

CENTRE FOR APPLIED MACROECONOMIC ANALYSIS

The Australian National University



CAMA Working Paper Series

October, 2008

NEWS AND EXPECTATIONS IN FINANCIAL MARKETS: AN EXPERIMENTAL STUDY

Gordon Menzies

University of Technology, Sydney
The Australian National University

Daniel Zizzo

University of East Anglia
The Australian National University

CAMA Working Paper 34/2008
<http://cama.anu.edu.au>

News and Expectations in Financial Markets: An Experimental Study

Gordon Douglas Menzies¹
University of Technology, Sydney
Centre for Applied Macroeconomic Analysis, Australian National University

Daniel John Zizzo²
University of East Anglia
Centre for Applied Macroeconomic Analysis, Australian National University

October 2008

Abstract

We consider an experimental setting where traders in stock markets or exchange rate markets receive one stylized piece of information at a time about the value of an asset. We find that having limited knowledge about the prior distribution of true asset values does not hamper the decision making by traders and markets. There is empirical support for the common modeling assumption of simplifying agent heterogeneity into two types, a rational one and a less rational one. A correspondence exists between the average degree of belief conservatism found with individual buying and selling prices and that observed with market prices.

Keywords: uncertainty, expectations, information, stock market, exchange rate market.

JEL Classification Codes: C91, D83, D84, F31, G12.

¹ Tel: +61-2-95147728; fax +61-2-95147711. Email address: gordon.menzies@uts.edu.au

² Tel: +44-1603-593668; fax: +44-1603-456259. Email address: d.zizzo@uea.ac.uk

We are grateful to the British Academy for financial support, to participants to presentations in Lyon and Sydney for helpful feedback, and to James Watson, Nicole Huang, Kei Tsutsui and Maya Elliott for research assistance. The usual disclaimer applies. The experimental instructions can be found in the appendix.

1. Introduction

This paper considers a simple experimental setting where traders in either stock markets or exchange rate markets receive one stylized piece of information at a time about the value of an asset. In financial markets, agents have only very limited information about the distribution from which the value of any given asset is drawn, requiring them to form expectations.

A key question we try to answer is whether the essential *uncertainty* (as opposed to risk, following Knight's, 1921, distinction) intrinsic in the information set available to traders makes them update their beliefs differently, resulting in a different trajectory of market prices. We isolate the effect of uncertainty by having a treatment where the distribution of true asset values (shares or foreign currency) is very limited, and one in which it is fully known..

This is of interest because both individual choice and market experiments, (as in, Ellsberg, 1961, and Sarin and Weber, 1993) show that subjects dislike ambiguity, i.e. Knightian uncertainty. While our paper is not about ambiguity aversion *per se*, this literature shows intrinsic ambiguity might matter when agents process, and trade on, information.

Identifying whether uncertainty matters is important for another, and more general, reason: if traders do not have sufficient information about the prior distribution of the value of the asset (including the standard deviation), Bayesian signal extraction (hereafter Bayes) is not feasible in principle on the part of rational expectations agents, unless they append 'educated guesses' about the prior distribution to their inference procedure. If they refrain from these auxiliary assumptions, the rational procedure is maximum likelihood estimation (MLE) based on an average of the signals received. Therefore, how we model expectation formation is potentially a function of how agents process the intrinsic ambiguity in the market environment.

Together with rational expectations based on either Bayes or MLE, we allow for traders and markets to respond sluggishly to the receipt of information. Sluggish belief adjustment is considered, among others, by Mankiw and Reis (2003), Carroll (2003) and Sims (2003), and finds justification in experimental evidence showing that at least in some settings agents over-utilize prior information (e.g., Camerer, 1995).

We model the degree of belief conservatism by assuming that agents hold *inferential expectations* (IE; Menzies and Zizzo, 2005, 2007), where ‘inferential’ refers to classical Neyman-Pearson statistical inference. Agents employ a classical statistical test of size α in order to infer whether or not to retain the prior belief (the null hypothesis) or switch to the alternative belief, which is assumed to be the RE value. The latter assumption implies that RE corresponds to $\alpha = 1$ as a corollary³ and, in keeping with hypothesis testing, a lower α value necessarily implies greater belief conservatism. IE ranges between being completely unresponsive to evidence ($\alpha = 0$) and being completely untainted by belief conservatism ($\alpha = 1$). Intermediate values of α create a profile of beliefs that attains RE from time to time, interspersed with a ‘no change’ status quo in between. If one can observe beliefs, or infer them from actions, one can ‘back out’ an α value corresponding to the closest profile, on a least squares criterion. IE can be considered as a bounded-rational ‘fast and frugal heuristic’ (Gigerenzer et al., 1999) employed due to attentional, information-gathering and information-processing costs.⁴

Since IE is tied to RE, a modelling choice needs to be determined on whether, when agents switch to the alternative hypothesis, they move to the Bayesian signal extraction RE

³ A Neyman-Pearson hypothesis test minimizes the probability of a type II error, given a fixed probability of a type I error. If the latter is unity, one is compelled to reject H_0 always, thereby setting the former to zero. Since H_0 is always rejected, H_1 : RE is always embraced.

⁴ Menzies and Zizzo (2005, 2007) provide experimental evidence in favour of belief conservatism as modelled by $\alpha < 1$. Using a test size α as a way of modelling belief conservatism draws inspiration on the scientific practices of scientists, including economists, relying on classical statistical inference and test size to form and revise their beliefs. For example, one author received criticism from a referee in a top-tier journal because significance was accorded to a test with a p -value of 0.052.

(Bayesian RE in what follows) or to the RE based on the average of the signals (MLE RE). As such, Bayesian RE and IE belong to the same class of model-consistent expectations, the only difference being in the value of α , while MLE RE and IE belong to a different class of such expectations, again differentiated only by the value of α .

Our analysis also includes a standard adaptive expectations (ADE) specification, with a weight β assigned to the previous period's observation in forming the current expectation. ADE are included as they are often purported to be good data descriptors despite a lack of theoretical microfoundation (e.g., Mankiw, 2001) or even any model consistency requirement.

Our key and surprising finding is that Bayesian models of expectation formation, whether RE or IE, tend to outperform the corresponding MLE models, *even* in the no prior distribution knowledge treatments.

Furthermore, about 1/3 to 1/2 of the sessions can be best described by assuming belief conservatism ($\alpha < 1$) with the remaining ones exhibiting none ($\alpha \approx 1$). This bimodality of market α values appears to be driven by a bimodality of individual agent types, with some agents being rational in updating their beliefs and others being conservative. We find that the market exhibits continual and rational updating of information (market $\alpha \approx 1$) if the *average* degree of belief conservatism of the individual traders is below a given threshold, which we quantify. This bimodality corroborates the practice in financial modelling of assuming two types of agents, such as rational traders and noise traders (e.g., Barberis and Thaler, 2003), and is amplified rather than reduced at the level of market prices.

Our key results are robust to the frame used (i.e. whether traders face an exchange rate market or a stock market), and we cannot find any evidence of agents in our preferred model (Bayesian RE) becoming more rational as they repeat the task. ADE fails to overperform (and often underperforms) Bayesian models at the market level.

Section 2 presents the theoretical setup and solutions under different expectations regimes. Sections 3 and 4 describe the experimental design and results, while section 5 concludes.

2. Theoretical Setup and Solutions for Different Expectation Assumptions

2.1 Setup

The theoretical setup is one in which agents must infer the value of a single parameter drawn from a distribution:

$$\lambda \sim (\Lambda, \sigma^2)$$

Agents are presented with a sequence of noisy signals $\lambda + e_t$ $e_t \sim (0, \sigma_e^2)$ which contain information about this unknown parameter:

$$x_1 = \lambda + e_1$$

$$x_2 = \lambda + e_2$$

.....

$$x_t = \lambda + e_t$$

and they must guess λ after each x_i is revealed. We now outline two RE solutions.

2.2 *The Maximum Likelihood RE Solution*

The simplest solution is simply to average the signals:

$$y_t^{MLE} = \frac{\sum_{i=1}^t x_i}{t} = \lambda + \frac{\sum_{i=1}^t e_i}{t} \quad (1)$$

Conditional on λ , this has all the desirable least-squares and maximum likelihood properties

$$E(y_t^{MLE} | \lambda) = \lambda$$

$$V(y_t^{MLE} | \lambda) = \frac{\sigma_e^2}{t}$$

and the Gauss-Markov theorem holds.⁵

Nevertheless, the estimator does not use all information. There is prior information on the distribution of λ , namely Λ and σ^2 .

2.3 The Bayesian Signal Extraction RE solution

Consider a class of solutions of the form:

$$\begin{aligned} y_1^{Bayes} &= \theta \Lambda + (1 - \theta) x_1 \\ y_t^{Bayes} &= \theta y_{t-1}^{SE} + (1 - \theta) x_t \end{aligned} \quad (2)$$

In other words, the first guess is based on a weighted average of the known mean of λ , namely Λ , and the signal, with the optimal weights to be determined presently. Thereafter, agents use a recursion with the same weights.

In general, optimal weights for (2) depend on the time horizon for which signals are available. But the solution requires numerical optimization, so we derive the optimal weight θ for the first period, and assume it prevails in the RE solution.⁶ Consider the first period:

$$\begin{aligned} y_1^{Bayes} &= \theta \Lambda + (1 - \theta) x_1 \\ &= \theta \Lambda + (1 - \theta) (\lambda + e_1) \\ &= (1 - \theta) e_1 + \lambda + \theta (\Lambda - \lambda) \\ \Rightarrow y_1^{Bayes} - \lambda &= (1 - \theta) e_1 + \theta (\Lambda - \lambda) \end{aligned}$$

We now seek θ that will minimize the expected value of the mean squared error (MSE).

When taking the expectation, we do not condition on λ , because we are seeking an optimal

⁵ That is, conditional on λ it is the minimum variance estimator among the class of all linear unbiased estimators.

⁶ The formula for y^{Bayes} in (2) can be written as $(1 - \theta)[e_t + \theta e_{t-1} + \dots + \theta^{t-1} e_1] + \lambda + \theta^t (\Lambda - \lambda)$. An

Optimal θ which minimize $\sum_{i=1}^t E(y_i^{Bayes} - \lambda)^2/t$, as in the text, requires a numerical solution, as it is a function of

t. When this expression is optimized for $t=8$, as in our experiment, the optimal theta is 0.71. In real-world settings, RE must sometimes resort to some ‘fast and frugal heuristics’ (Gigerenzer et al., 1999) as we do in the text.

weight that takes account of its distribution. We first derive the unconditional mean and variance ($V[.]$):

$$E(y_1^{Bayes} - \lambda) = E(1 - \theta)e_1 + \theta E(\Lambda - \lambda) = 0$$

$$\begin{aligned} V(y_1^{Bayes} - \lambda) &= V((1 - \theta)e_1) + V(\theta(\Lambda - \lambda)) \\ &= ((1 - \theta)^2 \sigma_e^2 + \theta^2 \sigma^2) \end{aligned}$$

We now minimize the MSE:

$$\begin{aligned} E[(y_1^{Bayes} - \lambda)^2] &= V(y_1^{Bayes} - \lambda) + \{E(y_1^{Bayes} - \lambda)\}^2 \\ &= ((1 - \theta)^2 \sigma_e^2 + \theta^2 \sigma^2) + 0 \end{aligned}$$

$$\frac{dE[(y_1^{Bayes} - \lambda)^2]}{d\theta} = -2(1 - \theta)\sigma_e^2 + 2\theta\sigma^2 = 0$$

$$\Rightarrow \theta = \frac{\sigma_e^2}{\sigma_e^2 + \sigma^2}$$

Our RE solution is called y^{Bayes} because this is a classic Bayesian signal extraction solution.⁷ We now outline our IE solution to this problem, before turning to the experimental results.

2.4 Inferential Expectations

In each period agents conduct a hypothesis test based on a previous guess of the parameter. The hypothesis test is of the following form.

From the first period until the first moment of rejection:

$$H_0: \lambda = \Lambda$$

$$H_1: \lambda \neq \Lambda$$

⁷ One subtle difference is that we do not require normal distributions, even though, in fact, we do use normal distributions for our experiment.

That is, under IE all agents start with a null belief about λ which is that it is equal to the mean of its distribution, namely Λ . In the experiment, Λ is unity. As they receive signals, they do the above hypothesis test based on the average of the signals. Since they conduct a hypothesis test assuming the null-hypothesis value for the signal, the test can be conditioned on that hypothesized λ , so it is a classic sample-mean hypothesis test, where only the volatility of e is relevant.

When the null is rejected, in period 's', the IE solution is the RE solution at that instant, y_s^{RE} . This value holds in every period thereafter, until the next rejection. We have:

$$H_0: y = y_s^{\text{RE}}$$

$$H_1: y \neq y_s^{\text{RE}}$$

This process continues indefinitely. Whenever the null is rejected, the IE solution jumps to the RE solution at the moment of rejection.

The test statistic for the hypothesis test is based on the average of the signals, since its properties for statistical inference are well known.⁸

$$Z_{\text{calc}} = \frac{\bar{x} - y^{\text{RE}}}{\left(\frac{\sigma_e}{\sqrt{t}} \right)} \sim N(0,1) \quad \bar{x} = \frac{\sum x_i}{t}$$

Reject H_0 if Z_{calc} is in a rejection region of size α .

Since we have outlined two possible RE solutions, we have to allow for the possibility that IE agents use either form of the RE solution at the moment of H_0 rejection. If agents are thought to use y^{MLE} , we call them MLE IE (and, naturally MLE RE if $\alpha=1$). If agents are thought to use y^{Bayes} , we call them Bayes IE (Bayes RE if $\alpha=1$). In the experiment, we vary the frame (the signal is next period's exchange rate, or, next period's share price) and we vary

⁸ If agents are given the true variances and they know the parent distribution of e is Normal, they may use the Normal distribution, even for a small sample size.

what agents know (sometimes they know the variance of λ and e and sometimes they do not). But in each case the IE solution will consist of periodic alignment with y^{Bayes} or y^{MLE} interspersed with a ‘no change’ default in between.

The calculation of α proceeds as follows. Given a profile of actual play (or, bids and offers at an individual level), we simulate a range of IE profiles incrementing α from zero to unity by small steps (0.001). Thus for every value of α there is a profile of IE play, and we can determine how ‘close’ the IE profile is to actual play by measuring the sum of squares.⁹ We then choose the value of α for the IE profile that minimizes this sum of squares.

This is illustrated below. On the left, the eight signals x_i based on a draw of $\hat{\lambda} = 0.76$ are shown.¹⁰ On the right, agents, who see the signals, would have different IE profiles for $\alpha=0$, 0.5 and 1.0. When $\alpha = 1$, we have plotted Bayes RE, corresponding to the Bayesian solution. When $\alpha = 0$, we have total insensitivity to evidence, and the belief never shifts from the null.

(Insert Figure 1 about here.)

Our algorithm repeated this for one thousand possible α profiles from 0.001 to 1.0 (naturally, this is not possible to plot) and selected the minimum sum-of-squares α .

2.5 Adaptive Expectations

Under Adaptive Expectations, agents follow $y_t^{\text{ADE}} = \phi y_{t-1}^{\text{ADE}} + \text{error}$ and there is no theoretical value for ϕ . It is obtained by Least-Squares when we analyze our results.

⁹ As discussed above, there are two possible IE profiles, Bayes IE and MLE IE, depending upon what is assumed to be RE (Maximum Likelihood or Bayesian). We therefore calculate the minimum-sum-of-squares error α for both Bayes and MLE IE.

¹⁰ The random variable λ has itself been draw from $N(\Lambda=1, \sigma_e^2=0.25^2)$.

3. Experimental Design

The experiment was conducted between January and June 2008 at the University of East Anglia.¹¹ Apart from the experimental instructions and a control questionnaire, the experiment was fully computerized. A total of 240 subjects participated in the 48 sessions: five subjects participated to each session, and they participated to one of four treatments (discussed below), giving a total of 12 independent observations per treatment. Subjects were randomly seated in the laboratory. Computer terminals were partitioned to avoid communication by visual or verbal means. Subjects read the experimental instructions and answered a control questionnaire before being allowed to proceed with the tasks. Experimental supervisors individually advised subjects with incorrect answers in the questionnaires.

The experiment lasted up to 2 hours and was divided into four *stages*, each divided in turn in 9 trading *periods*. The assets traded were either foreign currency (in an exchange rate frame) or shares (in a financial frame). Subjects did not know exactly what the intrinsic value λ of the asset was, though they knew that assets retained the same value λ throughout each stage though not across stages. They also received signals from the beginning of period 2 of each stage, as described below.

In the ‘limited prior knowledge’ treatments, subjects only knew that the asset value was drawn from a prior symmetric distribution with mean 1 (i.e., that on average the stock or foreign exchange could be transferred one-for-one into pounds at the end of the stage), but they did not know the exact shape or, crucially, standard deviation of the prior distribution. As a result, Bayes RE or Bayes IE was not feasible unless we assume that subject behaved as if they made auxiliary assumptions. Without these, MLE RE and MLE IE are the only feasible solutions.

In the ‘prior knowledge’ distribution, subjects were provided information about the prior distribution in the form of a ‘table of frequencies’ (effectively, a histogram - see the instructions in the appendix -); information was provided as a table of frequencies rather than

¹¹ The experimental instructions are provided in the appendix.

of probabilities since it is known that subjects process information more efficiently in terms of the former than in term of the latter (e.g. Hertwig and Gigerenzer, 1999). As a result, subjects could infer the standard deviation of the prior distribution, and Bayes RE and Bayes IE become feasible. The distribution actually used was the same in both the ‘limited prior knowledge’ and the ‘prior knowledge’ treatments. It was normal with mean 1 and standard deviation 0.25: only the extent of knowledge about the distribution differed across treatments.

Overall, the experiment had a 2×2 factorial design crossing the frame employed (stock market vs. exchange rate market) and the knowledge about the prior distribution. We had four treatments overall, AE (knowledge, exchange rate frame), BE (limited knowledge, exchange rate frame), AS (knowledge, stock market frame) and BS (limited knowledge, stock market frame).

(Insert Table 1 about here.)

At the beginning of each period subjects received 10 units of each of two assets (foreign and home currency in one frame, cash and shares in the other frame), with only imperfect information about their intrinsic value λ at which they would ultimately be converted in experimental points. The information that subjects had differed depending on the treatment, and is described below, but was common across subjects of any given session. Trade occurred according to a Walrasian clearinghouse auction mechanism, with each subject being asked to provide a buying price and a selling price. The excellent efficiency properties of Walrasian auctions are well-known, and close to that of double auctions (e.g., Holt, 1995).¹² Trade could only take place using the endowments received at the beginning of the relevant trading period. A generic experimental session is represented in the diagram below.

(Insert Figure 2 about here.)

¹² A Walrasian auction has the additional advantages of making certain that an equilibrium market price (exchange rate) would be formed (as bids and offers were required on the part of each trader), that there was a single value of the exchange rate per stage and that all trade was completed, with all bids and offers being elicited, within a shorter amount of time than a double auction (in an experiment that could already last over two hours, time was very much at a premium).

At the beginning of each stage a new value for the true value λ of the asset was drawn applicable to the stage, and subjects knew that. From the beginning of period 2 onwards, subjects received an independent noisy signal about the value of the asset of the kind discussed in section 2, i.e. of the form $\lambda + e_t$ $e_t \sim (0, \sigma_e^2)$. Subjects received a table of frequencies illustrating the (normal) distribution of the signal around its true value (see the appendix). In AE and AS, they were also shown the a table of frequencies of λ .

In all treatments, to help with understanding, subjects also received an example sheet showing examples of how the economy may work in each stage. At the end of each stage, the units of the asset obtained were converted into home currency (exchange rate frame) or cash (stock market frame) using the true value λ . At the end of the experiment, one period was chosen at random by the computer and each unit of home currency or cash earned in that period (whether directly or after the conversion using λ) was translated into U.K. pounds at the rate of 1 pound per unit. Mean experimental payment were 22-23 UK pounds (roughly 40 US dollars) per subject for between 1 ½ and 2 hours of work.

4. Experimental Results¹³

Subjects generally appeared to have a good understanding of the experiment, and asset prices were almost always in the 0.5 to 1.5 range with a mode around 1, which is what we would expect with our $N \sim (1, 0.25)$ distribution of true asset values (see Figure 3).

(Insert Figure 3 about here.)

Subjects bought asset units 62% of the times (between 59% and 64% depending on the treatment). When they did buy, they generally bought between 5 and 10 units in all treatments: the average positive amount bought was stable across treatments and was around 7 units (6.96). Figure 4 illustrates the dynamics of exchange rates, and, as an example, selling

¹³ All P values reported in this paper are two tailed.

prices from an AE treatment session. There is some evidence of hold – move patterns at the level of individual buying and selling prices while market prices are smoother in their adjustment process.¹⁴

(Insert Figure 4 about here.)

4.1 Estimating the degree of belief conservatism

We estimate α and β values that, for each expectation model and any given sample x , minimize the sum of squared errors between observed and predicted exchange rates in the given sample x . We call this α (or β) value the (error-minimizing) α (or β) by x . Typically there will be a range of α values with the same prediction, and so there will be an α_{\min} and an α_{\max} by x , respectively denoting the minimum and the maximum value of α within the error-minimizing range.

Our main attention is on Bayesian and MLE IE α and β values by session, estimated on the basis of market prices. It is useful, however, to compare these values with the average α (and β) values by individual subject, estimated on the basis of each subject's buying prices *or* selling prices. Table 2 provides information on both sets of data, and Figure 5 contains histograms on the distribution of α values.¹⁵

(Insert Figure 2 and Table 2 about here.)

While estimated β values tend to cluster around 1, average α values by session are around 0.6 in the case of Bayesian IE and 0.45 in the case of MLE IE. These averages, however, ignore the bimodality of the data as displayed in Figure 5. About 1/3 to 1/2 of the sessions have α values close to 0, displaying belief conservatism, whereas the rest have α values close to 1, implying continuous belief updating as in RE models.

¹⁴ Smoother market prices may be consistent with heterogeneous α values by individual subjects, in the same way in which in standard consumer theory a continuous aggregate demand schedule (e.g. a continuous market demand for cars, in Nicholson's, 2005) may follow from discontinuous (but heterogeneous) individual demand schedules (e.g., each consumer wishing to buy only one car).

¹⁵ β values are all clustered at values around or just below 1 (e.g., the β values by session range from 0.891 to 1.005).

The only treatment which appears to have lower α values is the AE treatment. If we compare MLE α values in the AE treatment with those in the other, we find that the difference is statistically significant (Mann-Whitney $P < 0.02$) except in relation to selling prices (*n.s.*). Bayesian α values appear more stable and analogous tests show $P < 0.05$ significance only for individual buying prices.¹⁶ That being said, nonparametric Kruskal Wallis tests or, say, Mann Whitney tests comparing treatments with exchange rate markets with treatments with stock markets, or treatments with limited knowledge with those with full knowledge of the prior distribution, are all statistically insignificant, and so there fails to be systematic evidence of an effect of providing more information on the degree of belief conservatism.

RESULT 1. The belief conservatism parameter α is either close to 0 (belief conservatism) or close to 1 (continuous belief updating). Market prices in between 1/3 and 1/2 of the sessions exhibit a degree of belief conservatism. Generally speaking, prior knowledge of the prior distribution and the type of market are not statistically significant factors, although there is some suggestive evidence of lower α values in the AE treatment, especially when computed using MLE.

Figure 5 and Table 2 also show a correspondence between this bimodality at the session level and one at the level of individual traders' buying and selling prices. Spearman correlation coefficients between α by session and average α by individual buying (selling) prices show a strong connection between the two, with ρ correlation coefficients between 0.637 and 0.744 (0.491 and 0.632) in relation to Bayes IE (MLE IE) depending on the

¹⁶ The P value is 0.089 in relation to Bayesian α_{\min} values based on market prices and above 0.1 in all other cases.

treatment and whether α_{\min} or α_{\max} is employed ($P < 0.001$ in all cases).¹⁷ Figure 6 and Table 3 help illustrating the connection between the two.

(Insert Figure 6 about here.)

They show that, when the average α by individual is below a lower threshold between 0.3 and 0.45 (depending on the specific measure), then it is almost always the case that the market price α by session is close to 0. Conversely, when the average α by individual is above an upper threshold between 0.4 and 0.6, it is generally the case that the market price α by session is close to 1. This correspondence between α values based on individual buying and selling prices, and α values based on market prices, is robust across treatments.

RESULT 2. The belief conservatism parameter α is bimodal in relation to individual buying prices and selling prices, with the two modes being either close to 0 (belief conservatism) or close to 1 (continuous belief updating). Prior knowledge of the prior distribution and the type of market are not statistically significant factors.

RESULT 3. There is an across treatments consistent correspondence between average α values based on individual buying or selling prices and the aggregate α values based on market prices. When the average α by individual buying or selling prices is above an upper threshold (between around 0.4 and 0.6), market prices will generally display α close to 1. When the average α by individual buying or selling prices is below a lower threshold (between around 0.4 and 0.6), market prices will typically display α close to 0.

¹⁷ Similarly, average β values by individual are positively correlated with market β values by session (Spearman $\rho = 0.702$ in relation to buying prices and 0.696 in relation to selling prices, both $P < 0.001$). Conversely, and plausibly, α and β values tend to be negatively correlated.

Another aspect of the relationship between α values based on market prices and those based on individual buying or selling prices is that α values based on market prices are more polarized around $\alpha = 0$ or $\alpha = 1$ than those based on individual buying or selling prices. In relation to market α values, the average distance from either 0 or 1 is between 0.021 and 0.135 depending on the treatment and on the specific measure used (Bayes or MLE, α_{\min} or α_{\max}); in relation to α values based on individual buying or selling prices, the average distance by session from either 0 or 1 ranges from 0.271 to 0.414, as each session is likely to have a mix of low and high α values type. The difference is always statistically significant in any corresponding pairwise comparison ($P < 0.005$ in a Wilcoxon test).

RESULT 4. The market aggregation process operates by polarizing market prices into a more bimodal pattern displaying either close to 0 (belief conservatism) or close to 1 (continuous belief updating) than what is observed at the level of individual buying and selling prices.

That being said, Table 2 above shows that, on average, Bayesian market α values are higher than those estimated on the basis of individual buying and selling prices: this was true in 32 sessions, against 16 where it was not true ($P = 0.03$), and is not especially restricted to any given treatment. The same statistically significant conclusions do not hold for MLE market α values, while β values are higher than the average β values by individual in 47 out of 48 session ($P < 0.001$).

RESULT 5. β values and Bayesian α values based on market prices tend to be higher than those by individual buying and selling prices.

We also estimated α (and β) values by stage. It is interesting to check whether market prices exhibited less belief conservatism as the experiment progressed. As shown by Table 4, there is no evidence that this is the case for average Bayesian IE α values, and, after stage 1, there is no evidence of any systematic change in average MLE α values as well.¹⁸

(Insert Table 4 about here.)

RESULT 6. There is no evidence of belief conservatism disappearing as traders gain more experience, at least after stage 1.

4.2 Comparing the goodness of fit of different classes of models

We now use the sum of squared errors (SSE) between predicted and observed exchange rates to compare the goodness of fit of the different classes of models: Bayesian RE or IE versus MLE RE or IE and versus ADE. Figure 7 compares the goodness of fit of different expectation models in predicting asset values.

(Insert Figure 7 about here.)

In each of the four treatments, there is a clear ranking, apart from adaptive expectations. The least close to actual market play is the MLE RE solution, which has the highest SSE in each treatment, followed by MLE IE. Next closest to actual play is the Bayesian RE solution, followed by its Bayesian IE equivalent. The superiority of IE over RE is unsurprising in the light of section 4.1, since it trivially follows from the estimated α being less than 1 in a number of sessions, implying a degree of belief conservatism.

The superiority of the Bayesian models over the MLE models is, however, surprising, as it applies equally to all treatments, including the BE and BS treatments with limited knowledge about the prior distribution. This superiority is unequivocally confirmed in

¹⁸ Similarly, there is no meaningful systematic trend in relation to β .

Wilcoxon tests ($P \leq 0.005$), and is also confirmed if we look at SSE in relation to individual buying or selling prices. Across treatment difference in SSE are never statistically significant in Kruskal Wallis or Mann Whitney tests.

RESULT 7. Bayesian expectation models have better predictive power than MLE expectation models in all treatments, including those with limited knowledge about the prior distribution.

The SSE in relation to ADE is lower than that in relation to Bayesian IE and RE in treatment AE, but not statistically significantly so. It is higher in the other treatments. Overall, at the market level the Bayesian models outperform ADE, marginally so in relation to Bayesian RE ($P = 0.055$ in a Wilcoxon test) and more clearly so in relation to Bayesian IE ($P = 0.03$) which, like ADE, has one estimated degree of freedom (α instead of β). These results contrast with the superiority of ADE over Bayesian RE at the level of individual buying and selling prices, whereas there are no statistically significant differences in goodness of fit between Bayesian IE and ADE at that level. There is a sense, therefore, that markets are ‘more rational’ (in terms of model consistency) than individuals.

RESULT 8. At the market level, Bayesian expectation models globally outperform ADE expectation models.

5. Conclusion

There are three main take home messages from this paper. The first is that explicit knowledge of the prior distribution of true asset values is not crucial for subjects to behave

(more) rationally. The second is that the Bayesian signal extraction framework describes behavior just as well without such knowledge as it does with it.

The third is that there is support for the modelling practice in behavioral finance and economics which assume two types of agents, perfect rational agents and some form of not equally rational agents. Our belief conservatism parameter α was bimodal – either close to 0 or 1 – for individual agents, and this bimodality translated into a bimodality of belief conservatism in market prices.

The effects of the market trading are particularly interesting in this regard. It appears to turn the somewhat bimodal distribution of individual bids and offers into a very stark bimodal market distribution. If the average α values by individual were above or below given thresholds, which we quantified as being approximately between 0.3 and 0.6, the market outcomes neatly separated into near-unity and near-zero. Our results suggest that simplifying agent heterogeneity into two groups should not be necessarily regarded as a ‘mere’ modelling device, though modellers are on stronger ground when they use this device for modelling market play than they are when they use it for modelling individual decisions.

Our take home messages are robust to whether a stock market frame or an exchange rate market frame is used, and so have a degree of generality.¹⁹ Interesting, allowing for estimating one degree of freedom in an adaptive expectation model did not enable adaptive expectations to outperform Bayesian signal extraction models. It is also interesting that, on average, according to our best performing Bayesian signal extraction framework, markets were ‘more rational’ than individuals, as illustrated by higher α values and by the better goodness of fit relative to adaptive expectations. While this result is in line with research on the disciplining effects of markets (e.g., Evans, 1997), it will be interesting to determine in future research why this result was obtained as opposed to the strong belief conservatism result at the market level found in the setup of Menzies and Zizzo (2007). Obviously, further experimental research is needed.

¹⁹ The AE treatment had suggestive evidence of greater belief conservatism than the other treatments, but, as we discussed, the evidence is not robust if we use the Bayesian signal extraction framework to compute α values.

REFERENCES

- BARBERIS, N. and THALER, R. (2003), "A Survey of Behavioral Finance" in Handbook of the Economics of Finance, eds. G. M. Constantinides, M. Harris and R. Stulz, Chapter 18, 1051-1121.
- CAMERER, C. (1995), "Individual Choice" in Kagel, J. H. and Roth, A. E. (eds.), *Handbook of Experimental Economics*, Princeton: Princeton University Press, 587-703.
- CARROLL, C. D. (2003), "Macroeconomic Expectations of Households and Professional Forecasters" *Quarterly Journal of Economics*, **118**, 269-298.
- ELLSBERG, D. (1961), "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, **75**, 643-669.
- EVANS, D. (1997), "The Role of Markets in Reducing Expected Utility Violations," *Journal of Political Economy*, **105**, 622-636.
- GIGERENZER, G., TODD, P. M. and THE ABC RESEARCH GROUP (1999), *Simple Heuristics That Make Us Smart*, Oxford: Oxford University Press.
- HERTWIG, R. and GIGERENZER, G. (1999), "The 'Conjunction Fallacy' Revisited: How Intelligent Inferences Look Like Reasoning Errors", *Journal of Behavioral Decision Making*, **12**, 275-305.
- HOLT, C. A. (1995), "Industrial Organization: A Survey of Laboratory Research," in: John H. Kagel and Alvin E. Roth (eds.), *The Handbook of Experimental Economics*, Princeton: Princeton University Press, 349-443.
- KNIGHT, F. H. (1921), *Risk, Uncertainty, and Profit*, Boston, MA: Hart, Schaffner & Marx.
<http://www.econlib.org/LIBRARY/Knight/knRUP1.html>
- MANKIW, N. G. (2001), "The Inexorable and Mysterious Tradeoff between Inflation and Unemployment," *Economic Journal*, **111**, C45-61.
- MANKIW, N. G. and REIS, R. (2003), "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, **117**, 1295-1328.
- MENZIES, G. D. and ZIZZO, D. J. (2005), "Inferential Expectations", Australian National

University Centre for Applied Macroeconomic Analysis Discussion Paper n. 12, June.

MENZIES, G. D. and ZIZZO, D. J. (2007), “Exchange Rate Markets and Conservative Inferential Expectations”, Australian National University Centre for Applied Macroeconomic Analysis Discussion Paper n. 2, January.

NICHOLSON, W. (2005), *Microeconomic Theory: Basic Principles and Extensions*, Mason, Ohio: Thomson.

SARIN, R. K. and WEBER, M. (1993), “Effects of Ambiguity in Market Experiments,” *Management Science*, **39**, 602-615.

SIMS, C. (2003), “Implications of Rational Inattention” *Journal of Monetary Economics*, **50**, 665–690.

Appendix – Experimental Instructions

Treatments AE and BE

Experimental Instructions

1. Introduction

This is an experiment in the economics of market decision making. The instructions are simple and, if you follow them carefully and make good decisions, you might earn a considerable amount of money which will be paid to you in cash. Please raise your hand if you have any questions at any point in the experiment.

The session is divided in four **stages**. Each stage is divided in 9 **periods**. At the start of each period, you receive an endowment of 10 **home currency** units and 10 **foreign currency** units. Market trading then occurs, according to the rules specified below.


2. Information on the Economy

The exchange rate tells you how much home currency you need to purchase one foreign currency unit. For example, if the exchange rate is 1.25, it means that you need 1.25 home currency units to buy 1 foreign currency unit, while an exchange rate of 1 means that you need exactly 1 home currency unit to buy 1 foreign currency unit. The higher the value of the exchange rate, the more the foreign currency is worth relative to the home currency.

At the end of each stage, after the final period trade is done, the foreign currency units you have obtained in each period are converted into home currency using the **end of stage exchange rate**. The **end of stage exchange rate** is randomly generated from a distribution with an **average end of stage exchange rate equal to 1**. That is to say, if you were to choose randomly the value of the end of stage exchange rate 1000 times, you would find that the average value of the exchange rate would be 1, i.e. that home currency and foreign currency

would be worth the same. How different the end of stage exchange rate is from 1 has been determined by chance.

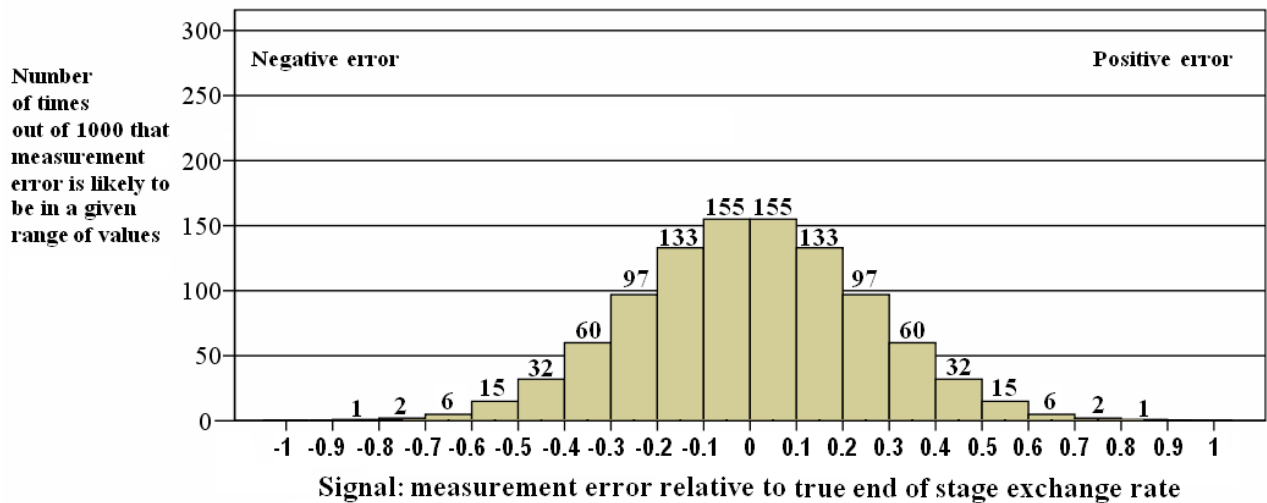
[Treatment AE only:]



The graph just means that, if you were to find 1000 occurrences of the end of stage exchange rate and were to count up the number of times out of 1000 that the end of stage exchange rate is in a given range of values, you would obtain the distribution in the picture (or something very close to it). Out of 1000 occurrences of the end of stage exchange rate, 155 are most likely to occur in the range between 1 and 1.1 (or 0.9 and 1), 133 in the range between 1.1 and 1.2 (or 0.8 and 0.9), 97 in the range between 1.2 and 1.3 (or 0.7 and 0.8), and so on.

[Both treatments:]

Once the end of stage exchange rate is picked randomly, it is fixed for the entire stage. You do not observe the true value of the end of stage exchange rate as such during the stage, but, from the beginning of period 2 of each stage, at the start of each period you receive a **signal** about its value for the given stage. The signal is equal to the true value of the end of stage exchange rate plus a **measurement error**. The size and direction of the measurement error has been determined by chance.



The graph just means that, if you were to find 1000 occurrences of the signal and were to count up the number of times out of 1000 that the measurement error is in a given range of

values, you would obtain the distribution in the picture (or something very close to it). Out of 1000 occurrences of the measurement error, 155 are most likely to occur in the range between 0 and 0.1 (or 0 and -0.1), 133 in the range between 0.1 and 0.2 (or -0.1 and -0.2), and so on.

As you move from one stage to the next, the economy is reset to its initial conditions and a new end of stage exchange rate is determined, with an average value of the end of stage exchange rate once again equal to 1.

The enclosed Examples Sheet provides examples both (a) of end of stage exchange rates drawn from the distribution with an average of 1 and (b) of the kind of signals that, given the measurement error, you may expect in the experiment for given values of end of stage true exchange rate.

3. Market Trading

A market opens up each period for foreign currency. The choices that you and the other market traders make determine the value of the exchange rate. You set:

- a **selling price** at which you would be willing to sell foreign currency, in exchange for home currency;
- a **buying price** at which you would be willing to buy foreign currency, in exchange for home currency.

The selling price must be higher than the buying price. Prices are always set in home currency units. Note that they need not be round numbers: for example, you could set 1.11 as selling price and 1.09 as buying price.

We ask you to make your choices in less than two minutes and thirty seconds. Please stick to this timetable, as it is essential for the session to finish in up to around two hours as scheduled.

Once everyone has entered their choices, a central computer determines the “market-clearing price” at which the total number of foreign currency sold is equal to the total number of foreign currency bought. The exchange rate for the period is this market-clearing price; if there is a range of market-clearing prices, the computer chooses the middle one as the exchange rate. Trade then occurs at this exchange rate: foreign currency is bought in exchange for home currency at the exchange rate for the period.

If you have stated a buying price higher than or equal to the exchange rate, then you are a buyer for the period: you and each of the other buyers buy an equal amount of the available supply, subject to none of you becoming bankrupt (if you risk becoming bankrupt, you simply buy as much as you can afford). If you have stated a selling price lower than or equal to the exchange rate, then you are a seller for the period: you and each of the other sellers sell an equal amount to buyers, as long as you have enough foreign currency to be able to do so (if you have not got enough, then you simply sell all the foreign currency you have). If your buying price is lower *and* your selling price is higher than the exchange rate, then you do not trade for the period; if everyone is in this situation, then no one trades for the period.

The computer will inform you about the exchange rate, about the amount of foreign currency that you have been able to buy or sell, and about the overall amount of home and foreign currency that you now own.

As you move from one period to the next, you can trade with a new endowment of 10 home currency units and 10 foreign currency units.

4. Session Earnings

At the end of the session one of the periods will be picked up at random. Your overall earnings for this period – namely the sum of your home currency units and of your foreign currency units, converted into home currency units at the corresponding end of stage exchange rate - will be converted into UK pounds at the rate of **1 pound (one pound)** per home currency unit.

6. Notes on the System

Your selling and buying prices are crucial in determining your home and foreign currency holdings on the basis of which you are paid at the end of the experiment. In choosing a selling price, make sure that you would only be willing to sell your foreign currency if the price were that high or higher. In choosing a buying price, make sure that you would only be willing to buy foreign currency if the price were that low or lower. Remember that prices with decimal places, such as 1.36 and 0.64, are perfectly acceptable.

How should you determine your buying prices and selling prices if you want to make money out of the experiment? It makes sense for you to consider the end of stage exchange rate. You would want to buy foreign currency if the exchange rate were below your best guess of the end of stage exchange rate, and you would want to sell foreign currency if the exchange rate were above your best guess of the end of stage exchange rate.

Before starting to take decisions, we ask you to fill the enclosed questionnaire, with the only purpose of checking whether you have understood these instructions. Raise your hand when you have completed the questionnaire.

Treatments AS and BS

Experimental Instructions

1. Introduction

This is an experiment in the economics of market decision making. The instructions are simple and, if you follow them carefully and make good decisions, you might earn a considerable amount of money which will be paid to you in cash. Please raise your hand if you have any questions at any point in the experiment.

The session is divided in four **stages**. Each stage is divided in 9 **periods**. At the start of each period, you receive an endowment of 10 units of **cash** and 10 **stock** market shares in a company. Market trading then occurs, according to the rules specified below.


2. Information on the Economy

The share price tells you how much cash you need to purchase one stock market share. For example, if the share price is 1.25, it means that you need 1.25 units of cash to buy 1 stock

share, while a share price of 1 means that you need exactly 1 unit of cash to buy 1 stock share. The higher the stock price, the more stock shares are worth relative to cash.

At the end of each stage, after the final period trade is done, the stock shares you have obtained in each period are converted into cash using the **end of stage share price**. The **end of stage share price** is randomly generated from a distribution with an **average end of stage share price equal to 1**. That is to say, if you were to choose randomly the value of the end of stage share price 1000 times, you would find that the average value of the share price would be 1, i.e. that a unit of cash and a stock share would be worth the same. How different the end of stage share price is from 1 has been determined by chance.

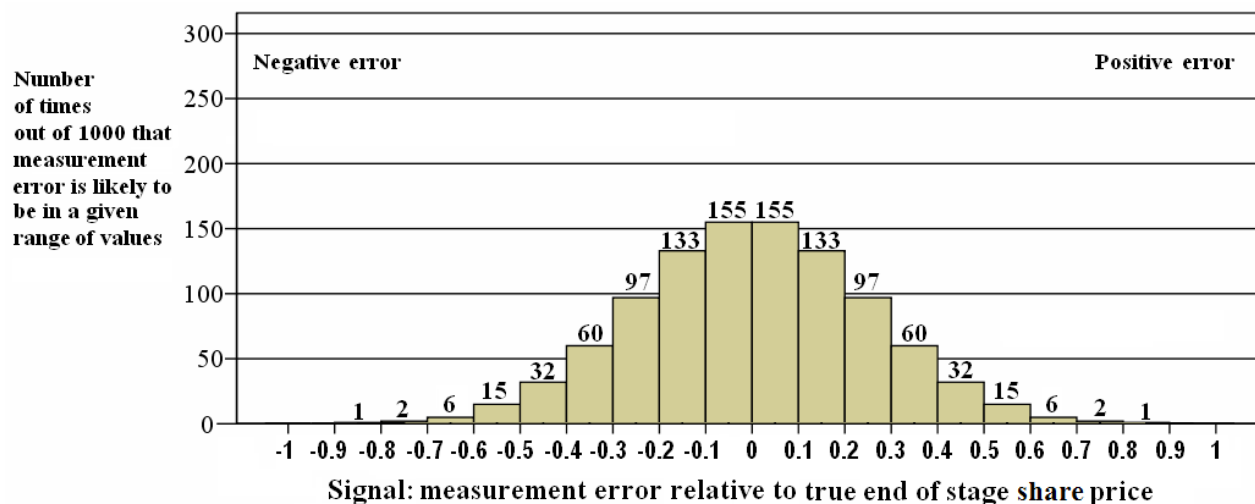
[*Treatment AS only:*]



The graph just means that, if you were to find 1000 occurrences of the end of stage share price and were to count up the number of times out of 1000 that the end of stage share price is in a given range of values, you would obtain the distribution in the picture (or something very close to it). Out of 1000 occurrences of the end of stage share price, 155 are most likely to occur in the range between 1 and 1.1 (or 0.9 and 1), 133 in the range between 1.1 and 1.2 (or 0.8 and 0.9), 97 in the range between 1.2 and 1.3 (or 0.7 and 0.8), and so on.

[*Both treatments:*]

Once the end of stage share price is picked randomly, it is fixed for the entire stage. You do not observe the true value of the end of stage share price as such during the stage, but, from the beginning of period 2 of each stage, at the start of each period you receive a **signal** about its value for the given stage. The signal is equal to the true value of the end of stage share price plus a **measurement error**. The size and direction of the measurement error has been determined by chance.



The graph just means that, if you were to find 1000 occurrences of the signal and were to count up the number of times out of 1000 that the measurement error is in a given range of values, you would obtain the distribution in the picture (or something very close to it). Out of 1000 occurrences of the measurement error, 155 are most likely to occur in the range between 0 and 0.1 (or 0 and -0.1), 133 in the range between 0.1 and 0.2 (or -0.1 and -0.2), and so on.

As you move from one stage to the next, the economy is reset to its initial conditions and a new end of stage share price is determined, with an average value of the end of stage share price once again equal to 1.

The enclosed Examples Sheet provides examples both (a) of end of stage share prices drawn from the distribution with an average of 1 and (b) of the kind of signals that, given the measurement error, you may expect in the experiment for given values of end of stage true share price.

3. Market Trading

A market opens up each period for stock shares. The choices that you and the other market traders make determine the share price. You set:

- a **selling price** at which you would be willing to sell stock shares, in exchange for cash;
- a **buying price** at which you would be willing to buy stock shares, in exchange for cash.

The selling price must be higher than the buying price. Prices are always set in units of cash. Note that they need not be round numbers: for example, you could set 1.11 as selling price and 1.09 as buying price.

We ask you to make your choices in less than two minutes and thirty seconds. Please stick to this timetable, as it is essential for the session to finish in up to around two hours as scheduled.

Once everyone has entered their choices, a central computer determines the “market-clearing price” at which the total number of stock shares sold is equal to the total number of stock shares bought. The share price for the period is this market-clearing price; if there is a range of market-clearing prices, the computer chooses the middle one as the share price. Trade then

occurs at this share price: stock shares are bought in exchange for cash at the share price for the period.

If you have stated a buying price higher than or equal to the share price, then you are a buyer for the period: you and each of the other buyers buy an equal amount of the available supply, subject to none of you becoming bankrupt (if you risk becoming bankrupt, you simply buy as much as you can afford). If you have stated a selling price lower than or equal to the share price, then you are a seller for the period: you and each of the other sellers sell an equal amount to buyers, as long as you have enough stock shares to be able to do so (if you have not got enough, then you simply sell all the stock shares you have). If your buying price is lower *and* your selling price is higher than the share price, then you do not trade for the period; if everyone is in this situation, then no one trades for the period.

The computer will inform you about the share price, about the amount of stock shares that you have been able to buy or sell, and about the overall amount of cash and stock shares that you now own.

As you move from one period to the next, you can trade with a new endowment of 10 units of cash and 10 stock shares.

4. Session Earnings

At the end of the session one of the periods will be picked up at random. Your overall earnings for this period – namely the sum of your cash and of your stock shares, converted into units of cash at the corresponding end of stage share price - will be converted into UK pounds at the rate of **1 pound (one pound)** per unit of cash.


6. Notes on the System

Your selling and buying prices are crucial in determining your cash and stock shares holdings on the basis of which you are paid at the end of the experiment. In choosing a selling price, make sure that you would only be willing to sell your stock shares if the price were that high or higher. In choosing a buying price, make sure that you would only be willing to buy stock shares if the price were that low or lower. Remember that prices with decimal places, such as 1.36 and 0.64, are perfectly acceptable.

How should you determine your buying prices and selling prices if you want to make money out of the experiment? It makes sense for you to consider the end of stage share price. You would want to buy stock shares if the share price were below your best guess of the end of stage share price, and you would want to sell stock shares if the share price were above your best guess of the end of stage share price.

Before starting to take decisions, we ask you to fill the enclosed questionnaire, with the only purpose of checking whether you have understood these instructions. Raise your hand when you have completed the questionnaire.

Figure 1. An Example of Market Dynamics



Note: λ is the true asset value and Bayesian α values are shown.

Figure 2. Timeline of Experimental Sessions






Figure 3. Histograms of Asset Values

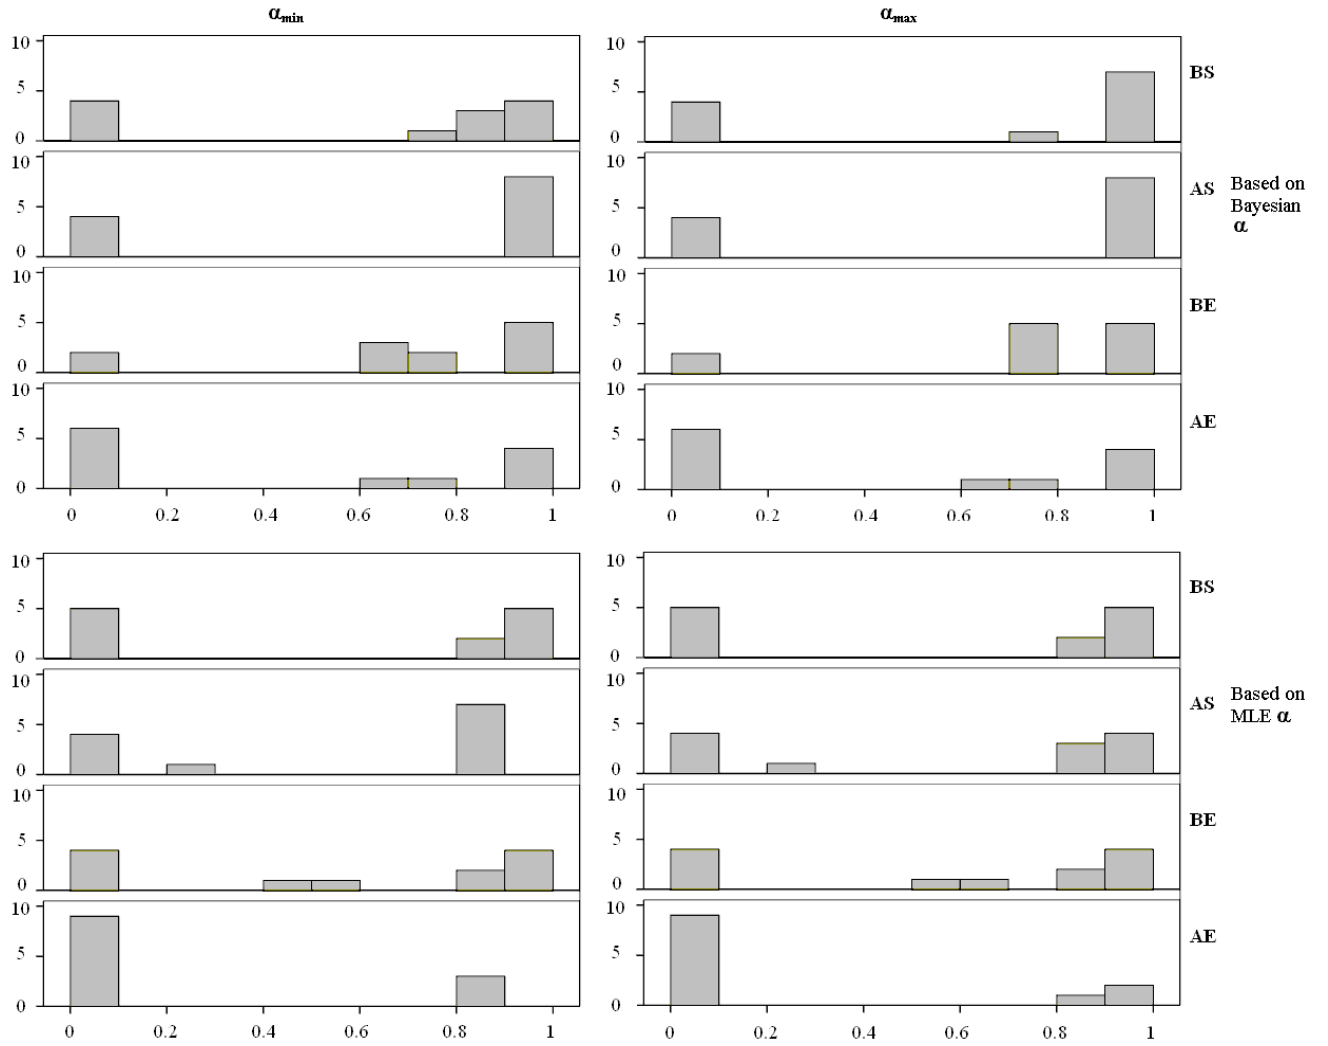
Figure 4. An Example of Market Asset Values and Predictions




Notes: the example is based on a treatment A session. ADE: adaptive expectations.

Figure 5. Histograms of α Values


(a) α Values Based on Market Prices



(b) Bayesian α Values Based on Individual Buying Prices and Selling Prices



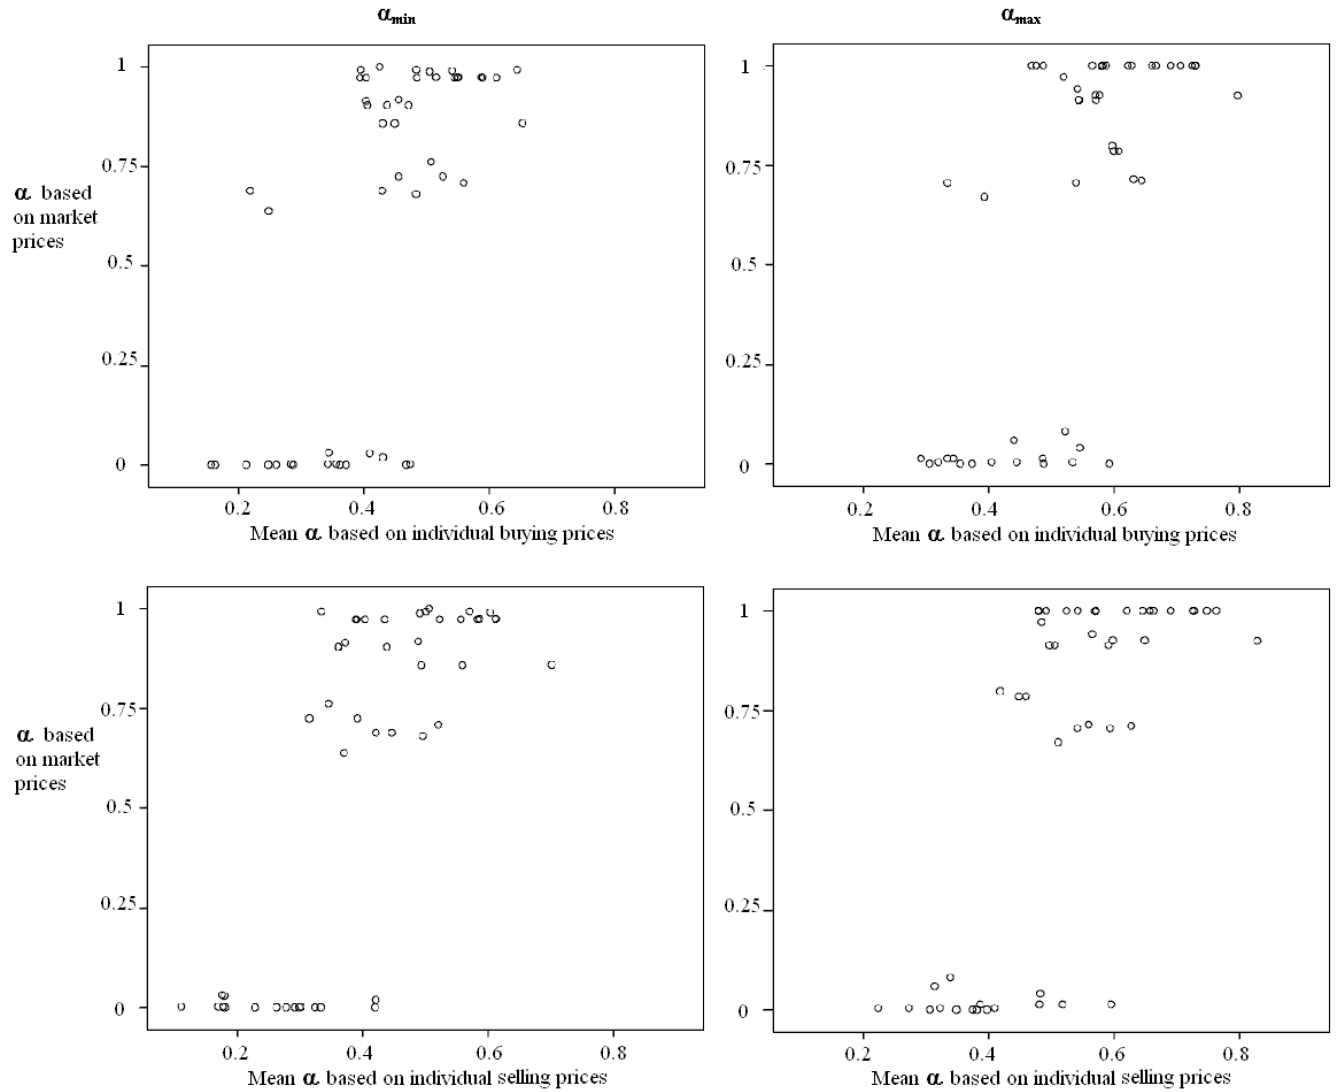
(c) MLE α Values Based on Individual Buying Prices and Selling Prices



Note: the unit of observation is the session in the case of α values based on market prices (panel a) and the individual subject in the case of α values by individual buying prices and selling prices (panels b and c).

Figure 6. Scatterplots Matching Mean α by Individual Buying and Selling Prices with α by Session Based on Market Prices

(a) Bayesian α Values




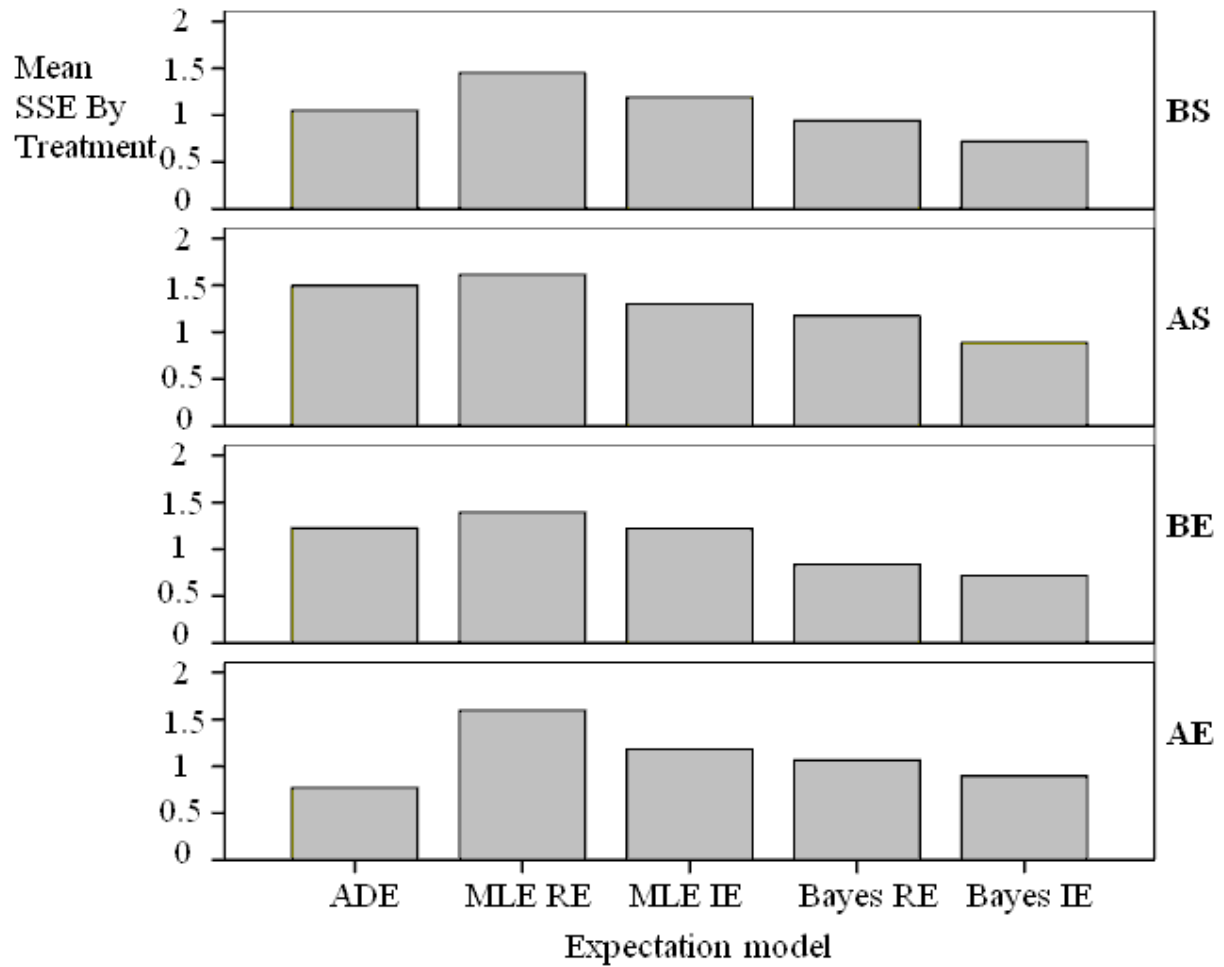
(b) MLE α Values

Figure 7. Goodness of Fit of Different Expectation Models



Notes: SSE: sum of squares error; ADE: adaptive expectations; MLE RE (IE): maximum likelihood estimator based rational expectations (inferential expectations); Bayes RE (IE): Bayesian signal extraction based rational expectations (inferential expectations);

Table 1. Experimental Treatments

Treatment AE	Treatment BE
Knowledge of Prior Distribution, Exchange Rate Frame	No Knowledge of Prior Distribution Exchange Rate Frame
Treatment AS	Treatment BS
Knowledge of Prior Distribution, Stock Market Frame	No Knowledge of Prior Distribution Stock Market Frame

Note: 12 sessions were run in each treatment.

Table 2. Estimated Mean α and β Values by Method and Treatment

		Treatment					
			AE	BE	AS	BS	Overall
Bayes	Market prices	α_{\min}	0.435	0.699	0.644	0.6	0.594
		α_{\max}	0.451	0.718	0.664	0.632	0.617
	Buying prices	α_{\min}	0.4	0.44	0.438	0.432	0.427
		α_{\max}	0.511	0.547	0.537	0.545	0.535
	Selling prices	α_{\min}	0.338	0.437	0.427	0.408	0.403
		α_{\max}	0.459	0.538	0.558	0.521	0.519
MLE	Market prices	α_{\min}	0.221	0.545	0.531	0.529	0.456
		α_{\max}	0.228	0.558	0.543	0.535	0.466
	Buying prices	α_{\min}	0.317	0.345	0.36	0.38	0.351
		α_{\max}	0.409	0.449	0.466	0.476	0.45
	Selling prices	α_{\min}	0.28	0.418	0.372	0.4	0.368
		α_{\max}	0.378	0.509	0.468	0.493	0.464
ADE	Market prices	β	0.987	0.971	0.965	0.974	0.974
	Buying prices	β	0.991	0.918	0.903	0.891	0.926
	Selling prices	β	0.932	0.899	0.884	0.904	0.905

Notes: α and β values were estimated by minimizing sum of squares error between actual and predicted values. Bayes refers to α values estimated based on Bayesian signal extraction. MLE refers to α values estimated based on maximum likelihood estimation. ADE: adaptive expectations estimated by OLS.

Table 3. Approximate Upper and Lower Thresholds on Correspondence between Market Price α Values and Average Individual α Values

			Threshold	
			Lower	Upper
Bayes	Buying prices	α_{\min}	0.4	0.5
		α_{\max}	0.45	0.6
	Selling prices	α_{\min}	0.3	0.4
		α_{\max}	0.4	0.6
MLE	Buying prices	α_{\min}	0.3	0.45
		α_{\max}	0.4	0.5
	Selling prices	α_{\min}	0.35	0.45
		α_{\max}	0.45	0.55

Note: The thresholds are approximate to the closest 0.05 value.

Table 4. Estimated Mean α and β Values By Stage

		Stage 1	Stage 2	Stage 3	Stage 4
Bayes	α_{\min}	0.439	0.505	0.442	0.457
	α_{\max}	0.545	0.618	0.563	0.58
MLE	α_{\min}	0.229	0.424	0.56	0.416
	α_{\max}	0.317	0.551	0.658	0.491
ADE	β	0.969	0.968	0.975	0.962