

Experiments on the Social Value of Public Information*

Donald J. Dale

John Morgan

Muhlenberg College

University of California, Berkeley

June 2012

Abstract

We present the results of laboratory experiments studying the “CNBC Effect”—situations where the social value of public information is negative. Payoffs depend on two factors: being right (i.e. matching an underlying, but unknown state variable) and coordinating with others, which has no social value. In our baseline treatment, individuals privately receive an informative signal. When we add a second, lower quality private signal, decisions improve. When the lower quality signal is public, however, (a) subjects strategically place inefficiently high weight on the public signal and (b) welfare (aggregate payoff) falls by 12% compared to the baseline. Welfare losses are due both to inefficient information use and greater random variation in choices.

JEL #s: C72, C92

Keywords: public information, CNBC effect, coordination

*We are grateful to Alessandro Pavan for helpful conversations. We also acknowledge the helpful comments of participants at the Bay Area Experimental Economics Conference for their feedback.

1 Introduction

Jim Cramer rants...a lot. Cramer, the star of the CNBC show *Mad Money*, offers impassioned stock recommendations during each show. The analysis for each stock is quite short, typically lasting for fewer than 2 minutes of airtime. Despite the paucity of analysis, Cramer's show is highly influential. With an audience of just under half a million people on average, the stocks touted by Cramer often experience unusually high volume in the period immediately following a recommendation.¹

Classical finance theory suggests that shows like Cramer's offers no benefit to investors. First, all of the information used in the recommendations is public. Second, since the broadcast is aired to so many people, any additional value from Cramer's unique use of this data should be competed away by investors. Appropriately, then, his show is titled *Mad Money*; based on classical models, investors would have to be mad to follow Cramer's advice.

John Maynard Keynes might, however, have a different view of the situation. In Keynes' view, successful stock picking was not just about observing discrepancies between a firm's fundamentals and its current selling price; it was also about anticipating the stocks others would find attractive. Keynes famously likened the situation to "beauty contests" offered in the London newspapers at the time. In the contests, photographs of various women were displayed, and the winner of the contest was the person who selected the most popular choices. Keynes' point was that, at least in the short run, there was a coordination aspect of successful stock picking apart from determining fundamental value.²

Thus, investors are pulled upon by two forces. They care about fundamentals as this will determine long-run stock returns. They also care about the perceptions of other investors as this will determine short-run stock returns. The seminal paper of Morris and Shin (2001)

¹Engelberg, *et al.* (forthcoming) provide evidence of 3% abnormal overnight returns for stocks picked by Cramer followed by frenzied short-selling the following day.

²The connection of stock trading to beauty contests appears in Keynes (1936, p. 156). Keynes' himself was largely a value investor, however. (Harrod, 1951).

formalizes the resolution of these competing forces in an elegant model.³ They postulate that an investor’s payoffs depend not only on how an individual’s choice relates to fundamentals (modeled as an underlying state variable), but also how it compares to other choices. While the first aspect has social value, the second, owing to the zero sum nature of trading gains and losses, does not. The surprising implication to emerge from their analysis is that additional *public* information can reduce payoffs on average—even when investors are fully rational. In other words, public information can have negative social value.⁴ The implications of this finding are profound: it argues against rules requiring transparency and disclosure, particularly in financial settings. For instance, it offers a ready justification for the secrecy of central banks.

The main intuition can be illustrated as follows: Suppose that each investor has private access to a high quality unbiased signal regarding the state variable, and this represents their sole information source. Obviously, investors act on this information and as a result, the average choice of investors, which one can roughly think of as the price of the stock, correctly encodes all available information. As the number of investors grows large, this average choice converges to the state variable. Adding a lower quality *public* signal to the mix alters investor choices. The beauty contest aspect of the stock picking game leads investors to overweight this signal relative to the efficient use of information. As a result, choices may become *less* correlated with fundamentals than in the absence of the public signal. If this strategic effect is large relative to the incremental information provided by the low quality public signal, then the average investor is made worse off. In other words, the model predicts that shows like *Mad Money* can, and will, influence investment decisions possibly driving stock prices away from fundamentals. In the model, the reaction to *Mad Money* is not the product of

³Their model gave birth to a vibrant theory literature examining these same two forces across a range of contexts. See, e.g. Angeletos and Pavan (2007), Dewan and Myatt (2008), and Myatt and Wallace (2012).

⁴Angeletos and Pavan (2007) point out that the absence of social returns from coordination is necessary for public information to have negative social value.

naivete or misperception of information quality among investors, rather it is a consequence of strategic calculations made under full rationality.

While this story is elegant and surprising, one wonders if it is, in fact, true as it requires a considerable degree of strategic sophistication on the part of investors. So far as we are aware, there have been no direct empirical tests of the phenomenon. Thus, to examine this question, we turn to controlled laboratory experiments. These have the advantage of allowing us precise control over the quality and type of signals investors receive. We can also observe the underlying state variable in the experiment. We can readily determine the impact of public signals on investor welfare.

Before proceeding, a note of caution is in order. The external validity of laboratory experiments is always subject to scrutiny. Our main research question is the extent to which actual behavior adjusts strategically to public information and the welfare consequences of this adjustment. To examine this, subjects participate in a stylized environment reflecting the competing forces faced by investors but without the particulars or context of an actual stock market. Since the effect is hypothesized to operate on the overall population, rather than a specialized population of sophisticated stock traders, our subject pool of undergraduates offers a reasonable proxy. The main power of the study is in the negation of the effect. Were we to find little evidence of public information shifting behavior in the close to ideal conditions of the laboratory, it would cast substantial doubt about the plausibility of the effect in practice. On the other hand, a positive finding suggests that the effect could be plausible, but says little about the magnitude of the effect in the “real world” outside the lab.

In this paper, we adapt almost exactly the Morris-Shin model for use in laboratory experiments. In the baseline treatment, subjects privately receive a single, reasonably precise, unbiased signal about the state. Equilibrium consists of choosing an action that is approximately equal to the signal. Actual subject play is reasonably close to the benchmark: On

average, subjects place about 90% weight on the signal. The remaining 10% is placed on the prior.

We then add a second private signal. This signal is generated the same way as the first but has a variance that is 4 times as large. Thus, the second signal is of lower quality than the first. Optimal signal processing consists of placing 80% weight on the high quality signal with the remaining 20% weight on the low quality signal. As a consequence of the additional information contained in the second signal, participants processing their signals optimally will choose an action closer to the underlying state, on average.

In the experiments, however, subjects tend to over-weight the low quality signal, placing about 26% weight on the low quality signal and 64% on the high quality signal with the remaining 10% on the prior. This reduces the welfare gains from the additional information. Even though welfare increases by about 7% compared to the baseline, the difference is not statistically different at conventional confidence levels.

Finally, we change the low quality signal from private to public. That is, a single low quality signal is drawn and observed by all players. While the additional information provided by the public signal is modest, a non-strategic planner could still make use of it to increase welfare. With strategic players, however, the publicness of the signal offers a chance to capitalize on the coordination aspect of payoffs. Theory suggests that, relative to the socially optimum, individuals will over-weight the public signal and welfare will fall. This is precisely what we observe in the experiment. While the degree of over-weighting of the public signal is *less* severe than theory predicts, it still leads to about a 12% drop in welfare, an amount that is both economically and statistically significant.

From a policy perspective, our results raise a number of thorny issues. In particular, they suggest that stock touts, such as Jim Cramer, are not merely a harmless diversion.⁵ More broadly, our results offer additional evidence of the potentially pernicious effects of

⁵In fact, Cramer may be talented in picking some stocks. Lim and Rosario (2010) offer evidence in favor of Cramer's long-run stock picking ability for small-caps.

government disclosure and transparency requirements. Specifically, in situations where coordination motives predominate, such as currency trading, requiring central banks to disclose their intentions can impede markets from functioning properly.

With respect to TV personalities, one might think that the media market itself might solve this problem. Unlike our experiment or the model on which it is based, investors in the real world are free to opt out of watching Jim Cramer and others of his ilk. Moreover, this would seem like the most sensible course of action for touts who reduce market efficiency. But this ignores the coordination aspect of payoffs. If other investors are watching *Mad Money*, then changing the channel merely puts the investor at a disadvantage. Thus, even if the quality of the analysis is low, shows that enjoy mass audiences can maintain their viewership through this novel network effect.

To sum up, the main contribution of the paper is to show that the negative value of public information postulated in the Morris-Shin model is not merely a theoretical curiosity. The introduction of a public signal shifts choice behavior and leads to both economically and statistically significant welfare losses despite the fact that all investors are better informed.

Related Literature

Obviously, our work is closely tied to Morris and Shin (2002). In section 2, we describe how we implemented their setting in the lab as well as deriving some additional results related to that adaptation. Our work is also somewhat related to the succeeding theory literature arising extending and generalizing the Morris-Shin framework.

From an experimental perspective, our work represents a variant of the familiar p -beauty contest (Nagel, 1995). Beauty contest experiments highlight the failure of most every version of equilibrium to accurately describe behavior. The important cognitive hierarchies literature (Camerer, Ho, and Chong, 2004 as well as Crawford and Costa-Gomez, 2006) largely arose as an attempt to explain behavior observed in these games. The addition of state dependent payoffs in our game dramatically effects equilibrium predictions and performance.

To see this, consider a version of the baseline treatment where subjects are only paid based on coordination. Under that circumstance, coordination on *any* number comprises an equilibrium. If, on the other hand, arbitrarily small weight were placed on matching the state variable, then the unique symmetric equilibrium is to choose an action equal to the signal. More broadly, our study offers a new variant in the substantial experimental literature on coordination games. Ochs (2001) offers a useful survey of this literature.

The spirit of our study is akin to the experimental voting literature. See Palfrey (2009) for a survey. Both literatures are motivated to use laboratory experiments to shed light on the behavioral relevance of sophisticated equilibrium reasoning in applied settings. They also share the feature that payoffs are state-dependent (at least in Condorcet jury voting settings) and also affected by aggregate behavior. The main divergence is in the models themselves. Strategy in the voting literature centers on the extent to which the voter incorporates information stemming from being pivotal, i.e. casting the decisive vote. Such pivotality considerations are entirely absent in our setting.

The remainder of the paper proceeds as follows: Section 2 outlines the design of the experiment, including our adaptation of the Morris-Shin model and its theory predictions. Section 3 reports on the results of the experiments. We find strong evidence of a strategic response to the publicness of the signal as well as the consequent welfare reduction. Finally, section 4 concludes.

2 Experiment

The experiment replicates the Morris-Shin theoretical framework. We study the effects of public information on choice behavior and resulting welfare in settings where individuals have state-dependent payoffs and coordination motives.

2.1 Theoretical Underpinnings

In each round of the interaction, we generated the underlying state (θ) from a uniform distribution on the interval $[0,10000]$.⁶ Each of the n participants simultaneously chose a location, a_i , on this interval. An individual's goal was to maximize the objective function

$$u_i = -(1 - r)(a_i - \theta)^2 - r(L_i - \bar{L})$$

where

$$\begin{aligned} L_i &= (a_i - \bar{a})^2 \\ \bar{L} &= \frac{1}{n} \sum_{j=1}^n L_j \\ \bar{a} &= \frac{1}{n} \sum_{j=1}^n a_j \end{aligned}$$

The first term in the objective function is the state-dependent (accuracy) portion of payoffs. The participant loses welfare according to the square of the distance between his selection and the underlying state. The second term represents the coordination motive. Here, a participant suffers a loss proportional to the square of the distance between the participant's selection and the mean of all participants' selections; this second term is normed to be zero-sum over all participants.⁷ The parameter r determines the relative weight placed on the coordination motive compared to the accuracy motive. Hence, the participant's goal is to choose an action close to the (unknown) underlying state and also close to the average of all participants' selections, with the relative weight of those two goals determined by r . When the coordination motive predominates (i.e., r is sufficiently high), the social value of public information turns negative. Hence, we chose $r = 4/5$ in the experiment.

⁶In our implementation, the actual number of choices available to the participants equal to the width of the slider bar in pixels on the PC display, which was approximately 1000.

⁷This specification differs slightly from Morris and Shin, but provides the same incentives as in their model. See Appendix B for details. The main advantage of this specification is that payoffs can be expressed purely as a function of the average of all the choices rather than as a vector of the choices.

In making choices, individuals receive one or more conditionally independent normally distributed signals. Some of these signals are private while others are public; that is, all individuals receive the same signal draw, and this fact is common knowledge.

Equilibrium

The analysis of equilibrium is quite tractable if one assumes, as Morris and Shin do, an improper prior (i.e. the state is uniformly distributed on $(-\infty, +\infty)$). For purposes of the experiment, it suffices to characterize equilibrium under three conditions: the case where each individual receives one private signal; the case where each individual receives two private signals with differing precision; and the case where each individual receives one private and one public signal whose precisions differ from one another.

The first case is trivial to analyze. Notice that, for when individual i receives signal s_i , $E[\theta|s_i] = s_i$, which is a consequence of placing no weight on the (improper) prior. Moreover, conditional on signal s_i , the expected value of any other individual j 's signal is $E[a_j|s_i] = s_i$. Thus, the coordination and accuracy motives are aligned in this case. As a consequence, equilibrium takes a particularly simple form. Morris and Shin show that, in the unique equilibrium,

$$a_i = s_i$$

for all i . That is, each individual simply chooses an action equal to the signal.

When there are two private signals with variance σ_1^2 and σ_2^2 , the analysis is similar since again the two motives are aligned. In that case

$$E[\theta|s_{1i}, s_{2i}] = \frac{\alpha s_{1i} + \beta s_{2i}}{\alpha + \beta}$$

where $\alpha = \frac{1}{\sigma_1^2}$ and $\beta = \frac{1}{\sigma_2^2}$ (i.e. the precision of the relevant signal). Again, the unique equilibrium is where

$$a_i = \frac{\alpha s_{1i} + \beta s_{2i}}{\alpha + \beta}$$

for all i . And furthermore,

$$E[a_j|s_{1i}, s_{2i}] = \frac{\alpha s_{1i} + \beta s_{2i}}{\alpha + \beta}$$

so that expected coordination payoffs are also maximized.

For these two cases, the equilibrium use of information also coincides with the efficient use of information. A planner directing individuals to make choices as a function of their realized signals can do no better than the equilibrium outcome.

The case where s_2 is a public signal is more involved. Now, there is conflict between the coordination motive, which dictates placing *nearly* all the weight on s_2 and the payoff associated with choosing an action close to the state, which dictates placing less weight on s_2 . The result is a “compromise” where excess weight (relative to the optimum were one simply trying to match $E[\theta|s_1, s_2]$) is placed on s_2 . Specifically, in the unique equilibrium, optimal actions are

$$a_i = \frac{\alpha(1-\rho)s_1 + \beta s_2}{\alpha(1-\rho) + \beta}$$

where $\rho = r \frac{(n-1)(n-2)}{n^2}$.

Notice that the degree of distortion (ρ) is proportional to r , the weight placed on the coordination motive. When this motive is completely absent ($r = 0$), the publicness of the signal is of no consequence—equilibrium actions are identical to the two private signals case. When $r = 1$, then, regardless of the precision of the private signal, in equilibrium it is ignored completely as the number of participants grows large.

When individuals receive a public signal, equilibrium and efficiency conflict. The efficient use of information would have individuals treat s_2 as private rather than public and place correspondingly less weight on this signal. Since the coordination motive is zero-sum, distortions created by this motive represent social losses, and, for this reason, public information can have negative social value.

2.2 Treatments

The experiment uses exactly the payoff, state, and signal structure described above. Each session comprised 100 rounds of interaction. Each round was one of three (randomly chosen)

strategic interactions, which varied in the number and character of the signals presented to the participants. In a baseline (single private signal) interaction, each participant received a unique, private signal of the underlying state. This signal was normally distributed around the underlying state, with a standard deviation of 833 (the entire interval spans 12 of these standard deviations). In a two private signal interaction, each participant received a signal as in the baseline and was also presented with an additional signal. This signal was also normally distributed around the underlying state, with a standard deviation of 1,667, or twice that of the first. In a public signal interaction, each participant received a private signal exactly as in the Type 1 interaction and also a common, public signal. This public signal was normally distributed around the underlying state with a standard deviation of 1,667. This signal was common to all participants within a round, and this commonality was common knowledge. All signals were truncated below at 0 and above at 10,000.

This treatment design permits sharp tests of the phenomena postulated by Morris and Shin. The only difference between the baseline and public signal interactions is the addition of a common, public signal; consequently, comparing participant welfare in these two treatments is a direct test of whether the social value of information is negative. Also, the only difference between the two private signal and public signal interactions is that whether the less precise signal is public or not. Comparing the within-subject weights given to the low quality signal under each treatment enables us to directly test the hypothesis that participants respond strategically to the ‘public-ness’ of the public signal.

2.3 Game Play

Figure 1 shows a screen capture of actual game play in a public signal interaction. The interval $[0,10,000]$ is represented by the slider bar that spans the top portion of the screen. Atop the slider bar is a set of tick marks; these function as a ruler, each denoting 833 spaces on the slider bar or 1 standard deviation of a signal. The black arrow is the individual’s

private signal, which is unbiased and distributed around the underlying state.⁸ This signal has a standard deviation of 833, or exactly one tick mark interval. The blue arrow is the public signal common to all participants for this round. This is also unbiased, with a standard deviation of 1,667 or two tick marks. We call this signal ‘half precision’ in the instructions and on the screen to remind the participants that this signal has a standard deviation that is twice that of the black signal. A baseline interaction would display only a black arrow, as that interaction does not involve a second signal. In a two private signal interaction, each participant would see a black arrow (her first signal) and a red arrow, representing her second, private signal with a standard deviation of 1,667 (‘half precision’).

After observing the signal or signals, the participant chooses an action with the mouse by clicking on the handle of the slider bar and dragging it to her desired location. When the participant has chosen an action, she clicks the onscreen Enter button. After all participants have made their selections, results and scores are computed and displayed to the participants.

Figure 2 shows the feedback provided to the participants after each round. After each round, the realization of the underlying state (the “Secret Spot”) is shown as a green arrow above the slider bar. The location of the mean of all participants’ actions is shown as a red arrow. We also display a yellow arrow that represents the location that, if chosen, would have earned that individual the maximum possible number of points in that round. This ex post optimum is simply the weighted average of the location of the underlying state and the average of all other participants’ actions, with the weights determined by the parameter r and the number of participants in the session.

We provided this additional feedback to help subjects think through the admittedly complex counterfactual analysis of which action they *should* have chosen. In earlier pilots, we simply provided a list of the choices of all subjects which, in principle, would permit the construction of this counterfactual. In practice, we found that subjects were simply

⁸We refer to the underlying state realization as the “Secret Spot” in both the instructions and the feedback screen.

overwhelmed by this data. Our summary signal offers a simple, graphical way to see that losses are proportional to the distance from this ex post optimum. One might worry that this feedback is “training” subjects to play the equilibrium. This seems doubtful. First, the calculation of the ex post optimum choice does not differ when a signal is public or private, whereas this difference is central to equilibrium play. Second, since the calculation is based on ex post realizations of the state and the choices made by all individuals, it does not coincide with equilibrium nor could a subject attempt to play this “strategy” as she lacks the data necessary to compute it at the time choices are made.

Also shown onscreen at this point are scores for the round. Each individual receives (an arbitrary) 1,000 points; since the objective function is a loss function, this keeps most point levels positive without distorting incentives. In the next line, the distance between the individual’s action and the underlying state (the “Secret Spot”) is calculated and points are debited appropriately. In the next line, the participant can gain or lose points according to the distance between the participant’s actions and the mean of all participants’ actions. Since this term sums to zero across all participants in a round, it is possible to net positive points from this term if the participant’s selection is particularly close to the mean or if other participants are quite distant from the mean. The total net points are calculated in the next line. Finally, the individual is informed of how many more points she could have earned if she had chosen the ex post optimum. Also onscreen at this point is a Next Round button. When all participants click the Next Round button, play continues to the next round, and points earned are added to the participant’s cumulative total onscreen. At the end of 100 rounds, total points were converted to dollars through a payoff factor that was calculated to pay the participants an average of \$15 if all participants played optimally. Actual payoffs averaged about \$10.50, and each session took approximately 1 hour.

Table 1: Participants by Session

Session	Participants
1	8
2	10
3	14
4	16
5	12

2.4 Sessions

We conducted 5 replications of the experiment, each with a random number of interactions for each of the three treatments. (In fact, we seeded the random number process identically for each trial, so the actual realizations of all random numbers are identical and consistent across treatments). Participants were recruited from the student body by announcement on the main campus website/bulletin board at a Northeastern liberal arts college. Table 1 summarizes the sessions and treatments.

3 Analysis

With these preliminaries in place, we are now in a position to report on the results of the experiment. We begin by studying the strategies employed by a representative player. The theory model has the advantage that equilibrium strategies are simply linear functions of the signals themselves; thus, the regression coefficients of a linear regression may be viewed as estimates of the equilibrium signal weights, at least theoretically.

This interpretation does, however, require several caveats: First, the theory assumes that individuals hold an improper prior over the state space or, equivalently, that they act as classical statisticians in forming beliefs conditional on receiving each signal. Second, the theory assumes that the strategy space is a continuum whereas, of necessity, the strategy space in any experiment is finite. To address the first concern, we rederived the model in a

setting where the state is uniformly distributed over a finite interval and where individuals hold proper prior beliefs. This analysis, contained in the appendix, shows that, for the parameters of the experiment, correctly specifying the prior has little effect on equilibrium actions except near the edges of the state space. With regard to the second discrepancy, our hope is that allowing subjects the choice over 10,000 possible actions is sufficient to approximate the continuum.⁹

Our main hypotheses are as follows:

Hypothesis 1: In making choices, the weight placed on the private signal is highest in the one signal treatment and lowest in the public signal treatment.

Hypothesis 2: In making choices, the weight placed on the low quality signal is higher when this signal is public than when it is private.

Hypothesis 3: Welfare (the average payoff of an individual in the experiment) is highest in the two private signal treatment and lowest in the public signal treatment.

All three hypotheses are derived analytically. We hoped that the 33 or so rounds of each of the three treatments would help ensure that the state and signal realizations would approximate average behavior and thereby ensure that the particular state and signal realizations were not driving the results. We also performed a number of analyses, described below, to help verify that this is indeed the case.

Table 2 presents mean values of the state for the three treatments. Under the null hypothesis that these data consist of i.i.d. draws from a uniform distribution, a Wilcoxon sum of ranks test is appropriate. Comparing the two private signals treatment with the private/public signal treatment yields a p-value of 0.48, which does not come close to sta-

⁹In the private signal treatments, the alignment of private and social incentives ensures that discreteness is of no consequence. Moreover, since payoffs are smooth and continuous and play is in linear pure strategies, discrete approximations of the continuum equilibrium can readily be shown to be ε -equilibria, where ε is inversely proportional to the size of the grid. For the parameters of the experiment, this amounts to fractions of a cent.

tistical significance. The same is true of all other pairwise comparisons.¹⁰ The table also presents signal quality, measured as the difference between the signal and the underlying state realization. Since the signals are conditionally normally distributed, the appropriate test statistic for comparing mean quality across treatments is a t-test, which reveals no significant differences at conventional levels.¹¹ We also compare the standard deviations of signal quality across treatments, which are also predicted to be identical. Here, the Levene test is appropriate. This too reveals no significant differences across treatments.¹²

These analyses provide strong evidence that, if we observe differences in behavior across treatments, they stem from the treatments themselves rather than differences in state or signal realizations.

3.1 Choice Behavior

In equilibrium, theoretical choice behavior is summarized by the weight placed on the high quality private signal. The remaining weight is placed on any other signal an individual might receive. For the parameters of the model, the predicted weight is 1 in the single private signal treatment. It falls to 0.8 when a second, low quality private signal is added. In the public signal treatment, the weight depends on the number of players, which varied across sessions. Here, the weight ranges from a low of 0.58 when there are 16 subjects to a high of 0.66 when there are 8 subjects. Regardless, the addition of a public signal depresses the weight on the private signal.

¹⁰Running that same test on the other pairwise comparisons yields p-values 0.60 (baseline versus private/public) and 0.68 (baseline versus two public signals).

¹¹For the higher precision signal comparing baseline to the two signal and the public signal treatment, respectively, yield p-values of 0.37 and 0.23. Comparing the two signal and public signal treatment yields a p-value of 0.82. The same comparison for the low quality signal yields a p-value of 0.99.

¹²For the high-precision signal, comparisons of baseline versus two signals yield a p-value of 0.49, versus public signal yields a p-value of 0.35. Comparing the two signal and the public signal treatments yields a p-value of 0.83 for the high precision signal and 0.30 for the low quality signal.

To examine these predictions empirically, we estimate

$$choice_{ijtk} = \beta_0 + \beta_{j1}s_{ijtk}^1 + \beta_{j2}s_{ijtk}^2 + \gamma X + \varepsilon_{ijtk} \quad (1)$$

where s^1 denotes the realization of the high quality private signal, s^2 denotes the realization of the lower quality signal.¹³ The subscripts i , j , t and k denote subject, treatment, round, and session, respectively. Recall that in the two private signals treatment, an independent realization of s^2 is privately observed by each subject. In the public signal treatment, one single realization of s^2 is generated in each round. This is revealed to all subjects and this fact is commonly known. The regressor X represents various controls described in detail in the tables below. We run this regression specification separately for each treatment and also stacked in one regression with treatment effects and interactions.

Of course, since subjects repeatedly interacted with one another in a given session, there is no reason to treat observations as independent. To account for this, in all specifications, we compute robust standard errors clustered by individual to correct for possible autocorrelation or heteroskedasticity of a participant's choices in a session. In some specifications we use individual level fixed and random effects to account for additional correlation in the choices of each participant in the panel.¹⁴

In principle, one should estimate separate coefficients for the public signal treatment that depend on the number of subjects. This suggests that, for the public signal estimates, we should use the regression

$$choice_{itk} = \beta_0 + \beta_1 s_{itk}^1 + \beta_2 s_{itk}^2 + \gamma_1 (s_{itk}^1 \times n_k) + \gamma_2 (s_{itk}^2 \times n_k) + \gamma_3 n_k + \varepsilon_{it}$$

¹³Our software generated a value for s^2 in the single private signal treatment, but did not display this signal to the participants. We kept these values in our analysis in order to be able to 'stack' treatments in a single regression. Obviously, coefficients on this variable in this treatment are predicted to be zero.

¹⁴The usual test for the validity of random effects is a Hausman test, which cannot be used for regressions with clustered standard errors. To deal with this issue, we reran all specifications without clustering and applied a Hausman test. It failed to reject random effects for all specifications.

where n_k denotes the number of subjects in session k . We performed this analysis (shown as Table A.1 in the appendix) and found that the interaction terms as well as the γ_3 coefficient were both statistically and economically insignificant. That is, the theoretical relationship between choice weights and the numbers of subjects is simply absent from the data.¹⁵ As a consequence, we pool all sessions and omit the relationship with the number of subjects in a given session from the estimates.

One Signal Treatment

Table 3 displays the results of this analysis. The various columns represent differing specifications of the controls and the error term. The first three rows display the coefficient estimates for the baseline, single private signal case. This specification includes the more precise signal, the less precise signal, and the constant term, which may be viewed as a measure of center bias. In the table, this last term has been normalized to indicate the weight placed on the center point in the distribution. Since subjects did not observe the less precise signal for this treatment, it is comforting that the coefficient estimates yield an economic and statistical zero coefficient.

Turning to the more precise signal, regardless of the specification, the coefficient estimate on s^1 is 0.881, which is significantly different from zero and, more importantly, also different from the theoretical benchmark of 1. This estimate implies that subjects tend to shade choices more toward the middle of the state space than the theory predicts.

One possible rationale for this behavior is that, owing to truncation of the support of the state space, for signal realizations near the endpoints of the distribution, it is optimal for a subject to make center-biased choices. To investigate this possibility, we re-estimated

¹⁵Recall that this theoretical relationship arises from differences in the marginal impact of a change in choice on the value of the average choice. Given the modest effect that this variation has on predicted choices and the limited amount of variation in the independent variable n , it is not altogether surprising that it is absent from the data.

equation (1) deleting any signal realizations that were greater than 8,000 or smaller than 2,000. This amounts to approximately 2.5 standard deviations from either endpoint. For these realizations, endpoint effects should, theoretically, be absent. The results of this exercise are displayed in columns B and D of Table A2. These regressions both produce a coefficient estimate 0.89 with miniscule standard errors. As before, this estimate is statistically different from both zero and one at the 1% significance level. Thus, while endpoint effects account for a fraction of biased behavior, considerable center-biasing remains.

Another explanation is that, despite the length of the interval over which the state variable is drawn, subjects continue to place significant weight on the prior. Unfortunately, the combination of a uniform prior and a normal signal does not lead to closed form solutions for the optimal linear weighting scheme. However, using a Monte Carlo simulation consisting of 500,000 observations and 1,000 iterations, we approximated it numerically.¹⁶ This analysis reveals that a subject optimally places a weight of 0.945 on s^1 . While this is closer to our estimates, the coefficients reported in Table 3 are still statistically significantly different from this value at the 1% level.

To summarize, the observed weight in the one signal treatment significantly differs from the theory (or the simple heuristic strategy of choosing the signal). Instead, subjects are more likely to choose actions closer to the center of the support. One can account for a portion of the observed center-biasing through the combination of endpoint effects and optimal weight given to the prior; however, this still leaves about 4 percentage points in weight (roughly one-third of the total deviation from the theory model) unaccounted for. One might speculate that this represents some form of aversion to extreme choices though given the number of choices (10,000) and the presence of the effect for signals well away from the edge of the choice space, this seems doubtful.

¹⁶This analysis is available upon request from the authors.

Two Private Signal Treatment

The second pair of rows of Table 3 displays the results for the two private signal treatment. Consistent with the theory, the addition of a second signal leads to a reduction in the weight placed on s^1 . Moreover, center-bias is now diminished as well. In our preferred specification given in column E , the center bias is 8.5%. The reduction in the weight on the higher quality signal, s^1 , is larger than theory predicts. Indeed, the difference between our coefficient estimate and the theory benchmark of 0.8 is significant at the 1% level. Alternatively, one can calculate the ratio of the weights placed on s^1 and s^2 ; that is, $\beta_1/(\beta_1 + \beta_2)$. Regardless of the weight placed on the prior, this ratio should remain at 0.8. By way of comparison, the coefficient estimates yield a ratio $\hat{\beta}_1/(\hat{\beta}_1 + \hat{\beta}_2) = 0.72$, which again places less weight on the high quality signal than theory predicts.

Why do individuals over-weight the less precise signal? One possibility is that subjects fail to account for the difference in the precision of the two signals. Under that hypothesis, we would expect that these subjects would choose the midpoint between the two signals as their guess. There is, however, little support for this hypothesis. We computed the fraction of choices where the observed weight on the higher quality signal was between 0.48 and 0.52 and found that only 3.9% of observations fall into this category.¹⁷ Not surprisingly, eliminating these observations from the regression alleviates over-weighting to a small extent, but does not fully account for the phenomenon.

An alternative explanation that subjects are simply uncomfortable with the high (80%) weight that Bayes' rule places on s^1 . Given the feedback subjects receive connecting their choices to the particular realization of the signals, learning to place the correct weight is apparently quite difficult. Moreover, a variety of experimental studies have shown an “end-point aversion” effect—subjects simply do not like to choose endpoints or near endpoints

¹⁷We omit observations where the two signals are identical since, for these observations, the weight placed on either signal is undefined. This will typically occur at endpoints of the state and signal distribution. Approximately 2.6% of observations are omitted for this analysis.

from choice sets. If one thinks of the gap between the two signals as the appropriate choice space, then underweighting the better signal is consistent with this type of effect.

In addition to underweighting the more accurate signal, subjects continue to be center biased, although less so than in the one signal treatment. This is consistent with Bayes' rule, which suggests that the greater the information contained in the signals, the less weight should be placed on the prior. To establish this effect formally, we tested the coefficient on center bias comparing Treatment 1 to Treatment 2. In all specifications, the coefficient on center bias is lower when a second signal is provided to the participants. A straightforward explanation of this effect is that the assumption that subjects place no weight on the prior, as specified in the theory model, is simply incorrect. Even with a uniform distribution over 10,000 points, as in the experiment, subjects continue to place weight on the midpoint of the distribution in making their choices.

Additional evidence for weight being placed on the prior comes from examining the fraction of observations where choices lay outside the interval between s^1 and s^2 but were closer to the center point. For example, if $s^1 = 6,000$ and $s^2 = 7,000$ but the choice was 5,500 then clearly, the subject was trying to locate closer to the center despite the signal realizations. Roughly 25% of observations fall into this category.

To summarize, there is considerable evidence that subjects give positive weight to the prior in the private signals treatments. As expected, the weight falls when subjects receive a second signal. The weight, however, is still excessive relative to the Bayesian optimum.

Public Signal Treatment

Finally, we turn to the public signal treatment. Here, we see a dramatic shift in the weight placed on the lower quality signal. The relevant coefficient for our preferred specification (column *E* of Table 3) shows that subjects place 34.3% weight on the lower quality signal, 57.8% on the higher quality signal with the remaining weight placed on the center. Directionally, this is entirely consistent with the theory—the coordination motive leads subjects

to overweight the public signal relative to its precision. Examining this in terms of relative weights (i.e. the weight on signal 2 versus that on signal 1), the relative weight on the lower quality signal is 37%; thus, excluding center bias, the level predictions of the theory perform quite well.

A formal test of the strategic effect arising from making a signal public may be seen by comparing the coefficient on the less precise signal in Treatment 2 against that of Treatment 3. Under the null hypothesis of no strategic effect, the coefficients on the interaction terms for the two signals treatment should be equal to the coefficients for the public signals treatment. Performing a chi-squared test of this hypothesis for our preferred specification (column *E*) yields a test statistic of 70.63, which is significant at the 1% level. In short, the null hypothesis is definitively rejected in favor of the alternative that subjects strategically reduce the weight on the private signal and raise it on the public signal. Testing this same hypothesis results in similar results for our other specifications.

In a way, this is surprising given the weak incentives toward equilibrium convergence. First, as we noted above, the feedback between choice and signals is far from direct. More importantly, given deviations from equilibrium play by other players, the incentives of each individual player are to slightly overweight the public signal relative to their peers, but not to jump all the way to the Nash equilibrium prediction. For example, if others ignore the strategic effects of the public signal, and hence weight it at 20%, an individual's optimal strategy is to weight the public signal at about 21%, but not to jump to the Nash level of 33% to 42% weight depending on the number of competitors. Nonetheless, subjects recognize the coordination value of the public signal and adjust their signal weights accordingly.

Paradoxically, the fact that the lower quality signal contains *some* information is the source of harm. If the public signal were totally uninformative then the unique equilibrium is one where all individuals ignore the public signal and simply act on the high quality private signal. In other words, were Jim Cramer truly a madman with no knowledge of the

stock market, and if this fact were common knowledge, his recommendations would cause no harm (and likely attract no viewers as well). It is the fact that Jim Cramer is *somewhat* knowledgeable that produces the strategic response. More broadly, in considering disclosure policy, the main danger lies in public disclosures that are somewhat informative, which leads to a crowding out of the weight placed on more informative private signals.

The impact of this response on welfare contains two competing forces. On the one hand, by overweighting the lower quality signal, the average guess of subjects is likely to be further from the underlying state leading to welfare losses. There is, however, a compensating force as well—for a given distance between the average guess and the “fundamentals,” it follows from Jensen’s inequality that welfare is higher if choices exhibit lower variance. By placing greater weight on the public signal, choices are more correlated with one another and this could possibly compensate for the loss of precision.

3.2 Welfare

Recall that the payoff of each individual is determined by two components: the difference between her choice and the average of all choices and the difference between her choice and the state. Since the former component is zero sum across all players in a given round, the average payoff in a given round is simply the average of the difference between choices and the underlying state. Thus, the welfare loss under treatment j occurring in round t of session k is given by

$$loss_{jtk} = \frac{1}{n_k} \sum_{i=1}^{n_k} (choice_{itk} - \theta_{tk})^2$$

Hypothesis 3 predicts that welfare will vary by treatment. A simple measure of this is to compute the welfare ordering by treatment of the average loss suffered by all individuals in a given session. In all sessions, save for session 2, the ordering is exactly that predicted by theory: welfare is highest with two private signals and lowest when there is a public signal. In session 2, welfare is highest in the baseline but remains lowest when there is a

public signal. A binomial test offers a simple statistical validation of the fact that welfare is lower with a public signal compared to the baseline. Such a test reveals a 1 in 32 chance of the observed welfare ordering under the null hypothesis of no treatment effect. Thus, we can reject the null in favor of the one-sided alternative suggested by hypothesis 3. A similar analysis places the welfare under two private signals above that when there is a public signal.

The other implication of hypothesis 3, that two private signals improves welfare over one signal, is broadly consistent with the data, but not statistically significant. Again performing a binomial test reveals a 5 in 32 chance of the occurrence we observe, which is not significant at conventional levels.

Of course, the simple welfare ordering ignores individual variation in welfare outcomes and does not allow us to measure the marginal effects of each treatment on welfare subject to appropriate controls. Additional evidence on the effects of public information on welfare can be deduced using the regression specification

$$loss_{jtk} = \beta_0 + \beta_j I_j + \gamma X + \varepsilon_{jtk}$$

where the subscripts and variables are as defined earlier. The loss can be interpreted as a measure of the degree to which the stock price varies from fundamentals or, more broadly, that the consensus forecast for a given event might diverge wildly from reality. As with the choice regressions, we use a variety of error specifications, fixed, and random effects to account for group interaction within a session. The results of this analysis are shown in Table 4. The table reports the estimated welfare loss under each treatment. For example, the row for the one signal treatment is simply the constant term of the regression. The rows for the other treatments represent the coefficients on the relevant dummy variable. Since welfare is reported in terms of losses, negative coefficients indicate welfare improvements compared to the baseline treatment while positive coefficients indicate additional welfare losses.

The Social Value of Private Information

In our preferred specification (column *A*), adding a second private signal improves welfare as indicated by the negative coefficient on the treatment dummy. While the implied loss reduction is economically significant, at 7%, it is not statistically significant ($p = 0.182$). Thus, the directional prediction of Hypothesis 3 is supported although the variability of outcomes does not permit precise estimation. The result is entirely intuitive—giving individuals more information when there is no conflict between private and social incentives leads to an improvement in payoffs.

The Social Value of Public Information

Finally, we come to main finding of the paper: The addition of a *public* signal is predicted to increase losses despite the fact that all individuals are better informed than in the one signal case. As Table 4 shows, this hypothesis is strongly supported in the data. The coefficient in Table 3 shows that this treatment increases losses by 12% compared to the one signal case. The magnitude of this coefficient is both economically and statistically significant at the 5% level ($p = 0.031$) in models A and C and at the 1% level in the random effects specification (Model B). Importantly, it is also consistent with the directional prediction of Hypothesis 3—additional public information leads to worse results.

Excess Welfare Losses

While the directional predictions of the theory are borne out, the level predictions are not. Table 5 offers an accounting of the key factors leading to the difference between the theoretically predicted losses and the actual losses. The first set of columns presents welfare results for the baseline treatment. Here, the first row represents actual loss in the experiment, averaged by session. The second row presents the welfare loss were a planner able to specify the strategy of the player under an improper priors model. For this treatment, the planner’s solution coincides with the equilibrium prediction, again under improper priors. Notice that

actual losses are about 32% greater than theory predicts. Moreover, once one accounts for the finite length of the state space, losses fall even further as evidenced by the column labeled "Equilibrium play, optimal center bias." The key to reconciling this difference between theory and actual losses is to incorporate variability in players' choices. The column labeled "Center Bias + Noise" computes losses under the hypothesis that all individuals play with the optimal amount of center bias but with the addition of an error term, which is normally distributed with a standard deviation equal to the choice error estimated in Column C of Table 3, which is about 500. Notice that this produces a figure close to the experienced welfare loss.

The second set of columns in Table 5 accounts for welfare in the two private signals treatment. Here, the first row presents actual loss in the experiment, again averaged by session. The second row shows the welfare result that would have been obtained if participants had used the improper prior but weighted their signals optimally (80% weight on the more precise signal) when choosing actions. Here, the difference between the theory prediction and the experienced loss is even more stark, amounting to 52%. One might think that this is a consequence of the suboptimal signal weights subjects use in the experiments. The fourth row, where the same amount of decision noise (500) is added to choices under optimal play, shows that variability alone is sufficient to reconcile the theory to the actual welfare loss.

The third set of columns in Table 5 accounts for welfare in the treatment in which participants received private and public signals. As before, the first row displays actual welfare outcomes. The second row displays the welfare result that would have occurred if participants had used the optimal nonstrategic weights on their signals under improper priors. Comparing this to its analog under the one signal treatment, one can see that significant reductions in welfare losses are possible with the addition of the private signal—the strategic effect created by the coordination motive undoes these gains. This may be readily seen in the third row, which shows higher welfare losses under equilibrium play when a private signal is present than when it is absent. Welfare losses under actual play are vastly higher than

equilibrium predictions. Finally, adding decision noise as estimated from Column C of Table 3 as we did in the other two treatments results in welfare outcomes remarkably similar to the actual outcomes.

To summarize, the hypothesized strategic effect of a public signal is borne out in the data. Adding a second private signal improves choices while adding a public signal reduces choice accuracy and increases losses. Despite the fact that signal variation accounts for the overwhelming fraction of choice variation, the remaining noise in choices is primarily responsible for why measured losses exceed the theory benchmark.

4 Conclusion

Today we live in the so-called knowledge economy. Thanks to gains in information technology and the internet in particular, more individuals gather more information more intensively than ever before. Moreover, the speed with which information is disseminated to the public is now nearly instantaneous thanks to social media like blogs and Twitter. Throughout most of history, the key problem decision makers had to overcome was lack of information. Now, many face the opposite problem—there is simply too much information to be processed while still making timely decisions.

Another important trend over the last 30 years has been the tightening and globalization of the supply chain. Now, more than ever production is a tightly choreographed dance by numerous players up and down the chain. Obviously, the need for coordination is paramount.

Together, these two trends mean that the conflict identified by Morris and Shin almost a decade ago—the conflict between the need to be right and the need to coordinate—is more important than ever. The democratization of publishing via the blogosphere and elsewhere implies that the “CNBC effect,” the prospect of a perfect storm where the addition of public data can reduce the welfare of all, is increasingly problematic...at least in theory. Identifying this effect in the field is challenging and, to the best of our knowledge, no convincing

study demonstrating the effect exists. Using laboratory experiments, we can easily create conditions where the perfect storm scenario envisaged by Morris and Shin, in fact, arises.

We found that the problem of public information is not merely theoretical, but real and perhaps even more pernicious than the models would suggest. We showed that, compared to the one signal case, the addition of a lower quality signal helps, but only if it is private. When this same signal is public, more information leaves everyone worse off on average—there are no winners and losers, only losers. The theoretical effect is strategic. When the signal is public, individuals give it excessive weight, owing to the need to coordinate, and this leads to worse decisions. This exact effect operates in the laboratory as well. Subject choices are more responsive to the low quality signal when it is public than when it is private. But we also observe a second effect—variation in choice behavior. While overall this variation is limited—over 95% of the variation in choices can be explained purely through the signals received by subjects—the remaining 5% leads to excessive losses relative to the theory predictions. Moreover, the complexity of the public signal environment increases the probability and magnitude of choice errors, thus exacerbating the negative social value of public information.

The policy implications of our work are clear and surprising. While there is a knee-jerk tendency to view increased information disclosure and transparency as unambiguously good for improving the functioning of markets, our results suggest that more caution is required. Some types of disclosures—public disclosures with limited accuracy—are potentially problematic. In a sense, there is a market failure here. Since these disclosures do contribute some new information to the public dialogue, they will not be driven out in the so-called marketplace for ideas. But the externalities that these disclosures impose, owing to coordination motives, will not be priced into the marketplace for ideas either. Thus, there is the potential for regulatory solutions to offer improvements.

Determining what those solutions are and the extent to which they are needed requires

identifying the magnitude of the problem outside of the laboratory. While we characterized the effect as a “perfect storm” which implies rarity, it is unclear whether the effect is, in fact, common or unusual. This remains for future research.

References

- [1] Angeletos, G.M. and A. Pavan (2007) “Efficient Use of Information and Social Value of Information,” *Econometrica*, 75(4), 1103-1142.
- [2] Camerer, C., T.-H. Ho, and J.-K. Chong (2004) “A Cognitive Hierarchy Model of Games” *Quarterly Journal of Economics* 119(3), 861-898.
- [3] Crawford, V. and M. Costa-Gomez (2006) “Cognition and Behavior in Two-Person Guessing Games: An Experimental Study.” *American Economic Review* 96, 1737-1768.
- [4] Hossain, T., D. Minor, and J. Morgan (forthcoming) “Competing Matchmakers: An Experimental Analysis,” *Management Science*
- [5] Dewan T. and D.P. Myatt (2008) “The Qualities of Leadership: Direction, Communication and Obfuscation,” *American Political Science Review*, 102, 351-368.
- [6] Engelberg, J., C. Sasseville, and J. Williams (forthcoming) “Market Madness? The Case of *Mad Money*” *Management Science*.
- [7] Harrod, R. F. (1951) *The Life of John Maynard Keynes*. London: Macmillan.
- [8] Keynes, J. M. (1936) *The general theory of interest, employment and money*. London: Macmillan.
- [9] Kogan, S. Behavioral models in experimental finance
- [10] Lim, B. and J. Rosario, (2010) “The Performance and Impact of Stock Picks Mentioned on *Mad Money*,” *Applied Financial Economics*, 20(14), 1113-24.
- [11] Morris, S. and H. S. Shin (2002) “The CNBC Effect,” *American Economic Review*, Vol. 92(5), 1521-34.

- [12] Myatt, D. P. and C. Wallace (2012) “Endogenous Information Acquisition in Coordination Games,” *Review of Economic Studies*, 79, 340-374.
- [13] Nagel, R. (1995) “Unraveling in Guessing Games: An Experimental Study,” *American Economic Review*, 85 (5), 1313-1326.
- [14] Palfrey, T. (2009) “Laboratory Experiments in Political Economy” *Annual Review of Political Science*, 12, 379-388.
- [15] Ochs, J. (1995) “Coordination Problems” in J. Kagel and A. Roth eds. *Handbook of Experimental Economics*. Princeton NJ: Princeton University Press, 195-262.

Figure 1: Screenshot of game play

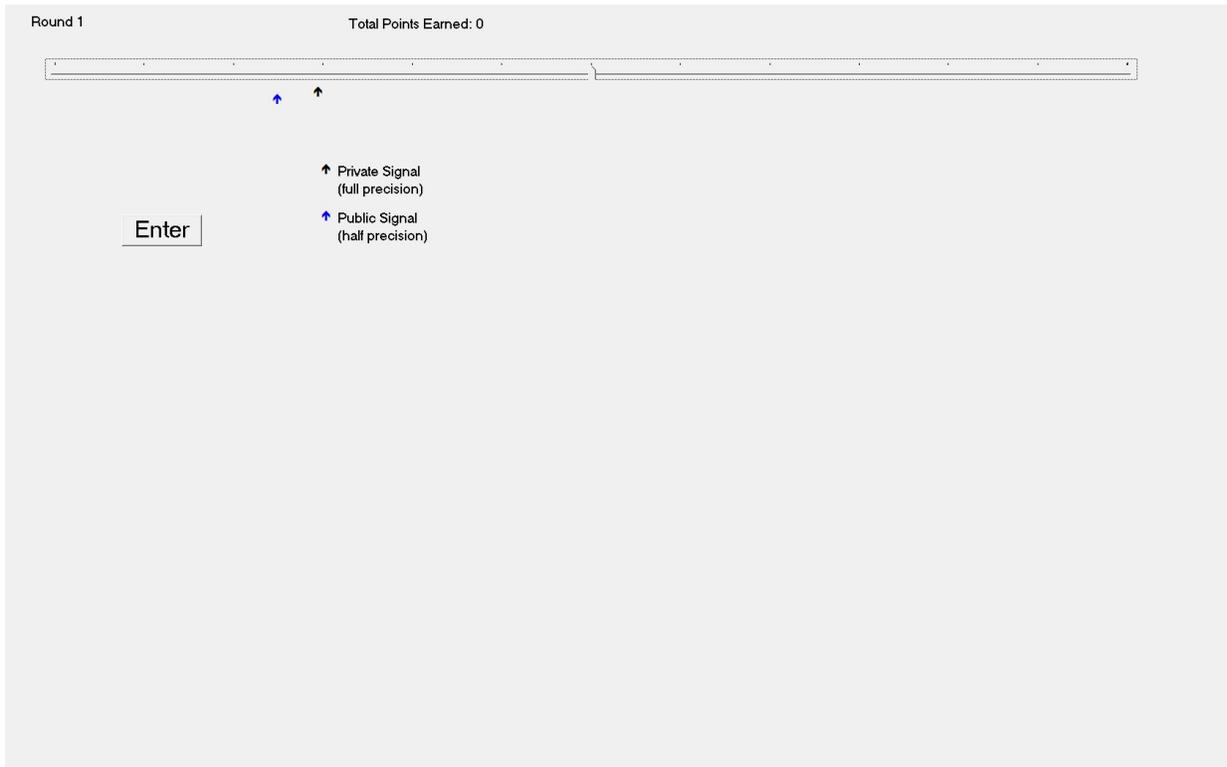
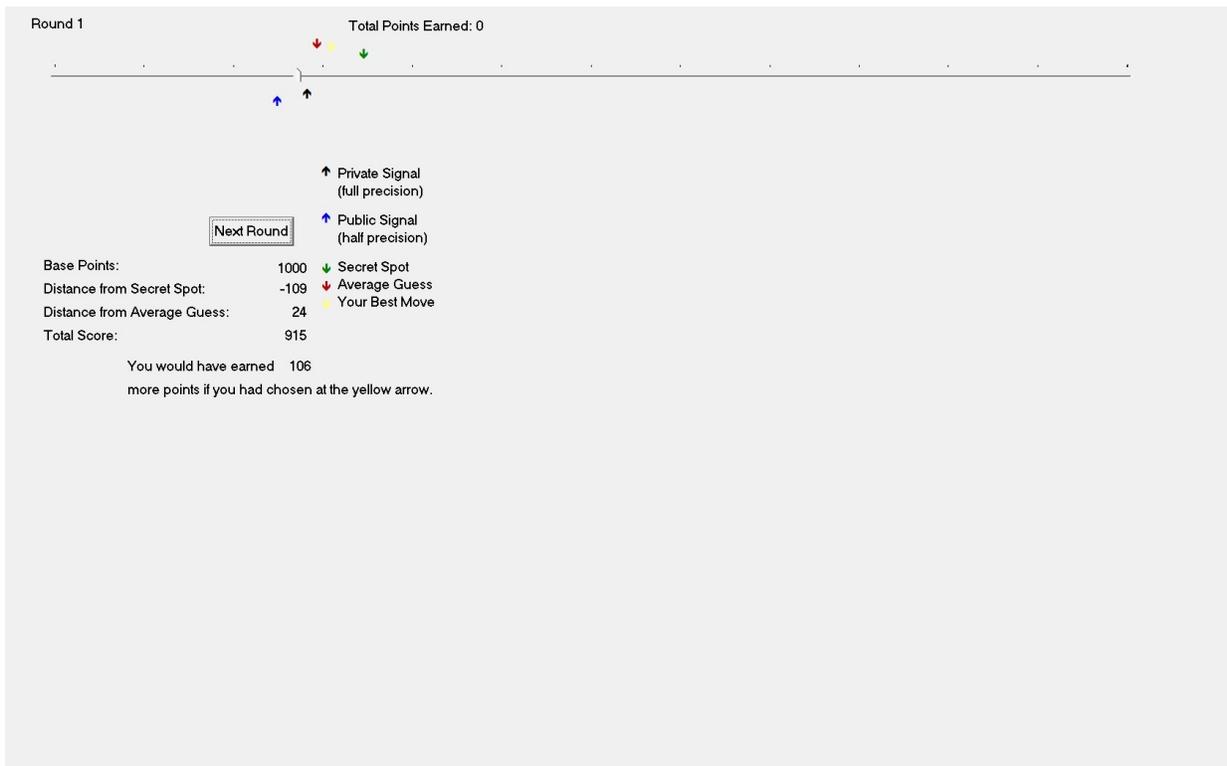


Figure 2: Screenshot of continuation of game play, presenting feedback



	State	More Precise Signal Quality	Less Precise Signal Quality
Baseline	4560 (2724)	-57.5 (777.6)	
Two Private Signals	4342 (2616)	-11.7 (801.9)	55.4 (1526.0)
Private/Public	4921 (3090)	-0.9 (809.4)	54.8 (1459.3)

Notes: This table presents realizations of signals and states aggregated across rounds and sessions, by treatment. The columns for signals indicate the mean difference between the signal realization and the underlying state variable. This is normally distributed with mean zero and standard deviation 833 for the more precise signal and 1667 for the less precise signal. The state column indicates the mean realization of the underlying state, which is uniformly distributed on $[0, 10,000]$. Standard deviations in parentheses.

Table 3: Testing Individual Strategy Choices

	Specification										
	A	B	C	D	E	F	G	H	I	J	K
Baseline Treatment:											
More precise (private) signal	0.881 (0.007)	0.881 (0.009)	0.881 (0.007)	0.881 (0.009)	0.881 (0.004)	0.881 (0.009)	0.881 (0.007)	0.881 (0.009)	0.881 (0.007)	0.881 (0.007)	0.881 (0.007)
Less precise (private) signal	-0.001 (0.004)	0.002 (0.005)	-0.001 (0.004)	0.002 (0.006)	-0.001 (0.004)	0.002 (0.006)	-0.001 (0.004)	0.002 (0.006)	-0.001 (0.004)	-0.001 (0.004)	-0.001 (0.004)
Center bias	0.120 (0.007)	0.120 (0.009)	0.120 (0.007)	0.120 (0.009)	0.120 (0.006)	0.120 (0.009)	0.120 (0.007)	0.120 (0.009)	0.120 (0.007)	0.120 (0.007)	0.120 (0.007)
Two Private Signals:											
More precise (private) signal	0.65 (0.013)	0.674 (0.02)	0.65 (0.013)	0.674 (0.02)	0.65 (0.008)	0.673 (0.02)	0.649 (0.013)	0.673 (0.02)	0.65 (0.013)	0.648 (0.013)	0.648 (0.013)
Less precise (private) signal	0.258 (0.011)	0.233 (0.018)	0.258 (0.011)	0.233 (0.018)	0.258 (0.007)	0.234 (0.018)	0.258 (0.011)	0.233 (0.018)	0.258 (0.011)	0.259 (0.012)	0.259 (0.012)
Center bias	0.085 (0.009)	0.087 (0.011)	0.085 (0.007)	0.087 (0.010)	0.085 (0.008)	0.086 (0.010)	0.085 (0.007)	0.087 (0.010)	0.085 (0.009)	0.107 (0.008)	0.084 (0.006)
Private/Public Signal:											
More precise (private) signal	0.578 (0.014)	0.591 (0.02)	0.578 (0.014)	0.591 (0.02)	0.578 (0.007)	0.59 (0.02)	0.578 (0.014)	0.591 (0.02)	0.578 (0.014)	0.578 (0.014)	0.578 (0.014)
Less precise (public) signal	0.343 (0.015)	0.326 (0.023)	0.343 (0.015)	0.326 (0.023)	0.343 (0.007)	0.326 (0.023)	0.343 (0.015)	0.326 (0.023)	0.343 (0.014)	0.343 (0.015)	0.343 (0.015)
Center bias	0.086 (0.009)	0.092 (0.011)	0.086 (0.007)	0.092 (0.010)	0.086 (0.007)	0.092 (0.011)	0.086 (0.007)	0.092 (0.010)	0.086 (0.009)	0.087 (0.009)	0.086 (0.007)
Individual Random Effects	No	No	No	No	Yes	Yes	Yes	Yes	No	No	No
Individual Fixed Effects	No	No	No	No	No	No	No	No	Yes	No	Yes
Individual Fixed Effects by Treatment	No	No	No	No	No	No	No	No	No	Yes	No
Demographic variables	No	Yes	No	Yes	No	Yes	No	Yes	No	No	No
Separate regressions by treatment	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes
R ²	0.96	0.96	0.96	0.96	0.96	0.96	N/A	N/A	0.96	0.96	N/A
N	6000	6000	6000	6000	6000	6000	6000	6000	6000	6000	6000

Notes: This table presents estimates of individual choice behavior as a function of the signals provided. All coefficients are significant at the 1% confidence level except the coefficient on the less precise signal in the baseline treatment; this signal was not displayed to participants in this treatment. The coefficient on the more precise signal in the Two Private Signals treatment is significantly greater than that of the Private/Public Signal treatment at the 5% confidence level in all specifications; likewise, the coefficient on the less precise signal is greater when it is public than when it is private in all specifications at the 5% confidence level. R² not reported for separate regressions since it takes on multiple values.

"Individual Fixed Effects by Treatment" includes categorical variables for each participant in each treatment -- that is, 180 fixed effects (60 participants x 3 treatments).

Demographic variables include gender, categorical variables indicating a major in economics or mathematics, the number of college-level mathematics courses taken and the number of college-level economics courses taken. These variables are interacted with treatment dummies and signal realizations.

Coefficients for center bias are simply the intercept coefficient for the regression transformed into a percentage by dividing by 5000, the expected value of the state ex ante. If participants are choosing a convex combination of their signals, then this coefficient will be zero; if participants are choosing the center of the choice space regardless of their signal realizations, then this coefficient will be 1.

Standard errors are clustered by participant.

	Specification		
	A	B	C
One Signal (Baseline)	824,527 (33,988)	824,527 (33,988)	824,527 (20,424)
Two Private Signals:	-57,134 (35,465)	-57,134 (35,465)	-57,134 (35,465)
Private/Public Signal	98,386 (30,033)	98,386 (30,033)	98,386 (30,033)
Session Random Effects	No	Yes	No
Session Fixed Effects	No	No	Yes
P-value of test: Baseline = Two Private	0.182	0.107	0.182
P-value of test: Baseline = Private/Public	0.031	0.001	0.031
R ²	0.02	0.02	0.02
N	500	500	500

Notes: This table presents estimates of aggregated welfare outcomes as a function of the treatment. The aggregation consists of combining the welfare outcomes of all participants in a given round for each session.

Standard errors are clustered at the session level.

	One signal	Two private signals	Private/Public signal
Actual welfare	810,130	764,246	910,009
Socially optimal signal weights, no center bias	611,567	501,762	553,632
Equilibrium play, optimal center bias	581,674	500,784	614,091
Noisy equilibrium play, optimal center bias	813,776	763,588	920,296

Notes: This table accounts for the sources of the welfare losses of the participants. For each treatment, the first row displays the average welfare loss of the participants. The second row displays the welfare loss from socially optimal signal weighting under an improper prior. The third row displays the welfare loss from equilibrium play with a proper prior. The fourth row displays the welfare loss from noisy equilibrium play with a proper prior. The noise specification assumes that choice is perturbed by a zero mean normal random variable with standard deviation equal to 500 units.