Bias in a Laboratory Simulation of a Signaling Game with Implications for the Influence of the News Media

Tim Groseclose
November 2011

Abstract

Cai and Wang (2005) conducted a laboratory simulation of the Crawford-Sobel "strategic information transmission" game. The researchers were most interested in observing the amount of information that senders in the game transmitted to receivers. Consequently, they did not examine what I call the *no-policy-bias* implication of the game. This implication is that—as the Nash equilibria to the game predict—the final policy that receivers choose, in expectation, should equal the policy that they would have chosen if the senders had sent no signal. That is, the senders should not be able to systematically fool the receivers. Contrary to the Nash equlibria, however, all versions of the experiment produced a policy bias. The results can be explained by a well-established empirical regularity documented by behavioral economists. This is that people tend to under-estimate the degree to which other people are strategic. I apply these results to the question of media effects. The Cai-Wang results—as well as the models of behavioral economists—suggest that real-world journalists should indeed be able to significantly affect the thoughts and behavior of real-world news consumers.

Perhaps the most famous and widely-applied signaling model of all time is Crawford and Sobel's (1982) "strategic information transmission" game. The game involves two players, a Sender and a Receiver. The Sender has information that the Receiver would like to know. The Sender sends a message, which might or might not reveal the information. After seeing the message, the Receiver chooses a policy for both players.

Because the players have different preferences, the Sender has an incentive to exaggerate or lie when sending a message. Consequently, the Receiver has an incentive to discount the message, possibly ignoring it completely.

Many parameterizations of the game possess, what I call, the *no-policy-bias* implication. That is, for instance, although the Sender may want to fool the Receiver into choosing a more liberal policy than he (the Receiver) would choose if he received no signal, in equilibrium she (the Sender) cannot do this.[1]

Cai and Wang (2005) have simulated the Crawford-Sobel game in a laboratory, and indeed the theory behind their parameterization possesses the no-policy-bias implication. Because Cai and Wang were most concerned with other implications of the game (specifically, the amount of information that was actually transferred relative to the amount that the theory predicts), they did not examine the no-policy-bias implication. The authors have graciously given me their data, and in this paper I examine that implication. Despite the theoretical predictions of the game, I find that, in practice, there

[1] An important aspect of the Crawford-Sobel game is that the Sender does not face a penalty for lying. Kartik (2009) and Ederer and Fehr (2009) have examined similar signaling games that *do* contain penalties for lying. Their results suggest that the no-policy-bias implication is robust. That is, when Crawford-Sobel-like games include a penalty for lying, the altered game still implies no policy bias. (Like me, however, Ederer and Fehr find that in practice—at least when it is played in a laboratory setting—their game *does* produce a policy bias. Further, like me, they find that the bias goes in the direction in which the Sender prefers.) Similar to results I present later, Kartik, Ottaviani, and Squintani (2007) show that when the Receiver is not fully rational, then Crawford-Sobel-type games *can* produce a policy bias.

was a significant policy bias. The Senders in the experiment were indeed able to systematically shift policy in the direction that they preferred.[2]

The results are consistent with a well-established empirical regularity within behavioral economics. This is that people tend to under-estimate the degree that other people are strategic.[3] More specific, such research is often consistent with what is sometimes called the *one-difference* principle. This is that the typical person tends to believe that the sophistication of other people is one degree lower than his own. As I show, if people act this way, then the Crawford-Sobel game indeed produces a policy bias. Further, once we adopt a simple parameterization of this regularity, it explains the Cai-Wang results much better than does a Nash-equilibrium model.

I relate these results to the question of "media effects." Namely, I argue that the Crawford-Sobel game is similar to the interaction between a journalist and a news consumer. For instance, if the journalist is, say, more liberal than the news consumer, then she has an incentive to report the news with a liberal bias. Her goal is to persuade the news consumer to vote in a more liberal manner than he would naturally—that is, if he received no news or the news he received was unbiased. However, the theory of the Crawford-Sobel game predicts that such a journalist cannot do this in equilibrium. That is, the theory of the game predicts that media effects will be nil. However, the results

---

[2] Among studies of legislatures (where the Sender is, say, a lobbyist or the median of a committee, and the Receiver is the median of the floor), the no-policy-bias implication is well known. For instance, what I call "policy bias" Krehbiel (1991, 75) calls a "distributional loss." He formally defines it as "the degree to which expected outcomes deviate from [the ideal point of the median floor member]." Gilligan and Krehbiel (1987) note that equilibrium play of their committee-floor signaling game has no distributional loss. (See, for example, page 311 or page 315 of their article.)

[3] Researchers who have examined the regularity include Nagel (1995), Stahl and Wilson (1995), Ho, et. al. (1998), Costa-Gomes, et. al. (2001), Crawford (2003), Costa-Gomes and Crawford (2004), Camerer, et. al. (2004), and Cai and Wang (2006). Cai and Wang (2006) provide an excellent review of such research.

from the Cai-Wang experiment suggest that the opposite should occur in practice—that real-world media effects should be significant.

Elsewhere (Groseclose, 2011), I introduce the concept of the *media λ*. This is a parameter between zero and one, which represents the degree to which the media can influence the thoughts and actions of ordinary citizens. As I show, results from the Cai-Wang experiment imply a *media λ* of approximately .32. Because of the pristine nature of the experiment and the concomitant fact that subjects will be *en guard* for biased messages, the results from the experiment likely understate the true *media λ*. I show that if (i) the Crawford-Sobel model indeed captures the strategic interaction between journalists and news consumers, and (ii) real-world journalists and real-world news consumers are slightly less sophisticated than the subjects of the Cai-Wang experiment—then the model implies a real-world media lambda of approximately .7—the value that I have previously estimated for real-world journalists and real-world consumers.

Results from the Cai-Wang Experiment

In the Cai-Wang experiment, a computer randomly chose a state of nature from the space {1, 3, 5, 7, 9}. This could be interpreted as the ideal policy preference of the Receiver. Only the Sender saw this number.

The Sender's preference was always greater than the Receiver's. Cai and Wang simulated four different versions of the game—where the Sender's preference was .5, 1.2, 2.0, and 4.0 units greater than the Receiver's preference. Both players were informed of this parameter, which Cai and Wang labeled *d*.

Play began with the Sender, who sent a message equal to 1, 3, 5, 7, or 9. After seeing the message, the Receiver chose a policy equal to 1, 2, 3, 4, 5, 6, 7, 8, or 9.

Players were paid according to how close the resulting policy was to their preference. (Specifically, let y be the policy that the Receiver chose, and let x_i be the preference of player i [where i = Sender or Receiver]. Player i's payoff [in points, which were eventually converted to money] was $110\text{-}10|y\text{-}\ x_i\ |^{1.4}$.)[4]

The following table summarizes the results of the experiment.

d = Diff. in Sender's and Receiver's Preference	Average Message	Average Policy[5]	Policy Bias
4.0	7.186	5.592	.592
2.0	5.894	5.282	.282
1.2	6.161	5.377	.377
.5	5.153	5.078	.078

[4] Cai and Wang's main conclusion was that the players "over-communicated" relative to the predictions from the theory. For instance, when d = 4.0, the only Nash equilibrium is a *babbling* one. That is, in equilibrium the Sender's message does not reveal any information about the Receiver's preference, and the Receiver's optimal strategy is to ignore the message. Accordingly, the theory of the game predicts that there should be zero correlation between: (i) the Receiver's preferences and the Sender's messages, (ii) the Sender's messages and the Receiver's policy choices, and (iii) the Receiver's preferences and the Receiver's policy choices. In practice, however, all three correlations were positive. Similar results occurred for the other versions of the experiment (when d was .5, 1.2, or 2.0). That is, for these versions correlations (i), (ii), and (iii) were usually larger than equilibrium behavior predicted.

[5] To calculate the averages in the table, I did not calculate the simple average from all the observations. Rather, I calculate (i) the average policy (or message or bias) associated for all the observations where the Receiver's preference was 1, (ii) the average policy where the Receiver's preference was 3, and so on. Next, I calculate the average of these five averages. This is identical to the simple-average method if the exactly one-fifth of the observations involved the Receiver having a preference of 1, exactly one-fifth involved the Receiver having a preference of 3, and so on. However, if, for instance, fewer than one-fifth of the observations involved a 1 preference for the Receiver, then those observations will be over-weighted in the average-of-averages method that I use. It is important to use the latter method instead of a simple average. To see why, note that when choosing a policy, Receivers will use two pieces of information: (i) the message they receive from the Sender, and (ii) their prior beliefs about their ideal preference. The latter factor influences Receivers to slant their policy choice toward 5, the expected value of their prior beliefs. This means that when the computer has drawn 1 as their true preference, the Receivers will tend to choose a policy that overshoots their true preference—that is, their policy choice will have a positive bias. The opposite will happen when the computer has drawn 9 as their preference. This means that if the computer randomly draws an inordinate number of low numbers, then the policy bias will be artificially high. And the policy bias will be artificially low if the computer draws an inordinate number of high numbers.

For each version of the game, a Nash equilibrium requires the average policy to be 5.0. This did not occur in practice. Thus, the experiment violated the no-policy-bias implication.[6]

Indeed, for each of the four versions of the game, policy was *greater* than 5.0—thus, the policy bias always occurred in the direction that the Sender preferred. In each version, the degree to which the Sender could pull policy—beyond the Nash-equilibrium level—was at least 14% of the amount that she preferred. That is, for instance, consider the case where $d = .5$. Here, the Sender wants the average policy to be 5.5—which is .5 units greater than Nash-equilibrium behavior predicts. Meanwhile, the average actual policy was .078 units greater than Nash equilibrium predicts—which is 15.6% of the degree of bias that the Sender preferred (i.e. 15.6% = .078/.5). In one case—where d = 1.2—the degree of bias was 31.4% of the amount that the Sender preferred.

Bounded-Rationality Strategies for the Crawford-Sobel Game

We now examine how bounded rationality—specifically, behavior consistent with the *one-difference* principle—can cause the Crawford-Sobel game to produce a policy bias. To do this, first, we must define a set of strategies. As Cai and Wang (2006) note, if the subject is a Sender, then the least-sophisticated strategy is for her always to report the truth—that is, to report the state of nature as her message. Call this strategy S_0. More

[6] The Cai-Wang experiment is the only experiment of which I am aware that (i) simulates the Crawford-Sobel game and (ii) implies no policy bias when subjects play Nash-equilibrium strategies. I am aware of only one other case—that by Dickhaut, McCabe, and Mukherji (1995)—where scholars have simulated the Crawford-Sobel game in a laboratory setting. The theory of the latter experiment, however, does not possess the no-policy-bias implication. That is, in the latter experiment—mainly because the Receiver is sometimes not allowed to choose his expected ideal policy—some of the Nash equilibria produce a policy bias. In the experiment, the Receiver's preferences were, with equal probability, 1, 2, 3, or 4. As an example of a Nash equilibrium that possesses a policy bias, suppose that the Sender adopts a babbling strategy. Thus, the expected ideal policy for the Receiver is 2.5. The experiment, however, only allowed the Receiver to choose an integer for an action. Thus, the Receiver is indifferent between choosing 2 or 3. If he randomized with equal probability over these two actions, then this would produce no policy bias. However, any other strategy would produce a policy bias. Even some non-babbling equilibria of the experiment imply a policy bias.

generally, define S_k as follows: Given that the sender observes the state of nature, *s*, she sends $s+dk$ as her message, or if that is not allowable (say, because it is above 9), then she reports the nearest allowable message to $s+dk$. For instance, suppose $d=2.0$, and $k=2$. Then for states of nature 1, 3, 5, 7, 9, she respectively reports 5, 7, 9, 9, 9.

Define R_k as the Receiver's optimal strategy under the assumption that his (Sender) opponent has adopted strategy S_k. For instance, under R_0 the Receiver assumes that the Sender always tells the truth. Thus, his optimal strategy is to choose policy equal to the message she reports.

As another example, continue to assume $d=2.0$, and consider R_2 —the Receiver's optimal strategy when he believes the Sender is playing S_2. Suppose he sees a message of 9. Note that the Sender sends this message when the state of nature is 5, 7, or 9. Thus, a message of 9 makes him believe that the expected state of nature is 7.0. Thus, 7.0 is his optimal policy. If he sees a message of 5 or 7, then his optimal policy is respectively 1 or 3. If he sees a message of 1 or 3, then this is an out-of-equilibrium event. Following Cai and Wang, I adopt the refinement that his policy choice is monotonic in the message. Thus, he chooses a policy of 1 if he sees either of these messages.

For strategies S_k and R_k , note that *k* roughly indicates the level of sophistication of the strategy. That is, for instance, under S_k the Sender exaggerates her message by *k* degrees. And under R_k the Receiver discounts the Sender's message by *k* degrees. The strategies are similar to Camerer, et. al.'s *cognitive hierarchy model*, which, as they note, "consists of iterative rules for players doing *k* steps of thinking (2004, 863)."

Let us now define types of subjects. If an individual is of type 0, then her strategy is R_0 if she is a Receiver and S_0 if she is a Sender. If an individual is of type *k* (where $k >$

0), then she believes that her opponent plays strategy R_{k-1} if he is a Receiver and strategy S_{k-1} if he is a Sender. Note how this definition adopts the *one-difference* principle. Each type believes that her opponent is one degree of sophistication lower than she is.

It can be shown that, for all $k > 0$, if the individual is of type k, then her optimal strategy as a Sender is S_k , and her optimal strategy as a Receiver is R_{k-1} . For instance, suppose $d=2.0$, the Sender is type 2, and the state of nature is 3.0. Note that this means that she prefers policy to be 5.0. Since she is of type 2, she assumes that her opponent's strategy is R_1. Thus, if she reports 7.0, then he will choose 5.0 as policy. Thus, 7.0 is her optimal message. Note that 7.0 = $s+2d$, which is her optimal message under S_2.

Now suppose that this same type-2 individual is a Receiver. Then she believes that her Sender opponent has adopted a strategy of S_1 . Recall that R_1 is defined as the optimal response to an S_1 strategy.

How Bounded Rationality Causes a Policy Bias in the Crawford-Sobel Game

To see how the above behavior can produce a policy bias, suppose as a very simple example, that half the subjects are type 1 and half are type 2. Suppose that d = 2.0, and suppose that the state of nature is 5.0.

The type-1 and type-2 Senders will respectively adopt strategies of S_1 and S_2 , which means their messages will respectively be 7.0 and 9.0.

The type-1 and type-2 Receivers will respectively adopt strategies of R_0 and R_1. The type-1 Receivers will be fully trusting. Thus, those who are paired with a type-1 Sender (respectively paired with a type-2 Sender) will choose a policy equal to 7.0 (respectively, 9.0). The type-2 Receivers who are paired with a type-1 Sender (respectively, a type-2 Sender) will choose a policy equal to 5.0 (respectively, 7.0).

It is easily shown that in this case—where the state of nature is 5.0—on average, the Receiver will choose a policy of 7.0. Note that this produces a policy bias of 2.0 units. It can similarly be shown that when the state of nature is 1, 3, 7, or 9, then on average the Receiver will respectively choose a policy of 3.0, 5.0, 8.0, or 8.0. Note that the average policy—across all five possible states of nature—will be 6.2. The average message will be 7.2, and the policy bias will be 1.2.

I have conducted another set of analysis that more accurately fits the Cai-Wang data. For this, I assume the following: Twenty percent of the subjects are type-3; 20% are type-2; 20% are type-1; 20% are type-0; and the remaining 20% are *completely non-strategic*. That is—say, because they don't understand the instructions of the experiment—they always choose 5.0 as the policy if they are the Receiver, and they always report a message of 5.0 if they are the Sender.

For each of the four versions of the experiment, the following table lists the expected message, policy, and bias that should occur under these assumptions. The table also lists the correlations that we should expect between: (i) the state of nature (s) and the Sender's message (m); (ii) the Sender's message (m) and the policy (p) chosen by the Receiver; and (iii) the state of nature (s) and the policy (p). The purpose is to demonstrate how well the above assumptions predict actual behavior—compared to the predictions of Nash equilibria. Accordingly, for each of the four versions of the game, the table first lists the results from actual play. Below that, in parentheses, it lists the prediction from the above assumptions. And below that, in brackets, it lists the prediction from Nash-equilibrium play.[7]

[7] The games often had multiple Nash equilibria in terms of the message that the Sender sends. This makes predictions for the Sender's message ambiguous. For these cases I list a "?" for the Nash-equilibrium

d	Ave. Message	Ave. Policy	Ave. bias	Correlations between: s-m	m-p	s-p
4.0	7.186	5.592	.592	.39	.54	.21
	(7.160)	(5.768)	(.768)	(.36)	(.65)	(.23)
	[?]	[5.000]	[.000]	[.00]	[.00]	[.00]
2.0	5.894	5.282	.282	.73	.79	.62
	(6.600)	(5.384)	(.384)	(.61)	(.78)	(.48)
	[?]	[5.000]	[.000]	[.50]	[.71]	[.71]
1.2	6.161	5.377	.377	.90	.92	.83
	(6.200)	(5.320)	(.320)	(.75)	(.84)	(.64)
	[?]	[5.000]	[.000]	[.75]	[.87]	[.87]
.5	5.153	5.078	.078	.92	.97	.88
	(5.320)	(5.256)	(.256)	(.86)	(.77)	(.89)
	[5.000]	[5.000]	[.000]	[1.00]	[1.00]	[1.00]

The bounded-rationality model predicts actual behavior better than does the Nash-equilibrium model. For instance, consider the predictions for the average policy bias. With the bounded rationality model, on average, the absolute difference of its prediction from the actual value was .128.[8] Meanwhile, the same figure from the Nash-equilibrium model was .332. That is, the Nash-equilibrium model was more than twice as inaccurate as the bounded-rationality model.

I conducted similar calculations for the correlations that the two models predict. With the bounded rationality model, on average, the absolute difference of its prediction from the actual value was .094. Meanwhile, the same figure from the Nash-equilibrium

---

prediction. For instance, suppose *d=2.0.* Then, if the equilibrium is most-informative, then the Sender must send one message when the state of nature is 1 and a different message when the state of nature is 3, 5, 7, or 9. In the first case, the natural prediction is that the Sender's message will be 1. However, in the second case there is no natural prediction—e.g. the sender could report "9" for each of the four states (3, 5, 7, or 9); or she could report "7" for each of the four states (or report "5" or "3"); or she could randomize between 3, 5, 7, and 9. Each of these possibilities forms a Nash equilibrium.

[8] To derive this figure, I used the numbers in the table. Specifically, I calculated (|.592-.768| + |.282-.384| + |.377-.320| + |.078-.256|) / 4 = .128.

model was .169. Thus, by this measure the Nash-equilibrium model was almost twice as inaccurate as the bounded-rationality model.

Note that, in the bounded-rationality model, I did not fine-tune the parameters. That is, I assumed that each of the five types (completely non-strategic, type 1, type 2, etc.) contained 20% of the subjects. Instead, I could have estimated percentages that would have optimized the accuracy of the model. Moreover, I constrained the model so that each subject acted as if her opponent had a sophistication level exactly 1.0 degrees less than her own. Instead, I could have substituted any number for 1.0.[9] By relaxing either assumption, I could have made the bounded-rationality model even more accurate.[10]

Signaling Models and Their Relevance to Media Effects

For many reasons, the Crawford-Sobel model is useful for analyzing media bias and its effects. To see this, first note that, similar to the model, a journalist (similar to the Sender) sends a message to a news consumer (similar to the Receiver). The message can be interpreted as the news story the journalist writes or produces. Second, usually the

---

[9] For instance, suppose a subject is type-2. Recall that she assumes her opponent will play strategies R_1 and S_1 . Now suppose we want her to treat her opponent as if he has a level of sophistication, say, 1.4 degrees lower than hers. To do that we would require her to believe that with probability .4 her opponent plays strategies R_0 and S_0 and with probability .6 her opponent plays strategies R_1 and S_1 .

[10] Cai and Wang conducted three bounded-rationality analyses on a fraction of their data—the fraction where $d = 4.0$. One involved fitting McKelvey and Palfrey's AQRE model to their data. This model predicted that the s-m, m-p, and s-p correlations would be .326, .400, and .178. The average absolute deviation of the predicted value from the actual value was .063. In another analysis Cai and Wang classified subjects into several types (e.g. those who usually adopted an S_0 strategy as Senders; those who usually adopted an R_1 strategy as Receivers; those who seemed to optimize, given the empirical distribution of the other subjects; and so on). Unlike my analysis, Cai and Wang did not constrain this model to be consistent with the *one-difference* regularity. That is, for instance, if a subject was classified as an S_2 sender, then he or she was not constrained to be an R_1 receiver. This model was separated into two separate "scenarios" (which differed on assumptions about the approximately 20% of the subjects who could not be classified). The scenario that predicted the three correlations best had an average absolute deviation of .023. The third analysis, which they called the "Crawford equilibrium," classified subjects into three types of Senders: (i) those who play S_1 (ii) those who play S_0, and (iii) those who are sophisticated. The analysis classified Receivers similarly. The analysis did not constrain subjects to be consistent with the *one-difference* principle. This analysis had an average absolute deviation of .325.

journalist will have spent days gathering facts and talking to witnesses and experts. Consequently, like the Sender in the model, she will often know more about her story's topic than does the average news consumer.[11] Third, just as the Sender in the model can lie about the true state of the world, a journalist can mislead readers by strategically reporting some facts, while omitting others.[12] Fourth, the news consumer is usually a voter. Thus, like a Receiver in the model, he helps to set policy.

If the interaction between a journalist and a news consumer is indeed like the Crawford-Sobel model, then the model makes a powerful prediction. Namely—given its no-policy-bias implication—it predicts that real-world media effects should be nil. That is, even if journalists report the news, say, with a liberal bias, then news consumers should be able to discount that bias. The end result is that news consumers should behave no more liberally than they would if the news were perfectly unbiased.

---

[11] For instance, in Groseclose (2011) I conduct a case study of an article that appeared in the *Los Angeles Times.* The focus of the article was the number of African Americans who would join 2006 freshmen class at UCLA. The article reported several facts, very few of which the general reader would have known before reading the article. These included: (i) The number of African Americans expected to enroll was only 96. (ii) That number was 20 fewer than the year before. (iii) A UCLA sociologist, Darnell Hunt, had written a report criticizing the UCLA admissions process. (iv) According to the report, one reason for UCLA's low black admission rate was that, unlike UC-Berkeley, it did not use a holistic system to evaluate applicants. (v) The county in which UCLA resides is 9.8% black. In addition, the journalist had interviewed at least six people who had expert or an insider's knowledge about the topic.

[12] Again, in Groseclose (2011) I provide an illustration of this point with the *Los Angeles Times* article. One of the general questions that the article addressed was the following: Was UCLA discriminating against African Americans in its admissions process? Alternatively, was UCLA discriminating *in favor* of African Americans—that is, granting affirmative action? As I discuss, by omitting several key facts, the article insinuated that UCLA was discriminating *against* African Americans. Such omitted facts included: 1) Among transfer students, black enrollment rose by 22 in 2006—thus among all undergraduate students, the net black enrollment rose slightly. 2) In order to boost minority admissions, UCLA had incorporated a *life challenge index*. The index boosted a student's chance for admission if, for example, she: (i) came from a poor family; (ii) went to an inferior high school; (iii) was raised by a single parent; or (iv) was a single parent herself. 3) Although the article noted that African Americans comprised 9.8% of Los Angeles County, it did not note that they comprised only 4.6% of the applicants to UCLA. 4) Of the six people quoted in the article, five had liberal political views and only one had conservative views. 5) Although the article noted the political views of the conservative ("Ward Connerly, the conservative former UC regent"), it did not do this for any of the five liberals.

The latter prediction is consistent with much of the early work examining media effects—especially that of Joseph Klapper (1960) and William McGuire (1986). The latter two scholars argued that the *minimal effects* hypothesis best describes the influence of the media.

The results of the Cai-Wang experiment, however, suggest that media effects should be significant. To illustrate this, I first review a concept that I introduce in Groseclose (2011)—the *media* λ.

To define the concept, I construct a simple model. Let x be the *natural views* of the average voter, and let m be the overall slant of the media. I assume that y, the views that we observe in the average voter, follows the equation,

$$y = (1-\lambda)x + \lambda m, \quad (1)$$

where λ is a parameter in [0,1]. The latter represents the effect of the media. I interpret x and m as exogenous variables; whereas I interpret y as endogenous—i.e., it is induced by x and m.

The variables y and m are directly observable. For instance, we can measure y by collecting survey data from voters. We can measure m by observing the content of news (see, for instance, Groseclose and Milyo, 2005, or Gentzkow and Shapiro, 2010, who provide numerical estimates of the slant of the media). Although x is not directly observable, it is precisely defined. In particular, it is the solution to two different thought experiments:

- I) Note, by (1), that if m=x, then y=x. That is, if the media were perfectly unbiased (m=x), then x equals the observed views of the average voter, y.
- II) If λ=0, then y=x. That is, if the media had no effect, then x equals the observed views of the average voter, y.

Further, x can be observed indirectly. Namely, if we have measurements of y, m, and λ, then we can use (1) to calculate x. Groseclose (2011), for instance, obtains the following estimates: $y \approx 50.4$, $x \approx 30.6$, $m \approx 58.9$, and $\lambda \approx .7$.

Note, from (1), that $dy/dm = \lambda$. Thus, if the media slant moves leftward by one (very small) unit, then the observed views of the average voter move leftward by λ units. In Groseclose (2011) I use this fact to estimate λ. Namely, I review research where—through an actual experiment or a natural experiment (e.g., Gerber, Karlan, and Bergan, 2009, and DellaVigna and Kaplan, 2007)—the slant of the media moved a small amount. I then observe the resulting change in the views of the average voter. My estimate of λ is the latter divided by the former.

We can conduct a similar exercise to estimate the media λ from the Cai-Wang results. First, consider the case where d=4.0. Note that the average message from the sender was 7.186. This can be interpreted as m, the average slant of the media. Now recall that the average state of nature was 5.0.[13] We can interpret this as x, the average natural (i.e. true) preference of the Receiver. Next, recall that the average policy chosen by the Receiver was 5.592. We can interpret this as y—the preference that was induced partly by the Sender's messages.[14]

Using some algebra, we can rewrite (1) as:

$$\lambda = (y - x) / (m - x).$$

[13] Recall note 4, where I weight observations so that this is necessarily true.
[14] Note that these interpretations are consistent with the thought experiments that I describe above. First, if the Sender always told the truth, then the average message, m, would equal *5.0* (i.e. m=x). If the Sender adopts such a strategy, and the Receiver knows this, then the Receiver would always set policy equal to the Sender's message (i.e. y=x). Second, if the Receiver always ignored the Sender's message (i.e. λ = 0), then he would optimally choose a policy of 5.0 each round. (i.e. y=x),

After substituting for *y*, *m,* and *x*, we obtain $\lambda = (5.592–5.0)/(7.186–5.0) = .271$. The following table lists the estimates of λ from the other versions of the game. It also lists the predicted λ from the above bounded-rationality model (where I assumed that 20% of the subjects were type-3, 20% were type-2, and so on).

d	estimate of λ	predicted estimate from bounded-rationality model
4.0	.271	.356
2.0	.315	.240
1.2	.325	.267
.5	.510	.800

Note that, across the four versions of the game, the median estimate and the median predicted estimate of λ are approximately .32.

What if Journalists and News Consumers are Less Strategic than the Subjects in the Cai-Wang Experiment?

Groseclose (2011) estimates that the media λ—for real-world journalists and real-world news consumers—is approximately .7. In this section I suggest that the Cai-Wang results, contingent on some reasonable assumptions, support that estimate.

For several reasons, it is reasonable to believe that the subjects in the Cai-Wang experiment were more strategic than real-world journalists and real-world news consumers. First, the subjects in the experiment—UCLA undergraduates—were likely more savvy at game-theoretic thinking than the average person. In fact, many of the subjects had previously taken a course in game theory.

Second, the Receivers in the experiment knew the exact direction and degree to which the Senders' preferences differed from their own. Translated to the real-world media, it would be as if every news consumer understood perfectly the degree to which a journalist's political views differed from his own on any news story he read or watched.

Third, for most rounds of the experiment, almost all of the Receivers had played a prior round as a Sender. Thus, they had experienced, first-hand, the incentives of Senders to send biased messages. Translated to the real world, it would be as if all news consumers had, at some point in their life, held jobs as journalists.

Fourth, after each round of the experiment, the Receivers learned the state of nature, and therefore they learned the precise degree to which the Sender's message was biased. They also learned the precise degree to which they were fooled by the message. Translated to the real world, it would be as if every newspaper reader were given a detailed report that explained all the biased aspects of every news story that he had ever read or watched.

Fifth, play was anonymous. Thus, the Senders in the experiment would face less shame for lying. Real-world journalists are almost never granted such anonymity.

Sixth, even if news consumers discount media bias—it is doubtful that journalists respond by exaggerating their bias to compensate for the discounting. That is, it is doubtful that many journalists' degree of sophistication is really type-2 or above. To see this, suppose, as an example, that—similar to the case study I examine in Groseclose (2011)—a journalist strongly favors affirmative action, and that, accordingly, she is willing to make it appear that a university is discriminating, say, to a weak degree, against black students. Further, she is willing to do this even though she knows that, in truth, the university is not discriminating. If her degree of sophistication is type-2 or higher, then her thought process must be something like the following: "My readers understand that I favor affirmative action. Therefore, they understand my incentives to misrepresent the truth. If my report suggests that the college discriminates only in a weak degree, then my readers will discount that and think that the college does not discriminate at all. Consequently, if I really want to change the attitudes of my readers, I need to misrepresent the truth even more—my story needs to suggest that the college discriminates to a strong degree." Although it is only conjecture, few journalists, I believe, think that strategically.

Now suppose we alter the bounded-rationality model to account for the above argument. Specifically, suppose we take all the subjects who had a sophistication level of at least type-1, and we make them one degree less sophisticated. This means that 20% are type-2; 20% are type-1; 40% are type-0; and 20% are spazzes.[15] Under this setup, it can be shown that each version of the game predicts a media λ of approximately .4.

Now suppose we alter the model again in this manner. This gives a distribution of subjects as: 20% type-1; 60% type-0; and 20% spazzes. Then, under this setup, each version of the game predicts a media λ of.8.

Note that, not counting the spazzes, the average type under the latter distribution is approximately one degree less sophisticated than the average type under the original distribution. Meanwhile, note that the resulting media-lambda prediction, .8, is approximately the value that I estimated for the real-world media, .7. Thus, if we are willing to believe that real-world journalists and real-world voters are, on average, one

[15] In the real-world context, a spaz news consumer is one who maintains his prior beliefs regardless of the news report he sees—or possibly because he receives no news report at all. A spaz journalist is one writes a story that confirms the reader's prior beliefs—regardless of the evidence she uncovers (i.e. regardless of what she learns about the state of the world).

degree less sophisticated than the Cai-Wang subjects, then my estimate and the Cai-Wang results are very consistent with each other.

Conclusion

As the results of the Cai-Wang experiment show, a systematic anomaly seems to occur in the way actors update their beliefs in some situations of asymmetric information. As this paper argues, the anomaly can be explained by bounded rationality. Specifically, consistent with many previous findings in behavioral economics, the typical subject seems to believe that his opponent is one level less strategic than he is. The anomaly has many implications for various real-world settings, one example of which involves the interaction of voters and the media.

References

Cai, Hongbin, and Joseph Tao-Yi Wang. 2005. "Overcommunication in Strategic Information Transmission Games," *Games and Economic Behavior,* 56: 7-36.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games, *Quarterly Journal of Economics*, 119: 861-98.

Costa-Gomes, M., Crawford, V. 2004. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study. Working paper. UCSD.

Costa-Gomes, M, Vincent Crawford, and B. Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study, *Econometrica*, 69: 1193-1235.

Crawford, Vincent. 2003. "Lying for Strategic Advantage: Rational and Boundedly-Rational Misrepresentation of Intentions," *American Economic Review*, 93: 133-49.

Crawford, Vincent, and Joel Sobel. 1982. "Strategic Information Transmission," *Econometrica*, 50: 1431-1451.

DellaVigna, Stefano, and Ethan Kaplan. 2007. "The Fox News Effect: Media Bias and Voting." *Quarterly Journal of Economics*, 122(August): 1187-1234.

DellaVigna, Stefano, and Mathew Gentzkow. 2010. "Persuasion: Empirical Evidence" *Annual Review of Economics*, Volume 2 (March), forthcoming.

Dickhaut, John, Kevin McCabe, and Arijit Mukherji. 1995. "An Experimental Study of Strategic Information Transmission," *Economic Theory*, 6(3): 389-403.

Ederer, Florian, and Ernst Fehr. 2009. "Deception and Incentives: How Dishonesty Undermines Effort Provision." Manuscript. http://www.anderson.ucla.edu/faculty/florian.ederer/dishonesty.pdf

Gentzkow, Mathew, and Jesse Shapiro. 2010. "What Drives Media Slant? Evidence from U.S. Newspapers," *Econometrica*. Forthcoming.

Gerber, Alan, Dean Karlan, and Daniel Bergan. 2009. "Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions," *American Economic Journal: Applied Economics*, 1 (2): 35-52.

Gilligan, Thomas, and Keith Krehbiel. 1987. "Collective Decision-Making and Standing Comittees: An Informational Rationale for Restrictive Amednment Procedures," *Journal of Law Economics and Organization* 3: 287-335.

Groseclose, Tim. 2011. *Left Turn: How Liberal Media Bias Distorts the American Mind.* New York: St. Martin's Press.

Groseclose, Tim, and Jeff Milyo. 2005. "A Measure of Media Bias," *Quarterly Journal of Economics.*

Kartik, Navin. 2009. "Strategic Communication with Lying Costs," *Review of Economic Studies,* 76: 1359-95.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani. 2007. "Credulity, Lies, and Costly Talk," *Journal of Economic Theory*, 134 (May): 93-116.

Klapper, Joseph T. 1960, *The Effects of Mass Communication*. New York: Free Press.

Krehbiel, Keith. 1991. *Information and Legislative Organization.* Ann Arbor, Michigan: University of Michigan Press.

McGuire, William J. 1986. "The Myth of Massive Media Impact: Savagings and Salvagings," *Public Communication and Behavior,* 1: 173-257.

McKelvey, Richard, and Thomas Palfrey. 1995. "Quantal Response Equilibria in Extensive-Form Games," *Experimental Economics*, 1: 9-41.

Stahl, D. and P. Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, 10:218-54.