



Signal extraction: experimental evidence

Te Bao¹ · John Duffy² 

Accepted: 12 October 2020 / Published online: 26 October 2020
© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

We report on an experiment examining whether individuals can solve a simple signal extraction problem of the type found in models with imperfect information. In one treatment, subjects must form point predictions based on observing both public and private signals, while in another they receive the same information but must decide on the weight to attach to each signal, which then determines their point prediction. We find that, at the aggregate level, signal extraction provides a good characterization of subjects' behavior in both treatments, but at the individual level, there is considerable heterogeneity in subjects' ability to perform signal extraction.

Keywords Signal extraction model · Belief updating · Heterogeneous expectations · Bayesian learning

1 Introduction

The process of extracting signals corrupted by noise is known as the signal extraction problem. Signal extraction is a particular type of linear filtering, known as the Wiener–Kolmogorov filter, that is applicable to settings where the sources of noise follow stationary processes.¹

¹ This method was developed simultaneously by Wiener (1941) and Kolmogorov (1941) during the Second World War with the aim of targeting radar-assisted anti-aircraft guns on incoming enemy aircraft

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s11238-020-09785-x>) contains supplementary material, which is available to authorized users.

✉ John Duffy
duffy@uci.edu

Te Bao
baote@ntu.edu.sg

¹ Division of Economics, School of Social Sciences, Nanyang Technological University, Singapore, Singapore

² Department of Economics, University of California, Irvine, USA

In economics, signal extraction is an important micro building block in characterizing the equilibria of models of imperfect information. Perhaps the most famous application of signal extraction is found in Lucas's (1972) "island model" where firms that are located on different islands (and thus lacking perfect information about the prices of other firms) have to determine the extent to which an increase in the price of their own product is due to an increase in the demand for their product or to a rise in the general level of prices.²

While the Lucas example is perhaps the most well-known application of signal extraction, there are many other economic applications involving the use of signal extraction problems. Fang and Moro (2011) show that statistical discrimination amounts to solving a signal extraction problem. Wolfers (2002) asks whether voters can assess the record of politicians apart from factors the politicians are unable to affect, which amounts to a signal extraction problem. Gabaix and Laibson (2017) show that agents can exhibit as-if discounting behavior when they face imperfect information and estimate the value of future events by applying a signal extraction model. Specifically, when signals about far distant future events are noisier, agents assign those signals less weight in their decision problem, and that partially mimics the properties of classical time preferences.

While the signal extraction model has been widely used, particularly in macroeconomic modelling (see the seminal works by Lucas 1972; Mills 1982; Sargent 1991 and Wallace 1992), we are not aware of any prior empirical testing of this fundamental building block used in economic models with imperfect information. Thus, the first contribution of this paper is to provide some evidence on the extent to which humans can solve signal extraction problems. In the experimental game theory literature, some researchers do study decision making by subjects who face a signal extraction problem, e.g., between two noisy signals, as in the global game experiments of Heinemann et al. (2004), Cornand and Heinemann (2008) and the generalized beauty contest game experiment of Cornand and Heinemann (2014). In that setting, subjects have to use one or more noisy signals to form expectations/beliefs about a payoff relevant state variable. However, those studies do not directly examine how well subjects are able to solve the signal extraction problem that they face; rather, they elicit a discrete *action choice* from subjects, who are players in an n -player game, (e.g., whether or not to attack a currency), which is theorized to depend on subjects' estimate of the state variable. Thus, they do not directly elicit subjects' beliefs or expectations regarding the state variable, or the decision weights they assign to each signal explicitly. Our experiment complements this type of research by providing a way to discern the effect of factors that influence the belief/expectation of payoff relevant state

Footnote 1 continued

(Pollock 2013). This was the first statistical filter and predecessor to many others, including the Kalman filter, which is relevant for filtering noise arising from non-stationary processes.

² In macroeconomics, the use of signal extraction to model imperfect information gave way to more tractable timing assumptions. As Sims (2003) observes, "much of subsequent rational expectations macroeconomic modeling has relied on the more tractable device (relative to signal extraction) of assuming an 'information delay,' so that some kinds of aggregate data are observable to some agents only with a delay, though without error after the delay."

variables prior to the choice of actions. In the experimental macroeconomics literature, Menzies and Zizzo (2012) study how subjects' exchange rate expectations adjust to monetary shocks, in a model where the rational expectations equilibrium relies on signal extraction to solve for the correlation coefficient between the exchange rate and the interest rate. The subjects in their experiment, however, submit bid/ask prices for the exchange rate and are not asked to infer the correlation coefficient directly. Therefore, the Menzies–Zizzo study is also not a direct test of the signal extraction model. We note that the signal extraction model can be considered a sub-case of Bayesian updating models. Previous experimental studies of Bayesian updating (e.g., Benjamin 2019; Massey and Wu 2005; Henckel et al. 2017) usually consider Bayesian updating only under a Bernoulli distribution. Signal extraction is Bayesian updating under a Gaussian distribution, and people may process information and make choices in a systematically different way when the signals are drawn from different distributions. While the Bernoulli distribution is often used in individual decision-making experiments, (e.g., where the aim is to guess which of two urns signals are being drawn from), the Gaussian distribution is more often used in macroeconomic and finance models (e.g. where the aim is to forecast prices or exchange rates).

A second contribution of this paper is to compare direct elicitation of subjects' beliefs/expectations in the face of a signal extraction problem with a setting where subjects instead decide how to weight information from different sources. Subjects' weighting choice determines their point prediction for the state variable. Thus, our experiment also makes a methodological contribution to the literature on experimental design for experiments examining how subjects process information from different sources.

To preview our results, we find that at the *aggregate* level, the predictions of the signal extraction model are a good fit to the experimental data, but that there is substantial heterogeneity at the individual level. We further find that the signal extraction model provides a better characterization of individuals' belief updating process when subjects are asked to weight their different information sources as opposed to the case where they are asked to make a point prediction forecast.

2 The signal extraction task

Suppose the state, θ , is a random variable. Agents desire to know the realization of this random state variable in each period t , θ_t , as it is payoff relevant, e.g., it represents the general price level. The distribution of this random variable (signal) is assumed to be known. For example, in our experiment we shall assume that $\theta_t \sim N(y, \frac{1}{\alpha})$. However, signals of the state variable are only observed with some noise. Specifically, each agent i only observes a noisy private signal of θ_t , s_t^i , e.g., the price of the good being sold on i 's own island. This signal distribution is known as well. In our experiment, we assume that $s_t^i \sim N(\theta_t, \frac{1}{\beta})$. Further, for simplicity, we assume that there is no covariance between the two signals.

The signal extraction problem for agent i is to find $E(\theta_t | s_t^i)$. Application of the Wiener–Kolmogorov filter yields the optimal solution:

$$E(\theta_t | s_t^i) = \frac{\alpha y + \beta s_t^i}{\alpha + \beta}$$

see, e.g., DeGroot (2004). Based on this theory, agent i 's expectation, conditional on the private signal, s_t^i , is a weighted average of the two signals, with the inverse of the variance of each signal serving as the weights. A further implication, given our assumption of normally distributed random variables, is that the posterior expectation lies between the prior expectation, y , and the noisy signal s_t^i , that is, there is *updating toward the signal*; see Chambers and Healy (2012).

We are interested in the extent to which individual agents can perform this optimal filtering task. We are further interested in whether the framing of the problem as a prediction task or as a weighting task matters for the accuracy of individual forecasts.

3 Experimental design

To investigate these questions, we design an individual-decision making experiment where subjects are repeatedly confronted with a signal extraction task and incentivized to correctly guess the random state variable given only a noisy signal of this variable. Our experiment consists of two treatments.

In *treatment P*, subjects *predict* the realization of θ_t in each of 15 periods, given their signal, s_t^i . That is, we directly elicit their “prediction” or expectation, $E(\theta_t^i | s_t^i)$.

In *treatment W*, subjects decide on a *weight*, $w_t^i \in [0, 1]$, to assign to their private signal, s_t^i , in each of 15 periods. They choose the weight by moving a slider bar in the experimental interface. Their implied price forecast is calculated as the weighted average of their private signal, s_t^i , and the public signal, y , the known mean of the distribution for θ . Thus, after choosing a weight, we calculate their prediction according to $E(\theta_t^i | w_t^i) = w_t^i s_t^i + (1 - w_t^i)y$, and this procedure is known to them. Subjects can view their implied prediction for θ_t in real time as they move the slider bar to select their weight. They can thus experiment with different weighting possibilities before finally settling on a choice they like, at which point they simply click “Submit” to enter that weighting choice. We chose this slider design for the *W* treatment since subjects are rewarded for the accuracy of the point prediction implied by their weight choice, and we wanted to make this difference clear to subjects.

Indeed, in both treatments, subjects are paid according to the accuracy of their forecasts for θ_t . Specifically, subject i 's payoff in experimental currency units (ECUs) in period t of either treatment is given by:

$$\pi_t^i = \frac{100}{1 + |E(\theta_t^i | \omega) - \theta_t^i|}$$

where $\omega = s_t^i$ or w_t^i depending on the treatment (P or W).

In both treatments, the distribution for the state variable is perfectly known. Specifically, we set $y = 10$ and $\alpha = 1$, so that $\theta \sim N(10, 1)$. The private signal distribution is also perfectly known: $s_t^i \sim N(\theta_t^i, \frac{1}{\beta_t})$. The variance of the private signal $\frac{1}{\beta_t}$ is the same for all subjects, and took on one of three values in different periods of the experiment:

$$\beta_t = \begin{cases} 1, & t \in [1, 5]. \\ 4, & t \in [6, 10]. \\ 0.25, & t \in [11, 15]. \end{cases}$$

That is, the variance of the private signal is 1 for periods 1–5, 0.25 for periods 6–10, and 4 for periods 11–15, and these values were perfectly known to subjects. Thus, the private signal is equally noisy as the public signal in the first 5 periods, less noisy than the public signal in the second 5 periods, and noisier than the public signal in the last 5 periods. The realized value of the state variable and the private signal, θ_t^i and s_t^i , are different for each individual, i , in each period. All the values were randomly generated before the experiment.

The sessions were run at the Laboratory of Experimental Economics at Nanyang Technological University (NTU), Singapore. 96 subjects (48 for treatment P and 48 for treatment W) were recruited via the ORSEE system. The subjects are undergraduate students of NTU. The session subjects were recruited for consisted of two parts, of which the signal extraction task reported on in this paper was the *first* part of the session.³ Payment for the first part (the 15 period signal extraction experiment) was based on Subjects' ECU earnings from one randomly chosen period among all 15 periods they played. Each subject's ECU earnings from the randomly chosen period was converted into Singapore Dollars at the known exchange rate of 100 ECU = 10 Singapore Dollars. The average payoff earned by subjects was 7 Singapore dollars for this signal extraction experiment and subjects also received a 3 Singapore dollars show-up payment. Subjects were given written instructions on the task they would face in each of the 15 periods (see the Electronic Supplementary Appendix for copies). These instructions made it clear to subjects what is meant by a random variable drawn from a normal distribution (as used in the experiment). To insure that subjects had a good understanding of the tasks they faced, after the instructions were read, we asked subjects some control questions to check their comprehension of the instructions. Subjects could only start the experiment if they had answered all of the control questions correctly.

4 Results

We divide our discussion of the results into aggregate level results, focusing on the mean or median behavior of all subjects of a treatment, and individual-level behavior.

³ The second part of the session was an asset pricing experiment which we do not report on in this paper, but for which subjects could earn additional money payments.

4.1 Deviations from theoretical predictions

We first consider the deviation of subjects' expectations from the theoretical predictions given by the signal extraction model, θ_t^i in each of our two treatments. Figure 1 plots the mean squared error in each period, t , $MSE_t = \frac{\sum_i (E(\theta_t^i | w_t^i) - \theta_t^i)^2}{N}$, (where N is the total number of subjects in each treatment) of the signal extraction model in treatment P or W. We did not elicit point predictions directly in treatment W, but for each weight that a subject i assigned to the private signal in treatment W, $w_t(i, t)$, the implied point prediction for θ_t^i is

$$E(\theta_t^i | w_t^i) = (1 - w_t^i)10 + w_t^i s_t^i.$$

In general, as Fig. 1 reveals, this aggregate MSE is not very large in either treatment, though it is found to be greater in treatment P as compared with treatment W.

We next consider individual MSEs in the two treatments. The MSE for individual i is given by $MSE^i = \frac{\sum_t (E(\theta_t^i | w_t^i) - \theta_t^i)^2}{15}$, that is, individual i 's average MSE over all 15 periods. Figure 2 shows the empirical cumulative distribution function (CDF) of these individual MSEs from the signal extraction model (where each subject i is an observation). We see clear evidence for first order stochastic dominance: the MSE for the P treatment stochastically dominates that of the W treatment. Indeed, the two distributions are significantly different according to a Kolmogorov–Smirnov test ($D = 0.4375, p = 0.000$).




Fig. 1 The MSE of the signal extraction model over the 15 periods of treatments P and W




Fig. 2 The empirical CDF of the MSE of the signal extraction model in treatments P and W

Still, the difference between the subjects' predictions and the prediction of the signal extraction model are very small in both the P and W treatments. The results of a t test on the individual level data (each subject is an independent observation, and the number of observations is 48) suggests that the average difference between the subjects' predictions and the predictions of the signal extraction model is not significantly different from 0 at the 5% level in both treatment P ($t = -0.1462, p = 0.8858$) and treatment W ($t = -0.712, p = 0.4795$).

For treatment P, the MSE of the signal extraction model is 1.4808 for the whole experiment, 1.0531 in the first 5 periods, 1.0762 in the second 5 treatments and 2.3131 in the final 5 periods. For treatment W, the MSE is 0.2604 for the whole experiment, 0.1313 in the first 5 periods, 0.0781 in the second 5 periods and 0.5719 in the final 5 periods.

The MSE in the final 5 periods, when the variance in the private signal is the highest, appears to be larger than the other two blocks of 5 periods in both treatments. In treatment P, the differences between the MSEs of the first and second 5 periods ($t = -0.0265, p = 0.9789$) and second and third 5 periods ($t = -1.2392, p = 0.2214$) are not significantly different from 0, while there is a significant difference between the MSEs in the first and third 5 periods ($t = -2.9860, p = 0.0045$). In treatment W, the difference between the MSE in the first and second 5 periods is not significant at the 5% level ($t = 1.8526, p = 0.0702$), while the difference is significant at 5% level between

the first and third 5 periods ($t = -3.8715, p = 0.0003$) and between the second and third 5 periods ($t = -4.3491, p = 0.0001$).

The above findings can be summarized as Result 1:

Result 1: *The signal extraction model provides a good description of subjects' predictions at the aggregate level. The difference between the average predictions of subjects is not significantly different from the theoretical predictions of the signal extraction model, and the average difference tends to be small (usually less than 10% of the true realized value of the variable). Deviations tend to be greater when subjects make point predictions (P) versus a weighting decision (W), and when the variance of the private signal is larger.*

4.2 Weight allocated to the private signal

For treatment W, we directly elicited the weight that subjects assigned to the private signal. This weight is not directly observable in treatment P, but we can calculate the implied weight for treatment P using the following equation:

$$w_t^i = \frac{E(\theta_t^i | w_t^i) - 10}{s_{i,t} - 10},$$

since $E(\theta_t^i | w_t^i) = (1 - w_t^i)10 + w_t^i s_t^i$.

Using the directly elicited or implied weights, Fig. 3 plots the median weight assigned to the private signal by subjects in treatments P and W in each period, along with the signal extraction model prediction.

Table 1 reports the average and median elicited or implied weight assigned to the private signal by subjects in all periods of the experiment and in each of the three 5-period subperiods.

In treatment P, the average implied weight is a very noisy measure; it lies outside of the interval [0, 1] for the whole sample and for the first and last 5 periods.⁴ Therefore, we focus our analysis on the median implied weight.

For the P treatment, the median implied weight assigned to the private signal is 0.5543 for all periods (not significantly different from 0.5, $z = 0.9270, p = 0.3587$); 0.5088 in the first 5 periods (not significantly different from 0.5, $z = 0.9679, p = 0.3381$), 0.8440 in the second 5 periods (not significantly different from 0.8, $z = -0.5582, p = 0.5794$) and 0.3100 in the third 5 periods (not significantly different from 0.2, $z = 1.5778, p = 0.1213$).

For the W treatment, the average elicited weight is 0.5225 for all periods (not significantly different from 0.5, $t = 1.3037, p = 0.1987$), 0.5399 in the first 5

⁴ We also considered only focusing on subjects whose implied weight choice was “rational” in the sense that it lied between 0 and 1, i.e., it satisfied the definition of updating towards the signal according to Chambers and Healy (2012), but it turns out that among the 720 decisions in treatment P, 131 (18.19%) decisions are associated with an implied weight $w_t^i < 0$, and 187 (25.97%) are associated with an implied weight $w_t^i > 1$. Among all the subjects, only 5 (10.42%) subjects made 0 implied weight decisions lying outside the [0, 1] interval over all 15 periods, and 32 subjects (66.67%) made irrational implied weight decisions more than 5 times. On average, the number of implied weight decisions that satisfy the definition of updating toward the signal is only 8.4, less than 60% of the total decisions. This result makes it difficult for us to simply remove the irrational decisions or subjects from the sample.




Fig. 3 The median elicited or implied weight assigned to the private signal in treatments P and W as compared with the prediction of the signal extraction model

periods (not significantly different from 0.5, $t = 1.9338, p = 0.0592$), 0.6433 in the second 5 periods (significantly different from 0.8, $t = -5.8873, p = 0.0000$) and 0.3843 in the final 5 periods (significantly different from 0.2, $t = -17.3405, p = 0.0000$). The *median* elicited weight assigned to the private signal is 0.5206 for all periods (not significantly different from 0.5, $z = 0.923, p = 0.3560$), 0.5318 in the first 5 periods (not significantly different from 0.5, $z = 1.744, p = 0.0812$), 0.6878 in the second 5 periods (significantly different from 0.8, $z = -4.831, p = 0.0000$) and 0.3422 in the final 5 periods (significantly different from 0.2, $z = 5.4260, p = 0.0000$). Differently from the P treatment, where subjects make a point prediction, subjects in the W treatment do not seem to have a consistent tendency to overweight the private signal. Instead, they appear to overweight the private signal when the theoretical prediction calls for a low weight, and they underweight the private signal when the theoretical prediction calls for a high weight. In other words, they seem to overweight/underweight predicted low/high weights to avoid making extreme predictions.⁵

We can summarize the above findings as Result 2:

Result 2: *The signal extraction model generally provides a good description of subjects' weighting of private versus public signals. In treatment P, the average implied weight assigned to the private signal often deviates from signal extraction*

⁵ Such behavior is reminiscent of the nonlinear weighting of probabilities that is a behavioral primitive of Prospect theory.

Table 1 The average and median implied or elicited weight assigned to the private signal

Treatment P			
Period	Average	Median	Theoretical prediction
All periods	1.458	0.554	0.5
First 5	3.07	0.509	0.5
Second 5	0.173	0.844	0.8
Third 5	1.132	0.31	0.2
Treatment W			
Period	Average weight	Median weight	Signal extraction model
All periods	0.5225	0.5206	0.5
First 5	0.5399	0.5318	0.5
Second 5	0.6433	0.6878	0.8
Third 5	0.3843	0.3422	0.2

model substantially, i.e., it lies outside of the [0, 1] interval. Yet, there is no significant difference between the median implied weights and the theoretical predictions of the signal extraction model. Using the average implied weight, there appears to be overweighting of the private signal in this treatment. In treatment W, the signal extraction model provides an unbiased prediction for both the average and the median weights chosen in the population. Finally, for both treatments, the deviation from the signal extraction model is greatest in the final 5 periods when the private signal is noisier than in the other two blocks of 5 periods. Subjects' behavior seems to be better described by the signal extraction model in treatment W than in treatment P.

4.3 Individual behavior over time

In this section we focus on individual behavior over time. We consider a simple linear panel data regression model for subjects' predictions in each period t :

$$E(\theta_t^i) = c_1y + c_2s_t^i + \epsilon_t^i, \tag{1}$$

where $y = 10$ is the public signal, and s_t^i is the private signal. We add a fixed effect for each individual subject in the regression to rule out the influence of idiosyncratic factors at the subject level. Notice that this panel data regression provides an alternative means of estimating the weight placed on the private signal in the P treatment, which here amounts to the estimate of the coefficient, c_2 .

The first 5 columns of Table 2 report results from the panel data regression estimation of (1) for each treatment, P or W, using the whole sample and for each block of 5 periods. The coefficient, c_2 , on the private signal in the current period, s_t^i , is significantly different from zero at the 1% level for all samples and subsamples.

Table 2 Fixed effects regression results for the behavioral prediction rule (1) $E(\theta_t^i) = c_1y + c_2s_t^i + \epsilon_t^i$ and (2) $E(\theta_t^i) = c_1y + c_2s_t^i + c_3(\theta_{t-1}^i - E(\theta_{t-1}^i)) + \epsilon_t^i$

Treatment P	All Pds (1)	First 5 (1)	Second 5 (1)	Third 5 (1)	All Pds (2)	First 5 (2)	Second 5 (2)	Third 5 (2)
y	0.415*** (0.0261)	0.363*** (0.0502)	0.143*** (0.0608)	0.476*** (0.0430)	0.436*** (0.028)	0.353*** (0.0540)	-0.298 (0.065)	0.502*** (0.0430)
s_t^i	0.590*** (0.0259)	0.645*** (0.0498)	0.979*** (0.0600)	0.532*** (0.0432)	0.567*** (0.0278)	0.658*** (0.0528)	1.017*** (0.0640)	0.500*** (0.0436)
$\theta_{t-1}^i - E(\theta_{t-1}^i)$					0.102* (0.0368)	0.148** (0.0626)	0.390*** (0.0684)	0.128** (0.0601)
Obs	720	240	240	240	576	192	192	192
Adj R^2	0.395	0.326	0.473	0.296	0.395	0.363	0.579	0.319
Treatment W	All Pds (1)	First 5 (1)	Second 5 (1)	Third 5 (1)	All Pds (2)	First 5 (2)	Second 5 (2)	Third 5 (2)
y	0.524*** (0.010)	0.441*** (0.017)	0.327*** (0.017)	0.606*** (0.018)	0.529*** (0.011)	0.439*** (0.019)	0.330*** (0.021)	0.613*** (0.019)
s_t^i	0.478*** (0.010)	0.562*** (0.017)	0.674*** (0.017)	0.393*** (0.018)	0.474*** (0.0114)	0.563*** (0.0184)	0.671*** (0.0208)	0.385*** (0.0198)
$\theta_{t-1}^i - E(\theta_{t-1}^i)$					-0.0055 (0.024)	-0.0112 (0.038)	0.065 (0.0454)	0.0123*** (0.0431)
Obs	720	240	240	240	576	192	192	192
Adj R^2	0.744	0.805	0.866	0.628	0.744	0.823	0.837	0.642

Standard errors are in parentheses

*** Stands for $p < 0.01$, ** for $p < 0.05$, and * for $p < 0.1$

After controlling for individual fixed effects, we still find overweighting of the private signal in the first ($0.645 > 0.5$), second ($0.979 > 0.8$) and third 5 periods ($0.532 > 0.2$) of the P treatment. In treatment W, controlling for fixed effects, we find that the private signal is over-weighted in the first ($0.562 > 0.5$) and third 5 periods ($0.393 > 0.2$), and under-weighted in the second 5 periods ($0.674 < 0.8$). Overall, the deviation of subjects' weight allocations from the signal extraction model predictions is again found to be smaller in treatment W as compared with treatment P.

We further consider evidence for learning behavior using a modified version of the linear, individual fixed effects regression model (1). Specifically, we suppose that even though the realization of θ_t^i is known to be *independent* across each period of our experiment, subjects nevertheless employ some type of adaptive learning behavior. That is, they incorrectly incorporate the immediate lagged forecast error into their prediction for the current period θ_t^i . To allow for this possibility, we estimate the following forecast rule:

$$E(\theta_t^i) = c_1 y + c_2 s_t^i + c_3 (\theta_{t-1}^i - E(\theta_{t-1}^i)) + \epsilon_t^i, \quad (2)$$

where, as before, $y = 10$, s_t^i is the private signal and $(\theta_{t-1}^i - E(\theta_{t-1}^i))$ is the lagged forecast error which captures subjects' tendency to adjust their expectations upwards/downwards when their expectation was too low/high in the last period. When running this estimation, we deleted the values of $(\theta_{t-1}^i - E(\theta_{t-1}^i))$ for periods 6 and 11 because they are the first period of a new block with a different variance of the private signal. Of course, we also do not have an error term in the first period as well.

Results from a fixed effects panel data regression of the modified specification (2) are reported in columns 6–9 of Table 2 for each treatment, P or W, using the whole sample (all Pds) and for each block of 5 periods. In general, the fit of this model seems a little better; the adjusted R^2 values are the same or better than for the model without the error term. The coefficient, c_2 , on the private signal in the current period, s_t^i , is again significant at the 1% level for all samples and subsamples. After controlling for individual fixed effects, the private signal continues to be overweighted in the first ($0.658 > 0.5$), second ($1.017 > 0.8$) and third 5 periods ($0.500 > 0.2$) in treatment P, and it is overweighted in the first ($0.565 > 0.5$) and third 5 periods ($0.385 > 0.2$), but underweighted in the second 5 periods ($0.671 < 0.8$) of treatment W. For the P treatment, the coefficient on the lagged error term, c_3 , is positive and significant at the 10% level for the whole sample, and at the 5% level for each block of 5 periods in treatment P. By contrast, in the W treatment, the coefficient on the lagged error term is not significantly different from zero for the whole sample and among the 5 period blocks, it is only significantly different from zero in the third block of 5 periods (albeit at the 1% level). This difference suggests that subjects may have been engaging in some type of adaptive learning behavior in the more complex treatment P, even though they should not be using their own past prediction errors in this experiment given that the realization of θ_t^i is i.i.d., and this fact was made clear to subjects in the instructions. Still, this

difference may help to explain why the deviation from the signal extraction model is larger in treatment P than in treatment W. When the lagged error term is included in the regression for treatment P, we observe that the estimated average weights assigned to the private signal, c_2 , get closer to the theoretical predictions for the whole sample and for the final 5 periods. Still, the added explanatory power of the forecast error term in the P treatment is somewhat limited, as the increase in the adjusted R^2 value from the model including the lagged error term is rather modest.

The above findings can be summarized as Result 3:

Result 3: *The signal extraction model also provides a good description of subjects' predictions at the individual level. As indicated by Result 2, subjects deviate less from the signal extraction model predictions in treatment W than in treatment P. There is a general tendency of subjects to overweight the private signal in treatment P (particularly in the first and third 5 periods) while subjects in treatment W over- (under-) weight the private signal when the theoretical prediction of the signal extraction model is low (high). Adaptive behavior (utilization of past forecast errors) has some explanatory power in treatment P, but not for treatment W.*

5 Conclusions

We have reported the first direct experimental evidence on the empirical relevance of the signal extraction model. Our experimental setup serves as a good platform for testing economic models that make use of signal extraction, e.g., Lucas (1972), Wallace (1992), Allen et al. (2006), Nimark (2008)) and Gabaix and Laibson (2017), among many others.

In general, we find that the signal extraction model provides a good description of agents' behavior at the *aggregate* level, which may be considered as a supportive evidence for applying signal extraction in macroeconomic modeling, using a representative agent framework. However, we also find that individuals may deviate from the predictions of the signal extraction model substantially at the disaggregated individual level, especially when they make point predictions as in our P treatment. We suspect that the comparatively better, individual-level performance of subjects in our W treatment stems from the fact that our W treatment puts constraints on the weighting of the two signals that are *not* present in the prediction (P) treatment. It would be of interest to think about extensions where we restricted the range of admissible predictions in the P treatment in a similar fashion to explore whether such restrictions also improved individual level performance in the P treatment. Alternatively, one might explore a reverse version of the W treatment where subjects use a slider to choose a point prediction and see the implied weight on the private signal. Finally, it would be of interest to compare individual subject's behavior when facing both the P and W treatments with the same variance for the private signal, to see if there is any consistency in their behavior. We leave all of these extensions to future research.

Acknowledgements We thank Frank Heinemann, Audrey Hu, Shaowei Ke, Rosemarie Nagel, and Songfa Zhong for their helpful comments, and Jiahua Zhu for excellent research assistance. Te Bao acknowledges financial support from a Tier 1 Grant from the Ministry of Education of Singapore (RG19/69), NTU-WebBank Joint Research Center (NWJ-2020-003) and the National Natural Science Foundation of China (No. 71803201, 71773013, and 71873149).

References

- Allen, F., Morris, S., & Shin, H. S. (2006). Beauty Contests and Iterated Expectations in Asset Markets. *Review of Financial Studies*, 19(3), 719–752.
- Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. In B. D. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of behavioral economics* (Vol. 2, pp. 69–185). Amsterdam: Elsevier.
- Chambers, C. P., & Healy, P. J. (2012). Updating toward the signal. *Economic Theory*, 50, 765–786.
- Cornand, C., & Heinemann, F. (2008). Optimal degree of public information dissemination. *The Economic Journal*, 118(528), 718–742.
- Cornand, C., & Heinemann, F. (2014). Measuring agents' reaction to private and public information in games with strategic complementarities. *Experimental Economics*, 17(1), 61–77.
- DeGroot, M. H. (2004). *Optimal statistical decisions*. Hoboken: Wiley.
- Fang, H. & Moro, A. (2011). Theories of statistical discrimination and affirmative action: A survey. In J. Benhabib, A. Bisin, & M. O. Jackson (Eds.), *Handbook of Social Economics* (Vol. 1, pp. 133–200). North-Holland.
- Gabaix, X., & Laibson, D. (2017). Myopia and discounting. NBER Working Paper No. 23254.
- Heinemann, F., Nagel, R., & Ockenfels, P. (2004). The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica*, 72(5), 1583–1599.
- Henckel, T., Menzies, G. D., Moffatt, P. G., & Zizzo, D. J. (2017). Sticky belief adjustment: A double hurdle model and experimental evidence. SSRN Working Paper No. 2954472.
- Kolmogorov, A. N. (1941). Interpolation and extrapolation. *Bulletin de l'Academie des Sciences de U.S.S.R., Series Mathematics*, 5, 3–14.
- Lucas, R. E, Jr. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2), 103–124.
- Massey, C., & Wu, G. (2005). Detecting regime shifts: The causes of under- and overreaction. *Management Science*, 51(6), 932–947.
- Menzies, G. D., & Zizzo, D. J. (2012). Monetary policy and inferential expectations of exchange rates. *Journal of International Financial Markets, Institutions and Money*, 22(2), 359–380.
- Mills, T. C. (1982). Signal extraction and two illustrations of the quantity theory. *American Economic Review*, 72(5), 1162–1168.
- Nimark, K. (2008). Monetary policy with signal extraction from the bond market. *Journal of Monetary Economics*, 55(8), 1389–1400.
- Pollock, D. S. G. (2013). Filtering macroeconomic data. In N. Hashimzade & M. A. Thornton (Eds.), *Handbook of research methods and applications in empirical macroeconomics* (pp. 95–138). Northampton: Edward Elgar.
- Sargent, T. J. (1991). Equilibrium with signal extraction from endogenous variables. *Journal of Economic Dynamics and Control*, 15(2), 245–273.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665–690.
- Wallace, N. (1992). Lucas's signal extraction model: A finite state exposition with aggregate real shocks. *Journal of Monetary Economics*, 30(3), 433–447.
- Wiener, N. (1941). Extrapolation, interpolation and smoothing of stationary time series. Report on the Services Research Project DIC-6037. Published in book form in 1949 by MIT Press.
- Wolfers, J. (2002). Are voters rational?: Evidence from gubernatorial elections. Stanford Graduate School of Business Working Paper No. 1730.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.