Special Issue: Experiments and Policy

Public Finance Review
1-23

DOI: 10.1177/1091142120960488
journals.sagepub.com/home/pfr
SAGE

Fake News, Information Herds, Cascades, and Economic Knowledge

Lazarina Butkovich[1], Nina Butkovich[1], Saba Devdariani[1], Charles R. Plott[1] and Han Seo[1]

Abstract

This article focuses on principles of information aggregation in the presence of false, public reports (fake news). The analysis explores news as having a public goods feature characterized by models of information and economic efficiency. The analysis is not tied to any particular theory about how or why unreliable news emerges. The reports could be purposeful deception, inten-tions to mislead or profit motivated responses to decision biases of readers. A well-known and widely studied "cascade" experiment is used to illustrate principles that provide links to standard economic models. News is modeled as an aggregation of a simple, fixed chain of decentralized observations and reports about an underlying, unknown state of nature. The personal value of an individual's decision depends on both the decision and the underlying state of nature. The information about the state used in the decision can reflect private observations or the "news" about the decisions of others. The experiments

[1] California Institute of Technology, Pasadena, CA, USA

Corresponding Author:
Charles R. Plott, California Institute of Technology, 1200 E. California Blvd., MC 228-77, Pasadena, CA 91125, USA.
Email: cplott@hss.caltech.edu

demonstrate that aggregated information is dependent on accumulated trust in news sources and has value as a special form of public goods.

Keywords
information aggregation, cascades, information source, information efficiency

While fake news is not a newly emergent phenomenon, it is currently the focus of discussions and concerns about the implications. This article addresses the subject of fake news from a narrow perspective of standard economic models and the possible efficiency-enhancing, benefits-producing properties of news. The perspective abstracts from strategic intentions and possible perception bias of news producers or various possible news preferences of consumers. In contrast to the broad discussions found in popular literature, the analysis is constructed from a very simple, often-studied experiment and a model that connects the phenomenon to principles found in traditional economics. News is modeled as observations about an underlying state (of nature) that cannot be observed directly. The theory rests on the possibility that the value of news reflects the fact that multiple, independent observations can support a guess about a common "state of the world" based on information available. As such, the value of news can exceed the value of the private observations of a single decision maker and thus has economic properties recognized as a "public good." A well-known laboratory economics experiment, typically used to illustrate "herding" or "cascading" decisions, is conducted to illustrate the model and explore its basic principles. The model is used to demonstrate the social, economic value of a well-functioning news system and how fake news can destroy that value.

This article is organized as follows. Information Aggregation and Efficiency from an Economics Perspective section is a discussion of the concepts and background theory. Experiments section is a discussion of the experiment, the experimental design, and the questions posed for testing. Results section is the statement of findings. Summary of Conclusions section is a summary of our contribution.

Information Aggregation and Efficiency from an Economics Perspective

The topic of fake news is very broad. By contrast, this article explores very simple and special case issues as an attempt to uncover some of

the basic principles that are at work. Of course, news is a deep part of the social fabric, and policies that shape its role, such as freedom of the press, are widely recognized as a cornerstone of a free society.[1] Complex and important issues abound but are not addressed here. Omitted important topics include the vehicles that carry news, demand for news, rumors, propaganda, misinformation as well as those who benefit from fake news and the tools from advertising and strategies they might use. The question is how fake news might emerge and the consequences of fake news in the absence of such promotions.

This article is focused on very simple cases in which the detailed features of general, widely used models can be specified, studied, and connected to related phenomena in economics. The theory reflects the work of Hayek (1945), "The Use of Knowledge in Society," and Hayek (1948) regarding how information is transmitted through market institutions and is applied widely to 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 and are the key features of the experiment. Within that general model, 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. In the Hayek model, the information becomes transferred and aggregated through publicly observed actions such as trades and trading prices in which good information and incentives to take actions are closely related. In a statistical sense, the information contained in scattered observations can become aggregated, pooled, and processed.[2] The resulting knowledge is a resource, a type of "public good," that can be used multiple times and can produce economic value to the user by preventing costly mistakes. Often the accumulated information is regarded as "common sense" and is easily detected in everyday events and life-shaping decisions. For example, people normally visit a restaurant because they like the food, and thus, the number of customers is used as a source of information about the quality of the restaurant. Or, the location of fishermen is used as 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.

The existence and growth of fake news phenomena have motivated warnings. For example, Michael Bloomberg suggests the possibility of direct damages to wealth through the effect of fake news on the stock market (https://www.bloomberg.com/view/articles/2017-10-23/why-

fake-news-is-so-harmful-to-investors). 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 suggestion is that the information contained in the news and the use of the news would deteriorate.

The traditional economic theory of information aggregation suggests two classes of principles that shape the dynamics that influence the growth or decline of the value of news. One principle, a principle of "information-revealing choices/behavior," points to the incentives of those whose actions can carry information. It is a type of "information-revealing invertibility." When incentives and information-revealing actions are aligned, the actions of others 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. Of course, experience contributes to the reliability of information and can lead to recognition of the absence of information-revealing invertibility when the information should be modified or ignored. 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 sightseers as opposed to fishermen, the diners and fishermen will ignore the actions suspecting 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.

A second principle applies when multiple observations or observers exist, and information can be accumulated with the possibility of multiple (independent) observations becoming aggregated in a statistical sense. It is a type of "principle of information aggregation." Each step rests on the hypothesis that those who make earlier decisions have sufficient confidence in the decisions of those who come before them. Such links are part of the aggregation process that pools otherwise separate observations. The structure of the underlying network and the pattern of observations available to individuals in the network have theoretical implications for news performance and could differ according to the news accumulation and delivery process. Certainly, any aggregation phenomena would be sensitive to the social structures that channel information exposure, the sequence of observations, what is observed, the independence of observations, and other institutional features. News from "social news" or gossip could differ from the

Table 1. Experimental Design: Dates, Conditions, and Periods.

Experiment Date	20180207	20180228	20180305	20180411
Number of subjects	10	10	10	10
Average total earnings per person (US$)	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	1–5	1–5	1–5	1–5
No fake news				
Paid practice (unused)	†	6–10	6–10	6–10
Data used	6–20	11–25	11–25	11–25
Fake known	21–35	26–40	26–40	26–40
Unknown fake	36–50	41–55	41–55	41–55

Note: EEPS = Experimental Economics and Political Science.
†For the first group (20180207), the first five paid practice rounds for "no fake news" were not conducted. For all of the other groups, these five practice rounds were conducted; therefore, to make comparisons, we haven't used data from these rounds from any group. So, we call the first actual round for group 20180207 the sixth round (to make it comparable with all the other groups).

"mainstream press." The aggregation of information and decisions might appear as a growth, cascading, or herding property that becomes stronger as the information builds even independent of accuracy. In the absence of an understanding of the role of the news, the social behavior could be interpreted as mindless mimicking, imitation, or conforming behavior as suggested by the terms "cascades" and "herds." However, the theory suggests the possibility that the news system is an information-building process in which the value of the public information provided increases with additional observation. In that sense, the basic theory of news is tangentially connected to the basic theory of public goods provision processes in which the news constitutes a "public good" shared in common by all as a source of uncertainty reduction.

Experiments

The experiments listed in table 1 are structured from the implications of the two principles and the background literature. The experiments take place in

a very simple, "one-dimensional" communication in which the predictions of the principles can be unambiguously specified. Each participant makes a decision under uncertainty. At the time of the decision, they have information from their private source and information implicit in the decisions of preceding participants, each of which gained relevant, but limited information. The reader will recognize the setting as the traditional cascading or herding experimental environment.[4]

The three conditions studied are all based on the same experimental setting. Under all settings, subjects are arranged in a randomized order. Two computerized 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. The experiment proceeds as a sequence of periods. One of these urns is chosen at random at the beginning of each period (50:50). The subject does not know which urn was chosen but is privately shown a random (independent) draw of one ball from the chosen urn. The subject gives a (public) report regarding the predicted chosen urn. The earnings of the subject depend on the subject's report and the actual urn. Following the draw and choice of the first subject, the ball is replaced, and a second draw is randomly drawn from the same urn and revealed to the second subject who makes a public decision. The process continues until all subjects have made a choice. At the end, all subjects have seen all previous decisions. After all the decisions are made, the actual urn from which all draws were made is publicly announced. After subjects have checked their earnings for the period, a new urn is chosen, subjects are randomly ordered, and a subsequent period is initiated.

Three experimental conditions are studied as summarized in table 1 and analyzed as suggested by the principles outlined above. The first condition (no fake news) is a simple case in which news sources can be trusted in the sense that the decision maker's incentives are clearly, similarly aligned. All subjects have an incentive to choose the urn representing the actual state (the urn from which the ball was drawn). In the no fake news condition, the subject earns a monetary reward (+US$1.50 in the experiment) if the report is correct and loses money (-US$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, isolated subjects with normal incentives report red urn if the ball is red and report white urn if the ball is white.

The second is a condition (known fake news) in which pattern (number but not the exact sources) of inaccurate news is known to all.[5] 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. 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 (+US$1.50 in the experiment) if the report is the urn not used and lose money (-US$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).

Notice that those who might create "fake news" produce incorrect information because they have different values conditional on the state of nature. They are not motivated to mislead or to fool others. Their purpose is not to cause others to make mistakes. Fraud and intentional deception are not a part of the motivations in the experiment. While such dimensions could be added, the addition would introduce theoretical complexities that, at this time, would obscure the basic principles that are the focus here, which are to explore how the consequences of fake news can emerge when such possibly dysfunctional motivations do not exist.

The third condition is an *unknown fake news* condition. All subjects know their own incentives but not the incentives of others. Subjects know that some subjects might have reverse incentives, but they do not know the number or location of subjects with reverse incentives. Thus, in the third condition, all subjects know about the possibility of fake news, but the pattern is not known, and sources are unknown. The reader of the news cannot reliably differentiate unreliable news from reliable news and knows nothing about the general reliability, that is, unknown fake news. Other features of the experiment are exactly the same as the other two conditions.

Experimental Procedures

Four experimental sessions were conducted on four different days. Each session used ten subjects for a total of forty subjects drawn from the Caltech student population. Each session started with five to ten practice periods and consisted of three experimental conditions. Subjects made fifteen reports under each of the conditions. The major features for all experiments are in table 1.

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 a view of any other computers or individuals; no communication of any kind was allowed except for questions to experimenters. After all the ten subjects read the 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. In general, the examples were technical but the following was included to emphasize the random nature of the drawing process and focus on the importance of their own draw, the draws of others, and the incentive structure: (If I have normal incentives) (I) if the first two 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, then given slight confidence in my choice over someone else, I should choose white; (IV) if the first two people chose red, but the first two people have reverse incentives, and I drew white, then I should choose white; and (V) if I have reverse incentives, and the first two people (normal incentives) chose white, and I drew white, then I should choose red. The use of the word "should" might suggest experimenter "demand effects," but the data across changing experimental conditions suggest that subjects were using their own judgments and that the instruction was successful eliminating confusion, misunderstandings, or misconceptions about the task, the role of randomness, and incentives.

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.

Measurements: Earnings and Efficiency

Subjects' earnings and thus the efficiency of the news system depended on the incentive structure and the decisions subjects made. The subjects 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 US$1.50 and loses US$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 US$1.50 and loses US$0.50 if the subject reports the correct urn (the urn from which the ball was drawn).

Efficiency measurements reflect 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 earnings depend on the 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 (fifteen rounds) is US$14.7 per person.[6] This means that in twelve of fifteen rounds, aggregating all the ten private signals would lead to the choice of the correct state of the world. For reference, there is 79 percent probability that six or more signals would be from the correct urn. Thus, 100 percent efficiency according to this measure is based on all available information even though potentially impossible to use due to the timing or incentives.[7]

Results

Four classes of results are reported. In all cases, the movements of all measurements are in the direction suggested by theory, but the small number of periods, given conditions, limits statistical significance.

The first result demonstrates that the principles work by comparing measurements of the wealth created by a news delivery system to an identical economic environment in which no news system exists and all agents

rely on private information. The second result is a demonstration that a reliable news source in the sense of no fake news creates social value through efficiency increases. The third result demonstrates that a news source with known biases toward fake news operates with equal efficiency as a news source with no fake news. The fourth and fifth results demonstrate that the introduction of unknown fake news removes the advantages of the news source, and the fifth result demonstrates that fake news can even do additional damage. Result 6 demonstrates that the results are consistent with a Bayes's law model of individual decisions with the weight shifting between public and private information when making decisions using new sources with varying reliability.

Result 1. When individuals have only their private sources of information/news (observations), they follow Bayes's law. They choose the state with the highest probability conditional on their information.

This result can be tested based on the reporting behavior of the subjects who make the first move. In our experiment, the first person, in each round, only observes a private draw and makes a decision solely based on this information. Since the prior belief about the urn being either color is "0.5," a person who follows the Bayes's rule should report the color of their private draw if they have normal preferences (and the opposite color of their draw if they have reverse preferences). The data show that 214 of 215[8] "first movers" actually report according to their private draw and preferences. Statistically, the probability that the choice of color matches (mismatches) the draw, if the person has normal (reverse) preferences when no other information exists, is essentially 1.

Result 2. Reliable news supports the creation of additional value compared to the case in which only private information exists, but the difference is not statistically significant. Relative to environments with no public news, the existence of public information increases earnings, efficiency, and the proportion of correct decisions relative to an environment with no public news, but the increase is not statistically significant at conventional levels.

Table 2 contains the data. From result 1, we conclude that in the absence of any public information, people use Bayes's rule and make decisions according to their private signal. Therefore, based on actual draws from the experiment, we can simulate the environment with no public

Table 2. Comparison of Private Information Condition and the No Fake News Condition.

Condition	Total Earnings (US$)	Earnings per Person per Round (US$)	Percentage Correct	Efficiency (Percentage)
Private information only	476	.79	64.7	80.9
No fake news	498	.83	66.5	84.7
Known pattern of fake news sources	528	.88	69	89.8
Unknown fake news patterns	456	.76	63	77.6

information, where each subject, regardless of their order in the experiment, only observes their private draw and makes the decision accordingly. Comparing this case to the treatment of reliable public information (no fake news) shows that people benefit from the existence of reliable public information in every aspect: more earnings, higher efficiency, and more accuracy. Having reliable public information comparing to only private information increases earnings from US$476 to US$498; the percentage of correct reports increases from 64.66 percent to 66.5 percent, while efficiency increases from 80.95 percent to 84.69 percent. We compared earnings between the two treatments using Wilcoxon signed-rank[9] two-sided test, and the difference in earnings between these two treatments is statistically weak with p value $= .322$, meaning that even though public information increases earnings, the difference is not statistically significant.

Result 3. If news sources have known biases, then individuals adjust for the biases. When biases exist and have known patterns, information aggregation tends to operate as if there were no biases by increasing efficiency, income, and percentage correct as compared to the condition when only private information exists.

Analysis from table 2 shows that, if people know about fake news, they can effectively adjust for it while making a final decision. To see the support of this property, we compare two treatments: no fake news and known fake news. For the two samples, we compare individual earnings.

We use the two-sided Wilcoxon signed-rank test[10] on the difference between the means of these two samples.[11]

Hypothesis 0: On average, people earn the same in the case of no fake news and known patterns of fake news.

Hypothesis 1: On average, people do not earn the same in the case of no fake news and known patterns of fake news.

We cannot effectively reject Hypothesis 0 against the alternative hypothesis, Hypothesis 1 (p value = .3982). The earnings tend to be the same under conditions of known patterns of fake news, and individuals adjust for the bias.

Overall, in the case of known patterns of fake news, the earnings increase from US$498 to US$528, the percentage of correct reports increases from 66.5% to 69%, and efficiency increases from 84.7% to 89.8%.

Result 4. The existence of possible fake news, or the existence of news with an unknown frequency of bias, removes the benefits of available information.

Previous results demonstrate that people benefit from the existence of reliable private information (result 2), and even if the information has known bias, people effectively adjust for it (result 3). Now, we examine the case when the information becomes unreliable. In the last treatment of our experiment (unknown fake news), people are told about the possibility of fake news, but they do not know the number of the identity of those with reversed preferences. Essentially, this treatment should make the public information (reports) a completely unreliable source, and the issue is the impact of the unreliability as compared to anticipated unreliability. We use the one-sided Wilcoxon signed-rank test.

Hypothesis 0: On average, people earn the same with unknown fake news as compared to known patterns of fake news.

Hypothesis 1: On average, people earn less with unknown fake news than known patterns of fake news.

We can reject Hypothesis 0 against the alternative hypothesis, Hypothesis 1 (p value = .01574). Therefore, the reliability (or known unreliability) of the information is important.

Unfortunately, the number of observations is not enough to make statistically significant inferences about the difference in earning separately for each order in the experiment. However, with naive analysis, there is still some indication that the order in the experiment has an effect. For comparison between unknown and known fake news, as expected, the difference in earnings between treatments is not statistically significant for individuals who move first in the sequence. Logically, since the first person only sees his or her private signal, the reliability of the public signal is completely irrelevant for them. In fact, the difference is not statistically significant up until the fifth member in the sequence. Moreover, the difference in earnings between the two treatments is again insignificant for the last mover in the sequence. There are two main forces at play here; the reliability of the information is more significant for people later in the sequence since they get more public information to rely on. On the other hand, the possibility of cascading also becomes more significant later in the sequence, and therefore, the public information might be less important for the last movers. Unfortunately, the lack of an effective measure of cascading as well as the limited number of observations makes it impossible to control.

Result 5. In the case of fake news with unknown reliability, people do even worse than in the case of only private information. System efficiency, incomes, and percentage correct were compared to the case of only private information.

The previous result 4 demonstrated a negative effect of unknown bias in public information on system efficiency and earnings. The next result illustrates that the reason for the negative impact is due to a tendency for people to simply ignore news that they think is unreliable and make decisions on the basis of private sources of information. To test for this property, we compare their performance under the condition of unknown fake news with the simulated case of only private information using the same sequence of private signals as in unknown fake news.

Hypothesis 0: On average, people earn the same with unknown fake news as compared to only private news.

Hypothesis 1: On average, people do not earn the same with unknown fake news as compared to only private news.

We can effectively reject Hypothesis 0 against our alternative hypothesis, Hypothesis 1 (p value = .005601), meaning that people do not just ignore unreliable public information. Moreover, we can do a one-sided test against:

Hypothesis 1′: On average, people earn less with unknown fake news than with only private news.

We can effectively reject Hypothesis 0 against our alternative hypothesis, Hypothesis 1′ (p value = .0028), meaning people do even worse with unreliable public information than without any public information. In the case of unknown fake news compared to only private information, earnings decrease from US$476 to US$456, the percentage of correct reports decreases from 64.66% to 63%, while efficiency decreases from 80.95% to 77.56%.

A deeper view of individual decisions can be obtained through Bayes's law and a model developed by Grether (1980, 1992). The following definitions are needed.

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's law, we obtain the following:

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

Taking logs and rearranging:

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

where Y_{it} is the belief about the state of the world, given x_{it} (from our data, we use private reports from the subjects adjusted with their preferences for the dependent variable). 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.

Table 3. Bayes's Logit Regression.

Variable	Coefficient	*SE*	*t* Value	*p* Value
Logit 1: No reverse				
Intercept	0.089671	.025670	3.493	0.000512
Private Signal	0.715789	.036873	19.412	<2e−16
Public Information	0.177257	.009097	19.484	<2e−16
Ratio	4.038			
Logit 2: Known reverse				
Intercept	0.01601	.03155	0.507	0.612
Private Signal	0.73137	.04582	15.963	<2e−16
Public Information	0.20128	.02060	9.769	<2e−16
Ratio	3.634			
Logit 3: Unknown reverse				
Intercept	0.02104	.01933	−1.089	0.2770
Private Signal	1.19115	.02814	42.326	<2e−16
Public Information	0.12284	.01187	10.351	<2e−16
Ratio	9.697			

We want to determine whether β and γ coefficients differ under different trial conditions (specifically comparing the earnings of news conditions "no reverse" with the "known patterns of 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 depending on the report.

Define s_{it} as the private signals each individual i receives at position t (1 or -1) and $S_{it} = \left(s_{i1}, s_{i2}, \ldots, s_{i(t-1)}\right)$ as the signals all individual from position 1 to $t-1$ received. Then, $\frac{P(S_{it}|A)}{P(S_{it}|B)} = \frac{P\left(\sum_{k=1}^{t-1} s_k|A\right)}{P\left(\sum_{k=1}^{t-1} s_k|B\right)}$. In other words, a 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 in the model.

With a simplifying 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[12] depending on the private signal and $\ln\left[\frac{P(a_{it}|A)}{P(a_{it}|B)}\right]$ would

range from −6.238 to 6.238. So, we can perform a linear regression after an appropriate change in variables.

Result 6. (i) Given the information contained in the public news, individuals do not weigh the public information enough relative to their private information. (ii) The relative weight on public information is reduced if the condition is changed from either no fake news or known patterns of fake news to the condition of unknown fake news.

Table 3 contains the result[13] of the regression using Bayes's law as expressed in equation (1) as a model. In the treatment of no fake news, individuals place more decision weight on the private sources of information (.72) than on the public sources of information (.18). Overall, the regression coefficients are similar for no reverse and known reverse condition, meaning that under those two conditions, the weights individuals give to the public and private information do not change between treatments. Specifically, β_1 and β_2 are statistically equivalent to each other and γ_1 and γ_2 are statistically equivalent to each other (Z-test, $\alpha = .05$). However, in the unknown reverse condition, the data demonstrate 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.

Summary of Conclusions

The experiments reported here draw on research found in the information aggregation literature from economics, political science, and finance. Information dispersed across many observers becomes aggregated in the form of a signal that can be valuable in the worlds of decision-making under uncertainty. The experiments demonstrate that fake news can undermine the foundation of the information aggregation 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 and that such reliance improved their income (although improvement was not statistically significant). The second experiment introduced reverse incentives that gave some subjects (the news sources) the incentive to produce false reports. In this experiment, subjects knew

very few 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 suspected to be biased toward x was properly translated to y. The result was that fake news had no effect.

The third experiment removed information about relative source reliability. The possible existence of reverse incentives and thus fake news was known, but its concentration was not known. In response to the uncertainty, the information use shifted away from public sources to private sources. The public news was used less. As a result, this third experiment removed the benefits of public news sources and information aggregation. The profits made by participants decreased and were comparable to the condition in which they had only private information. News was published, but information aggregation was lost.

According to this model, some of the benefits of public news and possibly some damages of unreliable 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 can emerge even in the absence of intentions of some to mislead others. Consequently, the emergence of fake news carries implications beyond the fact that some people might be naive in forming beliefs or that people might lie.

While this article is narrowly focused on the basic principles of standard models and makes no attempt to capture or model the many variables active in the news industry or possible biases of news consumers or producers, some speculation about the lesson is unavoidable. Broad attention to fake news is visible in current network news where 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 "tweetstorms" and news reported in related social media seem to amplify the messages and accusations with verbal, punishing attacks on reports and the reporters of facts with which they disagree. In the presence of such forms of punishing accusations, the traditional model suggests that fake news and the discussion it stimulates can damage the news system itself,[14] a tragedy of the commons, as the public loses confidence in the reliability of information delivered through the news and the social value provided by confidence in the news is lost.

APPENDICIES I

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.

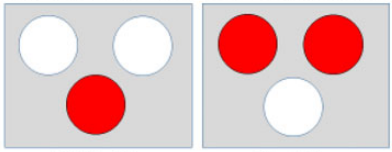

White Urn **Red Urn**

- 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:

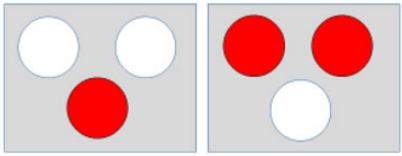

White Urn **Red Urn**

- **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

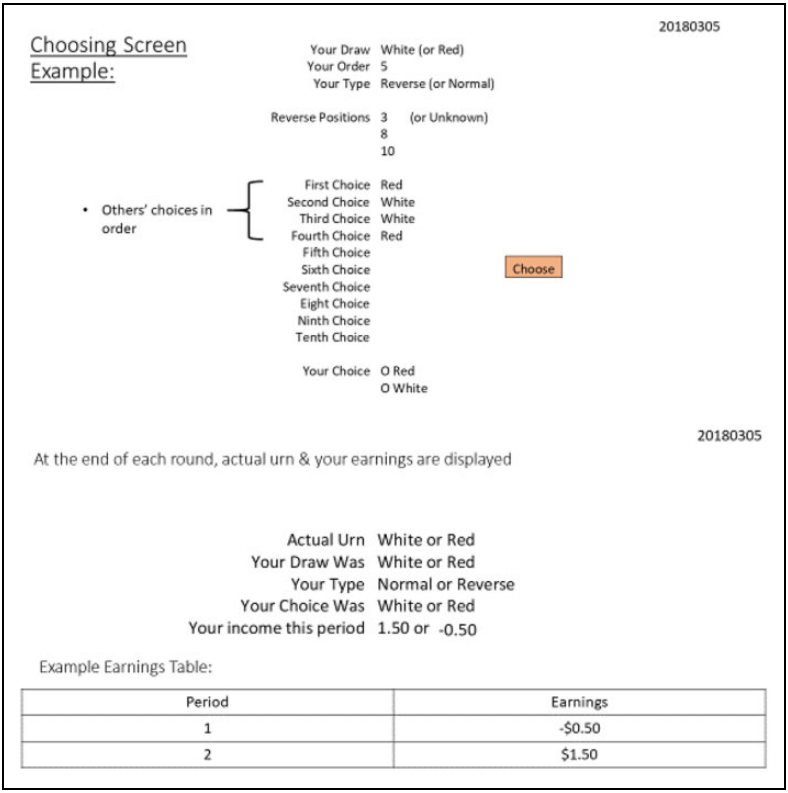

Authors' Note

All experimental programs were designed by Travis Maron.

Acknowledgments

The financial support of the John Templeton Foundation and the Rising Tide Foundation is gratefully acknowledged. The help of B. Atsavapranee is also gratefully acknowledged.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Funding was supplied by The John Templeton Foundation and The Rising Tide Foundation.

ORCID iD


Charles R. Plott https://orcid.org/0000-0001-8363-3628

Notes

1. 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.”
2. The phenomenon was first demonstrated in the market experiments by Plott and Sunder (1982, 1988).
3. The comments are her assessment of academic research as contained in her 2018 Arthur Miller Freedom to Write Lecture at the 14th Annual PEN World Voices Festival, April 14, 2018.
4. The experiment was first introduced by Anderson and Holt (1997). Information and efficiency measures were first recognized and developed by Hung and Plott (2001) who also replicated the Anderson and Holt’s results. The experiments are based on theoretical models developed by Banerjee (1992), Welch (1992) and by Bikhchandani, Hirshleifer, and Welch (1992). Willinger and Ziegelmeyer (1998), Kübler and Weizsäcker, (2004), Goeree et al (2007) and Ziegelmeyer et al. (2010) study cases in which participants receive different qualities of information and show that subjects with more accurate private signals correct inaccurate information aggregation. The delicate features of the aggregation and inference process were studied more recently by modeling individual beliefs about the beliefs of others when making choices. See an important step in understanding provided by Angrisani et al. (2018). Goeree et al (2007) and Kubler and Weizsacker (2004).
5. Early pilot experiments studied a treatment in which the identity of those with reverse incentives was known to all. Subjects just interpreted the choices for what they were and translated a report of one urn into a choice of the other if the choosing subject had reverse incentives. The pattern was so pronounced we chose not to use resources to gather data from the treatment.

6. From the actual draws used in the experiment, one of the fifteen 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.
7. All efficiencies were calculated by analyzing the actual draws people were given. Other efficiency measures are possible including efficiency relative to "completely rational" behavior of others. The measure is based on the assumption 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 three choose according to the majority signal of the first three people. If the first two choices disagree with the third person's draw, then the third person also chooses according to the first two 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. Because this measure is sensitive to the first three draws, it can exhibit variability of performance and efficiencies above 100%. Each such instance requires additional explanation, and thus, the measure has deficiencies as an explanatory and comparison tool.
8. For this result, we included data from practice round as well. The one person who did not follow their private signal was not from the practice round. Practice round data were excluded in all other tests.
9. We use signed-rank test throughout the analysis since the data we compare are paired with the order number in experiment. Each round is treated as independent observation (people are shuffled for order in experiment throughout the different rounds).
10. The Wilcoxon signed-rank test is often described as the nonparametric version of the paired *t* test. These two will essentially give the same results for large sample size.
11. The data are matched according to the order of the individual in the experiment.
12. If $d_{it} = 1$ (red ball), then we have $P(d_{it}|A) = \frac{2}{3}$ since red urn has two red and one white balls and $P(d_{it}|B) = \frac{1}{3}$ since white urn has one red and two white balls. Therefore, we have $\ln\left[\frac{P(d_{it}|A)}{P(d_{it}|B)}\right] = \ln\left[\frac{\frac{2}{3}}{\frac{1}{3}}\right] = \ln[2] \approx 0.693$. Similarly, when $d_{it} = -1$ (white ball) then $\ln\left[\frac{P(d_{it}|A)}{P(d_{it}|B)}\right] \approx -0.693$.
13. Each regression has 600 observations; unit of observation is subject decision.
14. The introduction of punishments of those whose reports differ from the reports of others can have a dramatic effect on the flow and content of news. Experiments similar to those reported here result in a substantial decrease in efficiency as those with information either modify their report to match the news reported by others and avoid the punishment or simply do not report. Those whose information and decision follow the first report simply conform by repeating

the content of the first report and ignoring their private information. The information aggregation process stops after the first decision as all subsequent reports tend to conform to the first reports. See Hung and Plott (2001).

References

Anderson, L. R., and C. A. Holt. 1997. "Informational Cascades in the Laboratory." *American Economic Review* 87:847–62.

Angrisani, M., A. Guarino, P. Jehiel, and T. Kitagawa. 2018. "Information Redundancy Neglect versus Overconfidence: A Social Learning Experiment." *American Economic Journal: Microeconomics*.

Banerjee, A. V.. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107:797–817.

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

Goeree, J. K., T. R. Palfrey, B. W. Rogers, and R. D. McKelvey. 2007. "Self-correcting Information Cascades." *The Review of Economic Studies* 74 (3): 733–62.

Grether, D. M.. 1980. "Bayes Rule as a Descriptive Model: The Representativeness Heuristic." *The Quarterly Journal of Economics* 95 (3): 537–57.

Grether, D. M.. 1992. "Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence." *Journal of Economic Behavior & Organization* 17 (1): 31–57.

Hayek, F. A. 1945. "The Use of Knowledge in Society." *The American Economic Review* 35 (4): 519–30.

Hayek, F. A. 1948. *Individualism and Economic Order*. Chicago: 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* 91:1508–20.

Kübler, D., and G. Weizsäcker. 2004. "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory." *The Review of Economic Studies* 71 (2): 425–41.

Plott, Charles R., and Shyam Sunder. 1982. "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-expectations Models." *Journal of Political Economy* 90:663–98.

Plott, Charles R., and Shyam Sunder. 1988. "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets." *Econometrica* 56 (5): 1085–118.

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

Willinger, M., and A. Ziegelmeyer. 1998. "Are More Informed Agents Able to Shatter Information Cascades in the Lab?" In *The Economics of Network*, edited by Patrick Cohendet, Patrick Llerena, Hubert Stahn, and Gisèle Umbhauer, 291–305. Springer, Berlin, Heidelberg.

Ziegelmeyer, A., F. Koessler, J. Bracht, and E. Winter. 2010. "Fragility of Information Cascades: An Experimental Study Using Elicited Beliefs." *Experimental Economics* 13 (2): 121–45.

Author Biographies

Lazarina Butkovich is an undergrad student at the California Institute of Technology.

Nina Butkovich is an undergrad student at the California Institute of Technology.

Saba Devdariani is a graduate student at the California Institute of Technology

Charles R. Plott is the William D. Hacker professor of economics and political science at the California Institute of Technology. He is a member of the National Academy of Sciences. He specializes in experimental economics and political science, with an interest in theory.

Han Seo is a graduate student at the California Institute of Technology. He is specializing in experimental economics, decision theory, and institutional design.