# Learning to Forecast Price[1]

Hugh Kelley[2] and Daniel Friedman

January 1999

[1] *Acknowledgments.* This work is supported by NSF grants SBR 9310347 and SBR 9617917. It benefitted from the comments of Jules Leichter, Dominic Massaro, Rachel Croson, Vai-Lam Mui and especially Arlington Williams, as well as participants at the Economics Science Association and Public Choice Society meetings at Tucson and New Orleans.

[2] Psychology Department, Indiana University Bloomington, Bloomington, IN 47405. For correspondence contact Hugh Kelley, email: hukelley@indiana.edu, Phone: 812-856-4678.

**Abstract**

We study human learning in a individual choice laboratory task called Orange Juice Futures price forecasting (OJF), in which subjects must implicitly learn the coefficients of two independent variables in a stationary linear stochastic process. The 99 subjects each forecast in 480 trials with feedback after each trial. Learning is tracked for each subject by fitting the forecasts to the independent variables in a rolling regression. Results include: (1) learning is fairly consistent in that coefficient estimates for most subjects converge closely to the objective values, but there is a mild general tendency toward over-response. (2) Typically learning is noticeable slower than the Marcet-Sargent ideal. Among the more striking treatment effects are a general tendency towards (3) over-response with high background noise and (4) under-response with asymmetric coefficients.

# 1 Introduction

Economists in recent years have begun to model how people might learn equilibrium behavior. Microeconomists following Binmore (1987) and Fudenberg and Kreps (1988) consider learning models with roots in Cournot (1838) and Brown (1951). Numerous laboratory studies test and refine the microeconomists' learning models; see Camerer (1998) for a recent survey. There is also a separate theoretical macroeconomics literature on learning following Marcet and Sargent (1989a,b,c) and Sargent (1994); see Evans and Honkapohja (1997) for a recent survey. Here the focus is on how people might learn to forecast relevant prices, and whether the learning process permits convergence to rational expectations equilibrium. We are not aware of any laboratory work intended to test and refine the learning models favored by macroeconomists.[1]

The work presented below examines human learning in an individual choice laboratory task called Orange Juice Futures price forecasting (OJF). The OJF task has a form and complexity similar to the forecasting tasks in macroeconomists' models: subjects must implicitly learn the coefficients of two independent variables in a linear stochastic process. We study stationary versions of the OJF task here in order to sharpen the evidence on learning per se, and leave for future work the nonstationary (or "self-referential") aspect of the macroeconomists' models that prices are endogenous and perhaps affected by individuals' learning processes.

The OJF task is based on the observation of Roll (1984) that the price of Florida orange juice futures depends systematically on only two exogenous variables, the local weather hazard and the competing supply from Brazil. The laboratory experiment consists of many independent trials in which human subjects forecast the OJF price after observing values of the two variables. After each trial the subject receives feedback in the form of the "actual" price generated from a linear stochastic model using the observed values of the two variables. We report results for 99 subjects, each

[1] We hasten to add that several important laboratory investigations have been inspired by other strands of macroeconomic theory. For example, Van Huyck et al (1997) and related work studies equilibrium convergence in coordination games, and Marimon and Sunder (1994) and related work studies sunspot equilibria in overlapping generations economies. Later in this introduction we discuss three laboratory studies of rational expectations equilibrium.

forecasting in 480 trials. Several treatments are varied across subjects, such as the noise amplitude and the relative impact of the two variables.

We are interested in two aspects of learning: consistency and speed. Roughly speaking, learning is consistent to the extent that subjects eventually respond correctly to the exogenous variables, and learning is speedy to the extent that subjects settle quickly into a systematic pattern of response to the variables. To measure learning speed and consistency, we introduce a rolling regression (or sequential least squares) technique inspired by Marcet and Sargent (1989a,b,c). The technique gives us trial-by-trial estimates of subjects' implicit coefficient values or responsiveness to the two exogenous variables. We deem learning to be consistent if these estimates converge by the last trial to the objective values, and say that there is under (or over) response if the absolute values of the coefficients are below (or above) the objective values. We measure learning speed mainly by comparing a subject's path of coefficient estimates to an ideal Bayesian (or Marcet-Sargent) path.

The OJF task is a continuous analogue of the discrete response Medical Diagnosis (MD) task studied intensively by psychologists such as Gluck and Bower (1988) and more recently by Kitzis et. al. (1998). The older psychological literature from Thorndike (1898) emphasizes reinforcement learning – actions that do well now are "reinforced" and chosen more frequently in the future. Naive reinforcement models do not extend naturally to our OJF task since it is not clear what reinforcement means in the context of continuous stimuli (weather and supply information) and continuous response (price forecast). The MD literature considers more sophisticated models of error-driven learning, including neural network or connectionist models and generalized discrete Bayesian models. The most striking finding of Kitzis et. al. (1998) is that a generalized Bayesian model outperforms alternative psychological models in the version of the MD task closest to the present OJF task. That paper also justifies reliance on least squares (as opposed to maximum likelihood) fitting techniques. The MD results encourage us to pursue rolling regression techniques in the OJF task.

There is also a related strand of experimental economics literature that examines rational expectations. Garner (1982) presents twelve subjects over 44 periods with a continuous forecasting task that implicitly requires the estimation of seven coefficients in a third order autoregressive linear stochastic model. He rejects stronger versions of rational expectations but finds some predictive power in weaker versions. Williams (1987) finds autocorrelated and adaptive forecast errors by traders in simple asset markets. Dwyer et al (1993) test subjects' forecasts of an exogenous random walk. They find excess forecast variance but no systematic positive or negative forecast bias.

Section 2 below describes our experiment. Section 3 presents the main results: (1) learning is fairly consistent in that most subjects' coefficient estimates converge closely to the objective values but there is a mild general tendency toward over-response. (2) Typically learning is noticably slower than the Marcet-Sargent ideal. Among the more striking treatment effects are a general tendency (3) towards over-response in the high noise treatment and (4) towards under-response in the asymmetric impact treatment.

Section 4 summarizes the results and discusses implications and extensions. Appendices A and B document the instructions to subjects and the identification of unresponsive subjects. Kelley and Friedman (1998) briefly summarize the recent MD results together with preliminary OJF results. Kelley (1998) reports additional OJF results, as described in section 4 below.

# 2 Laboratory Procedures

We induce the following linear stochastic relationship of price $p$ to contemporaneous values of two exogenous variables, $x_1$ and $x_2$:

$$p_t = a_1 x_{1t} + a_2 x_{2t} + e_t. \tag{1}$$

Subjects are told that $p$ refers to the local orange juice futures price relative to its normal level. They are also told that $x_1$ refers to the local weather hazard which could potentially destroy part of the domestic orange production, and that $x_2$ refers to the competing supply of oranges from Brazil. The realized price $p_t$ in trial $t$ depends on the realized value of $x_{1t} \in [0, 100]$ and its coefficient $a_1$ (approximately 0.4 in the baseline treatment), and on $x_{2t} \in [0, 100]$ and its coefficient $a_2$ (approximately -0.4 in the baseline treatment). The coefficient signs reflect the economic reality that loss of domestic crops tends to increase price and that increased foreign supply tends to decrease price. The noise term $e$ reflects the unpredictability of prices in field markets. Its value $e_t$ is drawn independently each trial from the uniform distribution on $[-v, v]$, where the (maximum) noise amplitude $v$ is a treatment variable (approximately 8 in the baseline treatment).

Subjects are instructed on the general nature of the task but are not specifically told the functional form or the coefficient values. Subjects are told that the experiment is a learning experience in which the goal is to learn the relationship between information (weather and competing supply) and the price of OJF. The instructions (attached as Appendix A) state in nontechnical language that the relationship is stable but subject to random events that are independent across trials. Treatments described in the next subsection are held constant for each subject and are varied across subjects.

*Subject Pool.* We have tested 99 undergraduates from the University of California at Santa Cruz, most of them from the pool of psychology students who need to fulfil a class requirement. Salient cash payments were offered in one treatment described below.

*Apparatus.* The experiment uses a graphics computer program written in C++, run on power Macintosh 7500/100 computers with full color monitors. Subjects in four sound dampened isolated testing rooms view controlled events on the monitor screen and respond via clicking the mouse on various icons on the display. See Figures 1.1-1.3 for examples of screen displays. This setup was chosen to minimize boredom and to eliminate the possibility of peer pressure.

*Stimuli.* The realized values for weather $x_{1t}$ and supply $x_{2t}$ are independently drawn each period from the uniform distribution on $[0, 100]$, so the variables are orthogonal. The noise term is independently drawn each period from a different uniform distribution, $U[-v, v]$. The realized values then are combined using equation (1) and chosen parameter values ($a_1, a_2$, and $v$) to produce a 480 trial sequence of prices. The same sequence of realized values and prices is used for all subjects in any given treatment condition.

*Method.* Each trial begins with the graphical presentation of the weather and supply values using two thermometer icons (labeled weather hazard and Brazilian supply) on the left side of the monitor display as in Figure 1.1 (A). Each thermometer is partially filled in red to indicate the realized value. Except in the no history treatment described below, the subject could also access (by clicking Previous Cases icon labelled (C) in Figure 1.1) the history of prices in previous trials with similar weather and supply levels, as in Figure 1.2 (D).[2]

[2]The history box (D) displays numerically the current realization of both variables, the number of previous trials for each variable whose realization is within 10 of the current

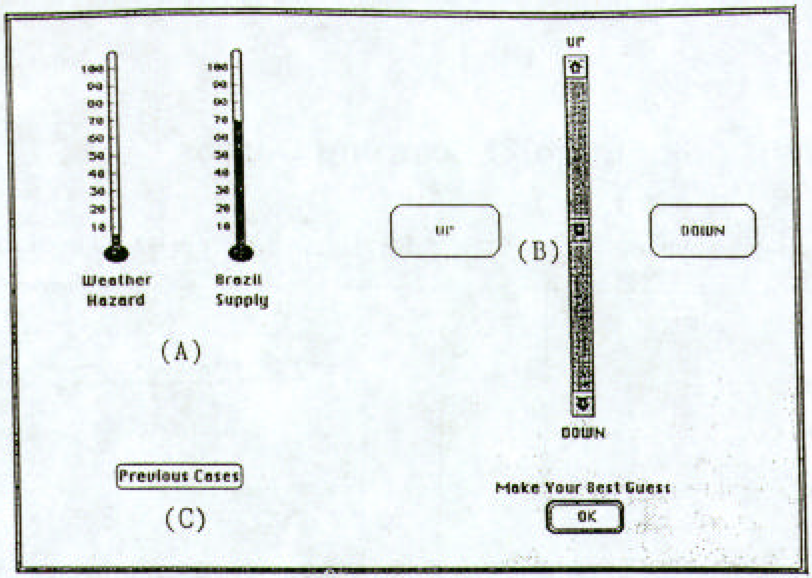


Figure 1.1

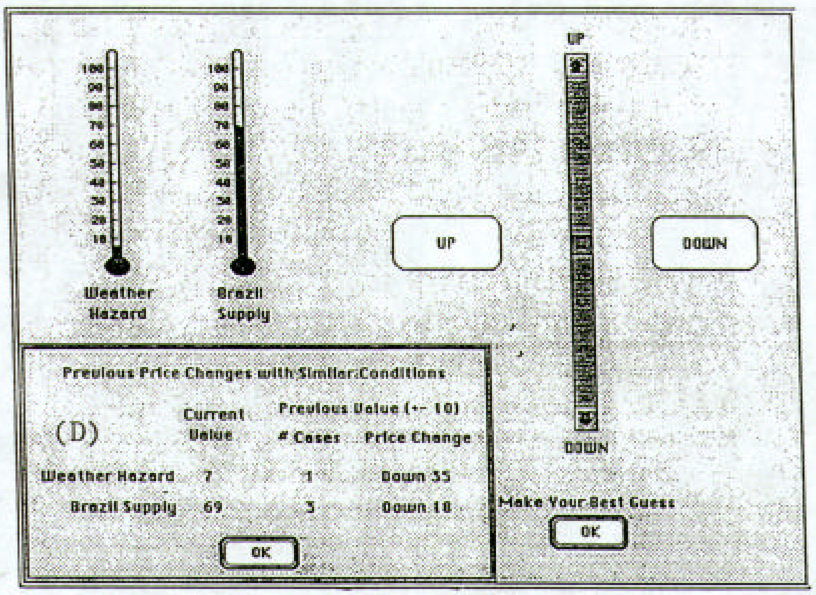


Figure 1.2

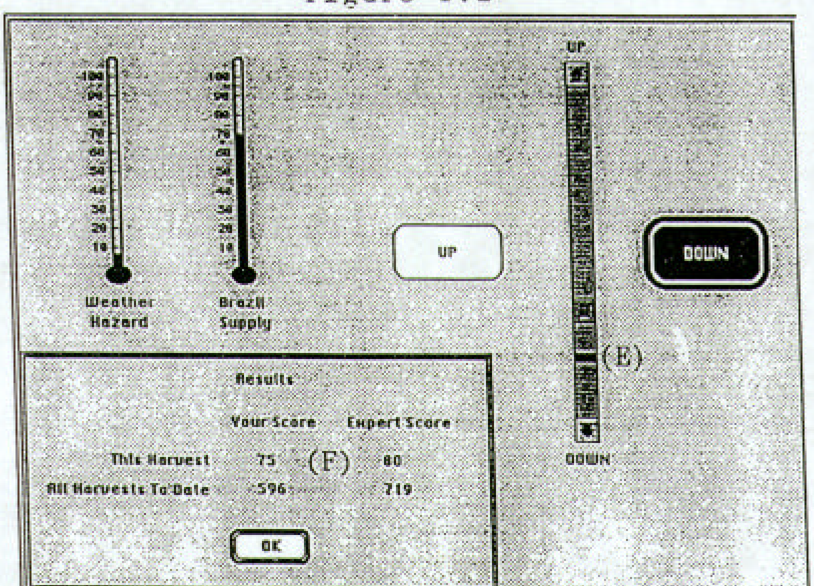


Figure 1.3

Subjects enter their forecast each period by moving slide (B) in Figure 1.1 up or down within the possible price range. After the price prediction is entered and confirmed, a blue line appears on the slide bar to indicate the actual price in that trial as in Figure 1.3 (E). Except in the no score treatment described below, the score box then appears as in Figure 1.3 (F).[3] After viewing the score box (if present) the subject advances to the next trial via a mouse click.

Each subject completes 480 self-paced trials. The session is broken into 3 blocks of 160 trials and subjects are permitted five minute breaks between blocks. Subjects generally finish in less than the allotted two hours.

## 2.1 Treatments.

*Baseline.* The baseline parameter values are $a_1 = 0.417$, $a_2 = -0.417$ and $v = 8.33$. The history and score boxes appear as described above. Each trial the score is calculated from the continuous price forecast $c$ and the realized price $p$ using the quadratic scoring rule $S(p, c) = A - B(p - c)^2$, with $A = 80$ and $B = 280$. Thus the maximum score (for a perfect forecast) is $A = 80$ points and the minimum is $-B = -200$ points. See Friedman and Massaro (1998) for a recent discussion of this scoring rule. The box also displays the "expert" score of a forecaster with nothing left to learn, i.e., the score earned by forecasting $a_1 x_{1t} + a_2 x_{2t}$ in trial $t$, using objective values of $a_i$. Subjects, of course, do not observe the expert forecast, just the expert score.

*Paid.* This treatment differs from baseline only in that subjects are paid according to their final scores. Each subject receives a $5.00 show up fee covering the first 30,000 points of final cumulative score. (Actual final scores always exceeded 30,000 with the top scores over 37,000.) Subjects also receive an additional dollar for each 700 points scored above 30,000. The median payment was about $15.00 with top payments about $16.50. Subjects are told the payment procedures on arrival.

*No Score.* This treatment differs from baseline only in that subjects do not have access to the Results or Score box icon and box (F).

realization, and the average realized price in those previous trials. The box remains on the screen until the subject clicks the OK icon.

[3]The score box (F) displays the subject's score on the current trial and the cumulative score through the current trial. The calculation is explained in section 2.1 below.

*No History.* This treatment differs from baseline only in that subjects do not have access to the Previous Cases or History icon and box (C) and (D).

*Asymmetric.* Here the coefficient values are $a_1 = 0.250$ and $a_2 = -0.583$. Thus the weather and the competing supply information no longer have equal (or symmetric) impact on OJF price.

*High Noise.* Here the noise amplitude is almost doubled, to $v = 14.3$, and the coefficient values $a_i$ are scaled to $\pm 0.357$, as described below. All other features are as in the baseline treatment.

## 2.2 Data Processing.

The baseline values of $a_i$ are scaled as follows. Begin with unscaled values $a_1^* = 0.5$ and $a_2^* = -0.5$. Given noise amplitude $v^*$, equation (1) implies that the unscaled price ranges from $p^* = 0.5(0) - 0.5(100) - v^* = -[50 + v^*]$ to $p^* = 0.5(100) - 0.5(0) + v^* = [50 + v^*]$. To fit in the screen's range $[-50, 50]$ we display the scaled price $p = \frac{50p^*}{[50+v^*]}$ . The scaled coefficients therefore are $a_i = \frac{50a_i^*}{[50+v^*]}$ and the scaled noise amplitude is $v = \frac{50v^*}{[50+v^*]}$. For the baseline noise value $v^* = 10$ we have $v = 8.33$ and $a_i = 0.833a_i^* = 0.417$. The scaled coefficients used in the high noise ($v^* = 20$) and asymmetric treatments are derived in a similar fashion.

For a given subsequence of trials $(p_t, x_{1t}, x_{2t})$, $t = t_0, ..., T$, we define the ideal Bayesian (or Least Squares or Marcet-Sargent) learner by regressing $p_t$ on the independent variables $x_{1t}$ and $x_{2t}$ via ordinary least squares (OLS). The regression over this subsequence of trials yields coefficient estimates $a_{1T}$ and $a_{2T}$. The subsequences we consider consist of trials 1 to 160, 2 to 161, ... ,320 to 480. Thus we obtain learning curves $a_{1T}$ and $a_{2T}$ for $T = 160, 161, ..., 480$, which can be interpreted as ideal subjective estimates of the objective values $a_1$ and $a_2$. We refer below to these as the M-S learning curves.

We use similar rolling regressions for human subjects. An actual subject may think of the task in various idiosyncratic ways — for example, he may believe that prices are serially correlated or that price is a nonlinear deterministic function of the exogenous variables, despite our instructions to the contrary. Nevertheless, the analyst can summarize the subject's beliefs by seeing how he responds to the current stimuli $x_{it}$, and can summarize the learning process by seeing how the subject's response changes with experience. Our approach therefore is to reconstruct implicit beliefs using equation (1) and subjects' actual responses.

The reconstruction proceeds as follows. Take the subject's actual forecast $c_t$ in trial $t$ as the dependent variable, and run rolling regressions as before on the realized values $x_{it}$, using a moving window of 160 consecutive trials with the last trial $T$ ranging from 160 to 480. Consistent and speedy learning is indicated by rapid convergence of the coefficient estimates $a_{iT}$ (as $T$ increases) to the objective values $a_i$. Obstacles to learning are suggested by slow convergence, convergence to some other value which represents over- or under-response, or divergence of the coefficient estimates.

Some details may be worth noting briefly. (1) In all the results reported below, the intercept coefficient $a_0$ is constrained to its objective value of zero. Excluding the intercept doesn't affect our main results but does it does reduce clutter and improve statistical efficiency. (2) In preliminary work we considered stretchable windows of data running from $t = 1$ to $T$, to capture fully the evidence available to the subject (or M-S ideal learner) in trial $T$. However, the entire learning curve then reflects the subject's initial response pattern as well as the recent response pattern. We concluded that learning curves would be more informative when estimated from a moving window that includes only the most recent responses. Of course, the recent responses already incorporate everything the subject has learned since the beginning of the session. (3) Lengthening a (non-stretchable) moving window reduces standard errors in the coefficient estimates, but also reduces the weight on the most recent responses. After a cursory investigation of preliminary data, we settled on length 160 as a reasonable compromise.

# 3 Results

Figure 2 presents a sample of learning curves in each treatment. Each panel of the Figure shows the objective coefficient values as a horizontal dotted line and shows the ideal M-S learning curves as thin continuous lines. The rolling regressions that generate the M-S curves seem to capture the price data quite well; typical $R^2$s ranged from 0.91 for the first 160 trial window of data to 0.93 for the last window. We were pleased to see that M-S learning is consistent and quite rapid, indeed virtually complete within the first 160 trials, as indicated by closeness of the dotted and thin continuous lines in every panel. The gap between the lines typically is about one standard error of the M-S coefficient estimate.

The heavy continuous lines in each panel of Figure 2 represent the learning curves for the highest scoring subject or the subject with the median score in each treatment. The corresponding rolling regressions again had typical $R^2$s above 0.90. The first two panels show moderate but persistent over-response to current weather and supply information, with implicit coefficient estimates lying closer to $\pm 0.45$ than to $\pm 0.42$ for both subjects in the baseline treatment. The next two panels suggest that the top scoring paid subject is right on target, but the median scorer tends to under-respond slightly. Over-response seems strongest with the top scoring subject in the no history treatment and the two subjects shown in the high noise treatment. The two subjects shown in the asymmetric treatment appear to under-respond in most trials.

To conserve space we do not show the learning curves for the other 87 subjects. Suffice it to say that subjects sometimes over-respond, sometimes under-respond, but typically are fairly close to the objective values. Subjects seem to update more slowly than the M-S ideal learner. The rest of this section will test these impressions more systematically.

## 3.1 Distribution of Scores

Figure 3 shows the distribution of the scores earned by subjects in each treatment. Forecasts often are quite good; in most treatments the highest score is close to 38,000, only a bit below the Marcet-Sargent ideal. The modal score and the median score usually are not very far behind. Mean scores are usually lower because the lowest scores are much lower, sometimes below 34,000. For comparison, we calculated scores in the baseline treatment for two sorts of zero intelligence agents or non-learners. An agent who always forecasted zero (the optimal uninformed forecast) would score 34,326 and an agent who always used last period's price as the forecast would earn 30,647.

Closer examination of the raw data raises questions about the motivation of the subjects with lowest scores. We found that these subjects generally stopped responding to the weather and Brazil supply information at some point during the session. Subjects who don't care about performance but seek only to finish quickly can do so by just clicking the "OK" icons in every trial, leaving the price forecast at the default value $c = 0$. We identified such behavior in 9 of the 99 subjects. Note that unthinking responses of $c = 0$ will bias the coefficient estimates towards 0, so it is potentially important to the data analysis to identify such "questionable" behavior. Appendix B lists the questionable subjects and the criteria used to identify them.

Do the treatments systematically affect performance? Figure 3 suggests that, compared to the baseline treatment, scores may be a bit higher for paid subjects and a bit lower in some of the other treatments. Standard Wilcoxon tests indicate significantly lower scores in the asymmetric (p-value = 0.002), and high noise (p=0.002) treatments, and no significant difference from baseline in the no score (p=0.71), and no history (p=0.56) treatments. The paid condition produced insignificantly higher scores than the baseline (p=0.15).

## 3.2 Distribution of Coefficient Estimates.

Figure 4 shows the final ($T = 480$) distribution of both coefficient estimates by treatment. Overall, the subjects seem to have it about right: the estimates center near the objective value and most of the estimates are not far away. Moreover, most of the outlying estimates are spurious under-responses from the 10 questionable subjects (denoted with asterisks (*)). The distributions seem tighter in the paid treatment than in the baseline, and perhaps a bit more dispersed in the last three treatments. More importantly, there may be a slight bias towards over-response in the high noise treatment and towards under-response in the asymmetric treatment.

Further analysis is required to explore these impressions. Table 1 classifies a final (T=480) coefficient estimate as objectively correct if its central 95% confidence interval contains the corresponding final value from the Marcet-Sargent simulation.[4] The estimate is classified as over- (or under-) response if the confidence interval lies entirely outside (or entirely within) the interval from zero to the (M-S) objective value. Overall, a plurality of estimates (71 of them) are classified as objectively correct, and there are about equal numbers of over-responses (59) and under-responses (48 plus 18 questionables).

[4]Note that this redefinition of the objective value uses the available sample information rather than unavailable population information to define the objective value. The original (population) definition differs by about .013, and would tend to shift the classifications very slightly towards over-response.

The main imbalances arise in the last two treatments. Under-response to the more important variable (labelled [Brazil] Supply in Figure 4) and over-response to the other variable are quite prevalent in the asymmetric treatment. In the high noise treatment, a majority of the non-questionable estimates for both coefficients are classified as over-response and none is classified as under-response.

The last column of the Table reports Wilcoxon p-values separately for each of the two coefficients in each treatment for the full sample (and in parentheses, for the reduced sample that excludes the 9 questionable subjects.) In three cases the tests reject (at the conventional p=0.05 level in the reduced sample) the null hypothesis that the estimates center at the objective value, in favor of the following one-sided alternatives. There is significant under-response to the Supply variable in the asymmetric treatment (p=0.00), and significant over-response to both variables in the high noise treatment (p=0.02,0.00). There is also marginally significant over-response to the Supply variable in the baseline treatment (p=0.08). The other cases of apparent under- and over-response do not produce significant results in this conservative test.

The impression of tighter distributions in the paid treatment is consistent with significant findings in other experiments (Smith and Walker, 1993) but turns out not to be significant in our data according to standard parametric and non parametric tests. Of course, the sample size is not large, and it may be worth noting that subjects with questionable motivation appeared in the baseline (and high noise and asymmetric) treatment but not in the paid treatment.

Table 1 also reports behavior observed halfway through the session, at T=240. Recall from Figure 2 the impression (confirmed in omitted figures for the other subjects) that modest but shrinking over-response is quite typical at this point. The Table shows that over-response at the halfway point indeed is somewhat more prevalent than at the end of the session, especially in the paid and high noise treatments.

### 3.3 Summary

Several general conclusions emerge from the data analysis. Our human subjects do not learn as fast as an ideal Bayesian (or Marcet-Sargent econometrician), but even at the half-way point (T=240) of the experiment, the coefficient estimates indicate that responses are not very far from the mark. The overall tendency is towards over-response to the current information $x_1$ and $x_2$, but this tendency almost disappears by the end (T=480) of the experiment. We conclude that learning in our experiment generally is reasonably rapid and very consistent.

The treatments have modest but detectable impacts. Compared to the baseline condition, fewer subjects in the paid treatment appear to have questionable motivation and the scores and coefficient estimates seem to have tighter distributions. Surprisingly, neither the no score treatment nor the no history treatment significantly impaired the the subjects' scores or accuracy of the estimated coefficients. The asymmetric treatment, however, significantly lowered scores and pushed subjects significantly towards under-response to the more important information and (insignificantly) towards over-response to the less important information. The high noise treatment had the strongest impact: lower scores and over-response to both information variables.

# 4 Discussion and Future Work.

Existing literature can easily give the impression that humans typically make very irrational choices in simple laboratory tasks; see Rabin (1998) for a recent thoughtful survey. In sharp contrast, our human subjects rather quickly learn highly rational behavior in a nontrivial forecasting task. What accounts for the divergent results?

In some ways our experiment makes it difficult for subjects to be rational. The task is challenging in that the target variable, price, is stochastic and contingent on two independent variables. Another challenging aspect of our experiment is that we used psychology pool subjects, unpaid in most treatments. Irrational behavior exhibited by such subjects in some tasks disappears when subjects drawn from other pools are offered salient payments (Friedman and Sunder, 1994). With the exception of 9 of 99 subjects whose motivation was questionable, our subjects behaved quite rationally.

But in other ways our experiment gives rationality its best shot. The basic task allows subjects to learn over a relatively long sequence of trials (480) in a stationary environment. Our laboratory setup encourages subjects to draw on relevant intuitions about price determination and avoids features that might suggest inappropriate heuristics. The visual interface encourages rapid and unbiased processing of information and feedback. If anything, the interface biases subjects towards under-response, since the default response is 0 and the subject must move the slide up or down from that point. The reduced sample used in some of the data analysis screened out the most egregious cases of default response, but perhaps some slight bias remains[5]. Arguably our setup is more representative of economically important field environments than the some of the setups used in laboratory studies that find irrational behavior.

The rational behavior is fairly robust. Performance was not significantly impaired in the no score and no history treatments, which eliminated useful feedback. Even in the asymmetric and high noise treatments, performance was still quite good. Kelley (1998) reports several additional robustness checks. Specifications designed to capture prior beliefs and non-linear responses detected some transient effects in many subjects, but for the most part these effects disappeared by the final trial. Tests allowing a non-zero intercept term ($a_0$) for the asymmetric weights treatment were also performed. The main effect of this additional parameter was to eliminate the marginally significant overresponse observed for the smaller stimuli. The underresponse observed for the larger stimuli remained significant. Responses remained fairly rational even in a treatment featuring a structural break.

An important extension of the work presented here, especially from the macroeconomics point of view, is to introduce self-referential price determination. Marcet and Sargent (1989abc) study several linear stochastic models where traders' expectations affect the actual price observed each period. They derive conditions on traders' learning processes (rolling regressions in essence) that ensure convergence of actual price to rational expectations equilibrium. It seems feasible to implement such economies in the laboratory and (given some stronger assumptions than needed in the present paper)

[5]The most questionable remaining subject is 009 in the no history treatment. He made very erratic choices until late in the session, spent no more time making choices than the screened subjects (about half as long as most remaining subjects), and earned almost as low a score as screened subjects. He was not screened out of the reduced sample because he entered mainly non-default responses, but his motivation is also questionable and his coefficient estimates indicate dramatic under-response. Indeed, the relevant test would indicate marginally significant over-response (to the second variable in the no history treatment, p=0.08) if this subject were screened out of the sample.

to extract estimates of subjects' learning processes. We conjecture that the empirical models introduced in the present paper will continue to do well in a more complex self-referential setting.

We see two main lessons in the present results. The discussion so far has emphasized the lesson that people can learn to make quite good forecasts. The other lesson is that some slight but systematic biases remain. In particular, even after 480 trials, subjects still tended to over-respond to news in the high noise environment. Slight individual biases might interact to produce economically important market biases (Akerlof and Yellen, 1985; Kelley, 1998). More theoretical and empirical work is needed to understand learning in self-referential, nonstationary environments.

# References

[1] Akerlof, G. and Yellen, J., 1985, Can Small Deviations from Rationality make Significant Differences to Economic Equilibria?, American Economic Review 75:4.

[2] Binmore, K., 1987, Modeling Rational Players, Part I, Economics and Philosophy 3, 179-214.

[3] Binmore, K., 1988, Modeling Rational Players, Part II, Economics and Philosophy 4, 9-55.

[4] Brown, G., 1951, Iterated Solution of Games by Fictitious Play, Activity Analysis of Production and Allocation (Wiley: New York).

[5] Camerer, Colin, 1998, Experiments on Game Theory, Draft manuscript, (Caltech Division of Humanities and Social Sciences).

[6] Cournot, A., 1838, Researches sur les Principes Mathematiques de la Theorie Richesses, English edition, N. Bacon, ed., 1897, Researches into the Mathematical Principles of the Theory of Wealth (MacMillan, NY).

[7] Dywer Jr., G., Williams, A., Battalio, R., and Mason, T., 1993, Tests of Rational Expectations in a Stark Setting, The Economic Journal 103, 586-601.

[8] Evans, George W. and Honkapohja, Seppo, 1997, Learning Dynamics, University of Oregon Economics working paper. To appear in J. Taylor and M. Woodford, eds., Handbook of Macroeconomics (Elsevier, NY).

[9] Friedman, D. and Massaro, D., 1998, Understanding Variability in Binary and Continuous Choice, Psychonomic Bulletin and Review, in press.

[10] Friedman, D. and Sunder, S., 1994, Experimental Methods: a Primer for Economists (Cambridge University Press).

[11] Fudenberg, D., and D. Kreps, 1988, A Theory of Learning, Experimentation, and Equilibrium in Games, mimeo (MIT).

[12] Garner, A., Experimental Evidence on the Rationality of Intuitive Forecasters, 1982, Research in Experimental Economics 2, 113-128 (JAI Press).

[13] Gluck, M. A. and Bower, G. H., 1988, From Conditioning to Category Learning: An Adaptive Network Model, Journal of Experimental Psychology: General 117, 225-244.

[14] Hull, C., 1943, Principles of Behavior, (New York Appleton-Century-Crofts).

[15] Kelley, H., 1998, Bounded Rationality in the Individual Choice Experiment, Unpublished Thesis, Economics Department, University of California Santa Cruz.

[16] Kelley, H. and Friedman, D., 1998, Learning to Forecast Rationally, Prepared for Charles Plott and Vernon Smith, eds., Handbook of Experimental Economics Results.

[17] Kitzis S., Kelley H., Berg E., Massaro D., Friedman D., 1998, Broadening the Tests of Learning Models, forthcoming Journal of Mathematical Psychology.

[18] Marcet, A. and Sargent, T., 1989a, Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models, Journal of Economic Theory 48, 337-68.

[19] Marcet, A. and Sargent, T., 1989b, Convergence of Least Squares Learning in Environments with Hidden State Variables and Private Information, Journal of Political Economy 97, 1306-22.

[20] Marcet, A. and Sargent, T., 1989c, Least Squares and the Dynamics of Hyperinflation, In W. Barnett, J. Geweke, and K. Shell, eds., Chaos, Complexity, and Sunspots (Cambridge University Press).

[21] Marimon, Ramon and Sunder, Shyam, 1993, Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence, Econometrica 61, 1073-1107.

[22] Rabin, M., 1988, Psychology and Economics, Journal of Economic Literature 34, 11-46.

[23] Roll, R., 1984, Orange Juice and Weather, American Economic Review 74, 861-880.

[24] Sargent, T., 1994, Bounded Rationality in Macroeconomics (Clarendon Press Oxford).

[25] Smith, V. and Walker, J., 1993, Monetary Rewards and Decision Cost in Experimental Economics, Economic Inquiry 31, 245-261.

[26] Thordike, E. L., 1898, Animal Intelligence: An experimental study of the associative processes in animals, Psychol. Monogr. 2.

[27] Van Huyck, J., Battalio R., and Rankin, F., 1997, On the Origin of Convention: Evidence From Coordination Games, Economic Journal 107 (442), 56-97.

[28] Williams, A. W., 1987, The Formation of Price Forecasts in Experimental Markets, Journal of Money Credit and Banking 19, 1-18.

# Appendix A: Instructions to Subjects

## ORANGE JUICE FUTURES EXPERIMENT
revised 5/98

GENERAL INFORMATION

In this experiment you will be asked to use information to make predictions. You will look at information on competing supply levels and on weather hazard and will predict orange juice futures prices. Orange juice price determination in this experiment is fictitious but basically similar to real life. Your job is similar to that of an investor who must use imperfect information to predict futures prices.

In this experiment, new information arrives each period (or harvest season) on (1) the weather hazard for the local orange crop and (2) the supply of oranges in the main competing region, Brazil (see label **A** at Figure 1.1). Each piece of information can take on a value from 0 to 100. A value of 0 for weather hazard means that there will be no loss of local production due to inclement weather and a value of 100 means likely massive damage to the local crop. Similarly, a value of 0 for supply means a very small Brazilian production and a value of 100 means the largest possible Brazilian crop.

Each period after viewing the information on weather and supply, you will enter your price prediction. Prices are measured within the range -100 (all the way 'DOWN', or 100 cents below the normal level) to +100 (all the way 'UP' or 100 cents above the normal level). For example, sliding the box (see Figure 1.1, **B**) to the topmost 'UP' position indicates that you believe that the current supply and weather conditions will result in a price 100 cents above the normal price. Likewise sliding the best guess box to the bottommost 'DOWN' position indicates that you believe the current crop conditions imply a price 100 cents below the normal price. Moving the box halfway up(halfway down) between the middle and top(bottom) predicts a price 50 cents price above(below) the normal level. Leaving the best guess box at its original position predicts exactly the normal price level.

READING CHARTS

Each period (or harvest season), you should first look at the information chart. You may be able to get useful additional information by clicking on the Previous Cases box. If it is present(see Figure 1.1, **C**) it will be under the chart symbols. When you click that box, a window will appear in the lower left corner of the screen (see Figure 1.2, **D**). The first column of the window lists the current information on competing supply and/or weather hazard. The second column lists the number of times so far in the experiment you have seen similar supply and weather conditions, i.e., within plus or minus 10. For example, in Figure 1.2 in all previous periods a weather hazard between 0 and 17 has occurred 1 time, and a supply between 59 and 79 has occurred 3 times. The third column gives the average price in these similar conditions. For example (see Figure 1.2, **D**), the current harvest's low weather hazard of (7) was associated with a price 35 cents below normal, and the somewhat high competing supply (69) was associated with a price 18 cents below normal. Click O.K. to leave the Previous Cases window.

After you have considered the relevant information, you enter your forecast by clicking the slide box and moving it to your chosen location on the ruler. After you have made your prediction the UP or DOWN box will be darkened if you predict a price different from the normal level, otherwise they will both remain light. Click on OK to submit your forecast. You will then be told the actual price that period. A blue bar will appear on the ruler to indicate the actual price (see Figure 1.3, **E**). You may then be given a numerical score for your prediction this harvest and a cumulative score for all harvests to date (see Figure 1.3, **F**). You will then get the information chart for the next period.

Your goal is to predict as accurately as possible each period. There will be many periods for you to predict. Work at your own pace. The whole experiment should take less than 2 hours. We ask that you do not take notes.

SCORING

Your score is the profit an investor makes when acting on your price prediction. Each harvest you earn points based on your prediction (between -100 and 100) and the actual price that harvest. Profit is higher the more accurate your forecast.(see Figure 1.3, **F**) For example if the actual price turns out to be 70 cents above normal, then your score is highest if your prediction was +70, a bit lower if your prediction was +60 or +80 and much lower if you predicted 0 or below.

USEFUL FACTS ABOUT PRICES IN THIS EXPERIMENT

You should not expect your forecasts to be exactly correct each period. The same supply and weather conditions can sometimes lead to a price increase and sometimes to a price decrease relative to the normal level. But if you properly use the average effects of weather and competing supply, your forecasts will usually be fairly accurate.

Each harvest period researchers collect available information about market conditions affecting orange juice. The information is distilled into the charts you see. The charts always record the available information correctly. The two pieces of information are independent in the sense that, for example, a high local weather hazard does not indicate a high or low Brazilian supply.

Each piece of information tends to be associated with higher or lower prices, but there is never certainty. An expert who completely understands the effects of competing supply and weather hazards typically earns much higher profits than a novice, but even the expert can't predict perfectly each period.

Feel free to ask the experimenter about anything in these instructions or in the experiment that is unclear to you.

# Appendix B: Identity of Questionable Subjects.

The reduced sample omits 9 of the 99 subjects. The omitted 9 usually earned the lowest scores in their particular treatment group. The criterion for omission was whether the subject actually responded to the stimuli, or always entered the default continuous response of 0 (corresponding to a normal price forecast, or no expected price change) in many consecutive trials. Here are the specifics.

| Subject # | Score | Treatment | Subject Characteristics |
|---|---|---|---|
| 10 | 32638.82 | Baseline | Virtually all responses are default ($c_t = 0$) for 50 to 200 consecutive trials. Second lowest score. |
| 20 | 33753.02 | High Noise | Virtually all responses are default ($c_t = 0$) for 50 to 200 consecutive trials. Lowest score in group. |
| 29 | 33071.24 | Asymmetric | Virtually all responses are default ($c_t = 0$) for 50 to 200 consecutive trials. Second lowest score. |
| 30 | 36294.49 | High Noise | Completely stopped responding early in experiment Sixth lowest score in group. |
| 34 | 31426.81 | Asymmetric | Virtually all responses are default ($c_t = 0$) for 50 to 200 consecutive trials. Lowest score. |
| 40 | 35851.27 | High Noise | Virtually all responses are default ($c_t = 0$) for 50 to 200 consecutive trials. Third lowest score. |
| 61 | 35558.64 | High Noise | Completely stopped responding. Over-response that moves to under-response. Second lowest score. |
| 74 | 32271.7 | Baseline | Virtually all responses are default ($c_t = 0$) for 50 to 200 consecutive trials. Lowest score. |
| 89 | 35965.6 | High Noise | Virtually all responses are default ($c_t = 0$) for 50 to 200 consecutive trials. Fourth lowest score. |

# Figure 2: Learning Curves

### 2.1.1 Top Scorer in Baseline treatment

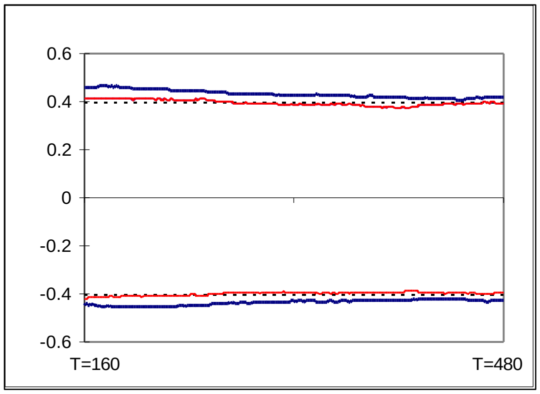


### 2.1.2 Median Scorer in Baseline treatment

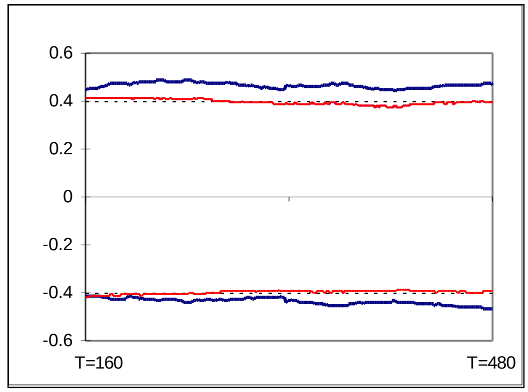


### 2.2.1 Top Scorer in Paid treatment

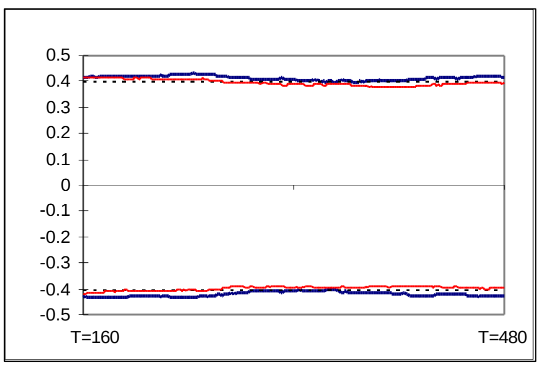


### 2.2.2 Median Scorer in Paid treatment

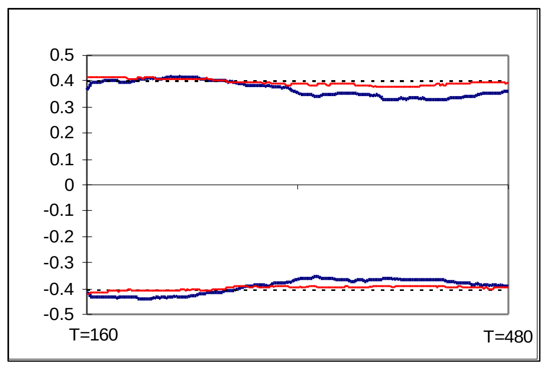


### 2.3.1 Top Scorer in No Score treatment

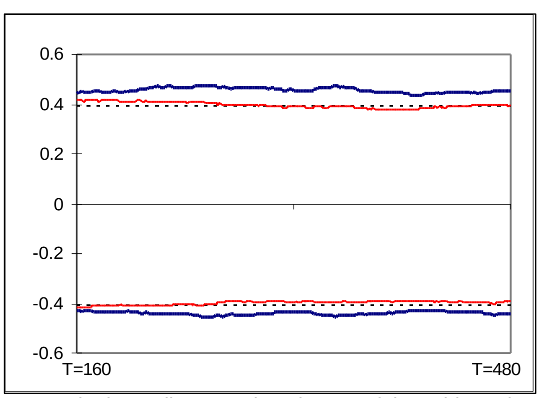


### 2.3.2 Median Scorer in No Score treatment

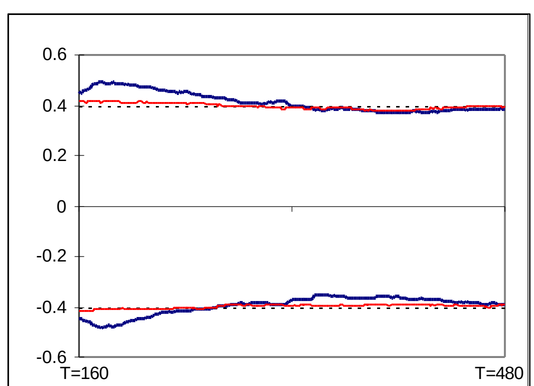


Note: The heavy lines graph estimates of the subject's implied coefficients $a_{1T}$ and $a_{2T}$; the width is approximately + - 1 standard error. The light lines graph the Marcet-Sargent (MS) ideal coefficients, and the dotted lines indicate the objective values a1 and a2.

**2.4.1 Top Scorer in No History treatment**

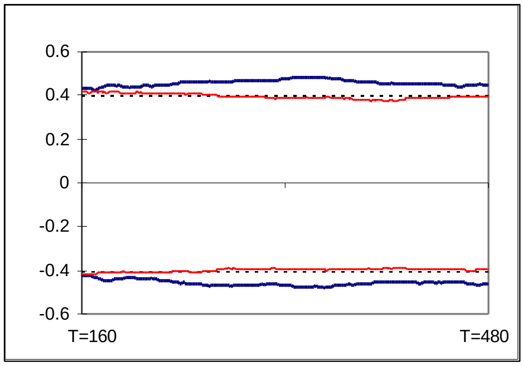


**2.4.2 Median Scorer in No History treatment**

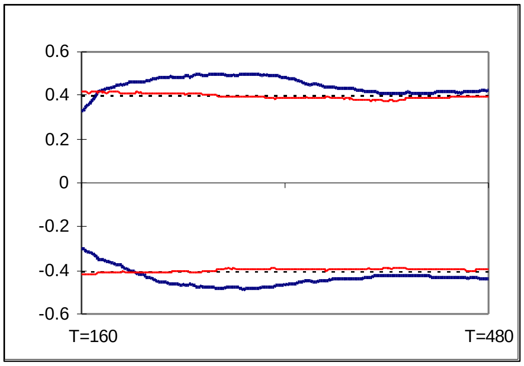


**2.5.1 Top Scorer in Asymmetric treatment**

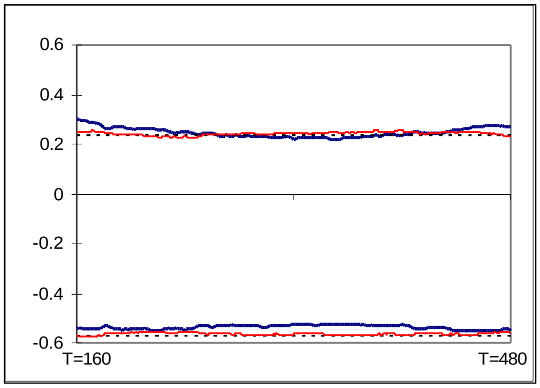


**2.5.2 Median Scorer in Asymmetric treatment**

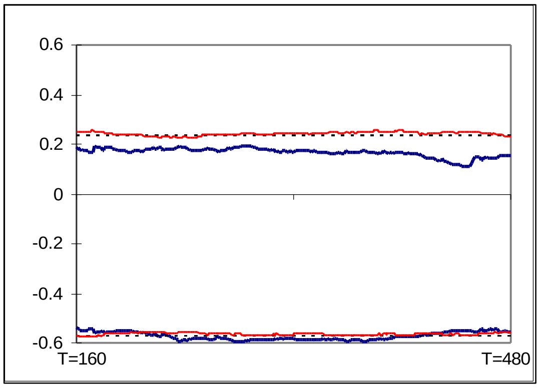


**2.6.1 Top Scorer in High Noise treatment**

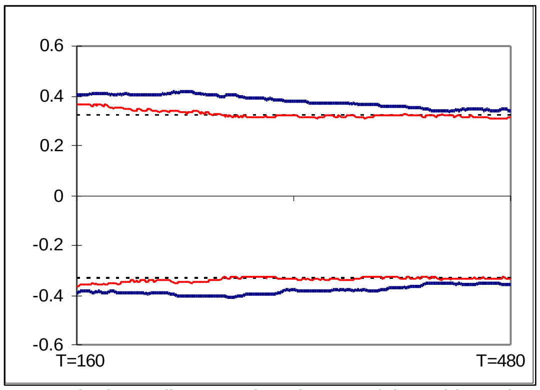


**2.6.2 Median Scorer in High Noise treatment**

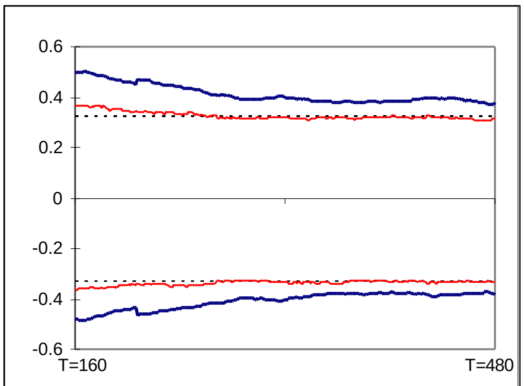


Note: The heavy lines graph estimates of the subject's implied coefficients $a_{1T}$ and $a_{2T}$; the width is approximately + - 1 standard error. The light lines graph the Marcet-Sargent (MS) ideal coefficients, and the dotted lines indicate the objective values a1 and a2.

# Figure 3: Distribution of Subjects' Scores

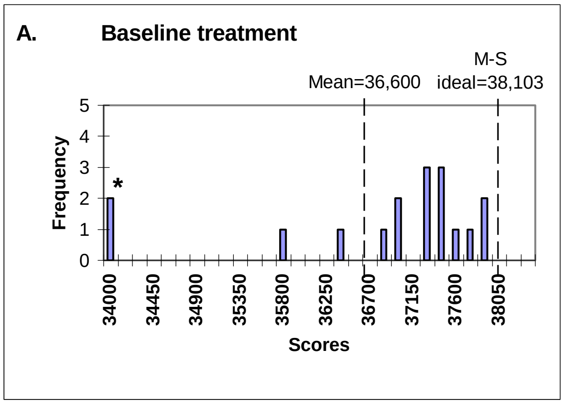


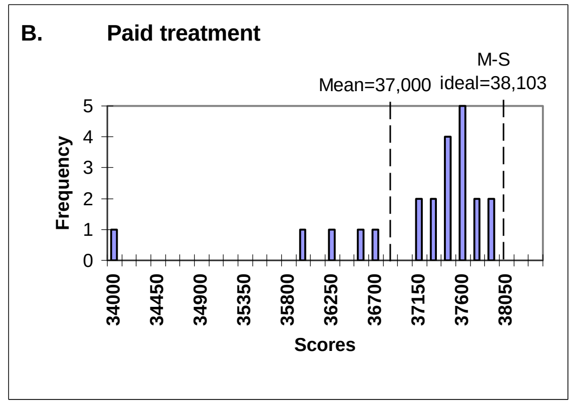


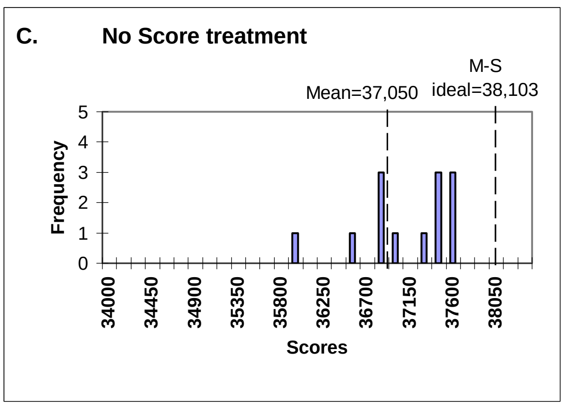


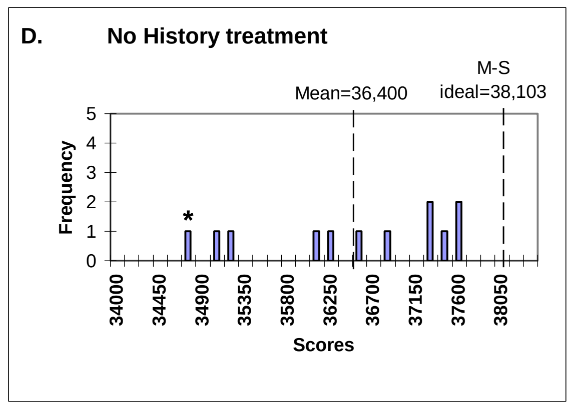


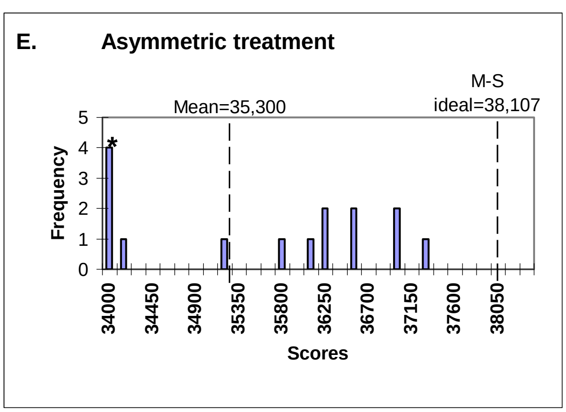


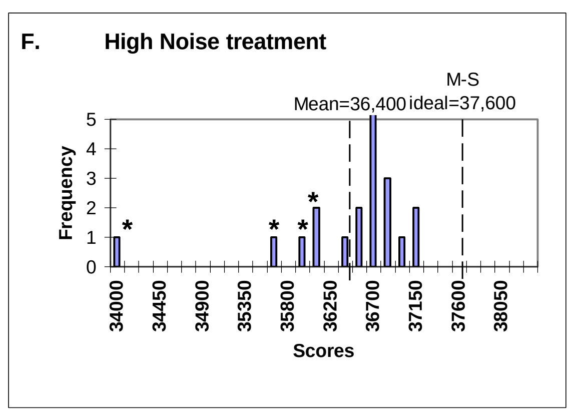

# Figure 4: Distribution of coefficient estimates at T = 480

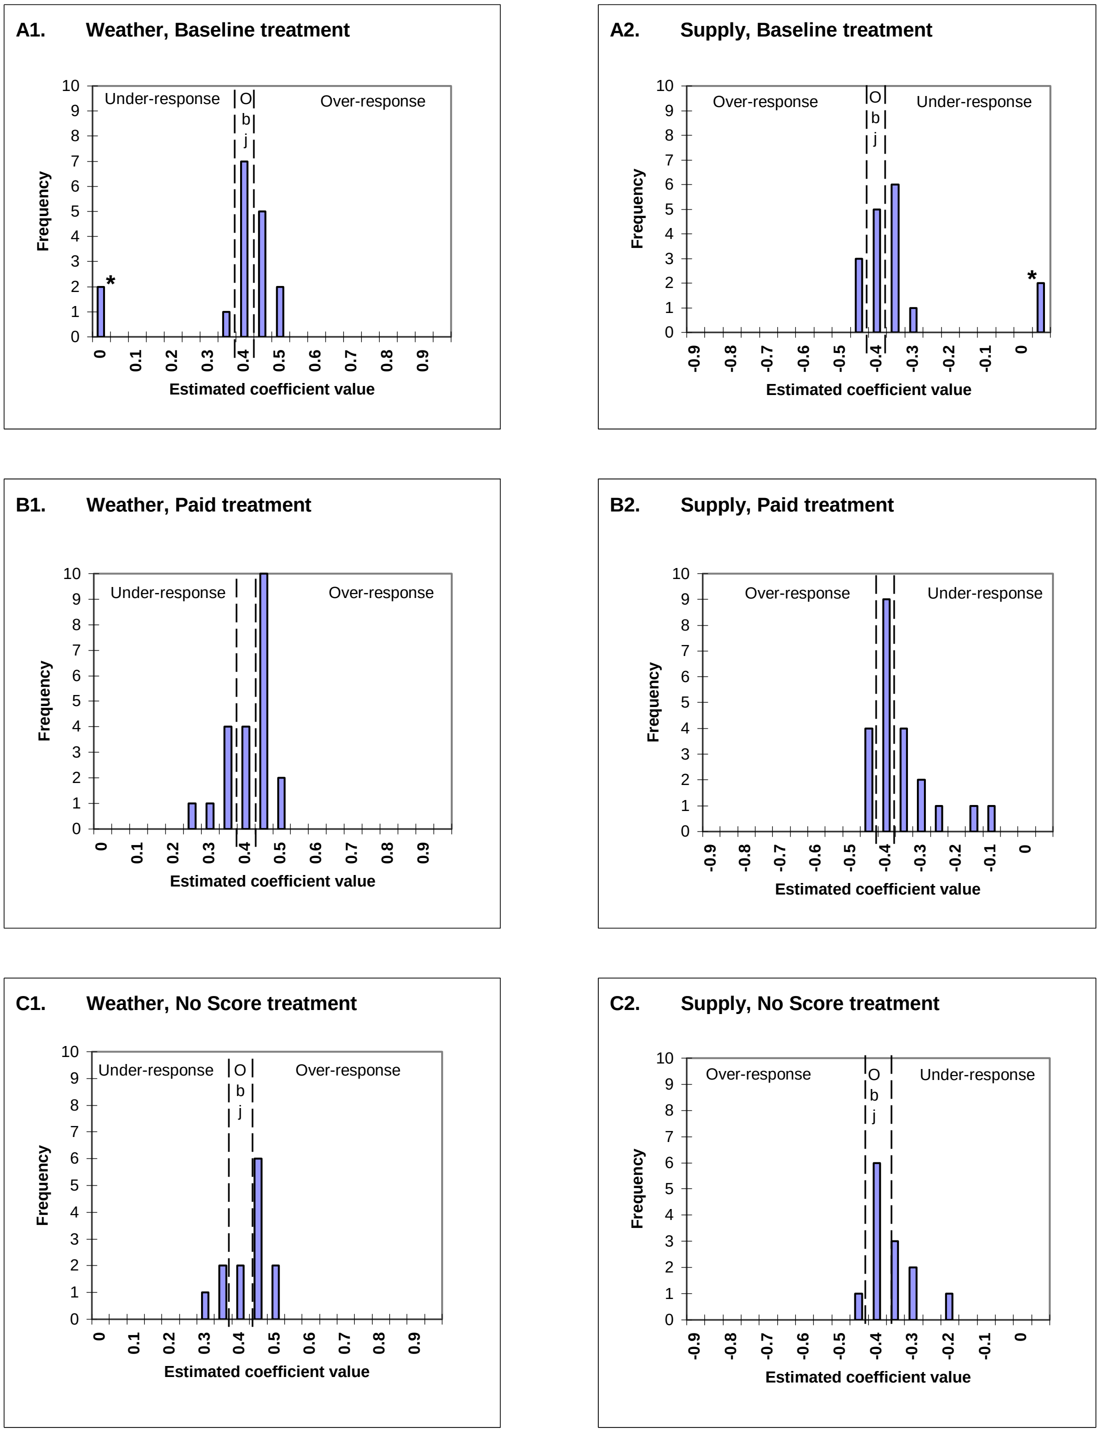


Note: The number of subjects are shown whose estimated coefficents $a_{iT}$ fall into the indicated ranges at T=480. "Obj" denotes the range in which the objective value $a_i$ falls, for i=1 (weather hazard) and i=2 (competing supply). Estimates smaller in absolute value are labelled "under-response" and larger estimates are labelled "over-response".

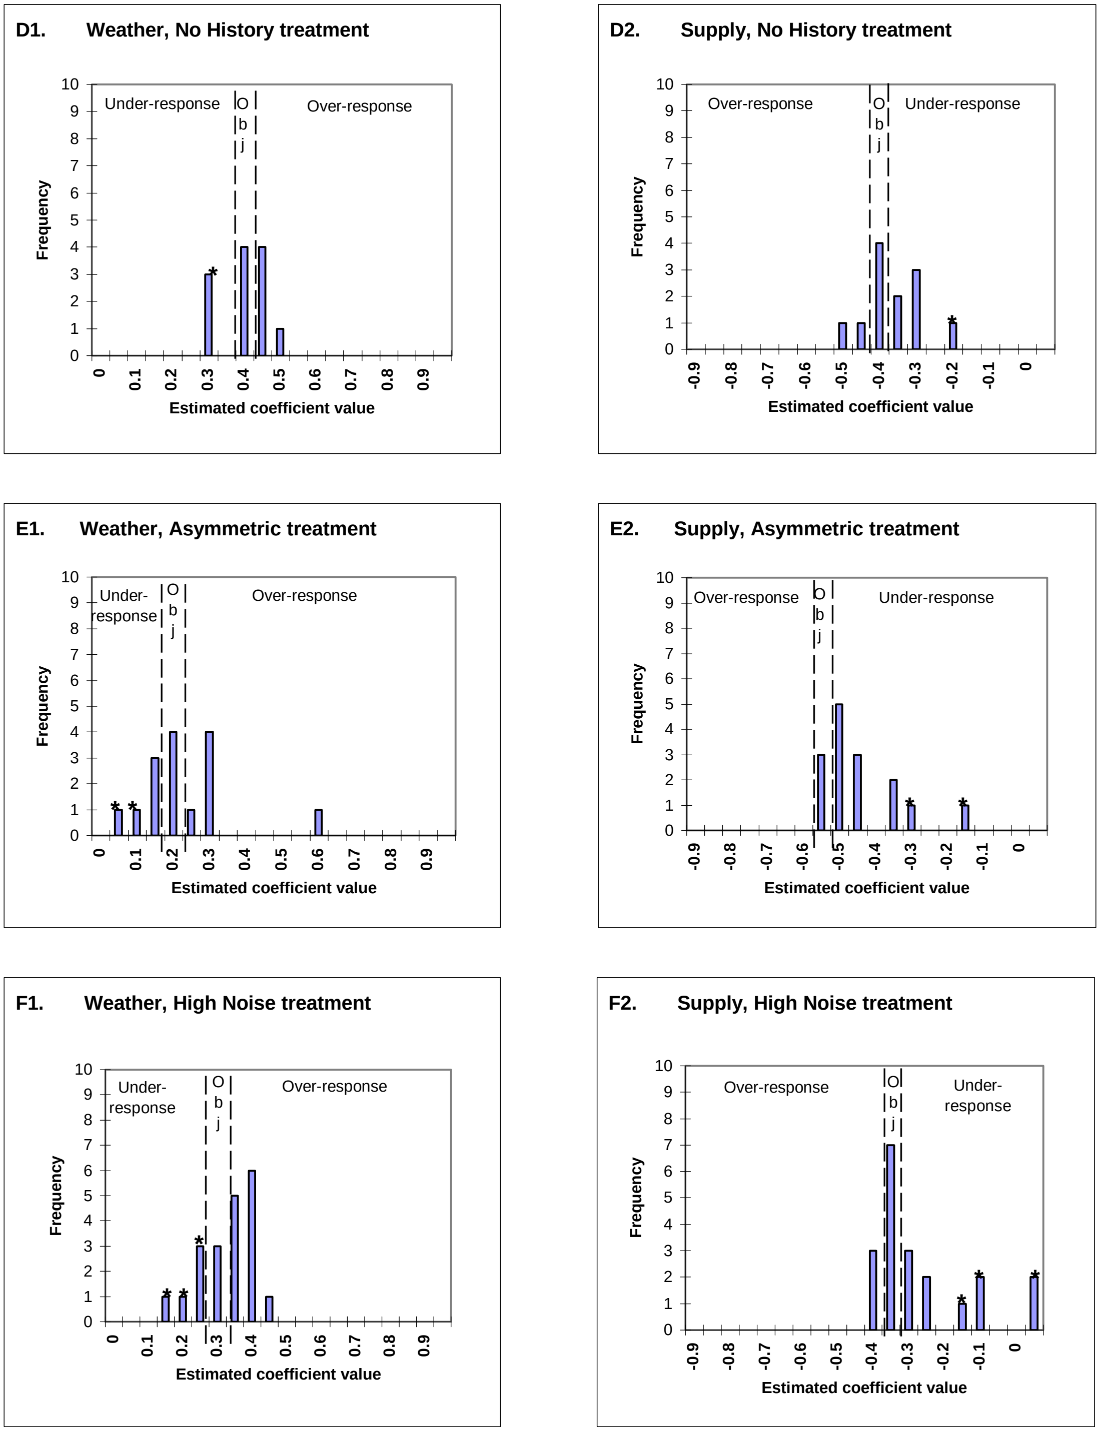


Note: The number of subjects are shown whose estimated coefficents $a_{iT}$ fall into the indicated ranges at T=480. "Obj" denotes the range in which the objective value $a_i$ falls, for i=1 (weather hazard) and i=2 (competing supply). Estimates smaller in absolute value are labelled "under-response" and larger estimates are labelled "over-response".