

Did You See What I Saw? Belief Revision when others' Information Source is Unknown *

Anthony M. Kwasnica[†]

Raisa Velthuis[‡]

Jared Williams[§]

August 22, 2014

Abstract

We conduct a series of forecasting experiments to examine how people update their beliefs upon observing information about others' beliefs and information when there is uncertainty about others' information source. We find that uncertainty over others' information source hinders information aggregation, particularly when people observe qualitatively similar information. Moreover, people tend to be overly skeptical of others' forecasts. Our findings suggest that people have difficulty recognizing that others can see information that is qualitatively similar, but distinct, from their own information, and that people underestimate others' forecast accuracy.

Keywords: **Bayesian Updating, Information Aggregation, Forecasting, Rational Expectations, Financial Analysts**

JEL Classification: G02, C91, D82

*We thank Robert Bloomfield, Peter Bossaerts, Joey Engelberg, Laura Field, Dave Haushalter, Steve Huddart, Shimon Kogan, Michelle Lowry, Hong Qu, and seminar participants at the University of Maryland, the University of Michigan, and Penn State University for their comments and suggestions. We also appreciate the financial support of the LEMA laboratory and Risk Management Department. All errors are our own.

[†]Smeal College of Business, Penn State University. 332 Business Building, University Park, PA. E-mail: kwasnica@psu.edu

[‡]Department of Finance, Smeal College of Business, Penn State University. 338A Business Building, University Park, PA. E-mail: velthuis@psu.edu

[§]Department of Finance, Smeal College of Business, Penn State University. 344 Business Building, University Park, PA. E-mail: jmw52@psu.edu

1 Introduction

Economists have long recognized that the information people observe can affect outcomes in financial markets and the macroeconomy. Much of this literature takes people’s information as exogenous; agents are simply endowed with information, or they passively observe informative signals (Admati, 1985). While there has been research on people’s decisions about acquiring information, most of this literature models *how much* information people will choose to acquire, or whether people will acquire any information at all (Grossman and Stiglitz, 1976). Recently, researchers have begun to model situations in which agents can choose *which* information to observe when there are multiple information sources. For example, Nieuwerburgh and Veldkamp (2009) develop a model in which investors are endowed with some information about their own country’s stocks and less information about foreign stocks. Investors then choose whether to learn more about domestic stocks or foreign ones. They show that in equilibrium, it is optimal for domestic investors to learn about their domestic stocks, and that their holdings will be concentrated in domestic stocks, thus offering a potential explanation for the home bias puzzle. Other examples of recent papers that consider people’s choice of which information to observe include Hellwig and Veldkamp (2009), Nieuwerburgh and Veldkamp (2010), and Myatt and Wallace (2012).

Though researchers have begun to rigorously model environments in which people choose what to learn about, little is known about how people actually learn from others’ actions when there is uncertainty over who has observed what information. The objective of this study is to help fill that void by examining the impact of information source uncertainty on decision making in a controlled laboratory environment. Outside of the laboratory, it is difficult to isolate the impact of uncertainty about others’ information sources; differences in preferences, unknown objective functions, multiplicity of equilibria, and behavioral biases can all affect people’s behavior.

Our experiments are designed to capture a pervasive aspect of many real-world settings in which people can update their beliefs upon observing information about others’ beliefs. Macroeconomic forecasters can observe other forecasters’ inflation and GDP growth forecasts; financial analysts can observe other analysts’ earnings forecasts; traders in financial markets can observe prices, which convey information about others’ beliefs about the stock’s value; meteorologists can observe others’ weather projections; sports bettors can observe pundits’ picks, etc. Given the vast amounts of available information and the limited time and resources to process the information, uncertainty about others’ information sources is undoubtedly present in all of these environments. Macroeconomic forecasters cannot analyze all public information about the economy; financial analysts and traders cannot consume all public information about a firm; sports bettors cannot monitor the status of every team or player in real-time. Hence, participants in these markets must restrict their attention to a subset of the available information, and they cannot know with certainty what information others have chosen to observe.

A person’s beliefs about others’ information sources should affect how he updates his beliefs upon observing others’ actions. Consider a macroeconomic forecaster who believes GDP growth will be high and observes that other forecasters are projecting high GDP growth. The incremental

informativeness of the others’ forecasts will depend on whether or not they are relying on the same sources of information as him—if everyone is optimistic for the same reason, then there is little incremental value in others’ opinions, but if everyone is optimistic based on different information, then he should become even more optimistic upon observing that others are optimistic. Our experimental design allows us to analyze how people update their beliefs in such situations, and thus, how information gets aggregated in these scenarios.

The context of our experiment is a financial forecasting exercise. In all three of our treatments, subjects are randomly paired with a partner and they are asked to provide two forecasts—an initial forecast and a revised forecast—for a hypothetical firm’s earnings. Prior to issuing their initial forecast, subjects observe either the firm’s revenue or the firm’s cost. Prior to issuing their revised forecast, subjects are given information about their partner—either the partner’s initial forecast, or whether the partner saw “good news” about the firm or “bad news.” While the two sources of information (revenue and cost) are different, they are equally informative. Subjects can observe good news about revenue (revenue is high) or cost (cost is low), and both types of good news result in the same conditional expectation for the firm’s earnings.¹ Similarly, subjects can observe bad news about revenue (revenue is low) or cost (cost is high), and both types of bad news result in the same conditional expectation for the firm’s earnings. Given this design, there are instances where information source uncertainty persists, and other instances where it is eliminated. More specifically, when a subject learns that he and his partner both observed good news, he cannot be sure whether they observed the same information, but if he observes good news and is informed that his partner observed bad news, it is possible for him to infer that their information sources must differ. By comparing behavior across our three treatments, each of which has the same unique symmetric Nash equilibrium prediction, we can examine how uncertainty about others’ information sources affects individual behavior and information aggregation.

We find that subjects tend to place too little weight on the information they observe about their partner when the partner’s information source remains uncertain. The fact that this effect is present in all our treatments allows us to eliminate various potential pathways by which such behavior could be rationalized. This bias, however, depends critically on the information outcomes in the experiments. When subjects observe information that is qualitatively different, e.g., a subject sees good news and his partner observes bad news, subject behavior is largely consistent with theory. When subjects observe information that is qualitatively similar, e.g., the subject and his partner both observe good news or bad news so that information source uncertainty remains, behavior is highly biased away from the equilibrium prediction in the direction of the subject’s own private information. In fact, the bias is so strong that many subjects act as if the information they are provided about their partner’s private information is completely uninformative. We believe that the persistence of such behavior across treatments and when information source uncertainty remains highlight a persistent behavioral bias that, to our knowledge, has not been previously documented.

Our paper is organized as follows. We describe our experiment and derive the equilibrium

¹Formally, $\mathbf{E}[\text{earnings} \mid \text{revenue is high}] = \mathbf{E}[\text{earnings} \mid \text{cost is low}]$.

predictions in Section 2. In Section 3, we describe our results, and in Section 4, we consider possible explanations for our findings. Section 5 concludes with a discussion of some potential implications of our findings.

2 Experimental Environment

We begin by describing the experimental environment and procedures in our baseline treatment (Treatment A).

All subjects participated in a forecasting decision process for 30 independent periods. In each period, subjects were randomly matched with a partner. Each pair of matched subjects faced the same task of forecasting a hypothetical firm’s earnings, X . The firm’s earnings were determined by its uncertain revenue, R , cost, C , and a random noise component, ϵ , as follows:

$$X = R - C + \epsilon.$$

Revenue was distributed according to $\Pr(R = \$20) = 0.5$ and $\Pr(R = \$10) = 0.5$, cost was distributed according to $\Pr(C = \$0) = 0.5$ and $\Pr(C = \$10) = 0.5$, and ϵ was distributed uniformly on $[-\$1, \$1]$. Revenue, cost, and ϵ were independently distributed across pairs and across periods, so each firm’s unconditional expected earnings was \$10:

$$\mathbf{E}[X] = \mathbf{E}[R] - \mathbf{E}[C] + \mathbf{E}[\epsilon] = \$15 - \$5 + \$0 = \$10.$$

Each independent experimental period proceeded in a series of two distinct stages: an initial forecast stage and a revised forecast stage. Subjects were compensated based on their accuracy in the randomly chosen compensation periods:

$$\text{Compensation} = \$6 - 0.04(f - X)^2, \tag{1}$$

where f is the subject’s forecast. Since the mean is the best predictor under squared loss, subjects could maximize their expected compensation by issuing forecasts equal to their conditional expectations of X .²

Subjects first chose whether to observe the firm’s revenue or cost.³ Subjects were then informed of their chosen components of the earnings realization in the following manner: they were told that they had observed either “good news” or “bad news,” and they were told the exact realization associated with such news. In the case of revenue, good news was associated with high revenue

²Since the object of this study was not compliance with a proper scoring rule, the experimental instructions provided guidance in this regard. Subjects were informed that they would maximize their expected compensation by issuing forecasts equal to what they think earnings would be on average given all the information they have observed. We formally verify the incentive properties of the scoring rule in Appendix A.

³The order of the buttons for selecting either revenue or cost was randomized on the computer screen to avoid potential issues with biases for subjects to select the top button. See Section 3.1 for a description of the frequency of choice of revenue or cost realizations.

($R = \$20$), and bad news was associated with low revenue ($R = \$10$). For cost, good news was associated with low cost ($C = \$0$), and bad news was associated with high cost ($C = \$10$). After observing the realization of the earnings component, subjects issued their initial forecasts. Given the information structure, the expectation of earnings in the initial stage depends only on whether the subject has observed good news or bad news (and not on whether revenue or cost were observed):

$$\mathbf{E}[X|\text{good news}] = \$15, \text{ and} \quad (2)$$

$$\mathbf{E}[X|\text{bad news}] = \$5. \quad (3)$$

After issuing an initial forecast, f_0 , the subject was shown the initial forecast of their randomly chosen partner but *not* whether the partner chose to observe revenue or cost. Each subject was then required to issue a revised forecast, f_R .⁴

At the end of each forecast period, subjects were informed of the actual earnings realization and their potential compensation (given by (1)) from both the initial and revised forecasting stages. They were not informed of the compensation, earnings component observed, or identity of their partner. Subjects were then randomly rematched with a partner and began the next period with a new, independent earnings realization. This proceeded for 30 periods, and the duration of the experiment was common knowledge among the subjects.

At the end of the 30 periods, one period for each subject was randomly chosen for compensation based upon their initial forecast for that period, and one period for each subject was random chosen for compensation based upon their revised forecast. We chose to pay for one random period for each forecasting stage to minimize potential period choice dependencies and the possibility of hedging across stages.

All sessions were conducted at the Pennsylvania State University in the Laboratory for Economics Management and Auctions. All experiments were conducted utilizing the z-Tree software (Fischbacher, 2007); a copy of the instructions provided to the subjects at the beginning of the experiment is included in Appendix B.

2.1 Equilibrium Prediction

It is easily verified that there exists a unique symmetric Nash equilibrium in our baseline treatment. In equilibrium, subjects randomly choose which component to observe (each with 50% probability), and their initial forecasts and revised forecasts are equal to the firm's expected earnings conditional on the information they have observed.

Deriving the optimal initial forecast is a straightforward task. Expected revenue is 15, and expected cost is 5. If a subject observes that revenue is high ($R = 20$) or that cost is low ($C = 0$), the firm's expected earnings conditional on the subject's information is 15 ($20 - 5 = 15 - 0 = 15$), and if a subject observes that revenue is low ($R = 10$) or that cost is high ($C = 10$), the firm's

⁴Subjects were not required to change their forecast but were required to enter a new (potentially similar) forecast number.

expected earnings conditional on the subject's information is 5 ($10 - 5 = 15 - 10 = 5$). It follows that a subject's optimal initial forecast perfectly reveals whether the subject observed good news or bad news but not necessarily whether the subject observed revenue or cost. More formally,

$$f_0^{\text{NE}} = \begin{cases} 15 & \text{if the subject observes good news } (R = 20 \text{ or } C = 0) \\ 5 & \text{if the subject observes bad news } (R = 10 \text{ or } C = 10) \end{cases} \quad (4)$$

The optimal forecast in the revised forecast stage is straightforward if the partner's qualitative information is dissimilar to the subject's, e.g., if the subject observes good news and the partner observes bad news. In this case, it is clear that the expected earnings conditional on the subject's information at the revised forecast stage is 10. The more interesting case arises when the subject and his partner observe qualitatively similar information. For example, consider the scenario in which the subject observes that revenue is high (good news) and he learns that his partner also observed good news. In this case, the subject cannot know whether his partner observed that revenue is high or that cost is low. The subject's optimal revised forecast depends on the relative likelihood of these two events. Let p denote the ex ante probability that the partner chose to observe revenue. It is straightforward to verify that

$$\Pr(\text{partner observed high revenue} \mid \text{subject and partner observe good news}) = \frac{2p}{1+p},$$

from which it follows that

$$\Pr(\text{partner observed low cost} \mid \text{subject and partner observe good news}) = \frac{1-p}{1+p}.$$

Hence, the subject's optimal revised forecast after learning that his partner also observed good news is given by

$$\begin{aligned} f_R^{\text{NE}} &= \mathbf{E}[\text{earnings} \mid \text{both subjects observe good news}] \\ &= \left(\frac{2p}{1+p} \right) 15 + \left(\frac{1-p}{1+p} \right) 20. \end{aligned} \quad (5)$$

To see why subjects must randomize over their information source in a symmetric Nash equilibrium, simply notice that it is in a subject's best interest to choose to observe the opposite piece of information as his partner, as this will maximize the likelihood that the subject and his partner observe qualitatively opposite news (one good and one bad), which is the scenario that minimizes the forecaster's expected mean squared forecast error in the revised forecast stage. Hence, if others' probability of observing revenue is $p > 0.5$, then a subject's expected compensation in the revised forecast stage is strictly higher if he chooses to observe cost than if he observes revenue. The $p < 0.5$ case is symmetric, so in any symmetric Nash equilibrium, $p = 0.5$ for all subjects. Therefore, it

follows from (5) that

$$f_R^{\text{NE}} = \begin{cases} 16\frac{2}{3} & \text{if the subject and his partner both observe good news} \\ 10 & \text{if the subject observes good news and his partner observes bad news (or vice versa)} \\ 3\frac{1}{3} & \text{if the subject and his partner both observe bad news.} \end{cases} \quad (6)$$

2.2 Additional Treatments

There are several reasons that have been identified in the literature as to why subject behavior might deviate from the equilibrium prediction in our baseline model. Our two other treatments allow us to examine whether these biases drive non-equilibrium behavior in our experiments.

Subjects' revised forecasts depend upon the information contained in the initial forecast of their randomly chosen partner. If their partner fails to issue an optimal (equilibrium) forecast, it may cause the subject to rationally discount this information. For example, Anderson and Holt (1997) find that in an information cascades experiment subjects may rely more on their own private signal due to the fact that initial subject choices might have been error prone. In order to control for this, we conducted a treatment (Treatment B) where the *qualitative* nature of the partner's information was revealed rather than their initial forecast. Subjects were informed of whether their partner observed good news (which can reflect either high revenue or low cost) or bad news (low revenue or high cost). In addition, in order to be as similar to Treatment A as possible, subjects were informed of the initial forecast that their partner would have issued if he were optimizing his expected compensation given the information he had observed. Since this qualitative information provides the same information as the equilibrium initial forecast in Treatment A, the revised forecast equilibrium prediction is identical.

The equilibrium in the initial information source selection stage prescribes that subjects play a mixed strategy. There is a substantial evidence that experimental subjects may fail to play mixed strategy equilibria (Erev and Roth, 1998). In Treatments A and B, the partner's choice of earnings component is unknown at the time of the revised forecast, since the primary interest of this study is on information source uncertainty. A subject's optimal revised forecast, however, relies upon his beliefs about what information his partner chose to observe (the p term in the above calculations). Any failure to play $p = 0.5$ or inconsistency between subjects' beliefs and these actual frequencies will result in revised forecasts that differ from the equilibrium prediction. In order to control for this potential source of deviation, we conducted a treatment (Treatment C) where, in addition to the change from Treatment B, we also randomly and independently assigned subjects to observe either revenue or cost, each with 50% probability. This probability was common knowledge to the subjects and their assignment to observe either cost or revenue was accomplished through an initial stage that was identical to the original initial stage except that the buttons allowing a choice of cost or revenue were not active. This treatment should be most conducive to efficient information aggregation and equilibrium play in the revised forecast stage since it eliminates both sources of strategic uncertainty in the earlier stages of the game. Table 1 contains a brief overview of the three

treatments that comprise our experiment.

[INSERT TABLE 1 HERE]

Although the treatments differ in their design, they all have the same equilibrium prediction given by (4) and (6). We formally verify this in Appendix A. When deriving the equilibrium, we assume that subjects are risk neutral. Risk aversion would not affect a subject’s optimal initial forecast or his optimal revised forecast when his news is qualitatively dissimilar from his partner’s, but it would affect the optimal revised forecast when his news is qualitatively similar to his partner’s—risk aversion would cause the optimal revised forecast in this scenario to be closer to 10.⁵ However, as long as a subject is not infinitely risk averse, his revised forecasts should differ from his initial forecast.

3 Results

We begin by discussing the experimental results in each of the three stages and provide a comparison with the equilibrium prediction. In Section 4, we discuss various models that might explain our data.

3.1 Choice of Information Source

In equilibrium, subjects should randomly choose to observe revenue or cost, each with 50% probability. In Figure 1, we plot the number of subjects (out of 60 subjects) who chose to observe revenue in n out of the 30 periods.

[INSERT FIGURE 1 HERE]

There is mixed evidence in support of the prediction that subjects randomize on their choice of information source in Treatments A and B. Only 20% of subjects choose to observe the same earnings component every period, whereas 25% of subjects choose to observe revenue in at least 10, but no more than 20, of the 30 periods. Just under half (40%) of subjects choose to observe revenue in at least 5, but no more than 25, of the 30 periods. In aggregate, subjects choose to observe revenue 57.1% of the time, which is not statistically significantly different than the 50% equilibrium prediction ($t = 1.50$).

⁵In the case where both subjects observe good news, the optimal revised forecast would be between 15 and $16\frac{2}{3}$, and in the case where both subjects observe bad news, the optimal revised forecast would be between $3\frac{1}{3}$ and 5.

3.2 Initial Forecasts

Recall from Section 2.1 that a subject’s optimal initial forecast is 15 (5) if he observes good (bad) news. Subjects generally issue initial forecasts that are close to the optimal forecast. We plot the distribution of subjects’ initial forecasts as a function of the four possible pieces of private information (high revenue, low revenue, high cost, and low cost) in Figure 2.

[INSERT FIGURE 2 HERE]

We report statistics on subjects’ initial forecasts, by Treatment, in Table 2. In all treatments, the majority of initial forecasts are *exactly* equal to the optimal initial forecast: 54% (63%, 52%) of subjects’ initial forecasts are optimal in Treatment A (B, C). These are reported in Panel A of Table 2. The average distance between f_0 and f_0^{NE} is 1.99 (1.80, 2.17) in Treatment A (B, C). Viewing f_0 as an estimate for f_0^{NE} , the mean squared error of f_0 is 11.72 (12.42, 12.97) in Treatment A (B, C). We report the t -statistics for differences in means across treatments in Panel B of Table 2. None of the differences are statistically significant, suggesting that there are no significant differences in subjects’ initial forecasts across treatments.⁶

[INSERT TABLE 2 HERE]

3.3 Revised Forecasts

In all three treatments, subjects receive relevant information between the time of their initial forecast and their revised forecast. Figure 2 and Table 2 reveal that subjects’ initial forecasts are generally accurate, so in Treatment A, subjects receive reasonably precise signals about whether their partners observed good news or bad news, and in Treatments B and C, subjects are directly informed whether their partner observed good news or bad news. Indeed, subjects should update their beliefs at least minimally between their initial forecast and their revised forecast, so their revised forecast should differ from their initial forecast. (See Proposition 3 in Appendix A.) However, we find that the majority (63%) of subjects’ revised forecasts in Treatment A are *exactly* equal to the subjects’ initial forecasts. In Treatments B and C, we directly inform subjects of their partner’s news quality (good news or bad news), so *their revised forecast should always differ from the initial forecast*,

⁶Much of our analysis on subjects’ revised forecasts is based on the differences in subjects’ behavior across treatments. Since the collection of subjects is different in each treatment, it is important that there are no significant differences in subjects’ rationality across the three treatments. To address this, we compute t -statistics for differences in means across treatments. Throughout the paper, we conduct our statistical analysis by treating the subject as the unit of observation. That is, for each variable, we compute the mean of the variable for each subject, and then compare the distribution of subject-level means across treatments, and our number of observations in each treatment is the number of subjects in that treatment.

regardless of what information the subject observes. However, we find that a large minority (42% and 40%) of revised forecasts are *exactly* equal to the initial forecast. These statistics are reported in the second column of Panel A in Table 3.

[INSERT TABLE 3 HERE]

As mentioned previously, subjects receive one of two types of information in the revised forecast stage. A subject observes qualitatively similar information if he and his partner both observed good (or bad) news; in this case, the information source of the partner remains uncertain. A subject observes qualitatively dissimilar information if he and his partner observe opposite types of news (good and bad or bad and good); in this case, all information source uncertainty is resolved since the same information source cannot provide different qualities of news in our setting. If information source uncertainty is an important factor in subject decision making, we hypothesize that behavior may systematically vary between these two situations. To examine this, we repeat the analysis from Table 3, except that for each treatment we partition the data based on whether the subject and partner observe qualitatively dissimilar private information (one good, one bad) or qualitatively similar private information (both good or both bad) in Tables 4 and 5, respectively.

[INSERT TABLES 4 AND 5 HERE]

When subjects and their partner observe qualitatively dissimilar information, revised forecasts rarely equal initial forecasts, especially in Treatments B and C (12% and 13%; second column of Table 4, Panel A). In contrast, when subjects observe qualitatively similar information, subjects tend to issue revised forecasts that *exactly* equal their initial forecasts, even in Treatments B and C (52% and 49%; second column of Table 5, Panel A).

Another measure of revision activity is the absolute distance between subjects' initial forecasts and revised forecasts, $|f_R - f_0|$. The equilibrium initial and revised forecasts are the same across all three treatments, so these differences should be the same across treatments according to standard theory. In Treatment A (B, C), the average distance between subjects' revised forecast and their initial forecast is 1.68 (2.89, 2.91), as shown in the third column of Panel A in Table 3. We report t -statistics for the differences in means across treatments in Panel B. Subjects revise their forecasts significantly more in Treatments B and C than in Treatment A ($t = 2.79$ and 3.12), contrary to the equilibrium prediction. Given the fact that many subjects' initial forecasts are suboptimal (Figure 2 and Table 2), it is not surprising that subjects respond less to their partner's forecast (Treatment A) than to the qualitative information that their partner observed (Treatments B and C); we discuss whether the magnitude of the difference in average forecast revisions across Treatments A and B/C can be rationalized in Section 4.⁷

⁷The differences in revision distance across Treatments B and C is insignificant ($t = 0.04$).

We report the average distance between the revised forecast and the equilibrium prediction, $|f_R - f_R^{\text{NE}}|$, in the fourth column of Tables 3-5.⁸ When information source uncertainty is resolved (Table 4), revised forecasts are closer to the equilibrium prediction than when information source uncertainty persists (Table 5). Comparing the prediction errors across treatments, we see that revised forecasts are closer to the equilibrium prediction when subjects are informed whether their partner saw good news or bad news (Treatments B and C) than when they are shown their partner’s initial forecast (Treatment A). This is true unconditionally (Table 3), when information source uncertainty is resolved (Table 4), and when information source uncertainty persists (Table 5).

Overall, our findings suggest that information source uncertainty has a significant effect on subject behavior. When it is present, subjects act as though new information is incrementally uninformative even though it is in fact incrementally informative; when the uncertainty is not present, subjects act as though they realize the information is incrementally informative. More generally, subject behavior aligns more closely with the equilibrium prediction when information source uncertainty is not present.

4 Discussion

Recall from Section 3 that subjects tend to issue revised forecasts that exactly equal their initial forecasts, and this tendency is strongest when their partner observes private information that is qualitatively similar to their own private information. More generally, subjects’ revised forecasts align more closely to the equilibrium when information source uncertainty is resolved than when it persists. In this section, we will discuss possible explanations for these findings based on previous research in economics and psychology.

4.1 Overconfidence

Overconfidence is a psychological bias that has received much attention in economics. Within financial economics, there is a large literature in both asset pricing (Daniel, Hirshleifer, and Subrahmanyam (1998), Odean (1998), and Scheinkman and Xiong (2003)) and corporate finance (Malmendier and Tate (2005), Malmendier, Tate, and Yan (2011) and Gervais, Heaton, and Odean (2011)). The “better than average effect” is a form of overconfidence in which the majority of people believe they are better than average at a given task. For example, Svenson (1981) documents that 82% of American drivers believe they are in the top 30% of drivers.⁹ See Odean (1998) for a more thorough review of the psychological evidence supporting overconfidence and the better than average effect.

Is subjects’ tendency of issuing revised forecasts that exactly equal their initial forecasts related to overconfidence? It is apparent from Figure 2 and Table 2 that the majority of subjects’ initial forecasts contain information that should be useful for other subjects when they issue revised

⁸Recall that f_R^{NE} is defined in (6).

⁹Of course, this finding is not necessarily the result of a cognitive bias; heterogeneity in opinions about what constitutes “proper driving” could also cause such a result.

forecasts: in all three treatments and for all four possible pieces of private information, the majority of subjects' initial forecasts are exactly equal to the optimal forecast. However, it is also clear that though the majority of subjects issue optimal initial forecasts, many of them do not. This complicates subjects' task in the revised forecast stage of Treatment A, as subjects do not know how rational their partner is. Hence, Treatment A captures a pervasive aspect of information aggregation in practice: people do not know whether other people are rationally processing the information that they have observed. If subjects exhibit a better than average effect, they might underestimate their partner's forecasting ability and largely ignore their partner's initial forecast. This might explain why subjects' revised forecasts are often equal to their initial forecasts in Treatment A.

To examine this, we define a measure (OPI_{NE}) to indicate whether subjects overweight their private information relative to the Nash equilibrium prediction. Suppose a subject and his partner both observe good news. If they both randomize over their information source (revenue versus cost), which is consistent with the equilibrium prediction that we are unable to reject (Section 3.1), then the optimal revised forecast is $16\frac{2}{3}$. A revised forecast that is less than $16\frac{2}{3}$ deviates from the predicted revised forecast in the direction of the subject's private information (i.e., 15), whereas if the revised forecast is greater than $16\frac{2}{3}$, the subject "overshot" the optimal forecast. That is, relative to the equilibrium prediction, subjects can be viewed as overweighting their private information if they issue revised forecasts less than $16\frac{2}{3}$, and underweighting their private information (i.e., straying too far from their optimal initial forecast) if they issue revised forecasts that are greater than $16\frac{2}{3}$.

Conversely, if a subject observes good news and his partner observes bad news, then the optimal revised forecast is 10. A revised forecast that is greater than 10 deviates from the predicted revised forecast in the direction of the subject's private information (i.e., 15), whereas if the revised forecast is less than 10, the subject "overshot" the optimal forecast. That is, relative to the Nash equilibrium prediction, subjects can be viewed as overweighting their private information if they issue revised forecasts greater than 10, and underweighting their private information (i.e., straying too far from their optimal initial forecast) if they issue revised forecasts that are less than 10. The cases in which subjects observe bad private information about the firm are symmetric.

Accordingly, we define the variable OPI_{NE} to determine whether a revised forecast overweightes the subject's private information relative to the Nash equilibrium prediction:

$$OPI_{NE} = \begin{cases} 16\frac{2}{3} - f_R & \text{if the subject and his partner both observe good news} \\ f_R - 10 & \text{if the subject observes good news and his partner observes bad news} \\ 10 - f_R & \text{if the subject observes bad news and his partner observes good news} \\ f_R - 3\frac{1}{3} & \text{if the subject and his partner both observe bad news} \end{cases} \quad (7)$$

Positive values of OPI_{NE} indicate that subjects overweight their private information relative to the equilibrium, whereas negative values indicate that they underweight their private information. If subjects' revised forecasts are equal to the equilibrium prediction plus noise, we should expect OPI_{NE} to have a mean that is not statistically different than 0.

We report the average values of OPI_{NE} for each of the three treatments in the fifth column of

Panel A in Table 3. The average value is 2.09 in Treatment A, which is highly statistically significant ($t = 8.87$). In Treatments B and C, the average value is smaller (0.51 and 0.68), but still highly statistically significant ($t = 2.92$ and 2.99). In Tables 4 and 5, we report the average values of OPI_{NE} separately for the cases in which information source uncertainty is resolved (Table 4) and when it remains (Table 5). When information uncertainty is resolved, the average value of OPI_{NE} is 0.09 and 0.10 in Treatments B and C, neither of which is significantly different than 0 ($t = 0.26$ and 0.67). Hence, revised forecasts are not systematically biased away from the equilibrium prediction when there is no information source uncertainty. When information uncertainty remains, the average value of OPI_{NE} is 0.65 and 0.87 in Treatments B and C, each of which is significantly positive ($t = 3.24$ and 3.22). Together, these findings indicate that revised forecasts systematically deviate from the equilibrium prediction in the direction of the expected earnings conditional on the subject’s private information, and that this deviation is driven by the instances when information source uncertainty persists.

The fact that OPI_{NE} is positive in Treatment A can be explained by the fact that subjects’ initial forecasts are quite noisy. Figure 2 and Table 2 show that many subjects’ initial forecasts are suboptimal, so it is rational for subjects to discount their partner’s initial forecasts. Moreover, even if subjects were certain they could infer their partner’s qualitative information from their partner’s initial forecast, we would still find a positive value of OPI_{NE} if subjects made mistakes when inferring their partner’s qualitative information from the partner’s initial forecast. Hence, positive values of OPI_{NE} in Treatment A do not imply that subjects are “biased” in weighting their private information more than the equilibrium prediction. It is far less clear why OPI_{NE} is significantly positive in Treatments B and C. Unlike Treatment A, it is not optimal for subjects to bias their revised forecasts away from the equilibrium prediction in these treatments; since subjects are directly informed whether their partner observed good news or bad news about the firm, their partner’s ability is irrelevant in these treatments. Moreover, since subjects’ private information is completely unrelated to their ability, there is no reason for subjects to exhibit overconfidence and overweight their private information vis-à-vis the information about their partners’ private information.

To determine whether subjects are optimally skeptical of their partner’s initial forecast in Treatment A, we propose a simple model that allows us to account for the potentially error prone nature of partners’ initial forecasts. While Anderson and Holt (1997) have examined similar models in the context of information cascades, the fact that a subject’s initial forecast is a continuous variable makes formulation of a quantal response type model analytically difficult. This analysis is restricted to the subset of our Treatment A data in which the partner’s initial forecast is consistent with optimal behavior.¹⁰ Formally, let A' denote the subset of our Treatment A data in which the subject’s partner issues an initial forecast equal to either 5 or 15. In other words, A' consists of our observations in which the subject’s partner *appears* (from the subject’s perspective) to issue an

¹⁰While it is unclear how a subject should interpret his partner’s initial forecast when it does not equal 5 or 15, it is clear that if his partner issues an initial forecast of 15 (5), then it is likely that his partner observed good (bad) news and issued an optimal initial forecast (recall Figure 2).

optimal initial forecast. With slight abuse of notation, we refer to this as “Treatment A’.”

Suppose a subject observes that revenue is high ($R = 20$) and learns that his partner’s initial forecast is 15. The forecast of 15 would be an optimal forecast for the partner only if the partner saw good news (high revenue or low cost), so the subject should infer that it is likely that his partner observed good news. However, it is also possible that his partner observed bad news and issued a suboptimal forecast. For example, the partner could be irrational, confused about the experiment, or have anti-social preferences and derive pleasure from misleading another subject.¹¹ If subjects randomize over their information source, and if the partner actually saw good news, then the subject’s optimal revised forecast is $16\frac{2}{3}$ (recall (6)). If, on the other hand, the partner had observed bad news and issued a suboptimal initial forecast, then the subject’s optimal revised forecast is 10. Taking the possibility that the partner issues suboptimal initial forecasts, the subject’s optimal revised forecast should be a weighted average of $16\frac{2}{3}$ and 10, where the weighting is based on the likelihood that the partner’s initial forecast was in fact optimal.

We determine a subject’s optimal revised forecast in Treatment A’ by estimating the proportion of initial forecasts that are consistent with Nash equilibrium, conditional on the initial forecast equaling 5 or 15. That is, we estimate $\rho \equiv \mathbf{Pr}(f_0 = f_0^{\text{NE}} | f_0 \in \{5, 15\})$. Note that $1 - \rho$ is the empirical probability that an apparently optimal forecast was actually not optimal, e.g., was 5 when it should have been 15 (or vice versa). Across the entire sample of all three treatments, ρ is 97.7%. Estimating ρ separately for each treatment, ρ is 95.2% (98.2%, 98.5%) in Treatment A (B, C).

Let f_R^{opt} be the optimal forecast accounting for this possible error. For observations in Treatment A in which the partner’s forecast is not equal to 5 or 15, f_R^{opt} is not clear so we restrict our attention to observations in A’. We define f_R^{opt} as

$$f_R^{\text{opt}} = \begin{cases} \rho 16\frac{2}{3} + (1 - \rho)10 & \text{if the subject observes good news and his partner's } f_0 = 15 \\ \rho 10 + (1 - \rho)16\frac{2}{3} & \text{if the subject observes good news and his partner's } f_0 = 5 \\ \rho 10 + (1 - \rho)3\frac{1}{3} & \text{if the subject observes bad news and his partner's } f_0 = 15 \\ \rho 3\frac{1}{3} + (1 - \rho)10 & \text{if the subject observes bad news and his partner's } f_0 = 5. \end{cases} \quad (8)$$

Since $\rho \approx 0.95$ in Treatment A’, we compute f_R^{opt} for each observation in A’ by setting $\rho = 0.95$, and (8) becomes

$$f_R^{\text{opt}} = \begin{cases} 16\frac{1}{3} & \text{if the subject observes good news and his partner's } f_0 = 15 \\ 10\frac{1}{3} & \text{if the subject observes good news and his partner's } f_0 = 5 \\ 9\frac{2}{3} & \text{if the subject observes bad news and his partner's } f_0 = 15 \\ 3\frac{2}{3} & \text{if the subject observes bad news and his partner's } f_0 = 5. \end{cases} \quad (9)$$

In Treatments B and C, subjects are informed whether their partner saw good news or bad news, so there is no reason for a subject to discount the qualitative information they are given about their

¹¹Of course, if the partner issued a suboptimal initial forecast, it would reduce the partner’s expected compensation since the partner has a monetary incentive to issue an optimal initial forecast.

partner's private information. We define f_R^{opt} in these treatments as

$$f_R^{\text{opt}} = f_R^{\text{NE}} \quad (\text{in Treatments B and C}). \quad (10)$$

We report the average distance between subjects' revised forecasts and their optimal revised forecast (2.58) for Treatment A' in the sixth column of Table 3, Panel A. Since $f_R^{\text{opt}} \equiv f_R^{\text{NE}}$ in Treatments B and C, we leave those entries blank in the sixth column.¹² As expected, this distance is less than the average distance between subjects' revised forecasts and the Nash equilibrium prediction for Treatment A (3.14, reported in the fourth column). In Panel B, we report t -statistics for the differences in distances across treatments. Even after accounting for the fact that subjects should partially discount their partner's forecast in A', we still see that they are further from their optimal forecast in Treatment A' than they are in Treatment B ($t = 1.73$). However, the difference between Treatment A' and C is not statistically significant ($t = 0.58$).

The variable OPI_{NE} (defined in (7)) was defined to determine whether subjects overweight their private information relative to the *equilibrium*. Since not everyone plays equilibrium strategies, we cannot use OPI_{NE} to determine whether subjects in Treatment A overweight their private information relative to the subject's *optimal* revised forecast. We define the variable OPI_{adj} to capture whether subjects overweight their private information relative to their optimal behavior. OPI_{adj} is equivalent to OPI_{NE} , except that in Treatment A', instead of comparing a subject's revised forecast to the equilibrium revised forecast, we compare the revised forecast to the subject's optimal revised forecast:

$$OPI_{\text{adj}} = \begin{cases} 16\frac{1}{3} - f_R & \text{if the subject and his partner both observe good news} \\ f_R - 10\frac{1}{3} & \text{if the subject observes good news and his partner observes bad news} \\ 9\frac{2}{3} - f_R & \text{if the subject observes bad news and his partner observes good news} \\ f_R - 3\frac{2}{3} & \text{if the subject and his partner both observe bad news} \end{cases} \quad (11)$$

We leave OPI_{adj} undefined for observations in Treatment A in which the partner's initial forecast is not equal to 5 or 15, and in Treatments B and C, we set OPI_{adj} equal to OPI_{NE} since subjects should not discount the information they are given about their subject's qualitative information in these treatments:

$$OPI_{\text{adj}} = OPI_{\text{NE}} \quad (\text{in Treatments B and C}). \quad (12)$$

The average value of OPI_{adj} in Treatment A' is 1.45, as reported in the seventh column of Table 3. The mean is highly statistically significant ($t = 4.19$), which suggests that relative to the optimal revised forecast, subjects place too much weight on their own private information and too little weight on the information contained in their partner's initial forecast. In Panel A of Tables 4 and 5, we report the average values of OPI_{adj} separately for the cases in which information source

¹²The values of $|f_R - f_R^{\text{opt}}|$ for Treatments B and C would be identical to the values of $|f_R - f_R^{\text{NE}}|$ for Treatments B and C that are reported in the fourth column.

uncertainty is resolved (Table 4) and when it remains (Table 5). When information uncertainty is resolved, the average OPI_{adj} is 0.86, which is not significantly different than 0 ($t = 1.62$), but it is significantly positive (1.61; $t = 4.37$) when information source uncertainty persists. These findings suggest that the bias of overweighting one’s private information is driven by the cases in which information source uncertainty persists.

In Panel B of Table 3, we report the t -statistics for the differences in OPI_{adj} across treatments.¹³ Recall that subjects overweight their private information in Treatments B and C (Panel A, fifth column). However, they overweight their private information significantly more (relative to the optimum) in Treatment A’ than in Treatments B or C ($t = 2.68$ and 1.87). This suggests that subjects are too skeptical of their partner’s forecast; that is, subjects underestimate the informativeness of their partner’s initial forecast. Hence, sequential forecasting introduces two inefficiencies relative to efficient information aggregation. First, if people misinterpret their private information, their forecast will have error, and this will cause errors in later forecasts that incorporate the initial forecasts. This is obvious. Our results suggest a second, less obvious, source of inefficiency: later forecasters will be overly skeptical of the earlier forecasts issued by others.

Summarizing, we find support for the better average effect in that subjects are overly skeptical of the informational content of their partner’s initial forecast (Treatment A’). However, overconfidence cannot explain why subjects overweight their private information in Treatments B and C, where subjects are directly told whether their partner observed good news or bad news. Hence, overconfidence does not fully explain our findings.

4.2 False Consensus Effect

The “false consensus effect” refers to people’s tendency to overestimate their similarity to others (Ross, Greene, and House, 1977). More specifically, people have a tendency to overestimate the likelihood that others engage in the same activities, have the same opinions and beliefs, and have the same preferences as them. See Williams (2013) for a more comprehensive review of the literature.

When subjects are given the choice of which earnings component to observe (Treatments A and B), a false consensus effect might cause subjects to naïvely assume that their partner chooses to observe the same information as them. If a subject believes his partner chooses the same information as him, he would not update his beliefs upon learning that his partner observed qualitatively similar information, because he would believe his partner saw the *same* information as him. That is, a false consensus effect could cause subjects to erroneously believe there is no incremental informational content provided in the revised forecast stage, especially when the partner’s qualitative information is consistent with the subject’s. More specifically, a false consensus effect predicts the positive OPI_{NE} and OPI_{adj} that we document in Table 3, and it also predicts the positive values of OPI_{NE} and OPI_{adj} to be driven by the cases in which the subjects’ information is qualitatively similar to his partner’s, which we document in Table 5.

¹³In Treatments B and C, $OPI_{adj} \equiv OPI_{NE}$, so we omit these tests in the seventh column, as the information would be identical to the tests reported in the fifth column.

However, the false consensus effect *cannot* explain anomalous behavior in Treatment C, where subjects know that the computer randomly (and independently) chooses each subject’s information source. In the third row of Panel B, Tables 3-5, we compute t -statistics for the differences in means of our forecast revision variables across treatments B and C. The fact that we find similar estimates of the likelihood of issuing revised forecasts that equal initial forecasts, average forecast revision distance, OPI_{NE} and OPI_{adj} in Treatments B and C suggests that the false consensus effect does not fully explain our results.

4.3 Cursedness

There are many settings in which people do not behave according to Nash equilibrium predictions. For example, in common value auctions, bidders tend to overbid (“the winner’s curse”), and in bilateral trade experiments, there is often trade when Nash equilibrium predicts there should not be any trade. Eyster and Rabin (2005) argue that this is because people underappreciate the link between others’ *information* and others’ *actions*. For example, in the context of common value auctions, the winner’s curse might arise because bidders do not take into account the fact that if they win the auction, everyone else must have bid less than them, in which case other bidders likely received low signals about the item’s value and the item probably has lower value than the winning bidder’s private signal suggested. If bidders do not take this into account, the value of the maximum bid will generally exceed the item’s value.

Eyster and Rabin (2005) extend the Nash equilibrium framework to allow for “cursedness.” In their model, χ represents the degree to which people are cursed. A value of $\chi = 0$ corresponds to the standard model, and $\chi = 1$ corresponds to people believing that there is *no* relationship between others’ actions and others’ information. A value of $\chi \in (0, 1)$ corresponds to people being aware of, but underestimating, the link between others’ actions and information.

More formally, each person correctly assesses the probability distribution of others’ actions conditional on his own information. In our Treatment A, this corresponds to the following: a subject who sees good news has correct beliefs about the distribution of others’ initial forecasts conditional on the news he observes. However, he does not fully appreciate the relationship between others’ forecasts and the information others’ observe. Basically, the subject acts as though with probability $1 - \chi$, his partner’s initial forecast is consistent with the information the partner observes, but with probability χ , the partner’s initial forecast is independent of the information the partner observes. An equilibrium occurs when each person’s behavior is optimal given the beliefs specified above.

We have already documented that in our revised forecast stage, people tend to overweight their own private information and underreact to the information content of their partner’s initial forecast. Hence, it seems likely that allowing for cursedness will improve the equilibrium prediction in Treatment A.

We show in the appendix that for every χ , there is a unique χ -cursed equilibrium in Treatment A. In it, subjects randomize over their choice to observe revenue or cost, each with equal probability. Their initial forecasts are the same as in the Nash equilibrium, specified in (4), and their revised

forecasts are given by

$$f_R^{\text{CE}} = \begin{cases} 16\frac{2}{3} - \frac{5}{3}\chi & \text{if the subject and partner both observe good news} \\ 10 + 5\chi & \text{if subject sees good news and his partner observes bad news} \\ 10 - 5\chi & \text{if subject sees bad news and his partner observes good news} \\ 3\frac{1}{3} + \frac{5}{3}\chi & \text{if the subject and partner both observe bad news} \end{cases} \quad (13)$$

Eyster and Rabin (2005) show that in many experiments, *for every* $\chi \in (0, 1]$, the χ -cursed equilibrium prediction fits the data better than the Bayesian Nash equilibrium prediction (which corresponds to $\chi = 0$). We follow that approach by computing the mean squared error and the mean absolute error of the χ -equilibrium in Treatment A, which we plot as a function of χ in Figure 3.

[INSERT FIGURE 3 HERE]

For $\chi = 0$, which corresponds to Nash equilibrium, the mean squared error is 17.14, and the mean absolute error is 3.14. As χ increases, the mean squared error decreases until it reaches its minimum of 14.23 at $\chi = 0.62$, after which it increases to 15.32 at $\chi = 1$. The mean absolute error decreases monotonically in χ until it reaches its minimum of 2.72 at $\chi = 1$. Under both measures of performance, the standard Nash equilibrium ($\chi = 0$) does worse than *every* cursed equilibrium with positive levels of cursedness. Hence, our setting is another example in which the cursed equilibrium predictions dominate the Nash equilibrium predictions for every level of cursedness.

However, Treatment A is not a complete success for cursed equilibrium. We partition the sample based on whether the subject and his partner observed qualitatively similar private information (i.e., both good news or both bad news) or qualitatively dissimilar information (one observes good news, one observes bad news), and we compute the prediction errors for each level of χ . The results are plotted in Figure 4.

[INSERT FIGURE 4 HERE]

Figure 4a shows the results for the sample in which the subject and partner observe qualitatively similar information, which constitutes the bulk of the sample (518 out of 660). In this subsample, prediction errors monotonically decrease as χ increases, regardless of whether the measure of error is mean squared error or mean absolute error, which is strong evidence in support of cursedness. In this subsample, the improved performance of cursed equilibria relative to Nash equilibrium is driven by subjects' tendency to act as though their partner observes *the same* information as them when the qualitative information is the same, which we discussed in Section 3; when subjects observe qualitatively similar information, $\chi = 1$ corresponds to subjects issuing revised forecasts that exactly equal their initial forecast (recall (13)).

However, in the subsample in which the subjects observe qualitatively dissimilar information (Figure 4b), the supporting evidence is much weaker. Mean absolute forecast errors monotonically *increase* as χ increases, i.e., incorporating cursedness decreases the prediction accuracy relative to Nash equilibrium. Mean squared forecast errors decline to a low of 12.81 at $\chi = 0.36$, and return to its $\chi = 0$ level of 15.99 at $\chi = 0.72$, after which it continues to rise to its maximum of 23.15 at $\chi = 1$. Hence, allowing for moderate levels of cursedness improves prediction in this sample when mean squared forecast error is the error measure.

Summarizing, cursed equilibrium generally performs better than the unique symmetric Nash equilibrium at predicting subject behavior in Treatment A. The superior predictive ability of cursedness appears to be due to subjects acting as though others observed the same information as them whenever they learn that their partner’s information is qualitatively similar to their own. While our findings are generally supportive of cursedness, it does not fully explain our findings. Specifically, Nash outperforms cursedness in Treatment A when mean absolute forecast error is the performance measure and subjects observe qualitatively dissimilar information. Moreover, we only analyze the performance of the cursed equilibrium prediction in Treatment A, because the relationship between others’ actions and others’ information is irrelevant in Treatments B and C. Hence, cursedness should not even be expected to apply, and it does not explain our findings in those treatments.

5 Conclusion

Through a collection of controlled laboratory experiments, we have demonstrated that people struggle to combine their private information with information about others’ information and beliefs. In particular, many subjects act as though there is no incremental value in their partner’s information when the partner’s information is qualitatively similar to their own information. This is despite the fact that for any realistic beliefs about the partner’s information source, a subject who observes good news should become more optimistic about earnings upon learning that the partner also observed good news. Moreover, people update their beliefs as though they underestimate their partner’s ability to rationally process information. The facts that (i) most initial forecasts are optimal, and (ii) most revised forecast are optimal when the subject can infer that his partner observed different information suggests our results are not driven by a lack of subject comprehension or a lack of control by the experimenter. We speculate that our results are applicable to many real world settings where people can gain information about whether others are optimistic or pessimistic, but where others’ information sources and ability are unknown.

Our experiments have obvious implications for the analyst forecasting literature. Our results suggest that financial analysts may have difficulty incorporating the information content of other analysts’ forecasts. More specifically, the typical analyst will likely underestimate other analysts’ ability and naïvely assume that other analysts are relying on the same information as him if the analyst and the other analysts are all optimistic (or pessimistic) about the firm’s future earnings. Empirically, this latter bias should yield testable predictions that are similar to a false consensus

effect, for which Williams (2013) finds supporting evidence.

Our findings also have important implications regarding how financial markets aggregate information. Two types of equilibrium concepts dominate this literature: difference of opinion (DO) and rational expectations (RE). In DO models, investors do not learn from price—they treat the equilibrium price as completely uninformative, despite the fact that it reveals useful information about others’ signals. In RE models, investors fully understand the distribution of everyone’s signals as well as the equilibrium pricing function that maps investors’ signals to the equilibrium price. They use this knowledge of their environment to update their beliefs about the asset’s value upon observing the equilibrium price—their posterior beliefs are a function of the equilibrium price and their private signal.¹⁴ Our experiments suggest that investors may learn from price in a manner that is between these two polar cases. In our experiments, subjects recognize that others’ signals are informative, but only when it is obvious that others observed *a different* signal. This suggests that a more appropriate model of information aggregation is one where investors learn from price (as in RE models) when the equilibrium price clearly reveals that others observed private signals that differ from their own. However, when the equilibrium price is somewhat consistent with an investor’s private signal, he naïvely assumes other investors observed the same information as him, so he treats the equilibrium price as incrementally uninformative (as in DO models). Choi and Williams (2013) develop a model of trading around earnings announcements in which investors learn from prices in this manner. They show that such learning naturally leads to post-earnings announcement drift and high trading volume around earnings announcements with extreme (high or low) earnings surprises and returns, consistent with the empirical evidence.

¹⁴Often, the equilibrium price makes their own private signal incrementally uninformative, and all agents ignore their own signals.

A Proofs

In this section we provide formal proofs for various claims made earlier in the paper.

Proof. **(Nash Equilibrium Prediction)**

Since subjects' compensation is a negative quadratic in their forecast errors, their objective is to minimize their expected squared forecast error.

To determine the symmetric Nash equilibrium behavior in our experiment, we start by solving for the optimal behavior in the initial forecast stage. Expected revenue is 15, and expected cost is 5. So if a subject observes that revenue is high ($R = 20$) or that cost is low ($C = 0$), the firm's expected earnings conditional on the subject's information is 15 ($20 - 5 = 15 - 0 = 15$), and if a subject observes that revenue is low ($R = 10$) or that cost is high ($C = 10$), the firm's expected earnings conditional on the subject's information is 5 ($10 - 5 = 15 - 10 = 5$). Hence, a subject's initial forecast perfectly reveals whether he saw good news or bad news: 15 reveals good news, and 5 reveals bad news.

Now consider the optimal behavior in the revised forecast stage. Note that in all three treatments, in equilibrium, subjects know whether their partner observed good news or bad news—in Treatments B and C, subjects are informed whether their partner observed good news or bad news, and in Treatment A, the subject is informed of his partner's initial forecast, which (in equilibrium) reveals whether he saw good news or bad news.

Consider an arbitrary subject. Without loss of generality, suppose the subject observes that revenue is high ($R = 20$). Let p denote the ex ante probability that his partner chooses to observe revenue. With probability $\frac{1-p}{2}$, his partner will observe that cost is high ($C = 10$). In this scenario, the partner's optimal initial forecast will be 5, and the subject's optimal revised forecast is $f_R = 10$. His expected squared loss is driven entirely by the random error term. It can easily be verified that the expected squared forecast error in this scenario is $\frac{1}{3}$:

$$\begin{aligned} \mathbf{E}[\text{squared loss}] &= \mathbf{E}[\epsilon^2] \\ &= \frac{1}{3}. \end{aligned} \tag{14}$$

With probability 0.5, cost will be low, in which case the partner will observe good private information (regardless of whether he observes revenue or cost). The subject's optimal revised forecast given his information is given by (5):

$$f_R = \left(\frac{2p}{1+p} \right) 15 + \left(\frac{1-p}{1+p} \right) 20. \tag{15}$$

Hence, his squared loss in this scenario is given by

$$(20 + \epsilon - f_R)^2 = \left(\frac{10p}{1+p} + \epsilon \right)^2, \tag{16}$$

so in this scenario,

$$\begin{aligned}\mathbf{E}[\text{squared loss}] &= \frac{(9p-1)^2 + (11p+1)^2 + (9p-1)(11p+1)}{3(1+p)^2} \\ &\equiv \ell_1.\end{aligned}\tag{17}$$

(We use the fact that if X is uniformly distributed over the interval $[a, b]$, then $\mathbf{E}[X^2] = \frac{a^2+b^2+ab}{3}$.)

The third scenario, which occurs with probability $\frac{p}{2}$, is that cost is high ($C = 10$) and the partner observes that revenue is high ($R = 20$). In this scenario, the subject's optimal revised forecast given his information is given by (5):

$$f_R = \left(\frac{2p}{1+p}\right) 15 + \left(\frac{1-p}{1+p}\right) 20.\tag{18}$$

Hence, his squared loss in this scenario is given by

$$(10 + \epsilon - f_R)^2 = \left(\frac{10}{1+p} + \epsilon\right)^2,\tag{19}$$

so in this scenario,

$$\begin{aligned}\mathbf{E}[\text{squared loss}] &= \frac{(9-p)^2 + (11+p)^2 + (9-p)(11+p)}{3(1+p)^2} \\ &\equiv \ell_2.\end{aligned}\tag{20}$$

Combining the likelihood of the preceding scenarios and the squared forecast error in each scenario (equations (14), (17), and (20)), the subject's expected squared forecast error is given by

$$\mathbf{E}[\text{squared loss}] = \frac{1-p}{6} + \frac{\ell_1}{2} + \frac{p\ell_2}{2}.\tag{21}$$

It can be verified that this function increases in p . In other words, subjects have incentive to choose the earnings component that their partner is less likely to observe. Hence, the only symmetric equilibrium that can exist has all subjects randomly choosing to observe revenue or cost, each with 50% probability.

Plugging $p = .5$ into (5) yields the revised forecasts given in (6).

□

Proof. (Cursed Equilibrium Prediction)

The proof that subjects randomize and issue optimal initial forecasts is analogous to the Nash equilibrium case. Suppose a subject sees good news. Then his partner will observe:

- the same good news with probability 0.5
- that the other earnings realization was good with probability 0.25
- that the other earnings realization was bad with probability 0.25.

From Lemma 1 of Eyster and Rabin (2005) (page 1630), we can compute the subjects' beliefs about his partner's information as a function of the partner's initial forecast and the subject's own private information in a χ -cursed equilibrium. Assume the subject sees good news and his partner's initial forecast is 15. Then, from Lemma 1, his beliefs about his partner's type is given by:

$$\begin{aligned}\Pr(\text{partner saw the same good news}) &= \left((1 - \chi) \frac{1}{.75} + \chi \right) \left(\frac{1}{2} \right) \\ &= \frac{1}{6}(4 - \chi)\end{aligned}\quad (22)$$

$$\begin{aligned}\Pr(\text{partner saw the other earnings component is good}) &= \left((1 - \chi) \frac{1}{.75} + \chi \right) \left(\frac{1}{4} \right) \\ &= \frac{1}{12}(4 - \chi)\end{aligned}\quad (23)$$

$$\begin{aligned}\Pr(\text{partner saw the other earnings component is bad}) &= \left((1 - \chi) \frac{0}{.25} + \chi \right) \left(\frac{1}{4} \right) \\ &= \frac{\chi}{4}.\end{aligned}\quad (24)$$

From (22)-(24), a subject's optimal revised forecast (given his beliefs) when he observes good news and his partner forecasts 15 is given by

$$\begin{aligned}f^R|_{\text{good news}, f_p^0 = 15} &= \frac{1}{6}(4 - \chi)(15) + \frac{1}{12}(4 - \chi)(20) + \frac{\chi}{4}(10) \\ &= 16\frac{2}{3} - \frac{5}{3}\chi.\end{aligned}\quad (25)$$

Now assume the subject sees good news and his partner's initial forecast is 5. Using Lemma 1, his beliefs about his partner's type (conditional on the subject's information) is given by:

$$\begin{aligned}\Pr(\text{partner saw the same good news}) &= \left((1 - \chi)0 + \chi \right) \left(\frac{1}{2} \right) \\ &= \frac{\chi}{2}\end{aligned}\quad (26)$$

$$\begin{aligned}\Pr(\text{partner saw the other earnings component is good}) &= \left((1 - \chi)0 + \chi \right) \left(\frac{1}{4} \right) \\ &= \frac{\chi}{4}\end{aligned}\quad (27)$$

$$\begin{aligned}\Pr(\text{partner saw the other earnings component is bad}) &= \left((1 - \chi) \frac{1}{.25} + \chi \right) \left(\frac{1}{4} \right) \\ &= 1 - \frac{3}{4}\chi.\end{aligned}\quad (28)$$

From (26)-(28), a subject's optimal revised forecast (given his beliefs) when he observes good news and his partner forecasts 5 is given by

$$\begin{aligned}f^R|_{\text{good news}, f_p^0 = 5} &= \frac{\chi}{2}(15) + \frac{\chi}{4}(20) + \left(1 - \frac{3}{4}\chi \right) (10) \\ &= 10 + 5\chi.\end{aligned}\quad (29)$$

The {subject observes bad news, partner forecasts 15} and {subject observes bad news, partner forecasts 5} cases are symmetric to (25) and (29), so it easily follows that

$$f^R|_{\text{bad news, } f_p^0 = 5} = 3\frac{1}{3} + \frac{5}{3}\chi \quad (30)$$

$$f^R|_{\text{bad news, } f_p^0 = 15} = 10 - 5\chi \quad (31)$$

Combining (25), (29), and (30)-(31) yields (13). □

In the rest of this section we establish some properties of incentives created by the scoring rule examined in the experiment. Throughout let G be the distribution of the random variable X , when there is the possibility of confusion we denote the expectation of X with respect to this distribution by $E_G(X)$. We assume agents have twice differentiable, strictly increasing, concave Bernoulli utility functions u with $u' > 0$ and $u'' \leq 0$. Upon receiving some information, agents are asked to provide a forecast f of X and are compensated according to the proper scoring rule:

$$s(f, x) = \alpha - \beta (f - x)^2$$

where $\alpha, \beta > 0$ and x is the realized value of the random variable. Note that the partial derivatives of the scoring function are given by

$$s_f(f, x) = -2\beta (f - x)$$

and

$$s_x(f, x) = 2\beta (f - x)$$

which differ only in sign.

We begin by establishing two rather obvious results related to the scoring rule.

Proposition 1. *If a decision maker is risk neutral, then the optimal forecast $f^* = E(X)$.*

Proof. If a decision maker is risk neutral they will want to select a forecast to maximize the expected value of the scoring rule yielding the following first order condition:

$$\begin{aligned} E(s_f(f^*, X)) &= 0 \\ E(-2\beta(f^* - X)) &= 0 \\ E(-2\beta f^*) - E(-2\beta X) &= 0 \\ -2\beta f^* + 2\beta E(X) &= 0 \\ f^* &= E(X). \end{aligned}$$

□

When there is no uncertainty, any risk averse decision maker should prefer to report the known value of X .

Proposition 2. *Consider a decision maker with preferences that are strictly increasing in payoffs ($u' > 0$). If the distribution of X is degenerate at x , then the optimal forecast is $f^* = x$.*

Proof. The agent will select a forecast to maximize the (certain) utility from the scoring rule yielding the following first order condition:

$$\begin{aligned} u'(s(f^*, x)) s_f(f^*, x) &= 0 \\ -u'(s(f^*, x)) 2\beta(f^* - x) &= 0 \end{aligned}$$

which is only satisfied if $f^* = x$. It is straightforward to verify that the second order condition at the optimal solution is given by $-2\beta u'(s(f^*, x))$ so $u' > 0$ ensures this is a maximum. □

Finally, we demonstrate that shifts in the distribution in terms of first order stochastic dominance will result in more optimistic forecasts by any risk averse decision maker. A distribution H is said to first order stochastically dominate G if for all x , $G(x) \geq H(x)$ meaning that higher values are ‘always’ more likely for the distribution H than G . We use f_H and f_G to denote the forecast made by the decision maker under the different distributions.

Proposition 3. *If H first order stochastically dominates G , then for any risk averse decision maker the optimal forecasts will be such that $f_H^* \geq f_G^*$.*

Proof. Prove by contradiction. Assume the forecasts are optimal but $f_G^* > f_H^*$. In order for each forecast to be optimal, they must satisfy the following first order conditions:

$$\begin{aligned} E_H(u'(s(f_H^*, X)) s_f(f_H^*, X)) &= 0 \\ E_G(u'(s(f_G^*, X)) s_f(f_G^*, X)) &= 0. \end{aligned}$$

Also, consider the derivative of the function inside the expectation with respect to x , which is given by

$$-u''(s(f, x)) (2\beta)^2 (f - x)^2 + u'(s(f, x)) 2\beta.$$

Since $u'' \leq 0$ and $u' > 0$, we have this function is strictly positive, or the function inside the expectation is an increasing function of x . Similarly, it follows that (taking the derivative of the function with respect to f) the function inside the expectation is a strictly decreasing function of f . Assuming optimality of f_H^* we have

$$E_H(u'(s(f_H^*, X)) s_f(f_H^*, X)) = 0$$

but since H first order stochastically dominates G and the function inside the expectation is increasing, it must be that

$$E_H (u' (s(f_H^*, X) s_f (f_H^*, X))) \geq E_G (u' (s(f_H^*, X) s_f (f_H^*, X)))$$

and, since $f_G^* > f_H^*$ and the function is strictly decreasing in f

$$E_G (u' (s(f_H^*, X) s_f (f_H^*, X))) > E_G (u' (s(f_G^*, X) s_f (f_G^*, X))).$$

However, putting these three inequalities together we have that

$$E_G (u' (s(f_G^*, X) s_f (f_G^*, X))) < 0$$

which contradicts the assumption that f_G^* is optimal. □

B Experimental Instructions

Instructions for Treatment A

You are about to participate in an experiment in economics of individual decision-making. If you follow these instructions carefully and make good decisions you will earn additional money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and the experimenter will answer it. We ask that you not talk with one another for the duration of the experiment.

The experiment will continue for a number of periods. Each period is independent of the other in the sense that the outcomes of one period do not directly influence the outcomes of another. In each period you will participate in a forecasting exercise described below.

How you earn money

This is an experiment on financial forecasting. Each period, you will forecast a firm's earnings and you will be compensated based on your forecast accuracy—the more accurate you are, the more money you will earn. Your earnings from any one forecast is given by:

$$\text{Compensation} = \$6 - 0.04(\text{Your Forecast} - \text{Actual Earnings})^2.$$

You can maximize your expected compensation by issuing forecasts that are equal to the expected (average) earnings given the information you have observed. For example, if you think it's 50% likely a firm's earnings are \$10, and 50% likely the firm's earnings are \$20, you would maximize your expected compensation by issuing a forecast equal to \$15 [because $.5(\$10) + .5(\$20) = \$15$.]

How firm earnings are determined

A firm's earnings is equal to its revenues minus its costs, plus a random noise term. In other words,

$$\text{Actual Earnings} = \text{Revenues} - \text{Costs} + \text{Noise}.$$

At the time of issuing your forecast, the firm's actual earnings will not be known to you. However, you will have the opportunity to obtain information that might help you issue a more accurate forecast.

Revenues are either \$10 or \$20, each with 50% probability. Costs are either \$0 or \$10, each with 50% probability. The noise term can take any value between -\$1.00 and \$1.00, and each value is equally likely. Note that on average, the noise term will equal \$0.

Revenues and costs are independent of each other. In other words, regardless of whether the firm's revenues are \$10 or \$20, there's a 50/50 chance its costs will be \$0, and regardless of whether

the firm's costs are \$0 or \$10, there's a 50/50 chance its revenues will be \$20. Revenues and costs are also independent across time. In other words, there's always a 50/50 chance costs will be \$0, and a 50/50 chance revenues will be \$20, regardless of previous periods' costs and revenues. The noise term is independent of revenues and costs. That is, whether revenues and costs are high or low does not affect whether the noise term is unusually high or low.

Practice Periods

In each of these periods, you will be shown the firm's revenues and its costs but not the noise. You will then be asked to provide your Initial Forecast for the firm's earnings given that information. This will continue for three periods. You will not be compensated for the decisions you make in these periods.

Compensation Periods

At the beginning of each period, you must choose whether to observe the firm's revenues or costs—you may observe one, but not both. After observing that information, you will be asked to place your Initial Forecast.

Revised forecast opportunity

Each period, you will be randomly assigned a partner. After issuing your Initial Forecast, you will be shown the Initial Forecast your partner issued that period. (Similarly, your partner will see the Initial Forecast that you issue that period.) After observing your partner's forecast, you will be asked to issue a Revised Forecast for the firm's earnings that period.

Your partner observes information about the same firm as you. In other words, the revenues and costs are the same for your firm as they are for his or her firm. (Subjects who are not your partner, e.g., your neighbor, may observe revenues or costs from a different firm.)

Note that you will generally be assigned a different random partner each period.

Results and earnings determination

At the end of the each period, the firm's actual earnings will be announced and your forecast earnings from both your Initial Forecast and Revised Forecast will be displayed.

After all the periods have been completed, one period will be randomly chosen as your Initial Forecast compensation period (i.e., you will be compensated based on the accuracy of your initial forecast that period), and one period will be randomly chosen as your Revised Forecast compensation period (i.e., you will be compensated based on the accuracy of the forecast you issue after observing your partner's initial forecast). Your total earnings will be the compensation from these two forecasts and the \$5 show up fee.

Instructions for Treatment B

You are about to participate in an experiment in economics of individual decision-making. If you follow these instructions carefully and make good decisions you will earn additional money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and the experimenter will answer it. We ask that you not talk with one another for the duration of the experiment.

The experiment will continue for a number of periods. Each period is independent of the other in the sense that the outcomes of one period do not directly influence the outcomes of another. In each period you will participate in the forecasting exercise described below.

How you earn money

This is an experiment on financial forecasting. Each period, you will forecast a firm's earnings and you will be compensated based on your forecast accuracy—the more accurate you are, the more money you will earn. Your compensation from any one forecast is given by:

$$\text{Compensation} = \$6 - 0.04(\text{Your Forecast} - \text{Actual Earnings})^2.$$

You can maximize your expected compensation by issuing forecasts that are equal to the expected (average) earnings given the information you have observed. For example, if you think it's 50% likely a firm's earnings are \$10, and 50% likely the firm's earnings are \$20, you would maximize your expected compensation by issuing a forecast equal to \$15 [because $.5(\$10) + .5(\$20) = \$15$.]

How firm earnings are determined

A firm's earnings are equal to its revenues minus its costs, plus a random noise term. In other words,

$$\text{Actual Earnings} = \text{Revenues} - \text{Costs} + \text{Noise}.$$

At the time of issuing your forecast, the firm's actual earnings will not be known to you. However, you will have the opportunity to obtain information that might help you issue a more accurate forecast.

Revenues are either \$10 or \$20, each with 50% probability. High revenues (\$20) are considered “good news,” and low revenues (\$10) are considered “bad news.” Costs are either \$0 or \$10, each with 50% probability. Low costs (\$0) are considered “good news,” and high costs (\$10) are considered “bad news.” The noise term can take any value between -\$1.00 and \$1.00, and each value is equally likely. Note that on average, revenues are \$15, costs are \$5, and the noise term is \$0.

Revenues and costs are independent of each other. In other words, regardless of whether the firm's revenues are good or bad, there's a 50% chance its costs will be good, and regardless of whether the firm's costs are good or bad, there's a 50% chance its revenues will be good. Revenues and costs are also independent across time. In other words, in every period there is a 50% chance

revenue will be good, and a 50% chance costs will be good, regardless of previous periods' costs and revenues. The noise term is independent of revenues and costs. That is, whether revenues and costs are high or low does not affect whether the noise term is unusually high or low.

Initial Forecasts

At the beginning of each period, you must choose whether to observe the firm's revenues or costs—you may observe one, but not both. After observing that information, you will be asked to place your Initial Forecast.

Revised forecast opportunity

Each period, you will be randomly assigned a partner. After issuing your Initial Forecast, you will be informed whether your partner observed “good news” or “bad news.” Recall that “good news” corresponds to either high revenues (\$20) or low costs (\$0), and “bad news” corresponds to either low revenues (\$10) or high costs (\$10). (Similarly, your partner will see whether the news you observed was “good” or “bad.”) You will not be told whether your partner chose to observe revenues or costs. After observing whether your partner observed good news or bad news, you will be asked to issue a Revised Forecast for the firm's earnings that period.

Your partner observes information about the same firm as you. In other words, the revenues and costs are the same for your firm as they are for his or her firm. (Subjects who are not your partner, e.g., your neighbor, may observe revenues or costs from a different firm.)

Note that you will generally be assigned a different random partner each period.

Results and earnings determination

At the end of the each period, the firm's actual earnings will be announced and your forecast compensation from both your Initial Forecast and Revised Forecast will be displayed.

After all the periods have been completed, one period will be randomly chosen as your Initial Forecast compensation period (i.e., you will be compensated based on the accuracy of your initial forecast that period), and one period will be randomly chosen as your Revised Forecast compensation period (i.e., you will be compensated based on the accuracy of the forecast you issue after observing your partner's initial forecast). Your total compensation will be the sum of your compensation from these two forecasts and the \$5 show up fee.

Instructions for Treatment C

You are about to participate in an experiment in economics of individual decision-making. If you follow these instructions carefully and make good decisions you will earn additional money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and the experimenter will answer it. We ask that you not talk with one another for the duration of the experiment.

The experiment will continue for a number of periods. Each period is independent of the other in the sense that the outcomes of one period do not directly influence the outcomes of another. In each period you will participate in the forecasting exercise described below.

How you earn money

This is an experiment on financial forecasting. Each period, you will forecast a firm's earnings and you will be compensated based on your forecast accuracy—the more accurate you are, the more money you will earn. Your compensation from any one forecast is given by:

$$\text{Compensation} = \$6 - 0.04(\text{Your Forecast} - \text{Actual Earnings})^2.$$

You can maximize your expected compensation by issuing forecasts that are equal to the expected (average) earnings given the information you have observed. For example, if you think it's 50% likely a firm's earnings are \$10, and 50% likely the firm's earnings are \$20, you would maximize your expected compensation by issuing a forecast equal to \$15 [because $.5(\$10) + .5(\$20) = \$15$].

How firm earnings are determined

A firm's earnings are equal to its revenues minus its costs, plus a random noise term. In other words,

$$\text{Actual Earnings} = \text{Revenues} - \text{Costs} + \text{Noise}.$$

At the time of issuing your forecast, the firm's actual earnings will not be known to you. However, you will obtain information that might help you issue a more accurate forecast.

Revenues are either \$10 or \$20, each with 50% probability. High revenues (\$20) are considered “good news,” and low revenues (\$10) are considered “bad news.” Costs are either \$0 or \$10, each with 50% probability. Low costs (\$0) are considered “good news,” and high costs (\$10) are considered “bad news.” The noise term can take any value between -\$1.00 and \$1.00, and each value is equally likely. Note that on average, revenues are \$15, costs are \$5, and the noise term is \$0.

Revenues and costs are independent of each other. In other words, regardless of whether the firm's revenues are good or bad, there's a 50% chance its costs will be good, and regardless of whether the firm's costs are good or bad, there's a 50% chance its revenues will be good. Revenues

and costs are also independent across time. In other words, in every period there is a 50% chance revenue will be good, and a 50% chance costs will be good, regardless of previous periods' costs and revenues. The noise term is independent of revenues and costs. That is, whether revenues and costs are high or low does not affect whether the noise term is unusually high or low.

Initial Forecasts

At the beginning of each period, the computer will randomly choose whether to display the firm's costs or revenues. It is equally likely that the computer will choose to display revenues as costs (each has 50% probability). After observing that information, you will be asked to place your Initial Forecast.

Revised forecast opportunity

Each period, you will be randomly assigned a partner. Your partner is also shown the firm's revenues or costs, but not both. Whether your partner is shown revenues or costs is independent of whether you are shown revenues or costs. That is, the likelihood that your partner observes revenues is 50% regardless of whether the computer shows you the firm's revenues or costs. After issuing your Initial Forecast, you will be informed whether your partner observed "good news" or "bad news." Recall that "good news" corresponds to either high revenues (\$20) or low costs (\$0), and "bad news" corresponds to either low revenues (\$10) or high costs (\$10). (Similarly, your partner will see whether the news you observed was "good" or "bad.") You will not be told whether your partner was shown the firm's revenues or costs. After observing whether your partner observed good news or bad news, you will be asked to issue a Revised Forecast for the firm's earnings that period.

Your partner observes information about the same firm as you. In other words, the revenues and costs are the same for your firm as they are for his or her firm. (Subjects who are not your partner, e.g., your neighbor, may observe revenues or costs from a different firm.)

Note that you will generally be assigned a different random partner each period.

Results and earnings determination

At the end of the each period, the firm's actual earnings will be announced and your forecast compensation from both your Initial Forecast and Revised Forecast will be displayed.


After all the periods have been completed, one period will be randomly chosen as your Initial Forecast compensation period (i.e., you will be compensated based on the accuracy of your initial forecast that period), and one period will be randomly chosen as your Revised Forecast compensation period (i.e., you will be compensated based on the accuracy of the forecast you issue after observing your partner's initial forecast). Your total compensation will be the sum of your compensation from these two forecasts and the \$5 show up fee.

References

- ADMATI, A. (1985): “A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets,” *Econometrica*, 53, 629–658.
- ANDERSON, L. R., AND C. A. HOLT (1997): “Information Cascades in the Laboratory,” *The American Economic Review*, 87, 847–862.
- CHOI, J., AND J. WILLIAMS (2013): “Earnings Announcement Returns, Trading Volume, and Price Drift When Investors Overestimate their Similarity to Others,” *Working Paper*.
- DANIEL, K., D. HIRSHLEIFER, AND A. SUBRAHMANYAM (1998): “Investor Psychology and Security Market Under- and Overreactions,” *The Journal of Finance*, 53, 1839–1885.
- EREV, I., AND A. E. ROTH (1998): “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria,” *The American Economic Review*, 88, 848–881.
- EYSTER, E., AND M. RABIN (2005): “Cursed Equilibrium,” *Econometrica*, 73, 1623–1672.
- FISCHBACHER, U. (2007): “z-Tree: Zurich Toolbox for Ready-made Economic Experiments,” *Experimental Economics*, 10(2), 171–178.
- GERVAIS, S., J. HEATON, AND T. ODEAN (2011): “Overconfidence, Compensation Contracts, and Capital Budgeting,” *The Journal of Finance*, 66, 1735–1777.
- GROSSMAN, S., AND J. STIGLITZ (1976): “Information and Competitive Price Systems,” *The American Economic Review (Papers and Proceedings)*, 66, 246–253.
- HELLWIG, C., AND L. VELDKAMP (2009): “Knowing what Others Know: Coordination Motives in Information Acquisition,” *The Review of Economic Studies*, 76, 223–251.
- MALMENDIER, U., AND G. TATE (2005): “CEO Overconfidence and Corporate Investment,” *The Journal of Finance*, 60, 2661–2700.
- MALMENDIER, U., G. TATE, AND J. YAN (2011): “Overconfidence and Early-Life Experiences: The Effect of Managerial Traits on Corporate Financial Policies,” *The Journal of Finance*, 66, 1687–1733.
- MYATT, D., AND C. WALLACE (2012): “Endogenous Information Acquisition in Coordination Games,” *The Review of Economic Studies*, 79, 340–374.
- NIEUWERBURGH, S. V., AND L. VELDKAMP (2009): “Information Immobility and the Home Bias Puzzle,” *The Journal of Finance*, 64, 1187–1215.
- (2010): “Information Acquisition and Under-Diversification,” *The Review of Economic Studies*, 77, 779–805.

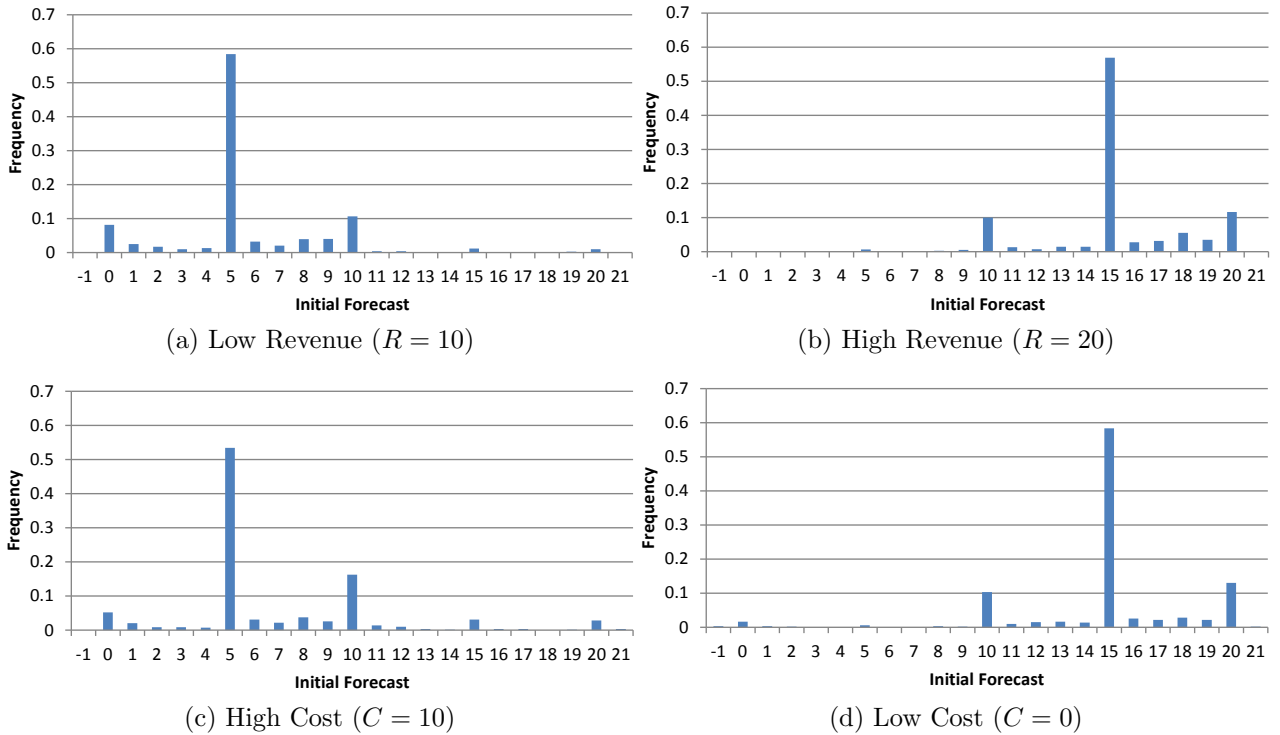
- ODEAN, T. (1998): “Volume, Volatility, Price, and Profit When All Traders Are Above Average,” *The Journal of Finance*, 53, 1887–1934.
- ROSS, L., D. GREENE, AND P. HOUSE (1977): “The “False Consensus Effect”: An Egocentric Bias in Social Perception and Attribution Processes,” *Journal of Experimental Social Psychology*, 13, 279–301.
- SCHEINKMAN, J., AND W. XIONG (2003): “Overconfidence and Speculative Bubbles,” *Journal of Political Economy*, 111, 1183–1219.
- SVENSON, O. (1981): “Are We all Less Risky and More Skillful than our Fellow Drivers?,” *Acta Psychologica*, 47, 143–148.
- WILLIAMS, J. (2013): “Financial Analysts and the False Consensus Effect,” *Journal of Accounting Research*, 51, 855–907.

Figure 1: Distribution of Subjects' Propensity to Choose to Observe Revenue




For each $n = 0, \dots, 30$, we plot the number of subjects (out of 60) from Treatments A and B who chose to observe revenue in n of the 30 periods.

Figure 2: Distribution of Initial Forecasts, by Private Information



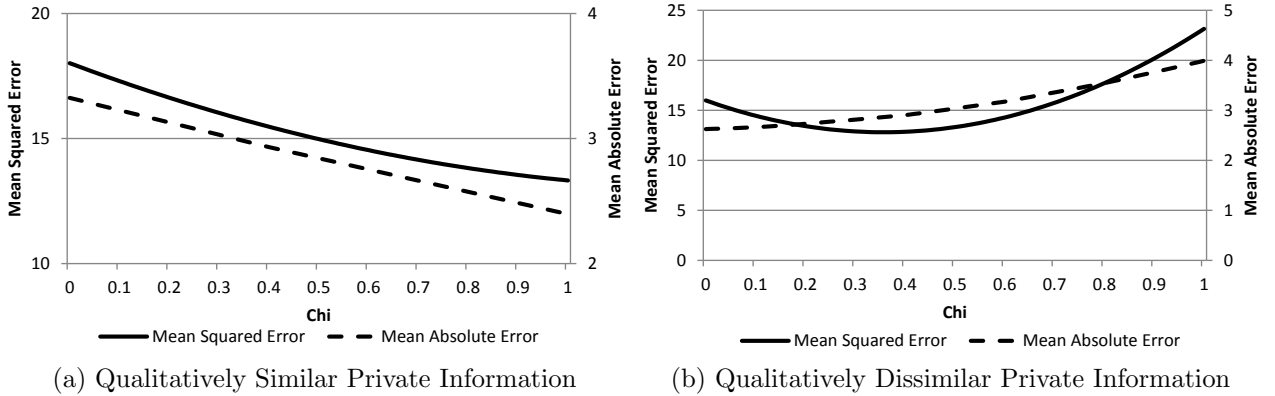
We plot the distribution of initial forecasts for each of the four possible pieces of private information: low revenue, high revenue, high cost, and low cost.

Figure 3: Prediction Errors of Cursed Equilibria, by χ



For each $\chi \in \{0, \frac{1}{100}, \dots, 1\}$, we plot the average squared prediction error, $(f_R - f_R^{\text{CE}})^2$, and the average absolute prediction error, $|f_R - f_R^{\text{CE}}|$, of the revised forecasts predicted by the χ -cursed equilibrium. The χ -equilibrium predicted revised forecast, f_R^{CE} , is defined in (13).

Figure 4: Prediction Errors of Cursed Equilibria, by χ and Information



(a) Qualitatively Similar Private Information

(b) Qualitatively Dissimilar Private Information

We partition Treatment A into two samples based on whether the subject and his partner observe qualitatively similar (Panel 4a) or qualitatively dissimilar (Panel 4b) private information. For each $\chi \in \{0, \frac{1}{100}, \dots, 1\}$ and for each subject, we compute the subject's average squared prediction error, $(f_R - f_R^{\text{CE}})^2$, and the average absolute prediction error, $|f_R - f_R^{\text{CE}}|$, of the revised forecasts predicted by the χ -cursed equilibrium. Then, for each χ , we compute the average of the subjects' average prediction errors. The χ -equilibrium predicted revised forecast, f_R^{CE} , is defined in (13).

Table 1: Overview of Treatments

	Treatment A	Treatment B	Treatment C
Number of Subjects	22	38	46
Number of Sessions	1	2	2
Subjects choose information source (revenues/costs)	Yes	Yes	No
Subjects observe partner's initial forecast	Yes	No	No
Subjects informed whether partner saw good/bad news	No	Yes	Yes

We present an overview of our three experimental treatments.

Table 2: Initial Forecasts

Treatment	$ f_0 - f_0^{\text{NE}} $	$(f_0 - f_0^{\text{NE}})^2$	$\mathbf{1}_{\{f_0 = f_0^{\text{NE}}\}}$
Panel A: Means			
A	1.99	11.72	0.54
B	1.80	12.42	0.63
C	2.17	12.97	0.52
Panel B: t -statistics (Difference of Means)			
B-A	-0.32	0.12	0.99
C-A	0.36	0.30	-0.16
C-B	0.81	0.12	-1.38

f_0 is a subject's initial forecast, and f_0^{NE} is defined in (4) and is the Nash equilibrium predicted initial forecast, which is also equal to the subject's compensation-maximizing initial forecast. In Panel A, we report means, and in Panel B, we report t -statistics for the difference in means across treatments. We conduct our statistical analysis by treating the subject as the unit of observation. That is, we compute means at the subject level, and then we compare the distribution of subject-level means across treatments, where our number of observations in each treatment is the number of subjects in that treatment.

Table 3: Revised Forecasts

Treatment	$\mathbf{1}_{\{f_R=f_0\}}$	$ f_R - f_0 $	$ f_R - f_R^{\text{NE}} $	OPI_{NE}	$ f_R - f_R^{\text{opt}} $	OPI_{adj}
Panel A: Means						
A	0.63	1.68	3.14	2.09***	.	.
B	0.42	2.89	2.03	0.51***	.	.
C	0.40	2.91	2.41	0.68***	.	.
A'	0.70	1.43	.	.	2.58	1.45***
Panel B: <i>t</i> -statistics (Difference of Means)						
B-A	-3.33	2.79	-3.53	-5.44	.	.
C-A	-3.50	3.12	-2.09	-3.85	.	.
C-B	-0.27	0.04	1.37	0.57	.	.
B-A'	-4.26	3.28	.	.	-1.73	-2.68
C-A'	-4.40	3.67	.	.	-0.58	-1.87

f_0 is the subject's initial forecast, f_R is the subject's revised forecast, f_R^{NE} is defined in (6) and is the revised forecast predicted by the unique symmetric Nash equilibrium, OPI_{NE} is defined in (7) and it represents the degree to which subjects overweight their private information relative to the symmetric Nash equilibrium, f_R^{opt} is defined in (9)-(10) and represents a subject's optimal revised forecast, and OPI_{adj} is defined in (11)-(12) and it represents the degree to which subjects overweight their private information relative to their optimal behavior (given other subjects' actual behavior). Treatment A' refers to the observations from Treatment A in which the subject's partner issues a forecast that appears to be optimal (from the subject's perspective). Our statistical methodology is described in Table 2. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 4: Revised Forecasts when Partner's Information is Qualitatively Dissimilar from the Subject's

Treatment	$\mathbf{1}_{\{f_R=f_0\}}$	$ f_R - f_0 $	$ f_R - f_R^{\text{NE}} $	OPI_{NE}	$ f_R - f_R^{\text{opt}} $	OPI_{adj}
Panel A: Means						
A	0.45	2.66	2.61	1.67***	.	.
B	0.12	4.94	0.91	0.09	.	.
C	0.13	4.94	1.48	0.10	.	.
A'	0.43	3.29	.	.	2.29	0.86
Panel B: <i>t</i> -statistics (Difference of Means)						
B-A	-4.96	4.61	-3.62	-3.68	.	.
C-A	-4.74	4.65	-2.31	-3.79	.	.
C-B	0.48	-0.55	1.28	0.24	.	.
B-A'	-4.44	3.47	.	.	-2.24	-1.42
C-A'	-4.30	3.44	.	.	-1.06	-1.37

This table is identical to Table 3 except that the analysis is conducted on a subset of the sample. Treatments A, B, and C are restricted to observations in which the subject and his partner observe qualitatively different private information, i.e., one observes good news and the other observes bad news. Treatment A' is restricted to observations in which a subject's private information is qualitatively opposed to the private information *implied* by the partner's initial forecast, i.e., the sample is restricted to observations for which one of the following conditions holds: (i) the subject observed good news and his partner's initial forecast is 5, or (ii) the subject observed bad news and his partner's initial forecast is 15. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 5: Revised Forecasts when Partner's Information is Qualitatively Similar to the Subject's

Treatment	$\mathbf{1}_{\{f_R=f_0\}}$	$ f_R - f_0 $	$ f_R - f_R^{\text{NE}} $	OPI_{NE}	$ f_R - f_R^{\text{opt}} $	OPI_{adj}
Panel A: Means						
A	0.68	1.41	3.28	2.21***	.	.
B	0.52	2.19	2.41	0.65***	.	.
C	0.49	2.25	2.71	0.87***	.	.
A'	0.77	0.93	.	.	2.66	1.61***
Panel B: <i>t</i> -statistics (Difference of Means)						
B-A	-2.24	1.54	-2.65	-4.40	.	.
C-A	-2.29	1.77	-1.75	-2.92	.	.
C-B	-0.25	0.07	1.03	0.66	.	.
B-A'	-3.22	2.43	.	.	-0.90	-2.55
C-A'	-3.20	2.76	.	.	-0.16	-1.59

This table is identical to Table 3 except that the analysis is conducted on a subset of the sample. Treatments A, B, and C are restricted to observations in which the subject and his partner observe qualitatively similar private information, i.e., both observe good news or both observe bad news. Treatment A' is restricted to observations in which a subject's private information is qualitatively consistent with the private information *implied* by the partner's initial forecast, i.e., the sample is restricted to observations for which one of the following conditions holds: (i) the subject observed good news and his partner's initial forecast is 15, or (ii) the subject observed bad news and his partner's initial forecast is 5. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.