Expectations and Social Influence
Dynamics in Economics

ExSIDE

ExSIDE Working Paper Series

No. 40- 2021
June 2021

Planar Learning to Forecast Market Games

Myrna Hennequin
Cars Hommes
Eva Levelt

www.exside-itn.eu

Marie Sklodowska-Curie Innovative Training Network "Expectations and Social Influence Dynamics in Economics (ExSIDE)"
This project has received funding from the European Union Horizon 2020 Research and Innovation action under grant agreement No 721846.

Planar Learning to Forecast Market Games

Myrna Hennequin, Cars Hommes & Eva Levelt

June 30, 2021

Abstract

In a system where expectations and realisations of a price feed back to each other, it has been found that the sign and strength of this feedback is an important predictor of the market stability. In this paper we contribute to the generalisation of this result to a two dimensional system, where the expectations and realisations of two prices affect each other. We conduct a Learning to Forecast Experiment and show that eigenvalues can be used as predictors of stability in such a higher dimensional framework. We investigate eigenvalues of positive real part, and found that complex eigenvalues with a polar angle of 45° lead to more stable dynamics than real eigenvalues with the same absolute value. For the real eigenvalues we find a change from stable to unstable dynamics inside the unit circle, which is in line with the findings from the one dimensional case. In order to reproduce the decisions being made in individual time steps, simple models like an adaptive rule are often sufficient. In order to reproduce long run dynamics, we develop a two dimensional Heuristic Switching Model. We use this model to predict the stability of systems with eigenvalues which we did not test experimentally.

This work has received funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skodowska-Curie grant agreement No 721846, Expectations and Social Influence Dynamics in Economics (ExSIDE).

Contents

1 Introduction **3**

2 Experimental set-up **5**

3 Aggregate dynamics **9**

4 Models of expectation formation **15**

5 Long run behaviour **21**

6 Robustness check: partial derivatives **27**

7 Conclusion **30**

A Instructions **34**

B Prediction error to penalty point conversion **36**

C Control Questions **37**

D Maps for the various treatments **37**

1 Introduction

The stability of equilibria in an economy depends on the way agents form expectations. Lucas (1986) points out that this stability is an empirical question:

I think an appropriate stability theory can be useful in weeding out these uninteresting equilibria ... But to be useful, stability theory must be more than simply a fancy way of saying that one does not want to think about certain equilibria. I prefer to view it as an experimentally testable hypothesis, as a special instance of the adaptive laws that we believe govern all human behavior. (Lucas, 1986, pp. S424 - 425).

Experiments with expectation feedback models have been run to determine stability of different equilibria. Subjects submit an expectation of some variable and these expectations determine the realisation of the variable. When the subjects then observe this realisation they (may) use this in forming their next expectation, creating a feedback loop. It has been found that the strength of feedback from expectations to realisations is an essential predictor of aggregate market behaviour (Sonnemans and Tuinstra, 2010). Positive feedback tends to result in large fluctuations and persistent deviations from the equilibrium, even when the absolute value is smaller than 1 . Negative feedback, however, often results in convergence to the equilibrium consistent with the rational expectations hypothesis, even when the absolute value is larger than 1 (Bao and Hommes, 2019).

The economy, however, is a higher dimensional system. The expectations people have of one variable tend to depend on various other variables, and the realisation of a variable depends on expectations of various others. In this work we investigate how the results on feedback strength generalise to higher dimensions. The higher dimensional generalisation of feedback strength is the eigenvalue. The eigenvalues determine the dynamical properties of the system when a model of expectation formation is assumed. We use the simple case of naive expectations. One would then expect in theory, that when this process of expectation formation is used, eigenvalues inside the unit circle lead to convergence. In particular, for real eigenvalues this would happen in a monotonic way and for complex eigenvalues in an oscillatory way. However, expectations are often not naive so dynamics may be different given the actual behaviour of participants. We use this experiment to find out if eigenvalues can still be used as predictors of stability.

When concerned with two dimensional models a first obvious question is whether, as in the one dimensional setting, it is possible for participants to learn the equilibrium of the market despite receiving only qualitative information. Then we ask if there are also eigenvalues inside the unit circle for which the aggregate market behaviour is unstable. In two dimensions, however, we also have the added possibilities of complex eigenvalues and saddle points. We have conducted experiments to study how aggregate market behaviour in a two dimensional model depends on the eigenvalues of the underlying expectation feedback system.

An important example of a two dimensional expectation feedback model in

macroeconomics is the New Keynesian (NK) model. Expectations of inflation and the output gap both determine the realisations of the two variables, which again affect future expectations. Several experiments have been conducted for the two-dimensional NK model, where subjects form expectations on inflation and output gap (Assenza et al. (2021), Hommes et al. (2019b)).[1]

Their findings suggest that designing monetary policy to temper the size of the eigenvalues of the system may help stabilize macro-economic variables. To better understand the role of eigenvalues, however, it is useful to systematically vary the eigenvalues in a simple abstract expectations feedback setting.

Anufriev et al. (2019) study the role of eigenvalues in the context of a planar p-guessing game, where the subjects predict one step ahead. In four treatments with four different linear models they examine the dynamics when the fixed point is a sink, a source or a saddle point with positive or negative leading eigenvalue. They only consider real eigenvalues. They give participants the precise laws of motion and only have interaction of the variables in one dimension.

In contrast, in our experiments we used a temporary equilibrium set-up where predictions are two steps ahead. This creates the possibility for rational bubble solutions as presented in Hommes et al. (2005). Our participants have only qualitative information about the market they are engaged in, mimicking real world conditions. Considering the important example of the NK model which only admits eigenvalues with positive real part we only examine this type of eigenvalues, and we examine both real and complex eigenvalues.

In our general LtFE set-up we examine whether eigenvalues can be used as predictors of market stability, both by comparing markets with different eigenvalues and by comparing markets with the same eigenvalues but with differences in the positive/negative feedback between variables. We find that eigenvalues can be used, but that simply focusing on the absolute value of the leading eigenvalue is not sufficient. The treatments with complex eigenvalues lead to more stable markets than those with real eigenvalues of the same absolute value. Explaining this stabilising effect of complex eigenvalues leads us to the development of a more sophisticated two dimensional Heuristic Switching Model (Brock and Hommes, 1997; Anufriev and Hommes, 2012) than has been used in the literature so far. We use this to predict the stability of markets with eigenvalues that we have not ourselves examined in the laboratory.

We also compare various models of expectation formation in their ability to reproduce the decisions made by participants at the individual level. We find that especially for the treatments with real eigenvalues, an adaptive model with only one parameter outperforms more involved options. For the complex treatments there is gain to be found in using multi parameter models, but the simpler adaptive model still performs very well.

This paper is structured as follows. We begin by presenting the design

[1] Kryvtsov and Petersen (2013), Arifovic and Petersen (2017) Pfajfar and Žakelj (2018), and Hommes et al. (2019a) also conduct LtFEs for the NK model. In these experiments the settings are more complicated which make them less suitable for consideration of the role of eigenvalues. For example, they have a second order model or apply expectational shocks with autocorrelations.

of the experiment we conducted. Section 3 compares the observed aggregate market behaviour in the different treatments. In Section 4 we investigate which learning models best replicate the decisions the participants made in the lab. Section 5 develops a new two dimensional HSM in order to be able to predict the qualitative market dynamics ex ante. We lay out the evidence for using eigenvalues rather than only positive/negative feedback as predictors in Section 6, and Section 7 concludes.

2 Experimental set-up

In the experiment individuals made decisions privately on a computer. They were randomly divided into groups of 6 and repeatedly submitted two numbers: a prediction for the price of some good x and a prediction for the price of some good y. The expectations of all members in the group determined what the next prices were. The participants did not know exactly how the prices were calculated, they only had qualitative descriptions of the relationships between predictions and realisations, namely whether the expectations feedback is positive or negative . An example of a chosen model, for eigenvalues $(0.95, 0.5)$ is:

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.725 & 0.225 \\ 0.225 & 0.725 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 0.65 \\ 0.65 \end{pmatrix}$$

Once the prices had been determined they were revealed to all participants before they made their next predictions. This process was repeated for fifty time-steps. The experimental instructions presented to the subjects can be found in Appendix A

Predictions were always made for time $t+1$, before the prices at time t were revealed. Subjects had all prices up till time $t-1$ available to them when forming their expectations. The first two sets of predictions therefore had to be made without any observed prices. To give some guidance, these initial predictions all had to be between 0 and 20. After the predictions of the first two time steps, the first prices were revealed. From then onwards the participants were only told that their predictions had to be non-negative. If, however, they attempted to submit predictions higher than 100, they were informed of an upper bound of 100.

Participants were rewarded based on the accuracy of their predictions. Earnings depended on both variables and every round equally. We determined the prediction error in each variable, added them up to find the total prediction error, and then use a hyperbolic pay-off function similar to Hommes et al. (2019a):

$$\frac{100}{1 + 0.5(|x - x_e| + |y - y_e|)}$$

One 'point' from the hyperbolic pay-off function was worth 1.2 eurocents. To ensure no one was under-compensated for their time (approximately 2 hours) we added a participation fee of €10.

Given the literature on loss aversion we framed the pay-off as a loss rather than a gain (Kahneman and Tversky, 1984). Participants were informed at the beginning of the experiment of the maximal earnings and were told that these earnings would shrink when prediction errors were made. The pay-off function is not affected by this framing but the instructions are different from those in other similar experiments. See Appendix A.3 for this part of the instructions.

We chose the function that determines the prices based on the predictions to be linear. Understanding first a linear model is helpful since non-linear models will usually admit a linear approximation around a steady state. We also chose a model without noise since a two-dimensional setting could already be expected to be hard to understand even in the deterministic limit. We aimed for a setting that was as simple as possible.

2.1 Treatments

We used five different treatments to study the effect of eigenvalues of the expectations feedback model on market stability. The treatments differ in the linear functions from predictions to prices. These are the functions f which determine the relationship:

$$(x_t, y_t) = f(\bar{x}^e_{t+1}, \bar{y}^e_{t+1})$$

Note that the realized prices x_t and y_t both depend on the average forecasts $\bar{x}^e_{t+1}$ and $\bar{y}^e_{t+1}$ of all individuals.

We looked only at eigenvalues with positive real part. For the real pairs, we kept the non-leading eigenvalue constant at 0.5. For the leading eigenvalue we then tried 0.7, 0.95 and 1.1. With the first two we aim to discover whether we see a change from stable to unstable behaviour before leaving the unit circle, as we do in the one dimensional case. The final eigenvalue allows us to test the new possibility of a saddle point. For the complex eigenvalues we want to be able to make a comparison with dynamics under real eigenvalues with the same absolute value. We look at absolute values 0.7 and 0.95, both at angle $\frac{\pi}{4}$. We will refer to the three treatments with real eigenvalues as $0.7real$, $0.95real$, $1.1real$ and the two treatments with complex eigenvalues as $0.7complex$ and $0.95complex$ respectively. This gives us five treatments, as illustrated in Figure 2.

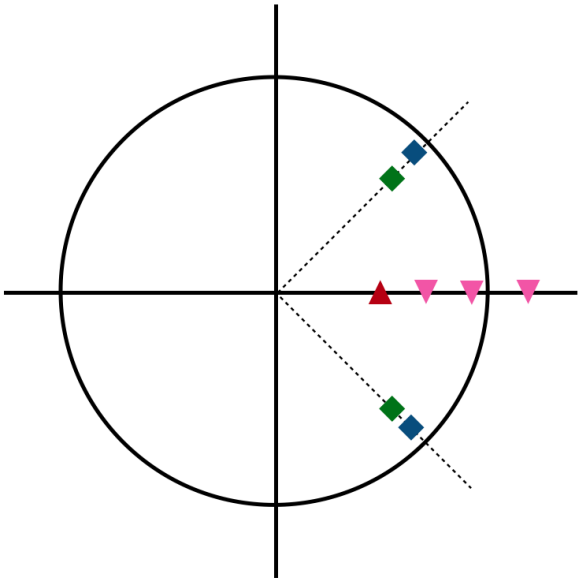

Figure 1: Eigenvalues in the five treatments. The red triangle at 0.5 pairs with the various pink ones (0.7, 0.95, 1.1) to form the real pairs, while the two green and two blue squares represent the complex pairs with the same absolute values of 0.7 and 0.95.

We chose the fixed point to be at $(13, 13)$ for all treatments. For the treatments with real eigenvalues we keep the eigenvectors constant as in Figure 2. By fixing this we have no further free parameters in the model and impose a consistent symmetry for simplicity. The dotted line, $x = y$, indicates the direction corresponding to the leading eigenvalue (so the solid line indicates the eigenvector corresponding to 0.5 every time). For the complex eigenvalues it is not possible to have these eigenvectors. We chose the models for the complex treatments to have one negative entry in the top-right position, since this setting arises in the linearised New-Keynesian model due to the negative feedback from the interest rate on the output gap (see Assenza et al. (2021).

Finally we performed a sixth treatment as a robustness check, $0.95 robust$, where the eigenvalues are the same as in $0.95 real$ but the model is different. This will be discussed in Section 6.

The models for all treatments can be found in Appendix D.

2.2 Limits

We limit predictions to $[0, 100]$, but for the $1.1 real$ treatment it is possible that the realised price falls outside that interval despite predictions being in it. To prevent frustration with participants we do not permit this, and change any price below 0 into 0 and above 100 into 100. In other words, given the functions

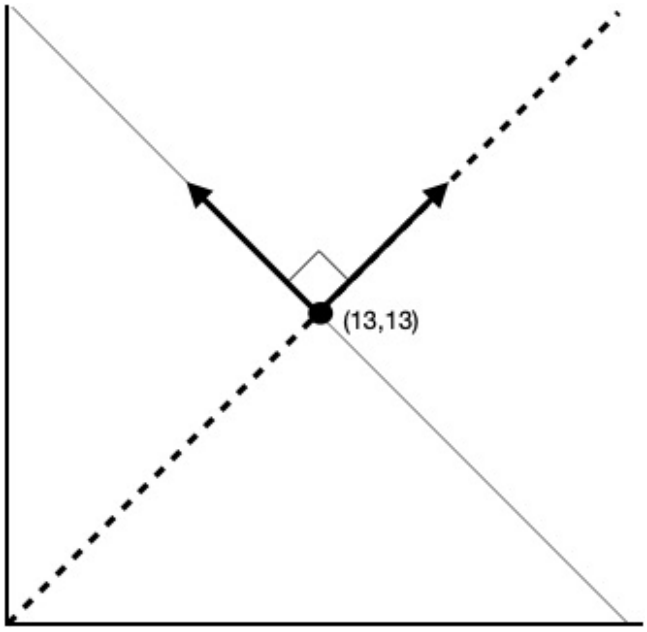

Figure 2: Eigenvectors kept fixed for all real treatments. The dotted line, $x = y$, indicates the direction corresponding to the leading eigenvalue. The solid line indicates the direction corresponding to eigenvalue 0.5.

f in Appendix D, the price is actually given by:

$$(x_t, y_t) = \min\Big(\max\big(f(\bar{x}^e_{t+1}, \bar{y}^e_{t+1}), 0\big), 100\Big)$$

This means that for the $1.1real$ treatment we create two extra equilibria: $(0, 0)$ and $(100, 100)$. These could be seen as proxies for divergence to $-\infty$ and ∞ respectively.

2.3 Implementation

The experiment was run in the CREED laboratory at the University of Amsterdam in April 2018 and April 2019. For the treatments $0.7real$, $0.95real$ and $1.1real$ we had six groups with six subjects each, and for the treatments $0.7complex$, $0.95complex$ and $0.95robust$ we had seven groups with six subjects each. This brought the total number of subjects to 234. No subject participated in more than one session. A session lasted about 2 hours in total. Earnings, including the €10 participation fee, averaged €33.24.

The experiment was programmed in oTree (Chen et al., 2016). Before the experiment began, participants read instructions on screen and answered a series of control questions to ensure they understood what they had read.

3 Aggregate dynamics

Figure 3 shows the market prices of x and y side by side, for all groups, sorted by treatment. A first, perhaps surprising result is that in a number of treatment, despite limited information, subjects are able to learn to coordinate on the steady state equilibrium. We see clear differences between the treatments. $0.7real$ gives for much more stable markets than $0.95real$ and $1.1real$. In the case of $0.95real$ we observe rather persistent oscillations in the various groups, albeit of different amplitudes. In treatment $1.1real$ the groups showed quite different dynamics. Some groups coordinated on the $(100, 100)$ equilibrium, others to $(0, 0)$, which can be seen as a proxy for divergence to $+\infty$ and $-\infty$ respectively. Others had oscillatory dynamics.[2] Notably, $0.7complex$ and especially $0.95complex$ seem to have more stable markets than their real counterparts.

Another way to look at the observed prices is by drawing phase plots. Figure 4 shows a typical example of a phase plot for each treatment. We see what we

[2]Looking at the different results for the unstable eigenvalue 1.1 these seem to be driven by individual participants making idiosyncratic predictions which threw off the rest of the group. In the group where things seemed to have stabilised at 100, one participant suddenly decided to put a very low prediction (potentially out of boredom) resulting in the price dropping and all participants changing their predictions. In the group where we see something close to convergence there is actually 1 participants putting huge fluctuations in (switching between 0 and 60) and the other participants switching between high values to try and understand this behaviour. Yet another group includes a participant who simply started predicting one value (far from the equilibrium). It is likely that these 'rogue' participants were so common in this treatment due to its heavily unstable behaviour. Only one other group exhibited this type of erratic individual behaviour, see next footnote.

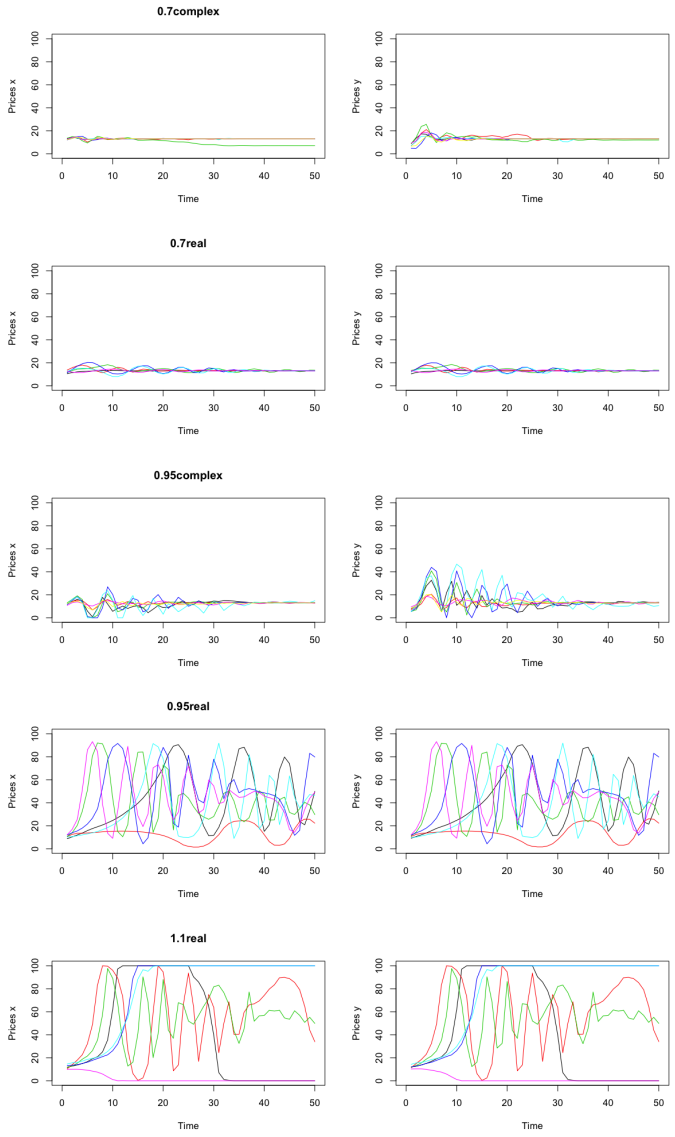

Figure 3: Market prices of x and y in each group, sorted by treatment

Table 1: Average RAD and RD across groups per treatment

	RAD	RD
0.7complex	0.036	0.002
0.7real	0.075	0.027
0.95complex	0.209	0.020
0.95real	2.012	1.917

would expect based on the eigenvalues and vectors: for the real treatments the fixed point is approached in a straight line following the $x = y$ direction (the stable/weakly unstable manifold), while in the complex treatment the fixed point is approached along a spiral. This result for the real treatment can be seen as a result of our choice for the manifolds, but it can also be that $x = y$ is simply a focal point. We will discuss this further in later sections, but it is important to note this for the further analysis where we see a much bigger difference between measures for x and for y for the complex treatments than for the real treatments. We also emphasise this result for the saddle point treatment $1.1real$, where the dynamics followed the unstable rather than the stable manifold, probably contributing to the lack of convergence to the interior equilibrium.

3.1 Measuring stability

In order to quantify the differences in stability we observe in Figure 3, we can use the bubble measures presented in Stöckl et al. (2010): the Relative Absolute Deviation (RAD) and Relative Deviation (RD) from the equilibrium $(13, 13)$.

$$RAD = \frac{1}{50}\sum_{t=1}^{50}\frac{|p_t - 13|}{13} \tag{1}$$

$$RD = \frac{1}{50}\sum_{t=1}^{50}\frac{p_t - 13}{13} \tag{2}$$

where p_t can be either x_t or y_t. We calculate the two measures for x and y separately and then take the average of those two.[3] Since there are three equilibria for the $1.1real$ treatment we do not apply these measures there. The averages per treatment can be found in Table 1.

[3] $0.7complex$, group 3 has been removed for further analysis since there was one 'rogue' subject in this group. The participant switched to entering only 0 and 100 after a certain number of rounds, even when the price stabilised near the equilibrium. As a result, this participant made much less money than all the others in her group. The participant clearly did not try to maximise her earnings by attempting to predict the prices as well as possible.

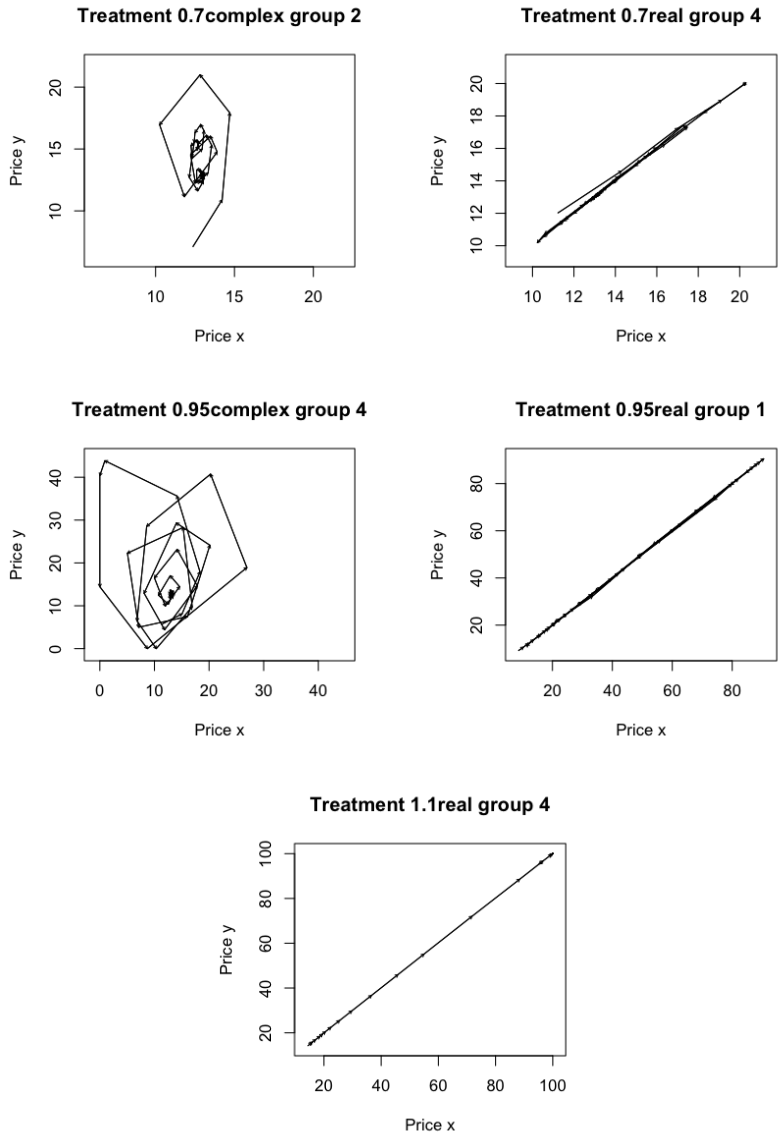

Figure 4: Phase plot showing the observed prices for one exemplary group per treatment. Note that the scales may differ per treatment.

Table 2: P values Wilcoxon test comparing average Relative Absolute Deviation (RAD) in groups per treatment

	0.7complex	0.7real	0.95complex
0.7real	0.132		
0.95complex	0.001	0.101	
0.95real	0.002	0.002	0.002

Table 3: P values Wilcoxon test comparing average Relative Deviation (RD) in groups per treatment

	0.7complex	0.7real	0.95complex
0.7real	0.310		
0.95complex	0.731	0.731	
0.95real	0.002	0.015	0.008

We see that all treatments have a positive RD, which means that the price is on average above the equilibrium. For $0.7real$, $0.7complex$ and $0.95complex$ the RD is very close to 0 indicating their stability on the fixed point. The RAD is also close to 0, but $0.95complex$ is larger here. The RAD and RD for $0.95real$ are much larger than for the other treatments, indicating its instability.

We use the Wilcoxon rank sum test to see whether these differences and similarities are statistically significant (see Tables 2, 3).

We see that the difference between $0.95real$ and all other treatments is indeed statistically significant at 2% for the RAD. Notably, this includes the difference between $0.95real$ and $0.95complex$ (0.002). $0.95complex$ on the other hand is not significantly different from the $0.7real$ treatment by either measurement (0.101, 0.731). We also see that $0.7real$ and $0.95real$ are different at a 5% level by both measures showing a big shift in stability within te unit-circle. $0.7complex$ and $0.95complex$ are different by the RAD (0.001), showing some shift in stability within the unit-circle for the complex values also.

3.2 Coordination

In order to get an idea of the level of coordination among participants we can look at the Dispersion Error. This is a measure of the heterogeneity in predictions among the participants in a group (Muth, 1961). Allowing for a brief period of learning, it is given by:

Table 4: Average dispersion error across groups per treatment

0.7complex	0.7real	0.95complex	0.95real	1.1real
0.577	0.222	12.616	30.080	12.948

Table 5: P values Wilcoxon test comparing average dispersion errors in groups per treatment

	0.7complex	0.7real	0.95complex	0.95real
0.7real	0.589			
0.95complex	0.051	0.181		
0.95real	0.041	0.041	0.234	
1.1real	0.041	0.065	0.345	0.818

$$\frac{1}{40}\frac{1}{6}\sum_{t=11}^{50}\sum_{i=1}^{6}(p_{it}^{e}-\bar{p}_{t}^{e})^{2} \tag{3}$$

If this measure is high, there will be high heterogeneity in the predictions, and little coordination. Of course, we will see a very low dispersion error as soon as a price has converged. Table 4 shows the average dispersion error in each treatment. We see that indeed, $0.7real$ and $0.7complex$ have the lowest dispersion error, which can probably be explained by the stability in these treatments. $0.95complex$ on the other hand has a much higher dispersion error. This is somewhat surprising, but links to the findings in Heemeijer et al. (2009). They found that their more stable treatments could have high dispersion errors, and that in fact this lack of coordination could lead to stabilisation since there is no coordination on trend-following heuristics and hence a bubble mentality.

We have again performed Wilcoxon rank sum tests in order to assess whether the differences we see in the dispersion errors are statistically significant. The results are reported in Table 5. We see that the dispersion error of $0.95complex$ is not significantly different from that of $0.95real$ (0.234) and $1.1real$ (0.345), indicating again how little coordination we saw in this relatively stable treatment.

4 Models of expectation formation

We begin by trying to find a model that can replicate the predictions made by the participants in the experiment in each time step. This means that we give the model the whole set of information available to the participant at time t, and we see how close the prediction the model makes given this information is to the prediction the participant made at time t.

4.1 Adaptive rules

A first model we can try to fit is some form of an adaptive model. By this we mean a model of the form:

$$\begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} = \begin{pmatrix} x^e_t \\ y^e_t \end{pmatrix} + \mathbf{A}^T \begin{pmatrix} x_{t-1} - x^e_t \\ y_{t-1} - y^e_t \end{pmatrix} \tag{4}$$

where $\mathbf{A}$ is a $2x2$ matrix. We specify three cases that we will estimate:

- The 'one parameter model', where $A_{11} = A_{22}$ and $A_{12} = A_{21} = 0$
- The 'two parameter model', where $A_{12} = A_{21} = 0$
- The 'four parameter model', where none of the entries are restricted

We estimate this through OLS.

4.1.1 One parameter adaptive rule

We begin by fitting a one-parameter model to our data. We estimate the parameter separately for each treatment. The first columns of Table 6 shows the results. We see that the estimated parameters are smaller for the complex treatments than for the real ones. Note that when $A = 1$, the participants always predict simply the most recently observed price. We call this the naive model. When $A = 0$ the participants consistently enter their initial prediction. A lower A would thus correspond roughly to more constant predictions, while a higher A would correspond to more weight on recent observations. This is a first explanation for the lower estimated parameter in the complex treatments, since these were quite stable and thus had quite constant predictions. It does not, however, explain the high estimate for $0.7real$. This can be explained by the fact that $0.7real$ has lower frequency oscillations than the complex treatments do. As a result, recent observations are more informative and receive a higher weight.

]

Table 6: Estimation results for adaptive rules. Stars indicate significance at 0.01 level (***), 0.05 level (**) or 0.1 level (*).

Treatment	One Parameter	AIC (1)	Two Parameter	AIC (2)	Four Parameter	AIC (4)
0.7complex	0.58(***)	12132.168	0.534(***)	12124.667	0.497(***)	11586.309
			0.601(***)		0.115(***)	
					0.655(***)	
					-0.378(***)	
0.7real	0.731(***)	11917.063	0.725(***)	11918.847	0.712(***)	11910.162
			0.736(***)		0.014	
					0.615(***)	
					0.134(***)	
0.95complex	0.645(***)	25154.28	0.651(***)	25156.102	0.628(***)	24442.575
			0.643(***)		0.155(***)	
					0.664(***)	
					-0.326(***)	
0.95real	0.939(***)	27516.181	0.939(***)	27518.181	1.59(***)	27495.478
			0.939(***)		-0.65(***)	
					0.34(**)	
					0.602(***)	
1.1real	0.831(***)	27009.56	0.837(***)	27011.231	1.204(***)	26979.885
			0.825(***)		-0.376(***)	
					0.955(***)	
					-0.133(**)	

4.1.2 Two parameter adaptive rule

Columns 3 and 4 of Table 6 show the results for the two parameter estimation. The first thing to note is that the estimated parameters for x and y in the real treatments are almost equal, and that these are also almost equal to the estimated parameter in the one-parameter case. This is simply explained by the fact that the observed x and y were almost always equal in these treatments. For the complex treatments we observe slightly bigger changes in the parameters, but no clear trend emerges. The AIC does not change much for any of the treatments, indicating that the added parameter does not increase the explanatory power of the model much. We seem to be able to conclude that adding this second parameter does not offer significant gains. This indicates that participants may not have used different methods for predicting x and y.

4.1.3 Four parameter adaptive rule

Column 5 and 6 of Table 6 finally show the results of an estimation with four parameters. For the real treatments we find some parameters with lower significance, and the AIC does not improve much. The added parameters again do not seem to add much explanatory power, which is unsurprising since x and y are generally equal for the real treatments. For the complex treatments we find that all parameters are significant at a 1% level, and that the AIC has dropped quite a lot. This model is thus an improvement to the one- and two-parameter models for the complex treatments. This indicates that for the complex treatments participants may have used information on x when predicting y and vice versa.

We will return to a comparison of the three types of adaptive rules in Subsection 4.3.

4.2 Heuristic Switching Model

A much more involved model of expectation formation, in which agents actually learn how to form their expectations as they go along, is the Heuristic Switching Model. This was introduced in Brock and Hommes (1997) and found to successfully explain dynamics observed in one-dimensional LtFEs in Anufriev and Hommes (2012).

The idea is that there are four heuristics that participants are choosing between to make their predictions in each time step.[4] We will use the heuristics used in Anufriev and Hommes (2012), namely the Adaptive heuristic (ADA), Weak Trend-following Rule (WTR), Strong Trend-following Rule (STR) and the Anchoring and Adjustment rule (LAA):

$$p^e_{1,t+1} = 0.65p_{t-1} + 0.35p^e_{1,t} \qquad \text{(ADA)}$$

[4]These 4 rules represent four different types of behaviour frequently observed in 1-D LtFEs: monotonic convergence, oscillatory convergence, explosive behaviour and persistent oscillations.

$$p^e_{2,t+1} = p_{t-1} + 0.4(p_{t-1} - p_{t-2}) \quad \text{(WTR)}$$

$$p^e_{3,t+1} = p_{t-1} + 1.3(p_{t-1} - p_{t-2}) \quad \text{(STR)}$$

$$p^e_{4,t+1} = 0.5(p^{av}_{t-1} + p_{t-1}) + (p_{t-1} - p_{t-2}) \quad \text{(LAA)}$$

First, a measure of the success realised by each strategy in the last time step is calculated. We use the same measure that is used to determine the pay-off of a prediction. We then include some memory of previous success by the heuristic, governed by memory parameter $\eta = 0.7$. This brings the performance measure of heuristic i to:

$$U_{t-1,i} = \frac{100}{1 + |p_{t-1} - p^e_{t-1,i}|} + \eta U_{t-2,i} \tag{5}$$

We then use a discrete choice model with asynchronous updating to determine what heuristics the participants will use in the next time step. Intensity of choice parameter β determines the extent to which participants will switch to the most successful rule, and parameter δ determines how many participants are switching strategies in an individual time step. The fraction of participants using strategy i to make their next prediction is then equal to:

$$n_{t,i} = \delta n_{t-1,i} + (1 - \delta)\frac{e^{\beta U_{t-1,i}}}{\sum_j e^{\beta U_{t-1,j}}} \tag{6}$$

We generally use the same form and parameters as in Anufriev and Hommes (2012), with one exception. For the Strong Trend Following Rule, which looks like:

$$p^e_{t+1} = p_{t-1} + \gamma(p_{t-1} - p_{t-2}) \tag{7}$$

we choose γ to be 1.8 rather than the usual 1.3. We found this to be a better fit both for the one step ahead forecasting we are considering here and for the fifty steps ahead learning that will be considered in Section 5.The original coefficient 1.3 was found in Hommes et al. (2005) in asset pricing experiments that had slightly different characteristics that would temper bubbles, such as a far from equilibrium stabilizing robot trader. This helps justify our choice for a higher parameter. We make the same choices as Anufriev and Hommes (2012) in initialising the model and start with 25% weight on each rule.

Now in order to generalise the HSM to two dimensions, the simplest way to proceed is to say that the agents only use information about x to predict x, and information about y to predict y. This is the approach used in Assenza et al. (2021). Participants can use one heuristic for predicting x and a different one for predicting y at any one time, but the parameters that determine the choice of heuristic for each variable will be the same. We will refer to this as the 'decoupled HSM'.

There are of course other ways to generalise the HSM to two dimensions. In particular, we could consider that participants use information about y when predicting x and vice versa, as in the four parameter adaptive model. We will

be developing such a model in Subsection 5.3 and we will refer to it as the Mixed HSM. For the mixed HSM we use different initial conditions, making direct comparisons here slightly muddied. We will return to the reasons for this in Section 5.

4.3 Comparing 1-step ahead forecasting models

Now that we have laid out our various adaptive and heuristic switching models, we will compare a number of learning models based on their performance as measured by the mean squared error in Table 7. In addition to the models named before, we will be considering the 'rational' model, where participants constantly predict the equilibrium $(13, 13)$.[5] We will also consider the naive model, where participants always predict the most recently observed price. When separate estimations were done for x and y, we here report the average of the two MSEs.

A first observation from Table 7 is that the rational model performs the worst of all considered models for every treatment, including the stable ones, with the exception of $0.7real$ HSM mixed. This is not too surprising as it has by far the least flexibility.

The naive model however, which is still very simple, outperforms some of the more tailored models for some of the treatments. In particular, the naive model often outperforms one or the other form of the HSM. It also performs only slightly worse than the various adaptive models for the real treatments.

For the adaptive models we see that the change from one to two parameters barely makes a difference for any of the treatments. This means that simply adding the flexibility of two different learning rules for x and for y does not help explain the observed predictions. But, the move from two to four parameters does make a bigger difference for the complex treatments. Thus, adding the flexibility of using both x and y to predict x and vice versa does help explain the observations in the complex treatments. For the real treatments this barely makes a difference since x and y are mostly equal.

The decoupled heuristic switching model performs relatively poorly for the more unstable treatments, and relatively well for the most stable treatments. This changes when we consider the mixed HSM that will be laid out in Section 5. We discuss the reasons for this in this later section as well.

[5] Note that in the case of the saddle point there are two other equilibria, but we here only consider the interior steady state.

Table 7: Mean Squared Errors for various models of expectation formation

	Rational	Naive	Adaptive, one par	Adaptive, two par	Adaptive, four par	HSM decoupled	HSM mixed
0.7complex	4.126	2.529	1.701	1.696	1.459	0.107	3.021
0.7real	6.972	1.853	1.602	1.602	1.596	0.237	27.474
0.95complex	52.643	32.845	23.343	23.342	19.676	9.511	13.200
0.95real	1,435.358	123.421	122.046	122.046	121.144	125.635	99.926
1.1real	2,998.534	113.765	106.024	106.015	104.979	103.194	103.323

4.4 Individual estimations

Finally, we have also estimated learning models for individual participants rather than for a treatment pool as a whole.

We estimated models of the form:

$$p_t^e = c + \sum_{k=1}^{3} \alpha_k x_{t-k} + \sum_{l=1}^{3} \beta_l y_{t-l} + \sum_{i=0}^{2} a_i x_{t-i}^e + \sum_{j=0}^{2} b_j y_{t-j}^e \tag{8}$$

which is a two dimensional version of the model estimated by Heemeijer et al. (2009). We used a model selection method called LASSO (Tibshirani, 1996) to determine what parameters are explanatory for individuals.

We found some patterns in the parameters that were deemed insignificant for the various participants and treatments. We look at the estimated models for the prediction of x. Labelling a parameter as unused when less than 6 people in a treatment used it in their predictions, we find that the third lag of the price of y is not used in any of the treatments except $0.95complex$ (in predicting x). The third lag of the price of x is not used in any of the treatments except $0.7complex$. The second lag on the prediction of both x and y are only used in the complex treatments, and the second lag of the price of y is only used in the complex treatments and $0.7real$. We conclude that more information is used to make predictions for the complex treatments than for the real treatments, both more information from the past and more 'mixed' information about both variables. This matches our lower adaptive parameters, which showed that in the complex treatments there was less weight on more recent observations. It also matches the success of the four parameter model for these treatments.

5 Long run behaviour

If our 1-step forecasting models in Section 4 performed perfectly, we would be able to use them to predict the aggregate dynamics of a market ex ante, without feeding the model information about observations in the laboratory past the initial stage. However, since our 1-step ahead forecasting models still make mistakes, these mistakes can add up to reproduce qualitatively different aggregate dynamics in the medium to long run from those observed. Since one of the goals of this type of research is to be able to predict the stability of equilibria that have not been tested in the lab, we are interested in predicting aggregate dynamics given only the initial conditions of the experiment. In this section we therefore investigate how various models perform qualitatively when we let them make predictions for 50 time steps, given only the initial conditions from the lab.

5.1 Successes and failures existing models

We will now consider to what extent the learning models we considered in Section 4 can qualitatively predict the aggregate market dynamics we observed.

5.1.1 Rational model

First of all, the Rational Expectations model would predict markets which are stable at $(13, 13)$. This may be a reasonably good description of the observed dynamics (albeit not the path of convergence) for $0.7real$ and $0.7complex$, and perhaps even the long run behavior (after period 30) of $0.95complex$, but it can not explain what we observe for $0.95real$ and $1.1real$, nor why we see differences between treatments to begin with.

5.1.2 Naive model

When we consider the naive model, and 'plug' this into the laws of motion in Appendix D, what we end up with is exactly the expectations feedback model for which we chose our eigenvalues and eigenvectors. We will thus see a change from stable to unstable dynamics exactly when the absolute value of the leading eigenvalue exceeds 1, which contradicts the observed instability of the $0.95real$ treatment. The naive model also predicts that convergence takes exactly as long for the complex treatments as their real counterpart. This too contradicts our findings, most clearly when comparing $0.95real$ to $0.95complex$.

5.1.3 Adaptive model

When simulating the adaptive model for 50 time steps, we can only replicate the oscillating dynamics that we found in several treatments for the complex treatments. We do see that the adaptive model fully stabilises on an equilibrium within 50 steps for exactly the same treatments that we found full stabilisation in the lab. However, in the treatments where we found instability ($0.95real$, $1.1real$) the adaptive model predicts one peak and then a slow but monotonic convergence towards the equilibrium, which does not match our qualitative findings.

5.1.4 Decoupled HSM

As argued by Anufriev and Hommes (2012) for one dimension, we find that the HSM can reproduce the change in stability when moving from $0.7real$ to $0.95real$. This happens because the Strong Trend-following Rule performs very well in the case of $0.95real$, amplifying an unstable market despite the eigenvalues being in the unit circle. For $0.7real$ there is less coordination at the start, in particular no coordination on the STR, reducing bubbles. The HSM needs to be initialised on the fractions for each heuristic and the dynamics can change quite a lot when these are altered. It would be hard to defend to choose different initial fractions for different treatments. This makes it impossible to reproduce the observed dynamics in all treatments using the decoupled HSM. Generally, the HSM underestimates how stable the $0.95complex$ treatment was, and overestimates how stable the $0.95real$ treatment was. We can reproduce large oscillations in the $0.95real$ treatment by making the initial fraction of strong trend

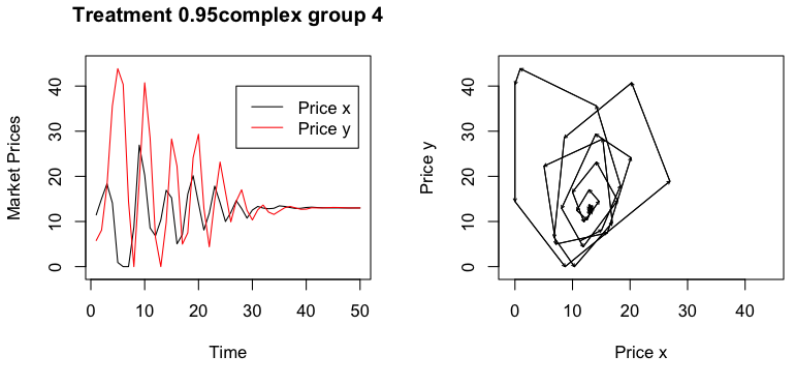

Figure 5: Example observed dynamics complex treatment

followers very large, but with such fractions the 0.95*complex* shows large persistent oscillations as well. To approximate the stability in the 0.95*complex* case we need to increase the fraction of the adaptive heuristic so much that the 0.95*real* treatment becomes very stable. Summarizing, the decoupled HSM can not simultaneously, i.e. with the same initial fractions, explain the unstable behavior in the 0.95*real* treatment and the stable behaviour in the 0.95*complex* treatment.

5.2 Evidence for a mixed model

The decoupled HSM came close to predicting all qualitative dynamics, but it underestimated the stability of 0.95*complex*. We may be able to reproduce this stability in the HSM by including some mixing, so by allowing y to be used when predicting x and vice versa. A mixed HSM would be conistent with the estimated individual rules in Section 4.4.

5.2.1 Why mixing may stabilise

The reason why mixing may aid in reproducing the stability is as follows. The spiralling dynamic we have in the complex treatments means that x and y are only moving in the same direction roughly half of the time, see Figure 5. When the two are moving in opposite directions, taking both trend-movements into account may temper bubble extrapolation. When subjects only focuses on one of the trends, the participant would simply extrapolate that trend. In the case of the two diverging trends, a form of mixed trend-following would take into account both the increasing trend of one and the decreasing trend of the other variable, and the overall result would be weaker trend-extrapolation.

For the real treatments, since here we found x and y to be equal almost all the time, mixing will not make a difference as long as the sum of the weights

on x and y does not change.

5.2.2 Evidence mixing

We found evidence of the participants mixing the history of x and y for complex treatments in Section 4. First of all, the four parameter adaptive model performed better than the two parameter model for the complex treatments. The difference between these two was the inclusion of mixing. We also find that in the four parameter adaptive model for the complex treatments, all variables have small p-values, indicating that allowing for four parameters contributes to the predictive power of the model.

In the individual estimations we find further evidence of mixing. For the complex treatments, we see that more information on both variables, including the other lagged variables, is used to make predictions than for the real treatments.

5.3 Proposed mixed HSM

One may argue that there are infinitely many ways in which we may adapt the HSM to include mixing. We choose the simplest adaptation of the decoupled HSM that would replicate the qualitative dynamics.

It turns out that keeping anchors as they are, and splitting adjustments equally across the two variables fits the data well. This means that when predicting x, one looks at the current price (or other anchor) of just x, but looks at the changes in both x and y to forecast their adjustments. The four rules are thus as follows:

$$p^e_{t+1} = 0.65p_{t-1} + 0.175x^e_t + 0.175y^e_t \quad \text{(ADA-m)}$$

$$p^e_{t+1} = p_{t-1} + 0.2(x_{t-1} - x_{t-2}) + 0.2(y_{t-1} - y_{t-2}) \quad \text{(WTR-m)}$$

$$p^e_{t+1} = p_{t-1} + 0.9(x_{t-1} - x_{t-2}) + 0.9(y_{t-1} - y_{t-2}) \quad \text{(STR-m)}$$

$$p^e_{t+1} = 0.5(p^{av}_{t-1} + p_{t-1}) + 0.5(x_{t-1} - x_{t-2}) + 0.5(y_{t-1} - y_{t-2}) \quad \text{(LAA-m)}$$

Now if we combine this with the right initial conditions, we can replicate the observed dynamics. For the initial conditions, we are mainly constrained by the large instability in the $0.95real$ case. The bigger we choose the initial fraction using STR, the better we replicate this instability. But doing this does not give wrong qualitative dynamics for the other treatments, as we saw occur in the decoupled HSM. We do however want to base our initial fractions on our observations from the lab. Since we found that about 70% of inititial predictions either increased or decreased from the first to the second time step, despite the fact that no new information came through, we conclude there was some ex ante trend extrapolation present. We thus place an initial weight of 0.7 on the STR and split the further weight equally across the three other heuristics.

Figure 6 shows for one exemplary group per treatment how the realised and simulated prices compare. These are representative of the fit we see in the

Table 8: Average fraction using each rule for fifty step ahead mixed HSM

	0.7complex	0.7real	0.95complex	0.95real	1.1real
ADA	0.683	0.090	0.610	0.033	0.137
WTR	0.096	0.147	0.153	0.023	0.236
STR	0.196	0.238	0.199	0.668	0.590
LAA	0.025	0.525	0.038	0.276	0.037

other groups, except for $1.1real$. For $1.1real$ the mixed HSM always predicts that the market will shoot to the $(100, 100)$ equilibrium while in some of our treatments we saw different behaviour, although there is agreement on the lack of convergence to the internal equilibrium.

5.4 Prevailing rules

Table 8 shows for the mixed HSM what fraction of agents uses what rule on average over time, over the two prices and over each market, per treatment. We see that the ADA rule prevails in the complex treatments and is barely used in the real ones. This is in line with the success of the four parameter adaptive model we saw in the one step ahead estimation. Since the ADA rule is stabilising this helps explain the stability observed in the complex treatments. In the real treatments the STR is common, and here we also see the LAA doing quite well. Since the LAA includes trend following as well as a strong anchoring, it fits both the unstable dynamics of $0.95real$ and the stable ones of $0.7real$.

5.5 Predicting out of sample

Now that we have found a model for expectation formation that can qualitatively replicate almost all of the behaviour we observed in the lab, we can use this model to predict stability of settings we have not studied in the lab. We have simulated the model for sets of eigenvalues across the complex unit square, using the same 'rules' for the positioning of eigenvectors, the equilibrium and the shape of the equations as we did for our treatments.[6] We then calculate the RAD for each set of eigenvalues to represent the stability. We sampled initial conditions randomly from the observations in the lab, and left the first 25 steps out of the calculation of the RAD to minimise the influence of the initial conditions. Figure 7 shows the results.

We find that a small polar angle is in fact destabilising (for large absolute value), and only larger angles stabilise. We find the asymmetry along the real

[6]Rather than forcing the upper right entry to be negative for the complex eigenvalue matrices, we fixed the constant in the equations to be equal to the constant we chose for our initial complex treatments, for computational reasons.

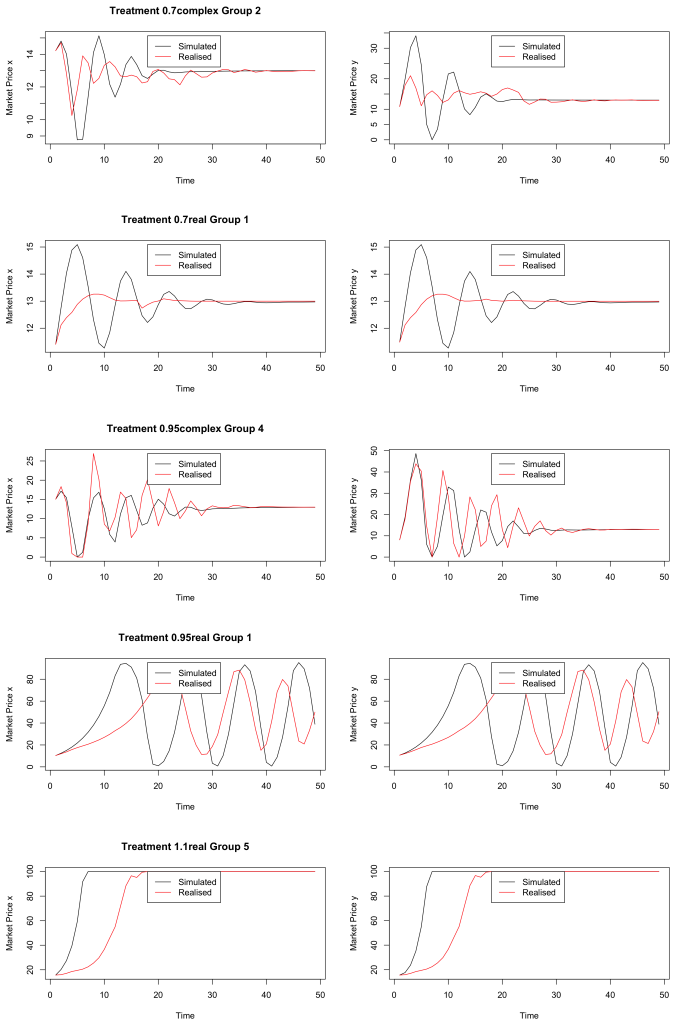

Figure 6: Realised prices and prices as simulated using the mixed HSM, 50 steps ahead, for both x and y for one exemplary group per treatment. Initial weight of 0.7 on STR, and 0.1667 on every other heuristic.

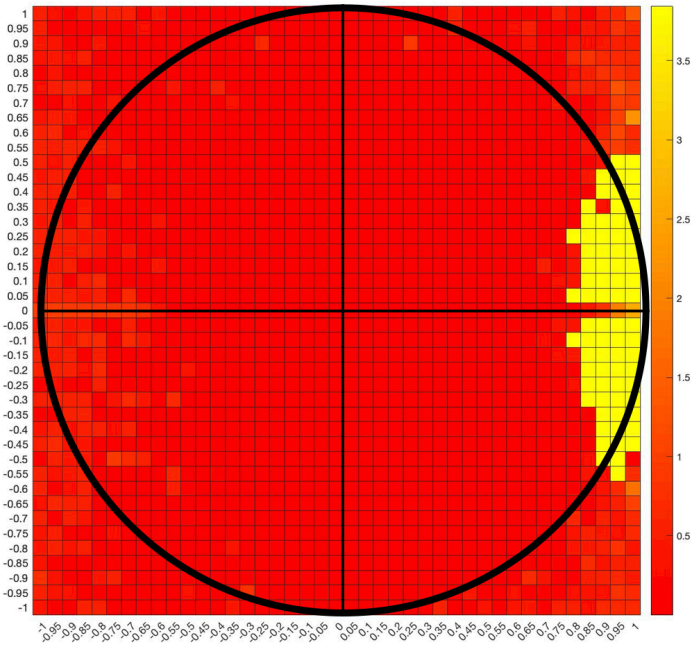

Figure 7: Heatmap showing the RAD simulated by the mixed HSM for different values on the complex unit square. The x-axis shows the real part, the y-axis the complex part of the (leading) eigenvalues. The unit circle is shown for clarity.

axis that we are familiar with from the one dimensional setting: negative real part generally leads to more stable markets than positive real part.

Figure 8 shows some eigenvalues from the literature displayed over the aforementioned heatmap. The dots in black are those from our experiment, the one in white is from Anufriev et al. (2019) and those in white with a black dot are from Hommes et al. (2019a). We find that our model correctly predicts the stability in these cases.

6 Robustness check: partial derivatives

In the treatments discussed so far, we found a change from unstable to stable dynamics when we went from real eigenvalues to complex eigenvalues with the same absolute value 0.95. It is possible, however, that this result was driven by the inclusion of a negative partial derivative. Equation (9) shows the model used for the treatment with real eigenvalues $(0.95, 0.5)$ and (10) the one for $0.95e^{\pm i\pi/4}$.

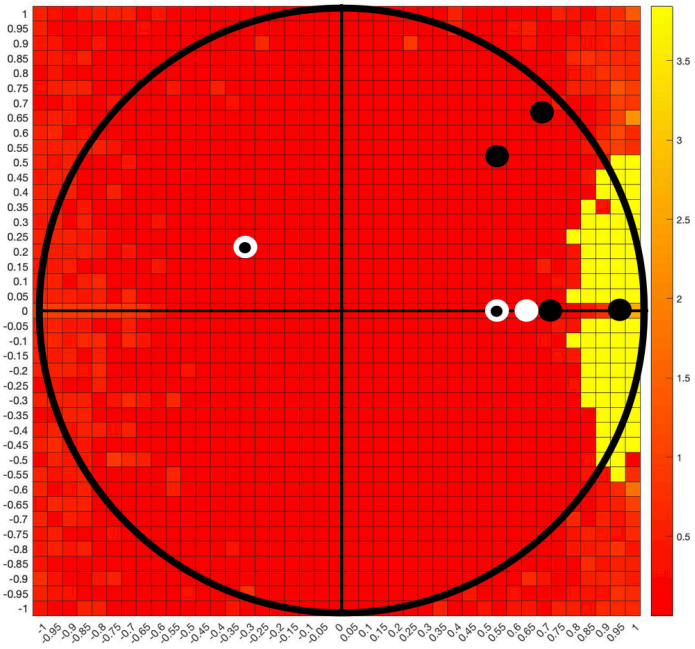

Figure 8: Heatmap showing the RAD simulated by the mixed HSM for different values on the complex unit square. The x-axis shows the real part, the y-axis the complex part of the (leading) eigenvalues. The unit circle is shown for clarity. The dots in black are those from our experiment, the one in white is from Anufriev et al. (2019) and those in white with a black dot are from Hommes et al. (2019a).

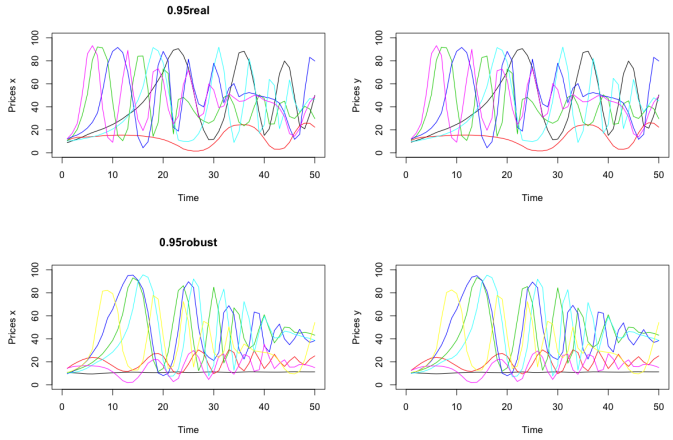

Figure 9: Market prices of x and y in each group, for the two treatments $0.95real$ (top panels) and $0.95robust$ (bottom panels).

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.725 & 0.225 \\ 0.225 & 0.725 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 0.65 \\ 0.65 \end{pmatrix} \tag{9}$$

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.651 & -0.420 \\ 1.076 & 0.693 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 10 \\ -10 \end{pmatrix} \tag{10}$$

Assenza et al. (2021) run 2D LtFEs in the New Keyensian model and argue that adding negative feedback through a more aggressive Taylor rule stabilises the NK system. One might therefore hypothesise that the signs of the partial derivatives rather than the eigenvalues should be used as predictors of stability. To test this hypothesis, we conducted an additional treatment where we had again the eigenvalues $(0.95, 0.5)$, but this time chose a model which included negative feedback, through one negative partial derivative, as in (10). If this would stabilise the dynamics it would suggest that eigenvalues are not sufficient as predictors, and we require the full picture of the various partial derivatives.

The treatment $0.95robust$, with real eigenvalues 0.5 and 0.95, was thus conducted with expectations feedback model that includes negative feedback :

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 1.15 & -0.2 \\ 0.65 & 0.3 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 0.65 \\ 0.65 \end{pmatrix} \tag{11}$$

6.1 Results

Let us first focus on the observed prices, shown in Figure 3. We could roughly classify the observed markets in the $0.95real$ treatment as twice persistent large oscillations, once persistent small oscillations, and three times large oscillations which seem to lead to convergence. In the $0.95robust$ we have four times large

Table 9: Bubble measures and average dispersion errors for 0.9real and 0.95robust

	0.95real	0.95robust
RAD	2.012	1.452
RD	1.917	1.332
Dispersion Error	30.080	15.011

oscillations which seem to lead to convergence, twice small oscillations which seem to lead to convergence, and one very stable market. Based on this qualitative evidence one could say that $0.95robust$ appears somewhat more stable, but it is not entirely clear. In Table 9 we see that the RAD and RD for $0.95robust$ are indeed smaller than for $0.95real$. We can perform a Wilcoxon rank-sum test to determine whether this difference is statistically significant. For both RAD and RD the p-value of this test is 0.45, meaning that the difference is not statistically significant. The difference between the dispersion errors we observe is also not significant, with a p-value of 0.366. We thus conclude that the differences between the $0.95real$ and $0.95robust$ treatments are not significant. We find evidence that eigenvalues rather than individual partial derivatives are good predictors of stability.

7 Conclusion

We have run Learning to Forecast Experiments in two dimensions to investigate the use of eigenvalues as predictors of stability. Despite participants only having qualitative information about the underlying feedback system, we found that in multiple treatments the markets converged to the equilibrium value. In particular, $0.7real$, $0.7complex$ and $0.95complex$ converged to the fixed point. This is especially notable since $0.95real$, which had the same absolute value as $0.95complex$, lead to very unstable markets. We find that complex eigenvalues with an angle of $\frac{\pi}{4}$ lead to more stable markets than real eigenvalues with the same absolute value.

The known models for 2D expectation formation we examined can not replicate this difference in stability between real and complex treatments. We present a version of the HSM where agents mix information on the two variables they are predicting when forming the adjustments to their anchors. This can explain stability for the complex treatments, where variables often move in opposite directions meaning the two adjustments offset each other.

We use this new model to make predictions for the stability of systems with eigenvalues that we did not test ourselves. We predict that eigenvalues with a small polar angle and a high absolute value lead to highly unstable markets,

while higher polar angles are stabilising. Since we used a generalised, low context set-up we hope that other 2D LtFEs in any context can form evidence for or against our model and test our predictions.

When trying to replicate only the expectations participants made in any individual time step, rather than for the whole duration of the experiment, the HSM does not necessarily outperform simpler models. While it does for some treatments, a simple one parameter adaptive model does not perform much worse over all and is a surprisingly good fit.

Finally we performed a robustness check to see if it is indeed the eigenvalue rather than all individual partial derivatives that should be used as predictors of stability. We find that the eigenvalue is indeed sufficient since two treatments with equal eigenvalues and different partial derivatives are not significantly different in our aggregate stability or coordination measures.

References

Anufriev, M., J. Duffy, and V. Panchenko
2019. Planar beauty contests. *Available at SSRN 3404334*.

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

Arifovic, J. and L. Petersen
2017. Stabilizing expectations at the zero lower bound: Experimental evidence. *Journal of Economic Dynamics and Control*, 82:21–43.

Assenza, T., P. Heemeijer, C. Hommes, and D. Massaro
2021. Managing self-organization of expectations through monetary policy: A macro experiment. *Journal of Monetary Economics*, 117(C):170–186.

Bao, T. and C. Hommes
2019. When speculators meet suppliers: Positive versus negative feedback in experimental housing markets. *Journal of Economic Dynamics and Control*, 107:103730.

Brock, W. A. and C. H. Hommes
1997. A rational route to randomness. *Econometrica: Journal of the Econometric Society*, Pp. 1059–1095.

Chen, D. L., M. Schonger, and C. Wickens
2016. otreean open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Heemeijer, P., C. Hommes, J. Sonnemans, and J. Tuinstra
2009. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic dynamics and control*, 33(5):1052–1072.

Hommes, C., D. Massaro, and I. Salle
2019a. Monetary and fiscal policy design at the zero lower bound: Evidence from the lab. *Economic Inquiry*, 57(2):1120–1140.

Hommes, C., D. Massaro, and M. Weber
2019b. Monetary policy under behavioral expectations: Theory and experiment. *European Economic Review*, 118:193–212.

Hommes, C., J. Sonnemans, J. Tuinstra, and H. Van de Velden
2005. Coordination of expectations in asset pricing experiments. *The Review of Financial Studies*, 18(3):955–980.

Kahneman, D. and A. Tversky
1984. Choices, values, and frames. *American psychologist*, 39(4):341.

Kryvtsov, O. and L. Petersen
2013. Expectations and monetary policy: experimental evidence. Technical report, Bank of Canada Working Paper.

Lucas, R. E.
1986. Adaptive behavior and economic theory. *Journal of Business*, Pp. S401–S426.

Muth, J. F.
1961. Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, Pp. 315–335.

Pfajfar, D. and B. Žakelj
2018. Inflation expectations and monetary policy design: Evidence from the laboratory. *Macroeconomic Dynamics*, 22(4):1035–1075.

Sonnemans, J. and J. Tuinstra
2010. Positive expectations feedback experiments and number guessing games as models of financial markets. *Journal of Economic Psychology*, 31(6):964–984.

Stöckl, T., J. Huber, and M. Kirchler
2010. Bubble measures in experimental asset markets. *Experimental Economics*, 13(3):284–298.

Tibshirani, R.
1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.

A Instructions

Welcome to this experiment! The experiment is anonymous, the data from your choices will only be linked to your station ID, not to your name. You will be paid privately at the end, after all participants have finished the experiment. After the main part of the experiment and before the payment you will be asked to fill out a short questionnaire. On your desk you will find a calculator and scratch paper, which you can use during the experiment.

During the experiment you are not allowed to use your mobile phone. You are also not allowed to communicate with other participants. If you have a question at any time, please raise your hand and someone will come to your desk. All participants will be randomly divided into groups of six people. The group composition will not change during the experiment. [7]

Your task in this experiment is to predict the prices of two different goods. The experiment consists of 50 time periods. In every period you will predict what the two prices will later be. The predictions you and the others in your group make determine what the prices actually become. The closer your predictions are to the realised prices, the higher your earnings in the experiment will be.

A.1 About the prices you are predicting

The two goods of which you are predicting the prices are called x and y. You are always making predictions for the price **in the next period**. When you make the prediction for the next period, you do not yet know the price in the current period. For example, the predictions made in period 5 are for the price in period 6. When you make the prediction for period 6, you do not yet know the price for period 5. The predictions you and the others in your group make for the price of x affect the realised prices of both x and y, and the predictions for the price of y affect the realised prices of both y and x. If the predictions for the price of x go up, this **increases** the realised price of x and **increases** the realised price of y (if nothing else changes). If the predictions for the price of y go up, this **increases** the realised price of x and **increases** the realised price of y (if nothing else changes). These effects are summarised in the table below. It shows the relationship between the average prediction of the price and the realised price. A positive relationship means that a higher prediction leads to a higher realisation, a negative one that a higher prediction leads to a lower realisation.

Table 10: Relationship between the predictions of a price and the realisation of that price

	Realisation of x	Realisation of y
Average prediction of x	+	+
Average prediction of y	+	+

[7] These first two paragraphs are verbatim Hommes et al. (2019b)

A.2 Your predictions

Your task, in each period of the experiment, is to predict the prices of goods x and y. At no point in the experiment can the prices (or your prediction of the prices) go below 0. When the experiment starts, you have to predict the prices for the first two periods. These predictions have to be in the range of 0 to 20. When all participants have made their predictions for the first two periods, the actual prices of x and y for period 1 are announced. Then period 2 begins. In period 2 you make a prediction for the prices in period 3. When all participants have made their predictions for period 3, the prices for period 2 are announced. This process repeats for 50 periods.[8] Therefore, when at a certain period t you make a prediction of the prices in period $t + 1$, the following information is available to you:

- The actual prices of both x and y up to and including period $t - 1$;
- Your predictions up to and including period t;
- Your points up to and including period t (see below).

A.3 Your earnings

Your earnings depend on the accuracy of your predictions. You start the experiment with 5000 points, which would be worth 50 euros at the end of the experiment. In every period you may lose some of these points, depending on how far your predictions are from the realised prices. To determine your points in each period, first your prediction errors in x and in y are calculated separately by finding the difference between your prediction and the realisation. Then, the two prediction errors are added to find your total error. For example: if your prediction of x was 3 and your prediction of y was 5 and the true prices turn out to be 8 for x and 6 for y, your prediction error in x is $|8 - 3| = 5$ and your prediction error in y is $|6 - 5| = 1$. Your total error is therefore $5 + 1 = 6$. Every period is worth 100 points. If you make no error, you keep all of these 100 points. If you make an error, points are taken away. The table and graph in section B show how many points you will have left in a period based on the error you make. In the example above, where your error is 6, you would have 25 points left in that round so you would have lost 75 points.

At the end of the experiment, your final payment will consist of 1 euro for every 100 points you have left from the initial 5000. You will also receive a show-up fee.

[8] This subsection closely follows **?**.

B Prediction error to penalty point conversion

Total prediction error	Points lost	Points left for this period	Total prediction error	Points lost	Points left for this period
0	0.0	100.0	25	92.6	7.4
1	33.3	66.7	26	92.9	7.1
2	50.0	50.0	27	93.1	6.9
3	60.0	40.0	28	93.3	6.7
4	66.7	33.3	29	93.5	6.5
5	71.4	28.6	30	93.8	6.3
6	75.0	25.0	31	93.9	6.1
7	77.8	22.2	32	94.1	5.9
8	80.0	20.0	33	94.3	5.7
9	81.8	18.2	34	94.4	5.6
10	83.3	16.7	35	94.6	5.4
11	84.6	15.4	36	94.7	5.3
12	85.7	14.3	37	94.9	5.1
13	86.7	13.3	38	95.0	5.0
14	87.5	12.5	39	95.1	4.9
15	88.2	11.8	40	95.2	4.8
16	88.9	11.1	41	95.3	4.7
17	89.5	10.5	42	95.5	4.5
18	90.0	10.0	43	95.6	4.4
19	90.5	9.5	44	95.7	4.3
20	90.9	9.1	45	95.7	4.3
21	91.3	8.7	46	95.8	4.2
22	91.7	8.3	47	95.9	4.1
23	92.0	8.0	48	96.0	4.0
24	92.3	7.7	49	96.1	3.9

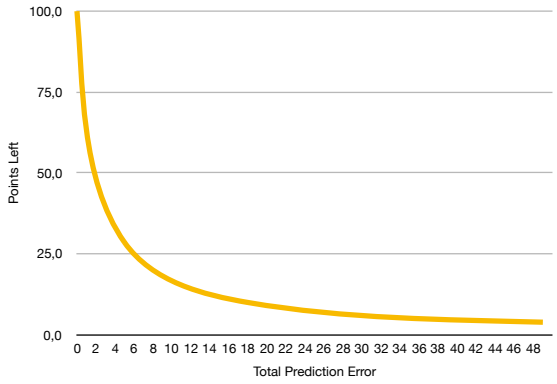

C Control Questions

- Suppose that in period 35 you make the following predictions: x will be 24 and y will be 7. Then the prices for x and y are revealed to be 20 and 13. Period 36 begins, you make another prediction, and the prices for x and y are now revealed to be 22 and 11. What is the prediction error for your predictions made in period 35? (Answer 6)
- Suppose everyone in your group, including yourself, predicts the price of y to stay constant and the price of x to go up. What effect will this have on the realised price? (Multiple choice, Answer: The price of x goes up, the price of y goes up.)
- Suppose you have 3500 points left when the prices for x and y are revealed to be 12 and 11. Your predictions for those prices were 8 and 15. How many points do you have left now? (Answer: 3420)
- Suppose that at the end of the experiment you have 2520 points left. How much is this worth in euros? (Answer: 21.00 euros)

D Maps for the various treatments

D.1 Treatment $0.7real$: eigenvalues (0.5, 0.7)

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.6 & 0.1 \\ 0.1 & 0.6 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 3.9 \\ 3.9 \end{pmatrix}$$

D.2 Treatment $0.95real$: eigenvalues (0.5, 0.95)

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.725 & 0.225 \\ 0.225 & 0.725 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 0.65 \\ 0.65 \end{pmatrix}$$

D.3 Treatment $1.1real$: eigenvalues (0.5, 1.1)

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.8 & 0.3 \\ 0.3 & 0.8 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} -1.3 \\ -1.3 \end{pmatrix}$$

D.4 Treatment $0.7complex$: eigenvalues $0.7e^{\pm i\pi/4}$

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.4354 & -0.2046 \\ 1.2147 & 0.5545 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 10 \\ -10 \end{pmatrix}$$

D.5 Treatment $0.95complex$: $0.95e^{\pm i\pi/4}$

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 0.6505 & -0.4197 \\ 1.0762 & 0.6930 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 10 \\ -10 \end{pmatrix}$$

D.6 Treatment $0.95robust$: eigenvalues (0.5, 0.95)

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} 1.15 & -0.2 \\ 0.65 & 0.3 \end{pmatrix} \begin{pmatrix} x^e_{t+1} \\ y^e_{t+1} \end{pmatrix} + \begin{pmatrix} 0.65 \\ 0.65 \end{pmatrix}$$