DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CALIFORNIA 91125

Fake News, Information Herds, Cascades and Economic Knowledge

Lazarina Butkovich
Nina Butkovich
Charles Plott
Han Seo
California Institute of Technology

SOCIAL SCIENCE WORKING PAPER 1442

July 2018

Fake News, Information Herds, Cascades and Economic Knowledge

L. Butkovich, N. Butkovich, C. Plott and H. Seo
California Institute of Technology
July 2018

Abstract

The paper addresses the issue of "fake news" through a well-known and widely studied experiment that illustrates possible uses of economics and game theory for understanding the phenomenon. Public news is viewed as an aggregation of decentralized pieces of valuable information about complex events. Success of news systems rests on accumulated investment in trust in news sources. By contrast, fake news involves cases in which news source reliability is not known. The experiment demonstrates how fake news can destroy both the investment in trust and also the benefits that successful news systems provide.

A. Introduction[1]

While fake news, propaganda and misinformation are not newly emergent phenomena the public's recognition of their growth is found in many sources. Major news networks, e,g, CNN, NBC, and Fox News, consistently publish remarks about a competitor's lack of reliability. Different sides of the political divide accuse each other of promoting fake news, and the "mainstream" media is accused of taking sides. Analyses find different reasons for fake news growth including underlying economics that reward "click baits", or an erosion of moral fiber and even espionage. Against that background, sources express concerns about the consequences of the tendency. Bloomberg suggests direct damages to wealth through an effect of fake news on the stock market.[2] Others have expressed concerns about long term consequences. For example, politician Hillary Clinton points to the possibility of systemic damage due to a leveling effect or "false equivalency" in which the unreliable news sources and reliable sources are regarded as equals.[3]

The underlying concerns about the consequences of fake news rest on more than intuition or impressions about the damage it can cause. Hints of scientific support appear to exist. Simple experimental exercises illustrate that basic principles of information processes that are known to be features of markets can also operate for public sources of information and broader decisions influenced by news. News from public sources can enhance efficiency and create economic benefits for the same reasons that improved information enhances market efficiency. The experiments also illustrate that the benefits that emerge from reliable news can be

[1] The insights of David Grether and the help of B. Atsavapranee, are gratefully acknowledged. All experimental programs were designed by Travis Maron. The financial support of the John Templeton Foundation is gratefully acknowledged.

[2] (Cite: https://www.bloomberg.com/view/articles/2017-10-23/why-fake-news-is-so-harmful-to-investors)

[3] The comments are her assessment of academic research as contained in her 2018 Author Miller Freedom to Write Lecture at the 14th Annual Pen World Voices Festival, April 14, 2018.

systematically disrupted by fake news. The experiments explore a specific basic principle that suggests how and why associated benefits are lost.

The experiments are narrowly focused and thus do not address many of the important issues associated with the reliability of news reports. Of course, news is a deep part of social fabric, and policies that shape its role, such as the broad reporting latitudes permitted by freedom of the press, are widely recognized as a cornerstone of a free society.[4]

The experiments reflect the perspective of F. Hayek, "The Use of Knowledge in Society"[5], which is now integrated with economic theory and finance as well as political theory. The basic principles apply to a world in which information is produced from decentralized sources of observation. Information originates locally as a product of the actions of those close to the facts and with interests aligned with the use of the facts. The information becomes transferred and aggregated through publicly observed actions, such as trades and trading prices in which good information and incentives to act quickly are closely related. In a statistical sense, the information contained in scattered observations can become aggregated, pooled, and processed.[6] The resulting knowledge is a resource, a type of "public good", that can be used multiple times and grants economic value by preventing costly mistakes.

Trust in the reliability of information sources is fundamental to the success of news systems. Typically, trust is derived from an alignment between incentives and the information revealing actions. For example, people normally visit a restaurant because they like the food, and thus, the number of customers can be a source of information about the quality of the restaurant. Or, the location of fishermen suggests information about the location of fish. Analogously, in markets, the upward movement of a stock price suggests the possibility of favorable earnings because those who have uncovered favorable facts have an incentive to buy the stock before others do. When incentives and information revealing actions are aligned, the actions can reveal information that can be trusted and used. On the other hand, if incentives and actions are not aligned, the information revealed can become a type of "fake news" that is misleading and can lead to costly mistakes. Experience contributes to an understanding about the reliability of information. For instance, the diners and fishermen will ignore information if it is known to be unreliable. That is, if customers are known to be paid by the owner to sit at a restaurant, or if boats are known to be populated by sight-seers as opposed to fishermen, the diners and fishermen will ignore the actions knowing that no information is carried by observed behavior. If incentives and actions are not aligned to produce reliable information, the use and value of public news suffer. Motives and motivation are important parts of the picture.

---

[4] The founders were very clear in their thinking about the matter. "Freedom of speech is a principal pillar of a free government: When this support is taken away, the constitution of a free society is dissolved," wrote Benjamin Franklin in *The Pennsylvania Gazette*. John Adams felt that "The liberty of the press is essential to the security of the state" and Thomas Jefferson held similar opinions holding that "Our liberty depends on the freedom of the press, and that cannot be limited without being lost."

[5]Hayek, F. A. (1945). The use of knowledge in society. The American economic review, 35(4), 519-530.; Hayek, F. A. (1948). Individualism and economic order. University of Chicago Press.

[6] The phenomenon was first demonstrated in market experiments by Plott and Sunder (1982, 1988).

A key feature of the model is a principle of “information revealing choices/behavior”. If an observer knows the motivations of a decision maker, then the observer can use the decisions made to make an inference about what the decision maker knew at the time of the decision. That is, if a decision maker’s incentives are known, then their choice of actions can give insights about the information on which the choice was made – a type of “invertability”. The fact that people can make this deduction, even though not always successfully, suggests the processes through which fake news finds its way to making an impact. If the bias of a source is known, an observer can compensate for the bias and extract the hidden information as if the bias did not exist. On the other hand, if the incentives are not known, then a foundation for collecting and using the information does not exist. The news carries no information aside from the fact that it cannot be trusted.

B. Behavioral Principles

A simple experiment is used to demonstrate four fundamental points. First and most basic are the facts that (i) trusted news is a key feature of information aggregation and (ii) the aggregated information creates economic value. The improved information leads to better decision-making that becomes directly translated into improved income, which can be measured in the experiment. (iii) Successful aggregation depends on the confidence placed by decision makers on the reliability of an information source. The confidence is derived from the decision makers’ understanding of the alignment of the source’s incentives and reports derived from actions taken by the source. Do the reports from the source accurately reflect the information held by the source? Do the actions taken by the source, possibly adjusted for known biases, accurately reflect what the source knows? (iv) Finally, if the confidence in the information source or the aggregation process is damaged then the value created by the information is lost. The confidence reflects a building process, a type of social investment in a “public good” resulting from time and experience.

Together the experiments illustrate how fake news can destabilize and erode the foundation of the information building process and cause a loss of the potential economic value. The experiments demonstrate how fake news can be economically damaging through an erosion of a major resource –knowledge. Such erosion could be damaging in other ways as well, e.g. socially.

C. The Experiment and Overview

Three experimental conditions are studied. The first condition is a simple case in which news sources can be trusted, i.e. no fake news. The second condition is a condition in which both the sources of inaccurate news and associated biases are known to all, i.e. known fake news. The third condition is one in which all know about the possibility of fake news but it cannot be objectively identified as such. The reader of the news cannot reliably differentiate unreliable news from reliable news, i.e. unknown fake news.

Naturally, the experiment is a very simple case where the principles can be observed. The three conditions studied are all based on the same set experimental setting.[7] Two urns have three balls each. One of the urns has two red balls and one white. Call this the red urn. The other urn has two white balls and one red. Call this the white urn. One of these urns is chosen at random at the beginning of a period (50:50). The subject does not know which urn was chosen but is shown a random draw of one ball from the chosen urn. The subject gives a (public) report regarding the predicted chosen urn, and the earnings of the subject depend on the report. Different conditions involve different reporting incentives that will be called "normal" or "reverse" as will be described next. All subjects know their own reporting incentives. All other subjects observe all previous reports, and the ball is replaced after each subject draws.

In the no fake news condition, subjects have incentives to report the actual urn used to make the draw, called "normal incentives". Of course, they do not know definitively but they have incentives to make a correct report and the incentives are common knowledge. In the no fake news condition, the subject earns a monetary reward (+$1.50 in the experiment) if the report is correct and loses money (-$0.50) if the report is incorrect. Since the urns are chosen with equal probability, the best report based on a single, isolated draw is the color of the revealed ball. That is, subjects with normal incentives report red urn if the ball is red and report white urn if the ball is white.

In the known fake news condition some subjects have normal incentives to report the actual urn used, and three randomly chosen subjects have incentives to report the urn NOT used. These three have "reverse" incentives and all subjects know which reports were made by subjects with reverse incentives. Given their information, those with reverse incentives can make their own determination about the actual urn and use that determination to form a report. Those with reverse incentives earn a monetary reward (+$1.50 in the experiment) if the report is the urn not used and lose money (-$0.50) if the report is the actual urn. Again, since the choice of the urn is 50/50, in the absence of additional information the best report is the color of the revealed ball for those with normal incentives (draw x and report x) but the opposite color for those with reverse incentives (draw x and report y).

In the unknown fake news condition subjects know that some subjects might have reverse incentives but they do not know the number or location of subjects with reverse incentives. All subjects know their own incentives but not the incentives of others. Other features of the experiment are exactly the same as the other two conditions.

[7] The experiment was first introduced by Anderson and Holt (1997). Information and efficiency measures were developed by Hung and Plott (2001) who also replicated the Anderson and Holt results. The experiments are based on theoretical models developed by Banerjee (1992) and by Bikhchandani, Hirshleifer, and Welch (1992). Willinger and Ziegelmeyer (1998) and Ziegelmeyer, Koessler, Bracht and Winter (2010) study cases in which participants receive different qualities of information and show that subjects with more accurate private signals correct inaccurate information aggregation.

In all conditions, all other subjects observe the report of the first subject but not the color of the ball that was shown to the first subject, and they do not know if the first subject reported correctly or not. A second subject is chosen at random. The experimenter uses the same urn as before, draws one ball at random and reveals it to the second subject who makes a guess about the urn and makes a public report on the urn. All other subjects observe the report. The ball is replaced in the urn. A third subject is chosen at random and shown a ball drawn at random from the same urn as the previous two subjects. The third subject reports on the urn, and all other subjects observe the report. The process continues until all subjects have made a report about the actual urn. Again, all subjects know that unless the incentives are reverse a subject earns $1.50 if the subject reported the actual urn but loses $0.50 if the subject reported the other urn. If the subject has reverse incentives the incentives are the opposite and in that sense the subject has an incentive to report "fake news".

Now, notice the information the third subject in the sequence has available as a result of the "news" of the previous two subject reports. If subject three believes the previous two subjects are rewarded by accuracy in the sense of reporting the correct urn and that the two other subjects want to make as much money as possible, then the reports carry information about the color of the drawn balls revealed to them. Thus, at the time of decision in the no fake news condition the third subject has information about three draws, the two previous subjects and his/her own draw. The optimal choice is dictated by the proportion of colors in the sample. If two or three are of the same color, then the best choice of urn is that color.

The value and productivity of the "news" can be theoretically computed and compared with the observed. If an individual has no private information and no news, then a natural model of choice would be random with 50% being correct and 50% being incorrect. The expected payoff would be $0.50 per period. If each individual has a private source of information, the private draw, then by using it as the basis of a decision they will choose the correct urn with 67% probability and the incorrect urn 33% of the time with an expected value of $0.83 (purely theoretical). If the calculation is based on the actual draws that were used in the experiment as opposed to the probabilities, then on average the value is $.79.

If an individual has access to news such as the reports of others, the information can be added to the information provided by the individual's private draw. In the case of no fake news if all individuals use the data from reports (the "news") then "herds" or "cascades" of decisions will be observed in which all decisions are eventually the same, which would lead to all making the correct choice approximately 75% of the times and incorrect 25% producing a per person expected value of, $1.01.[8] It is important to realize that such conformity is not accidental and

[8] The perfect cascading result was calculated by considering if each subject reports based on their draw unless the majority public information plus the private draw contradicts the private draw, in which case the subject would report based on the public information. All subjects consider the incentives behind each report. Over 15 rounds, the 10 subjects earn $151 total using this method, so the per person expected value per round is $1.01, approximately 75% correct and 25% incorrect.

is, indeed, beneficial as this value is higher than $0.79 that would be obtained if all subjects disregarded the public information.

D. Experimental Design and Procedures

Four experimental sessions were conducted on four different days. Each session used ten subjects for a total of forty subjects. Each session consisted of unpaid practice periods ranging from 5 to 10 and three experimental conditions. Subjects made fifteen reports under each of the conditions. The major features for all experiments are in Table 1.

Table 1 Experimental design: dates conditions, periods,				
experiment	20180207	20180228	20180305	20180411
Number of subjects	10	10	10	10
Average Total Earnings per person ($)	36.3	40.1	35.5	36.3
Experiment length	1.5 hours	1.5 hours	1.5 hours	1.5 hours
Subjects EEPS lab	Caltech students	Caltech students	Caltech students	Caltech students
Condition	periods	periods	periods	periods
Unpaid Practice fake news	1-5	1-5	1-5	1-5
No Fake News Paid Practice (unused) Data used	 * 6-20	 6-10 11-25	 6-10 11-25	 6-10 11-25
Fake known	21-35	26-40	26-40	26-40
Unknown fake	36-50	41-55	41-55	41-55

*For experiments 20180228, 20180305, and 20180411, the No Fake News condition has 5 extra practice rounds (rounds 6-10), and the data from these rounds were not used in the data analysis (they were treated as the other practice rounds).

Subjects were Caltech students recruited using the recruiting system of the Caltech Laboratory for Experimental Economics and Political Science (EEPS) and reported to the Caltech EEPS laboratory. Upon arriving at the laboratory, subjects were randomly seated at a station with screening partitions and a computer and were instructed to not talk or communicate.

All experiments were conducted in the same manner. When participants walked into the room, they were given colored PowerPoint instructions (see Appendix), a table to fill out during the experiment (about information such as their incentive type and per-round payoffs), and a writing utensil. Each participant was guided to a seat with a computer, without view of any other computers or individuals; no communication of any kind was allowed except for questions to supervisors. After all 10 people read instructions, the instructions were summarized (specifically, incentive types and how to use the program), and initial questions were answered individually. Then, several simple examples were shown on the board in the style of the program to be used. Examples included: [If I have normal incentives and if I think that others are reporting their draws accurately then (I) If the first 2 people chose red, and I drew red, then I should choose red; (II) If the first person chose white, the second chose red, and I drew red, then I should choose red; (III) If the first person chose red, and I drew white second, then given a slight confidence in my choice over someone else, I should choose white; (IV) If the first 2 people chose red, but the first two people have reverse incentives, and I drew white, then I should choose white. (V) If I have reverse incentives, and the first 2 people (normal incentives) chose white, and I drew white, then I should choose red.

After these examples, individual questions were answered. All programs were initiated for the practice rounds, during which people could ask final questions. After the practice rounds, no further questions were answered. People were reminded if parameters changed at the start of new periods. At the end of the experiment, participants calculated their total earnings, excluding the practice rounds, and were presented cash accordingly.

E. Measurements

Subjects' earnings and thus efficiency of the news system depended on the incentive structure and the decisions subjects made. The subject had incentives to make a correct report about either which urn was used to make the draw or the opposite, the urn that was NOT used to make the draw. The incentives differed across experimental conditions.
Normal incentive: If the subject reports the "correct urn", the urn from which the ball was drawn, the subject earns $1.50 and loses $0.50 if the report is not the correct urn.
Reverse incentive: If the subject reports the "incorrect urn", the urn from which the ball was not drawn, the subject earns $1.50 and loses $0.50 if the subject reports the correct urn (the urn from which the ball was drawn).

E.1. Efficiency Measurements
Typically, efficiency reflects the wealth produced by a process. In these types of experiments, it is the money earned by participants relative to the maximum that could have been earned. When earning depends on information the measurement must be adjusted to the information possibilities.

Complete information standard: The complete information standard reports efficiency relative to the hypothetical case in which there is complete public sharing of all draws before any choices are made. From the actual draws used in the experiment, it was calculated that if all subjects are informed of all draws before making a choice, the expected value for a condition (15 rounds) is $16.5 per person.[9] This means that in 12 out of 15 rounds, aggregating all the 10 private signals would lead to the choice of correct state of the world. For reference, there is 79% probability that six or more signals would be from the correct urn. Thus, 100% efficiency

[9] From the actual draws used in the experiment, 1 of the 15 rounds had a tie between the two colors. To account for this in the hypothetical complete information standard, half of the subjects gained money and half lost money.

according to this measure is based on all available information even though potentially impossible to use due to the timing or incentives. [10]

F. RESULTS

Four classes of results are reported and a test of individual's understanding of the instructions was performed.[11] The first result demonstrates that the principles work by measuring the wealth created by a news delivery system to an identical economic environment in which no news system exists. The second result is a demonstration that a news system that carries fake news can function well if the sources of fake news and the nature of the fake news are known in the sense that the motivation of those reporting the news is known. An understanding of the incentives of those reporting the news works as part of a correct interpretation. The third result demonstrates that fake news has a negative impact on the value of an information reporting process. If fake news sources cannot be distinguished from reliable news sources, then the benefits of news are lost. The system tends to revert to the case where no news is available. The final result focuses on the process by illustrating that the value of news is a product of an investment in learning that builds with experience, and in the presence of fake news the investment value becomes eroded.

Result 1 is the comparison between reliable news and no news made possible by Experiment 1. The appropriate measures from a technical computation when no news exists are in Table 2 and are presented with measures produced by Experiment 1.

Table 2 Comparison of private information condition and the no fake news condition

condition	Total earnings	Earnings per person per round	Percentage correct	Efficiency
Private Information Only	$476	$0.79	64.7%	77%
No fake news	$498	$0.83	66.5%	80%

[10] All efficiencies were calculated by analyzing the actual draws people were given. Other efficiency measures are possible. A measure based on "completely rational" behavior of others assumes that all other individuals use statistics properly and all assume that all others do as well. This means that in the absence of tied reports all individuals after the first 3 choose according to the majority signal of the first 3 people. If the first two choices disagree with the third person's draw, then the third person also chooses according to the first 2 people. If the first two choices disagree then the third person follows own signal. The fourth person can be placed in the same position as the third and in this case the analysis is repeated. This measure is sensitive to the first three draws and can exhibit efficiencies above 100%.

[11] The nature of the incentives dictates that the first two choosers in a sequence should report the color of their own draw or the opposite color if their incentives were reverse. All sessions produces a sample of 430 choices of the first two choosers of which 424 choices were consistent with their private draw and their types. Of the 6 inconsistent choices 2 were in in practice periods. Thus, the subjects understood the instructions.

Result 1. Reliable news creates additional value compared to the case when only private information exists. This occurs through a reporting and information aggregation process. If news from a public source is reliable (the absence of fake news) information becomes aggregated through a reporting process and the resulting knowledge becomes an additional, productive source of value creation.

Support. Table 2 demonstrates that in comparison with the identical system that has only private news sources, no public news, the existence of a reliable news source (no fake news) generates benefits in all dimensions of comparison. Total income goes up by 4.62%, the percent of correct (income earning) decisions goes up by 1.8% and the system efficiency goes up by 3%.

The second result compares behavior when everyone knows that no sources of fake news exist and when sources of fake news do exist and everyone knows both the sources of fake news and the biases.

Table 3 Comparison of reliable (no fake) news condition and known fake news sources condition

condition	Total earnings $675 maximum	Earnings per person	Percentage correct	Efficiency
No fake news	$498	$0.83	66.5%	80%
Known fake news sources	$529	$0.88	69%	85%

Result 2. If news sources have known biases, individuals adjust for the biases and information aggregation tends to operate as if there were no bias. Biased news can be processed and the unbiased news content extracted if the biases are fully recognized and understood.

Support. The support for the Result 2 is found in Table 3. The existence of fake news has no impact if people know the sources that are fake. If seeming irrationality or motives are understood, observers just modify the information according to the reliability of the information source. Here the comparisons are substantially along the lines of #1.
i. With known fake news numbers of correct decisions as comparable to no fake news. Compare 69% (fake news) vs. 66.5% (no fake news)
ii. Earnings are comparable between no fake news and known fake news- no statistical difference between earnings per person of no fake and known fake. Compare – not different with 83% confidence, z=0.9374
iii. information use in decisions is comparable – comparison of the Bayesian use of other's decisions as opposed to private information suggests similarities between no fake news and known fake news environments. Compared to the theoretically perfect information aggregation (cascade), no fake news/known fake news achieve 82.5% and 87.4% efficiencies

Of course, issues exist regarding the nature of the bias. Bias can exist in many forms including use of the language, illustrations, opinions advanced as reported facts, exposure frequency, location in a new source, etc. The question posed here is if a bias has an impact when it is obviously a bias. Is there a natural tendency to translate the message; remove the obvious bias and extract the accurate content as is possible in the simple experiment? Result 2 demonstrates that bias removal is clearly within a natural scope of individual human capabilities.

The third result studies the case in which biases are known to exist but the sources of fake news are unknown. The public information carries no obvious method of determining reports that represent fake news from reports that are accurate. The consequences are reductions in earnings and efficiency.

Table 4 Comparison of reliable (no fake) news condition and fake news (fake news sources unknown) condition

condition	Total earnings $X maximum	Earnings per person	Percentage correct	Efficiency
No fake news	$498	$0.83	66.5%	80%
Unknown fake news sources	$456	$0.76	63%	74%

Result 3. The effect of fake news is to substantially diminish or to terminate information aggregation and remove the benefits of the use of available information. Basically, the impact of fake news is to destroy the benefits of the news system. The system performance returns to the base condition in which a public news process does not exist referenced in Table 1. Individual decisions do not benefit from the information held by others.

Support. The support for Result 3 is found in Table 4. When compared to a system of reliable news (no fake news) the performance becomes degraded in all dimensions. Total earnings and average earnings fall. The percentage of correct decisions falls as does system efficiency. Additionally, compared to the baseline case of no public news, the earnings and the percentage of correct choices are lower.

The next result, Result 4, addresses the use of news sources and demonstrates that the impact of fake news works through a shift from decisions based on the aggregation of information found in the news to the less informative, private sources of information.[12] The result illustrates that in the absence of fake news the public news is used for decisions. However, in the presence of unidentifiable fake news the public news sources are abandoned in favor of news produced by private sources.

Bayes Law is used as a model of how decision data are incorporated into individual decisions (Grether, 1980, 1992). The following definitions are needed.

[12] Goeree et al (2007) demonstrate the agents tend to overweigh their private signals.

Let A be defined as the event that urn A is the actual urn, and let B be defined as the event that urn B is the actual urn. Let $x_{it} = (a_{it}, d_{it})$ be the information of individual i at position t, such that a_{it} is defined as the information (the number of A and B choices made by those ahead) that individual i has observed from individuals at positions previous to t, and d_{it} is defined as the private draw of individual i at position t. While a_{it} and d_{it} are correlated with each other, they are conditionally independent given a particular state of the world (A or B). Using Bayes Law, we obtain the following.

$$\frac{P(A|x_{it})}{P(B|x_{it})} = \frac{P(x_{it}|A)P(A)}{P(x_{it}|B)P(B)} = \frac{P(a_{it}|A)P(d_{it}|A)P(A)}{P(a_{it}|B)P(d_{it}|A)P(B)}$$

Taking logs and rearranging:

(1) $$Y_{it} \equiv \ln\left[\frac{P(A|x_{it})}{P(B|x_{it})}\right] = \alpha + \beta \ln\left[\frac{P(a_{it}|A)}{P(a_{it}|B)}\right] + \gamma \ln\left[\frac{P(d_{it}|A)}{P(d_{it}|B)}\right] + u_{it}$$

Where Y_{it} is the belief about the state of the world given x_{it}. Note that $ln\left[\frac{P(A)}{P(B)}\right] = 0$ is canceled out since P(A) = P(B) = ½ from the initial priors. From the data, we can apply equation (1) and find α, β, and γ.

We want to determine if β and γ coefficients differ under different trial conditions (specifically comparing the earnings of news conditions "no reverse" with the "known reverse" and "unknown reverse"). The key variable, d_{it}, the private information, is measured as +1 or -1 depending on the signal the subjects receive. Similarly, a_{it}, the public information in reports available to subject i at position t is measured by the difference in observed actions measured as (1 or -1) depending on the report.

Define s_{it} as the private signals each individual i receives at position t (1 or -1) and $S_{it} = (s_{i1}, s_{i2}, \cdots s_{i(t-1)})$ as the signals all individual from position 1 to t-1 received. Then $\frac{P(S_{it}|A)}{P(S_{it}|B)} = \frac{P(\sum_{k=1}^{t-1} s_k|A)}{P(\sum_{k=1}^{t-1} s_k|B)}$. In other words, sufficient statistic for calculating the posterior odds of the state of the world is simply the difference between the number of signals received and the order of the signals do not matter[13].

With the additional assumption that the subjects believe that previous choices were made in accordance with private signals, we can use the difference in the number of publicly observed choices as a proxy for a_{it}. Given such measurements, the variables are bounded: $\ln\left[\frac{P(d_{it}|A)}{P(d_{it}|B)}\right]$ would either be ±0.693 depending on the private signal and $\ln\left[\frac{P(a_{it}|A)}{P(a_{it}|B)}\right]$ would range from -6.931 to 6.931. So we can perform a linear regression after appropriate change in variables.

Result 4. (i) Individuals always place more decision weight on the private sources of information than on the public sources of information. (ii) The relative weight on public information is

[13] Grether and El-Gamal (1995, 1999) give clear evidence that in impact of order of the information arrival can be impacted by the availability of "heuristics".

reduced if the condition is changed from either no fake news or known fake news to the condition of unknown fake news.

Support. Table 5 contains the result of the regression using Bayes' law as expressed in (1) as a model.

Table 5: Bayes' linear regression

Variable	Coefficient	s.e.	Z-test	p-value
Logit 1: No reverse				
Intercept	-0.0897	0.0257	-3.49	0.0005
Public Information	0.1773	0.0091	19.48	<2e-16
Private Signal	0.7159	0.0369	19.41	<2e-16
Ratio	4.0380			
Logit 2: Known reverse				
Intercept	-0.0216	0.0261	-0.83	0.4070
Public Information	0.1808	0.0092	19.62	<2e-16
Private Signal	0.6822	0.0382	17.86	<2e-16
Ratio	3.7727			
Logit 3: Unknown reverse				
Intercept	-0.0210	0.0193	-1.09	0.2770
Public Information	0.1229	0.0119	10.35	<2e-16
Private Signal	1.1914	0.0282	42.33	<2e-16
Ratio	9.6981			

Overall, we observe that the regression coefficients are similar for 'no reverse' and 'known reverse' condition meaning that under those two conditions individuals give public and private information about the same weight when making their own decision. Specifically, β_1 and β_2 are statistically equivalent to each other and γ_1 and γ_2 are statistically equivalent to each other. However, in the 'unknown reverse' condition, the data demonstrates both a decreased influence of public information and an increased influence of private signal. $\beta_1 \approx \beta_2 > \beta_3$ and $\gamma_1 \approx \gamma_2 < \gamma_3$. This suggests that in the 'unknown reverse' condition, subjects tend to ignore the public information and stick with their private signal more.[14]

To check for robustness, of the Bayesian model we use a "Direct Measurement" method for comparing the relative importance of public information and private signal. Basically, we count the number of reports consistent and inconsistent with the information in an individual's own draw. Let u be the number of reports that a subject i observes which supports d, the private draw of i. Conversely, v is the number of reports observed by i that are the opposite of i's

[14] The intercept (α) is statistically equivalent to zero for both 'known reverse' and 'unknown reverse' condition. This is to be expected if the subjects do not have a bias for Red or White once public information and private signals are taken into consideration. The fact that the intercept is non-zero for the 'no reverse' condition is quite odd but the effect size is not that large since the bias corresponds to about 4% swing in probability.

private draw. Then we let the public information a(i,u,v) = u-v be a measure of the difference between reports in accordance with the private signal and against it.

In the Direct Measurement model the public information is based on the difference of observed choices between red(+1) and white(-1). The order is deemed to be irrelevant. The numbers are applied in a model to produce the probability of the color ball the individual will choose given the numbers the individual observes at the time of choice.

The measurement is taken from a logistic regression on the differences where we choose β and γ that best fits

(2) $$y = \begin{cases} 1 & if\ \beta \cdot a + \gamma \cdot d + \varepsilon > 0 \\ 0 & else \end{cases}$$

where a is the public information available to the individual at the time of decision, d is the private signal and y is the choice variable. ε is the error term distributed by standard logistic distribution and $y' = \beta \cdot a + \gamma \cdot d + \varepsilon$ is the latent variable which acts as an intermediate step towards measuring subjects' choices.

Note that d={-1,1} and a={x|-9 ≤ x ≤ 9, x is an integer}. The restriction follows from the number of subjects and the choices sequences used in the experiment. The individual in the kth choice slot has k-1 reports to view and there were 10 subjects in the experiment.

In the (2) the key variable of interest is γ/β ,which is the relative weight of the private signal relative to one unit of public information.

Table 6: Direct Measurement logit regression model

Variable	Coefficient	s.e.	Z-test	p-value
Logit 1: No reverse				
Public Information	1.1605	0.1097	10.58	<2e-16
Private Signal	3.313	0.2948	11.24	<2e-16
Ratio	2.855			
Logit 2: Known reverse				
Public Information	1.0852	0.1035	10.49	<2e-16
Private Signal	2.9449	0.2613	11.27	<2e-16
Ratio	2.714			
Logit 3: Unknown reverse				
Public Information	1.1717	0.1475	7.946	1.92e-15
Private Signal	5.0822	0.422	12.043	<2e-16
Ratio	4.337			

The completed regression produces a binary prediction of the subjects' choice based on the coefficients. In other words, the estimated parameters for γ and β are given as $\hat{\beta}$ and $\hat{\gamma}$. Then for all sets of private signal and publicly observed information (a,d), if $\hat{\beta} \cdot a + \hat{\gamma} \cdot d > 0$, subjects are predicted to choose red (+1), if $\hat{\beta} \cdot a + \hat{\gamma} \cdot d < 0$, then subjects are predicted to choose white(-1).

Figure 1 demonstrates that the direct measurement model's prediction agrees with the actual subjects' behavior. The horizontal axis displays the net public information. i.e., {the number of observed public reports that agree with the subject's private signal} minus {the number of observed public reports that disagree with the subject's private signal}. The stacked bar graph shows for each level of public information, the number of the subjects' choices that were correctly predicted by the model in blue and that were incorrectly predicted in grey.

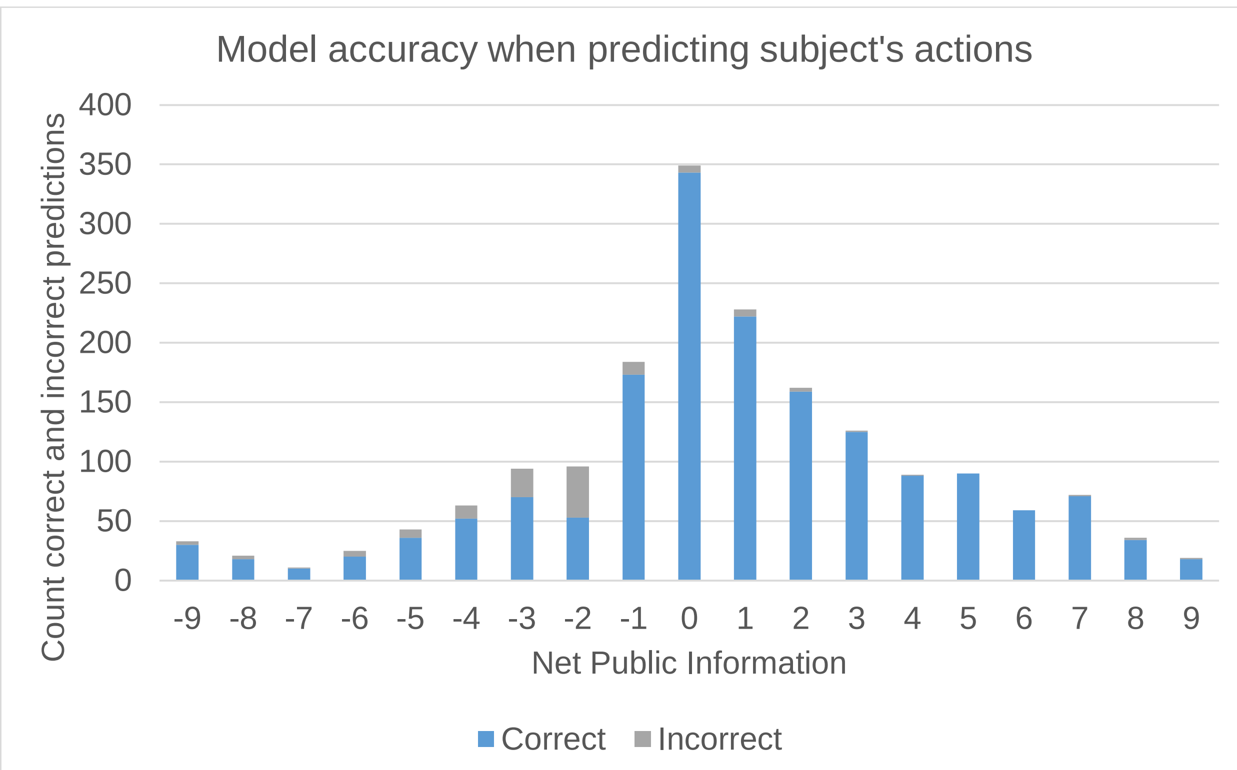

Figure 1: Count of correctly and incorrectly predicted actions according to model

The model correctly predicts 92.8% of subjects' choices pooled across all settings. Most of the "incorrect" prediction occurs when the subjects observe a net of -2 or -3 public information that goes against their private signal.

Furthermore, the model can be used to augment the binary prediction by adding measurements of how the choice probability changes according to the different levels of public information encountered. The figures show the actual proportion of subjects choosing according to their private signals (blue solid line) in contrast to the choice probability model derived from the logit regression (pink dashed line). The horizontal axis displays the net public information as before.

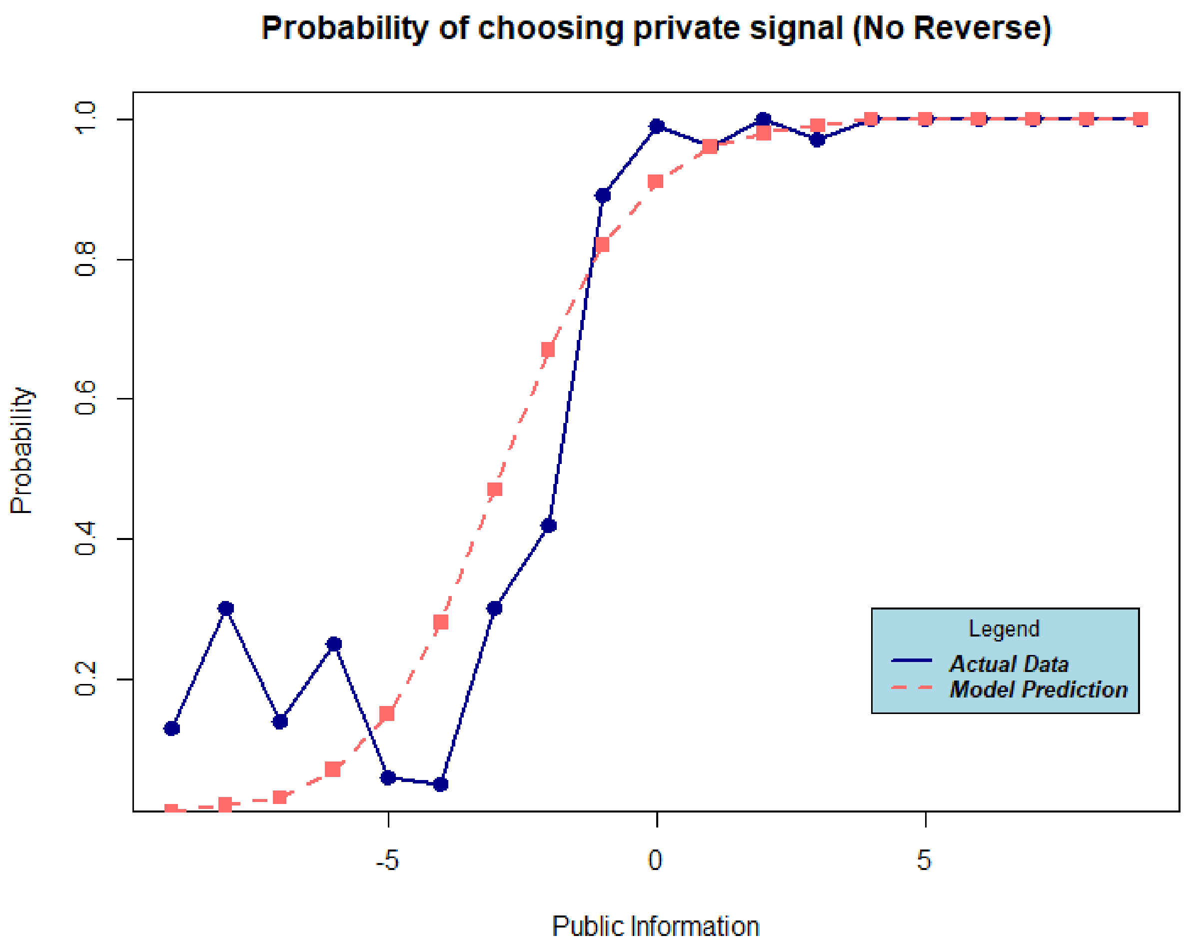

Figure 2: Probability of choosing in accordance with private signal (no reverse) – Data and Logit model prediction

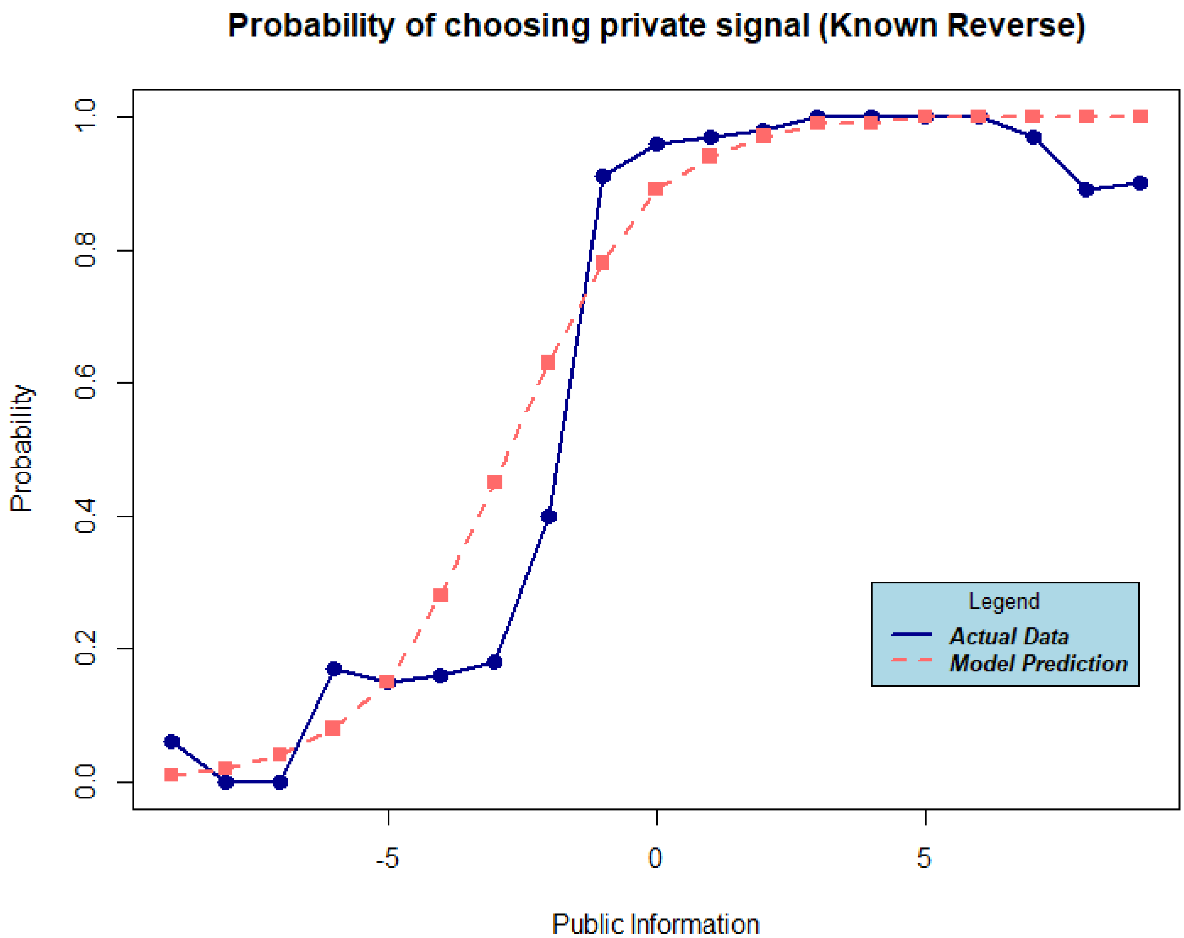

Figure 3: Probability of choosing in accordance with private signal (known reverse) – Data and Logit model prediction

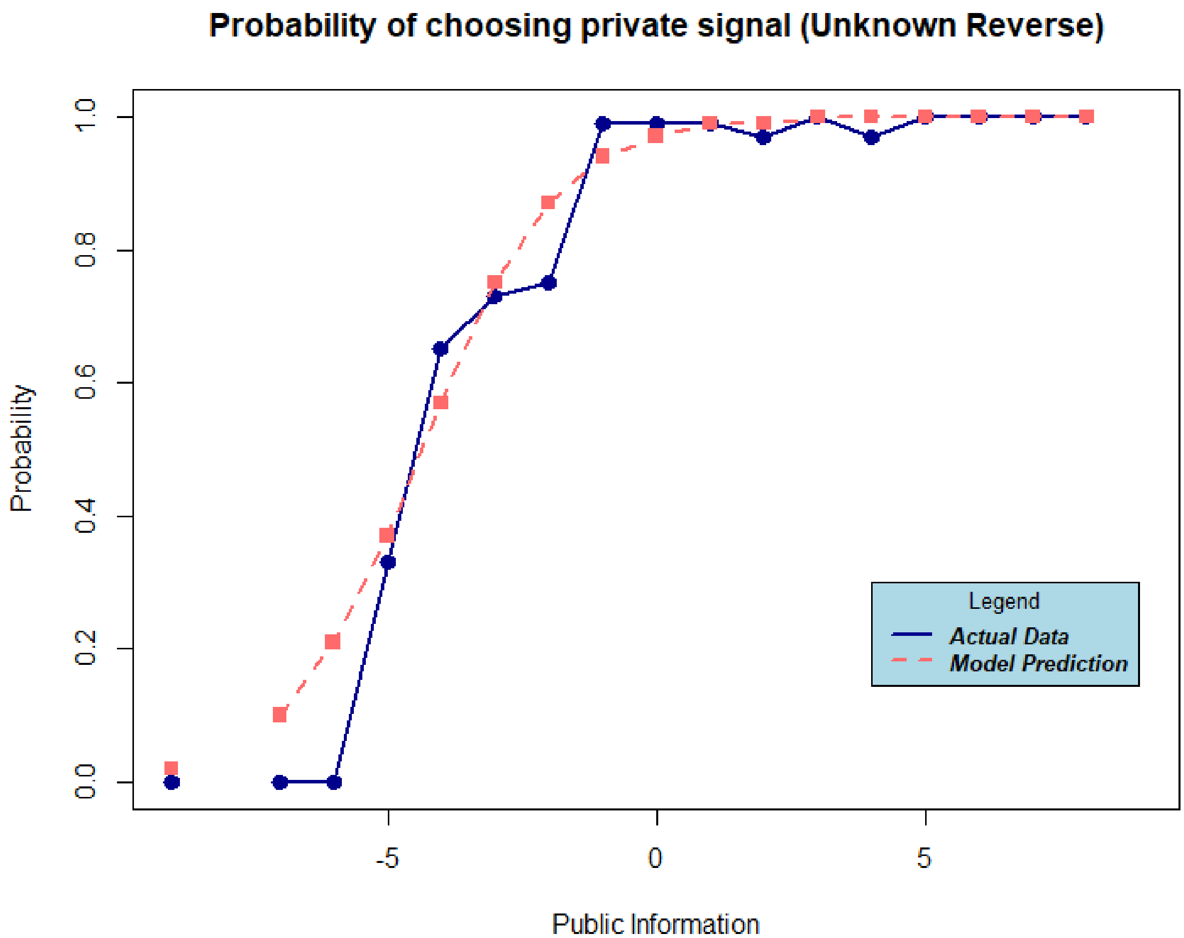

Figure 4: Probability of choosing in accordance with private signal (Unknown reverse) – Data and Logit model prediction

We can observe that the regression model fits the data well across all settings. This provides additional validity to our usage of 'direct measurement' model.

Specifically, the ratio of γ/β tells us the answer to the question "How strong should the public information be in order for the majority of subjects to go against their private signals?". In the 'no reverse' and 'known reverse' case, the answer is somewhere between 2 and 3 while for the 'unknown reverse' case, a signal strength of 4~5 is required.

While the result itself is intuitive in that people would put more emphasis on their private signal in the face of uncertainty, the difference between the settings is more pronounced than the numbers simply suggest.

In the 'no reverse' and 'known reverse' case, 21% of the subjects chose to go against their private signal while in the 'unknown reverse' case, that number drops to 7%, 1/3 of the value.

The reason for this big change is that the 'unknown reverse' setting has two effects. The first one, as mentioned, is that subjects put more emphasis on their private signals. The second one is that the introduction of reverse incentives itself inherently reduces the probability of a strong public information, even by chance. In the 'no reverse' case even if everybody were naïve and relied on their own private signals, the fourth subject would observe -3 public signal with 12% probability. In the 'unknown' reverse case, this is reduced to 4%. The lack of cascade in the 'unknown reverse' case is a crucial feature.

G. Summary of Conclusions

Current news providers appear to be engaged in a war with each devoting resources to illustrate that the other side is guilty of producing fake news. Each provides evidence that the other side does not produce reliable news and that the other side uses subtle tools to avoid being detected. The hostilities can extend beyond challenges of credibility and involve forms of punishment such as social sanctions, attempts to create economic harm[15] or even physical confrontations. Conflicts of this sort have properties that game theorists call a "war of attrition" that is wasteful as are all wars, and ends only when a "winner" emerges and collects any resources that remain. In the case of fake news, the damage might be to the news system itself as the public loses confidence in the reliability of information delivered through the news. Basically, the experiment suggests that the fears and negative prognostications about the consequences of fack news could be justified.

The experiment reported here draws on research found in the information aggregation literature from economics and finance. Information dispersed across many observers becomes aggregated in the form of a signal that can be valuable in worlds of decision making under uncertainty. The experiments demonstrate that fake news can undermine the foundation of the process when source reliability is unknown.

Three experiments were conducted. The first involved no fake news and in this case the experiments demonstrated that subjects learned to rely on public news sources because such reliance improved their income. The second experiment introduced reverse incentives for some news sources that gave the incentive to produce false reports. In this experiment subjects knew which reporters had reverse incentives and as a consequence subjects adjusted and translated the report so the proper news was extracted. The report from a source known to be biased toward x was properly translated to y. The result was that fake news had no effect.

[15] Hung and Plott (2001) studied the experiment reported in this paper but imposed social pressures for conformity. People were punished with a monetary fine for reports that diverged from the reports of others. The experiment demonstrates that incentives for conformity have the capacity to remove almost all benefits of news. The very first report had an impact on all subsequent reports as subsequent reports conformed and avoided punishment by matching the content of the first report. Thus the first report had the impact of blocking the information content of additional reports. Even the first reports became modified as the reporters attempt to conform to the anticipated substance of subsequent reports.

The third experiment removed information about source reliability. Incentives of all individuals were unknown but the possible existence of reverse incentives was known. Information use shifted away from public sources to private sources. The advantage of information aggregation was lost. As a result, this third experiment removed the advantages of public news sources and information aggregation. The profits made by participants decreased and were comparable to if they had only private information.

The lesson here is that some of the advantages of public news are derived from known principles of behavior found operating in many places in the economy. In part these principles depend on a trusted connection between reporters' incentives and the information they are capable of reporting. Fake news destroys that relationship and consequently carries potentially important implications beyond the fact that some people lie.

H. References

Anderson, L.R., Holt, C.A. (1997). Informational Cascades in the Laboratory. *American Economic Review,* 87, 847-862.

Banerjee, A.V. (1992). A Simple Model of Herd Behavior. *Quarterly Journal of Economics,* 107, 797-817.

Bikhchandani, S., Hirshleifer, D., Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades*. Journal of Political Economy,* 100, 992-1026.

Goeree, J. K., Palfrey, T. R., Rogers, B. W., & McKelvey, R. D. (2007). Self-correcting information cascades. *The Review of Economic Studies*, *74*(3), 733-762.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics*, *95*(3), 537-557.

Grether, D. M. (1992). Testing Bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior & Organization*, *17*(1), 31-57.

Grether,D.M. and Mahmoud A. El-Gamal (1995), "Are people Bayesian? Uncovering behavioral strategies", *Journal of the American Statistical Association*90 (1995): 1127-1145

Grether,D.M. and Mahmoud A. El-Gamal (1999), "Changing Decision Rules: Uncovering Behavioral Strategies Using Estimation/Classification {EC), M. J. Machina et al. (eds.), *Beliefs, Interactions and Preferences in Decision Making* , Springer Science+Business Media, New York.

Hayek, F. A. (1945). The use of knowledge in society. The American economic review, 35(4), 519-530.

Hayek, F. A. (1948). Individualism and economic order. University of Chicago Press.

Hung, Angela and Charles Plott (2001). "Information Cascades: Replication and an Extension to Majority Rule and Conformity Rewarding Institutions." *American Economic Review,* 91, 1508-1520.

Kübler, D., & Weizsäcker, G. (2004). Limited depth of reasoning and failure of cascade formation in the laboratory. *The Review of Economic Studies*, *71*(2), 425-441.

Plott, Charles R. and Sunder, Shyam (1982). "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models," *Journal of Political Economy*, 90, 663-698.

Plott, Charles R. and Sunder, Shyam (1988). "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, 56(5), 1085-1118.

Welch, Ivo (1992). "Sequential Sales, Learning, and Cascades," *Journal of Finance*, 47(2), 695-732.

Willinger, M., & Ziegelmeyer, A. (1998). "Are more informed agents able to shatter information cascades in the lab?". In *The Economics of Networks* (291-305). Springer, Berlin, Heidelberg.

Ziegelmeyer, A., Koessler, F., Bracht, J., & Winter, E. (2010). Fragility of information cascades: an experimental study using elicited beliefs. Experimental Economics, 13(2), 121-145.

SECTION I APPENDICES
INSTRUCTIONS
DATA APPENDIX

0=normal
INSTRUCTIONS POWERPOINT:

20180305

- At the start of each round, an urn is randomly chosen. The same actual urn is used throughout a period.

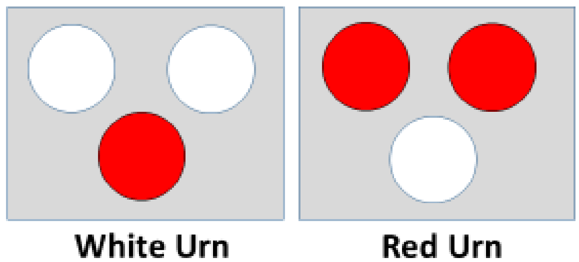

- A ball is randomly selected from the actual urn. The color of the drawn ball is revealed to the next person in the sequence and no one else. The ball is returned to the urn.
- The individual chooses an urn. The individual's choice of urn is revealed to all others but the ball drawn is not revealed to others.
- The process is repeated until all have made a choice.
- The actual urn is revealed and all are paid for the round based on their choice and the actual urn from which all draws were made

20180305

Incentives:

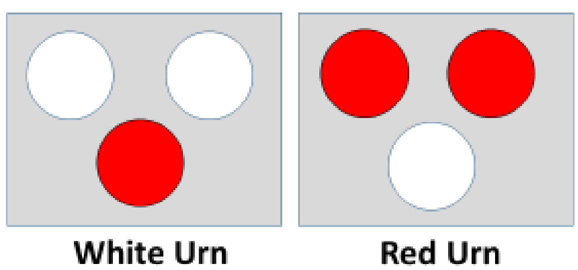

- **Normal Incentive** – choose the actual urn: you get $1.50
  – choose the other urn: you lose $0.50
- **Reverse Incentive** – choose the actual urn: you lose $0.50
  – choose the other urn: you get $1.50

20180305

Choosing Screen Example:

Your Draw	White (or Red)
Your Order	5
Your Type	Reverse (or Normal)
Reverse Positions	3 (or Unknown) 8 10
First Choice	Red
Second Choice	White
Third Choice	White
Fourth Choice	Red
Fifth Choice	
Sixth Choice	
Seventh Choice	
Eight Choice	
Ninth Choice	
Tenth Choice	
Your Choice	O Red O White

- Others' choices in order

Choose

20180305

At the end of each round, actual urn & your earnings are displayed

Actual Urn	White or Red
Your Draw Was	White or Red
Your Type	Normal or Reverse
Your Choice Was	White or Red
Your income this period	1.50 or -0.50

Example Earnings Table:

Period	Earnings
1	-$0.50
2	$1.50