \sigma^2$ is assumed to be equal and constant for all types. In the case of zero supply of outside shares the market clearing price is given by

$$P_t = \frac{1}{1 + r} \sum_{h=1}^H n_{h,t} E_{h,t}[P_{t+1} + Y_{t+1}], \quad (9)$$

where $n_{h,t}$ denotes the time-varying fraction of trader type h and $E_{h,t}[P_{t+1} + Y_{t+1}]$ denotes the beliefs about the future price and the future dividend by investor type h .

The dividend process follows a geometric random walk with drift:

$$\log Y_{t+1} = \mu + \log Y_t + \nu_{t+1}, \quad \nu_{t+1} \sim IID(0, \sigma_\nu^2). \quad (10)$$

Investors have correct beliefs about dividends and estimate the constant growth rate $g \equiv e^{\mu + \frac{1}{2}\sigma_\nu^2}$ by averaging over $\log(\frac{Y_{t+1}}{Y_t})$. Therefore the agents have model-consistent

beliefs about the exogenous dividend process: $E_{i,t}[Y_{t+1}] = (1 + g)Y_t$. This assumption has the convenient feature that the model can be written in deviations from a benchmark fundamental.

In the special case where *all* agents have rational expectations about prices, the price equals its RE *fundamental value* given by the discounted sum of all future expected dividends²:

$$P_t^* = \frac{1 + g}{r - g} Y_t. \quad (11)$$

Hence, under RE the price-to-dividend ratio is constant and given by

$$\frac{P_t^*}{Y_t} = \frac{1 + g}{r - g} \equiv \delta^*. \quad (12)$$

Fig. 3 illustrates the S&P500 stock market index, the price-to-dividend ratio $\delta_t \equiv P_t/Y_t$ and the fundamental value P_t^* . The S&P500 index clearly exhibits *excess volatility*, a point already emphasized in the seminal paper of Shiller (1981). Boswijk et al. (2007) estimated a 2-type model using yearly S&P500 data from 1871-2003. More recently Hommes and Veld (2013) updated the estimation of the 2-type model using quarterly data 1950Q1-2012Q3, using the updated data set of Shiller (2006).

In deviations from the fundamental value $x_t \equiv \delta_t - \delta^*$, the 2-type model can be rewritten as:

$$x_t = \frac{1}{R^*} (n_{1,t} E_{1,t}[x_{t+1}] + n_{2,t} E_{2,t}[x_{t+1}]), \quad R^* \equiv \frac{1 + r}{1 + g}. \quad (13)$$

Notice that the asset pricing model has *positive expectations feedback*, that is, realized price deviation increases (decreases) when (average) expected deviation increases (decreases). The simplest form of heterogeneity occurs when belief types are linear in the last observation:

$$E_{h,t}[x_{t+1}] = \phi_h x_{t-1}. \quad (14)$$

²This solution is known as the Gordon model (Gordon, 1962).




Figure 3: Top panel: Time series of S&P500, its fundamental value P_t^* in (11) and the one-step ahead forecast of the 2-type model; Second panel: price-to-dividend ratio δ_t , its (constant) fundamental δ^* in (12) the one-step ahead forecast of the 2-type model; Third panel: estimated fraction $n_{1,t}$ of fundamentalists, and Fourth panel: the corresponding time varying market sentiment ϕ_t in (21)

Two types, $h = 1, 2$, are sufficient to capture the essential difference in agents' behavior: fundamentalists believe the price will return to its fundamental value ($0 \leq \phi_1 < 1$) and chartists believe that the price (in the short run) will move away from the fundamental value ($\phi_2 > 1$).

The fractions of the two types are updated with a multinomial logit model as in Brock and Hommes (1997), with *intensity of choice* β :

$$n_{h,t+1} = \frac{e^{\beta U_{h,t}}}{\sum_{j=1}^H e^{\beta U_{j,t}}}. \quad (15)$$

The performance measure $U_{h,t}$ is a weighted average of past profits $\pi_{h,t}$ and past fitness $U_{h,t-1}$, with memory parameter ω :

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

with profits, up to a constant factor, given by

$$\pi_{h,t} = z_{h,t-1}R_t = \phi_h(x_{t-2} - Rx_{t-1})(x_t - Rx_{t-1}). \quad (17)$$

The econometric form of the endogenous strategy switching model is an AR(1)- model with a time-varying coefficient:

$$R^*x_t = n_t\phi_1x_{t-1} + (1 - n_t)\phi_2x_{t-1} + \epsilon_t \quad R^* = \frac{1+r}{1+g}, \quad (18)$$

where ϵ_t is an i.i.d. error term. Combining equations (15), (??) and (16), fractions depend nonlinearly on past realisations:

$$n_{1,t} = (1 + \exp[\beta(\phi_1 - \phi_2) \sum_{j=0}^{t-4} [\omega^j (1 - \omega)x_{t-3-j}(x_{t-1-j} - R^*x_{t-2-j})]])^{-1}, \quad (19)$$

$$n_{2,t} = 1 - n_{1,t}. \quad (20)$$

The estimated parameter values in Hommes and in't Veld (2013) are:³

³In this estimation the intensity of choice $\beta = 1$ has been fixed. Estimating the β -parameter is hard and yields non-significant results, probably due to the relatively small sample size. At the same time the coefficients ϕ_1 and ϕ_2 are significantly different from each other and therefore β is non-zero. See Hommes and Veld (2013) for bootstrap analyses and an extensive discussion.

- $\phi_1 = 0.953$: type 1 therefore are *fundamentalists*, expecting mean reversion of the price towards its fundamental value;
- $\phi_2 = 1.035$: type 2 are *trend extrapolators*, expecting the price deviation from fundamental to increase by 3.5% per quarter;
- $\omega = 0.816$: implying almost 20% weight is given to the most recent profit observation and about 80% to past profitability.

Define the *market sentiment* as

$$\phi_t = \frac{n_t \phi_1 + (1 - n_t) \phi_2}{R^*} \quad (21)$$

Figure 3 also shows time series of estimated fractions of fundamentalists and the market sentiment. The fraction of fundamentalists varies considerably but gradually (due to memory) over time, with values between 0.25 and 0.9 until the 1990s, and more extreme values ranging from close to 0 to almost 1 after the dot com bubble. The switching model offers an intuitive explanation of the dot com bubble as being triggered by economic fundamentals (good news about a new internet technology) subsequently strongly amplified by trend-following behavior. Estimates of the market sentiment ϕ_t vary between 0.96 and 1 until the 1990s, showing near-unit root behavior. During the dot com bubble the market sentiment ϕ_t exceeds 1 for several quarters and therefore the market is temporarily explosive. During the financial crisis the market is mainly dominated by fundamentalists indicating that the financial crisis has been reenforced by fundamentalists who expected a correction of asset prices back to fundamentals.

In this behavioral asset pricing model with heterogeneous beliefs, agents switch between a mean-reversion and a trend-following strategy based upon realized profitability. Strategy switching driven by profitability leads to an almost self-fulfilling equilibrium with bubbles and crashes triggered by shocks (“news”) to economic fundamentals amplified by endogenous switching between trend-following and fundamentalist’s strategies.

4 Laboratory Experiments

In this section we consider Learning-to-Forecast Experiments (LtFE) with human subjects; a survey of these experiments is given in Hommes (2011). Subjects have to forecast a price, whose realization depends endogenously on their average forecast. The main goal of these experiments is to study how individual expectations are formed, how these interact and which structure emerges at the aggregate level. Will agents coordinate on a common forecast and will the price converge to the rational expectations benchmark or will other, almost self-fulfilling behaviorally rational equilibria arise?

As already noted in Muth's classical paper introducing rational expectations, a crucial feature for aggregation of individual expectations, is whether the deviations of individual expectations from the rational forecast are *correlated* or not. To quote Muth (1961, p.321, emphasis added):

“Allowing for cross-sectional differences in expectations is a simple matter, because their aggregate affect is negligible as long as the deviation from the rational forecast for an individual firm is not strongly correlated with those of the others. Modifications are necessary only if the correlation of the errors is large and depends systematically on other explanatory variables”.

Laboratory experiments are well suited to study correlation of individual expectations in a controlled environment. It turns out that the type of expectations feedback, positive or negative, is crucial. In the case of positive (negative) feedback, an increase (decrease) of the average forecast, causes the realized market price to rise (fall). Positive feedback seems particularly relevant in speculative asset markets. If many agents expect the price of an asset to rise they will start buying the asset, aggregate demand will increase and so, by the law of supply and demand, the asset price will increase. High price expectations then become self-fulfilling leading to high realized asset prices. In markets where the role of speculative demand is less important, e.g. in markets

for non-storable commodities, negative feedback may play a more prominent role. For example in a supply-driven commodity market, if many producers expect future prices to be high they will increase production which, according to the law of supply and demand, will lead to a lower realized market price.

Heemeijer et al. (2009) investigate how the expectations feedback structure affects individual forecasting behaviour and aggregate market outcomes by considering market environments that *only* differ in the sign of the expectations feedback, but are equivalent along all other dimensions. The realized price is a linear map of the average of the individual price forecasts $p_{i,t}^e$ of six subjects. The (unknown) price generating rules in the *negative* and *positive* feedback systems were respectively:

$$p_t = 60 - \frac{20}{21} \left[\left(\sum_{i=1}^6 \frac{1}{6} p_{i,t}^e \right) - 60 \right] + \epsilon_t, \quad \text{negative feedback} \quad (22)$$

$$p_t = 60 + \frac{20}{21} \left[\left(\sum_{i=1}^6 \frac{1}{6} p_{i,t}^e \right) - 60 \right] + \epsilon_t, \quad \text{positive feedback} \quad (23)$$

where ϵ_t is an exogenous random shock to the pricing rule. The *only* difference between (22) and (23) is the sign of the slope of the linear map, $20/21 \approx +0.95$ resp. $-20/21 \approx -0.95$ ⁴.

Heemeijer et al. (2009) consider positive and negative feedback systems with small IID shocks $\epsilon_t \sim N(0, 0.25)$. Here we focus on the more recent experiments of Bao et al. (2012), with large permanent shocks to the fundamental price level. More precisely, these shocks have been chosen such that, both in the negative and positive feedback treatments, the fundamental equilibrium price p_t^* changes over time according to:

$$\begin{aligned} p_t^* &= 56, & 0 \leq t \leq 21, \\ p_t^* &= 41, & 22 \leq t \leq 43, \\ p_t^* &= 62, & 44 \leq t \leq 65. \end{aligned} \quad (24)$$

⁴In both treatments, the absolute value of the slopes is 0.95, implying in both cases that the feedback system is stable under naive expectations.

The purpose of these experiments was to investigate how the type of expectations feedback may affect the speed of learning of a new steady state equilibrium price, after a relatively large unanticipated shock to the economy.

Figure 4 shows for positive and negative feedback the average price behavior (top panels), realized prices in all groups (middle panels) and an example of individual forecasts in a positive as well as a negative feedback group (bottom panels). Aggregate behaviors under positive and negative feedback are strikingly different. Negative feedback markets tend to be rather stable, with price converging quickly to the new (unknown) equilibrium level after each unanticipated large shock. In contrast, under positive feedback prices are sluggish, converging only slowly into the direction of the fundamental value and subsequently overshooting it by large amounts.

Figure 5 reveals some other striking features of aggregate price behavior and individual forecasts. The top panel shows the time variation of the median distance to the RE benchmark price over all (eight) groups in both treatments. For the negative feedback treatment, after each large shock the distance spikes, but converges quickly back (within 5-6 periods) to almost 0. In the positive feedback treatment after each shock the distance to the RE benchmark shows a similar spike, but falls back only slowly and does not converge to 0. The bottom panel shows how the *degree of heterogeneity*, that is, the median standard deviation of individual forecasts, changes over time. For the positive feedback treatment after each large shock heterogeneity decreases very quickly and converges to (almost) 0 within 3-4 periods. Under positive feedback, individuals thus coordinate expectations quickly, but they all coordinate on the “wrong”, i.e., a non-RE price. In the negative feedback treatment heterogeneity is more persistent, for about 10 periods after each large shock. Persistent heterogeneity stabilizes price fluctuations and after convergence of the price to its RE fundamental individual expectations coordinate on the correct RE price.

One may summarize these results in saying that in the positive feedback treatment individuals quickly coordinate on a common prediction, but that coordination on the

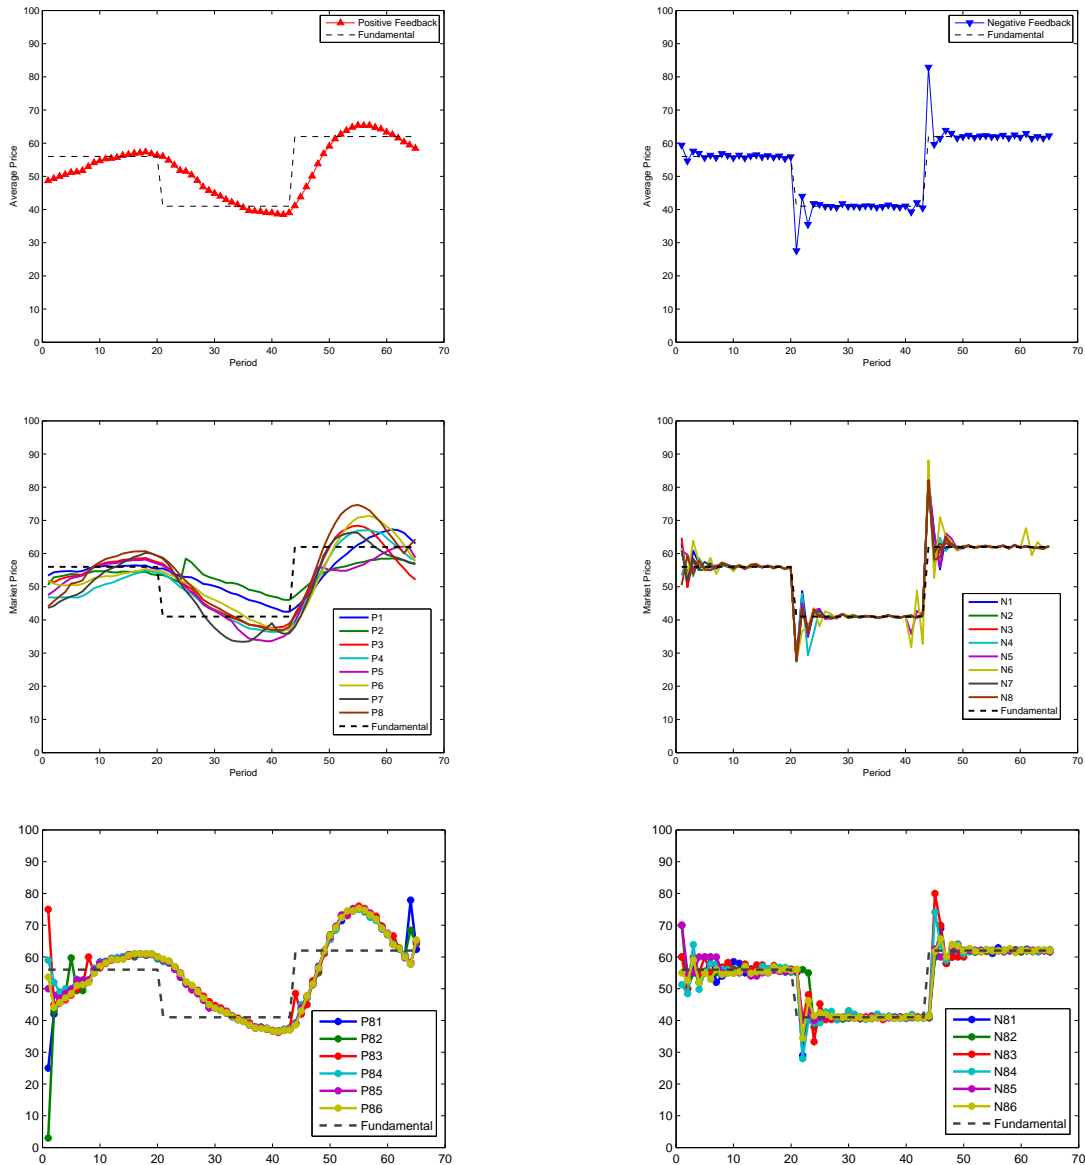


Figure 4: Positive feedback (left panels) and negative feedback (right panels) experimental data. Top panels: The average realized price averaged over all eight groups; Middle panels: the market prices for eight different groups; Bottom panels: predictions of six individuals in group P8 (left) and group N8 (right) plotted together with fundamental price (dotted lines).




Figure 5: Positive/Negative feedback markets with large shocks. These plots illustrate price discovery (top panel) and coordination of individual expectations (bottom panel). The top panel shows the median absolute distance to RE fundamental price, while the bottom panel shows the median standard deviation of individual predictions. In positive feedback markets coordination is quick, but on the “wrong”, i.e. non-RE, price.

“wrong” non-fundamental price occurs. As a result price behavior is very different from the perfect rational expectations equilibrium price. On the other hand, in the negative feedback treatment coordination is much slower, heterogeneity is more persistent but price convergence is quick. Stated differently, positive feedback markets are characterized by quick coordination and slow price discovery, while negative feedback markets are characterized by slow coordination, more persistent heterogeneity and quick price discovery. Notice also that under positive feedback, coordination on a non-RE-fundamental price is *almost self-fulfilling*, with small individual forecasting errors. The positive feedback market is thus characterized by coordination on almost self-fulfilling equilibria with prices very different from the perfectly rational self-fulfilling equilibrium⁵. Similar results have been obtained in laboratory experiments in other market settings, including a New Keynesian macro framework (Adam, 2007; Pfajfar and Zakelj, 2009; Assenza et al., 2012) and in a Lucas asset pricing model (Asparouhova et al., 2013).

Heuristic Switching Model (HSM)

Bao et al. (2012) fit a simple, behavioral heuristics switching model to the experimental data to explain both individual micro and aggregate macro behavior. Their model is a modification of Anufriev and Hommes (2012), who fitted an extension of the heterogeneous expectations model of Brock and Hommes (1997) to asset pricing experiments. The key feature of the model is that the subjects choose between a number of simple heuristics depending upon their relative performance. To keep it

⁵Wagener (2013) uses the same experimental data and shows weak individual rationality (i.e. unbiased forecast errors without autocorrelations) for both the negative and positive feedback treatments, but strong rationality (i.e. prices converge to the homogeneous REE price) only under negative feedback.

simple, the model only uses four rules⁶. An adaptive expectation (ADA) rule:

$$p_{t+1,1}^e = p_t^e + 0.85(p_t - p_{t,1}^e). \quad (25)$$

A contrarian rules (CTR) given by:⁷

$$p_{t+1,2}^e = p_t - 0.3(p_t - p_{t-1}). \quad (26)$$

A trend extrapolating rule (TRE) given by:

$$p_{t+1,2}^e = p_t + 0.9(p_t - p_{t-1}). \quad (27)$$

The coefficients of the first three rules are the medians of the estimated individual linear rules in Bao et al. (2012). The fourth rule is called an anchoring and adjustment heuristic (A&A) (Tversky and Kahneman, 1974):

$$p_{t+1,4}^e = 0.5(p_t^{av} + p_t) + (p_t - p_{t-1}). \quad (28)$$

The rule uses a time varying anchor, $0.5(p_t^{av} + p_t)$, which is the average of the last price and the sample mean of all past prices, and extrapolates the last price change $p_t - p_{t-1}$. Because of its flexible time-varying anchor, the A&A rule was successful in explaining persistent oscillations in asset pricing experiments (Anufriev and Hommes, 2012).

Subjects switch between these rules depending upon their relative performance. The performance of heuristic h , $h \in \{1, 2, 3, 4\}$ is measured by the squared prediction error⁸:

$$U_{t,h} = -(p_t - p_{t,h}^e)^2 + \eta U_{t-1,h}, \quad (29)$$

⁶Anufriev et al. (2013) fit a HSM with two heuristics, adaptive expectations versus a trend-following rule, to the positive-negative expectations feedback experiments of Heemeijer et al. (2009).

⁷Anufriev and Hommes (2012) used two different trend following rules in their model, a weak and a strong trend following rule, to describe asset pricing experiments with positive feedback. Because of the negative feedback treatment in Bao et al. (2012), one trend-following rule was replaced by a contrarian rule, i.e. with a negative coefficient (-0.3) which is able to detect (short run) up and down price oscillation characteristic for negative feedback markets.

⁸The participants in the experiment were also paid according to the quadratic prediction errors.

and $n_{h,t}$ is the fraction of the agents using heuristic h in the whole population. The parameter $\eta \in [0, 1]$ shows the relative weight the agents give to past errors compared to the most recent one. When $\eta = 0$, only the most recent performance is taken into account, and when $\eta > 0$, all past errors matter for the performance.

$$n_{t,h} = \delta n_{t-1,h} + (1 - \delta) \frac{\exp(\beta U_{t-1,h})}{\sum_{i=1}^4 \exp(\beta U_{t-1,i})}. \quad (30)$$

The parameter $\delta \in [0, 1]$ reflects the inertia with which participants stick to their rule. When $\delta = 1$, the agents simply do not update. When $\delta > 0$, each period a fraction of $1 - \delta$ participants updates their weights. The parameter $\beta \geq 0$ represents the “sensitivity” to switch to another strategy. The higher the β , the faster the participants switch to more successful rules in the recent previous periods. When $\beta = 0$, the agents will put equal weight on each rule. When $\beta = +\infty$, all agents who update switch to the most successful rule.

Figure 6 shows realized market prices and the one-period ahead simulated market prices (top panels), together with the evolution of the fractions of the four strategies of the heuristics switching model (bottom panels) for a typical group of the negative feedback (left panels) and the positive feedback treatment (right panels). The heuristics switching model matches the aggregate behavior of both positive and negative feedback quite nicely and provides an intuitive, *behavioral explanation* why these different aggregate patterns occur. In the negative feedback market, trend following strategies perform poorly and the contrarian strategy quickly dominates the market (more than 70% within 20 periods) enforcing quick convergence to the RE benchmark after each large shock. In contrast, in the positive feedback treatment, the trend following strategy performs well and dominates the market (with more than 50% trend-followers after 10 periods). The survival of trend following strategies in the positive feedback markets causes persistent deviations from the RE steady states, overreaction and persistent price fluctuations.

The difference in aggregate behavior in these experiments is thus explained by the fact that *trend-following rules are successful in a positive feedback environment*




Figure 6: Experimental and simulated prices using HSM model in one typical group from the positive (top left, group $P8$) and negative feedback treatment (bottom left, group $N8$) respectively. Experimental data (blue squares) and one-step ahead simulated prices from the HSM model (red circles) almost overlap. The right panels show the evolution of the four market heuristics in the positive (top right) and negative feedback treatments (bottom right). The trend following rule dominates in the positive feedback markets, while the contrarian rule dominates in the negative feedback markets. Parameters are: $\beta = 0.4, \eta = 0.7, \delta = 0.9$, as in Anufriev and Hommes (2012).

amplifying price oscillations and persistent deviations from the rational equilibrium benchmark price, while the same trend-following rules are driven out by the contrarian rule in the case of negative feedback. Coordination of individual expectations on trend-following rules and almost self-fulfilling equilibria in a positive expectations feedback environment has a large aggregate effect with realized market prices deviating significantly from the perfectly self-fulfilling RE benchmark.

5 Conclusions

Behaviorally rational expectations are characterized by an intuitive, plausible and easily observable form of consistency between forecasts and market realizations. Behaviorally rational expectations are almost, but not necessarily exactly, self-fulfilling. Parsimony is an important and attractive feature of behavioral rationality as simplicity makes coordination on an almost self-fulfilling equilibrium more likely as a description of aggregate macro behavior. Agents make mistakes, as their beliefs are only an approximation of complex reality. But in equilibrium, the mistake becomes self-fulfilling and it is not easy for agents to improve upon their individual forecasting. We have presented a simple theoretical case study -a first-order consistent expectations equilibrium in a macro NKPC-, an empirical case study -a 2-type asset pricing model with (almost) self-fulfilling bubbles and crashes amplified by agents' switching between mean-reverting and trend-following beliefs-, and an experimental case-study -positive and negative feedback experiments. Positive expectations feedback is a common feature of these case-studies, leading to coordination of individual expectations on almost self-fulfilling equilibria very different from exactly self-fulfilling equilibria. A simple and intuitive behavioral heuristics switching model explains these data well both at the micro and macro level.

References

- [1] Adam, K., (2007), Experimental Evidence on the Persistence of Output and Inflation, *Economic Journal*, 117, 603-635.
- [2] Adam, K., Marcet, A., 2011. Internal rationality, imperfect market knowledge and asset prices, *Journal of Economic Theory* 146 (2011) 12241252.
- [3] Anufriev, M., and Hommes, C.H., (2012), Evolutionary Selection of Individual Expectations and Aggregate Outcomes in Asset Pricing Experiments, *American Economic Journal: Microeconomics* 4(4), 35-64.
- [4] Anufriev, M., Hommes, C.H. and Philipse, R., (2013), Evolutionary selection of expectations in positive and negative feedback markets, *Journal of Evolutionary Economics* 23, 663-688.
- [5] Asparouhova, E., Bossaerts, P., Roy, N. and Zame, W. (2013), 'Lucas' in the laboratory, Caltech and UCLA Working Paper, May 2013.
- [6] Assenza, T., Heemeijer, P., Hommes, C.H. and Massaro, D. (2012), Individual expectations and aggregate macro behavior, CeNDEF Working Paper, University of Amsterdam.
- [7] Bao, T., Hommes, C.H., Sonnemans, J. and Tuinstra, J. (2012), Individual expectations, limited rationality and aggregate outcomes, *Journal of Economic Dynamics and Control* 36, 1101-1120.
- [8] Beshears, J., Choi, J.J., Fuster, A., Laibson, D. and Madrian, B.C. (2013), What Goes Up Must Come Down? Experimental Evidence on Intuitive Forecasting, *American Economic Review* 103, Paper and Proceedings, May 2013, pp. 570-574.
- [9] Boswijk, H.P., Hommes, C.H. and Manzan, S. (2007), Behavioral heterogeneity in stock prices, *Journal of Economic Dynamics and Control* 31, 1938-1970.

- [10] Branch, W.A., Evans, G.W., 2010. Asset return dynamics and learning. *Review of Financial Studies* 23(4), 1651-1680.
- [11] Branch, W.A., Evans, G.W., 2011. Learning about risk and return: a simple model of bubbles and crashes. *American Economic Journal: Macroeconomics* 3(3), 159-191.
- [12] Brock, W.A. and Hommes, C.H. (1997), A Rational Route to Randomness, *Econometrica* 65, 1059-1095.
- [13] Bullard, J., 1994. Learning equilibria. *Journal of Economic Theory* 64, 468-485.
- [14] Bullard, J., 2006. The learnability criterion and monetary policy. *Federal Reserve Bank of St. Louis Review* 88, 203-217.
- [15] Evans, G.W., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- [16] Evans, G.W., Honkapohja, S., 2013. Learning as a rational foundation for macroeconomics and finance, forthcoming in *Rethinking Expectations: The Way Forward for Macroeconomics*, Roman Frydman and Edmund S. Phelps (eds.), Princeton University Press, 2013.
- [17] Fuster, A., Laibson, D. and Mendel, B., 2010. Natural expectations and macroeconomic fluctuations. *Journal of Economic Perspectives* 24, 67-84
- [18] Gigerenzer, G. and Selten, R. (Eds.), 2001, *Bounded rationality. The adaptive toolbox*, MIT Press, Cambridge.
- [19] Gigerenzer, G., Todd, P.M. and the ABC Research Group, 1999, *Simple heuristics that make us smart*, Oxford University Press, Oxford.
- [20] Gordon, M. (1962), *The investment financing and valuation of the corporation*, Irwin, Homewood, IL.

- [21] Grandmont, J.-M., (1998), Expectation formation and stability in large socio-economic systems, *Econometrica* 66, 741-781.
- [22] Heemeijer, P., Hommes, C.H., Sonnemans, J., Tuinstra, J., 2009. Price stability and volatility in markets with positive and negative expectations feedback. *Journal of Economic Dynamics & Control* 33, 1052-1072.
- [23] Hommes, C.H., 2011. The heterogeneous expectations hypothesis: some evidence for the lab. *Journal of Economic Dynamics & Control* 35, 1-24.
- [24] Hommes, C.H., (2013), Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems, Cambridge University Press, Cambridge.
- [25] Hommes, C.H., Rosser, J.B., 2001. Consistent expectations equilibria and complex dynamics in renewable resource markets. *Macroeconomic Dynamics* 5, 180-203.
- [26] Hommes, C.H., Sorger, G., 1998. Consistent expectations equilibria. *Macroeconomic Dynamics* 2, 287-321.
- [27] Hommes, C.H. and Veld, D., in't, (2013), Behavioral heterogeneity and the financial crisis, CeNDEF Working Paper, University of Amsterdam, July 2013.
- [28] Hommes, C.H. and Zhu, M. (2013), Behavioral learning equilibria, *Journal of Economic Theory*, forthcoming.
- [29] Lansing, K.J., 2009. Time-varying U.S. inflation dynamics and the new Keynesian Phillips curve. *Review of Economic Dynamics* 12, 304-326.
- [30] Lucas, R., 1972. Expectations and the neutrality of money. *Journal of Economic Theory* 4, 103-124.
- [31] Muth, J.E., 1961. Rational expectations and the theory of price movements. *Econometrica* 29, 315-335.

- [32] Pfajfar, D. and B. Zakelj, Experimental Evidence on Inflation Expectation Formation, Discussion Paper, 2009.
- [33] Sargent, T.J., 1993. Bounded Rationality in Macroeconomics. Oxford University Press Inc., New York.
- [34] Shiller, R.J., (1981) Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421-436.
- [35] Shiller, R.J., 2000, *Irrational Exuberance*, Princeton University Press, Princeton.
- [36] Tversky, A. and Kahneman, D. (1974), Judgment under uncertainty: heuristics and biases, *Science* 185, 1124-1131.
- [37] Woodford, M., (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, Princeton, NJ.