

Monetary Policy Rules in a Non-Rational World: A Macroeconomic Experiment

Felix Mauersberger*

Barcelona GSE and Universitat Pompeu Fabra

This draft: February 14, 2017

Columbia Economics Discussion Paper Series No. 1617 01

Abstract

This paper introduces a new learning-to-forecast experimental design, where subjects in a virtual New-Keynesian macroeconomy based on Woodford (2013) need to forecast *individual* instead of *aggregate* outcomes. This approach is motivated by Preston (2005) and Woodford (2013) who note that forecasting individual outcomes is consistent with the microfoundations of the New-Keynesian model. Using this design, I analyze the impact of different interest rate rules on expectation formation and expectation-driven fluctuations. Even if the Taylor principle is fulfilled, instead of quickly converging to the REE, the experimental economy exhibits persistent purely expectation-driven fluctuations not necessarily around the REE. Only a particularly aggressive monetary authority achieves the elimination of these fluctuations and quick convergence to the REE. To explain the aggregate behavior in the experiment, I develop a “noisy” adaptive learning approach, introducing *endogenous* shocks into a simple adaptive learning model. This model is compared to a reinforcement model, which is developed to match individual behavior.

*I am in particular very grateful for the guidance by Davide Debortoli, Jordi Gali, Rosemarie Nagel and Michael Woodford. Furthermore, I would like to thank Jess Benhabib, John Duffy, George Evans, Albert Marcet, Luba Petersen and Alan Taylor as well as the participants of the Cognition and Decision Laboratory Meeting at Columbia University, the Monetary Economics Colloquium at Columbia University, the Macro Lunch at New York University, the CREI Macroeconomic Breakfast, the Barcelona GSE Summer Forum conference and the Tinbergen Institute Workshop in Behavioral Macroeconomics for their feedback. I would furthermore like to thank the Barcelona GSE for financial support for this project.

1 Introduction

A large empirical literature following Taylor (1993) documents that interest rate setting by central banks can be described by simple rules. The validity of policy recommendations regarding interest rate rules crucially depends on how agents form their expectations regarding future economic conditions. The standard approach of modeling expectations as model-consistent or rational expectations can be considered as unnecessarily strong, as this presumes a correct understanding of the model, knowledge of all parameters and common knowledge of rationality, i.e. that all agents know that all other agents are rational. The central policy recommendation for the design of interest rate rules both under rational expectations and in the macroeconomic learning literature is considered to be the "*Taylor principle*", implying that interest rates should actively respond to inflation. Under rational expectations, the reason for the desirability of the Taylor principle is that in New-Keynesian frameworks it guarantees local determinacy of the rational expectations equilibrium (REE) and thus avoids the emergence of multiplicity.

A weak requirement for the adoption of rational expectations is often whether agents at least asymptotically learn to form model-consistent or rational expectations. (see e.g. Bullard and Mitra (2002); Evans and Honkapohja (2003, 2006), Marimon and Sunder (1994)) A theoretical learning literature addresses this question using particular hypotheses regarding the process of expectation formation. One assumption that is frequently made is that agents in the model behave like econometricians. (Marcet and Sargent, 1989; Sargent, 1994; Evans and Honkapohja, 2001) In the context of a New-Keynesian model, Bullard and Mitra (2002) investigate stability under least square learning (E-stability¹) of the REE based on the *reduced-form* equations, i.e. the "dynamic IS" equation, the "Phillips curve" and the interest rate rule. They find that the Taylor principle is necessary and sufficient for E-stability. This finding has, however, already been questioned by the adaptive learning literature, since Orphanides and Williams (2007) recommend a more aggressive monetary policy in the presence of learning, if agents' misperceptions of the economy's steady state levels are taken into account.

Yet, many applications in the theoretical learning literature (e.g. Evans and Honkapohja (2001), Bullard and Mitra (2002)) rely on the *reduced-form* equations of the New-Keynesian model. This means treating the problem of aggregation in a simple manner by solving the *rep-*

¹Marcet and Sargent (1989) show a one-to-one correspondence between E-stability and learnability under least square learning.

representative agents' maximization problem and then replacing the expectations operator in the log-linearized solutions, i.e. the “dynamic IS” equation and the “Phillips curve”, by a single non-rational form of expectations (e.g. Bullard and Mitra (2002)) or an aggregate over individual heterogeneous expectations. (e.g. Adam (2007), Assenza et al. (2014), de Grauwe (2011)). Preston (2005), Woodford (2013) and Kurz et al. (2013) instead suggest a more *structural-form* approach² by solving the *individual* maximization problem, specifying a possibly non-rational form of individual forecasting of individual variables and then aggregating. It is apparent that an aggregate of individual expectations of *individual* variables, e.g. $\int_{j=0}^1 \hat{E}_t^i c_{t+1}^i$ is generally not the same as an aggregate of individual expectations of an *aggregate* variable, e.g. $\int_{j=0}^1 \hat{E}_t^i c_{t+1}$. Furthermore, Preston (2005) notes that the law of iterated expectations holding at individual level does not imply that it holds at the aggregate level.³ While the theoretical learning literature relies on particular assumptions regarding agents' expectation formation behavior, the nature of belief formation is eventually an empirical matter. As belief formation may be different when individuals think about aggregates rather than market outcomes of their individual household or firm unit, to what extent the simplification of the reduced form is adequate ultimately becomes an empirical question.

This paper addresses this question and investigates the implications of learning in the structural form of the New-Keynesian model on interest-rate setting. Furthermore, it characterizes the manner of expectation formation as observed in a structural form setup. To obtain data on expectation formation, I introduce a new learning-to-forecast laboratory experiment⁴ based on the structural form of the heterogeneous-expectations New-Keynesian framework in Woodford (2013). In this type of experiment, subjects act as forecasters in virtual markets or economies in which outcomes are generated by the computer conditional on subjects' beliefs. My design is without any exogenous shocks so that all shocks in the experiment come from agents' expectations.

The methodology of a laboratory experiment is adopted due to several distinct advantages. As opposed to real macroeconomic time series where the REE cannot easily be observed, in the laboratory the REE is controlled by the experimenter, which enables to easily examine whether learning dynamics converge to the REE. Secondly, surveyed inflation expectations depend on

²One could alternatively label the reduced form “representative agent” model and the structural form “heterogeneous agent” model. However, in the context of this paper these would be confusing labels, as the beliefs of the multiple subjects in the laboratory are generally diverse.

³Branch and McGough (2009) and Kurz et al. (2013) derive the reduced-form New-Keynesian equations under diverse beliefs rely on other strong behavioral assumptions on the expectations operator, whose empirical validity raises questions.

⁴Laboratory experiments are an increasingly popular methodology to study questions related to macroeconomics. See Duffy (2014) for a survey.

many uncontrollable factors so that it can be immensely challenging to isolate and identify the effect of monetary policy. Thirdly, in the laboratory the experimenter can incentivize subjects by paying them according to their forecasting performance, while respondents in surveys do not have any incentive to accurately contemplate their beliefs. Fourthly, a laboratory setting is a less complex environment than reality, as the experimenter limits the task for subjects and controls the available information. If subjects are not able to rationally form expectations in a stable setting, then rational expectations are unlikely to be a reasonable assumption in a more complicated setting, being subject to structural changes, regime switches, trade shocks etc.. Furthermore, simplifying the task of forecasting for subjects with respect to reality should render it easier for the researcher to detect regularities in subjects' behavior.

The experimental test about monetary policy in a *reduced-form* New-Keynesian model has been delivered by Pfajfar and Žakelj (2014) and Assenza et al. (2014). These authors likely used the reduced form, because they were the first to implement a New-Keynesian economy in the laboratory and this likely represents the simplest way to do so. Both corroborate the theoretical literature that relies on least square learning and find convergence to the REE under satisfaction of the Taylor principle, while without satisfying the Taylor principle non-convergence patterns and high fluctuations are observable.⁵

My experimental results differ from Pfajfar and Žakelj (2014) and Assenza et al. (2014) and are more congruent with the theoretical results of Orphanides and Williams (2007) and Ferrero (2007), as I do not find evidence that Taylor-rule reaction coefficients outside the unit circle guarantee learning dynamics converging to the REE in a structural New-Keynesian model. If the Taylor principle is barely satisfied, I instead find persistent fluctuations not necessarily around the steady state. Since without any exogenous shocks all fluctuations are purely expectation-driven, it is desirable to eliminate these and ensure convergence. To do so, my results and the subsequently developed “noisy” adaptive model suggest that the monetary authority should adopt a reinforced Taylor principle reacting more strongly to undesired inflation, with a reaction coefficient on inflation around 3. Furthermore, I show that the fluctuations in a suitable sample of US macroeconomic data are matched better by the experimental economies using the structural form than by the ones using the reduced form.

To understand the differences observed in the experimental economies under different Taylor

⁵Arifovic and Petersen (2015) use a similar design to examine policies to escape expectation-driven liquidity traps.

rules, I consider two general approaches to modeling learning: The first approach is using a parsimonious model to explain the *aggregate* behavior in the experiment, while the second approach is considering a more sophisticated model to replicate *aggregate and individual* behavior. I conduct agent-based computational economic (ACE) path simulations to investigate whether the dynamics in the experimental economies are predicted by different learning models.

The starting point for the first approach, i.e. modeling aggregate behavior, is exploring a frequently assumed hypothesis in the macroeconomic learning literature proposed by Marcet and Sargent (1989) and Evans and Honkapohja (2001), postulating that agents' forecasting behavior resembles an econometrician. While subjects certainly do not literally apply econometric techniques, least square forecasters can alternatively be interpreted as adjusting their forecasts in the direction of the last observed forecast error and attaching decreasing weight to new observations. This is the intuition why least square learning can be considered a special case of *adaptive learning*. (Sargent (1994) for details) I also consider an alternative case of adaptive learning, which is *constant gain learning*, meaning that agents discount past observations. While adaptive learning predicts the different convergence patterns of weakly active and strongly active monetary policy, in the absence of shocks the convergence is monotonic so that the fluctuations remain unexplained.

A combination of the macroeconomic and the microeconomic learning literature, which I call *noisy adaptive learning*, improves the fit to the second moments of the experimental data: similarly to some applications in the microeconomic learning literature, see e.g. Fudenberg and Harris (1992); Binmore et al. (1995); Nagel and Vriend (1999); Anderson et al. (2004), idiosyncratic cognitive shocks are introduced in the learning process, which is taken from the macroeconomic literature as constant-gain adaptive learning. A novel feature is that the shock variance is endogenous with a specification similar to the conditional heteroskedasticity models in the financial econometrics literature. (Engle, 1982; Bollerslev, 1986).

Following the second approach, i.e. to fit both *aggregate and individual* behavior, I consider the hypothesis by Brock and Hommes (1997) and Anufriev and Hommes (2012), where agents choose among a finite set of possible forecasting models with endogenous probabilities of using each model. I find that the functional form for the switching probabilities of the reinforcement model by Roth and Erev (1998) provides a better fit to the data than the one specified by Brock and Hommes (1997) or Anufriev and Hommes (2012). The reason is the assumption by Brock and Hommes (1997) that agents consider fictitious play, i.e. the ceteris-paribus payoff for forecasting

rules that were not used. Fictitious play may be dubious in the context of my experiment, where agents' own forecasts have a relatively large influence on individual market outcomes.

2 Model

A former approach has been to implement complete structural versions of the New-Keynesian model in the laboratory (see e.g. Petersen (2015)), where subjects assume the roles of households and firms and make similar decisions to agents in the model such as production, consumption and labor supply decisions. Yet, in such a framework, it is challenging to isolate the role of expectations.

Hence, I take a different approach and base my experiment on the linearized heterogeneous expectations New-Keynesian model by Woodford (2013).⁶ Expectations represent a well-behaved probability measure, but can be heterogeneous across agents and need not necessarily be model-consistent. For the experiment, I only make minimal changes from Woodford (2013):

- Although Woodford (2013) presents the model with an exogenous preference shock and a markup shock, in the model used for the experiment there will be no exogenous shocks. Since the minimum state-variable solution of the model is thus a constant, it will create an easier learning environment for agents. Shocks are introduced by individuals deviating from the rational expectations outcome.
- While Woodford (2013) considers a log-linearized model around a zero inflation steady state, I consider a model of full price indexation as first used by Yun (1996), i.e. in the periods between price re-optimizations firms mechanically adjust their prices according to steady state inflation $\bar{\pi}$. The reason is threefold: firstly, most central banks target a medium-run inflation above zero. Secondly, assuming that subjects start with a random guess, zero would be the expected initial outcome, so that it is convenient to specify the equilibrium of one group to be non-zero. Thirdly, a non-zero inflation steady state provides additional margin to prevent interest rates from attaining the zero lower bound.

⁶Based on the idea that the dynamics should be consistent with exact optimization behavior independently of the size of the expectational error, Hommes et al. (2015) introduce a non-linear New-Keynesian model into the laboratory. For simplicity, the input obtained from subjects are point forecasts. As the variables that are to be forecast are not independent, determining behavior in a non-linear model would ideally not only require conditional means of future variables but also other moments of the conditional distribution. Hence, a more complex description of subjects' probability beliefs would be needed.

The model is presented in more detail in Woodford (2013) and Appendix 9.1. Below an abridged version.

2.1 The demand side

This cashless economy is populated by a continuum of households $i \in (0, 1)$ maximizing

$$\max_{C_T^i, h_T^i(j)} \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} [U(C_T^i) - \Upsilon(h_T^i(j))] \quad (1)$$

where \hat{E}_t^i is an arbitrary subjective (not necessarily rational) expectations operator of household i given the information set in period t , which however satisfies the law of iterated expectations such that $\hat{E}_t^i \hat{E}_{t+1}^i x_{t+k} = \hat{E}_t^i x_{t+k}$. U is utility, C_t^i is a composite consumption good, Υ is the disutility of labor supply and $h^i(j)$ is the household i 's labor supply for the production of good j . There is an equal number of households supplying labor for each type of good.⁷

As a union negotiates the wage on behalf of all households, the individual household has to supply the hours of work demanded by the firm at the given wage. Thus, the household's only decision is its (real) consumption expenditure, C_t^i . There is one single riskless one-period bond, whose holdings in a period t by household i can be denoted B_t^i . One main difference of this heterogeneous agent model to the representative agent model is that households can have non-zero asset holdings. A log-linearised approximation to the consumption function takes the form:

$$\hat{c}_t^i = (1 - \beta)\hat{b}_t^i + \sum_{T=t}^{\infty} \beta^{T-t} \hat{E}_t^i \{ (1 - \beta)(\hat{Y}_T - \hat{\tau}_T) - \beta\sigma(\hat{i}_T - \hat{\pi}_{T+1}) + (1 - \beta)s_b(\beta\hat{i}_T - \hat{\pi}_T) \} \quad (2)$$

where the $\hat{\cdot}$ -superscripts denote log-deviations from the steady state, b_t the value of the maturing bond holdings deflated by the price level at period $t - 1$, Y aggregate output, τ the level of lump-sum taxes deflated by the current period price level, i the nominal interest rate set by the central bank, π the current level of inflation and s_b the steady-state level of government debt divided by the steady state level of output. (2) can be rewritten in its recursive form as

$$\hat{c}_t^i = (1 - \beta)\hat{b}_t^i + (1 - \beta)(\hat{Y}_T - \hat{\tau}_T) - \beta[\sigma - (1 - \beta)s_b]\hat{i}_t - (1 - \beta)s_b\hat{\pi}_t + \beta\hat{E}_t^i v_{t+1}^i \quad (3)$$

⁷Woodford (2003, p. 144ff.) shows the equivalence to a representative agent supplying labor for each type of good.

where

$$v_t^i \equiv \sum_{T=t}^{\infty} \beta^{T-t} \hat{E}_t^i \{ (1-\beta)(\hat{Y}_T - \hat{\tau}_T) - [\sigma - (1-\beta)s_b](\beta\hat{i}_T - \hat{\pi}_T) \} \quad (4)$$

I adopt this notation from Woodford (2013), since the advantage is that individuals only need to forecast a single variable. Using the goods market clearing condition $\hat{Y}_t = \int \hat{c}_t^i di$, aggregate demand can be obtained as

$$\hat{Y}_t = (1-\beta)\hat{b}_t + v_t - \sigma\hat{\pi}_t \quad (5)$$

Woodford (2013) assumes for simplicity that government expenditure is an exogenous disturbance, but, since I introduce the model into the laboratory without any shocks, in my setup, the government merely uses the taxes to service the debt it has accumulated. Hence, the government's flow budget constraint is given as

$$\hat{b}_{t+1} = \beta^{-1}[\hat{b}_t - s_b\hat{\pi}_t - \hat{\tau}_t] + s_b\hat{i}_t \quad (6)$$

(4) implies the recursive form

$$v_t^i = (1-\beta)v_t + \beta(1-\beta)(\hat{b}_{t+1} - \hat{b}_t) - \beta\sigma(\hat{i}_t - \hat{\pi}_t) + \beta\hat{E}_t^i v_{t+1}^i \quad (7)$$

where $v_t = \int v_t^i di$ is the average value across agents of the expectational variable defined in (4).

I follow Woodford (2013) in assuming that expectations are Ricardian so that

$$b_t = \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} [\hat{\tau}_T - s_b(\beta\hat{i}_T - \hat{\pi}_T)] \quad (8)$$

This is a strong assumption. However, it simplifies the model considerably and this assumption is frequently made in non-RE analyses. (See Woodford (2013) for a detailed discussion.) In the context of this experiment, this can be interpreted as the computerized household understanding that government cannot forever accumulate debt. Under this assumption, aggregate demand is independent of the supply of public debt and can be more compactly written as

$$\hat{Y}_t = \bar{v}_t - \sigma\hat{\pi}_t \quad (9)$$

where $\bar{v}_t = v_t + (1 - \beta)\hat{b}_t$ is the aggregate of a subjective variable \bar{v}_t^i , which is the variable that one group of subjects, labeled as “household advisors”, needs to forecast in the experiment, and which can be defined simply as

$$\bar{v}_t^i \equiv \sum_{T=t}^{\infty} \hat{E}_t^i \{(1 - \beta)\hat{Y}_T - \sigma(\beta i_T - \pi_T)\} \quad (10)$$

(10) together with (2) and (8) imply

$$\bar{v}_t^i = \hat{c}_t^i + \sigma \hat{\pi}_t - (1 - \beta)(b_t^i - b_t) \quad (11)$$

The model was log-linearized near a steady-state in which all households have the same level of bond holdings, so that $b_t^i - b_t$ is small by assumption. Since this term is additionally multiplied by a small factor $1 - \beta$, the following approximate relationship holds

$$\bar{v}_t^i \approx \hat{c}_t^i + \sigma \hat{\pi}_t \quad (12)$$

As c_t^i represents real expenditure, nominal expenditure can be written (in log deviation) as $\hat{c}_t^i + p_t$. To make the variable stationary (as there will be a unit root in the price level in equilibrium), one can define a real measure of expenditure in terms of the previous period price level p_{t-1} as

$$\hat{c}_t^i + p_t - p_{t-1} \quad (13)$$

Since $\sigma = 1$, (13) is equal to

$$\hat{c}_t^i + \underbrace{\sigma(p_t - p_{t-1})}_{=\hat{\pi}_t} \approx \bar{v}_t^i \quad (14)$$

so that \bar{v}_t^i is close to real expenditure. Likewise (9) implies

$$\bar{v}_t = \hat{Y}_t + \sigma \hat{\pi}_t \quad (15)$$

so that the aggregate of \bar{v}_t^i is likewise closely related to the real output gap (expressed in terms of the price level in period $t - 1$.)

(10) is consistent with a recursive form of

$$\bar{v}_t^i = (1 - \beta)\bar{v}_t - \beta\sigma(\hat{i}_t - \hat{\pi}_t) + \beta\hat{E}_t^i \bar{v}_{t+1}^i \quad (16)$$

which is the data-generating process in the experiment.

The computerized central bank is specified as to adopting a Taylor rule, which satisfies the zero lower bound:

$$i_t = \max(0, \bar{i} + \phi_\pi(\pi_t - \bar{\pi})) \quad (17)$$

2.2 The supply side

Woodford (2013) assumes Calvo (1983) price-setting such that a fraction $0 < \alpha < 1$ of goods prices are exogenously held fixed in any period. Producers engage in Dixit-Stiglitz monopolistic competition, which means that each firm sets the price for a good that it alone produces. Under full price indexation, the log-linearized approximation of the inflation dynamics is given by

$$\hat{\pi}_t = (1 - \alpha)\hat{p}_t^* \quad (18)$$

where, for each firm $j \in (0, 1)$ that is chosen to re-optimize its price in period t , p_t^{*j} is the amount by which the firm would choose to set the log price of its good higher than p_{t-1} . $p_t^* = \int p_t^{*j} dj$ is the average value of this variable across all firms that are chosen by the Calvo mechanism to reoptimize prices in period t . The solution to firm j 's maximization problem takes the form:

$$\hat{p}_t^{*j} = (1 - \alpha\beta) \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \hat{E}_t^j \{ p_T^{opt} - p_{t-1} - (T - t + 1)\bar{\pi} \} \quad (19)$$

where p_T^{opt} is the single-period optimal log price, being the same for each firm, as they face the same labor market and aggregate conditions. The recursive form of (19) is:

$$\hat{p}_t^{*j} = (1 - \alpha\beta)(p_t^{opt} - p_{t-1} - \bar{\pi}) + \alpha\beta(\hat{E}_t^j \hat{p}_{t+1}^{*j} + \hat{\pi}_t) \quad (20)$$

Suppose that the union, setting the wage on behalf of the households, pursues the objective that

the intertemporal optimality condition holds at the aggregate so that:

$$\frac{\Upsilon_h(h_t)}{u_{C,t}(C_t)} = \frac{w_t}{P_t} \quad (21)$$

By log-linearising (21), one obtains

$$\hat{\omega}_t = \hat{\vartheta}_t - \hat{u}_{c,t} \quad (22)$$

where ω_t is the log real wage, ϑ_t is the log of the (common) marginal disutility of labor and $u_{c,t}$ is the (log) aggregate marginal utility of additional real income across households. Hence

$$\hat{\omega}_t = \hat{\vartheta}_t + \sigma^{-1} \hat{c}_t = \hat{\vartheta}_t + \sigma^{-1} \hat{Y}_t \quad (23)$$

Given that $\hat{m}c_t = \hat{\omega}_t - m\hat{p}n_t$ and since both $\hat{\vartheta}_t$ and $m\hat{p}n_t$ can be expressed as functions of labor hours and thus as output (being determined by the market clearing condition for the aggregate goods market), one can summarize p_t^{opt} as

$$p_t^{opt} = p_t + \xi \hat{Y}_t \quad (24)$$

Under the assumption of constant returns to scale in the aggregate and a disutility function of labor of

$$\Upsilon(h_t^i(j)) = \frac{h_t^i(j)^{1+\varphi}}{1+\varphi} \quad (25)$$

it is easy to see that $\xi \equiv \sigma^{-1} + \varphi$. We obtain, by using (24) in (20):

$$\hat{p}_t^{*j} = (1 - \alpha)\hat{p}_t^* + (1 - \alpha\beta)\xi\hat{Y}_t + \alpha\beta\hat{E}_t^j\hat{p}_{t+1}^{*j} \quad (26)$$

2.3 From the structural to the textbook reduced form

Only under the hypothesis that all expectations are identical, subjective expectations $\{\hat{E}_t^i\}$ and $\{\hat{E}_t^j\}$ can be replaced by the single expectations operator \hat{E}_t and the system reduces to

$$\bar{v}_t = -\sigma(i_t - \pi_t) + \hat{E}_t \bar{v}_{t+1} \quad (27)$$

$$\pi_t = \kappa \hat{Y}_t + \beta \hat{E}_t \pi_{t+1} \quad (28)$$

where $\kappa \equiv \frac{(1-\alpha)(1-\alpha\beta)}{\alpha} \xi$. This represents a reduced-form system similar to the ones in Woodford (2003), Galí (2008) or Walsh (2010), which has been used in the experiments by Pfajfar and Žakelj (2014) and by Assenza et al. (2014).

2.4 System

By integrating (7) over i and (26) over j and by inserting the resulting equations back into (7) and (26), one obtains the system upon which the experimental economy in the new design is based

$$\begin{bmatrix} \bar{v}_t \\ p_t^* \\ \bar{v}_t^i \\ p_t^{*j} \end{bmatrix} = \begin{bmatrix} (1-\alpha)\alpha\beta\sigma(\phi_\pi - 1) \\ (1-\beta) + \phi_\pi\sigma\xi(1-\alpha)(1-\alpha\beta) \\ \alpha\beta\sigma(\phi_\pi - 1)(1-\alpha) \\ \alpha(1-\beta) + \phi_\pi\sigma\xi(1-\alpha)(1-\alpha\beta) \end{bmatrix} \Omega \bar{\pi} + A \begin{bmatrix} \int_{i=0}^1 \hat{E}_t^i \bar{v}_{t+1}^i di \\ \int_{j=0}^1 \hat{E}_t^j p_{t+1}^{*j} dj \\ \hat{E}_t^i \bar{v}_{t+1}^i \\ \hat{E}_t^j p_{t+1}^{*j} \end{bmatrix} \quad (29)$$

with $\Omega \equiv \frac{1}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)}$ and

$$A \equiv \begin{bmatrix} \frac{\alpha + \sigma \xi (1-\alpha)(1-\alpha\beta)}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)} & \frac{\alpha\beta\sigma(\phi_\pi - 1)(\alpha - 1)}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)} & 0 & 0 \\ -\frac{\xi(\alpha\beta - 1)}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)} & \frac{\alpha\beta}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)} & 0 & 0 \\ \frac{\alpha(1-\beta) + (1-\alpha)(1-\alpha\beta)\sigma\xi(1-\beta\phi_\pi)}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)} & \frac{\alpha\beta\sigma(\phi_\pi - 1)(\alpha - 1)}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)} & \beta & 0 \\ -\frac{\xi(\alpha\beta - 1)}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)} & \frac{-\alpha\beta(\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta) - 1)}{\alpha + \phi_\pi \sigma \xi (1-\alpha)(1-\alpha\beta)} & 0 & \alpha\beta \end{bmatrix} \quad (30)$$

2.5 Calibration

I follow Clarida et al. (2000), using $\sigma = 1$, $\kappa = 0.3$ and $\beta = 0.99$. I set $\bar{\pi} = 2$, since this is a commonly chosen inflation target by central banks. For the structural form, I need to specify two more parameters. Following Rotemberg and Woodford (1997), I specify $\alpha = 0.66$, which implies $\xi = 1.76$ to be consistent with $\kappa = 0.3$.

Parameter	Value	Comment
β	0.99	implies quarterly risk-free real rate of 0.01
κ	0.3	
α	0.67	Rotemberg and Woodford (1997)
ξ	1.76	consistency with $\kappa \equiv \frac{(1-\alpha)(1-\alpha\beta)}{\alpha} \xi = 0.3$
σ	1	elasticity of intertemporal substitution
$\bar{\pi}$	2	commonly chosen inflation target

Table 1: **Calibration**

With this calibration, the coefficient matrices become⁸

Structural-form system with $\phi_\pi = 0.5$

$$A^{\phi_\pi=0.5} = \begin{bmatrix} 1.1304 & 0.1435 & 0 & 0 \\ 0.7826 & 0.8609 & 0 & 0 \\ 0.1404 & 0.1435 & 0.99 & 0 \\ 0.7826 & 0.2009 & 0 & 0.66 \end{bmatrix} \quad (31)$$

Structural-form system with $\phi_\pi = 1.5$

$$A^{\phi_\pi=1.5} = \begin{bmatrix} 0.8966 & -0.1138 & 0 & 0 \\ 0.6207 & 0.6828 & 0 & 0 \\ -0.0934 & -0.1138 & 0.99 & 0 \\ 0.6207 & 0.0228 & 0 & 0.66 \end{bmatrix} \quad (32)$$

Structural-form system with $\phi_\pi = 3$

$$A^{\phi_\pi=3} = \begin{bmatrix} 0.6842 & -0.3474 & 0 & 0 \\ 0.4737 & 0.5211 & 0 & 0 \\ -0.3058 & -0.3474 & 0.99 & 0 \\ 0.4737 & -0.1389 & 0 & 0.66 \end{bmatrix} \quad (33)$$

2.6 Discussion of the different systems

The linear systems imply a unique rational expectations steady state with $\bar{v}^i = \bar{v} = 0$ and $p^{*j} = p^* = \bar{\pi} = 2$. For the REE to be determinate, the eigenvalues of the Jacobian matrix A need to lie inside the unit circle. (Blanchard and Kahn, 1980) If monetary policy conforms with the

⁸The systems are shown for the case where the zero lower bound is not binding.

Taylor principle and the ZLB is not binding, the absolute values of all eigenvalues of coefficient matrix A in systems (32) and (33) lie within the unit circle. However, if agents do not know the underlying model but have to learn the coefficients, they implicitly also need to learn the eigenvalues of the system. Assenza et al. (2014, p.30) note that, in the reduced-form system, where both eigenvalues depend on the reaction coefficient ϕ_π , the perceived law of motion can potentially be rendered stationary by reacting more strongly to inflation, as this can push the eigenvalue sufficiently far away from the unit circle.

Yet, this conclusion becomes questionable, if one notes that the Jacobian matrix A in the structural-form systems (32) and (33) has two eigenvalues that are constants equal to $\alpha\beta$ and β and thus independent of ϕ_π . Since α and β are constants within the unit circle under a standard calibration, the rational expectations equilibrium is determinate. Yet, if individuals have to learn the coefficients, especially $\beta = 0.99$, implying near-unit root behavior, a problem arises, since even very small misperceptions by individuals might give rise to a perceived law of motion that has an eigenvalue outside the unit circle and which thus results in explosive or oscillatory paths. As a coefficient of $\beta = 0.99$ appears in the equation (7) of the households, one would, in particular, expect unstable dynamics for \bar{v} .

3 New experimental design

The experiment was conducted in Spanish in the BES laboratory at Pompeu Fabra University (Barcelona, Spain) in April and May 2015. The whole experiment was computerized with z-tree (Fischbacher, 2007). Most subjects were undergraduates studying Business Administration, Economics, Engineering, Humanities, Management or Medicine. The experiment consisted of 50 periods. Subjects were recruited using the software ORSEE (Greiner, 2004) and every subject was only allowed to participate in one experimental session.

I first investigated the replicability of the experiments by Assenza et al. (2014) and Pfajfar and Žakelj (2014) in this laboratory and thus ran two groups with a reduced-form design based on the “DIS equation” (27) and the “Phillips curve” (28). Appendix 9.2 shows that the result of these replications conforms with the previous results of Assenza et al. (2014) and Pfajfar and Žakelj (2014), i.e. that a Taylor rule coefficient ϕ_π just outside the unit circle guarantees quick

convergence to the REE.

The new experimental design based on a *structural-form* New-Keynesian model consists of⁹

- computerized households
- computerized firms
- a computerized central bank
- 6 HUMAN household advisors being asked to submit $\hat{E}_t^i \bar{v}_{t+1}^i$
- 6 HUMAN firm advisors being asked to submit $\hat{E}_t^j p_{t+1}^{*j}$

Hence, there were two possible roles to which subjects were randomly assigned: an “advisor to firms” or an “advisor to households”. They keep the same role for all 50 periods. As opposed to the design of Pfajfar and Žakelj (2014) and Assenza et al. (2014),¹⁰ who ask subjects to forecast *aggregate* outcomes, subjects in my experiment forecast *individual* outcomes.

3.1 Role 1: Household advisor

In each period, advisors to households are asked to submit $\hat{E}_t^i \bar{v}_{t+1}^i$. An alternative would be asking them to forecast inflation, their own optimal consumption and bond holdings, but forecasting several variables at the same time complicates the task for subjects.¹¹

3.2 Role 2: Firm advisor

In each period, advisors to firms are asked to submit $\hat{E}_t^j p_{t+1}^{*j}$, i.e. the forecast of the deviation of a firm’s own optimal price in the next period (t+1) from the general price level in the current period (t). Since in the model only the forecasts of the subsample of firms matters that is chosen by the Calvo mechanism to reoptimize their prices, subjects were informed that the firm that seeks their advice can be different in each period. This way the forecasts of all subjects could be considered

⁹By assumption, fiscal policy has no effect on the beliefs of computerized households and aggregate output is entirely independent of the paths of both public debt and taxes. Hence, mentioning the existence of a government to agents might only create confusion and does not provide any benefits.

¹⁰See Appendix 9.2 for more details.

¹¹Since leaving variables in abstract terms would make the instructions too complicated, \bar{v}_t was labeled for participants as the deviation of real expenditure (out of the current labor income) from a household’s “usual expenditure”. See Section 2 for the explanation why \bar{v}_t is closely related to a household’s real expenditure. “Usual expenditure” is supposed to be a more intuitive description of the long-run steady state and was explained to be participants to be “average expenditure per period that a household can expect to be able to afford over time.”

for the inflation dynamics. Due to identical production technologies of all firms, the optimal price setting $p_t^{*j_1}$ of a particular firm j_1 in period t can be compared to the forecast of this optimal price setting made in the previous period for a possibly different firm j_2 , $\hat{E}_{t-1} p_t^{*j_2}$. Since deflation beyond -100% is impossible, -100% was introduced as a natural lower bound. To be consistent across groups, -100% was also set as a lower bound for the household advisors.¹²

3.3 Aggregate outcomes

Due to a continuum of agents in the model, the influence of one person should be infinitesimally small. Since it is not feasible to conduct the experiment with a very large number of subjects, aggregate outcomes are proxied by medians so that in particular

- p_t^* is the median price change of firms that are paired with subjects
- \bar{v}_t is the median of \bar{v}_t^i across households advised by subjects
- $\int_{i=0}^1 \hat{E}_t^i \bar{v}_{t+1}^i di$ is the median expectation of all household advisors
- $\int_{j=0}^1 \hat{E}_t^j p_{t+1}^{*j} dj$ is the median expectation of all firm advisors

A further advantage of aggregating with the median rather than the average is to avoid that aggregate variables are driven by a single subject. Hence, employing the median for aggregating has been done by past learning-to-forecast experiments like Kryvtsov and Petersen (2015), Petersen (2015) and Arifovic and Petersen (2015).

3.4 Information provided to subjects

3.4.1 Beginning of the experiment

All subjects received the same instructions,¹³ which informed them about their own role, the role of the other group and I also gave qualitative information about the economy. The qualitative information included explanations of the macroeconomic variables and a verbal description of equations (7), (26) and the Taylor rule (17). Subjects did not know the coefficients in the equations or the steady states of the model.¹⁴ Apart from the fact that this has become standard for complex

¹²For technical and operational reasons, I also had to introduce an upper bound, which was set to 1000.

¹³The full instructions are available upon request.

¹⁴Full price indexation in this context could also be interpreted as the firms not chosen by the Calvo mechanism adjusting their prices on average by the long-run steady state.

experimental games¹⁵ as this one, this approach can be justified in particular as follows: firstly, it is not realistic that agents have fixed rules in mind when making decisions, but that these rules need to be extracted from possibly long histories of past data. Secondly, since agents are provided with a calculator, any algebraically inclined participants could theoretically calculate the coefficients after few periods, if they were of interest. Thirdly, two central questions of the adaptive macroeconomic learning literature can be considered to be: firstly, whether REEs are learnable solely from a sufficiently long history of data; secondly, since the macroeconomic literature commonly assumes that agents in the model know its structural equations, it addresses the question how agents inside the model learn the structure of the economy. These questions also lie at the core of this study.¹⁶

3.4.2 Information at the beginning of period t

In every period t, individuals observe their individual outcomes and payoffs as well as the macroeconomic variables - output gap, inflation and interest rate - up to period t-1, as well as all their predictions including the one for period t made at time t-1. The information was both shown in a history table and depicted graphically. (see Figure 1 for an exemplary screenshot)

3.5 Payoff

Participants are rewarded according to the absolute distance of their forecast to the actual optimal outcome in the next period. The payoff function is of the same functional form as in Adam (2007), Pfajfar and Žakelj (2014) and Assenza et al. (2014):¹⁷

$$Score = \frac{100}{1 + |\text{Forecast error}|} \quad (34)$$

where the forecast error is $\hat{E}_t^i \bar{v}_{t+1}^i - \bar{v}_{t+1}^i$ for the advisors to households and $\hat{E}_t^j p_{t+1}^{*j} - p_{t+1}^{*j}$ for the advisors to firms. Subjects were informed about the functional form of the payoff function and were provided with a graph and a table in the instructions, illustrating the relationship between the forecast error and the payoff. Hence, the scores were based on how close participants' forecasts

¹⁵See e.g. Nagel and Vriend (1999), Lei and Noussair (2002), Hommes et al. (2005), Hommes et al. (2008), Heemeijer et al. (2009), Bao et al. (2013), Pfajfar and Žakelj (2014)

¹⁶Arifovic and Petersen (2015) deviate from this approach and provide agents with the model equations, but their focus is on comparing different policies in the laboratory and whether agents learn the REE is not a central question in their study.

¹⁷This payoff function is to be preferred to a quadratic distance metric. Quadratic distance functions become very flat for small forecast errors so that subjects have little incentives to think about small fluctuations in the variables that they need to forecast.

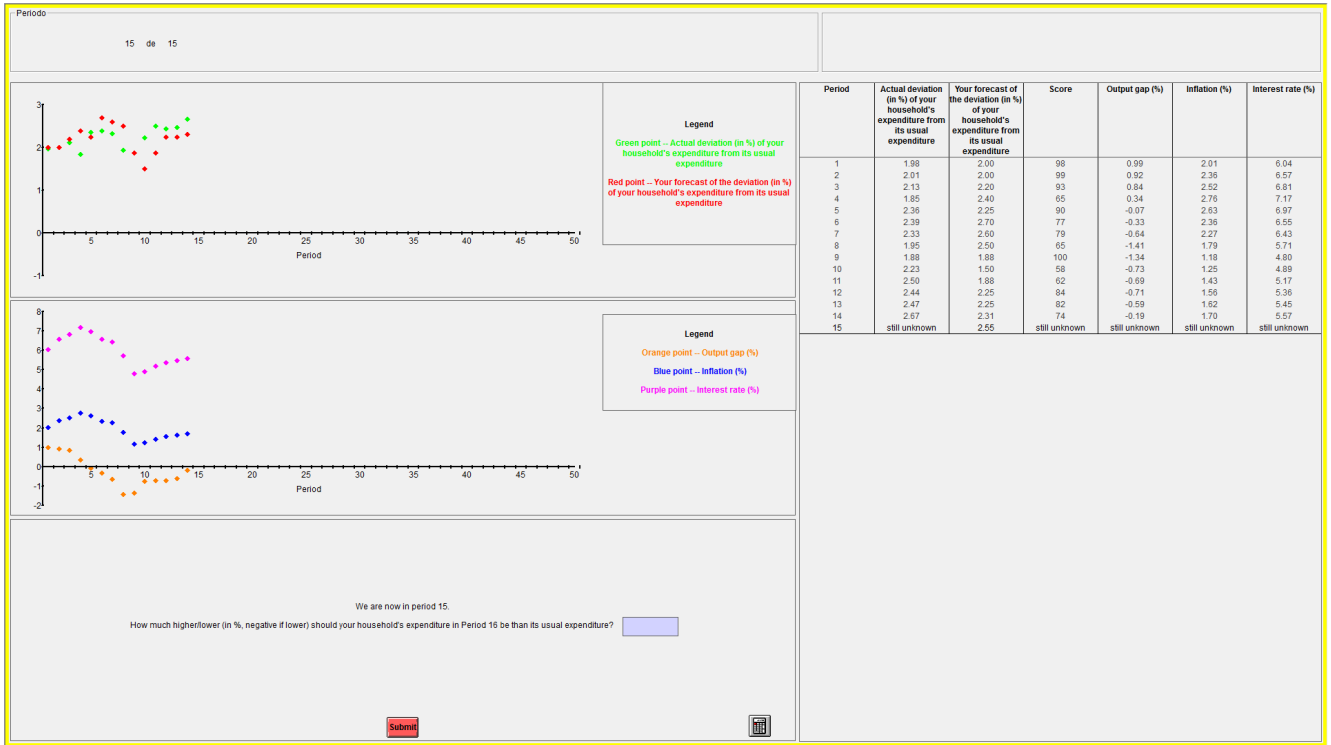


Figure 1: Computer screen for a household advisor with graphed time series and corresponding history table

were to the actual outcomes in the next period. The payoff units were fictitious and converted to euros according to a fixed exchange rate which was 1 euro per 150 points. In addition, participants received a fixed show-up fee of 4 euros.¹⁸ Average earnings per session ranged between 15 and 25 euros.

4 Treatments

Treatment 1 is based on the structural model with (31) and $\phi_\pi = 0.5$ so that the Taylor principle is not fulfilled.¹⁹

Treatment 2 is based on the structural model with system (32) and $\phi_\pi = 1.5$ so that the Taylor principle is satisfied. $\phi_\pi = 1.5$ is the original coefficient found by Taylor (1993) and, moreover, the coefficient used by Pfajfar and Žakelj (2014) and Assenza et al. (2014).

Treatment 3 uses the structural model with system (33) and $\phi_\pi = 3$. The motivation for this treatment is twofold: Firstly, Clarida et al. (2000) find that US monetary policy in periods of

¹⁸The fixed payoff in treatment 1 had to be adjusted to 14 euros to guarantee an adequate reward for two-and-a-half-hours in the experiment.

¹⁹I refrain from using ϕ_π close to 1, since as $\phi_\pi \rightarrow 1$ the (solved) equation for households becomes $v_t^i = 0.01 \int_{i=0}^1 \hat{E}_t^i \bar{v}_{t+1}^i di + 0.99 \hat{E}_t^i v_{t+1}^i$. Hence, there would be disproportionate weight on the individual forecast, while the median forecast would hardly matter.

macroeconomic stability reacts more strongly to inflation with empirical estimates of ϕ_π ranging from 2.15 to 3.13, depending on the subsample and the planning horizon of the central bank.²⁰ The second reason is the lacking evidence for convergence with $\phi_\pi = 1.5$. The treatments are summarized in Table 2, also showing the number of observations (groups) collected for each treatment:

Treatment	ϕ_π	no. observ.	avg. earnings p (v) in points
1	0.5	2	319 (373)
2	1.5	6	2214 (2537)
3	3	5	2905 (2537)

Table 2: **Treatments**

5 Results

5.1 Aggregation over groups

Figure 2 shows the averages of inflation and output gap over the different experimental groups. The following observations stand out:

Observation 1. Not fulfilling the Taylor principle (Treatment 1 with $\phi_\pi = 0.5$) leads to divergence.

Observation 2. With the Taylor principle barely satisfied (Treatment 2 with $\phi_\pi = 1.5$), inflation tends to remain above the RE steady state and there is no evidence that the economy converges within 50 periods.

Observation 3. With particularly aggressive monetary policy (Treatment 3 with $\phi_\pi = 3$), the economy more likely converges to the RE steady state within 50 periods and tends to be similarly stable to the reduced-form experiment.

5.2 Results at the group level

Analyzing single experimental groups could be important, since large heterogeneity between them is an indicator that the economic dynamics under this particular monetary policy regime exhibit high uncertainty.

²⁰Target horizons play no role in this experimental study. One reason why central banks adopt a target horizon is that monetary policy affects macroeconomic variables with some lag. (Clarida et al. (2000), p.160) However, by design, in the experimental economy, monetary policy contemporaneously affects macroeconomic variables so that there is no good reason to engage in forward-looking behavior for the central bank.

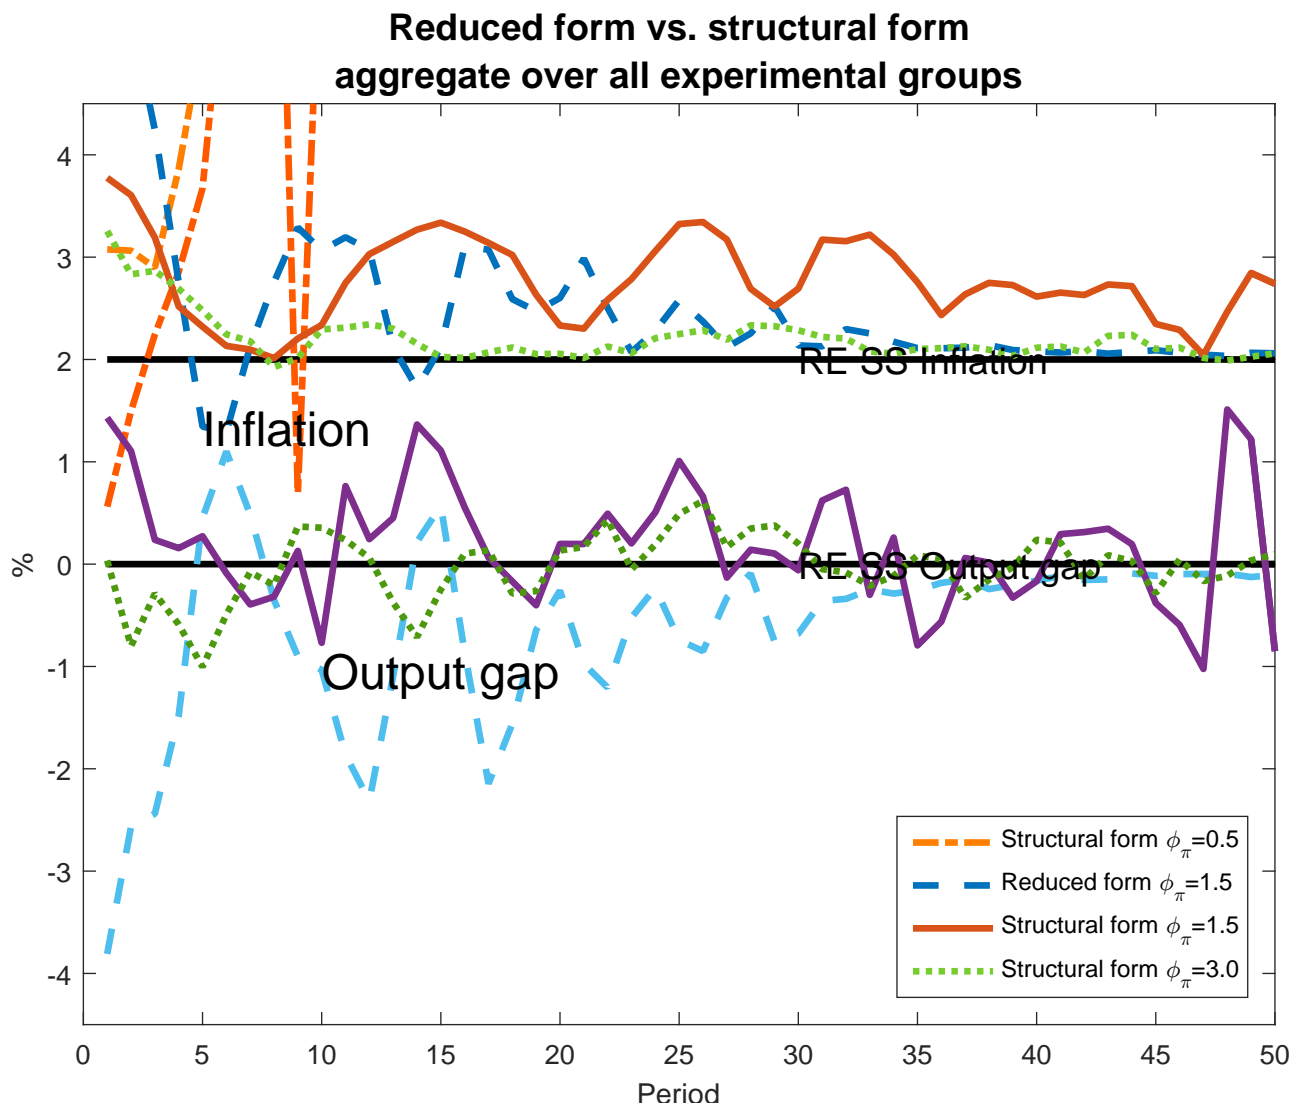


Figure 2: Aggregation (by mean) over experimental groups

Observation 4. For low Taylor rule coefficients, i.e. $\phi_\pi = 0.5$ or $\phi_\pi = 1.5$, there is considerable inter-group heterogeneity, while the heterogeneity tends to vanish with more aggressive monetary policy, i.e. $\phi_\pi = 3$.

Figure 3 shows that the dynamics differ between the two groups in treatment 1 with $\phi_\pi = 0.5$: in group 1, aggregate outcomes follow an explosive path and never return to the vicinity of the REE. In group 2, the path of macroeconomic variables displays extreme fluctuations. Bootstrapped t-tests strongly reject the null hypothesis that inflation equals its target of 2 % for both groups. (p-values: 0.0000 (Gr. 1); 0.0000 (Gr. 2))

Also figure 4 shows considerable heterogeneity between the six groups in treatment 2 with $\phi_\pi = 1.5$. Bootstrapped t-tests reject the null hypothesis that inflation equals its target of 2 % in the last ten periods of the experiment for five out of the six groups. (p-values: 0.0411 (Gr. 1); 0.0006 (Gr. 2); 0.0257 (Gr. 3); 0.0229 (Gr. 4); 0.9306 (Gr. 5); 0.0002 (Gr. 6))




Figure 3: Outcomes with $\phi_\pi = 0.5$

In contrast, figure 5 shows less distinct heterogeneity for the five groups of treatment 3 with $\phi_\pi = 3$, as the outcomes are all fairly stable and close to the RE steady state. This is similar to the reduced-form experimental design. (Appendix 9.2.3 for the two reduced-form groups) Bootstrapped t-tests do not reject the null hypothesis that inflation equals its target of 2 % in the last ten periods of the experiment for three out of the five groups. (p-values: 0.0002 (Gr. 1); 0.0005 (Gr. 2); 0.9938 (Gr. 3); 0.1324 (Gr. 4); 0.0726 (Gr. 5))

In some groups, the output gap exhibits some instability and displays persistent fluctuations. Yet, these fluctuations are considerably lower than in the sessions with $\phi_\pi = 1.5$.

Observation 5. In treatment 3 with $\phi_\pi = 3$, inflation fluctuations tend to be countercyclical with respect to output fluctuations.

This stylized fact can be explained by the higher Taylor rule, reducing the positive feedback in the system.

5.3 Evaluation of the results

Table 3 shows the mean quadratic distance from the REE of each group. The highest distance from the RE steady state is achieved, when the Taylor principle is not satisfied. However, this at best implies that the Taylor principle is necessary for learnability but not that it is sufficient. Treatment 2 with $\phi_\pi = 1.5$, where the Taylor principle is satisfied, still does not provide evidence for convergence to the RE steady state.

Only if the monetary authority adopts a Taylor principle with $\phi_\pi = 3$, the quadratic distance

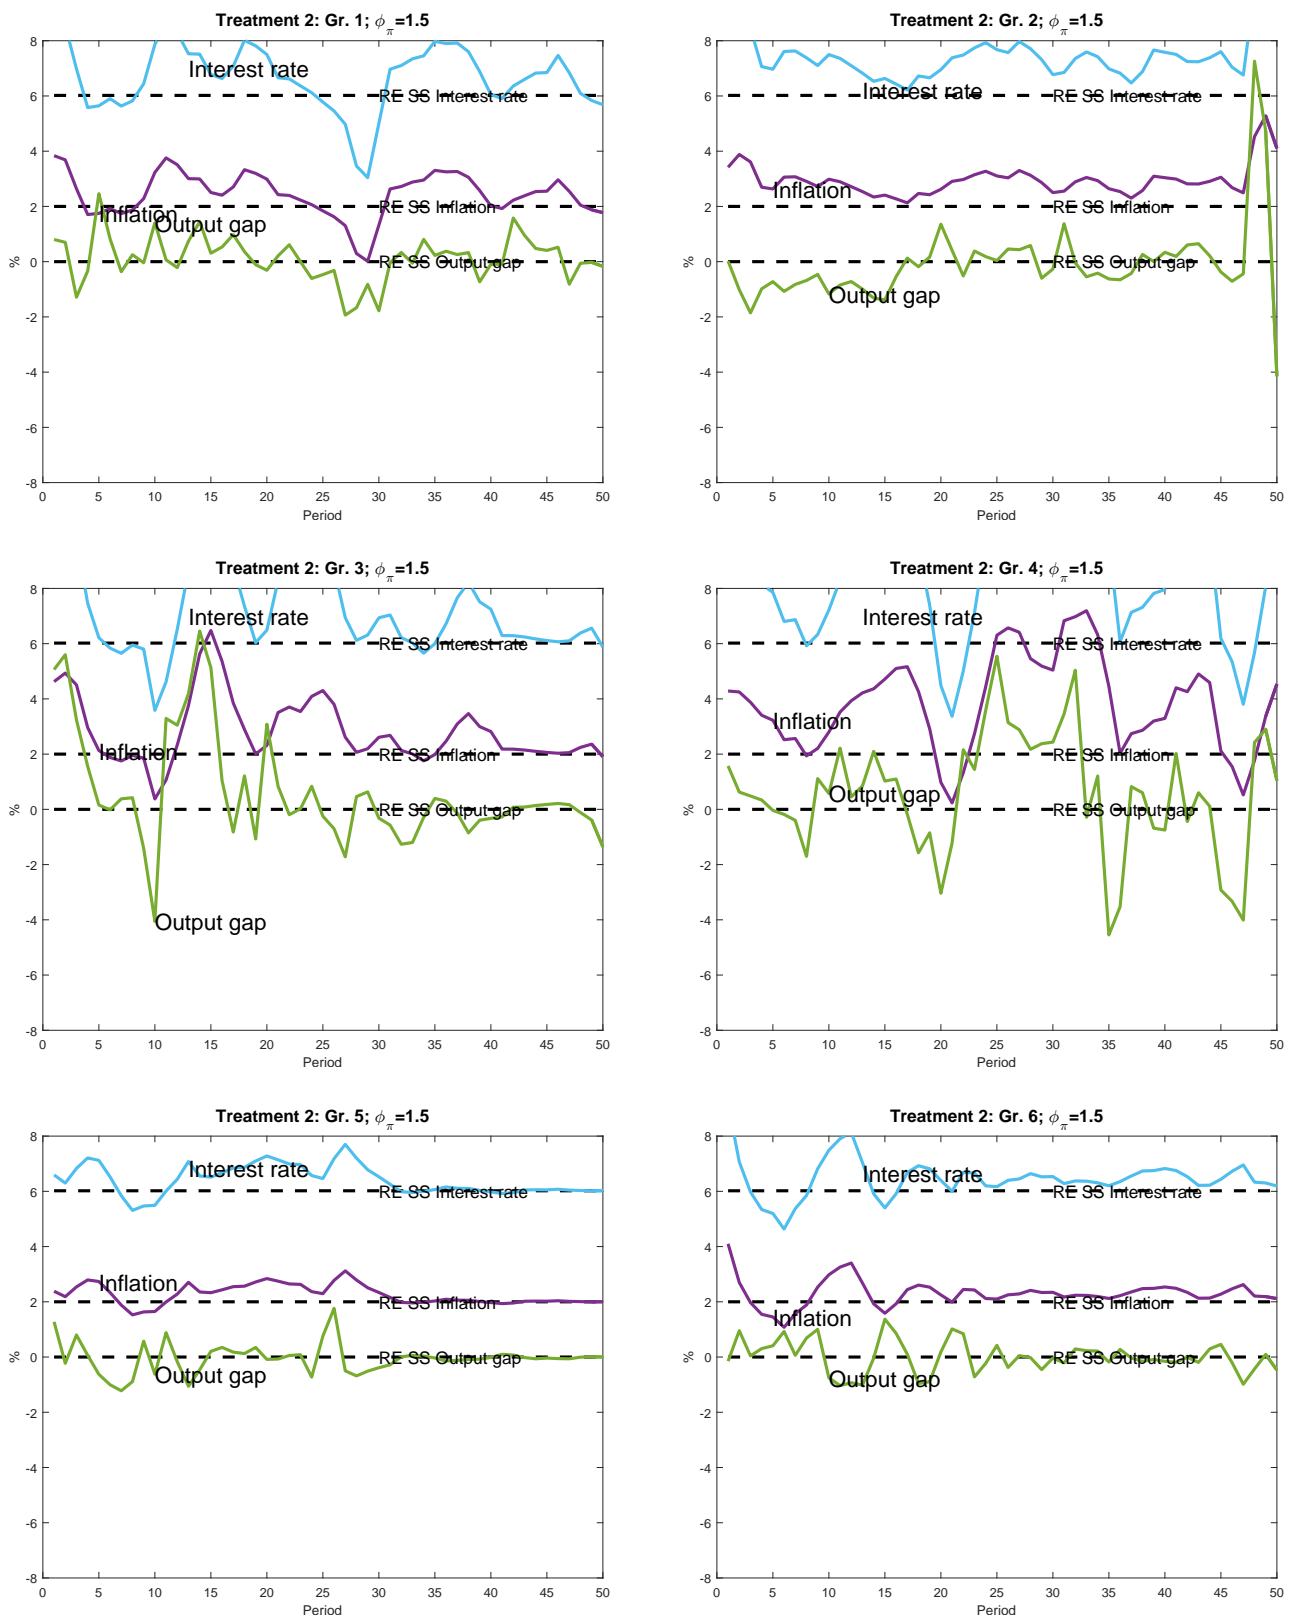


Figure 4: Outcomes with $\phi_\pi = 1.5$

from the REE becomes sufficiently small so that in this case the REE provides a good approximation to actual economic outcomes. A Wilcoxon-Mann-Whitney rank-sum test shows a significant difference for the average quadratic difference between inflation in treatments 2 and 3 (p-value: 0.0285), while for the output gap the difference between treatments 2 and 3 exhibits borderline

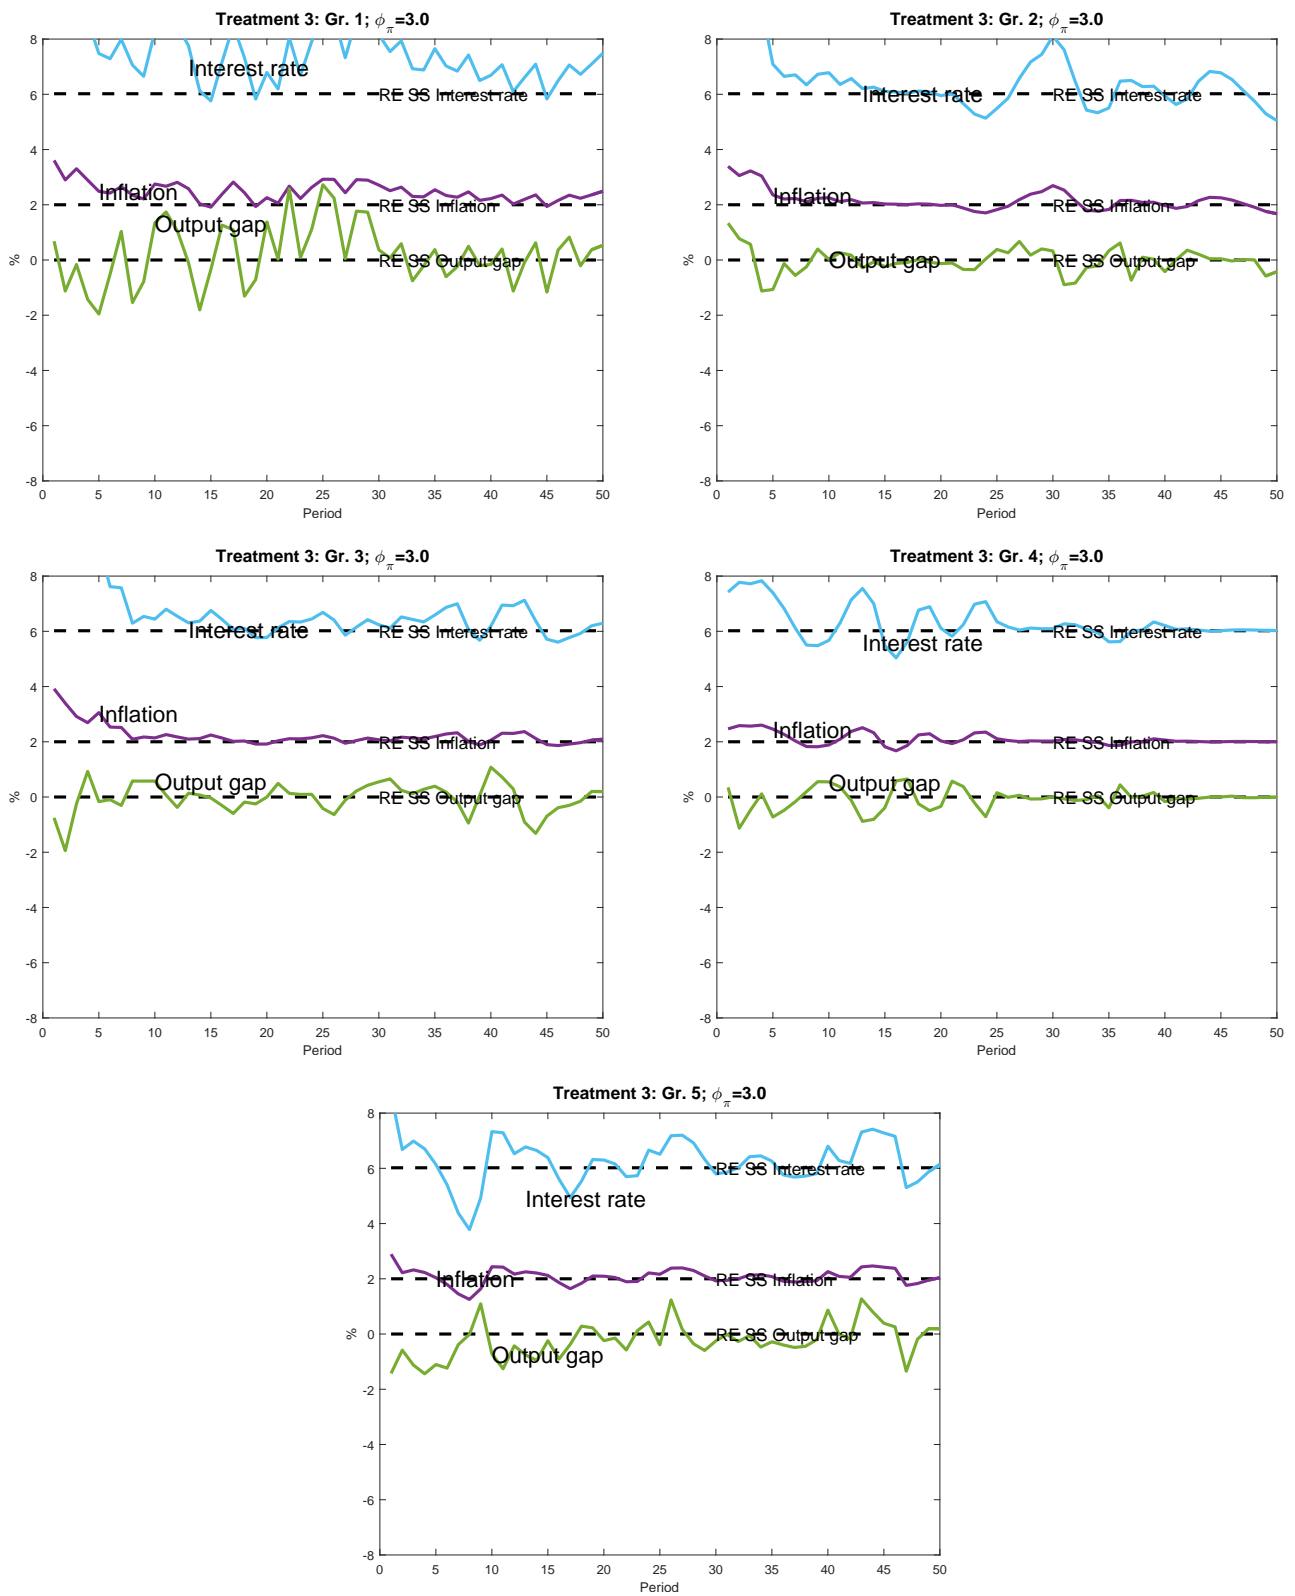


Figure 5: Outcomes with $\phi_\pi = 3$

significance (p-value: 0.1003).

Treatment	ϕ_π	Group	Inflation	Output gap
1	0.5	1	76524.54	154036
1	0.5	2	37769.78	19311.5
2	1.5	1	0.81	0.68
2	1.5	2	1.21	2.35
2	1.5	3	2.11	4.42
2	1.5	4	6.21	4.98
2	1.5	5	0.19	0.29
2	1.5	6	0.32	0.33
3	3	1	0.35	1.22
3	3	2	0.16	0.21
3	3	3	0.19	0.30
3	3	4	0.05	0.15
3	3	5	0.08	0.47

Table 3: Average quadratic distance from the RE steady state

5.4 Comparison to real macroeconomic data

The inquisitive reader might wonder how these experimental results of the structural-form experiments are consistent with previous empirical results using real data. Most prominently, Taylor (1999) finds that the economy is in fact stabilized under reaction coefficient of $\phi_\pi = 1.5$. This section argues that the experimental results are not inconsistent with empirical findings with macroeconomic data.

In a first step, I replicate Taylor's (1999) results by using quarterly data from 1987 Q1 to 1997 Q3 of the federal funds rate i_t , the inflation rate as the percentage change of the GDP deflator over the past four quarters, π_t and an output gap measure \hat{y}_t using detrended output with a Hodrick-Prescott filter (Hodrick and Prescott, 1997) and smoothing value 1600. The data were downloaded from the Federal Reserve, as Taylor's (1999) original datasource is outdated and unavailable. The OLS results are consistent with the original results:

$$i_t = 1.76 + 1.53\pi_t + 0.76\hat{y}_t + u_t \quad R^2 = 0.68 \quad (35)$$

(2.61) (6.32) (3.81)

where u_t is the residuals and the numbers in brackets presents t-ratios. However, as shown in Figure 6, Taylor's sample starts in a period, where the economy is stabilized with inflation being 2.10 %. Since at the start of the experiment in period zero no previous information is available, the natural benchmark for the experiment is a sample starting with a period of a substantial structural change. Many economists believe that the largest change in the policy stance of the

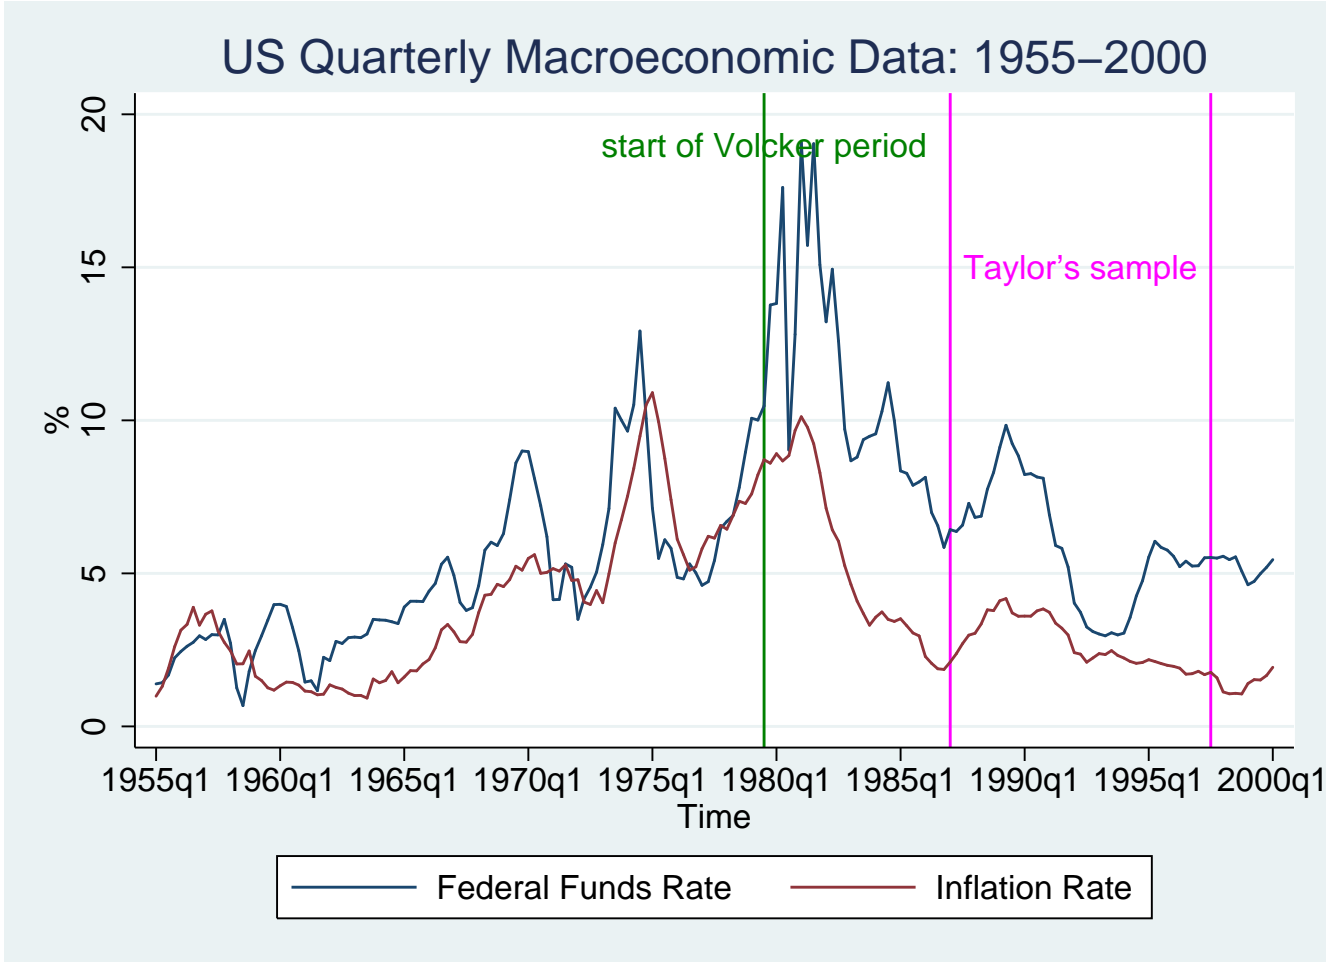


Figure 6: US quarterly macroeconomic data from 1955 to 2000

Federal Reserve occurred when Paul Volcker was appointed Chairman in 1979, whose policy was continued by Alan Greenspan (Clarida et al., 2000). Figure 6 also illustrates that the inflation dynamics resemble the dynamics of the experimental data, as there is first a downward trend and more stability in later periods. I first examine whether the monetary policy regime during the whole Volcker-Greenspan era could be described as using a Taylor rule with reaction coefficient $\phi_\pi = 1.5$. The OLS-regression above for the sample from 1979 Q3 to 1997 Q3 yields:

$$i_t = 2.70 + 1.38 \pi_t + 0.20 \hat{y}_t + u_t \qquad R^2 = 0.77 \qquad (36)$$

(6.38) (14.71) (1.21)

where the output-gap is insignificant. The hypothesis that the coefficient on inflation is 1.5 could

not be rejected (p-value: 0.2080) However, it would be misleading to compare the experimental data to the whole sample because of different initial values: the values in 1979 Q3 are not matched by the first period values of the experiment. 1987 Q1 corresponds to 31 periods after 1979 Q3 so that a valid comparison is between Taylor’s sample and the experimental data from period 31 onwards. Following the early macroeconomic literature which compares simulation results of simple macroeconomic models to real data, the second moments are compared. The results are displayed in table 4 and show that the Structural form model performs better than the reduced form model. Another question that may arise is why the monetary policy regime before 1979 that

	Standard deviation of		
	Inflation	Output gap	Interest rate
<u>US quarterly data</u>			
1987:1-1997:3	<i>0.78</i>	<i>0.94</i>	<i>1.94</i>
<u>Experiments (average over all experimental groups periods 31-50)</u>			
Reduced form	0.07	0.1	0.11
Structural form	0.64	1.07	0.95

Table 4: Comparison between macroeconomic data 1987:1-1997:3 and experimental data periods 31-50

was found to be passive by Taylor (1999) and Clarida et al. (2000) is not accompanied by the hyperinflationary episodes found in treatment 1 ($\phi_\pi = 0.5$). Figure 6 clearly shows that inflation was on an increasing trend pre-1979. Secondly, many measures in the real world to prevent hyperinflation are beyond the scope of the model used for this experiment. Examples include establishing anti-inflation credibility through exchange rate pegs, which may have been relevant for the US before it left Bretton Woods in 1971, and the role of institutional factors in anchoring expectations (Sargent, 1986, 1994).

6 Modeling expectation formation

One main objective of this paper is investigating whether the experimental data support an alternative formal modeling device that can be used instead of rational expectations. A quantitative model of expectation formation should ideally be able to forecast the evolution of macroeconomic time series under different coefficients, a different number of players or for a different time length than the ones used in the experimental sessions. Hence, the learning parameters should be stable

over treatments.

There are two general approaches to modeling learning: The first approach is finding a parsimonious model with few parameters that is relatively easy to track and that predicts behavior well *on average*. This view tends to be taken in the macroeconomic learning literature. *Adaptive learning*, a frequently taken approach in macroeconomic contexts, tends to rely on few parameters and on the weak assumption that agents adjust their beliefs in the direction of the observed outcomes. (Marcet and Sargent, 1989; Sargent, 1994; Evans and Honkapohja, 2001) Adaptive learning could be expected to correctly predict the aggregate patterns of non-convergence for 1.5 and convergence for 3 within 50 periods, as faster convergence for a higher Taylor rule coefficient was shown by Orphanides and Williams (2007) and Ferrero (2007) in a reduced-form New-Keynesian framework.

The second approach is building up *behavioral models of heterogeneous expectations*, which are not only supposed to fit aggregate outcomes but also individual outcomes. A way how individual heterogeneity has previously been introduced into learning is that agents choose among a finite set of forecasting models, each of which has an endogenous probability of being chosen. (Brock and Hommes, 1997; Anufriev and Hommes, 2012)

To assess the empirical performance of the learning models, I create *agent-based computational economic (ACE)* path simulations, which is a common methodology to understand the aggregate behavior in laboratory experiments with human subjects. (See Duffy (2006) for a detailed survey) A path simulation uses initial behavior (e.g. the first two observations) from the experimental sessions with the human subjects and subsequently calculates the experimental outcomes for fictitious computerized subjects that are endowed with the learning rules given in the models. The criterion for model assessment is how well, in terms of mean square error to the experimental data, a simulated path over 50 periods created by the model can replicate the time series of the different experimental groups. Additionally, following the idea of Kydland and Prescott (1982), I consider whether the model simulations are able to match a set of aggregate and individual statistics. For the sake of conciseness, the average of the outcomes and other statistics over experimental groups are reported.

6.1 Modeling average behavior: Noisy adaptive learning

Appendix 9.3 shows that adaptive learning both with least squares and with constant gain qualitatively predicts large differences in the convergence patterns between $\phi_\pi = 0.5$, $\phi_\pi = 1.5$ and

$\phi_\pi = 3$. Yet, it would merely predict fluctuations in the presence of exogenous noise. Since in this experimental economy there is no exogenous noise and all shocks come from the subjects, an intriguing question to explore is whether the fit of an adaptive learning model to the data can be improved by specifying noise in subjects' learning processes. This would allow for more randomness in individuals' behavior and thus also address the conceptual concern of heuristic-switching models that subjects' behavior may not be accurately described by a discrete, finite set of rules.

This section develops a simple model that has three building blocks: an adaptive rule as a benchmark; shocks representing randomness and an *endogenous* variance. The model thus adopts and combines concepts from different kinds of literatures: adaptive learning similarly to Marcet and Sargent (1989), Sargent (1994) and Evans and Honkapohja (2001), randomness in behavior as in the microeconomic learning literature (see e.g. Anderson et al. (2004)) and conditional heteroskedasticity as in the financial econometrics literature (Engle, 1982; Bollerslev, 1986). While the exposition is for \bar{v}^i , the procedure is exactly analogous for p^{*j} .

6.1.1 The adaptive specification

There are several possible specifications for adaptive learning:²¹

$$\hat{E}_t^i \bar{v}_{t+1}^i = \hat{E}_{t-1}^i \bar{v}_t^i + \gamma_t (\bar{v}_{t-1}^i - \hat{E}_{t-1}^i \bar{v}_t^i) \quad (37)$$

$$\hat{E}_t^i \bar{v}_{t+1}^i = \hat{E}_{t-2}^i \bar{v}_{t-1}^i + \gamma_t (\bar{v}_{t-1}^i - \hat{E}_{t-2}^i \bar{v}_{t-1}^i) \quad (38)$$

(37) would be the specification, using the standard assumption in the macroeconomic learning literature that agents' perceived law of motion corresponds to the minimum-state variable solution, giving a perceived law of motion corresponding to a constant.²² (37) is the recursive form of agents' estimate of this constant, using observations until period $t-1$.

(38) is plausible in an environment, where agents need to make two-period ahead forecasts such as in the experiment, since \bar{v}_t^i is a function of $E_t^i \bar{v}_{t+1}^i$ and thus cannot be known at the time of

²¹For the sake of parsimony and because section 6.2.1 later demonstrates no evidence that individuals use aggregate variables, I assume that subjects merely consider their individual outcome \bar{v}_t^i in their adaptive learning mechanism.

²²The specification as it is commonly used in the adaptive learning literature (Marcet and Sargent, 1989; Evans and Honkapohja, 2001) would be

$$\hat{E}_t^i \bar{v}_{t+1}^i = \hat{E}_{t-1}^i \bar{v}_t^i + \gamma_t (\bar{v}_t^i - \hat{E}_{t-1}^i \bar{v}_t^i)$$

However, this specification would be inconsistent with the information structure in the experiment, as \bar{v}_t^i is unknown at the time where agents need to submit $\hat{E}_t^i \bar{v}_{t+1}^i$.

being asked to submit $E_t^i \bar{v}_{t+1}^i$. (38) states that agents adjust their forecasts in the direction of the last *observed* forecast error, $\bar{v}_{t-1}^i - E_{t-2}^i \bar{v}_{t-1}^i$. Thus, it corresponds to an explicit specification of the “directional learning hypothesis” (Selten and Stoecker, 1986; Selten, 1998), which has frequently been used in the microeconomic learning literature (see e.g. Nagel (1995); Anderson et al. (2004)) and which proposes that agents tend to shift their decisions in the direction of a best response to recent outcomes.²³

Another degree of freedom is the gain sequence, as one could adopt a decreasing gain scheme such as $\gamma_t = (t - 1)^{-1}$, a constant gain learning scheme such as $\gamma_t = \bar{\gamma}, \forall t$ or an endogenous gain as for instance proposed by Marcet and Nicolini (2003).

When introducing exogenous noise with endogenous variance as described below, it turns out that with constant gain (38) provides a better fit to the data ($R^2 = 0.27$) than (37) ($R^2 = 0.01$). Appendix 9.6 shows that a decreasing gain delivers a much worse fit for Treatment 2 ($\phi_\pi = 1.5$) than constant gain, and endogenous gain only yields improvements for Treatment 1 ($\phi_\pi = 0.5$) but does not seem preferable, as it renders the model considerably more intractable.

6.1.2 Introducing randomness

I first investigate how forecasting behavior of individuals in the experiment differs from a forecaster using the adaptive rule (38). I define the forecast of this adaptive forecaster, $\hat{E}_t^{i,ADA} \bar{v}_{t+1}^i$, as²⁴

$$\hat{E}_t^{i,ADA} \bar{v}_{t+1}^i = \hat{E}_{t-2}^{i,ADA} \bar{v}_{t-1}^i + \gamma(\bar{v}_{t-1}^i - \hat{E}_{t-2}^{i,ADA} \bar{v}_{t-1}^i) \quad (39)$$

and the deviation of the actual forecast submitted by the subject from the adaptive forecast, being the shock to adaptive learning, as

$$\eta_t^i \equiv \hat{E}_t^i \bar{v}_{t+1}^i - \hat{E}_t^{i,ADA} \bar{v}_{t+1}^i \quad (40)$$

Figure 7 shows the distributions for η_t^i over the range -20 to 20.²⁵

Observation 6. As shown in figure 7, η_t^i is approximately normally distributed.

²³The non-recursive form of (38) would imply that agents merely consider the even or the odd outcomes respectively. This can be reconciled by the behavioral assumption of *imperfect recall*. (See e.g. Piccione and Rubinstein (1997), Molavi et al. (2015)) I assume that agents treat $\bar{v}_{t-1}^i - \hat{E}_{t-2}^i \bar{v}_{t-1}^i$ as a sufficient statistic for all information available to them without considering how $\hat{E}_{t-2}^i \bar{v}_{t-1}^i$ was formed.

²⁴The gain $\bar{\gamma}$ is calibrated as 0.2, being the maximum likelihood estimate obtained in Section 6.1.3.

²⁵This range comprises over 90 % of all observations in treatments 2 and 3, but less than 50 % of all observations for treatment 1.

Observation 7. The variance of the distribution of η_t^i differs considerably across treatments. The lower ϕ_π , the more likely extreme deviations from an adaptive rule occur.

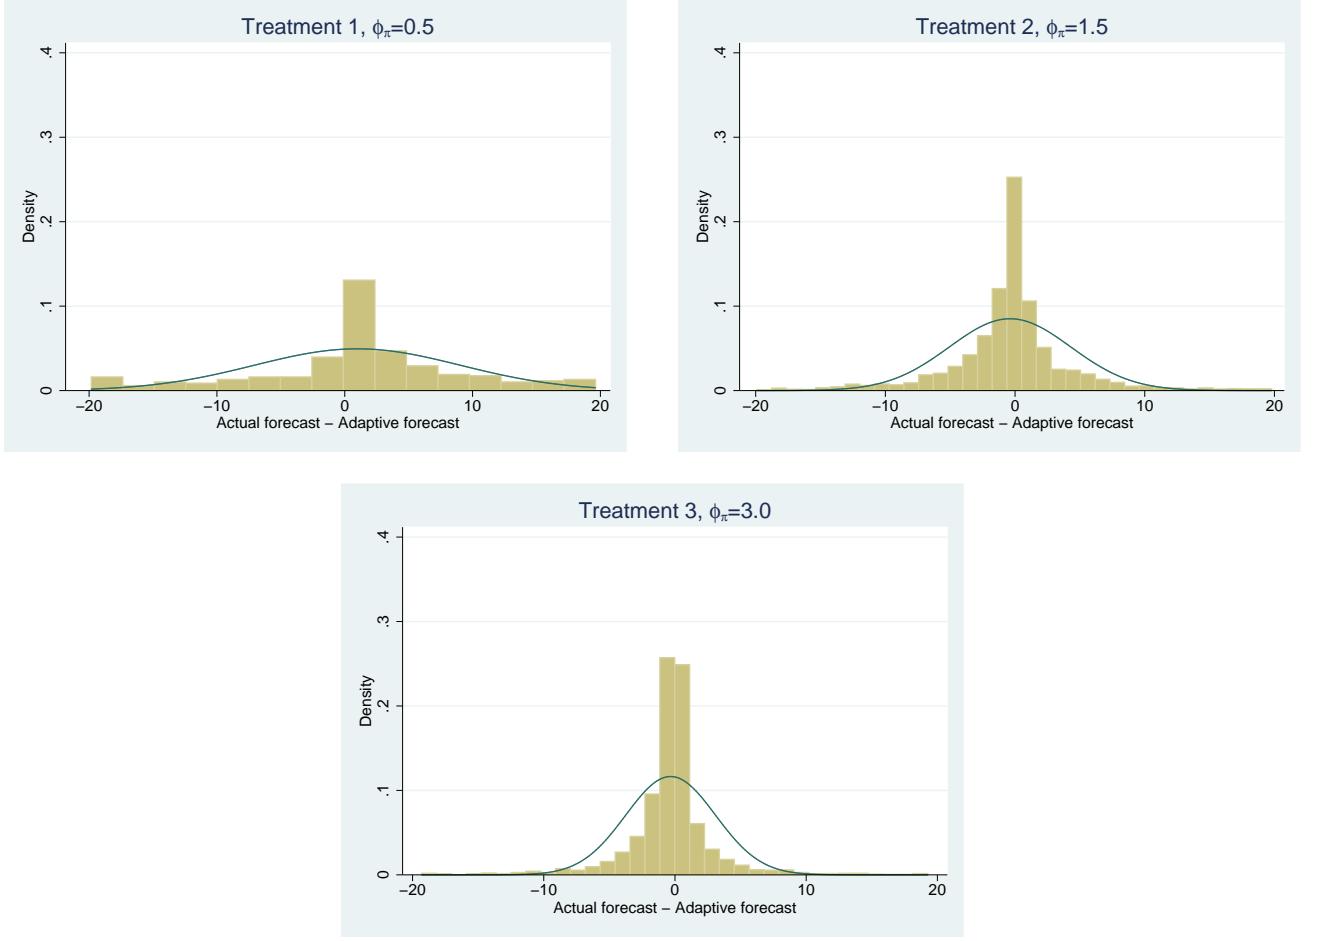


Figure 7: Deviations from adaptive rule with a constant gain of 0.2

Motivated by observations 6 and 7, I specify the shock $\eta_t^i \sim N(0, \sigma_{t,i}^2)$ as a random draw from a normal distribution with an *endogenous* variance that depends on the forecast errors that would have been achieved with an adaptive rule:²⁶

$$\sigma_{t,i}^2 = \omega \sum_{j=1}^{t-1} (1 - \omega)^{j-1} (\bar{v}_{t-j}^i - \hat{E}_{t-j-1}^{i,ADA} \bar{v}_{t-j}^i)^2 \quad (41)$$

$$= (1 - \omega) \sigma_{t-1,i}^2 + \omega (\bar{v}_{t-1}^i - \hat{E}_{t-2}^{i,ADA} \bar{v}_{t-1}^i)^2 \quad (42)$$

(42) postulates that if an adaptive rule performs well (poorly), subjects follow an adaptive rule

²⁶I considered the following alternative way of introducing shocks

$$\begin{aligned} \text{adaptive specification: } \hat{E}_t^i v_{t+1}^i &= \hat{E}_{t-2}^i v_{t-1}^i + \gamma(v_{t-1}^i - \hat{E}_{t-2}^i v_{t-1}^i) + \eta_t^i \\ \text{shock variance: } \sigma_{t,i}^2 &= (1 - \omega) \sigma_{t-1,i}^2 + \omega (v_{t-1}^i - \hat{E}_{t-2}^i v_{t-1}^i)^2 \end{aligned}$$

This approach is considerably less successful in fitting the experimental data no matter whether (37) or (38) is used as the main specification, as σ^2 would have a unit root. See the proof in the appendix 9.7.

more (less) closely, as deviations from an adaptive rule tend to be small (large). In other words, if an agent performs poorly despite adapting her forecast in the direction of the observed error, this agent engages in experimentation by “jumping” to another point in space.

Hence, this approach can be considered a more tractable simplification of the heuristic-switching model by Brock and Hommes (1997) and Anufriev and Hommes (2012). The difference from Anufriev and Hommes (2012), who model deviations from an adaptive rule as switching to different forecasting rules, is that I consider *random* deviations from the adaptive rule. Noisy adaptive learning not only predicts the different dynamic patterns in treatment 2 ($\phi_\pi = 1.5$) and treatment 3 ($\phi_\pi = 3$) but was also successfully applied to other setups: Appendix 9.2.4 shows that noisy adaptive learning reproduces the convergence result of the reduced-form setup and Mauersberger et al. (2017) show that noisy adaptive learning well describes the learning dynamics in the experimental oligopoly market of Nagel and Vriend (1999).

Noisy adaptive learning is also related to endogenous variance models in the financial econometrics literature (Engle, 1982; Bollerslev, 1986). In fact, (42) could be considered a GARCH(1,1) specification, as it includes the first lag of the variance itself and the last observed (i.e. the first lag of the) forecast error, which is the explanatory variable in the specification of the conditional mean of a subject i 's forecast.

6.1.3 Estimation

To simulate this model, $\hat{E}_0^{i,ADA} \bar{v}_1^i$ and $\hat{E}_0^{i,ADA} \bar{v}_2^i$ have been initialized as $\hat{E}_0^i \bar{v}_1^i$ and $\hat{E}_0^i \bar{v}_2^i$, i.e. the first two expectations submitted by each subject in the experiment. In period 2, $\hat{E}_2^{i,ADA} \bar{v}_3^i$ is calculated according to (39) and the variance is initialized according to (42) as $\sigma_2^{2,i} = \omega(\bar{v}_1^i - \hat{E}_0^i \bar{v}_1^i)^2$. Following a standard approach in the experimental (see e.g. Stahl (1996), Roth and Erev (1998), Camerer and Ho (1999)) and macroeconomic literature (see e.g. McGrattan et al. (1997)), the parameters γ , and ω are estimated by maximum likelihood using the individual data of all treatments, which gives parameter estimates of $\hat{\gamma} = 0.20$ and $\hat{\omega} = 0.59$.²⁷ Figure 8, depicting the experimental data and the mean of 6,000 replications of the simulation, shows that the noisy learning model captures the differences in the speed of convergence between treatment 2 ($\phi_\pi = 1.5$) and treatment 3 ($\phi_\pi = 3$) and the explosive patterns in treatment 1 ($\phi_\pi = 0.5$.) Appendix 9.2.4 shows that noisy

²⁷Estimating the model under the hypothesis of an endogenous variance gives different gain parameters from the baseline case. One explanation is given by Engle (1982), showing that maximum likelihood, jointly estimating the parameters of the mean and variance specifications, gives more efficient parameter estimates.

learning can also replicate the convergence in the reduced-form design with $\phi_\pi = 1.5$. Single replications with the noisy learning model are shown in Appendix 9.5.2.

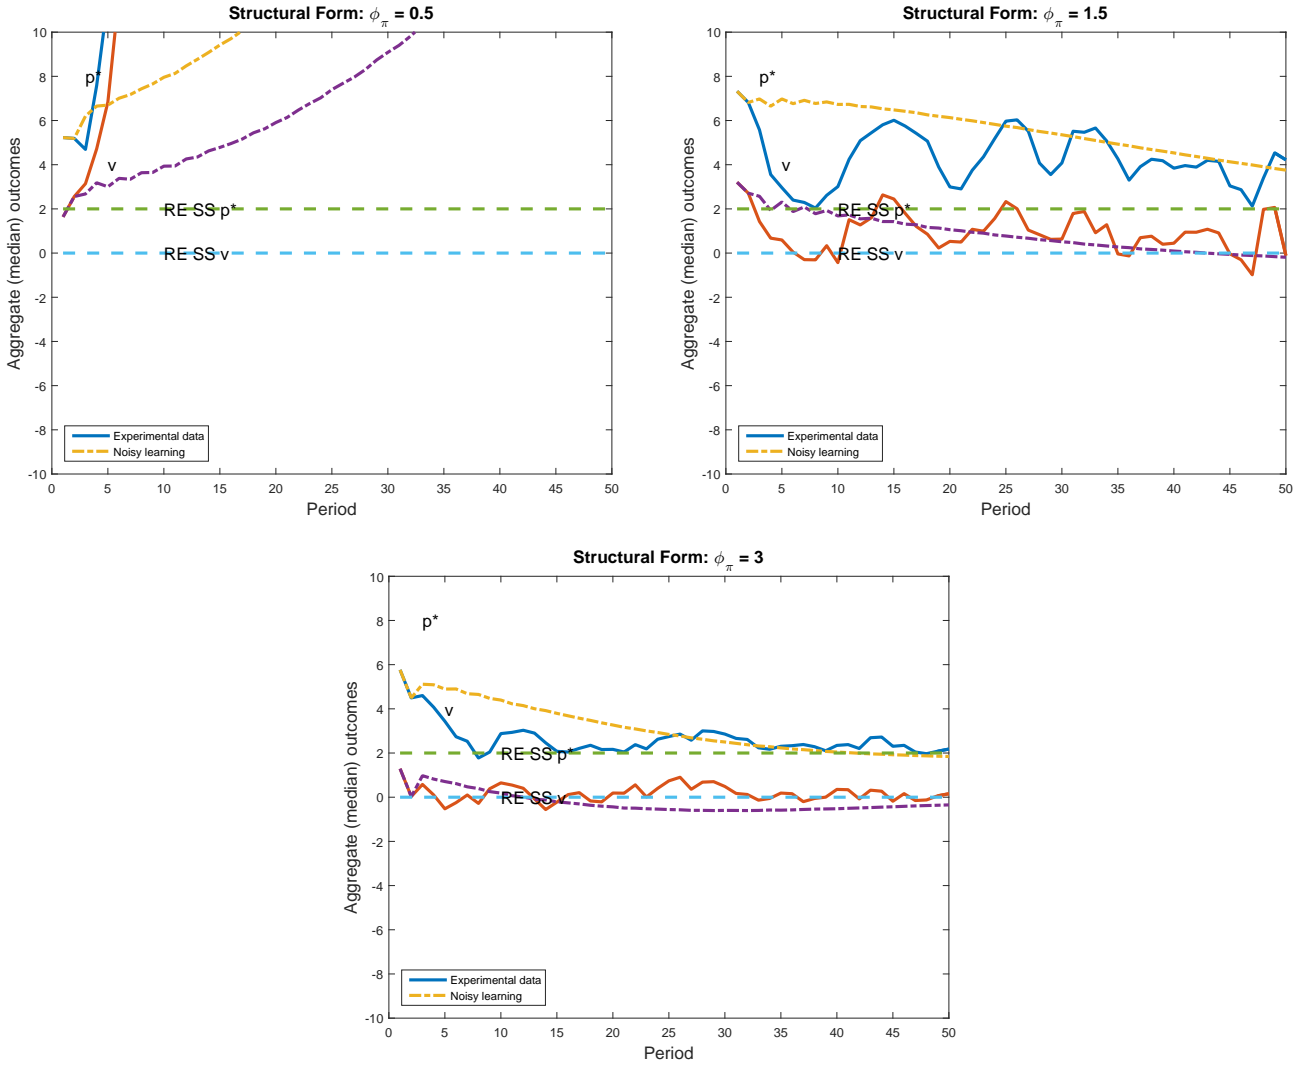


Figure 8: Actual (mean over all experimental groups) outcomes and simulated (mean over 6,000 replications) aggregated outcomes using noisy learning

6.2 Modeling average and individual behavior: reinforcement

Appendix 9.4.1 shows that the heuristic-switching model by Brock and Hommes (1997) predicts too fast convergence for $\phi_\pi = 1.5$. Hence, it is worthwhile to reconsider the single components of this model: The steps are 1. evaluating which information individuals use in this more complex setup than the reduced form (Section 6.2.1); 2. investigating which forecasting rules individuals use (Section 6.2.2); 3. describing the model of endogenous switching (Sections 6.2.3 and 6.2.4) in which these rules are used.

6.2.1 Which information do individuals use?

The approach taken to analyze which information individuals use is running OLS-regressions on all information available to subjects. One challenge to these regressions is that in the structural form, by construction, the individual outcomes are perfectly linear functions of both the individual forecasts and the aggregate outcomes, which results in strong multicollinearity when using regressions. This is not only apparent for regressions at individual level (sample size: 50) but even for regressions on individual data at group level (sample size: 600=12 subjects·50 periods), forecaster type level (sample size: 1800 / 1500 / 600=6 / 5 / 2 groups · 6 subjects·50 periods) and even treatment level (sample size: 3600 / 3000 / 1200=6 / 5 / 2 groups · 12 subjects·50 periods), where estimation requires omitting regressors due to collinearity. Providing enough variation in the data to disentangle the effect of different variables required pooled panel vector autoregressions (VARs) for both forecasting groups (firms and households) over several treatments:

$$\begin{aligned} \hat{E}_t^i x_{t+1}^i = & \alpha + \sum_{k=1}^3 \beta_k x_{t-k}^i + \sum_{l=1}^3 \gamma_k x_{t-l}^{i,av} + \sum_{m=1}^3 \mu_m \hat{E}_{t-m} x_{t-m+1}^i + \sum_{n=1}^3 \nu_n \pi_{t-n} \\ & + \sum_{o=1}^3 \omega_o \bar{v}_{t-o} + \sum_{p=1}^3 \xi_p i_{t-p} + \eta_{t,i} \end{aligned} \quad (43)$$

so that subjects' forecasts of their individual variables $x_t^i = \{\bar{v}_t^i, p_t^{*j}\}$ are regressed on the first three lags of all potential available information: the individual outcome, the sample mean of the individual outcomes $x_{t-l}^{i,av}$ until time $t-l$, the individual forecast, inflation, aggregate expenditure and the interest rate.²⁸ Standard errors were clustered at the subject level. Besides the baseline pooled OLS-regression over all treatments, I provide the following robustness checks: a) inclusion of individual fixed effects²⁹ and b) omitting the unstable treatment 1 ($\phi_\pi = 0.5$). Table 5 reports the results.

In all specifications, the means of the past outcomes are highly significant. Due to the high collinearity, F-tests for joint significance have been conducted on the null hypothesis that the aggregate variables have no predictive power for individual forecasting behavior. Only the second lag of the interest rate is significant at the 5 % level in regression 2 and the third lag exhibits

²⁸Panel regressions often include time fixed effects to capture omitted factors that are common in the period. Since in this application all information that subjects receive is controlled for, the interpretation of these time dummies would become dubious.

²⁹Nickell (1981) shows that the fixed effects estimator is inconsistent when the strict exogeneity assumption is violated like in this application with a dynamic panel. However, as the number of time periods is large, the bias is small.

VARIABLES	(1) Tr. 1-3	(2) Tr. 1-3	(3) Tr. 2-3	(4) Tr. 2-3
<i>Dep. variable: Forecast</i>				
Outcome (lag 1)	-0.0640	0.0523	0.955**	0.939**
Outcome (lag 2)	-0.224	-0.212	-1.213	-1.080
Outcome (lag 3)	0.0705	0.150	0.322	0.484
Mean past outcome (lag 1)	9.962***	8.534***	11.29***	9.745***
Mean past outcome (lag 2)	-12.59***	-10.49**	-14.14***	-12.34***
Mean past outcome (lag 3)	3.061	2.263	3.114*	2.612*
Forecast (lag 1)	0.432***	0.345**	-0.535	-0.493
Forecast (lag 2)	0.337	0.291	1.321	1.177
Forecast (lag 3)	-0.0685	-0.152	-0.323	-0.470
Inflation (lag 1)	0.360	-0.272	3.410	2.979
Inflation (lag 2)	-0.0679	0.191	-6.918	-6.092
Inflation (lag 3)	0.543	-0.0490	4.675*	3.551
Aggregate \bar{v} (lag 1)	-0.0291	-0.0208	-0.781	-0.773
Aggregate \bar{v} (lag 2)	0.216	0.216	1.818	1.724
Aggregate \bar{v} (lag 3)	-0.0887	-0.0750	-0.763*	-0.809*
Interest rate (lag 1)	-0.0948	0.985	-0.414	-0.127
Interest rate (lag 2)	-0.790	-1.327**	0.395	0.0787
Interest rate (lag 3)	-0.898*	0.195	-0.784	-0.255
Constant	12.28**	9.167**	3.693*	3.590
Fixed effects	NO	YES	NO	YES
Observations	7,332	7,332	6,204	6,204
R^2	0.795	0.605	0.803	0.743
Number of subjects	156	156	132	132

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 5: Regression

borderline significance in regression 1. This, however, seems a non-credible result.

An F-test indicates that the remaining aggregate variables are jointly significant in regression 1 (p-value: 0.0000). This result is not robust if fixed effects are included, as the hypothesis that all aggregate variables except for the significant second-order lag of the interest rate could not be rejected in regression 2 (p-value: 0.2494). Once the unstable treatment 1 ($\phi_\pi = 0.5$) is removed, the aggregate variables are all jointly highly insignificant (p-value in regression 3: 0.3393; p-value in regression 4: 0.3462) Lags one to three of the individual forecasts and outcomes, however, are jointly highly significant independently of the specification.

Eyeballing the data raises concern that behavior cannot be adequately captured by linear forecasting models. A Ramsey RESET test, which detects potential misspecification and structural

breaks, provides strong evidence that this general linear model is misspecified for each specification (p-value: 0.0000), a result which remains robust when the insignificant variables are removed. (p-value: 0.0000)

6.2.2 Individual regressions

Forecasting rules are solely based on individual variables for three reasons: first, there is no compelling evidence in the pooled regression that aggregate variables play any role in determining forecasting behavior; secondly, the inclusion of aggregate and individual variables would cause collinearity in regressions at the subject level; thirdly, the assumption that subjects use heuristics solely based on the variable they are asked to forecast is frequently made in learning-to-forecast experiments. (See e.g. Assenza et al. (2014), Anufriev and Hommes (2012)).

The rules considered are the ones that have been found to be important in previous learning-to-forecast experiments and are described in more detail in table 6:

Rule	Description
adaptive	ADA $x_{t+1}^e = x_t^e + \alpha(x_{t-1} - x_t^e)$
trend-following	TR $x_{t+1}^e = x_{t-1} + \alpha(x_{t-1} - x_{t-2})$
anchor & adjustment	LAA $x_{t+1}^e = 0.5(x_{t-1}^{av} + x_{t-1}) + \alpha(x_{t-1} - x_{t-2})$

$x_t = \{\bar{v}_t^i, p_t^{*j}\}$

Table 6: **Set of heuristics**

The linear model that best describes each subject on average is determined by comparing OLS-regression outcomes for each rule and for each individual. Since in all three rules, one coefficient had to be estimated, R^2 could be used to determine the best linear model for each subject.³⁰ Figure 9 shows considerable heterogeneity in forecasting behavior both across treatments and across roles (household vs. firm advisor.)

Observation 8. In all treatments, households engage more in adaptive behavior than firms.

Observation 9. The impact of the trend-following rule becomes smaller, as the Taylor rule coefficient is increased.

³⁰The use of R^2 is equivalent to the use other measures, such as sum of squared errors (SSE) or mean square error (MSE).

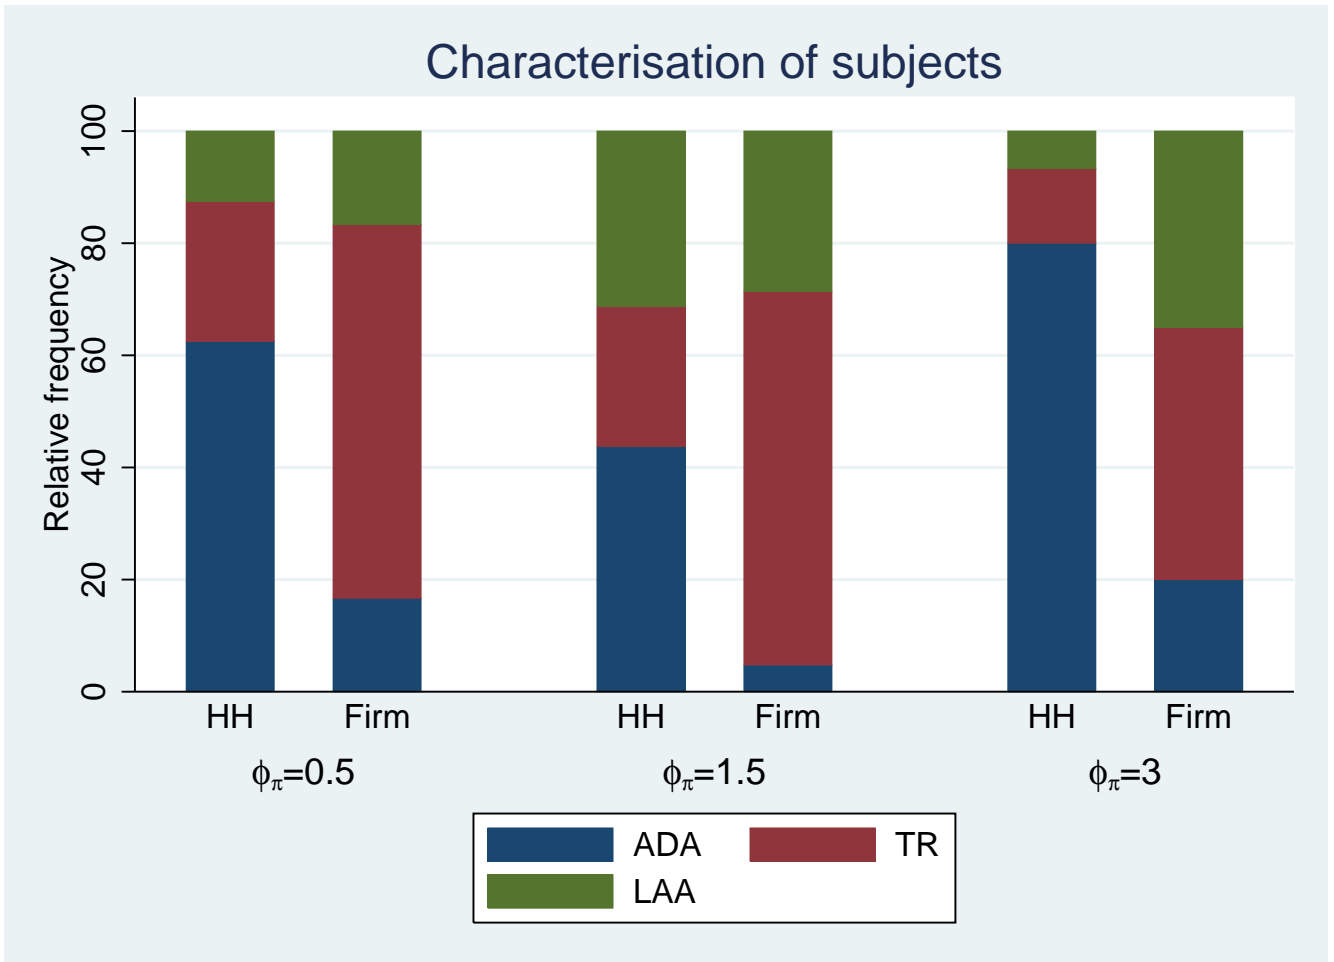


Figure 9: Characterization of the subjects

6.2.3 Endogenous switching

The analysis in Section 6.2.2, assuming that individuals can throughout be represented by the same forecasting model, represents a simplistic view, as previous learning-to-forecast experiments (Anufriev and Hommes, 2012; Bao et al., 2013; Pfajfar and Žakelj, 2014; Assenza et al., 2014) document that switching between different models provides a more accurate description of their decisions. Therefore, Brock and Hommes (1997) and Anufriev and Hommes (2012) endogenize the switching probability by linking it to the past performance of different forecasting rules.

Their model presumes a finite set of forecasting rules H . Given the evidence in 6.2.2, I use the same four heuristics (adaptive, weak trend-following, strong trend-following, anchor and adjustment) as Anufriev and Hommes (2012).³¹ However, one heuristic is added, as it is frequently observed:

Observation 10. *Unchanged behavior*, which means submitting exactly the same forecast as in

³¹I follow Anufriev and Hommes (2012) in obtaining the coefficients of the heuristics from subject-level regressions. While the coefficients for trend-following rules are similar to reduced-form experiments, the gain coefficient in the adaptive rule differed considerably between the structural and the reduced form. While the median gain across all subjects of type ADA in the reduced form replication corresponded roughly to 0.65, which is the estimates of Anufriev and Hommes (2012), in the structural the median gain was approximately 0.3.

the previous period so that $x_{t+1}^e = x_t^e$, is observed 151 times (12.58 % of all observations) in treatment 1 ($\phi_\pi = 0.5$), 364 times (10.11 % of all observations) in treatment 2 ($\phi_\pi = 1.5$) and 216 times (7.20 % of all observations) in treatment 3 ($\phi_\pi = 3$).

Unchanged behavior is a special case of adaptive (ADA) forecasting with $\alpha = 0$. The precise specifications of the five heuristics are given in table 7. While the heuristic-switching model

Rule		Description
adaptive	ADA	$x_{1,t+1}^e = 0.3x_{t-1} + 0.7x_{1,t}^e$
weak trend-following	WTR	$x_{2,t+1}^e = x_{t-1} + 0.4(x_{t-1} - x_{t-2})$
strong trend-following	STR	$x_{3,t+1}^e = x_{t-1} + 1.3(x_{t-1} - x_{t-2})$
anchor & adjustment	LAA	$x_{4,t+1}^e = 0.5(x_{t-1}^{av} + x_{t-1}) + (x_{t-1} - x_{t-2})$
unchanged	UC	$x_{5,t+1}^e = x_{5,t}^e$

$x_t = \{\bar{v}_t^i, p_t^{*j}\}$

Table 7: **Set of heuristics**

correctly predicts the explosive behavior for $\phi_\pi = 0.5$, the model does not accurately predict the distinction between $\phi_\pi = 1.5$ and $\phi_\pi = 3$ as it is observed in the experimental treatments, because the heuristic-switching model predicts relatively fast convergence for $\phi_\pi = 1.5$ to the RE steady state, while in the experiment, in particular for p^* , there is no evidence for convergence to the RE steady state within 50 periods.

Observation 11. The different game structure in the structural form as compared to the reduced form makes it less plausible that subjects consider “fictitious play.”

Heuristic-switching models assume that subjects fix the outcomes in their minds and consider the fictitious *ceteris-paribus* payoffs that unchosen strategies would have yielded. Yet, in the structural form, subjects are informed that their outcome depends not only on aggregate behavior but also on their individual forecasts. From a conceptual point of view, it is therefore not reasonable to hold the outcome fixed and consider the payoff of an alternative strategy. Thus, the rule to which subjects switch may be determined by experimentation rather than by fictitious play, an idea which is referred to as “reinforcement learning” (Roth and Erev, 1998). If this argument is correct, a reinforcement mode should provide a more accurate fit than heuristic-switching.

6.2.4 Specification for reinforcement

There are two related modeling approaches for reinforcement: firstly, introducing a parameter that governs fictitious play into Anufriev and Hommes (2012); secondly, using the original specification

by Roth and Erev (1998).

Table 6.2.4 illustrates the similarities and differences between Anufriev and Hommes (2012) and Roth and Erev (1998). Both models consist of an equation updating the performance of each heuristic h for an individual i . The crucial difference is that Anufriev and Hommes (2012) consider fictitious play so that they do not distinguish whether the heuristic h has been played or not, while Roth and Erev (1998) merely update the performance of heuristics that have been played. To the best of my knowledge, this is the first application of Roth and Erev (1998) to data from a “learning to forecast” experiment.

The performance measure of heuristic h is then used to calculate the probability of playing this particular heuristic h , where the probability of heuristics not being applied is not updated in Roth and Erev (1998). An adaptation of Anufriev and Hommes (2012) to my setup is the individual-superscript. This is unnecessary for the original setups for which the model has been developed, where agents forecast the same aggregate outcomes and $n_{h,t}$ thus depends on identical outcomes being common knowledge. In my setting, subjects forecast individual outcomes, whose realizations are observed in addition to aggregate data.³² I use the same calibration for the learning parameters as Anufriev and Hommes (2012) and Roth and Erev (1998).³³

An idea would be introducing a parameter λ , governing fictitious play, into Anufriev and Hommes (2012). Yet, it turns out that the approach by Roth and Erev (1998) yields a better quantitative fit than Anufriev and Hommes (2012) even when the λ -parameter is introduced. (Details for the simulation of the heuristic-switching model in Appendix 9.4.1). Appendix 9.2.4 shows that even for the reduced-form model reinforcement yields a lower mean square error than the heuristic-switching model for both aggregate variables.

The simulation is initialized by the first two individual outcomes for 6 individuals, two initial attractions for each strategy, which have been set to 0.2 to give each strategy equal weight. With the initial values and the initial probabilities, the outcomes in periods 3 and 4 can be computed.

³²Thus, there are two possible approaches: the first one is interpreting $n_{h,t}^i$ as the *probabilities* of agents using a particular rule, so that $\hat{E}_t^i v_{t+1}^i = v_{h,t+1}^{i,e}$ with probability $n_{h,t}^i$. The second approach is interpreting $n_{h,t}^i$ as the *shares* that agent i attaches to each heuristic h so that $\hat{E}_t^i v_{t+1}^i = \sum_{h=1}^H n_{h,t}^i v_{h,t+1}^{i,e}$. While the latter approach gives deterministic forecasts and outcomes, the former gives stochastic simulations so that each time the simulation is executed the simulated path is different. There is little difference in mean square error over repeated simulations between these approaches.

³³These parameters cannot be estimated by maximum likelihood, as heuristic-switching and reinforcement merely choose between five points for each subject, which would give a likelihood of zero if actual forecasts in the experiment differ from the forecasts predicted by the model. Anufriev and Hommes (2012) find these parameters by “experimentation.”

Heuristic-switching
(Brock and Hommes, 1997)
Anufriev and Hommes (2012)

Reinforcement
(Roth and Erev, 1998)

Updating of performance	
$U_{h,t}^i = \eta U_{h,t-1}^i + \text{Payoff}_{t-1}^i$ Calibration: $\eta = 0.7$	$q_{h,t}^i = \begin{cases} (1 - \phi)q_{h,t-1}^i + \text{Payoff}_{t-1}^i & \text{if rule h used} \\ 0 & \text{otherwise} \end{cases}$ $\phi = 0.1$
probabilities	
$n_{h,t}^i = \delta n_{h,t-1}^i + (1 - \delta) \frac{\exp(\beta U_{h,t-1}^i)}{\sum_h \beta U_{h,t-1}^i}$ Calibration: $\delta = 0.9$ $\beta = 0.4$	$p_{h,t}^i = \frac{q_{h,t}^i}{\sum_h q_{h,t}^i}$
<p><u>idea:</u> reconsider relative weight of fictitious play λ:</p> $U_{h,t-1}^i = \eta U_{h,t-1}^i + (\lambda + (1 - \lambda) \mathbb{1}(\text{h played})) \text{Payoff}_{t-1}^i$	
Anufriev and Hommes (2012): $\lambda = 1$	$\lambda = 0$ I follow Roth and Erev
<p>$\mathbb{1}(\cdot)$ denotes the indicator function, being 1 if the statement in brackets is true and 0 otherwise</p>	

Table 8: Reinforcement

From period 5, each individual expectation and outcomes is fully determined by the simulations. For the starting values of each experimental group, 6,000 replications of the simulation have been conducted. Finally, I average over all experimental groups. Examples of single replications are given in Appendix 9.5.1. Figure 10 depicts the average outcomes over 6,000 replications, showing that the reinforcement model only converges slowly in treatment 2 ($\phi_\pi = 1.5$).

7 Model comparison

7.1 Criteria for model assessment

7.1.1 Aggregate statistics

I use a quadratic loss function between the aggregate outcomes \bar{v} and p^* in each experimental group and the aggregate outcomes \bar{v}^M and p^{*M} created by the simulated paths described above. This procedure is standard both in experimental economics (see e.g. Roth and Erev (1998)) and in the macroeconomic learning literature (see e.g. Orphanides and Williams (2007)), as a central focus in the both literatures are long-run phenomena such as convergence. While the exposition is for \bar{v} , the procedure is exactly analogous for p^* . The mean square error (MSE) by each model

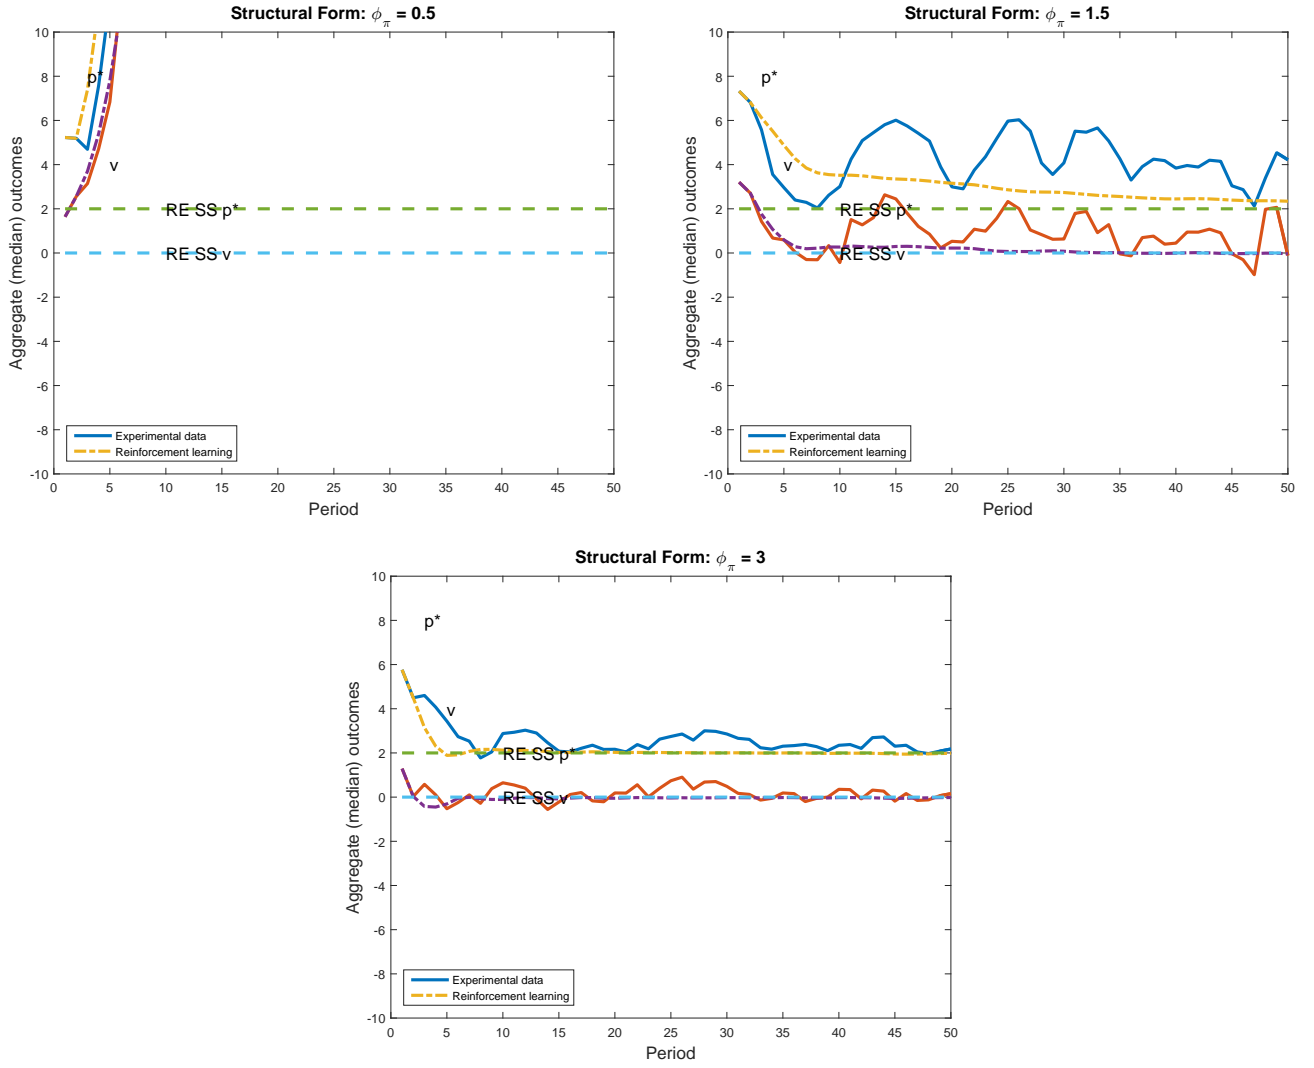


Figure 10: Actual (mean over all experimental groups) outcomes and simulated (mean over 6,000 replications) aggregated outcomes using reinforcement

M is:

$$MSE^M = \frac{1}{4} \sum_{g=1}^4 \frac{1}{48} \sum_{t=3}^{50} (\bar{v}_{t,g} - v_{t,g}^M)^2 \quad (44)$$

where t is period and g is group subscript. Thus, $\bar{v}_{t,g}$ is the aggregate expenditure variable in group g in time period t , while $\bar{v}_{t,g}^M$ is the forecast of the aggregate variable for period t by model M . The particular $\bar{v}_{t,g}^M$ that minimizes the MSE is the conditional expectation $\mathbb{E}(\bar{v}_{t,g}^M | [\bar{v}^1, \dots, \bar{v}^6]_{t=1,2}, [p^{*1}, \dots, p^{*6}]_{t=1,2})$ (see e.g. Hamilton (1994); p. 72) Since noisy learning and reinforcement are stochastic models and the conditional expectation of reinforcement is particularly challenging to evaluate analytically, I use the ensemble average of a large number of replications ($N = 6,000$), being a good approximation of the conditional expectation due to the weak law of large numbers. To ensure

comparability, this procedure has been adopted for both models. Specifically:

$$MSE^M \equiv \frac{1}{4} \sum_{g=1}^4 \frac{1}{48} \sum_{t=3}^{50} \left(\bar{v}_{t,g} - \frac{1}{N} \sum_{n=1}^N \bar{v}_{t,g,n}^M \right)^2 \quad (45)$$

where the n-subscript denotes the n-th replication. Due to the self-referential nature of the system, the models exhibit path dependence so that initial conditions can have a large impact on the dynamic behavior of the outcomes. Hence, I discard the first 30 periods, as this is standard for simulations in the learning literature to attenuate the effect of initial conditions. See for example Orphanides and Williams (2007).

In addition, I compare other statistics predicted by the path simulations of the learning models to those in the experiment:

- first moments (mean)
- second moments (standard deviation)
- mean square distance (MSD) to REE

7.1.2 Individual statistics

As a particularly interesting question to ask is how well the learning models capture the underlying individual behavior, a set of individual statistics is reported:

- mean square error (MSE) averaged over all individuals
- mean square distance (MSD) from REE averaged over all individuals
- an index of intra-period dispersion, calculated as $\frac{1}{T} \sum_{t=1}^{50} \text{Std. dev.}(v_t^i)$ for household advisors and analogously $\frac{1}{T} \sum_{t=1}^{50} \text{Std. dev.}(p_t^{*j})$ for firm advisors

7.2 Results

The results for noisy adaptive learning and reinforcement are reported in table 9 for all periods and in table 10 for periods 30-50. The complete statistics for all learning models, including least square learning, constant gain and heuristic-switching, are in Appendix 9.8. The simulation result that is closer to the experimental data is bold-faced.³⁴ The following results stand out:

³⁴Investigating significance is not conclusive in this context, as the sample size is very large due to a large number of replications. An extremely large sample size leads to a standard error close to zero, which leads to extremely

Observation 12. Treatment 1 ($\phi_\pi = 0.5$): While noisy adaptive learning qualitatively predicts divergence, reinforcement provides a better quantitative fit.

Observation 13. Treatment 2 ($\phi_\pi = 1.5$): Noisy adaptive learning provides a better quantitative fit than reinforcement.

Observation 14. Treatment 3 ($\phi_\pi = 3$): While statistics for the whole sample indicate that reinforcement is the better model, this conclusion may be driven by initial conditions. For periods 30-50 in treatment 3 ($\phi_\pi = 3$), it largely depends on the statistic of interest which model yields a better fit.

Hence, the answer to the question which model should be used for expectation formation depends on the context: If the researcher looks for a tractable model with few parameters, noisy adaptive learning may be a good choice. The reason is that it qualitatively predicts the divergent path for passive monetary policy and, despite its parsimony in terms of parameters, provides a good quantitative fit for the active monetary policy treatments. If the research focus is, on the other hand, for instance on examining differences in expectation formation for a divergent path as opposed to paths that stay in the vicinity of the REE, a more sophisticated model such as reinforcement may need to be used.

Measures	$\phi_\pi = 0.5$			$\phi_\pi = 1.5$			$\phi_\pi = 3.0$		
	<i>Data</i>	<i>Reinf.</i>	<i>Noisy</i>	<i>Data</i>	<i>Reinf.</i>	<i>Noisy</i>	<i>Data</i>	<i>Reinf.</i>	<i>Noisy</i>
<u>Aggregates</u>									
MSE: v	0.00	149,720.10	271,444.44	0.00	4.98	5.03	0.00	0.65	1.12
MSE: p^*	0.00	305,299.15	517,036.07	0.00	12.05	13.37	0.00	0.92	1.74
Mean: v	377.43	485.33	8.89	0.96	0.29	0.91	0.17	-0.03	-0.24
Mean: p^*	562.20	725.75	15.90	4.31	3.22	5.60	2.63	2.17	3.11
MSD from REE: v	268,390.11	443,216.51	140.60	5.86	2.90	3.93	0.71	0.61	1.69
MSD from REE: p^*	514,324.42	925,942.38	317.16	16.27	8.98	18.20	1.52	0.91	3.40
Standard dev.: v	311.47	402.38	6.36	1.80	1.16	1.48	0.71	0.63	1.18
Standard dev.: p^*	434.66	569.56	9.09	2.55	1.82	1.62	0.96	0.93	1.41
<u>Individuals</u>									
MSE: v^i	0.00	142,536.97	346,757.07	0.00	4,867.40	7,050.73	0.00	5,017.29	6,889.12
MSE: p^{*j}	0.00	310,652.46	507,851.64	0.00	17.45	25.19	0.00	5.04	9.53
MSD from REE: v^i	343,486.43	674,365.20	552.14	6,826.17	351.39	94.03	6,603.38	136.98	204.94
MSD from REE: p^{*j}	505,904.75	1,320,941.26	392.91	24.96	9.61	29.15	9.25	4.08	12.39
Dispersion of v^i	278.19	219.71	14.87	39.43	21.12	7.93	30.31	13.10	7.36
Dispersion of p^{*j}	160.10	42.83	5.88	2.52	1.49	3.16	1.69	1.13	2.74

MSE=mean square error; MSD=mean square distance

Table 9: Measures all periods

low p-values even if the difference is minimal.

Measures	$\phi_\pi = 0.5$			$\phi_\pi = 1.5$			$\phi_\pi = 3.0$		
	<i>Data</i>	<u>Reinf.</u>	<u>Noisy</u>	<i>Data</i>	<u>Reinf.</u>	<u>Noisy</u>	<i>Data</i>	<u>Reinf.</u>	<u>Noisy</u>
<u>Aggregates</u>									
MSE: v	<i>0.00</i>	201,402.64	329,792.12	<i>0.00</i>	4.95	4.96	<i>0.00</i>	0.35	0.79
MSE: p^*	<i>0.00</i>	383,002.42	557,155.69	<i>0.00</i>	12.09	10.26	<i>0.00</i>	0.53	0.47
Mean: v	<i>494.84</i>	847.00	14.19	<i>0.72</i>	0.00	0.12	<i>0.09</i>	-0.03	-0.50
Mean: p^*	<i>707.07</i>	1247.83	24.29	<i>4.09</i>	2.49	4.54	<i>2.34</i>	1.98	2.09
MSD from REE: v	<i>343,633.59</i>	847,755.38	279.97	<i>4.98</i>	3.77	1.95	<i>0.35</i>	0.69	0.98
MSD from REE: p^*	<i>588,244.14</i>	1,735,919.83	640.78	<i>14.57</i>	8.60	10.94	<i>0.52</i>	1.12	0.61
Standard dev.: v	<i>241.03</i>	196.44	5.72	<i>1.44</i>	0.52	0.98	<i>0.55</i>	0.33	0.67
Standard dev.: p^*	<i>272.92</i>	188.03	7.20	<i>1.92</i>	0.65	0.98	<i>0.52</i>	0.35	0.57
<u>Individuals</u>									
MSE: v^i	<i>0.00</i>	176,738.81	426,232.66	<i>0.00</i>	1,984.51	1,808.37	<i>0.00</i>	1,839.98	2,139.97
MSE: p^{*j}	<i>0.00</i>	400,773.41	584,959.77	<i>0.00</i>	16.04	14.71	<i>0.00</i>	2.11	3.33
MSD from REE: v^i	<i>444,173.06</i>	823,199.31	740.93	<i>1,831.77</i>	2,476.59	88.56	<i>2131.86</i>	2,058.74	87.67
MSD from REE: p^{*j}	<i>616,582.76</i>	1,713,221.67	694.30	<i>18.85</i>	9.66	17.40	<i>2.18</i>	1.50	3.68
Dispersion of v^i	<i>360.54</i>	299.81	16.44	<i>25.06</i>	21.80	7.87	<i>20.43</i>	15.09	7.08
Dispersion of p^{*j}	<i>221.88</i>	46.21	6.20	<i>1.91</i>	0.84	1.87	<i>0.90</i>	0.48	1.52

MSE=mean square error; MSD=mean square distance

Table 10: Measures periods 30-50

8 Conclusion

This study has shown that due to different feedback structures, the structural-form New-Keynesian model does not produce the same results as the reduced-form New-Keynesian model. This implies that analysts or researchers should be careful, if they draw policy implications based on empirical analysis with the reduced form. More generally, relying upon RE models for policy analysis might not be constructive, if the observed process of expectation formation exhibits bounded rationality, as conclusions and policy implications that are based upon RE might be misleading.

This study provides an important example, since the experimental results indicate that barely fulfilling the Taylor principle can at best “prevent the worst” but exhibits no evidence of convergence to the REE even within the considerable time frame of 50 periods. A novel feature of this paper is the application of the reinforcement model by Roth and Erev (1998) to data generated by a “learning-to-forecast” experiment. Furthermore, a theoretical contribution is the noisy adaptive learning model, which explains both the slow convergence for a Taylor rule coefficient of 1.5 in the structural form and the expectation-driven fluctuations. Consistently with the experimental results, both adaptive and noisy learning suggest that the monetary authority should adopt a reinforced Taylor principle with a larger reaction coefficient to ensure stability of the REE.

Laboratory experiments are one example of providing empirical foundations for actual behavior, but they are not the end of exploring bounded rationality. One needs to test their external validity and find a good connection to decisions of real firms and households. This is crucial in order to verify the policy recommendations implied by the experimental data and thus certainly

a potential direction of further research. Yet, laboratory experiments are a good starting point for several reasons. Firstly, they provide the closest link to the theoretical models, as important assumptions from theory can easily be implemented in the laboratory. Secondly, they can offer intuition and hints, pinpointing in which direction existing models could be extended and what to look for in real data.

References

- Klaus Adam. Experimental Evidence on the Persistence of Output and Inflation. *Economic Journal*, 117(1995):1–48, 2007.
- Simon P. Anderson, Jacob K. Goeree, and Charles A. Holt. Noisy directional learning and the logit equilibrium. *Scandinavian Journal of Economics*, 106(3):581–602, 2004.
- Mikhail Anufriev and Cars Hommes. Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(225408):35–64, 2012.
- Jasmina Arifovic and Luba Petersen. Escaping Expectations-Driven Liquidity Traps : Experimental Evidence. *Simon Fraser University Working Paper*, 2015.
- Tiziana Assenza, Peter Heemeijer, Cars Hommes, and Domenico Massaro. Managing Self-organization of Expectations through Monetary Policy : a Macro Experiment. (225408):1–53, 2014.
- Te Bao, John Duffy, and Cars Hommes. Learning, forecasting and optimizing: An experimental study. *European Economic Review*, 61:186–204, 2013.
- KG Binmore, L Samuelson, and R Vaughan. Musical chairs: modeling noisy evolution. *Games and Economic Behavior*, 11(1):1–35, 1995.
- Olivier J. Blanchard and Charles M. Kahn. The solution of linear difference models under rational expectations. *Econometrica*, 48(5):1305–1311, 1980.
- Tim Bollerslev. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 31(1):307–327, 1986.

- William A. Branch and Bruce McGough. A New Keynesian model with heterogeneous expectations. *Journal of Economic Dynamics and Control*, 33(5):1036–1051, 2009.
- William A Brock and Cars H Hommes. A Rational Route to Randomness. *Econometrica*, 65(5): 1059–1095, 1997.
- James Bullard and Kaushik Mitra. Learning about monetary policy rules. *Journal of Monetary Economics*, 49:1105–1129, 2002.
- Guillermo A. Calvo. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(1978):383–398, 1983.
- Colin Camerer and Teck-hua Ho. Experienced-Weighted Attraction Learning in Normal Form Games. *Econometrica*, 67(4):827–874, 1999.
- Richard Clarida, Jordi Galí, and Mark Gertler. Monetary Policy Rules And Macroeconomic Stability: Evidence And Some Theory. *Quarterly Journal of Economics*, 115(February):147–180, 2000.
- Paul de Grauwe. Animal spirits and monetary policy. *Economic Theory*, 47(2-3):423–457, 2011.
- John Duffy. Agent-Based Models and Human Subject Experiments. In Leigh Tesfatsion and Kenneth L. Judd, editors, *Handbook of Computational Economics*, volume 2, chapter 19, pages 949–1011. Elsevier, Amsterdam, 2006. ISBN 9780444512536.
- John Duffy. Macroeconomics: a survey of laboratory research. *Handbook of experimental economics* 2, pages 1 – 97, 2014.
- Robert F Engle. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50(4):987–1007, 1982.
- George W. Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.
- George W Evans and Seppo Honkapohja. Expectations and the Stability Problem for Optimal Monetary Policies. *Review of Economic Studies*, 70:807–824, 2003.
- George W Evans and Seppo Honkapohja. Monetary policy, expectations and commitment. *Scandinavian Journal of Economics*, 108(0136848):15–38, 2006.

- George W. Evans, Seppo Honkapohja, and Noah Williams. Generalized stochastic gradient learning. *International Economic Review*, 51(1):237–262, 2010.
- Giuseppe Ferrero. Monetary policy, learning and the speed of convergence. *Journal of Economic Dynamics and Control*, 31(9):3006–3041, 2007.
- Urs Fischbacher. Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10:171–178, 2007.
- D Fudenberg and C Harris. Evolutionary dynamics with aggregate shocks. *Journal of Economic Theory*, 57:420–441, 1992.
- Jordi Galí. *Monetary Policy, Inflation and the Business Cycle*. Princeton University Press, Princeton, NJ, 2008. ISBN 978-0-691-13316-4.
- Ben Greiner. An Online Recruitment System for Economic Experiments. *Forschung und Wissenschaftliches Rechnen*, 63:79–93, 2004.
- James D Hamilton. Time Series Analysis. *Prentice Hall New Jersey 1994*, SFB 373(Chapter 5): 837–900, 1994.
- Peter Heemeijer, Cars Hommes, Joep Sonnemans, and Jan Tuinstra. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control*, 33:1052–1072, 2009.
- Robert J. Hodrick and Edward C Prescott. Postwar U.S. Business Cycles: An Empirical Investigation. *Journal of Money, Credit and Banking*, 29(1):1, feb 1997.
- Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk van de Velden. Coordination of Expectations in Asset Pricing Experiments. *Review of Financial Studies*, 18(3):955–980, aug 2005.
- Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk van de Velden. Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization*, 67(1):116–133, jul 2008.
- Cars Hommes, Domenico Massaro, and Isabelle Salle. Monetary and Fiscal Policy Design at the Zero Lower Bound – Evidence from the lab. *working paper, University of Amsterdam*, 2015.

- Oleksiy Kryvtsov and Luba Petersen. Expectations and Monetary Policy: Experimental Evidence. *Working Paper*, 2015.
- Mordecai Kurz, Giulia Piccillo, and Howei Wu. Modeling diverse expectations in an aggregated New Keynesian Model. *Journal of Economic Dynamics and Control*, 37(8):1403–1433, 2013.
- Finn Kydland and Edward C. Prescott. Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370, 1982.
- Vivian Lei and Charles N Nousseir. An Experimental Test of an Optimal Growth Model. *American Economic Review*, 92(3):411–433, jun 2002.
- Albert Marcet and Juan Pablo Nicolini. Recurrent hyperinflations and learning. *American Economic Review*, 93(5):1476–1498, 2003.
- Albert Marcet and Thomas J Sargent. Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory*, 48(2):337–368, 1989.
- Ramon Marimon and Shyam Sunder. Expectations and learning under alternative monetary regimes: an experimental approach. *Economic Theory*, 4(1):131–162, jan 1994.
- Felix Mauersberger, Rosemarie Nagel, and Nicolaas J. Vriend. Learning in an oligopoly market. *mimeo*, 2017.
- Ellen R. McGrattan, Richard Rogerson, and Randall Wright. An Equilibrium Model of the Business Cycle With Household Production and Fiscal Policy. *International Economic Review*, 38(2):267–90, 1997.
- Fabio Milani. Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics*, 54(7):2065–2082, 2007.
- Fabio Milani. Learning and time-varying macroeconomic volatility. *Journal of Economic Dynamics and Control*, 47:94–114, 2014.
- Pooya Molavi, Alireza Tahbaz-Salehi, and Ali Jadbabaie. Foundations of non-Bayesian Social Learning. *Working Paper*, (October), 2015.
- Rosemarie Nagel. Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326, 1995.

- Rosemarie Nagel and Nicolaas J Vriend. An experimental study of adaptive behavior in an oligopolistic market game. *Journal of Evolutionary Economics*, 9(1):27–65, 1999.
- Stephen Nickell. Biases in Dynamic Models with Fixed Effects. *Econometrica*, 49(6):1417–1426, 1981.
- Athanasios Orphanides and John C. Williams. Inflation scares and forecast-based monetary policy. *Review of Economic Dynamics*, 8(2 SPEC. ISS.):498–527, 2005a.
- Athanasios Orphanides and John C. Williams. The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic Dynamics and Control*, 29(11):1927–1950, 2005b.
- Athanasios Orphanides and John C. Williams. Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics*, 54(5 SPEC. ISS.):1406–1435, 2007.
- Luba Petersen. Do Expectations and Decisions Respond to Monetary Policy? *Journal of Economic Studies*, 43(6):1047–1062, 2015.
- Damjan Pfajfar and Emiliano Santoro. Heterogeneity, learning and information stickiness in inflation expectations. *Journal of Economic Behavior and Organization*, 75(3):426–444, 2010.
- Damjan Pfajfar and Blaž Žakelj. Experimental evidence on inflation expectation formation. *Journal of Economic Dynamics and Control*, 44:147–168, jul 2014.
- Michele Piccione and Ariel Rubinstein. On the Interpretation of Decision Problems with Imperfect Recall. *Games and Economic Behavior*, 20(1):3–24, jul 1997.
- Bruce Preston. Learning about Monetary Policy Rules when Long-Horizon Expectations Matter. *International Journal of Central Banking*, 1:81–126, 2005.
- Julio J Rotemberg and Michael Woodford. Optimization-Based Framework for the Evaluation of Monetary Policy. In *NBER Macroeconomics Annual*, volume 12, pages 297–346. 1997. ISBN 0262024357.
- Alvin Roth and Ido Erev. Prediction how people play games: Reinforcement learning in games with unique strategy equilibrium. *American Economic Review*, 88(4):848–881, 1998.

- Thomas J. Sargent. The Ends of Four Big Inflations. In *Rational Expectations and Inflation*, volume I, chapter 2, pages 41–97. Harper & Row, New York, 1986. ISBN 9781400847648.
- Thomas J. Sargent. *Bounded Rationality in Macroeconomics: The Arne Ryde Memorial Lectures*. Clarendon Press, 1994.
- Reinhard Selten. Features of experimentally observed bounded rationality. *European Economic Review*, 42(3-5):413–436, 1998.
- Reinhard Selten and Rolf Stoecker. End behavior in sequences of finite Prisoner’s Dilemma supergames A learning theory approach. *Journal of Economic Behavior & Organization*, 7(1): 47–70, 1986.
- Dale O Stahl. Boundedly Rational Rule Learning in a Guessing Game. *Games and Economic Behavior*, 16:303–330, 1996.
- John B. Taylor. Discretion versus policy rules in practice: two critical points. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214, 1993.
- John B. Taylor. *A Historical Analysis of Monetary Policy Rules*, volume I. 1999. ISBN 0226791246.
- Carl E. Walsh. *Monetary Theory and Policy*. MIT Press, Cambridge, 3rd edition, 2010.
- Michael Woodford. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, 2003.
- Michael Woodford. Macroeconomic analysis without the rational expectations hypothesis. *Annual Review of Economics*, 5(1):303–346, 2013.
- Tack Yun. Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of Monetary Economics*, 37:345–370, 1996.

9 Appendix

9.1 Model in Woodford (2013)

9.1.1 The demand side

This cashless economy is populated by a continuum of households, indexed by i , seeking to maximize the present discounted value of expected utility

$$\hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} [U(C_T^i) - \Upsilon(h_T^i(j))] \quad (46)$$

where C_t^i is a composite consumption good and \hat{E}_t^i is an arbitrary subjective (not necessarily rational) expectations operator of household i given the information set in period t , which however satisfies the law of iterated expectations such that $\hat{E}_t^i \hat{E}_{t+1}^i x_{t+k} = \hat{E}_t^i x_{t+k}$. $h^i(j)$ is the amount of labor supplied by household i for the production of good j . The second term in the brackets is to be interpreted as the total disutility of labor supply. There is an equal number of households supplying labor for each type of good.³⁵

The household has no choice but supplying the hours of work demanded by the firm at the given wage, being negotiated by a union on behalf of all households, so that $H^i(t) = \int_0^1 h^i(j) = H_t, \forall i$. Thus, a household has a single decision each period, which is the amount to individually spend on the composite consumption good, C_t^i , defined as:

$$C_t^i \equiv \left[\int_0^1 c_t^i(j)^{\frac{\epsilon-1}{\epsilon}} \right]^{\frac{\epsilon}{\epsilon-1}} \text{ with associated price index } P_t \equiv \left[\int_0^1 P_t(j)^{1-\epsilon} \right]^{\frac{1}{1-\epsilon}} \quad (47)$$

where $c_t^i(j)$ is the (real) expenditure of household i on good j . One main difference of this heterogeneous agent model to the representative agent model is that households can have non-zero asset holdings. There is one single riskless one-period bond in the economy and the household's law of motion of bond holdings can be written as

$$B_{t+1}^i = (1 + i_t) [B_t^i + W_t H_t + \int_{j=0}^1 \tilde{\Pi}_t(j) dj - \int_{j=0}^1 p_t(j) c_t^i(j) dj - T_t] \quad (48)$$

³⁵Woodford (2003, p. 144ff.) shows that under this assumption, this is equivalent to a representative agent supplying labor for each type of good.

where i_t denotes the nominal interest rate on bonds held at the end of period t , B_t^i the household's nominal bond holdings carried into period t , W_t is the nominal wage, $\tilde{\Pi}_t(j)$ is nominal profits of firm j (distributed in equal dividend shares to the households), $p_t(j)$ is the price of good j and T_t denotes net lump-sum taxes. Since each firm's profits are given by

$$\tilde{\Pi}_t(j) = p_t(j)y_t(j) - W_t H_t(j) \quad (49)$$

we have

$$\int_{j=0}^1 p_t(j)y_t(j)dj = W_t H_t + \int_{j=0}^1 \tilde{\Pi}_t(j)dj \quad (50)$$

Optimal allocation of household expenditure across differentiated goods yields the set of demand equations

$$c_t^i(j) = C_t^i \left(\frac{P_t(j)}{P_t} \right)^{-\epsilon} \quad (51)$$

which further implies

$$\int_{j=0}^1 p_t(j)c_t^i(j)dj = P_t C_t^i \quad (52)$$

$$\int_{j=0}^1 p_t(j)y_t(j)dj = P_t Y_t \quad (53)$$

where Y_t is aggregate demand for the composite consumption good defined in (47). Standard analysis shows that household intertemporal optimality is given by the consumption Euler equation:

$$\frac{1}{1 + i_t} = \beta \hat{E}_t \left[\frac{P_t}{P_{t+1}} \frac{U_{C^i,t+1}(C_{t+1}^i)}{U_{C^i,t}(C_t^i)} \right] \quad (54)$$

To work with a predetermined indicator of bond holdings, we define $b_t^i \equiv \frac{B_t^i}{P_{t-1}\Pi}$. The structural relations of the model are subsequently log-linearized around a deterministic steady state in which (a) the inflation steady state $\bar{\pi}$ is set by the monetary authority, $i_t \approx \ln(1 + \bar{i}) = \bar{\pi} - \ln \beta$ (as implied by (54))³⁶, $b_T^i = \frac{B_T^i}{P_{T-1}} = \bar{b}$, $Y_T = \bar{Y}$, $C_T = \bar{C}$ and $\bar{\tau}_T \equiv \frac{T_T}{P_T} = \bar{\tau}$ for all $T \geq t$, (b) all

³⁶Assenza et al. (2014) and Pfajfar and Žakelj (2014) neglect the term $\ln \beta$.

subjective expectations are correct. By log-linearising (54) and (48), we obtain

$$\hat{c}_t^i = \hat{E}_t^i \hat{c}_{t+1}^i - \sigma(\hat{i}_t - \hat{E}_t^i \hat{\pi}_{t+1}) \quad (55)$$

$$\hat{b}_{t+1}^i = s_b(\hat{i}_T - \beta^{-1} \hat{\pi}_t) \beta^{-1} (\hat{b}_t^i + (\hat{Y}_t - \hat{\tau}_t) - \hat{c}_t^i) \quad (56)$$

where $s_b \equiv \frac{\bar{b}}{\bar{Y}}$. “ \hat{x}_t ” (in minuscules) denotes the deviation from the steady state in natural logarithms of any variable X at time t ; except for nominal bond holdings, where $\hat{b}_t^i \equiv \frac{b_t^i - \bar{b}}{\bar{Y}}$ is written in terms of steady state output, and the interest rate, where $\hat{i}_t = \ln \frac{1+i_t}{1+i}$ is used.

Solving (55) forward at time t and substituting the result into the also forward-solved equation resulting from (56) gives:

$$\hat{c}_t^i = (1 - \beta) \hat{b}_t^i + \sum_{T=t}^{\infty} \beta^{T-t} \hat{E}_t^i \{ (1 - \beta)(\hat{Y}_T - \hat{\tau}_T) - \beta \sigma (\hat{i}_T - \hat{\pi}_{T+1}) + (1 - \beta) s_b (\beta \hat{i}_T - \hat{\pi}_T) \} \quad (57)$$

(57) can be rewritten in its recursive form under internally consistent expectations of the household as

$$\hat{c}_t^i = (1 - \beta) \hat{b}_t^i + (1 - \beta)(\hat{Y}_T - \hat{\tau}_T) - \beta [\sigma - (1 - \beta) s_b] \hat{i}_t - (1 - \beta) s_b \hat{\pi}_t + \beta \hat{E}_t^i v_{t+1}^i \quad (58)$$

where

$$v_t^i \equiv \sum_{T=t}^{\infty} \beta^{T-t} \hat{E}_t^i \{ (1 - \beta)(\hat{Y}_T - \hat{\tau}_T) - [\sigma - (1 - \beta) s_b] (\beta \hat{i}_T - \hat{\pi}_T) \} \quad (59)$$

The advantage of this notation is that individuals only need to forecast a single variable. Using the goods market clearing condition $\hat{Y}_t = \int \hat{c}_t^i di$, aggregate demand can be obtained as

$$\hat{Y}_t = (1 - \beta) \hat{b}_t + v_t - \sigma \hat{\pi}_t \quad (60)$$

Woodford (2013) assumes for simplicity that government expenditure is an exogenous disturbance, but, since I introduce the model into the laboratory without any shocks, in my setup, the government merely uses the taxes to service the debt it has accumulated. Hence, the government’s flow

budget constraint is given as

$$\hat{b}_{t+1} = \beta^{-1}[\hat{b}_t - s_b \hat{\pi}_t - \hat{\tau}_t] + s_b \hat{i}_t \quad (61)$$

(4) implies the recursive form

$$v_t^i = (1 - \beta)v_t + \beta(1 - \beta)(\hat{b}_{t+1} - \hat{b}_t) - \beta\sigma(\hat{i}_t - \hat{\pi}_t) + \beta\hat{E}_t^i v_{t+1}^i \quad (62)$$

where $v_t = \int v_t^i di$ is the average value across agents of the expectational variable defined in (4).

I follow Woodford (2013) in assuming that expectations are Ricardian so that

$$b_t = \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} [\hat{\tau}_T - s_b(\beta \hat{i}_T - \hat{\pi}_T)] \quad (63)$$

This is a strong assumption. However, it simplifies the model considerably and this assumption is frequently made in non-RE analyses. (See Woodford (2013) for a detailed discussion.) In the context of this experiment, this can be interpreted as the computerized household understanding that government cannot forever accumulate debt. Under this assumption, aggregate demand is independent of the supply of public debt and can be more compactly written as

$$\hat{Y}_t = \bar{v}_t - \sigma \hat{\pi}_t \quad (64)$$

where $\bar{v}_t = v_t + (1 - \beta)\hat{b}_t$ is the aggregate of a subjective variable \bar{v}_t^i , which is the variable that one group of subjects, labeled as “household advisors”, needs to forecast in the experiment, and which can be defined simply as

$$\bar{v}_t^i \equiv \sum_{T=t}^{\infty} \hat{E}_t^i \{(1 - \beta)\hat{Y}_T - \sigma(\beta \hat{i}_T - \pi_T)\} \quad (65)$$

9.1.2 The supply side

We assume Calvo (1983) price-setting such that a fraction $0 < \alpha < 1$ of goods prices are exogenously held fixed in any period. Producers engage in monopolistic competition, which means that each firm sets the price for a good that it alone produces. Under full price indexation, the

log-linearized approximation of the inflation dynamics is given by

$$\hat{\pi}_t = (1 - \alpha)\hat{p}_t^* \quad (66)$$

where, for each firm j that is chosen to re-optimize its price in period t , p_t^{*j} is the amount by which the firm would choose to set the log price of its good higher than p_{t-1} . $p_t^* = \int p_t^{*j} dj$ is the average value of this variable across all firms that are chosen by the Calvo mechanism to reoptimize prices in period t . The firm j 's maximization problem is

$$\max_{P_t^{*j}} \sum_{T=t}^{\infty} \alpha^{T-t} \hat{E}_t^j \{ Q_{t,T} (P_t^{*j} Y_T(j) \bar{\Pi}^{T-t} - \Psi(Y_T(j))) \} \quad (67)$$

subject to (51), where $Y_T(j)$ denotes output in period T for a firm j , $\Psi(\cdot)$ is the nominal cost function and $Q_{t,T}$ is the stochastic discount factor, describing how a unit of income in each state and at date T is valued in the present:

$$Q_{t,T} = \beta^{T-t} \frac{P_t U_C(Y_T)}{P_T U_C(Y_t)} \quad (68)$$

Denoting $MC_{T|t} = \frac{\psi_t}{P_T}$ as the real marginal cost, the log-linearized first-order condition is:

$$\hat{p}_t^{*j} = (1 - \alpha\beta) \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \hat{E}_t^j \{ p_T^{opt} - p_{t-1} - (T - t + 1)\bar{\pi} \} \quad (69)$$

where $p_T^{opt} \equiv \hat{m}c_t + p_T$ is the single-period optimal log price, being the same for each firm, as they face the same labor market and aggregate conditions. Rewriting (19) in its recursive form gives:

$$\hat{p}_t^{*j} = (1 - \alpha\beta)(p_t^{opt} - p_{t-1} - \bar{\pi}) + \alpha\beta(\hat{E}_t^j \hat{p}_{t+1}^{*j} + \hat{\pi}_t) \quad (70)$$

Suppose that a union sets the wage on behalf of the households under the objective that at that wage, a marginal increase in labor demand would neither increase nor decrease average perceived utility across households, if for each household the marginal utility of additional wage income is weighted against the marginal disutility of additional work. Hence, the optimality condition for

the union is

$$\frac{\Upsilon_h(h_t)}{u_{C,t}(C_t)} = \frac{w_t}{P_t} \quad (71)$$

By log-linearising (71), one obtains

$$\hat{\omega}_t = \hat{\vartheta}_t - \hat{u}_{c,t} \quad (72)$$

where ω_t is the log real wage, ϑ_t is the log of the (common) marginal disutility of labor and $u_{c,t}$ is the (log) aggregate of the marginal utility of additional real income. Since

$$\hat{u}_{c^i,t} = -\sigma^{-1}\hat{c}_t^i \quad (73)$$

we have

$$\hat{\omega}_t = \hat{\vartheta}_t + \sigma^{-1}\hat{c}_t = \hat{\vartheta}_t + \sigma^{-1}\hat{Y}_t \quad (74)$$

Given that $\hat{m}c_t = \hat{\omega}_t - \hat{m}pn_t$ and since both $\hat{\vartheta}_t$ and $\hat{m}pn_t$ can be expressed as functions of labor hours and thus as output (which is determined by the market clearing condition for the aggregate goods market), one can summarize p_t^{opt} as

$$p_t^{opt} = p_t + \xi\hat{Y}_t \quad (75)$$

Assuming constant returns to scale in the aggregate and a disutility function of labor of

$$\Upsilon(h_t^i(j)) = \frac{h_t^i(j)^{1+\varphi}}{1+\varphi} \quad (76)$$

it is easy to see that

$$\xi \equiv \sigma^{-1} + \varphi \quad (77)$$

We obtain, by using (75) in (20):

$$\hat{p}_t^{*j} = (1 - \alpha)\hat{p}_t^* + (1 - \alpha\beta)\xi\hat{Y}_t + \alpha\beta\hat{E}_t^j\hat{p}_{t+1}^{*j} \quad (78)$$

9.2 Reduced form

9.2.1 Systems

In matrix form, the system comprising the “dynamic IS” equation (27) and the “Phillips curve” (28) can be rewritten as

$$\begin{bmatrix} \bar{v}_t \\ \pi_t \end{bmatrix} = \Psi d\bar{\pi} + \Psi C \begin{bmatrix} \hat{E}_t\bar{v}_{t+1} \\ \hat{E}_t\pi_{t+1} \end{bmatrix} \quad (79)$$

with $\Psi \equiv \frac{1}{(\alpha + \phi_\pi\sigma\xi - \alpha\phi_\pi\sigma\xi - \alpha\beta\phi_\pi\sigma\xi + \alpha^2\beta\phi_\pi\sigma\xi)}$ and

$$d \equiv \begin{bmatrix} \alpha\beta\sigma(\phi_\pi - 1) \\ \alpha(1 - \beta) + \phi_\pi\sigma\xi(1 - \alpha)(1 - \alpha\beta) \end{bmatrix}; C \equiv \begin{bmatrix} \alpha + \sigma\xi(1 - \alpha)(1 - \alpha\beta) & -\alpha\beta\sigma(\phi_\pi - 1) \\ \xi(\alpha - 1)(\alpha\beta - 1) & \alpha\beta \end{bmatrix} \quad (80)$$

With the calibration given in Table 1, being the same parameter as in Pfajfar and Žakelj (2014) and Assenza et al. (2014), we have

$$C = \begin{bmatrix} 0.8966 & -0.3414 \\ 0.2069 & 0.6828 \end{bmatrix} \quad (81)$$

9.2.2 Experimental design

I follow the existing design based on the reduced-form (Pfajfar and Žakelj, 2014; Assenza et al., 2014), which consists of

- a computerized central bank
- 6 HUMAN output gap forecasters being asked to submit $\hat{E}_t\tilde{Y}_{t+1}$
- 6 HUMAN inflation forecasters being asked to submit $\hat{E}_t\pi_{t+1}$

Since the objective of this experiment is to replicate previous findings, I follow Assenza et al. (2014) and Pfajfar and Žakelj (2014) and elicit the simple averages of the forecasts of different

individuals. Since I expected convergence in the reduced form, I used a different exchange rate (1 euro per 250 points) in order to save resources.

9.2.3 Results with $\phi_\pi = 1.5$

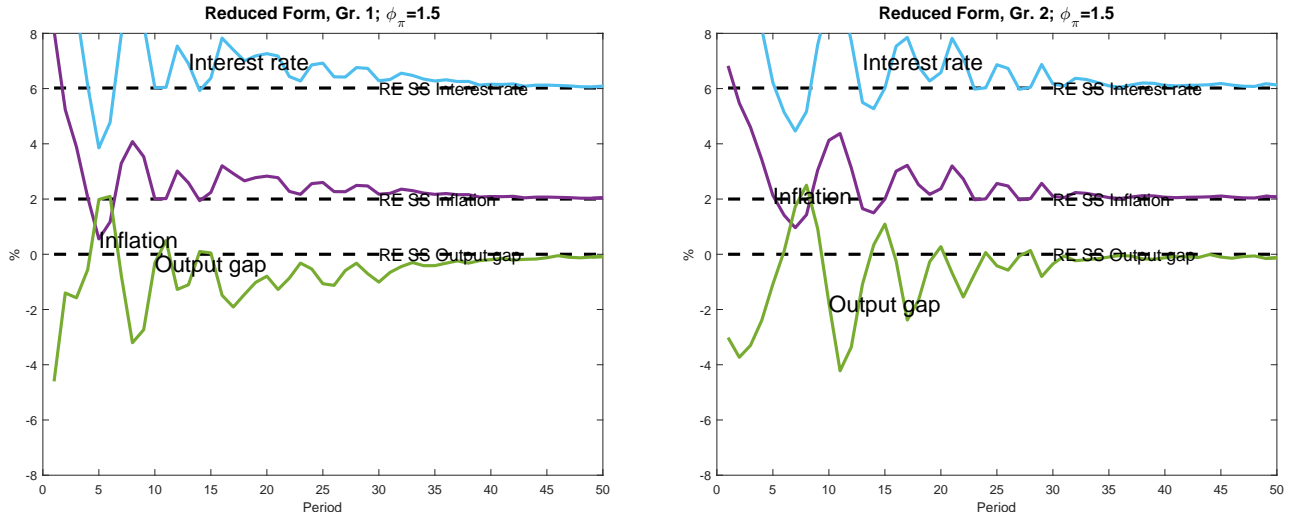


Figure 11: Reduced-form outcomes

Observation 15. In the reduced form with $\phi_\pi = 1.5$, the economy converges to the REE after 30 periods.

Figure 11 depicts the macroeconomic outcomes for the reduced form. Both groups converge to the REE, although out-of-equilibrium behavior at the beginning differs across groups: While group 1 converges to the REE after about 15 periods, group 2 takes about 20 periods. Due to Preston’s [2005] criticism on the internal validity of this approach and since the paramount result that in both groups the experimental economy converges to the REE is consistent with Assenza et al. (2014) and with Pfajfar and Žakelj (2014), no further sessions with the reduced form were conducted and the paper only entails a minimalistic analysis of this experiment.

9.2.4 Learning models and reduced form

The same learning models that have been applied to the structural form were also applied to the reduced form. The following points are worth highlighting:

- Since the heuristic-switching model was originally developed for setups like the reduced form, it was used in its original form. Hence, the simulation outcomes with the heuristic-switching model are deterministic.

- When fitting an adaptive rule to subjects' data, the gain coefficient differs considerably between the structural and the reduced form. Thus I allow the following parameters to differ across setups: the adaptive-rule coefficient in the heuristic-switching model and reinforcement (structural form: 0.3; reduced form: 0.65), the gain in constant gain learning (structural form: 0.11; reduced form: 0.80) as well as the gain parameters in noisy learning (structural form: $\hat{\gamma} = 0.20$, $\hat{\omega} = 0.62$; reduced form: $\hat{\gamma} = 0.71$, $\hat{\omega} = 0.23$.) Conceptually, the setup-varying gain is defensible, since in a setup, where the same outcome needs to be forecast, there is a stronger coordination motive so that making strong adjustments in the direction of the forecast error is sensible. In the structural form, however, as subjects are informed that their forecasts have a direct impact on their outcomes, it is reasonable to make cautious adjustments.³⁷ Rather than endogenizing the gain, this simplistic specification has the advantage that it keeps the model tractable. Furthermore, appendix 9.6 shows that an endogenizing the gain makes a low marginal contribution.

Figure 12 and tables 11 and 12 report the average results for 6,000 repeated simulations.

Reduced form: $\phi_\pi = 1.5$ all periods						
Measures	Experiment	HSM	OLS	CG	Reinforcement	Noisy
MSE: \bar{v}	<i>0.00</i>	0.81	0.64	0.61	0.31	0.88
MSE: inflation	<i>0.00</i>	0.64	1.77	0.71	0.34	0.99
Standard deviation: \bar{v}	<i>0.68</i>	0.94	0.61	0.57	0.65	1.05
Standard deviation: inflation	<i>1.05</i>	1.10	0.82	1.12	0.96	1.31
Mean: \bar{v}	<i>-0.09</i>	-0.31	-1.04	-0.18	-0.27	-0.34
Mean: inflation	<i>2.51</i>	2.22	3.40	2.23	2.31	2.31
Mean squared distance from REE: \bar{v}	<i>0.47</i>	0.93	1.82	0.39	0.53	1.43
Mean squared distance from REE: π	<i>1.35</i>	0.42	1.89	0.32	0.21	1.06

Table 11: Measures all periods

9.3 Adaptive learning

9.3.1 Least square learning

In specifying the functional form or Perceived Law of Motion (PLM) of agents' regression, I follow the standard in the adaptive learning literature by assuming that this functional form corresponds

³⁷Given that the structural form is a more sophisticated experimental game than the reduced form and may thus be closer to reality, the hypothesis that the gain is smaller in the structural than in the reduced form can also be supported by the fact that the empirical learning literature based on surveyed expectation data and real macroeconomic data also finds small gains. See for example Orphanides and Williams (2005a), Orphanides and Williams (2005b), Milani (2007), Orphanides and Williams (2007) and Pfajfar and Santoro (2010).

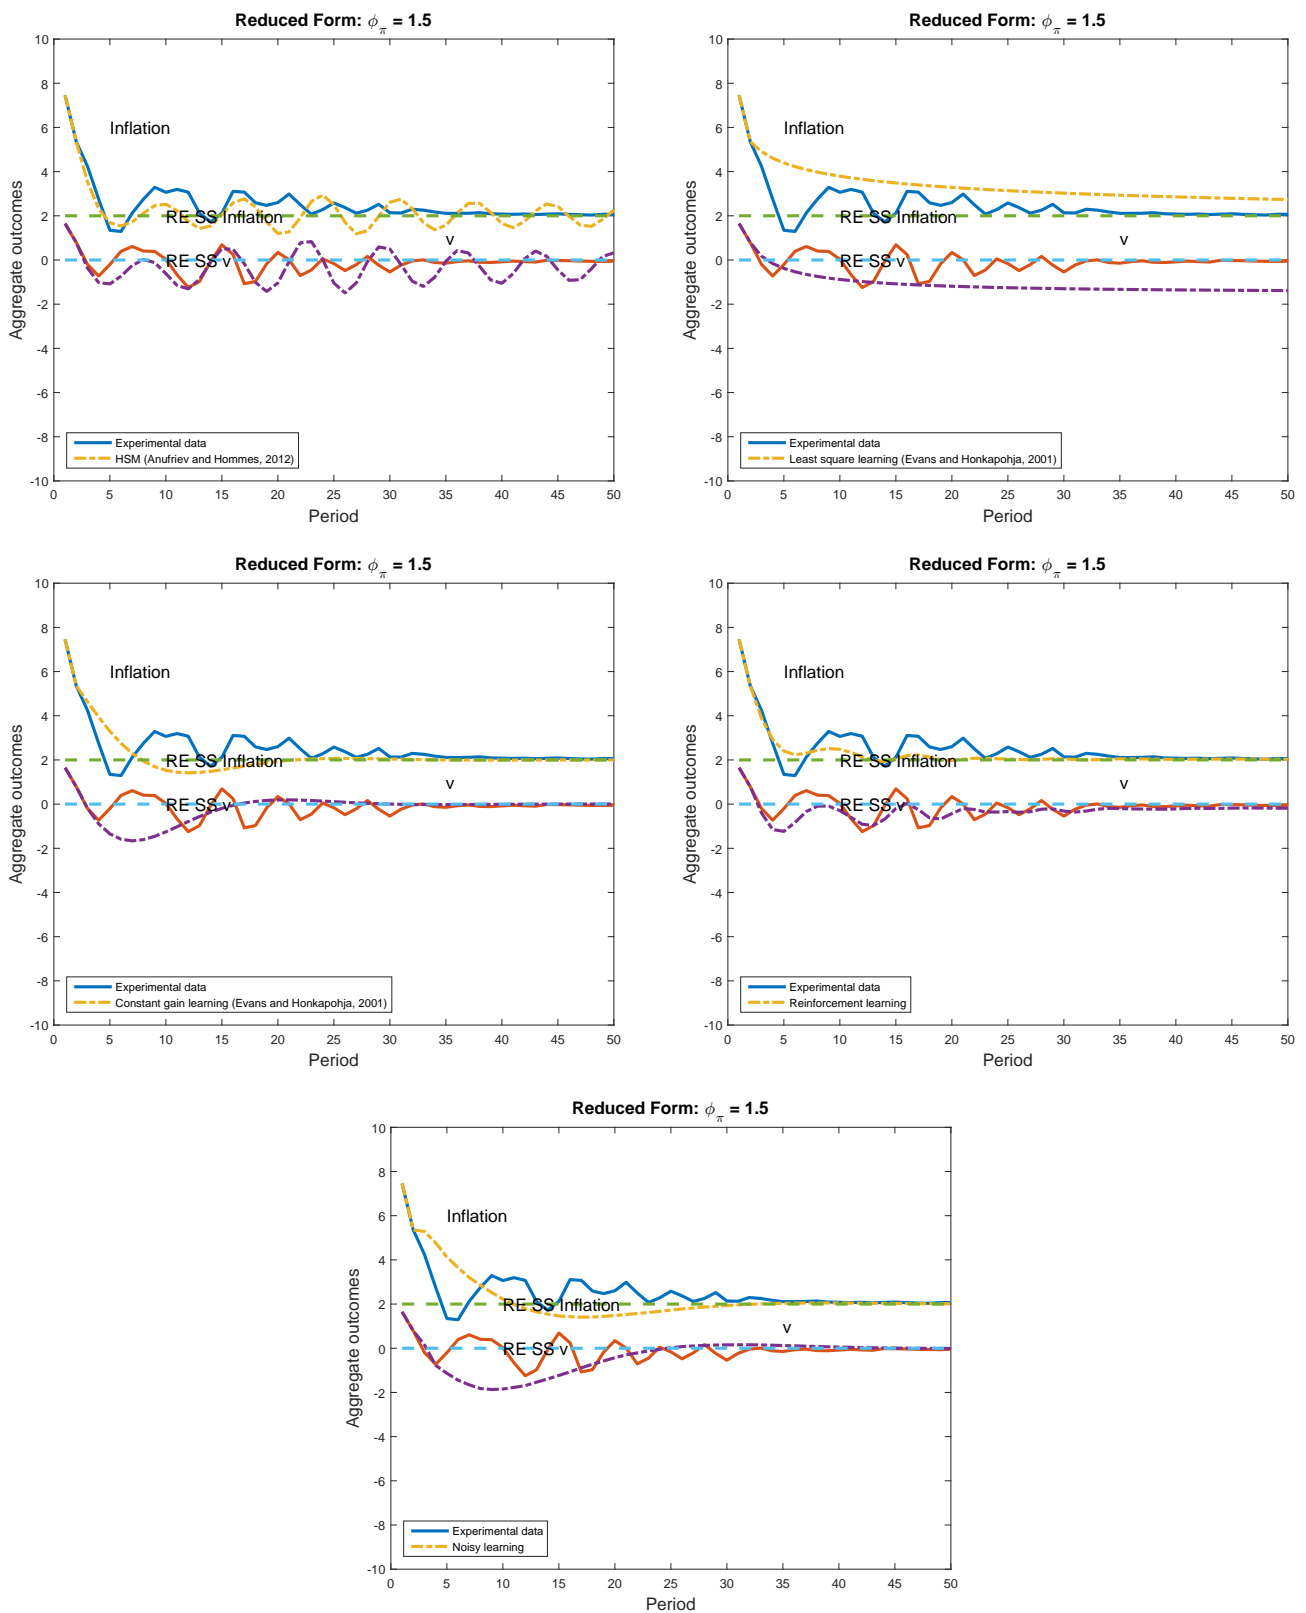


Figure 12: Actual (mean over all experimental groups) outcomes and simulated (mean over 6,000 replications) aggregated outcomes

to the functional form of the REE. Hence I assume that the PLM is a constant so that the least square forecast of agents for period $t+1$ corresponds to the sample mean of the past observed

Reduced form: $\phi_\pi = 1.5$ periods 30-50						
Measures	Experiment	HSM	OLS	CG	Reinforcement	Noisy
MSE: \bar{v}	<i>0.00</i>	0.58	0.46	0.03	0.09	0.06
MSE: inflation	<i>0.00</i>	0.23	1.14	0.02	0.04	0.02
Standard deviation: \bar{v}	<i>0.12</i>	0.61	0.04	0.01	0.09	0.19
Standard deviation: inflation	<i>0.07</i>	0.47	0.09	0.01	0.04	0.13
Mean: \bar{v}	<i>-0.09</i>	-0.30	-1.35	-0.01	-0.21	0.07
Mean: inflation	<i>2.11</i>	2.04	2.87	2.00	2.03	2.02
Mean squared distance from REE: \bar{v}	<i>0.03</i>	0.53	2.18	0.00	0.11	0.07
Mean squared distance from REE: π	<i>0.02</i>	0.23	0.85	0.00	0.03	0.03

Table 12: Measures from period 30 to 50

values of this variable up to period t-1:

$$E_t^i \bar{v}_{t+1}^i = (t-1)^{-1} \sum_{s=1}^{t-1} \bar{v}_{t-s}^i \quad (82)$$

or written recursively

$$E_t^i \bar{v}_{t+1}^i = E_{t-1}^i \bar{v}_t^i + (t-1)^{-1} (\bar{v}_{t-1}^i - E_{t-1}^i \bar{v}_t^i) \quad (83)$$

Similar to the simulation with the heuristic-switching model, I base the simulation on the two initial values of the outcomes \bar{v}_t^i, p_t^{*j} and the submitted forecasts that generated these outcomes. From period 3, expectations and the corresponding outcomes are generated by least square learning. Figure 13, showing the outcomes for aggregate \bar{v} and p^* respectively simulated by least square learning, indicates that least square learning outperforms the heuristic-switching model in predicting the different convergence patterns on average between Treatment 2 ($\phi_\pi = 1.5$) and Treatment 3 ($\phi_\pi = 3.0$).

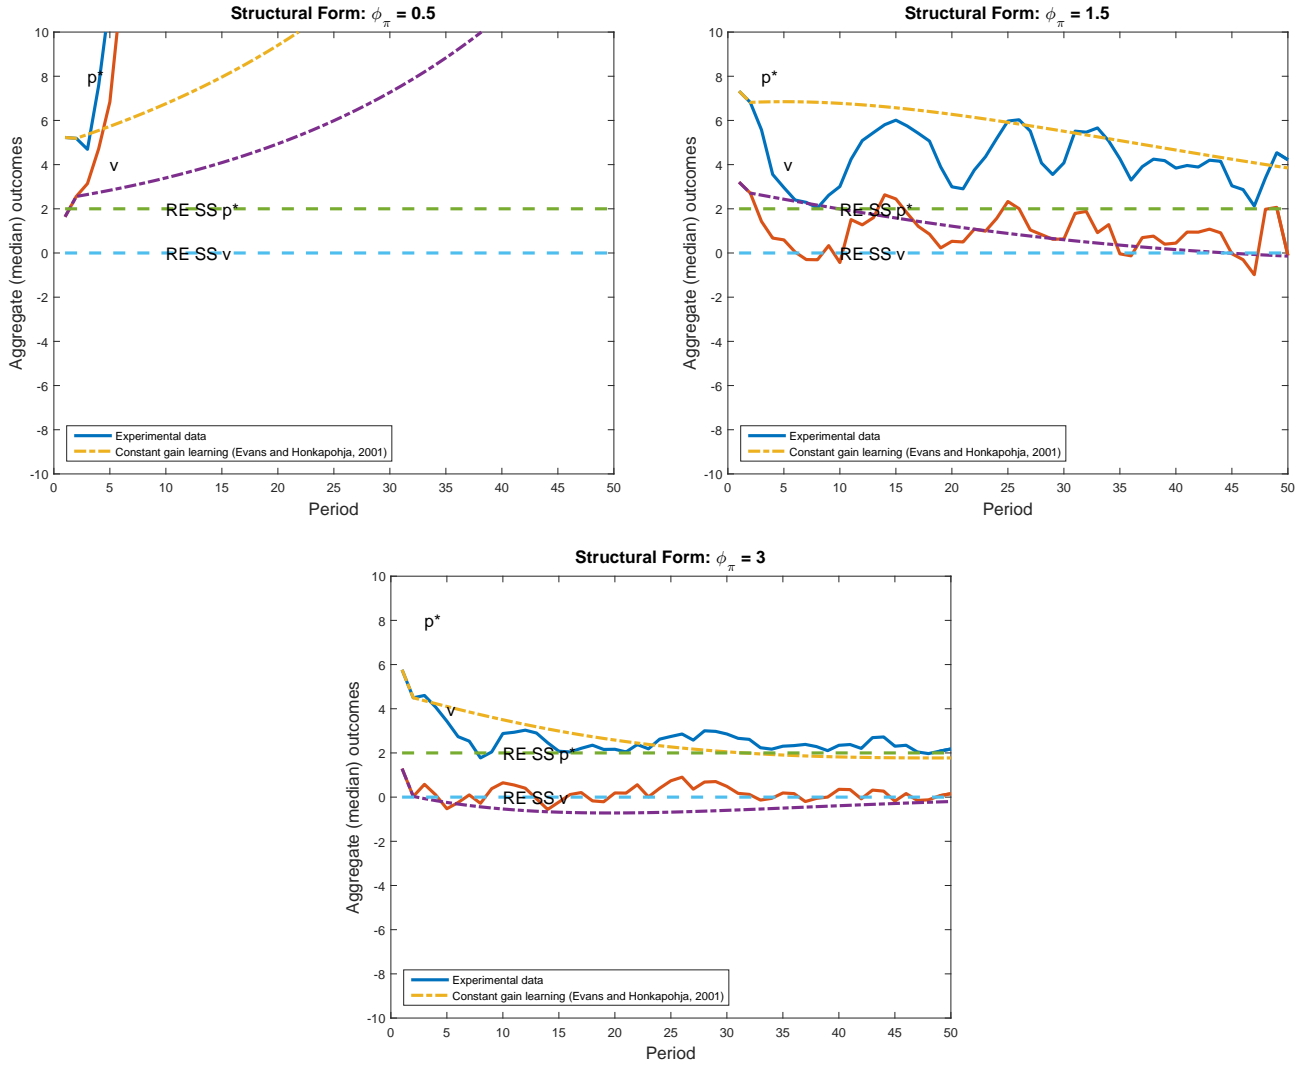


Figure 13: Actual (mean over all experimental groups) outcomes and simulated (mean over 6,000 replications) aggregated outcomes using least squares

9.3.2 Constant gain learning

A similar learning mechanism to recursive least square learning is constant gain learning, which geometrically discounts old observations instead of attaching equal weight to all observations. This mechanism can be viewed as approximately optimal, if agents suspect living in an unstable environment with parameter drift as shown by Evans et al. (2010). The updating equation, which would be analogous to the least square case in (83), would be

$$E_t^i \bar{v}_{t+1}^i = E_{t-1}^i \bar{v}_t^i + \gamma(\bar{v}_{t-1}^i - E_{t-1}^i \bar{v}_t^i) \quad (84)$$

Assuming that the learning parameter γ is stable across subjects and treatments yields an estimate of 0.11.

Figure 14, depicting the simulations of the constant-gain model with these estimates, shows that the constant gain version also predicts the difference between treatment 2 ($\phi_\pi = 1.5$) and treatment 3 ($\phi_\pi = 3$).

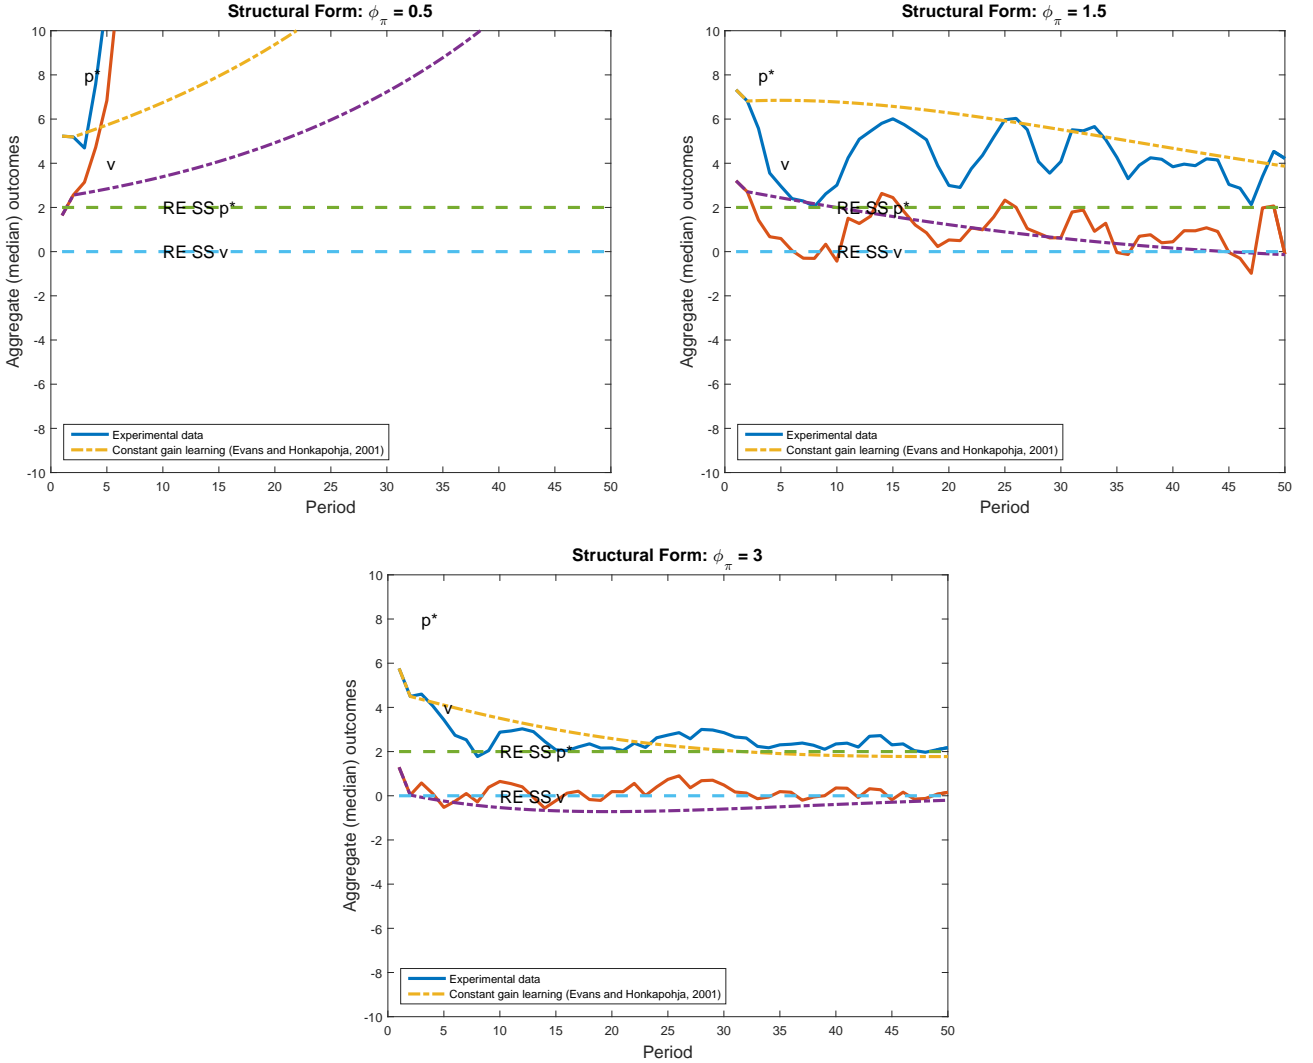


Figure 14: Actual (mean over all experimental groups) outcomes and simulated (mean over 6,000 replications) aggregated outcomes using constant gain

9.4 Heuristic-switching model

9.4.1 Simulation

The simulations are initialized by two initial values of the outcomes \bar{v}_t^i, p_t^{*j} respectively and initial weights $n_{t,h}^i$ which have been set to 0.25 for each subject and each heuristic h. With the initial values and the initial weights, the outcomes in periods 3 and 4 can be computed. From period 5, each individual expectation and outcomes is fully determined by the simulations. For the starting values of each experimental group, 6,000 replications of the simulation have been conducted.

For each replication the median outcomes and other the other interesting statistics have been calculated and subsequently averaged. Finally, I average over all experimental groups. Figure 15 depicts the averages of aggregate \bar{v} and p^* over 6,000 simulation replications for each group.

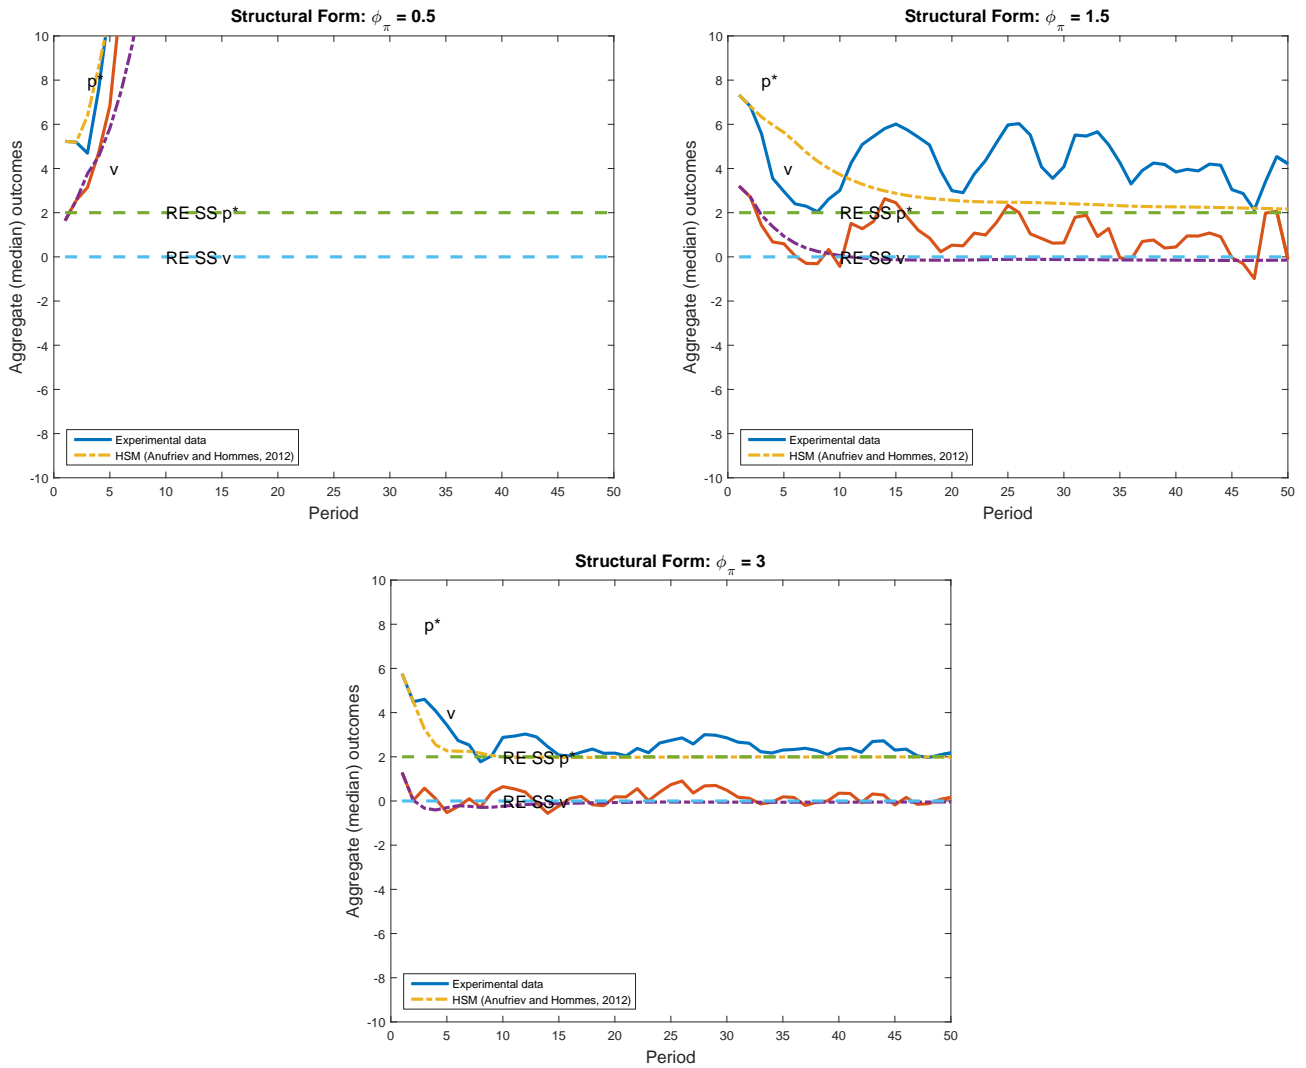


Figure 15: Actual (mean over all experimental groups) outcomes and simulated (mean over 2000 replications) aggregated outcomes using HSM

9.5 Single replications

9.5.1 Reinforcement

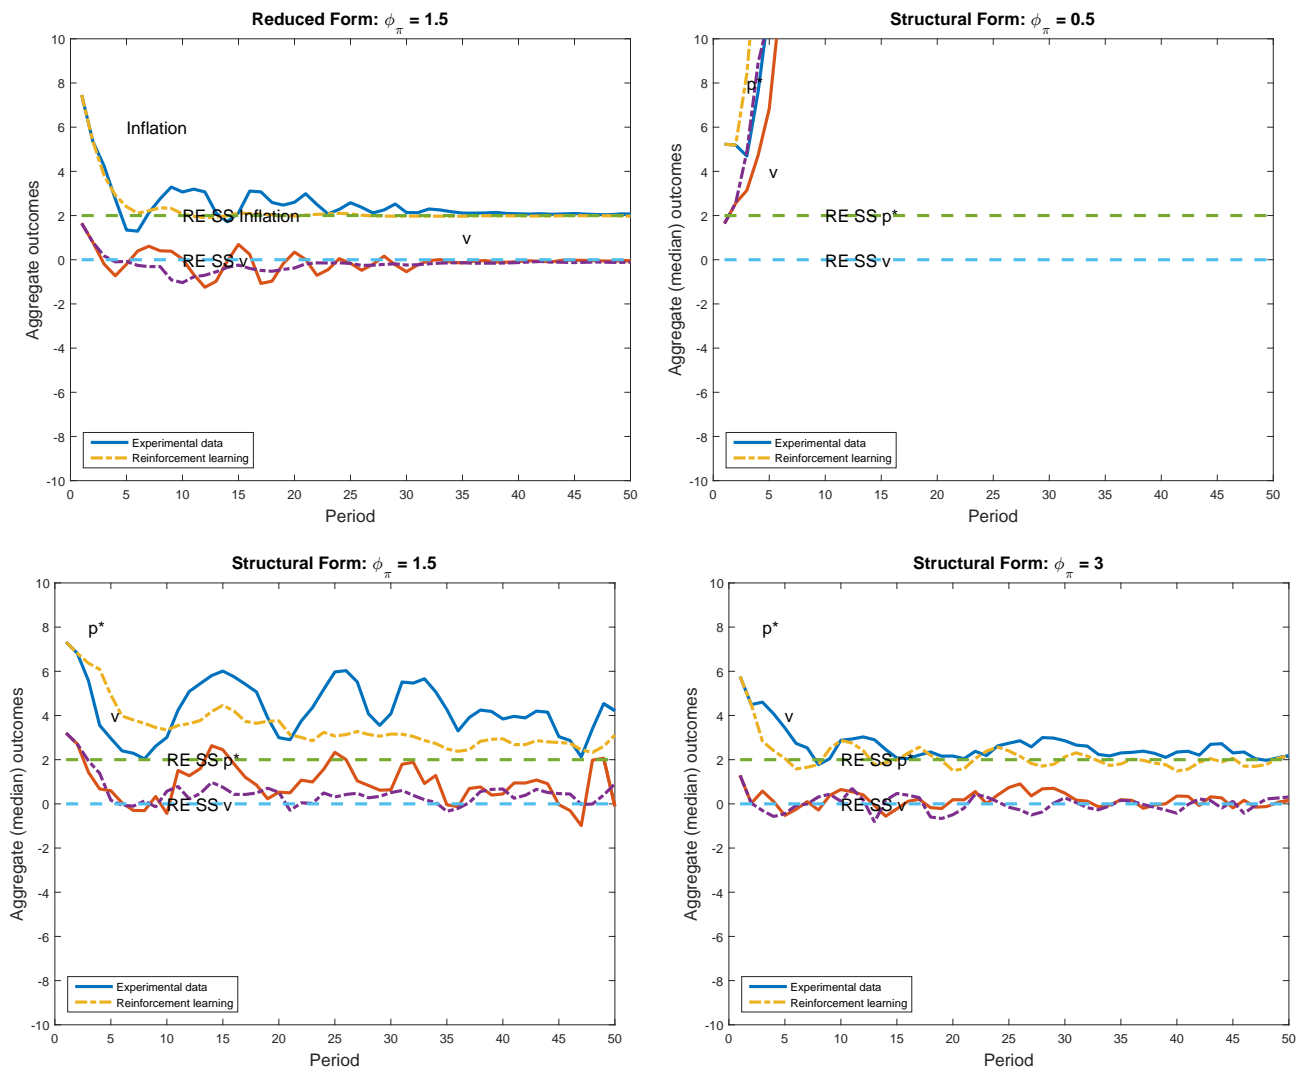


Figure 16: Actual (mean over all experimental groups) outcomes and simulated (one replication) aggregated outcomes using reinforcement

9.5.2 Noisy learning

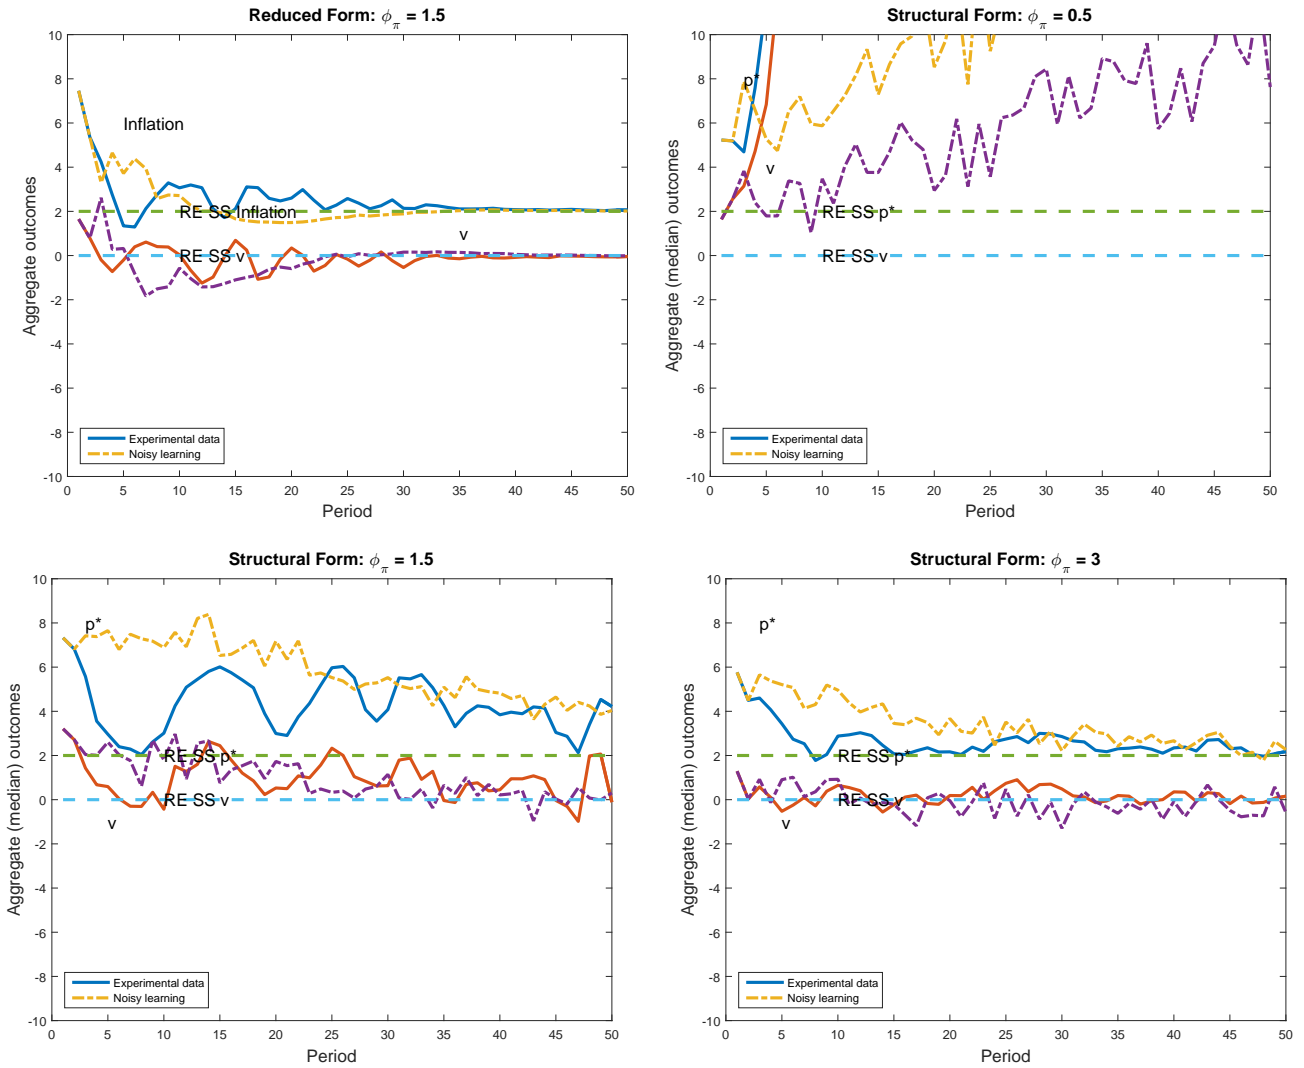


Figure 17: Actual (mean over all experimental groups) outcomes and simulated (one replication) aggregated outcomes using noisy learning

9.6 Endogenous gain

An intriguing question is whether other gain algorithms than constant gain improve the fit of noisy adaptive learning to the experimental data. The exposition is for (38), since changing the gain scheme γ_t in (37) does not lead to a better fit. I consider both decreasing gain and a more sophisticated endogenous gain algorithm by Marcet and Nicolini (2003).

9.6.1 Decreasing gain

A plausible mechanism is a gain of $(t - 1)^{-1}$, which could be interpreted as subjects making high adjustments at the beginning of the experiment, which are, however, decreasing, as subjects gain

more experience. Under this mechanism, the only parameter that needs to be estimated, the variance coefficient, is given by approximately $\hat{\omega} = 0.58$.

9.6.2 Marcet and Nicolini (2003)

Marcet and Nicolini (2003) relate the gain to the past observed forecast errors, which is, however, at the expense of tractability. Since subjects directly observe the profit as a non-linear transformation of the forecast error, I specify their endogenous gain algorithm as

$$\gamma_t = \begin{cases} \frac{1}{\bar{\gamma}^{-1}+h} & \text{if } \frac{\sum_{k=t-J}^{t-1} \text{Profit}_k^i}{J} \geq v \\ \bar{\gamma} & \text{if } \frac{\sum_{k=t-J}^{t-1} \text{Profit}_k^i}{J} < v \end{cases} \quad (85)$$

where h denotes the number of periods since the last switch to decreasing gain, J is the window length of past forecast errors considered, v is an arbitrary cutoff point and $\bar{\gamma}$ is the threshold gain once subjects switch to a constant gain algorithm. Milani (2014) assumes that v is endogenously given by the mean of the past forecast errors. I further generalize this mechanism by assuming that

$$v = \frac{\sum_{l=t-W}^{t-1} \text{Profit}_l^i}{W} \quad (86)$$

so that subjects have a certain window length of their profits as a benchmark. Estimating the parameters of noisy learning with endogenous gain using maximum likelihood gives $\hat{\gamma} = 0.24$, $\hat{\omega} = 0.59$, $\hat{J} = 1$ and $\hat{W} = 3$.³⁸

Table 14, reporting several measures, shows that decreasing gain considerably worsens the fit, while the fit with the endogenous gain leads to similar results to the constant gain case. The latter is, however, preferable, for at least two reasons: first, it retains tractability; secondly, since constant gain noisy learning and endogenous gain noisy learning are both stochastic models whose parameters are estimated using maximum likelihood, a direct likelihood-comparison is possible. The likelihood of endogenous gain learning is -25,611, while the likelihood of constant gain learning is -25,542 so that constant gain learning yields a higher likelihood despite having fewer parameters.

³⁸Using an unrestricted maximum likelihood yields an estimate for ω outside the unit circle. Hence, J and W were chosen as to obtain the maximum likelihood that satisfies the restriction that ω lies inside the unit circle.

Measures	$\phi_x = 0.5$				$\phi_x = 1.5$				$\phi_x = 3.0$			
	<i>Data</i>	Noisy (constant gain)	Noisy (decreasing gain)	Noisy (endogenous gain)	<i>Data</i>	Noisy (constant gain)	Noisy (decreasing gain)	Noisy (endogenous gain)	<i>Data</i>	Noisy (constant gain)	Noisy (decreasing gain)	Noisy (endogenous gain)
Aggregates												
MSE: v	0.00	271,444.44	275986.45	269551.22	0.00	5.03	6.16	5.46	0.00	1.12	1.21	1.12
MSE: p^*	0.00	517,036.07	527092.08	513705.65	0.00	13.37	20.52	15.84	0.00	1.74	1.25	1.17
Mean: v	377.43	8.89	4.42	9.71	0.96	0.91	1.51	0.79	0.17	-0.24	-0.48	-0.35
Mean: p^*	562.20	15.90	8.99	17.47	4.31	5.60	6.73	5.65	2.63	3.11	3.30	2.77
MSD from REE: v	268,390.11	140.60	26.90	178.94	5.86	3.93	5.74	3.95	0.71	1.69	1.52	1.52
MSD from REE: p^*	514,924.42	317.16	57.09	414.36	16.27	18.20	29.33	19.86	1.52	3.40	2.86	2.42
Standard dev.: v	311.47	6.36	2.29	7.47	1.80	1.48	1.21	1.49	0.71	1.18	1.02	1.09
Standard dev.: p^*	434.66	9.09	2.36	10.66	2.55	1.62	1.10	1.73	0.96	1.41	0.96	1.32
Individuals												
MSE: v^i	0.00	346,757.07	352,235.36	345,435.07	0.00	7,050.73	7,037.07	7,067.46	0.00	6,889.12	6,867.46	6,877.79
MSE: p^{*j}	0.00	507,851.64	517,248.50	504,815.72	0.00	25.19	33.53	26.33	0.00	9.53	10.70	8.21
MSD from REE: v^i	343,486.43	552.14	180.46	256.17	6,826.17	94.03	83.12	82.58	6,603.38	204.94	83.73	38.87
MSD from REE: p^{*j}	505,904.75	392.91	82.36	368.98	24.96	29.15	51.78	38.40	9.25	12.39	16.10	12.56
Dispersion of v^i	278.19	14.87	10.44	14.20	39.43	7.93	7.08	7.90	30.31	7.36	6.27	7.34
Dispersion of p^{*j}	160.10	5.88	4.76	5.69	2.52	3.16	3.48	2.92	1.69	2.74	3.05	2.41

MSE=mean square error; MSD=mean square distance

Table 13: Measures - different versions of noisy learning - all periods

Measures	$\phi_x = 0.5$				$\phi_x = 1.5$				$\phi_x = 3.0$			
	<i>Data</i>	Noisy (constant gain)	Noisy (decreasing gain)	Noisy (endogenous gain)	<i>Data</i>	Noisy (constant gain)	Noisy (decreasing gain)	Noisy (endogenous gain)	<i>Data</i>	Noisy (constant gain)	Noisy (decreasing gain)	Noisy (endogenous gain)
Aggregates												
MSE: v	0.00	329,792.12	338,843.03	326,336.02	0.00	4.96	6.92	5.44	0.00	0.79	1.00	0.68
MSE: p^*	0.00	557,155.69	577,102.62	551,316.62	0.00	10.26	21.15	12.29	0.00	0.47	0.78	0.58
Mean: v	494.84	14.19	5.01	15.96	0.72	0.12	1.27	-0.04	0.09	-0.50	-0.65	-0.38
Mean: p^*	707.07	24.29	9.98	27.33	4.09	4.54	6.56	4.37	2.34	2.09	2.91	1.90
MSD from REE: v	343,633.59	279.97	33.15	368.04	4.98	1.95	5.00	2.14	0.35	0.98	1.26	0.81
MSD from REE: p^*	588,244.14	640.78	72.56	855.52	14.57	10.94	28.74	10.82	0.52	0.61	1.70	0.49
Standard dev.: v	241.03	5.72	2.19	6.81	1.44	0.98	1.05	1.04	0.55	0.67	0.73	0.63
Standard dev.: p^*	272.92	7.20	2.08	8.58	1.92	0.98	0.90	1.06	0.52	0.57	0.62	0.51
Individuals												
MSE: v^i	0.00	426,232.66	436,684.21	423,425.91	0.00	1,808.37	1,771.93	1,789.17	0.00	2,139.97	2,136.85	2,143.16
MSE: p^{*j}	0.00	584,959.77	603,807.08	579,332.33	0.00	14.71	31.84	16.06	0.00	3.33	8.18	3.08
MSD from REE: v^i	444,173.06	740.93	162.79	832.62	1,831.77	88.56	63.56	96.73	2,131.86	87.67	44.54	96.00
MSD from REE: p^{*j}	616,582.76	694.30	105.86	919.18	18.85	17.40	45.69	15.95	2.18	3.68	11.15	2.61
Dispersion of v^i	360.54	16.44	10.02	17.18	25.06	7.87	6.85	8.10	20.43	7.08	5.88	7.30
Dispersion of p^{*j}	221.88	6.20	4.35	6.63	1.91	1.87	2.92	1.75	0.90	1.52	2.56	1.26

MSE=mean square error; MSD=mean square distance

Table 14: Measures - different versions of noisy learning - periods 30-50

9.7 An alternative version of noisy learning?

A natural question is why the noisy learning model should not be formulated the following way:

$$\hat{E}_t x_{t+1} = \hat{E}_{t-2} x_{t-1} + \gamma(x_{t-1} - \hat{E}_{t-2} x_{t-1}) + \eta_t \quad (87)$$

with $\eta_t \sim N(0, \sigma_t^2)$ and the variance being endogenous

$$\sigma_t^2 = (1 - \omega)\sigma_{t-1}^2 + \omega(x_{t-1} - \hat{E}_{t-2} x_{t-1})^2 \quad (88)$$

Proposition 16. σ_t^2 has a unit root in this case.

Proof. Assume the converse, i.e. that σ_t^2 is stationary so that the long-run variance $\mathbb{E}(\sigma_t^2) \equiv \sigma^2$ is well-defined. Evaluate (87) at $t-2$ so that we have

$$\hat{E}_{t-2}x_{t-1} = \hat{E}_{t-4}x_{t-3} + \gamma(x_{t-3} - \hat{E}_{t-4}x_{t-3}) + \eta_{t-2} \quad (89)$$

Using (89) in (88) and using the expectations operator gives

$$\underbrace{\mathbb{E}(\sigma_t^2)}_{=\sigma^2} = (1 - \omega) \underbrace{\mathbb{E}(\sigma_{t-1}^2)}_{=\sigma^2} + \omega \left(\underbrace{\mathbb{E}(\eta_{t-2}^2)}_{=\mathbb{E}(\sigma_{t-2}^2)=\sigma^2} + \text{TERMS...} \right) \quad (90)$$

This shows that the long-run variance $\mathbb{E}(\sigma_t^2) \equiv \sigma^2$ is not well-defined, which is a contradiction. \square

9.8 Tables for all learning models

9.8.1 All periods

<u>Measures</u>	$\phi_\pi = 0.5$					
	<u>Data</u>	<u>HSM</u>	<u>OLS</u>	<u>CG</u>	<u>Reinforcement</u>	<u>Noisy</u>
<u>Aggregates</u>						
MSE: v	<i>0.00</i>	366,507.77	264,495.34	263,408.09	149,720.10	271,444.44
MSE: p^*	<i>0.00</i>	554,832.90	504,139.91	501,642.93	305,299.15	517,036.07
Mean: v	<i>377.43</i>	557.42	5.98	7.09	485.33	8.89
Mean: p^*	<i>562.20</i>	765.16	11.12	12.84	725.75	15.90
MSD from REE: v	<i>268,390.11</i>	630,093.47	50.17	84.95	443,216.51	140.60
MSD from REE: p^*	<i>514,324.42</i>	1,123,272.73	111.03	197.29	925,942.38	317.16
Standard dev.: v	<i>311.47</i>	501.38	0.74	1.85	402.38	6.36
Standard dev.: p^*	<i>434.66</i>	649.56	0.91	2.54	569.56	9.09
<u>Individuals</u>						
MSE: v^i	<i>0.00</i>	374,015.88	338,145.43	336,934.82	142,536.97	346,757.07
MSE: p^{*j}	<i>0.00</i>	558,014.57	495,119.52	492,604.48	310,652.46	507,851.64
MSD from REE: v^i	<i>343,486.43</i>	614,143.10	84.83	120.99	674,365.20	552.14
MSD from REE: p^{*j}	<i>505,904.75</i>	1,120,158.18	132.25	223.09	1,320,941.26	392.91
Dispersion of v^i	<i>278.19</i>	78.59	4.91	4.91	219.71	14.87
Dispersion of p^{*j}	<i>160.10</i>	22.67	2.98	3.21	42.83	5.88

Table 15: All periods for $\phi_\pi = 0.5$

	$\phi_\pi = 1.5$					
<u>Measures</u>	<u>Data</u>	<u>HSM</u>	<u>OLS</u>	<u>CG</u>	<u>Reinforcement</u>	<u>Noisy</u>
<u>Aggregates</u>						
MSE: v	<i>0.00</i>	5.74	5.97	6.44	4.98	5.03
MSE: p^*	<i>0.00</i>	14.54	19.79	18.18	12.05	13.37
Mean: v	<i>0.96</i>	0.13	0.97	1.04	0.29	0.91
Mean: p^*	<i>4.31</i>	3.06	5.92	5.71	3.22	5.60
MSD from REE: v	<i>5.86</i>	0.72	2.90	3.88	2.90	3.93
MSD from REE: p^*	<i>16.27</i>	3.83	24.25	22.01	8.98	18.20
Standard deviation: v	<i>1.80</i>	0.87	0.55	0.83	1.16	1.48
Standard deviation: p^*	<i>2.55</i>	1.54	0.62	1.14	1.82	1.62
<u>Individuals</u>						
MSE: v^i	<i>0.00</i>	7,090.74	6,741.29	6,736.73	4,867.40	7,050.73
MSE: p^{*j}	<i>0.00</i>	19.26	26.08	26.49	17.45	25.19
MSD from REE: v^i	<i>6,826.17</i>	51.54	30.50	31.54	351.39	94.03
MSD from REE: p^{*j}	<i>24.96</i>	10.06	34.36	34.95	9.61	29.15
Dispersion of v^i	<i>39.43</i>	6.09	5.30	5.29	21.12	7.93
Dispersion of p^{*j}	<i>2.52</i>	1.22	2.36	2.54	1.49	3.16

Table 16: All periods for $\phi_\pi = 1.5$

	$\phi_\pi = 3.0$					
<u>Measures</u>	<u>Data</u>	<u>HSM</u>	<u>OLS</u>	<u>CG</u>	<u>Reinforcement</u>	<u>Noisy</u>
<u>Aggregates</u>						
MSE: v	<i>0.00</i>	0.74	1.45	1.30	0.65	1.12
MSE: p^*	<i>0.00</i>	0.89	0.81	1.08	0.92	1.74
Mean: v	<i>0.17</i>	-0.07	-0.60	-0.45	-0.03	-0.24
Mean: p^*	<i>2.63</i>	2.17	2.41	2.62	2.17	3.11
MSD from REE: v	<i>0.71</i>	0.14	0.55	0.45	0.61	1.69
MSD from REE: p^*	<i>1.52</i>	0.20	0.46	1.13	0.91	3.40
Standard deviation: v	<i>0.71</i>	0.41	0.45	0.46	0.63	1.18
Standard deviation: p^*	<i>0.96</i>	0.75	0.87	1.08	0.93	1.41
<u>Individuals</u>						
MSE: v^i	<i>0.00</i>	6,840.40	6,569.41	6,569.64	5,017.29	6,889.12
MSE: p^{*j}	<i>0.00</i>	4.46	7.34	8.76	5.04	9.53
MSD from REE: v^i	<i>6,603.38</i>	26.50	17.00	16.95	136.98	204.94
MSD from REE: p^{*j}	<i>9.25</i>	4.65	9.60	13.14	4.08	12.39
Dispersion of v^i	<i>30.31</i>	4.94	4.34	4.34	13.10	7.36
Dispersion of p^{*j}	<i>1.69</i>	0.98	2.25	2.43	1.13	2.74

Table 17: All periods for $\phi_\pi = 3.0$

9.8.2 Periods 30-50

<u>Measures</u>	$\phi_\pi = 0.5$					
	<u>Data</u>	<u>HSM</u>	<u>OLS</u>	<u>CG</u>	<u>Reinforcement</u>	<u>Noisy</u>
<u>Aggregates</u>						
MSE: \bar{v}	<i>0.00</i>	526,997.96	337,958.89	334,883.59	201,402.64	329,792.12
MSE: p^*	<i>0.00</i>	689,847.60	573,762.16	566,270.92	383,002.42	557,155.69
Mean: \bar{v}	<i>494.84</i>	1,050.05	7.39	10.97	847.00	14.19
Mean: p^*	<i>707.07</i>	1,389.42	13.41	19.13	1,247.83	24.29
MSD from REE: \bar{v}	<i>343,633.59</i>	1,287,328.10	71.82	165.35	847,755.38	279.97
MSD from REE: p^*	<i>588,244.14</i>	2,216,424.87	164.33	396.96	1,735,919.83	640.78
Standard dev.: \bar{v}	<i>241.03</i>	119.60	0.30	1.93	196.44	5.72
Standard dev.: p^*	<i>272.92</i>	122.23	0.47	3.03	188.03	7.20
<u>Individuals</u>						
MSE: \bar{v}^i	<i>0.00</i>	507,459.91	436,485.68	433,011.44	176,738.81	426,232.66
MSE: p^{*j}	<i>0.00</i>	714,967.94	600,927.47	593,717.51	400,773.41	584,959.77
MSD from REE: \bar{v}^i	<i>444,173.06</i>	1,235,593.52	106.63	204.42	823,199.31	740.93
MSD from REE: p^{*j}	<i>616,582.76</i>	2,211,486.29	177.38	407.29	1,713,221.67	694.30
Dispersion of \bar{v}^i	<i>360.54</i>	94.59	4.79	4.74	299.81	16.44
Dispersion of p^{*j}	<i>221.88</i>	12.41	2.20	1.61	46.21	6.20

Table 18: Periods 30-50

	$\phi_\pi = 1.5$					
<u>Measures</u>	<u>Data</u>	<u>HSM</u>	<u>OLS</u>	<u>CG</u>	<u>Reinforcement</u>	<u>Noisy</u>
<u>Aggregates</u>						
MSE: v	<i>0.00</i>	5.32	6.28	5.75	4.95	4.96
MSE: p^*	<i>0.00</i>	13.43	20.41	16.75	12.09	10.26
Mean: v	<i>0.72</i>	-0.15	0.55	0.19	0.00	0.12
Mean: p^*	<i>4.09</i>	2.27	5.46	4.68	2.49	4.54
MSD from REE: v	<i>4.98</i>	0.18	1.46	0.64	3.77	1.95
MSD from REE: p^*	<i>14.57</i>	0.50	21.30	14.33	8.60	10.94
Standard deviation: v	<i>1.44</i>	0.21	0.07	0.24	0.52	0.98
Standard deviation: p^*	<i>1.92</i>	0.32	0.14	0.59	0.65	0.98
<u>Individuals</u>						
MSE: v^i	<i>0.00</i>	1,834.55	1,808.21	1,814.29	1,984.51	1,808.37
MSE: p^{*j}	<i>0.00</i>	17.01	24.46	19.92	16.04	14.71
MSD from REE: v^i	<i>1,831.77</i>	46.56	27.67	26.40	2,476.59	88.56
MSD from REE: p^{*j}	<i>18.85</i>	0.88	24.94	16.24	9.66	17.40
Dispersion of v^i	<i>25.06</i>	5.71	5.23	5.18	21.80	7.87
Dispersion of p^{*j}	<i>1.91</i>	0.50	1.73	1.27	0.84	1.87

Table 19: Periods 30-50

$\phi_\pi = 3.0$						
<u>Measures</u>	<u>Data</u>	<u>HSM</u>	<u>OLS</u>	<u>CG</u>	<u>Reinforcement</u>	<u>Noisy</u>
<u>Aggregates</u>						
MSE: v	<i>0.00</i>	0.37	0.92	0.67	0.35	0.79
MSE: p^*	<i>0.00</i>	0.53	0.62	0.68	0.53	0.47
Mean: v	<i>0.09</i>	-0.05	-0.60	-0.40	-0.03	-0.50
Mean: p^*	<i>2.34</i>	1.99	1.99	1.86	1.98	2.09
MSD from REE: v	<i>0.35</i>	0.02	0.43	0.21	0.69	0.98
MSD from REE: p^*	<i>0.52</i>	0.02	0.09	0.07	1.12	0.61
Standard deviation: v	<i>0.55</i>	0.08	0.03	0.14	0.33	0.67
Standard deviation: p^*	<i>0.52</i>	0.08	0.06	0.12	0.35	0.57
<u>Individuals</u>						
MSE: v^i	<i>0.00</i>	2,116.61	2,119.46	2,118.82	1,839.98	2,139.97
MSE: p^{*j}	<i>0.00</i>	2.02	5.58	4.01	2.11	3.33
MSD from REE: v^i	<i>2,131.86</i>	23.87	16.43	16.01	2,058.74	87.67
MSD from REE: p^{*j}	<i>2.18</i>	0.14	4.19	2.37	1.50	3.68
Dispersion of v^i	<i>20.43</i>	4.65	4.30	4.25	15.09	7.08
Dispersion of p^{*j}	<i>0.90</i>	0.27	1.67	1.22	0.48	1.52

Table 20: Periods 30-50