

Diverging Opinions*

James Andreoni

Tymofiy Mylovanov

University of California, San Diego
and NBER

University of Pennsylvania

June 28, 2011

Abstract

People often see the same evidence but draw opposite conclusions, becoming increasingly polarized over time. More surprisingly, such disagreements persist even when they are commonly known. In this paper, we derive a simple model and present an experiment showing that opinion polarization can emerge trivially when one-dimensional opinions are formed from two-dimensional information. Contrary to the theoretical prediction, however, when subjects are given sufficient information to reach agreement, disagreement persists. Our analysis shows that people discount information when it is filtered through the actions of others, but not when it is presented directly, indicating that common knowledge of disagreement may be possible because people are overly skeptical of the decision-making skills of others.

Keywords: polarization of opinions, Bayesian beliefs, updating of beliefs

*We are deeply grateful to Nageeb Ali, Julia L. Evans, Eran Hanany, Navin Kartik, Peter Klibanoff, Georg Noldeke, Justin Rao, Matthew Rabin, Joel Sobel, Leeat Yariv, and Muhamet Yildiz for the many helpful comments and conversations. We gratefully acknowledge financial support from the National Science Foundation, grant 1024683, and from the German Science Foundation (DFG) through SFB/TR 15 “Governance and the Efficiency of Economic Systems.” We are thankful to Ben Cowan and Megan Ritz for expert research assistance.

1 Introduction

We see many examples in the world around us in which individuals observe the same information but draw opposite conclusions, and in which additional information only results in increased polarity. This is true even among educated experts, such as on a divided supreme court, or academics locked in ideological struggles. We observe diverging opinions about sundry issues, such as gun control, social welfare benefits, affirmative action, the war in Iraq, and the death penalty. Differences of opinions can be a cause of speculative trade in financial markets, inefficient delays in bargaining, and political polarization.

How can two people see the same information and draw opposite conclusions? How can additional public information persistently draw two people into a stronger disagreement? And why can disagreement persist after it becomes commonly known and individuals could learn from opinions of others? If the individuals share common prior beliefs and are Bayesian, their posterior beliefs after additional information should converge (Blackwell and Dubins 1962). If, in addition, there is common knowledge of rationality, any remaining disagreement should disappear after individuals are allowed to share their beliefs (Aumann 1976, Geanakoplos and Polemarchakis 1982).

In this paper, we explore these questions in a controlled experimental environment and demonstrate the possibility of disagreement—and its persistence—despite the fact that subjects are provided sufficient information for disagreement to vanish. In the data, persistence of disagreement can be explained by a bias that causes individuals to underestimate the precision of others' information relative to that of their own. We quantify the degree of this bias in our experiment.

We show theoretically, and demonstrate experimentally, that when the dimensionality of information exceeds the dimensionality that describes the true state of nature, it is trivial to generate examples where opinions can differ and diverge. We explore the case in which the state of nature is one dimensional, while the uncertainty about the state of nature is two dimensional. In the experimental group, individuals

observe a private signal about one dimension followed by a sequence of public signals about the other dimension. Both dimensions are important for identification of the state of nature and different private signals can induce diverging posterior beliefs: Public information resolves uncertainty on one dimension and exaggerates the impact of private signals about the other dimension on the posterior beliefs. In the control group, the private information on one dimension is followed by the sequence of public signals providing information about *both* dimensions. In this environment, common public information eventually overwhelms any difference in beliefs due to differential private signals and disagreement in beliefs disappears.

Consider the following illustration. A large country is considering whether to intervene in an uprising against a dictator in a distant small country. What is public opinion on the matter – should the big country aid the rebels or stay out? Over the years, the public has paid differential attention to news about the dictator. Sometimes the dictator has not been friendly to the big country, but at times he seems more moderate than others in the region. Forming an opinion on this policy decision requires two pieces of information: how friendly is the dictator to the big country, and will the rebels be more friendly or even worse than the dictator? Whether one evaluates public information about the rebels as being better or worse than the dictator will depend on the private view one has gained over the years about the disposition of the dictator. As everyone forms more confident beliefs about the attitudes of the rebels—and even agreeing on what those attitudes are—it may not lead to agreement on the policy. One group with a dim view of the dictator may simply becoming more certain that the rebels are better than the dictator (thus favoring intervention) while a group with a more optimistic view of the dictator may be more certain that rebels are clearly worse (thus opposing intervention). As long as the debate focuses on the new information, the public can come more and more into *disagreement* on the policy, even while they come into greater *agreement* about the meaning of the new information. Unless everyone can also reach agreement about their privately held

beliefs about the dictator, agreement on the policy can become impossible, and may become persistent.

While this example is simple, it is meant to highlight what we believe to be the most interesting and vexing aspect of disagreement. In particular, while the emergence of differences of opinion is important to understand, the persistence of diverging opinions once disagreement is commonly known is perhaps an even deeper puzzle.

How do we address the question of the persistence of differing views in the presence of common awareness of disagreement? In the final round in each session of our experiment, we provide subjects with information about the actions of others in all of the previous rounds. If the subjects believe the reasoning of others is similar to their own, that is, if there is common knowledge of rationality, this information is sufficient to infer their private information and should eliminate disagreement. Surprisingly, we find that, despite giving subjects the common information they need to reach full agreement, a sizable minority of our subjects maintain their opposing views.

Why do people remain in disagreement? Agreement requires two ingredients. One is sufficient rationality, that is, individuals must be reacting to their own information so that their choices reveal what they know. Second, agreement requires a common (or at least sufficient) knowledge of this rationality. We find evidence that both of these may be missing. First, we look at choices in round 1, when individuals should still maintain common priors, being indifferent about the true state. Nonetheless, we see that about 20% of the sample erroneously disagrees and favors one point of view. Moreover, while other errors tend to diminish as the experiment progresses, the fraction making this type of error is nearly constant. One may interpret disagreement in this case as evidence of erroneous or nonrational choices. Next, we look at the final round where information about disagreement is made public and, under common knowledge of rationality, should be sufficient to eliminate disagreement. Here we find that individuals weigh their own information more than twice that of the five others

in their group. When we look separately at those who err by disagreeing in round 1, we find that these people weigh their own information more than 10 times that of others, putting virtually no stock in public information. This indicates a different type of error, that is, a failure of some individuals to learn from each other. This error is quite large and for a nontrivial minority of the population.¹ Setting aside the subjects who make systematic errors, we find that individuals still put 50% more weight on their own information than they do on the information revealed through the actions of others, although this difference is not statistically significant.

In our experiment, subjects are instilled with a common objective prior. Furthermore, there is no interdependence among subjects' actions – each subject faces an individual decision problem.

It is important to note that inducing common priors about the fundamentals has two methodological advantages in our context. First, it makes it easier to control subjects' information in the laboratory: in the experiment, we simply show the subjects the urns, representing possible states, from which signals are drawn and then measure their beliefs to validate our controls. The alternative of non-common priors would require either deception or selection of subjects based on the beliefs acquired outside the lab. Second, the common prior assumption allows for a clear test of the common knowledge of rationality among subjects. Without this assumption, we would not be able to interpret the observation that disagreement may continue to persist despite sufficient information that should eliminate it.

In summary, the results in this paper suggest the following story. The initial disagreement may arise and continue to increase in the light of commonly observed new evidence because the public has different private views of the world, formalized and induced in our experiment by an observation of a private signal. Furthermore,

¹An important paper by Weizsacker (2010) comes to a strikingly similar conclusion, but in a much different context from that studied here. Independently of our study, he shows in a meta-analysis of games with information cascades that a significant share of subjects do not conform to rational expectations, and put about twice the weight on their own information as they do on others. The confluence of his results with our own potentially bolsters both findings.

as confirmed by our experiment, even if there is sufficient information to infer the models of others, the public may fail to do so and the disagreement can persist and become common knowledge.

If our multi-dimensional-information theory is correct, it changes the focus of the puzzle about belief polarization. Rather than ask what are the causes of polarization, this suggests that instead we should ask why they may have different views of the world and why they don't share all the information that shaped their individual views? What keeps them from communicating enough to draw their "world views" together? That is, why would they doubt rationality of others?

The remainder of the paper is organized as follows. Next we briefly review the background literature. The model is presented in section 3. Section 4 describes our experimental design. The experimental results are provided in section 5. Section 6 concludes. Some proofs are in the appendix.

2 Background

Much of the evidence on diverging opinions comes from psychological studies.² In the seminal experiment by Lord, Ross, and Lepper (1979) subjects who were selected because of differing views on the death penalty were pulled further apart after reading the same essay about the death penalty. This type of result has been replicated in numerous studies.³ The important novelty in our paper is to show that polarization might occur in an environment with an instilled objective prior and objective

²We refer the reader to the surveys of the literature by Barberis and Thaler (2003), Gerber and Green (1999), Hirshleifer (2001), Narasimhan, He, Anderson, Brenner, Desai, Kuksov, Messinger, Moorthy, Nunes, Rottenstreich, Staelin, and and (2005), and Rabin (1998).

³Similar results are obtained by Houston and Fazio (1989) and Schuette and Fazio (1995) in the context of capital punishment, Katz and Feldman (1962) and Sigelman and Sigelman (1984) in the context of presidential debates, Kinder and Walter R. Mebane, Jr. (1983) in the context of evaluation of the state of economy, and Sears (1968) in the context of the credibility of the source of factual information. Nickerson (1998) provides a survey of related evidence; additional references can also be found in Gerber and Green (1999). Finally, a recent study by Westen, Blagov, Harenski, Kilts, and Hamann (2006) finds further support for the effect of prior political attitudes on the interpretation of available evidence in an fMRI study.

information and that it can persist even after it becomes commonly known.

The existing explanations of belief polarization include heterogeneous prior beliefs (Dixit and Weibull 2007, Acemoglu, Chernozhukov, and Yildiz 2009), non-Bayesian updating caused by, e.g., confirmatory bias (Rabin and Schrag 1999)⁴ and ambiguity aversion (Zimper and Ludwig 2007, Baliga, Hanany, and Klibanoff 2011), differential private information (Kondor 2011), and memory constraints (Wilson 2005).

In our model, private signals are used to interpret the implication of the public signals and different private signals induce distinct interpretations. A similar idea is present in Acemoglu, Chernozhukov, and Yildiz (2009) and Kondor (2011)⁵. These papers consider Bayesian models in which there is uncertainty about the state of nature and individuals have different priors or different private information that determines interpretation of public signals.

Our paper complements this body of work, first, by offering a very simple and transparent example of a Bayesian environment with diverging opinions, second, by experimentally demonstrating the possibility of disagreement—and its persistence—in an environment with common priors, and, third, by testing the assumption of common knowledge of rationality. The essential contribution is to demonstrate that persistence of disagreement in our experiment is due to failure of common knowledge of rationality.

The theme of this paper is related to that of Cripps, Ely, Mailath, and Samuelson (2008) (CEMS) and Sethi and Yildiz (2009). CEMS provide conditions under which individuals who privately learn the value of a parameter will also learn it commonly. In a model with heterogeneous priors and private information, Sethi and Yildiz (2009) provide conditions under which private information is aggregated through repeated communication.

In our model and experiment, the only source of disagreement is differential

⁴See also Gerber and Green (1999) for a review from the political science perspective and Nickerson (1998) from the psychological perspective. Eil and Rao (2011), however, show that much of the evidence on confirmatory bias is conflated with a good-news/bad-news effect.

⁵See also Kandel and Pearson (1995) and Kim and Verrecchia (1997)

private information. An alternative reason for disagreement and its persistence might be a conflict of preferences among individuals as in Santos-Pinto and Sobel (2005).

Finally, there is a large literature in psychology on the possible reasons for diverging opinions. We refer the reader to the literature reviewed in an earlier version of this paper (Andreoni and Mylovanov 2010).

3 Model

The model presented here is made as simple as possible. While it can be generalized to signal distributions other than the one considered below (c.f., Section 3.5), the current model is sufficient to clarify the features that generate diverging opinions. In our model, private and public signals are complements: the value of either public or private signals alone is zero.⁶ To see the intuition for our results imagine that two players play one shot of matching pennies. There are two outside observers who receive noisy information about the players' moves and are asked to bet on the winner of the game. Imagine that the observers receive different private information about the move of player 1. This information is not helpful in determining the winner of the game and hence does not affect their opinions. We now let them observe a public signal about the move of player 2. Together with private information, this signal is valuable. Furthermore, the observers with different private information will now diverge in their opinions about who is the more likely winner of the game. Thus, arrival of public information can cause divergence of opinions.

3.1 Environment

The state of nature $\theta = (\alpha, \beta)$ is a realization of a random variable $\tilde{\theta} = (\tilde{\alpha}, \tilde{\beta})$, where $\tilde{\alpha}, \tilde{\beta} \in \{0, 1\}$. All states are equally likely. There are two Bayesian agents, each can

⁶Boergers, Hernando-Veciana, and Krähmer (2009) study signals that are complements and substitutes and show that complementary signal structures are *not* non-generic.

take an action $a \in \{Even, Odd\}$. The payoff of an agent is

$$\begin{aligned} u(Even, \theta) &= \begin{cases} 1, & \text{if } \theta \text{ equals } (0, 0) \text{ or } (1, 1); \\ 0, & \text{otherwise;} \end{cases} \\ u(Odd, \theta) &= \begin{cases} 1, & \text{if } \theta \text{ equals } (1, 0) \text{ or } (0, 1); \\ 0, & \text{otherwise,} \end{cases} \end{aligned} \quad (3.1)$$

independently of the action taken by the other agent.

Agents do not know the state and observe two signals, $\tilde{a}, \tilde{b} \in \{0, 1\}$, that are distributed independently conditional on the state with $\Pr(\tilde{a} = \alpha | \alpha) = p_\alpha > 1/2$ and $\Pr(\tilde{b} = \beta | \beta) = p_\beta > 1/2$.

There are infinitely many periods, $t = 0, 1, \dots$. In period zero, the agents privately observe independent realizations of signal \tilde{a} . Starting from the first period, the agents commonly observe a realization of signal \tilde{a} or \tilde{b} in each period. We will consider two settings: in one all public signals are \tilde{b} , whereas in the other public signals are of both types, \tilde{b} in odd periods and \tilde{a} in even periods.

3.2 Disagreement about the optimal action

Let *type 0* and *type 1* denote the agents who observe private signal $a = 0$ and $a = 1$ respectively. After observing $a = 1$, *type 1* believes that $\alpha = 1$ is more likely, while his beliefs about β are unaffected. However, (3.1) implies that this type is indifferent about which action to take. A similar argument applies to type 0. Hence, although different signals in the first period might lead to distinct beliefs about α , they cannot create a disagreement about the optimal action.

Nevertheless, private signals can be used to interpret future signals about the other dimension and can lead to disagreement. We say that type 0 and type 1 *disagree* about the optimal action if they strictly prefer different actions.⁷ Imagine that in the second period both types observe $b = 1$. Now, type 1 believes that the

⁷We say that the different types *weakly disagree* about the optimal action if one type is indifferent about which action to take and the other type believes that one of the actions is more likely to be optimal.

state $\theta = (1, 1)$ is more likely and the optimal course of action is $a = \text{Even}$, while type 0 disagrees.

The independence of the signals conditional on the state and their binomial distribution implies that the agent's posterior beliefs depend only on the difference between the number of realizations of different signals, but not their order. Let t_a and t_b be the number of respective public signals and k_a and k_b be the number of realizations of the corresponding signals equal to one. We define

$$\begin{aligned}\delta_a^0 &= 2k_a - t_a - 1, \\ \delta_a^1 &= 2k_a - t_a + 1, \\ \delta_b &= 2k_b - t_b.\end{aligned}$$

Remark 1. If $\delta_b = 0$, both types believe that both actions are equally likely to be optimal.

Remark 2. Different types disagree about the optimal action if and only if (i) $\delta_b \neq 0$ and (ii) $\delta_a^1 = 1$ and $\delta_a^0 = -1$.

3.3 Public signals \tilde{b}

If all public signals are \tilde{b} , we have $\delta_a^1 = 1$ and $\delta_a^0 = -1$ for any number of signals. Then, the probability of disagreement between types 0 and 1 is equal to $1 - \Pr(\delta_b = 0)$.

If the number of public signals is odd, then $\delta_b \neq 0$, in which case different types disagree with probability one. If, however, the number of public signals \tilde{b} is even, $t_b = 2N$, $N > 0$, the probability of $\delta_b = 0$ is equal to

$$\Pr(\delta_b = 0) = \frac{(2N)!}{(N!)^2} (p_\beta(1 - p_\beta))^N.$$

This expression is decreasing in N and converges to zero as $N \rightarrow \infty$. Hence,

Proposition 1. *If all public signals are \tilde{b} , the probability of disagreement is*

1. *one if the number of public signals is odd;*

2. positive and increasing in N if the number of public signals is even.

As a measure of intensity of disagreement, we now consider the absolute value of the difference between the beliefs about the optimal action conditional on different private signals. Let

$$q^1(\text{Even}|\delta_b) = \Pr(\beta = 1|\delta_b) \Pr(\alpha = 1|a = 1) + (1 - \Pr(\beta = 1|\delta_b))(1 - \Pr(\alpha = 1|a = 1))$$

denote the probability that type 1 assigns to the event that the optimal action is $a = \text{Even}$, conditional on δ_b . Define $q^0(\text{Even}|\delta_b)$ analogously for type 0.

Definition *The absolute value of the disagreement between the beliefs about the optimal action is*

$$\Delta(\delta_b) = |q^1(\text{Even}|\delta_b) - q^0(\text{Even}|\delta_b)| = (2p_\alpha - 1) \left| 2 \frac{p_\beta^{\delta_b}}{p_\beta^{\delta_b} + (1 - p_\beta)^{\delta_b}} - 1 \right|.$$

The following proposition states that the expected value of disagreement, conditional on the realized state, is increasing in the number of signals.

Proposition 2. *For any $n > 0$, the expected absolute value of disagreement conditional on the realized state satisfies*

$$E \{ \Delta(\delta_b) | t_b + 1, \alpha, \beta \} = E \{ \Delta(\delta_b) | t_b, \alpha, \beta \}, \quad \text{if } t_b = 2n + 1, \quad (2)$$

$$E \{ \Delta(\delta_b) | t_b + 1, \alpha, \beta \} > E \{ \Delta(\delta_b) | t_b, \alpha, \beta \}, \quad \text{if } t_b = 2n. \quad (3)$$

3.4 Public signals \tilde{a} and \tilde{b}

We now turn to the setting in which agents observe signal \tilde{b} in odd periods and signal \tilde{a} in even periods.

The probability of disagreement in this environment is non-monotone, due to discreteness of the signals, but it decreases after every four periods and converges

to 0 as the number of signals becomes large. This is because the uncertainty about dimension α becomes small as more signals are realized. We can formally state this as follows. Let $z(t_a, t_b)$ denote the probability of disagreement, where $t_a > 0$ and $t_b > 0$ are respectively the number of public signals \tilde{a} and \tilde{b} .

Proposition 3. *Let public signals be of both types. Then,*

$$z(t_a, t_b) = \Pr(\delta_a^1 = 1|t_a)(1 - \Pr(\delta_b = 0|t_b))$$

and

$$z(t_a, t_b) > z(t_a + 2, t_b + 2).$$

Furthermore, $\lim_{t_a, t_b \rightarrow \infty} z(t_a, t_b) = 0$.

We now consider the absolute value of the difference between the beliefs about the optimal action conditional on different private signals. Let

$$q^1(\text{Even}|\delta_b, \delta_a^1) = \Pr(\beta = 1|\delta_b) \Pr(\alpha = 1|\delta_a^1) + (1 - \Pr(\beta = 1|\delta_b))(1 - \Pr(\alpha = 1|\delta_a^1))$$

be the probability that type 1 assigns to the event that the optimal action is $a = \text{Even}$, again conditional on δ_b . Define $q^0(\text{Even}|\delta_b, \delta_a^0)$ analogously for type 0.

The absolute value of the difference between beliefs about the optimal action of different types is equal to,

$$\Delta(\delta_b, \delta_a^1, \delta_a^0) = |q(\text{Even}|\delta_b, \delta_a^1) - q(\text{Even}|\delta_b, \delta_a^0)|.$$

The expected value of $\Delta(\delta_b, \delta_a^1, \delta_a^0)$ is non-monotone. Furthermore, it does not have to decrease every four periods. For example, if the signals about α are not very informative, i.e., p_α is close to $1/2$, while the signals about β are sufficiently informative, e.g., $p_\beta = 3/5$, then the expected value of $\Delta(\delta_b, \delta_a^1, \delta_a^0)$ will be *increasing* in the initial periods.

Nevertheless, after sufficiently many public signals the uncertainty about dimension α vanishes and, as a result, the expected value of $\Delta(\delta_b, \delta_a^1, \delta_a^0)$ converges to 0.

Proposition 4. *Let public signals be of both types. Then,*

$$\lim_{t_a, t_b \rightarrow \infty} E \{ \Delta(\delta_b, \delta_a^1, \delta_a^0) | t_a, t_b, \alpha, \beta \} = 0.$$

Proof. The result follows from Theorem 1 in Freedman (1963). □

3.5 More general environments

Our results above can easily be derived in more general environments, and do not depend on specific prior beliefs, signals being binary, or even on the shapes of the underlying distributions. The essential feature of the model is simply that the optimal action depends on relative values of different dimensions of the information space and, as such, contains at least one fewer degree of freedom. While individuals may agree on what the evidence indicates on several dimensions, they may disagree on the implications of this evidence if they have differential private information on other dimensions.

Consider for instance, two members of an electorate with differing private information on the incumbent politician that are each learning about the challenger at the same time and with equal precision. They can both agree on a distribution of beliefs about an index of the challengers quality, say $x \in [0, 1]$. However, if ones private information on the incumbents quality index, y_1 , is skewed toward 1 while the others index, y_0 , is skewed toward 0, more precise information on the challenger may not draw the two sides together as long as $Ey_0 < Ex < Ey_1$. Note, to get this result we did not need to specify the dimensionality of the information or the probability distributions on information, only that the ultimate choice (which candidate is better) has at least one fewer dimensions than the information used to make that choice (the relative quality of various candidates).

3.6 Agreeing to disagree

If there is common knowledge of rationality and the agents hold common priors, then disagreement cannot persist: Common knowledge of disagreement is impossible

(Aumann 1976) and communication of posterior beliefs is sufficient to achieve common beliefs (Geanakoplos and Polemarchakis 1982). Aumann, of course, meant his theorem to be a critique, rather than a defense, of the common prior and the common knowledge of rationality assumptions (Aumann (1976), pp. 1237-1238).

Imagine the situation in subsection 3.3 in which m individuals have observed their own private signal \tilde{a} and a series of public signals \tilde{b} . Based on the public signals, each has taken an action that, under individual rationality, should be consistent with their beliefs. Moreover, imagine sufficient public draws have been made such that these actions would fully reveal each person's private information, if she is rational. Suppose that individuals then can "communicate" in the sense that they see each other's private actions. In theory, this should lead to full agreement about the optimal action, even if subjects disagreed prior to this. This is a test we perform in the experiment described in the next section.

What if disagreement persists in the light of this style of communication? How can we make sense of this? To retain falsifiability of the model, we must maintain the assumption of individual rationality (otherwise any outcome can be made consistent with the model). Since the fundamentals of the joint probability distribution are straightforward, it seems that in this context, abandoning the assumption of common priors about the fundamentals would be unsatisfying. The remaining avenue is to relax common knowledge of rationality.

As an illustration, consider the following simple example. There are two individuals and imagine that each individual is rational and, furthermore, she believes that the other individual is rational with probability $\rho \in [0, 1)$ and makes random choices with complementary probability. Furthermore, these beliefs are common knowledge among individuals. That is, while everyone is rational, there is no common knowledge of rationality. Then, at the extreme $\rho = 0$, the agents will be unable to learn anything from each other's actions. For the other values of ρ , the agents' beliefs will moderate toward each other, but not agree.

Note that this explanation does not rely on non-common priors about the fundamentals. In the model above and in the experiment to follow, we have controlled for common priors. This indicates to us that many interesting questions now open up for understanding diverging opinions and persistence of disagreement.

4 Experimental Design

The experiment was conducted using undergraduate subjects in 8 sessions involving 6 subjects each. Each session involved three sets. In each set, we randomly selected one of the four urns, A , B , C , and D , corresponding respectively to the four states of the world $(0, 0)$, $(1, 1)$, $(0, 1)$, and $(1, 0)$ in our model. Figure 1 illustrates how the urns were presented to the subjects.

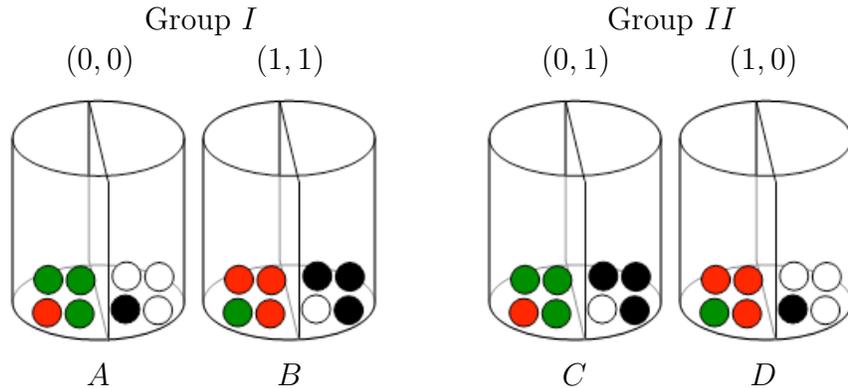


Figure 1: Content of the urns

There were two compartments in each urn. One compartment had red and green balls. The urns A and C (states $(0, 0)$ and $(0, 1)$) had three green and one red balls. The urns B and D (states $(1, 1)$ and $(1, 0)$) had one green and three red balls. A random draw from this compartment is equivalent to a signal $\tilde{\alpha}$ with the support $\{Green, Red\}$, whose distribution is given by $p_\alpha = \Pr(Red|\alpha = 1) = \Pr(Green|\alpha = 0) = 3/4$.

The other compartment contained white and black balls. The urns A and D (states $(0, 0)$ and $(1, 0)$) had three white and one black balls. The urns B and C

(states $(0, 1)$ and $(1, 1)$) had one white and three black balls. A random draw from this compartment is equivalent to a signal \tilde{b} with the support $\{White, Black\}$, whose distribution is given by $p_\beta = \Pr(White|\beta = 1) = \Pr(Black|\beta = 0) = 3/4$.

In every set, each subject observed a total of 15 draws with replacement from the selected urn. First, each subject observed one private draw from the compartment containing red and green balls (signal \tilde{a}). After that, subjects commonly observed 14 public draws. There were two types of sets: joint and separate. In the joint set, public draws were equally likely to be made from either of the compartments (signals \tilde{a} and \tilde{b}).⁸ In the separate set, public draws were made only from the compartment containing white and black balls (signal \tilde{b}).

The urns were divided in two groups. Group 1 consisted of urns A and B (action *Even*) and Group 2 consisted of urns C and D (action *Odd*). To infer subjects' beliefs, subjects placed bets on which Group they thought the urn was in. There were 16 rounds of bets for each set. First, subjects could place bets after each one of the 15 draws. In addition, after the bets on the 15th draw, the total cumulative numbers of bets on each group by all the participants were revealed, and the subjects could make bets one more time. The purpose of the 16th round of bets was to allow subjects to update their beliefs based on the information contained in the aggregate of all bets. If subjects are all perfect risk-neutral Bayesians and this is common knowledge, then there should be no disagreement on the round 16 bets.

In each round of bets, subjects could place from 0 to 9 bets on *each* of the groups of urns. That is, after each draw they could simultaneously make up to 18 bets, 9 or

⁸In the 8 sessions of the experiment there were a total of 12 joint sets. In the first three of these sets, the separator between the compartments was removed and the public draws were made from the common pool containing balls of all four colors. Nonetheless, some of the sequences contained majority of balls from one of the compartments blurring the contrast between joint and separate sets. Therefore, in the remaining 9 sets, we did not remove the separator and instead alternated draws between the two compartments. Unfortunately, we overlooked this problem during the initial design of the experiment. Comparing the behavior of the two subsamples, however, finds no discernible effect on choices. Furthermore, note that this issue biases our experiment *against* finding a difference between joint and separate sets and hence strengthens our results.

fewer bets on Group 1 and 9 or fewer bets on Group 2.

In the end of each set, one of the 16 rounds of bets was selected at random to determine the earnings of the subjects in that set. The subjects were given 10 points for every successful bet in this round, that is, the bet on the group which contained the urn used in this round. The bets made in this round also entailed costs. The first bet made on Group 1 in this round cost one point. The cost of each additional bet on Group 1 was one point more than the cost of the previous bet on Group 1. That is, the n th bet cost n points. Similarly, the 1st bet made on Group 2 in this round cost 1 point and the n th bet cost n points. If individuals are risk neutral Bayesian payoff maximizers, then bets should be revealing of their beliefs about the probabilities, that is, this is a proper scoring rule. For instance, a subject who thinks the likelihood is 0.35 that Group 1 is the true state, and 0.65 that it is Group 2, should place 3 bets on Group 1 and 6 bets on Group 2. Total bets across the two groups should always be 9 or 10.

To our knowledge, this paper's scoring rule is unique. It turns out however, that the incentives of our rule are precisely those of the Quadratic Scoring Rule. To see this, let b_i be the bets on state i and p_i beliefs that i is the true state. In our design, if the marginal cost of another bet on state i is b_i and the expected payoff is $10p_i$, then one should clearly stop betting when $b_i > 10p_i$.

To see this is the same as the Quadratic Scoring Rule, write the expected payoff of our task:

$$\mathbb{E}\pi = 10(p_1b_1 + p_2b_2) - \sum_{j=0}^{b_1} j \, dj - \sum_{k=0}^{b_2} k \, dk.$$

Rewrite this as a continuous choice of b 's by replacing the sums with integrals:

$$\begin{aligned} \mathbb{E}\pi &= 10(p_1b_1 + p_2b_2) - \int_{j=0}^{b_1} j \, dj - \int_{k=0}^{b_2} k \, dk \\ &= 10(p_1b_1 + p_2b_2) - b_1^2/2 - b_2^2/2 \end{aligned}$$

This is precisely the quadratic scoring rule. Rather than giving our subjects the quadratic function, as prior researchers have done, we gave them the first order con-

dition of the quadratic scoring rule, which is a simple linear problem that subjects hopefully find more tractable.

There are eight possible permutations of the sequence of three sets, each of which could be either joint or separate: (separate, separate, separate), (joint, separate, separate), etc. We conducted eight sessions with three sets each, one session for every possible permutation of sets. Hence, there were 12 separate sets and 12 joint sets. In one set, each of the six subjects made 16 rounds of bets on two groups. There was a total of 48 subjects, 6 subjects in each of eight sessions. We obtained a total of 4,608 observations ($16 \text{ rounds} \times 3 \text{ sets} \times 6 \text{ subjects} \times 8 \text{ sessions} \times 2 \text{ groups}$).

To guarantee that earnings were non-negative, subjects were endowed with 45 points in each set. Subjects kept track of the draws, their bets, and their earnings in each round using a computer interface. The information about the rules of the experiment and the content of the urns was known to the participants.

The experiment lasted about one hour. We paid US \$1 for each 10 points earned in the experiment. The subjects were anonymously paid their cumulative earnings in cash at the end of the experiment. The earnings averaged \$20.52 (standard deviation \$3.98), ranging from \$9 to \$27. Subjects' instructions are available in the appendix.

5 Experimental Results

In the next two subsections, we present our evidence for Propositions 1 and 3 on the frequency of disagreement and for Propositions 2 and 4 on the expected value of disagreement. In the analysis, we excluded two sets, one joint set and one separate set, in which all subjects observed the same private information.⁹ Our model predicts the overall pattern in the data well, with several important exceptions. Subjection 5.3 will examine our predictions about choices in the first round of each set, and subsection 5.4 will discuss our predictions for round 16, when we reveal the sum of

⁹We do *not* exclude any subjects from the analysis, even the subjects who, regardless of the realized information, have placed the same combination of bets in all rounds in all sets.

all prior bets by all players.

5.1 Probability of disagreement

Figure 2 depicts the frequency of the disagreement, for both theoretical and observed and both joint and separate sets, over the first 15 rounds.¹⁰ The observed frequency has an increasing trend in separate sets and a decreasing trend in joint sets, in accordance with the theoretical predictions. At the same time, the observed frequency of disagreement in the first round is significantly larger than the theoretical value of zero. We will consider the bets made by the subjects in the first period in detail in Section 5.3.

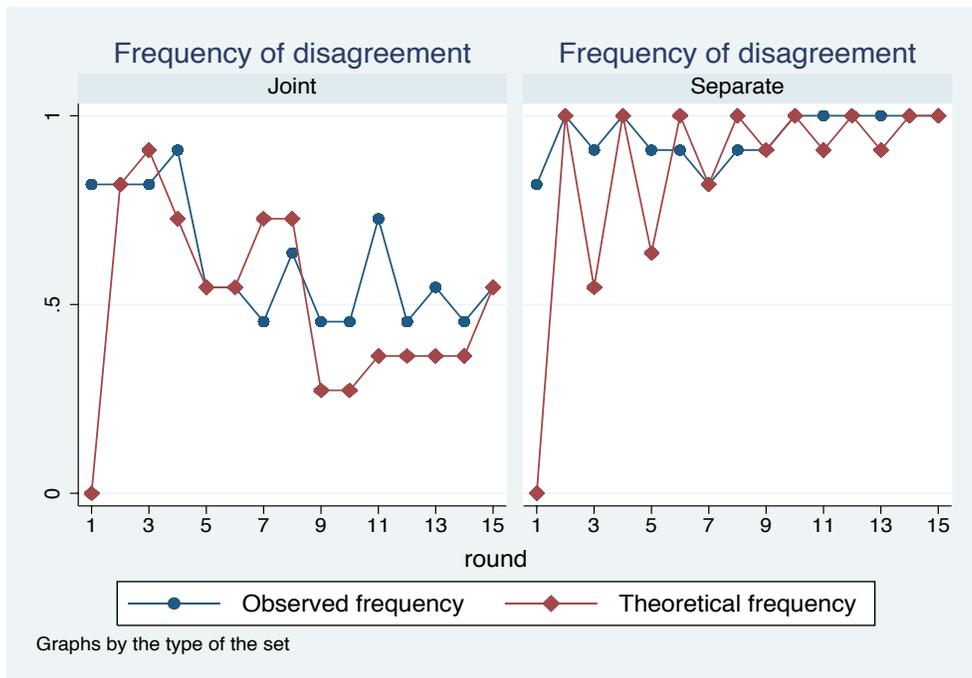


Figure 2: The expected and observed frequency of disagreement. In *joint* sets public signals are of both types, in *separate* sets public signals are *Black* and *White*.

How is Figure 2 generated? To determine the theoretical frequency of disagree-

¹⁰We exclude the 16th round because in this round the information observed by the subjects is the cumulative number of bets made by all subjects in the previous rounds. The Bayesian model alone cannot determine the combination of bets that maximize the expected payoff; this combination depends on the beliefs of the subject about how the other subjects make their bets.

ment, for each subject we calculate Bayesian beliefs about whether the urn belongs to Group I given the subject's information. The theoretical frequency of disagreement is the frequency with which Bayesian beliefs of the subjects with different private information disagree in a given round in a given type of set.

There are multiple ways to define the observed frequency of disagreement. We have chosen the following definition. First, for each subject we calculate the difference between the bets made on Group I and on Group II. If the difference between the bets is larger than one, we say that the subject prefers the group on which she places more bets. Otherwise we say that the subject is indifferent.¹¹

Next, we determine the group preferred by the majority of subjects who observe the same private information as follows: we exclude the subjects who are indifferent between the groups, and find the group preferred by the majority of the remaining subjects with strict preference. We say that the subjects are, on average, indifferent about which group to bet on if equal numbers of subjects strictly prefer different groups or if all subjects are indifferent. The observed frequency of disagreement is, then, the frequency with which the subjects who observe different private information prefer different groups: that is, the (majority of the) subjects with one private signal prefer one group and the subjects with the other private signal either prefer the other group or are, on average, indifferent about which group to bet on.¹²

¹¹One alternative we have considered is to define a subject to be indifferent between groups if and only if she places the same number of bets on both groups. The disadvantage of this definition is that it may incorrectly classify subjects as not indifferent. Imagine that a risk-neutral subject believes that both groups are equally likely. Then, she is willing to pay up to 5 points for a bet on each of the groups. Because, in our experiment, the 5th bet costs 5 points, the subject is indifferent between placing 4 or 5 bets on each of the groups. Hence, the following combinations of bets are consistent with the belief that both groups are equally likely: (4,4), (5,5), (4,5), and (5,4).

¹²In order to count, the direction of disagreement does not have to coincide with the theoretical prediction, although it does in almost every case. Deviations, therefore, work against hypothesis. Also, note that our definition of preference for a group applies regardless of whether the bets made by the subject in question are consistent with payoff-maximizing behavior. Hence, our conclusion that the observed frequency of disagreement is close to the theoretical frequency of disagreement *does not imply* that our theoretical model can explain the bets of the subjects.

5.2 Expected value of disagreement

Our first task is to describe how we determine the theoretical and observed values of disagreement. To determine the theoretical value of disagreement, for each subject we calculate Bayesian beliefs about whether the urn belongs to Group I for each subject given the subject's information. The theoretical value of disagreement is the absolute value of the difference of Bayesian beliefs, multiplied by 10, about the event that the urn belongs to Group I for subjects with different private information in a given round, averaged over all sets of a given type. We determine the observed value of disagreement as follows. First, for each round in each set and for each group we calculate the absolute value of the difference between the average bets made by the subjects who observe different private information. The average of this value in a given round over both groups and all sets of a given type is the observed value of disagreement.

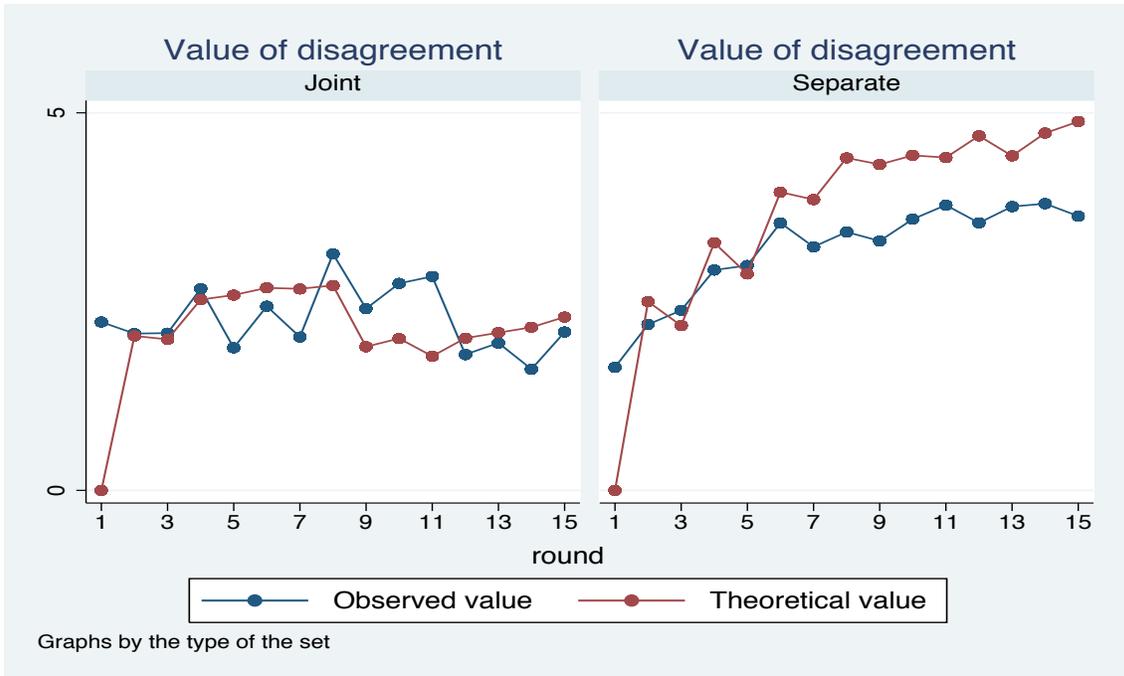


Figure 3: The expected and observed value of disagreement. In *joint* sets public signals are of both types, in *separate* sets public signals are *Black* and *White*.

Figure 3 depicts the expected absolute value of disagreement, both theoretical and observed, for the first 15 rounds, averaged for each of the joint and separate sets. The observed value has an increasing trend in separate sets, in accordance with the theoretical predictions. In joint sets, both the observed and the theoretical value of disagreement do not increase. As with the frequency of disagreement, the observed value of disagreement in the first round is larger than the theoretical value of zero. We explore this deviation more closely next.

5.3 First round

In the first round a decision maker has information on only one dimension of the information, therefore a Bayesian decision maker should believe that both groups are equally likely. As a result, subjects with different private information should not disagree.¹³

In our experiment, this means that in the first round subjects' bets should be symmetric, that is, they should place 4 or 5 bets on each state. Given this, there are three types of errors a subject can make. The first two respect symmetry, but bet either more than 10 bets in total, or fewer than 8, but bets on each state do not differ by more than 1. Notice, however, that these two errors will not cause disagreement, but just result in suboptimal payment outcomes for the subjects. The third kind of error is to show a clear favorite in round 1, that is, to have bets that are asymmetric and that differ by two or more. This asymmetric error will cause disagreements, even in round 1 where we predict none to exist. Figure 4 shows the proportions of both correct and incorrect choices on round 1.

Figure 4 shows three interesting patterns. First, the fraction of individuals making correct bets is increasing over the sets. Second, the number of symmetric

¹³Recall that to calculate the frequency of disagreement, we say that a subject is indifferent between the groups if her bets on the groups differ by at most one. If the difference between the bets on different groups is more than or equal to two, we say that the subject prefers the group on which she places majority of her bets.

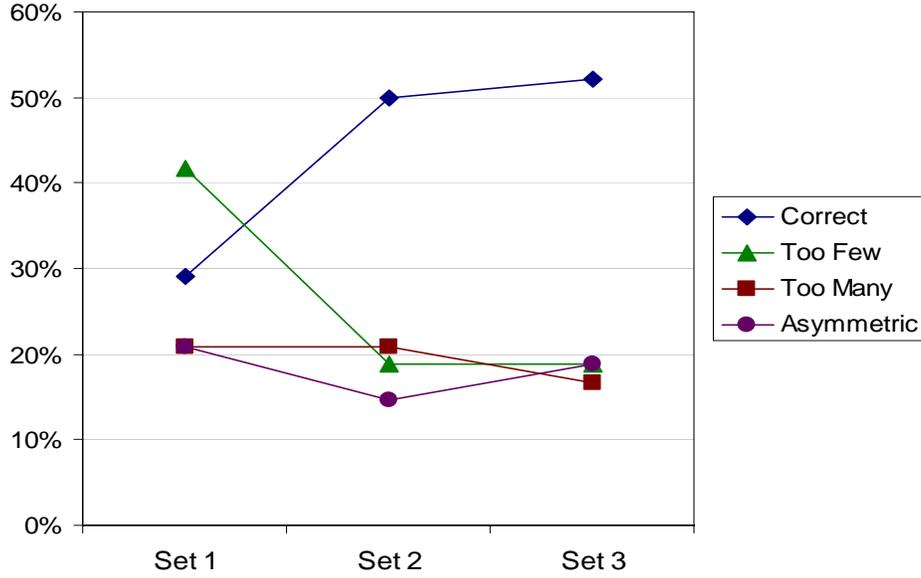


Figure 4: Errors in first round betting. Errors of too many or two few both respect a neutral prior but make the wrong number of bets, while asymmetric bets show a clear preference despite the evidence suggesting indifference.

errors (both too many and two few bets) is going down over time. Throughout the experiment, however, around 80% of the subjects correctly agree on round 1 by placing equal bets on both states. Third, about 20% of subjects make the error of placing asymmetric bets in round 1, and this fraction remains steady throughout the experiment.

This means that the first component necessary for agreement in round sixteen is missing for a at least 20% of the subjects. That is, this fraction seems to incorrectly disagree with the remaining subjects even when their private information is not sufficient to cause disagreement.

It is also easy to show that risk aversion cannot explain differences in betting if people believe both states are equally likely.¹⁴

¹⁴Note that risk aversion implies $\mathbb{E}\pi = p_1 u(10b_1 - C(b_1, b_2)) + (1 - p_1)u(10b_2 - C(b_1, b_2))$ where $C(b_1, b_2) = b_1^2/2 + b_2^2/2$. Since we *do not* restrict $b_1 + b_2 = 1$, the two first-order conditions are

$$\begin{aligned}
 p_1 u'_1 \cdot (10 - b_1) - (1 - p_1) u'_2 \cdot b_1 &= 0 \\
 -p_1 u'_1 \cdot (b_1) + (1 - p_1) u'_2 \cdot (10 - b_2) &= 0,
 \end{aligned}$$

5.4 Round Sixteen

The data from round 16 allows us to test a fundamental prediction of the model: If there is common knowledge of rationality, then when people can aggregate others' information they should no longer disagree.

Recall that in round 16, the last round in a set, no balls were drawn. Instead, the subjects were informed about the total number of bets on each of the groups made in the previous rounds by all subjects in the current set. Then, the subjects placed their bets once again. Of course, if all subjects were Bayesian decision makers, then it would be impossible to have common knowledge of disagreement.¹⁵

We found, contrary to the prediction, that in round 16 there was disagreement in 9 out of 24 sets. The more important question, however, is how frequently in the last round subjects prefer the group on which there was a majority of bets in all previous rounds? We found that there were 16 disagreements with the majority over all 144 individual observations. Separating these out by set, we find 2 disagreements in set 1, 5 in set 2, and 9 in set 3.

Why would subjects not place most of their bets on the group that received the majority of bets in the previous periods? One possibility is that the subjects expect the information about the majority of bets to be a noisy signal and therefore might place a higher weight on their own opinion. We explore this hypothesis in Table 1 where we regress the difference between one's own bets on Group I and Group II in round 16 on the cumulative differential in own bets and others' bets in the prior 15 rounds, with standard errors clustered for each subject. The assumption of common knowledge of rationality would lead to a prediction that that subjects should base

which implies

$$\begin{aligned}\frac{p_1}{1-p_1} &= \frac{u'_2}{u'_1} \frac{b_1}{10-b_1}, \text{ and} \\ &= \frac{u'_2}{u'_1} \frac{10-b_2}{b_2}\end{aligned}$$

Notice, if $p_1 = p_2 = 1/2$, then regardless of risk aversion, $b_1 = b_2 = 5$ is the only solution.

¹⁵More precisely, this follows from Theorem 3 in Nielsen, Brandenburger, Genakoplos, McKelvey, and Page (1990).

their round 16 bets on the total rounds 1-15 bets, including their own and others, on each group. That is, as long as all subjects are assumed to be rational, then each prior bet is an equally valid piece of information, whether it is information seen directly or inferred from the actions of others. This means that in the regressions in Table 1 we should see coefficients on own and the others' differentials to be equal—all information is treated the same. We see from column 1, however, that own experience receives more than twice the weight given other's experience (0.039 versus 0.015, $p < .01$). Nevertheless, in our experiment the majority of bets indicated the winning group in 21 out of 24 sets.¹⁶

What is the root of this over-weighting of one's own information? This is explored in column 2 of Table 1. Here we interact own and others' bet differentials with the three types of errors we saw in the prior subsection. We see that subjects who commit errors in round 1 that respect symmetry of bets have no differential impact on the weights given to own and other's experience. However, those who make asymmetric bets in round 1 tend to put weight on their own experience ($0.028 + 0.046 = 0.074$) that is more than 10 times what they put on the experience of others ($0.019 - 0.012 = 0.007$) (while economically small, the estimate of 0.007 is nonetheless statistically significant ($t = 1.92$)). This suggests that 20% of the sample will virtually ignore the information on others in determining their round 16 bets.

¹⁶It is interesting to note that the magnitudes of this effect mirror those of Weizsacker (2010) who finds, in a meta-analysis of experiments on information cascades, subjects appear to weigh their own information about twice that of others when, according rational expectations, the weights should be equal. As in our studies, this is also related to the degree of errors among others, but on average results in losses for a significant share of subjects. This comparison across games, we believe, underscores the potential value of our finding.

Table 1: Regressions of Round 16 Difference in Group I and Group II Bets on Own and Other's Cumulative Differentials from rounds 1-15.

Coefficient, Standard Errors, and p -values.

Independent Variable	(1)	(2)
Own Differential	0.039*** 0.010 ($p = 0.000$)	0.028* 0.016 ($p = 0.095$)
Others' Differential	0.015*** 0.003 ($p = 0.000$)	0.019*** 0.004 ($p = 0.000$)
Set 2	3.734*** 0.710 ($p = 0.000$)	3.924*** 0.803 ($p = 0.000$)
Set 3	1.396* 0.831 ($p = 0.100$)	1.594* 0.849 ($p = 0.067$)
<i>Interactions with Round 1 Errors:</i>		
Own Differential		
× Too Few		0.014 0.024 ($p = 0.561$)
× Too Many		0.007 0.020 ($p = 0.714$)
× Asymmetric		0.046** 0.018 ($p = 0.016$)
Others' Differential		
× Too Few		-0.006 0.005 ($p = 0.235$)
× Too Many		0.001 0.005 ($p = 0.799$)
× Asymmetric		-0.012*** 0.004 ($p = 0.007$)
Observations	144	144
Clusters	48	48
R^2	0.764	0.787

Notes: Standard errors clustered at individual level. Significance level: * 10% , ** 5%, *** 1%.

What about those who make correct bets in round 1? These subject also overweight their own information, although the bias is not statistically significant. As shown in column 2 of Table 1, own experience has a coefficient of 0.028 versus a 0.019 coefficient on the bets of others, $t = 0.46$, thus suggesting a slight but insignificant

bias towards own informations.¹⁷ This indicates that the differential weighting of own and others' bets found in column (1) is due those making errors of asymmetry in rounds 1. This suggests a bias in learning from others, but that the bias tends to be concentrated in fraction of the population (about 20% in our sample) who make systematic departures in Bayesian inference at the outset of each set and who deeply discount the information of others.

6 Conclusions

This paper is concerned with polarization of individual opinions that occurs as an optimal response to additional public information and persistence of such polarization even after the public information becomes sufficient to remove any disagreement.

We present a simple environment in which private opinions can diverge in response to additional public information. The important feature of this environment, which we believe is relevant in practice, is that that the information is multidimensional and the optimal action depends on the *relative* value of information on different dimensions.

We demonstrate in our experiment that polarization may become commonly known and persist even after there is sufficient information for it to disappear. This result suggests that polarization of opinions can be a lasting phenomenon and communication or debate might have limited effectiveness in aggregating information and attaining agreement.

Finally, persistence of polarization is sensitive to the source of information. In our experiment, polarization persists in the experimental treatment because subjects undervalue the information of others when it must be gleaned from their choices. By contrast, in the control group, in which the information intended to resolve disagree-

¹⁷Similar findings exist in both the psychology and finance literatures. Krueger and Dunning (1999) for instance show that errors that lead people to make imprecise probabilistic assessments also interfere with their ability to judge their own relative performance. Chen and Jiang (2006) show that financial analysts overweight their private information when issuing forecasts.

ment is provided directly by experimenter, polarization disappears.

The fact that disagreement depends on the means by which the information is provided raises a number of interesting questions about individual cognition, the nature of inference, and the importance of debate. For instance, why do people appear to systematically put too much weight on the information they receive directly, rather than indirectly through another's actions? Why do people tend to believe they are more rational decision makers than others? Finally, are there ways of structuring debates to overcome these biases so that individuals can share information, incorporate it into posteriors, and willfully revise their opinions? Our results indicate that these are difficult yet valuable topics for future research.

Appendix 1: Proofs omitted in the text

Proof of Proposition 2. Assume that the realized state is $(\alpha, \beta) = (1, 1)$. The proof for the other cases is analogous. For a given (α, β) , define

$$\gamma(t_b) = \frac{E\{\Delta(\delta_b)|t_b, \alpha, \beta\}}{2p - 1}.$$

We have

$$\begin{aligned} \gamma(2n+1) &= \sum_{i=0}^{2n+1} \frac{(2n+1)!}{(2n+1-i)!i!} p^{2n+1-i} (1-p)^i \left| 2 \frac{p^{2n+1-i} (1-p)^i}{p^{2n+1-i} (1-p)^i + p^i (1-p)^{2n+1-i}} - 1 \right| \\ &= \sum_{i=n+1}^{2n+1} \frac{(2n+1)!}{(2n+1-i)!i!} v(i, 2n+1), \quad (n \geq 0) \end{aligned}$$

where

$$v(i, l) = p^i (1-p)^{l-i} - p^{l-i} (1-p)^i, \quad (i \geq 0, l \geq 1).$$

Similarly,

$$\gamma(2n) = \sum_{i=n}^{2n} \frac{(2n)!}{(2n-i)!i!} v(i, 2n), \quad (n > 0).$$

Note that for all $n \geq 1$,

$$v(n, 2n) = 0, \tag{A1}$$

$$v(i, 2n-1) = v(i, 2n) + v(i+1, 2n). \tag{A2}$$

Before we proceed with the proof, recall the following useful fact

$$\frac{N!}{(N-i)!i!} = \frac{(N-1)!}{(N-i)!(i-1)!} + \frac{(N-1)!}{(N-1-i)!i!}, \quad (0 < i < N, N \geq 2). \tag{A3}$$

We now prove (2). First, it is straightforward to check that $\gamma(2) = \gamma(1)$. Now,

let $n \geq 1$. Using (A3), we can write

$$\begin{aligned}
\gamma(2n) &= \sum_{i=n}^{2n-1} \left[\frac{(2n-1)!}{(2n-1-i)!i!} + \frac{(2n-1)!}{(2n-i)!(i-1)!} \right] v(i, 2n) + v(2n, 2n) \\
&= \sum_{i=n}^{2n-1} \frac{(2n-1)!}{(2n-1-i)!i!} v(i, 2n) + \sum_{i=n-1}^{2n-2} \frac{(2n-1)!}{(2n-1-i)!i!} v(i+1, 2n) + v(2n, 2n) \\
&= \sum_{i=n}^{2n-1} \frac{(2n-1)!}{(2n-1-i)!i!} (v(i, 2n) + v(i+1, 2n)) + \frac{(2n-1)!}{n!(n-1)!} v(n, 2n).
\end{aligned}$$

Then, it follows from (A1) and (A2) that

$$\gamma(2n) - \gamma(2n-1) = 0.$$

Next, we prove (3). Using (A3),

$$\begin{aligned}
\gamma(2n+1) &= \sum_{i=n+1}^{2n} \left[\frac{(2n)!}{(2n-i)!i!} + \frac{(2n)!}{(2n+1-i)!(i-1)!} \right] v(i, 2n+1) + v(2n+1, 2n+1) \\
&= \sum_{i=n+1}^{2n} \frac{(2n)!}{(2n-i)!i!} v(i, 2n+1) + \sum_{i=n}^{2n-1} \frac{(2n)!}{(2n-i)!i!} v(i+1, 2n+1) \\
&\quad + v(2n+1, 2n+1) \\
&= \sum_{i=n+1}^{2n} \frac{(2n)!}{(2n-i)!i!} (v(i, 2n+1) + v(i+1, 2n+1)) + \frac{(2n)!}{n!n!} v(n+1, 2n+1).
\end{aligned}$$

Then, from (A1) and (A2), we get

$$\gamma(2n+1) - \gamma(2n) = \frac{(2n)!}{n!n!} v(n+1, 2n+1) > 0.$$

□

Proof of Proposition 3. We provide a proof of the first part of the proposition for the case of $t_a = t_b = 2n$, $n \geq 1$. (The argument for the remaining cases is analogous.) Because t_a and t_b are even, the probability of disagreement is given by

$$z(t_a, t_b) = z(n) = \Pr(\delta_a^1 = 1|t_a)(1 - \Pr(\delta_b = 0|t_b)). \quad (\text{A4})$$

Set $q_k = p_k(1 - p_k)$ where $k = \alpha, \beta$, and note that $q_k \in [0, 1/4]$. Then, (A4) can be rewritten as

$$z(n) = \frac{(2n)!}{(n!)^2} q_\alpha^n \left(1 - \frac{(2n)!}{(n!)^2} q_\beta^n \right).$$

It follows that

$$\begin{aligned} z(n+1) - z(n) &= \frac{(2n+2)!}{(n+1)!^2} q_\alpha^{n+1} \left(1 - \frac{(2n+2)!}{(n+1)!^2} q_\beta^{n+1} \right) - \frac{(2n)!}{(n!)^2} q_\alpha^n \left(1 - \frac{(2n)!}{(n!)^2} q_\beta^n \right) \\ &= q_\alpha^n \frac{(2n)!}{(n!)^2} g(n), \end{aligned}$$

where

$$g(n) = \left(\frac{(2n+2)(2n+1)}{(n+1)^2} q_\alpha \left(1 - \frac{(2n+2)!}{(n+1)!^2} q_\beta^{n+1} \right) - \left(1 - \frac{(2n)!}{(n!)^2} q_\beta^n \right) \right).$$

Before we proceed with the proof, note that for any $q_\beta \in [0, 1/4]$ and any $n \geq 1$

$$\left(1 - \frac{(2n+1)^2}{(n+1)^2} q_\beta \right) q_\beta^n \leq \frac{n^n (n+1)^{n-1}}{(2n+1)^{2n}} \quad (\text{A5})$$

Furthermore,

$$\frac{n^n (n+1)^{n-1}}{(2n+1)^{2n}} \leq \frac{1}{4^n (n+1)} \quad (\text{A6})$$

Now, we have

$$\begin{aligned} g(n) &\stackrel{q_\alpha < 1/4}{<} \frac{2n+1}{2(n+1)} \left(1 - \frac{(2n+2)!}{(n+1)!^2} q_\beta^{n+1} \right) - \left(1 - \frac{(2n)!}{(n!)^2} q_\beta^n \right) \\ &= -\frac{1}{2(n+1)} + \frac{(2n)!}{(n!)^2} \left(1 - \frac{(2n+1)^2}{(n+1)^2} q_\beta \right) q_\beta^n \\ &\stackrel{(\text{A5}), (\text{A6})}{\leq} \frac{1}{2(n+1)} \left(\frac{2(2n)!}{4^n (n!)^2} - 1 \right) \leq 0. \end{aligned}$$

We now turn to the limit result. It follows from (A4) that if $t_a = 2n$, then the probability of disagreement is bounded from above by $\Pr(\delta_a^1 = 1|t_a) = \frac{(2n)!}{(n!)^2} q_\alpha^n$. The bound converges to 0 as $n \rightarrow \infty$. Similarly, if $t_a = 2n+1$, the probability of disagreement is bounded from above by $\Pr(\delta_a^1 = 2|t_a) + \Pr(\delta_a^1 = 0|t_a) = \frac{(2n+1)!}{n!(n+1)!} q_\alpha^n$, which also converges to 0 as $n \rightarrow \infty$. \square

References

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2009): “Fragility of Asymptotic Agreement under Bayesian Learning,” .
- ANDREONI, J., AND T. MYLOVANOV (2010): “Diverging Opinions,” *working paper*.
- AUMANN, R. J. (1976): “Agreeing to Disagree,” *The Annals of Statistics*, 4(6), 1236–1239.
- BALIGA, S., E. HANANY, AND P. KLIBANOFF (2011): “Polarization and ambiguity,” *working paper*.
- BARBERIS, N., AND R. THALER (2003): “A survey of behavioral finance,” in *Handbook of the Economics of Finance*, ed. by G. Constantinides, M. Harris, and R. M. Stulz, vol. 1 of *Handbook of the Economics of Finance*, chap. 18, pp. 1053–1128. Elsevier.
- BLACKWELL, D., AND L. DUBINS (1962): “Merging of Opinions with Increasing Information,” *The Annals of Mathematical Statistics*, 33(3), 882–886.
- BOERGERS, T., A. HERNANDO-VECIANA, AND D. KRÄHMER (2009): “When are signals complements or substitutes?,” *working paper*.
- CHEN, Q., AND W. JIANG (2006): “Analysts’ weighting of private and public information,” *The Review of Financial Studies*, 19(1), 319–355.
- CRIPPS, M. W., J. C. ELY, G. J. MAILATH, AND L. SAMUELSON (2008): “Common learning,” *Econometrica*, 76(4), 909–933.
- DIXIT, A. K., AND J. W. WEIBULL (2007): “Political Polarization,” *Proceedings of the National Academy of Science of the United States of America*, 104(18), 7351–7356.

- EIL, D., AND J. RAO (2011): “The good news - bad news effect: Asymmetric processing of objective information about yourself,” *American Economic Journal: Microeconomics*, 3(2), 114–138.
- FREEDMAN, D. A. (1963): “On the Asymptotic Behavior of Bayes’ Estimates in the Discrete Case,” *The Annals of Mathematical Statistics*, 34(4), 1386–1403.
- GEANAKOPOLOS, J. D., AND H. M. POLEMARCHAKIS (1982): “We can’t disagree forever,” *Journal of Economic Theory*, 28(1), 192–200.
- GERBER, A., AND D. GREEN (1999): “Misperceptions about Perceptual Bias,” *Annual Review of Political Science*, 2, 189–210.
- HIRSHLEIFER, D. (2001): “Investor Psychology and Asset Pricing,” *Journal of Finance*, 56(4), 1533–1597.
- HOUSTON, D. A., AND R. H. FAZIO (1989): “Biased processing as a function of attitude accessibility: Making objective judgments subjectively,” *Social Cognition*, 7(1), 51–66.
- KANDEL, E., AND N. D. PEARSON (1995): “Differential interpretation of public signals and trade in speculative markets,” *Journal of Political Economy*, 103(4), 831–872.
- KATZ, E., AND J. J. FELDMAN (1962): “The Debates in Light of Research,” in *The Great Debates*, ed. by S. Kraus, pp. 173–223. Bloomington: Indiana University Press.
- KIM, O., AND R. E. VERRECCHIA (1997): “Pre-announcement and event-period private information,” *Journal of Accounting and Economics*, 24, 395–419.
- KINDER, D. R., AND WALTER R. MEBANE, JR. (1983): “Politics and Economics in Everyday Life,” in *The Political Process and Economic Change*, ed. by B. S. Frey, and K. R. Monroe, pp. 141–180. Algora Publishing.

- KONDOR, P. (2011): “The more we know on the fundamental, the less we agree on the price,” *working paper*.
- KRUEGER, J., AND D. DUNNING (1999): “Unskilled and unaware of it: how difficulties in recognizing one’s own incompetence lead to inflated self-assessments,” *Journal of Personality and Social Psychology*, 77(6), 1121–1134.
- LORD, C. G., L. ROSS, AND M. R. LEPPER (1979): “Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence,” *Journal of Personality and Social Psychology*, 37(11), 2098–2109.
- NARASIMHAN, C., C. HE, E. ANDERSON, L. BRENNER, P. DESAI, D. KUKSOV, P. MESSINGER, S. MOORTHY, J. NUNES, Y. ROTTENSTREICH, R. STAELIN, AND G. W. AND (2005): “Incorporating Behavioral Anomalies in Strategic Models,” *Marketing Letters*, 16(3), 361–373.
- NICKERSON, R. S. (1998): “Confirmatory Bias: A Ubiquitous Phenomenon in Many Guises,” *Review of General Psychology*, 2(2), 175–220.
- NIELSEN, L. T., A. BRANDENBURGER, J. GENAKOPLIS, R. MCKELVEY, AND T. PAGE (1990): “Common Knowledge of an Aggregate of Expectations,” *Econometrica*, 58(5), 1235–39.
- RABIN, M. (1998): “Psychology and Economics,” *Journal of Economic Literature*, 36(1), 11–46.
- RABIN, M., AND J. L. SCHRAG (1999): “First Impressions Matter: A Model Of Confirmatory Bias,” *The Quarterly Journal of Economics*, 114(1), 37–82.
- SANTOS-PINTO, L., AND J. SOBEL (2005): “A Model of Positive Self-Image in Subjective Assessments,” *American Economic Review*, 95(5), 1386–1402.

- SCHUETTE, R. A., AND R. H. FAZIO (1995): “Attitude Accessibility and Motivation as Determinants of Biased Processing: A Test of the MODE Model,” *Personality and Social Psychology Bulletin*, 21(7), 704–710.
- SEARS, D. O. (1968): “Political Behavior,” in *The Handbook of Social Psychology*, ed. by G. Lindzey, and E. Aronson, chap. 5, pp. 315–458. Reading, MA: Addison-Wesley, 2 edn.
- SETHI, R., AND M. YILDIZ (2009): “Public disagreement,” *working paper*.
- SIGELMAN, L., AND C. K. SIGELMAN (1984): “Judgments of the Carter-Reagan Debate: The Eyes of the Beholders,” *The Public Opinion Quarterly*, 48(3), 624–628.
- WEIZSACKER, G. (2010): “Do We Follow Others When We Should? A Simple Test of Rational Expectations,” *American Economic Review*, 100(5), 2340–60.
- WESTEN, D., P. S. BLAGOV, K. HARENSKI, C. KILTS, AND S. HAMANN (2006): “Neural Bases of Motivated Reasoning: An fMRI Study of Emotional Constraints on Partisan Political Judgment in the 2004 U.S. Presidential Election,” *Journal of Cognitive Neuroscience*, 18(11), 1947–1958.
- WILSON, A. (2005): “Bounded Memory and Biases in Information Processing,” *working paper*.
- ZIMPER, A., AND A. LUDWIG (2007): “On attitude polarization under Bayesian learning with non-additive beliefs,” *Journal of Risk and Uncertainty*, 39(2), 181–212.